CENTRO DE INVESTIGACIÓN Y DOCENCIA ECONÓMICAS, A.C.

ESTIMATING THE ELASTICITY OF INTERTEMPORAL SUBSTITUTION FOR MEXICO: A SYNTHETIC COHORT PANEL APPROACH

TESINA

QUE PARA OBTENER EL TÍTULO DE

LICENCIADA EN ECONOMÍA

PRESENTA

MARIANA RUIZ ROLDÁN

DIRECTOR DE LA TESINA: DR. JUAN RAMÓN HERNÁNDEZ GONZÁLEZ

CIUDAD DE MÉXICO 2024

Acknowledgements

Quiero expresar mi más profundo agradecimiento a todas las personas que han hecho posible la elaboración de esta tesina.

A mi familia, por su constante apoyo a lo largo de mis estudios. En especial, a mi mamá, Sandra, por asegurarse de que siempre tuviera todo lo necesario para llegar hasta aquí.

A mi asesor, Juan Ramón, por la orientación, enseñanzas y oportunidades que me ha brindado. No hay palabras suficientes para expresar mi gratitud y admiración hacia él.

A mis lectores, por sus valiosos comentarios y sugerencias que enriquecieron este trabajo.

A mis amigos cideítas, Alan y Rodrigo, por ser excesivamente pacientes y buenos conmigo, así como por enseñarme más de lo que sola hubiera aprendido.

Finalmente, al CIDE, por permitirme conocer a personas tan diversas y fascinantes, por mostrarme el valor del conocimiento y por despertar en mí el deseo de seguir aprendiendo, pues esta carrera apenas comienza.

Summary

This dissertation estimates the Elasticity of Intertemporal Substitution (EIS) for Mexico, a crucial parameter for economic policy, as it measures how households adjust their consumption in response to changes in their assets' expected returns. Using a synthetic cohort panel constructed from ENIGH surveys spanning 1996 to 2022 and employing the Generalized Method of Moments (GMM), it addresses aggregation bias and weak identification issues commonly inherent in time-series data and interest rate predictions, unlike previous attempts for Mexico.

The findings show variable EIS estimates across different specifications and GMM estimators. However, they consistently reveal a statistically significant responsiveness of consumption to income growth, suggesting a prevalent income effect. This may reflect underlying liquidity constraints or impulsive consumption patterns akin to hand-to-mouth behavior. Moreover, incorporating additional instruments and controls not only refines EIS estimates but also consistently points to a statistically significant EIS typically below one.

Contents

List of Figures

List of Tables

1 Introduction

The elasticity of intertemporal substitution (EIS) in consumption is a key parameter in dynamic models of macroeconomics and finance, as it captures households' consumption responsiveness to changes in their assets' expected returns. To fully grasp its meaning and relevance, it is necessary to first understand the dual forces that come into play when interest rates rise: the income and substitution effects.

Higher interest rates increase future cash flows and overall wealth, prompting households to potentially spend more today—this is the income effect. Conversely, higher interest rates also raise the opportunity cost of current consumption, encouraging households to save more to benefit from future higher returns—this is the substitution effect. The prevailing effect depends on the value of the EIS. At a critical EIS of unity, both effects offset each other; while at a high EIS, above unity, the substitution effect dominates, indicating households are more sensitive to shifts in monetary or tax policies.

Furthermore, the EIS plays a pivotal role in determining the magnitude of consumption growth. Within the Solow model framework, it influences how variations in the interest rate, depreciation rate, and intertemporal discount factor affect this. For instance, a higher EIS implies that households are more inclined to defer consumption, leading to faster economic expansion over time.

A natural question that arises in any given economy is whether the EIS is below or above unity, and implicitly, what its value might be. While standard general equilibrium macroeconomic models are calibrated with an EIS above unity, empirical estimates show considerable variation (see Havranek et al. (2015) and Thimme (2017) for a deeper exploration of EIS heterogeneity), with some findings at odds with these assumptions. For instance, the long-run risk asset pricing model proposed by Bansal and Yaron (2004), which has the potential to solve the equity premium puzzle, assumes an EIS of 1.5. This leads to cyclical price-dividend ratios, a low and stable risk-free rate, and a significant equity premium. Similarly, other long-run risk asset pricing models, like those by Drechsler and Yaron (2010) and Ai (2010), consider an EIS above unity, e.g., 2.

Following the widely used approach of Hall (1988), who estimated the EIS by examining the consumption Euler equation, and building on the methodology of Attanasio and Weber (1995), who adopted a synthetic cohort panel approach to circumvent the aggregation bias inherent in time series data, I endeavor to estimate the EIS for Mexico. This is particularly pertinent given the limited recent research in this area within the country. To my knowledge, the only prior attempt is by Arrau and van Wijnbergen (1991). Although they apply the model of Epstein-Zin, which decouples the EIS from the risk aversion coefficient unlike the expected utility framework with CRRA utility functions, their EIS estimates carry certain caveats.

The first and most relevant caveat is that even though their EIS estimate of 1.2 is statistically different from zero, they cannot conclude whether it differs from 1 due to high imprecision. This could stem from their use of seasonally unadjusted time series data and aggregate consumption data that do not distinguish between durable and non-durable goods. Additionally, the nature of aggregate data does not exclude households facing liquidity constraints, for whom the Euler equation may not hold. The second caveat relates to the period under study: their estimates are based on quarterly data spanning from 1980 to 1990.

With this in mind, my objective is to leverage more recent microdata to address the challenges posed by aggregation biases. Nonetheless, I will rely on the time-additive utility function to derive the EIS (expected utility framework), where it is the inverse of the Arrow-Pratt measure of relative risk aversion (RRA)—a critical parameter for determining the net effect of higher future aggregate payoffs on their present value. In this context, highly risk-averse households might also exhibit a low EIS. However, since an infinitely high RRA does not align with observed risk-taking behavior, it is important to decouple the direct linkage between EIS and RRA. Therefore, this work's findings do not automatically reflect risk aversion levels.

To maintain simplicity and comparability with previous studies predominantly employing CRRA utility functions, I also rely on the Generalized Method of Moments (GMM) estimator proposed by Hansen and Singleton (1982). This is robust to weak identification driven by weak instruments for either consumption growth or interest rates.

2 Literature Review

Overall, most efforts to estimate the EIS have concentrated on either US or UK data, providing powerful insights into potential paths for uncovering this parameter.

Starting with Hall (1988), we can trace the initial attempts to estimate the EIS. Unlike earlier studies that based their EIS estimation on consumption or savings functions—often obtaining values as high as 1—he relies on the idea that consumers plan to change their year-to-year consumption based on their expectations of real interest rates, similar to the now broadly used Euler equation for consumption. Hall highlights the difficulty of defining consumption or savings functions, as the relationship between consumption, interest rates, and income depends on a broader context that can change despite agents always aiming to maximize the same utility function.

By applying some transformations to the variables of interest, consumption and real interest rates, extracted from the Euler equation, he concludes that US aggregate consumption in the twentieth century does not imply a positive value for the EIS. Most of his estimates are quite precise, suggesting it is unlikely to be much above 0.1 and may well be zero. Subsequent EIS estimates have been very heterogeneous, but many still point out the possibility that the EIS—at least for the US—may indeed be statistically zero.

More recently, Yogo (2004) argues that conventional Instrumental Variables (IV) methods result in the empirical puzzle that the EIS is statistically less than 1, while its reciprocal (the RRA coefficient under the expected utility framework) is not different from 1. Hence, he employs various techniques robust to weak instruments, as well as tests and confidence intervals, ranging from k-class estimators—including Two-Stage Least-Squares (TSLS), Limited Information Maximum Likelihood (LIML), and Fuller-k (Fuller, 1977)—to continuously updated GMM. This approach allows him to solve the puzzle: weak instruments are relevant in practice and conventional t-tests can lead to misleading evidence. Even so, the EIS appears to be less than 1 and not significantly different from zero for eleven advanced economies. It is worth noting that these methods are not a definitive solution to weak instruments, as confidence intervals remain uninformative when identification is poor. However, they help prevent erroneous inferences.

These studies relied on aggregate data when estimating the EIS. Nonetheless, many researchers have pointed out the biases these may introduce and have stressed the advantages of using microdata instead, as recommended by Ascari et al. (2021). The main is a greater likelihood of identifying a statistically significant EIS, as the econometrician has more control over the aggregation process and can account for sociodemographic or household-specific variables that would otherwise remain unobserved or averaged out. For this reason, the following literature review focuses on studies that have estimated the EIS using microdata, emphasizing the potential gains and issues associated with this approach.

2.1 Data structure implications

Attanasio and Weber (1993) compare EIS estimates using aggregate data versus cohort-average data—concentrating on households headed by individuals born between 1930 and 1940 from the UK—to ascertain if different data structures alter the principal conclusions. In general, they discover that the EIS values derived from aggregate data are consistently lower than those from cohort data, ranging from 0.3 in the former to 0.8 in the latter. The 95% confidence intervals include unity only in cohort-based estimates.

Their analysis indicates that using aggregate data often leads to the rejection of the model's implications, such as the Permanent Income Hypothesis (PIH) or the Life Cycle Hypothesis (LCH)—both of which presume that individuals smooth consumption over their lifetime—and the overidentification restrictions that validate the exogeneity of instruments for the real interest rate. They argue that controlling for labor-supply and sociodemographic factors reduces the sensitivity of consumption growth to income in cohort data, but not in aggregate data. This suggests the presence of liquidity constraints and the inappropriateness of imposing uniform preferences across different cohorts in aggregate analyses.

A further issue with using aggregate data is that it typically involves the arithmetic mean of consumption rather than the geometric mean, which can be derived from microdata. Although the estimates of EIS do not differ significantly, the geometric mean—from cohort-average data—yields more precise estimates and avoids the higher order serial correlation observed in aggregate analysis. Attanasio and Weber (1995) further reinforce these findings, showing that time series data often reject the overidentification restrictions due to an incorrect aggregation process that overlooks non-linearities, resulting in an inadequately determined

EIS.

