

Cross-Country Heterogeneity in Intertemporal Substitution*

Tomas Havranek^{a,b}, Roman Horvath^{b,c}, Zuzana Irsova^b, and Marek Rusnak^{a,b}

^aCzech National Bank

^bCharles University, Prague

^cIOS, Regensburg

October 14, 2013

Abstract

We collect 2,735 estimates of the elasticity of intertemporal substitution in consumption from 169 published studies that cover 104 countries during different time periods. The estimates vary substantially from country to country, even after controlling for 30 aspects of study design. Our results suggest that income and asset market participation are the most effective factors in explaining the heterogeneity: households in rich countries and countries with high stock market participation substitute a larger fraction of consumption intertemporally in response to changes in expected asset returns. Micro-level studies that focus on sub-samples of rich households or asset holders also find systematically larger values of the elasticity.

Keywords: Elasticity of intertemporal substitution, consumption, meta-analysis, Bayesian model averaging

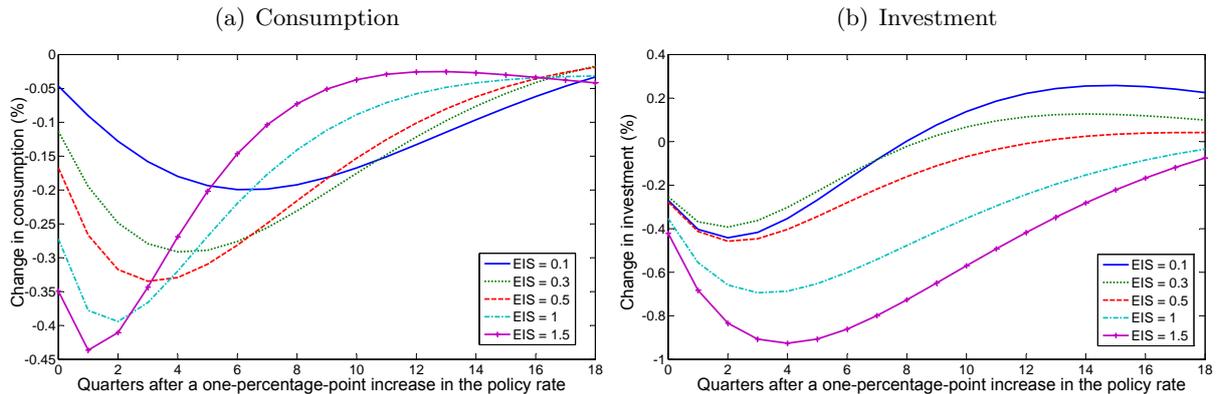
JEL Codes: C83, D91, E21

*An online appendix with data, code, and a list of studies included in the meta-analysis is available at meta-analysis.cz/substitution. Corresponding author: Tomas Havranek, tomas.havranek@ies-prague.org. Havranek, Irsova, and Rusnak acknowledge support from the Grant Agency of Charles University in Prague (grant #554213). We thank Jan Babecky, Michal Bauer, Iftekhar Hasan, Jiri Schwarz, and seminar participants at CERGE-EI and the Czech National Bank for their helpful comments. The views expressed here are ours and not necessarily those of our employers.

1 Introduction

The elasticity of intertemporal substitution in consumption (EIS) reflects households' willingness to substitute consumption between time periods in response to changes in the expected real interest rate. Therefore it represents a crucial parameter for a wide range of economic models involving intertemporal choice, from modeling the behavior of aggregate savings and the impact of fiscal policy to computing the social cost of carbon emissions, and has been estimated by hundreds of researchers. Figure 1 illustrates how the elasticity matters for the modeled effects of monetary policy: we use the popular model of Smets & Wouters (2007), vary the calibrated value of the EIS, and for different values of the EIS plot the impulse responses of consumption and investment to a one-percentage-point monetary policy shock. It is apparent that the modeled development of these aggregates depends strongly on the value of the elasticity of intertemporal substitution.

Figure 1: The elasticity of intertemporal substitution matters



Notes: The figure shows simulated impulse responses to a one-percentage-point increase in the monetary policy rate. We use the popular model developed by Smets & Wouters (2007) and vary the value of the elasticity of intertemporal substitution while leaving all other parameters calibrated at the posterior values from Smets & Wouters (2007). For the simulations we use Matlab code from The Macroeconomic Model Data Base (Wieland *et al.*, 2012).