Subsequent studies by Alan et al. (2019) assess the estimation of the EIS using the linearized Euler equation. They estimate six life-cycle models that vary by individual impatience rates, the potential for zero income occurrences, and the type of credit constraints. By numerically solving each model—based on an isoelastic utility function with a relative risk aversion coefficient of 4 (EIS of 0.25)—they generate consumption functions for a representative consumer over 40 periods and simulate consumption trajectories for 10,000 ex ante identical individuals. These simulations form the basis for a series of Monte Carlo experiments replicating three data structures: an ideal 40-period long panel, a commonly encountered 14-period short panel, and a synthetic cohort panel constructed as an alternative to either long or short panels.

Their results show that the long panel is the most effective in accurately recovering the EIS, due to sufficient data variation, although measurement error in consumption data can lead to more inflated and less precise estimates. Conversely, estimates from the short panel result in lower EIS values and are less efficient. Estimations using synthetic cohort panel data exhibit fewer validity issues compared to available short panels, whereas the standard instruments—such as lags of consumption growth, income, and interest rates—prove less effective for explaining interest rates. While the EIS estimates from synthetic cohort panel data closely mirror those of the long panel in terms of average bias, the precision significantly diminishes.

Ultimately, these studies highlight the advantages of employing microdata over aggregate or time-series data, particularly underscoring the benefits of adopting a synthetic cohort approach when real panels are unavailable, as is the case of Mexico.

2.2 Liquidity constraints and other sources of EIS heterogeneity

As previously noted, a significant limitation of using aggregate data to estimate the EIS is its potential to overlook individuals facing liquidity constraints, such as hand-to-mouth consumers. Nonetheless, there remains an ongoing debate about whether microdata can effectively address this concern, especially in light of the frequent rejection of overidentification restrictions in aggregate analyses. The evidence on this issue is varied.

Zeldes (1989) examines the Euler equation for two distinct groups: those with liquidity constraints and those without. He finds that the equation does not apply to the constrained group, suggesting that EIS estimates should focus on individuals free from liquidity constraints. Conversely, Runkle (1991) argues that the rejection of the Permanent Income Hypothesis (PIH) in aggregate data is not necessarily due to liquidity constraints but rather to aggregation bias: households likely base their economic expectations on individual past experiences rather than on aggregate data predictions.

Continuing the discussion on the variability of EIS across different groups, Vissing-Jørgensen (2002) emphasizes the importance of considering limited asset market participation. Using data from the US Consumer Expenditure Survey (1980-1996) and a synthetic cohort approach, she estimates the EIS for shareholders, non-shareholders, bondholders, and non-bondholders. Her findings indicate a higher EIS for shareholders, especially among the wealthiest, compared to non-shareholders, with EIS values of 0.4 versus nearly zero, respectively.

A similar pattern is observed for bondholders, with an EIS of approximately 0.8—reaching 1.6 for the top tier of bondholders—versus a statistically insignificant EIS for non-bondholders. For households possessing savings accounts but no bonds or stocks, she identifies a negligible and statistically indistinguishable EIS from zero. This suggests that while these assets typically track Treasury rates, the minimal financial wealth of these families means they do not engage in sophisticated intertemporal optimization, or they might be net lenders or borrowers who are operationally constrained.

Similarly, Attanasio et al. (2002) use the UK Family Expenditure Survey to estimate EIS for shareholders and non-shareholders, relying on a synthetic cohort panel approach based on the predicted probability of stock ownership from 1980 to 1995. They report an EIS close to 1 for shareholders, whereas the structural parameters—including the EIS—for non-shareholders appear implausible.

In a study together, Vissing-Jørgensen and Attanasio (2003) employ the data from Vissing-Jørgensen (2002) to estimate the EIS based on three different Euler equations—two including bonds and stocks, and a third considering human capital returns. They find that the EIS is less than 1 for the overall household data but greater than 1 when focusing solely on shareholders.

Lastly, Guvenen (2006) studies a real business cycle model by introducing two sources of heterogeneity: limited participation in asset markets and different EIS among households. He demonstrates that the inconsistency between EIS estimates is largely a consequence of assuming a representative agent. In this model, limited participation leads to wealth inequality—consistent with the asymmetric distribution of data observed in the US—implying that the properties of aggregate variables related to wealth are mostly determined by the high elasticity of shareholders, who own most of the economy's capital. Conversely, consumption is relatively uniformly distributed among households, also in line with US data, suggesting that aggregate consumption primarily reflects the low EIS of the majority.

Given all this evidence, one of the principal aims of this work is to build a synthetic cohort panel, in the absence of a sufficiently long panel, to control for both household-specific factors and the presence of liquidity constraints that might prevent estimating a precise or statistically significant EIS. It is important to mention that, due to the lack of detailed information on the assets held by each household, the best assumption is that only households from the first income quintile are credit constrained. This, of course, may not be entirely accurate.

3 Model

In the standard intertemporal optimization problem, a representative household h with time-separable preferences maximizes its expected lifetime utility derived from consumption. This utility is discounted by a “pure" discount factor β, and the household’s choices are constrained by a feasible allocation of resources across different periods.

The optimization problem can be formally expressed as:

$$\max_{\{C_s^h\}} \quad E_t \left[\sum_{s=t}^{T} U(C_s^h, z_s^h, v_s^h) \beta^{s-t} \right]$$

subject to the period-by-period budget constraint:

$$A_{s+1}^h = (1 + R_{s+1})(A_s^h + Y_s^h - C_s^h), \quad s = t, t+1, \ldots, T-1$$

where the terminal condition must satisfy:

$$A_T^h \geq 0$$

Here, E_t represents expectations based on information known at time t, I_t; C_s^h denotes the household’s non-durable consumption in period s; z_s^h includes observable variables such as demographic or labor supply factors; and v_s^h captures unobservable influences like taste shocks or measurement errors that potentially affect the marginal utility derived from consumption. Y_s^h, A_s^h and R_s are the household’s income, non-human wealth, and the one-period real interest rate, respectively.

If the household optimally plans its consumption and has no liquidity constraints, the first-order condition yields the consumption Euler equation:

$$E_t \left[\beta \frac{U'(C_{t+1}^h, z_{t+1}^h, v_{t+1}^h)}{U'(C_t^h, z_t^h, v_t^h)} (1 + R_{t+1}) \right] = 1 \tag{1}$$

which suggests that the household allocates consumption between periods such that it cannot expect to make itself better off by reducing its consumption today in exchange for more consumption later. A widely adopted specification for utility in this context assumes a Constant

Relative Risk Aversion (CRRA) utility function:

$$U(C,z,\nu) = \frac{C^{1-\gamma}}{1-\gamma}\exp\left(\theta z + v\right) \tag{2}$$

where γ is the Arrow-Pratt measure of RRA. Substituting equation (2) into (1) gives:

$$E_t\left[\beta\left(\frac{C^h_{t+1}}{C^h_t}\right)^{-\gamma}\exp(\theta\Delta z^h_{t+1} + \Delta v^h_{t+1})(1+R_{t+1})\right] = 1 \tag{3}$$

This equilibrium condition can provide orthogonality conditions crucial for estimating the parameters of the utility function and testing overidentifying restrictions. Considerations like labor supply choices within z help model preferences optimally, even when dealing with institutional constraints or corner solutions.

Ignoring the presence of taste shocks or measurement error (v_t) simplifies the estimation, provided sufficient orthogonality conditions exist. Under rational expectations, these expectations are assumed to be orthogonal to all available information at time t (I_t), simplifying their empirical implementation.

However, preference heterogeneity or the presence of v complicates parameter estimation due to the non-linear characteristics of the Euler equation (3). While using aggregate time series data might be feasible under the assumption of a representative consumer without significant measurement error, this condition is often not met in practical scenarios. Consequently, it is advisable to employ models that are linear in parameters, which can more effectively handle the inevitable v in individual-level data, irrespective of its source.

An alternative approach is the log-linear approximation of the Euler equation. This is obtained by taking the natural logarithm of both sides of (3), followed by using the second-order Taylor approximation for $\ln(1+x)$:

$$\ln\left(\frac{C^h_{t+1}}{C^h_t}\right) = \frac{1}{\gamma}\left(k_t + \theta\Delta z^h_{t+1} + \ln(1+R_{t+1}) + \Delta v^h_{t+1} + u^h_{t+1}\right) \tag{4}$$

Assuming the variables involved are log-normal, then:

$$\begin{aligned} k_t =& \ln(\beta) + \gamma^2 \left(\text{var}_t \left(\ln \left(\frac{C_{t+1}^h}{C_t^h} \right) \right) + \text{var}_t \left(\ln(1 + R_{t+1}) \right) \right) \\ & - 2\gamma \, \text{cov}_t \left(\ln \left(\frac{C_{t+1}^h}{C_t^h} \right), \ln(1 + R_{t+1}) \right) \end{aligned} \tag{5}$$

Where the t subscripts indicate the second moments are conditional on information available at time t, I_t. If the conditional distribution of the relevant variables deviates from log-normality, k_t will also include higher conditional moments. Since these moments are generally not observable, it is useful to rewrite (4) as:

$$\ln \left(\frac{C_{t+1}^h}{C_t^h} \right) = \frac{1}{\gamma} \left(\bar{k} + \theta \Delta z_{t+1}^h + \ln(1 + R_{t+1}) \right) + e_{t+1}^h \tag{6}$$

In this formulation, $\bar{k}$ includes the natural logarithm of the discount factor β and the unconditional mean of the second and higher moments of consumption growth and real interest rates. The residual e_t captures expectational errors u_t, unobserved heterogeneity v_t, and the deviations of k_t from $\bar{k}$. Now, however, only the parameters γ and θ from (3) can be estimated, generally using Instrumental Variables (IV) or Generalized Method of Moments (GMM). Instruments for IV or GMM estimation should be variables known to the household at time t and should be exogenous to the decision-making process regarding consumption.

Obviously, equations (4) and (6) serve as approximations to (3), nevertheless, several studies have investigated the specific conditions under which these yield consistent estimates of θ and γ. Attanasio and Low (2004), for instance, show that the log-linear approximation of the Euler equation does not introduce bias as long as the innovations in the conditional variance of consumption, included in $\bar{k}$ and determined endogenously, are not correlated with the instruments. They also find that grouping households reduces the variability of idiosyncratic income shocks, thereby improving the efficiency of estimates. Additionally, they observed that estimates using the non-linear GMM approach tend to be less reliable compared to those obtained from the linear model, especially when the discount factor β is high, mirroring situations akin to those faced by liquidity-constrained households.

Furthermore, Attanasio and Low underscore that including the conditional variance of consumption growth to account for precautionary savings does not substantially change the EIS estimates, though it increases the standard errors, attributable to weak instruments for higher-order terms. Complementary findings by Gomes and Ribeiro (2015) indicate that the impact of precautionary savings on consumption is constrained by low risk levels. More recently, Alan et al. (2019) have confirmed that a second-order approximation of the Euler equation does not improve estimates, primarily due to weak instruments, and that non-linear GMM approaches are adversely affected by issues related to handling measurement error in consumption data.

Based on these insights, the chosen structural equation for this analysis is:

$$\Delta \ln(C^h_{t+1}) = \text{constant} + \sigma \ln(1 + R_{t+1}) + \theta' \Delta z^h_{t+1} + e^h_{t+1} \tag{7}$$

where $\sigma = \frac{1}{\gamma}$ is the EIS, and z includes demographics and labor-supply variables.

4 Data

4.1 Building a synthetic cohort panel

Thus far, I have emphasized numerous advantages of employing micro-level data over time-series or aggregate data, with the primary benefit being the control of aggregation bias. However, the lack of a sufficiently long panel dataset for Mexico has guided my methodology towards the approach of Attanasio and Weber (1995), who developed a panel by averaging household observations from multiple cross-sectional datasets.

In Mexico, the Encuesta Nacional de Gasto e Ingreso de los Hogares (ENIGH) serves as a rich source of consumption expenditure and sociodemographic data across a wide array of households, enabling the creation of representative cohorts over time. Administered biennially by INEGI since 1992—and initially conducted every four years starting in 1984—this survey provides a substantial temporal dimension for analysis. Therefore, I have selected ENIGH surveys from 1996 to 2022, encompassing a total of 15 surveys.

Due to methodological changes across different ENIGH series, I include only those household variables that can be consistently measured across the surveys. These include quarterly consumption expenditure—–which accounts for monetary outlays on non-durable items like food, clothing, education, leisure, transportation, transfers to third parties, health, rent, and other personal goods—–and current income, encompassing wages, rents, transfers, and other sources of income. Additionally, I consider variables like the number of family members and earners, the age and years of education of the household head, the number of family members by age group, and the general minimum wage.

All financial figures, including income and consumption, are adjusted to constant prices using the Mexican Consumer Price Index (INPC), with 2018 as the base year. This adjustment ensures that the analysis reflects real economic terms, removing the distortions caused by inflation.

Continuing with the tradition established by Attanasio and Weber (1995), my analysis includes cohorts composed of households headed by individuals born between 1962 and 1975. This age grouping ensures that the youngest cohort head is 21 years old in 1996—the beginning of the period under analysis—and the oldest cohort head is 60 years old in 2022, the concluding year. Accordingly, one cohort averages observations from households with heads born in 1965, while another averages those from households with heads born in 1966, and so forth. Table 6 in the Appendix details the average number of households per cohort. With fourteen cohorts consistently observed over thirteen periods, excluding the initial years to accommodate lagged variables, I manage a balanced panel comprising 182 observations.

The method of grouping by birth year rather than by age significantly enhances the study of life cycle behavior by allowing for the concurrent tracking of cohorts as they age. This avoids potential misinterpretations that might arise from age-based grouping in the presence of cohort effects.

Regarding the 3-month interest rate, I use the Mexican Treasury Bill rate, Cetes, considered a risk-free benchmark. Since the ENIGH surveys do not provide details on whether households actively participate in the asset market or the specific interest rates they encounter, employing Cetes serves as a neutral choice, under the assumption that market rates exhibit a similar dynamic. Without other means to identify asset-holding households, I control for liquidity constraints by excluding those in the lowest income quintile of each survey, presuming that households above this threshold may invest in Cetes, the risk-free asset.

4.2 Life cycle behavior

A key limitation of using a synthetic cohort panel approach, as opposed to a full panel, is that it does not track individual households throughout their entire life cycle. Instead, it captures the average consumption behavior of homogeneous groups as they age.

Figures 1 and 2 display the average real (non-durable) log-consumption and average log-income over time for all cohorts, respectively, including younger and older cohorts not used in the core estimation to depict a holistic view of life cycle behavior. Each colored segment in these figures represents the mean log-consumption or income for a specific cohort as it progresses through different life stages.

Notably, both consumption and income exhibit a hump-shaped trajectory and tend to move in tandem, reflecting a strong correlation between these two. This pattern challenges the life cycle hypothesis by suggesting that households may not be smoothing their consumption as expected throughout their lives. Furthermore, consumption tends to decline more sharply than income towards the end of the life cycle, possibly because the income data includes transfers that older household heads continue to receive.

Figure 1

Average Logarithmic Consumption by Household Head Age and Cohort

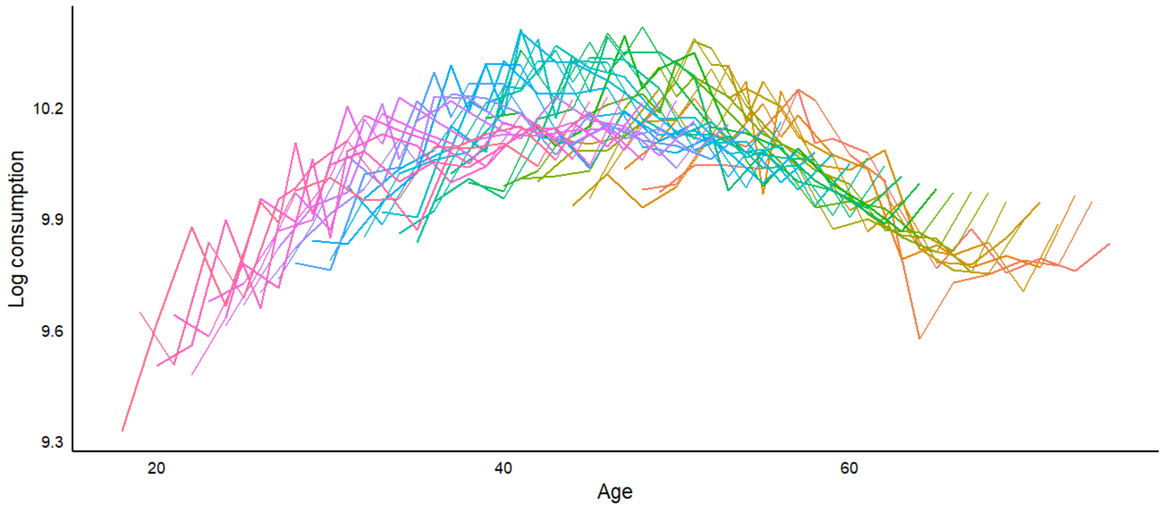

Note: The log of non-durable consumption is reported in real terms, adjusted to 2018 prices. Cohorts with household heads born between 1947-1961 and 1976-1978 were considered but not included in the estimation. **Source:** Own elaboration based on data from the ENIGH surveys, 1996-2022.

Figure 2

Average Logarithmic Income by Household Head Age and Cohort

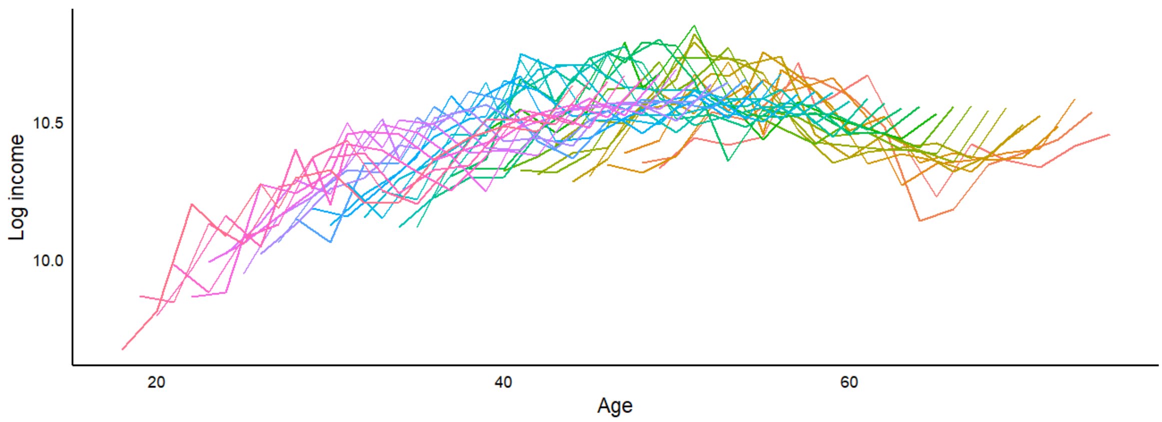

Note: The log of income is reported in real terms, adjusted to 2018 prices. Cohorts with household heads born between 1947-1961 and 1976-1978 were considered but not included in the estimation. **Source:** Own elaboration based on data from the ENIGH surveys, 1996-2022.

Despite these observations, Attanasio and Weber (1995) caution that the hump-shaped behavior alone is not sufficient to dismiss life cycle hypothesis, as other sociodemographic factors, such as family composition, might significantly influence consumption decisions. To illustrate, Figure 3 plots the evolution of family size composition for all cohorts throughout their lifespans. Each colored segment in this figure delineates the progression of family size for each cohort. As seen, family size also follows a hump-shaped pattern, closely aligning with the fluctuations in consumption decisions.[1]

Figure 3

Average Logarithmic Family Size by Household Head Age and Cohort

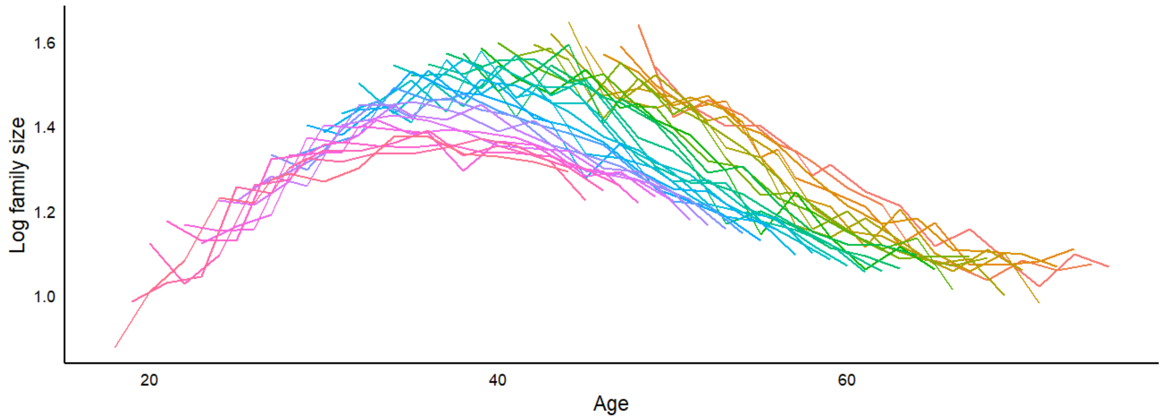

Note: Cohorts with household heads born between 1947-1961 and 1976-1978 were considered but not included in the estimation. **Source:** Own elaboration based on data from the ENIGH surveys, 1996-2022.

Additionally, several effects like life cycle, cohort, and business cycle might influence consumption. By regressing average cohort consumption on cohort dummies and a fifth-degree age polynomial, following Attanasio and Weber (1995), both cohort and life cycle effects can be isolated, allowing the residuals to be interpreted as time effects. This approach helps to distinguish the influences on consumption, attributing observed trends to a combination of cohort and age effects.

Figure 4

Average Cohort Consumption Over Time (Age and Cohort Effects Removed)

Note: Consumption is reported in real terms, adjusted to 2018 prices. Cohorts with household heads born between 1947-1961 and 1976-1978 were considered but not included in the estimation. **Source:** Own elaboration based on data from the ENIGH surveys, 1996-2022.

[1] However, to conclusively determine if consumption behavior correlates with income—and thereby either support or refute the Life Cycle or Permanent Income Hypotheses—further estimation is required.

As illustrated in Figure 4, there is remarkable synchronization among cohorts' consumption, with noticeable peaks in 2004 and 2006, and a significant decline during the 2020 crisis. Figure 5 further shows that average consumption aligns closely with cohort-specific consumption patterns.

Figure 5
Detrended Aggregate Consumption

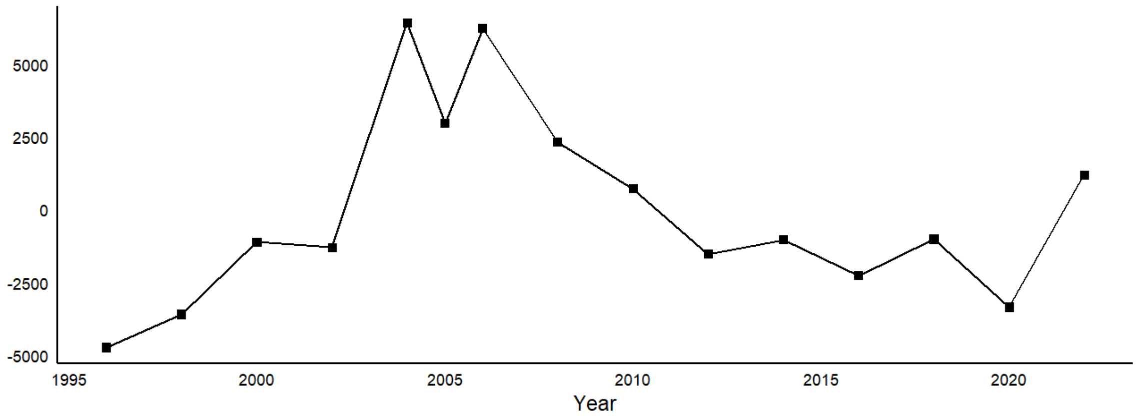

Note: Consumption is reported in real terms, adjusted to 2018 prices. **Source:** Own elaboration based on data from the ENIGH surveys, 1996-2022.

5 Econometric Analysis

5.1 Introducing Generalized Method of Moments[2]

Consider the following classical linear regression model:

$$y = X\beta_0 + u \tag{8}$$

As specified in equation (7), let $y = \Delta \ln C$ represent an $NT \times 1$ vector of the changes in log consumption for all N cohorts over T periods. The $NT \times (k+r+1)$ matrix $X = [\mathbf{1}, \ln(1+R), \Delta z]$ includes the explanatory variables that capture the effects of the real interest rate ($\log(1+R)$), other household variables (Δz), and a column of ones ($\mathbf{1}$) to account for the constant term. Here, k denotes the number of endogenous variables, including the interest rate and household choice variables at time $t+1$, and r represents the number of exogenous variables, such as the age of the household head. The vector u is an $NT \times 1$ vector, including both an expectational component and an MA(1) measurement error component for each cohort and time period. The $NT \times (m+r+1)$ matrix Z comprises r exogenous regressors directly from X, m valid instruments, and a column of ones ($\mathbf{1}$) to account for the intercept. These instruments typically consist of lags of the endogenous variables within X.

Under the assumption of instrument exogeneity, i.e., the errors u are uncorrelated with all the exogenous regressors, the $(k+r+1) \times 1$ parameter vector β_0 in (8) can be estimated by solving $E\left[(y - X\beta_0)'Z\right] = 0$, which constitutes a system of $m+r+1$ equations involving the $k+r+1$ unknown elements of β_0.

When the errors are homoskedastic and the system is exactly identified ($m = k$), these population moments are replaced by their sample moments, and the system of equations $(y - Xb_0)'Z = 0$ can be solved for b_0, where its value serves as the IV estimator of β_0. However, when the system is overidentified ($m > k$), the equations in the system cannot be simultaneously satisfied by the same value of b_0 due to the presence of more equations than unknowns.

An approach to solving the problem of estimating β_0 in such cases of overidentification is to use the GMM estimator, which balances the need to satisfy each equation by minimizing

[2] This section draws heavily on Attanasio and Weber (1995).

a quadratic form involving all these, known as moment conditions. Specifically, let H be an $(m + r + 1)$ symmetric positive semidefinite weight matrix and $\hat{\beta}_0$ the estimator that minimizes:

$$\min_{b_0} \quad (y - Xb_0)'ZH^{-1}Z'(y - Xb_0) \tag{9}$$

Upon taking the derivative of the objective function with respect to b_0, setting the result expression to zero, and rearranging, the efficient estimator for β_0 is obtained:

$$\hat{\beta}_0 = P_{xz}^{-1}X'ZH^{-1}Z'y \tag{10}$$

where $P_{xz} = X'ZH^{-1}Z'X$ and the asymptotic variance-covariance matrix for β_0 is given by $\Sigma_{\beta_0} = P_{xz}^{-1}$. In cases of homoskedastic errors, the efficient β_0 estimator is the Two Stage Least Squares (TSLS) estimator, which is achieved by setting $H = Z'Z$ in (10). Conversely, in the presence of heteroskedastic errors, the efficient estimator becomes the GMM estimator, obtained by substituting $H = Z'\Omega Z$ in (10), as follows:

$$\hat{\beta}_0^{GMM} = P_{xz}^{-1}X'Z(Z'\Omega Z)^{-1}Z'y \tag{11}$$

Here, $P_{xz} = X'Z(Z'\Omega Z)^{-1}Z'X$ and Ω is an $NT \times NT$ block matrix. Each block of the main diagonal is a $T \times T$ matrix representing the variance-covariance matrix of the residuals for one cohort, with the diagonal and the band surrounding the diagonal differing from zero to reflect the MA(1) structure of each cohort's residuals due to measurement error. The off-diagonal blocks of Ω represent the correlation of the residuals between different cohorts and are assumed to be diagonal, reflecting only constant contemporaneous correlation.

For (11) to yield a feasible estimator, an estimate of Ω is necessary. This is achieved through a preliminary round of consistent estimates obtained by initially using the identity matrix in place of Ω, analogous to deriving β_0 from the TSLS estimator, as previously discussed. Following the approach of Attanasio and Weber (1995), the residuals from this first round are used to construct an estimate of $Z\hat{\Omega}Z$ that is robust to heteroskedasticity of unknown form. The estimate is computed using the following expressions:

$$Z'\hat{\Omega}Z = P_0 + \alpha_1 P_1 + \alpha_2 P_2, \quad 0 \leq \alpha_i \leq 1, \quad i = 1, 2; \tag{12}$$

where:

$P_0 = \frac{1}{N}\sum_{j=1}^{N}\frac{1}{T}\sum_{t=1}^{T} z_{j,t}z'_{j,t}\hat{u}^2_{j,t}$;

$P_1 = \frac{1}{N}\sum_{j=1}^{N}\frac{1}{T-1}\sum_{t=2}^{T}\left(z_{j,t}z'_{j,t-1}\hat{u}_{j,t}\hat{u}_{j,t-1} + z_{j,t-1}z'_{j,t}\hat{u}_{j,t}\hat{u}_{j,t-1}\right)$;

$P_j = \frac{1}{N}\sum_{j=1}^{N}\sum_{i=1}^{N}\frac{1}{T}\sum_{t=1}^{T}\left(z_{j,t}z'_{i,t}\hat{u}_{i,t}\hat{u}_{j,t} + z_{i,t}z'_{j,t}\hat{u}_{j,t}\hat{u}_{i,t}\right)$.

Here, $z_{j,t}$ is the $(k+r+1)\times 1$ vector of instruments for cohort j in period t; T, the number of periods each cohort is observed, is uniformly 13 across the balanced panel. N, the number of cohorts, is 14. The ad-hoc weights α_1 and α_2 ensure the estimated variance-covariance matrix is positive definite. Specifically, $\alpha_1 = 1 - \dfrac{1}{T+1}$, as suggested by Newey and West (1987) for addressing first-order autocorrelation, and α_2 is set to 1. However, empirically setting both weights to 1 does not significantly alter the point estimates.

Finally, the Sargan (Hansen) test, which tests the null hypothesis that all the overidentifying restrictions are valid against the alternative that some or none hold, relies on the GMM J-statistic given by:

$$J^{GMM} = \frac{1}{NT}(Z'\hat{u}^{GMM})'\hat{H}^{-1}(Z'\hat{u}^{GMM}) \tag{13}$$

where $\hat{u}^{GMM}$ are the residuals from equation (8), estimated with the feasible, efficient GMM from the second (or final) step estimation. $\hat{H}$ is the weight matrix used in the second step, specifically $\hat{H} = Z'\hat{\Omega}Z$. This statistic follows a χ^2 distribution with $m-k$ degrees of freedom. Up to now, the GMM estimator outlined here was the two-step GMM estimator proposed by Hansen and Singleton (1982). Nevertheless, the iterative GMM procedure can be extended beyond the first one by reestimating the matrix $\hat{H}$ as described in equation (12), but using the residuals obtained from the two step GMM estimator, β_0^2.

In this extended procedure, a new estimator, β_0^3, is constructed employing the updated $\hat{H}$ matrix. This process can be iterated further, yielding a sequence of estimators β_0^j. Specifically, during each iteration j, residuals are obtained using the estimated parameters from the previous iteration to refine the estimation. The matrix $\hat{H}$ is recomputed with these updated residuals, leading to a new estimator β_0^{j+1}. The iterative process continues until β_0^{j+1} converges to a stable value or until a predetermined number of iterations is reached.

5.2 How GMM estimators circumvent potential estimation issues

As previously discussed, in scenarios devoid of persistent household-specific effects, measurement error in consumption, or aggregate shocks not captured by interest rates—to name a few—it is feasible to directly estimate the non-linear Euler equation (3) using a Non-linear GMM estimator. Alternatively, the linearized Euler equation (4) can be estimated using TSLS, which assumes that the term v_t^h is independent both across households within each period and over time. However, since such ideal conditions are rarely met, it is crucial to demonstrate how the GMM estimator can address these issues.

First, as Attanasio and Weber (1995) noted, constructing synthetic cohort panels from cross-sectional data introduces a measurement error in levels, resulting in an MA(1) structure in first differences. For example, if a sample for a given period includes an exceptionally wealthy household, this will induce a positive measurement error in the consumption growth at time t, followed by a negative error at time $t+1$. Consequently, the error term e_{t+1}^h in (7) consists of two components: a white noise component reflecting expectational errors and an MA(1) component with a coefficient of -1.

Furthermore, the variance of the measurement error component of the residuals changes with cell size, as shown in Table 6 in the Appendix, leading to significant heteroskedasticity. Thus, when estimating standard errors, both heteroskedasticity and an MA(1) structure are accommodated, as demonstrated in matrices P_0 and P_1. Contemporaneous correlation among the residuals of different cohorts, which is influenced by some cohorts meeting the age restriction before or after a specified year, is also considered in the P_j matrix.

A second potential issue involves the persistence of household or cohort-specific effects. Although less likely since each cohort is defined by the household head's year of birth and updated annually with different households, persistent effects might still emerge if households within a cohort share additional common components beyond the head's year of birth. To address this, one can treat persistent cohort-specific effects as a form of autocorrelation. For instance, it is possible to allow not just for first-order correlation in each cohort (as in matrix P_1), but also for higher orders of correlation, specifically T, the number of times each cohort appears in the sample. However, allowing for more than first-order autocorrelation in each cohort and adjusting for this using Newey and West (1987) ad-hoc weights results in minimal changes in standard errors estimates, suggesting that this is not a significant concern.

Conversely, in the presence of fixed cohort effects, it is appropriate to demean all household-specific variables considered in the Euler equation (7) by the cohort mean before estimation. Nevertheless, this adjustment typically has little impact on the estimation results. Additionally, in scenarios involving both household-specific effects and measurement error in consumption, Runkle (1991) recommends differencing the interest rates and other household variables in the Euler equation and accounting for first-order autocorrelation when computing the residuals for each cohort. This, in fact, is the approach adopted here, though, the other approaches provide no evidence of cohort-specific effects concerns, aligning with expectations given the data construction.

Finally, testing for the presence of aggregate shocks not explained by interest rates employs a straightforward method, crucial for short-term analyses. Chamberlain (1984) highlights that the sample average of the orthogonality condition, $E(e^h_{t+1}|I_t)$ converges to zero as the number of time periods increases but not if the cross-sectional observations increase while time periods remain constant. If aggregate shocks are present, using time-specific dummy variables as instruments for estimating (7) becomes problematic due to their correlation with these shocks.

To test for the presence of aggregate shocks, I initially apply the simplest specification of the Euler Equation, using only the interest rate as the explanatory variable, and instrument it with a standard set including lags of inflation, nominal interest rates, consumption growth, and income growth (as specified in the Table 2 first column). I then re-estimate this model, adding time-dummy variables as instruments. The critical test involves comparing the J-statistics of these two specifications, which should follow a Chi-squared distribution with degrees of freedom equal to the number of added instruments: the difference—indicative of aggregate shocks presence—is minor, at 2.96, with a p-value of 0.98, suggesting no significant aggregate shocks.

5.3 Selecting instrument sets

Given that consumption decisions and interest rates are determined simultaneously at time $t+1$, endogeneity poses a significant concern due to potential reverse causality between these variables and the influence of omitted variables affecting both. This issue can be addressed by instrumenting with variables known at time t or strictly exogenous factors, such as the age or education level of the household head——an attribute that remains largely unchanged during the observed phase of their life cycle.

Likewise, measurement error in consumption introduces an MA(1) structure in the residuals. This suggests that second lags, or even further lags, of other variables might be more appropriate as valid instruments than first lags. However, since the households used to construct the variables in the estimated equation differ from those used in constructing the first-lagged instruments, there are no validity concerns with using first-lagged instruments. Moreover, due to the biennial frequency of each ENIGH survey, employing second or more distant lags could introduce weak identification issues. For example, the relevance of four-year-old interest rates in explaining current rates is not the same as in panels with more frequent observations.

Building on this, when estimating the Euler equation (7), it is advisable to include additional control variables that may also influence consumption, such as income growth and family size growth at time $t+1$: critical determinants of household consumption and observable in ENIGH surveys. Nevertheless, like the other variables discussed, these too are endogenous and may require the same set of instruments used for explaining interest rates.

Hence, I consider four basic sets of instruments, where each subsequent set introduces an additional instrument to the preceding set. This allows for the evaluation of whether the explanation of the endogenous variables, similar to a first-stage estimate in a TSLS setting, improves with the inclusion of extra instruments.

The first set encompasses the most basic and standard instruments typically employed when household-specific variables are not available: lags of inflation, the interest rate, consumption growth, and income growth. The inclusion of the latter is particularly pertinent here as it may serve as a more effective predictor for income dynamics observed at the micro-level. The second set expands on this by adding both first lags of the number of earners and family size growth, the latter of which is also used as a control in some specifications. The third set introduces the household head's age, and the fourth set further adds the years of education of the household head.

Table 1
F-statistics for First-Stage Regressions of Real Interest Rate, Income Growth, and Family Size Growth on Instrument Sets

	Real Interest Rate	Income Growth	Family Size Growth
Set 1:	64.06	8.12	—
inflation, nominal interest rate, consumption and income growth	[0.591]	[0.155]	
Set 2:	45.75	8.14	8.02
+ family size growth, earners	[0.610]	[0.218]	[0.215]
Set 3:	39.53	6.941	6.894
+ head age	[0.613]	[0.218]	[0.217]
Set 4:	50.03	6.362	11.381
+ head education	[0.698]	[0.227]	[0.344]

Note: R^2 values shown in brackets. All regressions adjust consumption (non-durable) and income growth for inflation using INPC, base year 2018. Interest rates are based on Cetes rates. Data from 1996-2022 ENIGH surveys, excluding the lowest income quintile. $NT = 182$ observations for each regression.

Table 1 summarizes these sets and provides the F-statistics for first-stage regressions of key variables in the Euler equation across the different instruments. Although these statistics alone may not definitively indicate the best instrument set, they are useful for identifying potential weak identification issues within the GMM framework.

Table 7 in the Appendix shows the Stock and Yogo (2005) critical values for the different number of instruments used for testing the null hypothesis that the bias of TSLS as a fraction of OLS bias is greater than 5%, 10%, 20%, and 30%. It is noteworthy that there are no concerns regarding the relevance of instruments for explaining interest rates, and the TSLS relative bias is not greater than 5% across all instrument sets considered. However, adding household-specific variables does not enhance the prediction of these rates, as they are common across households, i.e., aggregate or shared.

Conversely, for household-specific variables such as income growth or family size growth, more concerns about instrument relevance arise. For none of the sets is the null hypothesis of weak instruments by Stock and Yogo (2005) rejected, since the TSLS bias as a fraction of OLS bias is greater than at least 20%. This calls for the use of an estimator robust to weak instruments, e.g. GMM, as IV estimates may not be reliable.

Interestingly, the more extensive sets, 3 and 4, do not significantly improve predictions for income growth and may even introduce some noise. Nonetheless, the fourth set performs better for explaining family size growth than the third set, which only adds the household head's age and may not be more effective than the second set for explaining all variables.

For the most basic specification, where no independent variables other than the real interest rate are considered, the standard set may be the optimal choice. Table 8 in the Appendix offers a more detailed comparison of the impact of switching between these sets on the validity of instruments for each variable of interest. Overall, the findings consistently suggest that sets 2 and 4 might be suitable for scenarios requiring predictions of income growth and interest rates, while set 4 could be preferred when including family size growth as well.

6 Results

Tables 2 and 3 present eight specifications for estimating the EIS (the coefficient on the real interest rate), using four different data sets and adding controls such as income growth and family size growth. Tables 4 and 5 replicate these specifications using the Iterative GMM estimator instead of the Two-Step GMM, because in a well-specified model, conclusions about the validity of instruments, as indicated by the Sargan Criterion, should not change significantly (Davidson & MacKinnon, 2004). Figure 6 plots the confidence intervals of the EIS for each specification, allowing comparisons between the results obtained with different GMM estimators for the same specification.

The first specification, which uses no controls in the structural equation and a standard set of instruments, typically yields a high, yet imprecise EIS within the Two-Step GMM framework, while the Iterative GMM estimator results in an even higher and more precise EIS: 2.42 (0.320) versus 0.904 (0.485), where the standard errors are in parentheses.

The second specification, still employing the standard set of instruments but incorporating household income growth as a control in the Euler equation, leads to statistically significant EIS values with both the Two-Step and Iterative GMM estimators. The Iterative GMM estimator yields higher EIS values: 2.279 (0.389) versus 0.753 (0.325), with standard errors in parentheses.

These results highlight several typical issues when additional household variables are not available. First, consistent with EIS estimates using aggregate data and the Two-Step GMM, the EIS appears statistically insignificant without further controls, even when more aggregate data is available but lacks specific household variables beyond lagged income growth as an instrument, which may not be more informative for explaining interest rates. Second, both specifications (1) and (2) may not robustly estimate the EIS, as evidenced by significant shifts in EIS confidence intervals depending on the chosen GMM estimator. Additionally, the Sargan Criterion results from the second specification using the Iterative GMM estimator reject the hypothesis of instrument validity. Therefore, if income growth is included as a control in the Euler equation, it is necessary to use additional instruments for this variable.

Table 2
Two-step GMM estimates of Euler Equation for Consumption
$\Delta \ln(C_{t+1}) = constant + \sigma \ln(1 + R_{t+1}) + \theta' z_{t+1} + \epsilon_{t+1}$

	(1)	(2)	(3)	(4)
Real interest rate	0.904	0.753	0.758	0.664
	(0.485)	(0.325)	(0.289)	(0.359)
Income growth		0.524	0.889	0.939
		(0.079)	(0.072)	(0.081)
Family size growth				0.065
				(0.125)
Sargan Criterion	2.47	3.76	3.75	3.41
(*p-value*)	(0.48)	(0.15)	(0.43)	(0.33)
+ instruments	+ 2	+ 2	+ 4	+ 4
$(m - k)$				

Note: Standard errors are reported in parentheses under each coefficient. The first two columns use a standard instrument set, which includes lags of inflation, nominal interest rate, consumption growth, and income growth. The third and fourth columns add lags of family size growth and number of earners to this set. Both consumption (non-durable) and income growth are adjusted for inflation using the INPC with 2018 as the base year. Interest rates are based on the Mexican Treasury Bill rate, Cetes. The analysis excludes households from the lowest income quintile, covering the 1996-2022 ENIGH surveys. $NT = 182$ observations for each specification.

The third specification expands on this by including two household variables as instruments: the lags of family size growth and the number of earners. This adjustment leads to a notable increase in the impact of income growth on consumption growth, rising from 0.524 in the second specification to 0.889 in the third, based on the Two-Step GMM estimator. Also, the third specification appears more robust, as Iterative GMM estimations maintain consistent conclusions: the EIS remains statistically significant and crucially in neither case can the hypothesis of instrument validity be rejected.

Table 3
Two-step GMM estimates of Euler Equation for Consumption (Continued)
$\Delta \ln(C_{t+1}) = constant + \sigma \ln(1 + R_{t+1}) + \theta' z_{t+1} + \epsilon_{t+1}$

	(5)	(6)	(7)	(8)
Real interest rate	0.733	0.543	0.385	0.327
	(0.284)	(0.353)	(0.223)	(0.250)
Income growth	0.892	0.960	0.955	0.981
	(0.072)	(0.082)	(0.065)	(0.068)
Family size growth		0.116		0.211
		0.922		(0.066)
Sargan Criterion	3.76	3.39	4.06	3.80
(*p-value*)	(0.58)	(0.49)	(0.66)	(0.57)
+ instruments	+ 5	+ 5	+ 6	+ 6
$(m - k)$				

Note: This table presents additional specifications for the Two-step GMM estimates. Columns (5) and (6) build on the instrument set from Columns (3) and (4), which include lags of inflation, nominal interest rate, consumption growth, income growth, family size growth, and number of earners, by incorporating the household's head age. Columns (7) and (8) further extend this set by adding the household head's years of education. $NT = 182$ observations for each specification.

Specification (4) continues to use the second set of instruments but adds family size growth as an explanatory variable in the Euler equation. This can be viewed as a robustness check for the third specification in both GMM settings. Adding more controls sustains the conclusions regarding the impact of the real interest rate (EIS) and income growth on consumption growth, though it reduces the magnitude of the EIS. With the Two-Step GMM estimator, the EIS decreases from 0.758 (0.289) in the third specification to 0.664 (0.359) in the fourth, and with the Iterative GMM, it decreases from 0.935 (0.345) to 0.836 (0.401). Likewise, in specification (4), regardless of the chosen GMM estimator, the conclusions about the impact of family size growth remain consistent: it is not statistically significant in any case.

Table 4

Iterative GMM estimates of Euler Equation for Consumption

$\Delta \ln(C_{t+1}) = constant + \sigma \ln(1 + R_{t+1}) + \theta' z_{t+1} + \epsilon_{t+1}$

	(1)	(2)	(3)	(4)
Real interest rate	2.421	2.279	0.935	0.836
	(0.320)	(0.389)	(0.345)	(0.401)
Income growth		0.059	0.966	0.978
		(0.483)	(0.075)	(0.084)
Family size growth				0.042
				(0.119)
Sargan Criterion	3.41	8.45	4.30	3.85
(*p-value*)	(0.33)	(0.01)	(0.36)	(0.27)
+ instruments	+ 2	+ 2	+ 4	+ 4
$(m - k)$				

Note: Standard errors are reported in parentheses under each coefficient. The first two columns use a standard instrument set, which includes lags of inflation, nominal interest rate, consumption growth, and income growth. The third and fourth columns add lags of family size growth and number of earners to this set. Both consumption (non-durable) and income growth are adjusted for inflation using the INPC with 2018 as the base year. Interest rates are based on the Mexican Treasury Bill rate, Cetes. The analysis excludes households from the lowest income quintile, covering the 1996-2022 ENIGH surveys. $NT = 182$ observations for each specification.

Specifications (5) and (6) extend specifications (3) and (4) by including the age of the household head as an additional instrument. The inclusion of this instrument has minimal impact on the coefficients for interest rates and income growth in both the Two-Step GMM and Iterative GMM settings, as seen when comparing specification (3) with (5), and specification (4) with (6). The latter comparison shows that while there are no major shifts prior to the addition of an extra instrument, adding both more instruments and controls tends to slightly reduce the EIS value. However, these specifications may not necessarily outperform the earlier specifications (3) and (4), as an ANOVA analysis indicates (see Table 8 in the Appendix).

Table 5
Iterative GMM estimates of Euler Equation for Consumption (Continued)
$\Delta \ln(C_{t+1}) = constant + \sigma \ln(1 + R_{t+1}) + \theta' z_{t+1} + \epsilon_{t+1}$

	(5)	(6)	(7)	(8)
Real interest rate	0.992	0.905	0.449	0.354
	(0.331)	(0.399)	(0.218)	(0.237)
Income growth	0.974	0.980	1.183	1.174
	(0.076)	(0.084)	(0.063)	(0.065)
Family size growth		0.034		0.142
		0.120		(0.064)
Sargan Criterion	4.2	3.77	5.09	5.07
(*p-value*)	(0.51)	(0.43)	(0.53)	(0.40)
+ instruments	+ 5	+ 5	+ 6	+ 6
$(m - k)$				

Note: This table presents additional specifications for the Iterative GMM estimates. Columns (5) and (6) build on the instrument set from Columns (3) and (4), which include lags of inflation, nominal interest rate, consumption growth, income growth, family size growth, and number of earners, by incorporating the household's head age. Columns (7) and (8) further extend this set by adding the household head's years of education. $NT = 182$ observations for each specification.

Specifications (7) and (8) expand on the previous specifications by adding the household head's years of education as an instrument, significantly reducing the EIS value. With the Two-Step GMM estimator, the EIS is recorded at 0.385 (0.223) in specification (7) and 0.327 (0.250) in specification (8). Similarly, with the Iterative GMM, the EIS estimates are 0.449 (0.218) and 0.354 (0.237), respectively. These findings suggest that controlling for the household head's years of education may result in an EIS that is not statistically different from zero. Furthermore, unlike the coefficients on family size growth, which were not statistically significant in previous specifications that included this variable in the Euler equation, the coefficient becomes statistically significant in specification (8). This change may be attributed to the inclusion of the household head's years of education, which could improve the prediction of family size growth, as evidenced in Tables 1 and Table 8 in the Appendix.

Figure 6
EIS Confidence Intervals by GMM Estimator

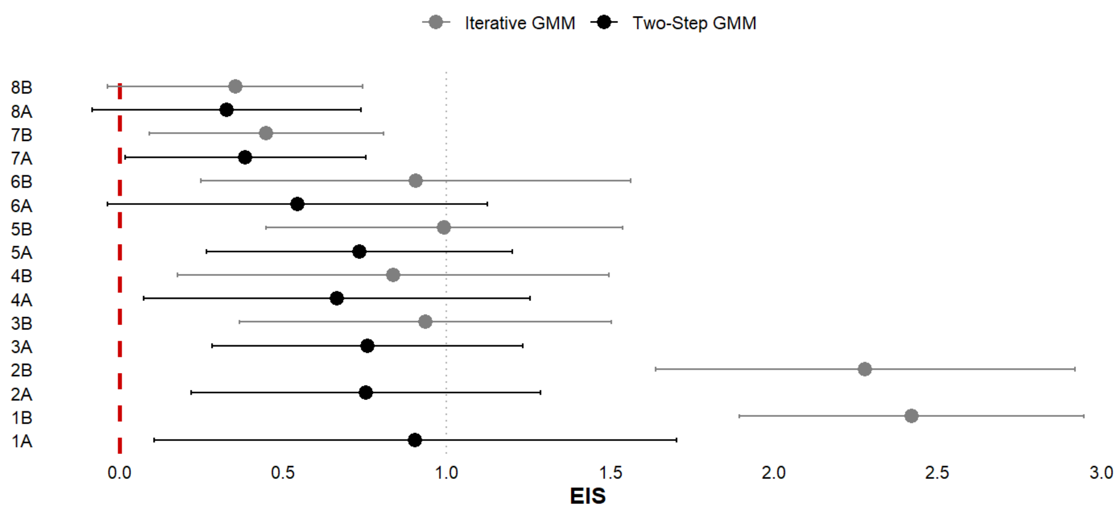

Note: 90% confidence intervals for the EIS are reported. Specifications "A", which consider various controls and instument sets, are presented in Tables 2 and 3. Specifications "B" are detailed in Tables 4 and 5. **Source:** Own elaboration based on data from the ENIGH surveys, 1996-2022.

Despite these findings, caution is warranted. Cohorts generally display similar years of education, ranging from 5-7. Yet, in the initial ENIGH surveys used for estimation (1998 and 2000), all respondents reported having more than this average, which could be due either to a reporting bias in these waves or more educated household heads being overrepresented. It is unclear if the method of questioning changed across ENIGH surveys, which could influence how education data was collected.

In this context, it is reasonable that education might more accurately explain the commonly shared real interest rate than the instrument sets 3 and 4, which do not include it as an instrument (as shown in Table 8 in the Appendix), because it reflects differences across time periods. However, this introduces the potential for a confounder that could simultaneously influence both interest rates and education levels, an aspect not accounted for in the current analysis. When this factor is controlled for, the education of the household head may prove less effective as an instrument for explaining interest rates. This is particularly relevant given that the average education level among the households in the study is relatively uniform.

Therefore, education could be endogenous and might not serve as a valid instrument in this analysis.

Additionally, results from specifications (3)-(6), which exclude education as an instrument, tend to produce more consistent coefficients, regardless of whether an additional instrument is used or a different estimator is applied.

Interestingly, in almost all specifications (except for the second one, where income growth needs to be instrumented by additional household-specific variables beyond household income growth), the coefficient on income growth is very stable and statistically significant. Wald tests for specifications (3)-(8) confirm that it is not statistically different from 1 or less, as expected, across specifications and between estimators (see Table 9 in the Appendix).

Overall, these results offer insights into household consumption behavior. First, the significant sensitivity of consumption to income growth challenges the Life Cycle and Permanent Income hypotheses, which poses that households do not react substantially to changes in their current income as they tend to smooth consumption over their lifetimes. This finding is also consistent with the observed hump-shaped behavior shared between consumption growth and income growth in Graphs 1 and 2. Furthermore, even when controlling for family size growth, the sensitivity of consumption to income growth remains evident.

It is important to note that all households considered in constructing the synthetic cohort panel are at a stage of life where their consumption behavior is likely stable——they are neither too young, with typically rising consumption, nor too old, where consumption might decline more significantly than income. Thus, the observed high sensitivity of consumption to income could indicate the presence of liquidity constraints, aligning with the average income levels of all cohorts in the sample. These cohorts fall between the 52nd and 66th income percentiles. Despite excluding the lowest quintile before averaging household observations to obtain cohort data, many households may still face credit constraints. Their quarterly real income ranges from 37,435 to 81,166 MXN.

Without additional information to exclude households that may engage in the asset market by lending or borrowing, the analysis is narrowed to consider only households in the highest income quintile before constructing the synthetic cohort panel. Unfortunately, limiting the sample to the wealthiest households (90th percentile) leads to an insufficient number of observations for building representative cohorts. Figure 7 in the Appendix shows the income distribution.

Moreover, repeating estimations with households solely in the top income quintile introduces further challenges in estimating the EIS. This narrowing of the sample reduces variation in the instrument sets, making the EIS more unstable across both specifications and estimators. In addition, some specifications now reject the hypotheses of instrument validity, suggesting that other instruments might be necessary for estimating the EIS under this scenario. However, one conclusion remains consistent: the sensitivity of consumption to income growth is still statistically significant, albeit less than 1.

Second, the inclusion of additional controls and instruments consistently leads to a lower EIS. Nonetheless, specification (3)——which is preferred over specifications (4), (6) and (8) due to potential weak identification issues in explaining family size growth or concerns about the endogeneity of the household head's education——indicates a statistically significant EIS. The confidence intervals for the EIS from this specification, whether using the Two-Step GMM or Iterative GMM, include 1, preventing a definitive conclusion on whether it is higher or lower than unity. Notably, specifications (4) and (6) that control for family size tend to be more imprecise (possibly due to weak instruments for explaining this), displaying higher standard errors, suggesting that the true EIS might well exceed 1.

Indeed, instrumenting for additional household factors generally reduces the EIS value, and the significant coefficient on income may still indicate a predominant income effect over a substitution effect. For instance, estimating the Euler equation with both real interest rates and income growth as explanatory variables, and instrumenting these with the third set of instruments plus a squared term for the household head's age, yields an EIS of 0.749 (with a standard error of 0.201 and a J-statistic p-value of 0.49) using the Iterative GMM estimator, which typically reports higher EIS values. In this context, including more instruments and household-specific variables has the potential to enhance the precision of EIS estimates.

7 Conclusions

In general, the estimation of the Elasticity of Intertemporal Substitution (EIS) appears highly sensitive to both the chosen specification and the GMM estimator used. In finite samples like the one in this study, different GMM estimators can yield different conclusions. However, relatively robust estimates and a J-statistic that does not reject the null hypothesis of instrument validity suggest that the model may be well-specified. The preferred specifications yield statistically significant EIS values, indicating that including more household-specific variables as controls or instruments could refine the precision of these estimates, potentially leading to an EIS value less than one. This would suggest a predominance of the income effect over the substitution effect.

Among all well-specified models and instrument sets, one consistent finding emerges: the significant influence of income growth on consumption growth, which underscores the likelihood of liquidity constraints within the sample, challenging the assumption that the Euler equation holds for these households. This observation also contradicts the theory that households—or at least Mexican—, on average, smooth consumption over their life cycles. Additionally, there may be very impatient households——akin to hand-to-mouth consumers——who prefer (or have no other option than) immediate consumption over saving, even when interest rates rise, consistent with a dominant income effect or an EIS less than one. Unfortunately, the level of individual impatience is not directly observable in this study, as linearizing the Euler equation places the discount factor, among other terms, in the intercept.

Moreover, more sensitive EIS estimates from households in the top income quintile raise questions about whether the Cetes interest rate truly reflects the rates considered by these individuals. In an ideal scenario, distinguishing between households that engage in lending/borrowing and those that do not, and identifying the interest rates they consider, would allow for a more precise estimation of the EIS and help determine if asset holders exhibit a larger EIS than non-asset holders or non-bondholders. Nonetheless, the findings presented should be interpreted as reflective of the average Mexican household, typically situated within the 50th to 60th income percentiles, who likely considers the risk-free interest rate.

For future research, employing the available data to estimate the EIS using the Continuous Updating GMM estimator——recommended by Hansen et al. (1996) for its superior finite sample performance compared to the Two-Step or Iterative GMM——may be beneficial. Furthermore, exploring utility functions that specifically incorporate Epstein and Zin (1989) preferences could yield a higher estimated EIS without imposing a rigid relationship between this parameter and the risk aversion coefficient.

In sum, while the results are not definitive, there is suggestive evidence that the income effect is dominant: consumption is highly responsive to changes in current income, and households may not significantly alter their consumption behaviors in response to changes in interest rates.

8 References

Ai, H. (2010). Information quality and long-run risk: Asset pricing implications. *The Journal of Finance*, *65*(4), 1333–1367.

Alan, S., Atalay, K., & Crossley, T. F. (2019). Euler equation estimation on micro data. *Macroeconomic Dynamics*, *23*(8), 3267–3292.

Arrau, P., & van Wijnbergen, S. (1991). Intertemporal substitution, risk aversion, and private savings in mexico. *WPS 682*.

Ascari, G., Magnusson, L. M., & Mavroeidis, S. (2021). Empirical evidence on the euler equation for consumption in the us. *Journal of Monetary Economics*, *117*, 129–152.

Attanasio, O. P., Banks, J., & Tanner, S. (2002). Asset holding and consumption volatility. *Journal of Political Economy*, *110*(4), 771–792.

Attanasio, O. P., & Low, H. (2004). Estimating euler equations. *Review of Economic Dynamics*, *7*(2), 406–435.

Attanasio, O. P., & Weber, G. (1993). Consumption growth, the interest rate and aggregation. *The Review of Economic Studies*, *60*(3), 631–649.

Attanasio, O. P., & Weber, G. (1995). Is consumption growth consistent with intertemporal optimization? evidence from the consumer expenditure survey. *Journal of Political Economy*, *103*(6), 1121–1157.

Bansal, R., & Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, *59*(4), 1481–1509.

Chamberlain, G. (1984). Chapter 22 panel data. Elsevier.

Davidson, R., & MacKinnon, J. G. (2004). Econometric theory and methods. Oxford University Press.

Drechsler, I., & Yaron, A. (2010). What's Vol Got to Do with It. *The Review of Financial Studies*, *24*(1), 1–45.

Epstein, L. G., & Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, *57*(4), 937–969.

Fuller, W. A. (1977). Some properties of a modification of the limited information estimator. *Econometrica*, *45*(4), 939–953.

Gomes, F. A. R., & Ribeiro, P. F. (2015). Estimating the elasticity of intertemporal substitution taking into account the precautionary savings motive. *Journal of Macroeconomics*, *45*, 108–123.

Guvenen, F. (2006). Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective. *Journal of Monetary Economics*, *53*(7), 1451–1472.

Hall, R. E. (1988). Intertemporal substitution in consumption. *Journal of Political Economy*, *96*(2), 339–357.

Hansen, L. P., Heaton, J., & Yaron, A. (1996). Finite-sample properties of some alternative gmm estimators. *Journal of Business and Economic Statistics*, *14*(3), 262–280.

Hansen, L. P., & Singleton, K. J. (1982). Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica*, *50*(5), 1269–1286.

Havranek, T., Horvath, R., Irsova, Z., & Rusnak, M. (2015). Cross-country heterogeneity in intertemporal substitution. *Journal of International Economics*, *96*(1), 100–18.

Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, *55*(3), 703–708.

Runkle, D. E. (1991). Liquidity constraints and the permanent-income hypothesis: Evidence from panel data. *Journal of Monetary Economics*, *27*(1), 73–98.

Stock, J. H., & Yogo, M. (2005). Testing for weak instruments in linear iv regression. In *Identification and inference for econometric models: Essays in honor of thomas rothenberg* (pp. 80–108). Cambridge University Press.

Thimme, J. (2017). Intertemporal substitution in consumption: A literature review. *Journal of Economic Surveys*, *31*(1), 226–257.

Vissing-Jørgensen, A. (2002). Limited asset market participation and the elasticity of intertemporal substitution. *Journal of Political Economy*, *110*(4), 825–853.

Vissing-Jørgensen, A., & Attanasio, O. P. (2003). Stock-market participation, intertemporal substitution, and risk-aversion. *The American Economic Review*, *93*(2), 383–391.

Yogo, M. (2004). Estimating the elasticity of intertemporal substitution when instruments are weak. *The Review of Economics and Statistics*, *86*(3), 797–810.

Zeldes, S. P. (1989). Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy*, *97*(2), 305–346.

9 Appendix

Table 6
Cohort Definition

Cohort	Year of Birth	Average Cell Size
1	1962	739
2	1963	703
3	1964	765
4	1965	749
5	1966	794
6	1967	719
7	1968	811
8	1969	738
9	1970	826
10	1971	728
11	1972	801
12	1973	756
13	1974	747
14	1975	701

Note: Data from 1996-2022 ENIGH surveys, excluding the lowest income quintile.

Table 7
Stock and Yogo Critical Values for the Weak Instrument Test Based on TSLS Bias, 5% significance level

Number of Instruments	5%	10%	20%	30%
4	16.85	10.27	6.71	5.34
5	18.37	10.83	6.77	5.25
6	19.28	11.12	6.76	5.15
7	19.86	11.29	6.73	5.07
8	20.25	11.39	6.69	4.99

Note: Taken from Stock and Yogo (2005). These critical values are for testing the null hypothesis that the bias of TSLS as a fraction of OLS bias is greater than the value in the columns for each number of instruments. For instance, to ensure that TSLS relative bias is not greater than 10% when 4 instruments are considered, the F-statistic in Table 1 first row must be greater than 10.27.

Table 8
First-Stage Regression Comparative Analysis of Different Instrument Sets Across Endogenous Variables in Euler Equation

Comparison	Real Interest Rate	Income Growth	Family Size Growth
1 *vs* 2	4.322	7.062	—
	(0.014)	(0.001)	
2 *vs* 3	1.469	0.003	0.327
	(0.227)	(0.952)	(0.568)
2 *vs* 4	25.093	1.012	17.05
	(0.000)	(0.365)	(0.000)
3 *vs* 4	48.318	2.021	33.711
	(0.000)	(0.156)	(0.000)

Note: This table presents F-statistics and corresponding p-values (shown in parentheses) from ANOVA comparisons of nested models in first-stage regressions. Each model comparison assesses the inclusion of additional instruments in explaining the endogenous variable. Definitions of instrument sets are provided in Table (1). All regressions adjust consumption (non-durable) and income growth for inflation using INPC, base year 2018. Interest rates are based on Cetes rates. Data from 1996-2022 ENIGH surveys, excluding the lowest income quintile. $NT = 182$ observations for each regression.

Table 9
Wald Tests for Specifications (3)-(8) Using Two-Step GMM and Iterative GMM Estimators:
Testing the Null Hypothesis That Variables Are Equal to 1

	Interest Rate	Income Growth	Interest Rate & Income Growth
(3) - Two-Step GMM	0.699	2.33	5.139
	(0.403)	(0.126)	(0.076)
(3) - Iterative GMM	0.035	0.209	0.426
	(0.850)	(0.647)	(0.807)
(4) - Two-Step GMM	0.868	0.548	3.297
	(0.351)	(0.459)	(0.192)
(4) - Iterative GMM	0.168	0.0628	0.550
	(0.681)	(0.801)	(0.759)
(5) - Two-Step GMM	0.880	2.225	5.389
	(0.348)	(0.135)	(0.067)
(5) - Iterative GMM	0.000	0.114	0.158
	(0.982)	(0.735)	(0.924)
(6) - Two-Step GMM	1.668	0.226	3.797
	(0.196)	(0.634)	(0.149)
(6) - Iterative GMM	0.056	0.052	0.271
	(0.812)	(0.818)	(0.873)
(7) - Two-Step GMM	7.57	0.450	9.367
	(0.005)	(0.501)	(0.009)
(7) - Iterative GMM	6.382	8.31	15.20
	(0.011)	(0.003)	(0.000)
(8) - Two-Step GMM	7.209	0.073	8.03
	(0.007)	(0.785)	(0.018)
(8) - Iterative GMM	7.375	7.107	13.439
	(0.006)	(0.007)	(0.001)

Note: Wald tests for the null hypothesis that the column variable is equal to 1. p-values are in parentheses under the Wald statistics. Specifications (3)-(8) are detailed in Tables 2-5.

Figure 7
Log-transformed Income Distribution by Cohort

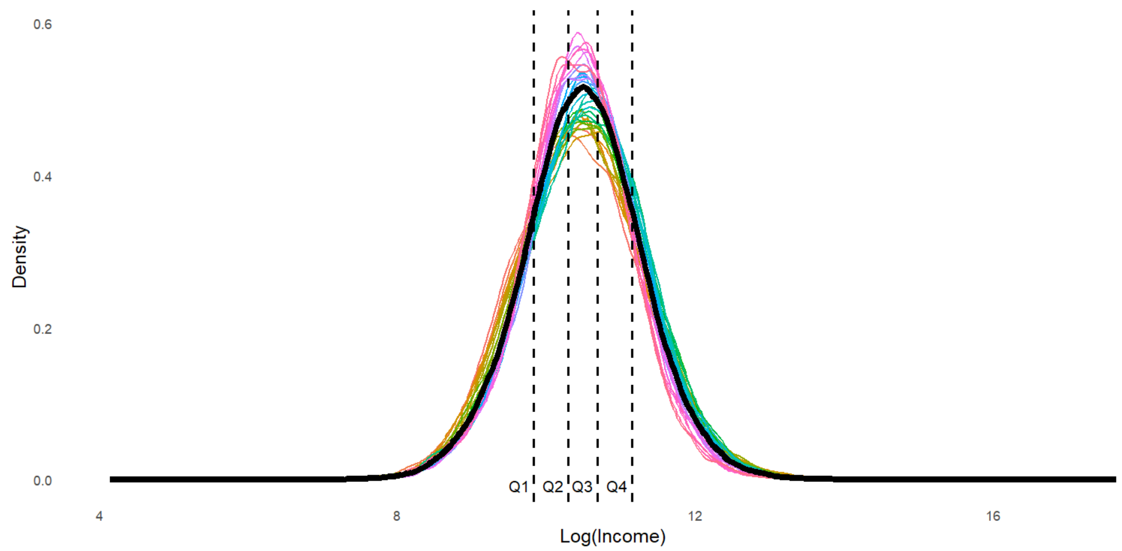

Note: Average log-transformed income distribution is displayed in black. Dashed lines represent each income quintile. **Source:** Own elaboration based on data from the ENIGH surveys, 1996-2022.