The figure shows impulse responses for the EIS calibrated between 0.1 and 1.5, and in the literature we indeed encounter such large differences in calibrations of the elasticity. The most cited empirical study estimating the elasticity, Hall (1988), who concludes that the EIS is not likely to be larger than 0.1, has influenced many researchers. Some studies use a value of 0.2 (Chari *et al.*, 2002; House & Shapiro, 2006; Piazzesi *et al.*, 2007), or a value of 0.5 (Jin, 2012; Trabandt & Uhlig, 2011; Rudebusch & Swanson, 2012), or a value of 2 (Ai, 2010; Barro, 2009;

Colacito & Croce, 2011), to name but a few recent examples of different calibrations. The reason for the different calibrations is differences in the results of empirical studies on the EIS. For example, the standard deviation of the estimates reported by the 33 studies in our sample which were published in the top five general interest journals is 1.4, outliers excluded. Most commentators would agree with Ai (2010, p. 1357), who starts his discussion of calibration by noting that “empirical evidence on the magnitude of the EIS parameter is mixed.”

In this paper we collect 2,735 estimates of the elasticity of intertemporal substitution reported in 169 studies and review the literature quantitatively using meta-analysis methods. Meta-analysis, which has been employed in economics by Card & Krueger (1995), Ashenfelter *et al.* (1999), Stanley (2001), Disdier & Head (2008), and Chetty *et al.* (2011), among others, allows us to examine systematically the influence of methodology on the results. In this framework we can address the challenge put forward by an early survey of the empirical evidence from consumption Euler equations (Browning & Lusardi, 1996, p. 1833): “It is frustrating in the extreme that we have very little idea of what gives rise to the different findings. (...) We still await a study which traces all of the sources of differences in conclusions to sample period; sample selection; functional form; variable definition; demographic controls; econometric technique; stochastic specification; instrument definition; etc.”

While controlling for differences in methodology, we focus on explaining country-level heterogeneity. The studies in our sample provide us with estimates of the EIS for 104 countries, and we show that the mean values reported for the countries vary substantially. We build on the literature that explores the heterogeneity in the EIS at the micro level. For example, Blundell *et al.* (1994) and Attanasio & Browning (1995) suggest that rich households tend to show a larger elasticity of intertemporal substitution, and we examine whether GDP per capita is associated with the mean EIS reported for the country. Mankiw & Zeldes (1991) and Vissing-Jorgensen (2002) find a larger elasticity for stockholders than for non-stockholders, and we explore the relationship between stock market participation and the elasticity of intertemporal substitution at the country level. Bayoumi (1993) and Wirjanto (1995), among others, indicate that liquidity-constrained households show a smaller EIS, and we examine whether ease of access to credit helps explain the cross-country variation in the elasticity. More details on factors potentially causing heterogeneity in the EIS are available in Section 3.

The mean estimate of the elasticity of intertemporal substitution reported in empirical studies is 0.5, but we show that cross-country differences are important. Since it is often unclear which aspects of methodology should matter for the magnitude of the estimated EIS, we include all 30 that we collect and employ Bayesian model averaging (Raftery *et al.*, 1997) to deal with the resulting model uncertainty. Our findings suggest that a larger EIS is associated with higher per capita income of the country, and especially with higher stock market participation. According to our baseline model, a 10-percentage-point increase in the rate of stock market participation is associated with an increase in the EIS of 0.24. Moreover, wealth and asset market participation are also important at the micro level: studies estimating the EIS using a sub-sample of rich households or asset holders find on average an EIS larger by 0.21.

The remainder of the paper is structured as follows. Section 2 explains how we collect data from studies estimating the elasticity. Section 3 discusses the reasons for including variables that may explain the differences in the reported estimates of the EIS. Section 4 describes the results, while Section 5 provides robustness checks. Appendix A lists mean values of the EIS reported for various countries and summary statistics of all variables used in our analysis. Appendix B provides diagnostics on Bayesian model averaging. An online appendix with data, code, and a list of studies included in the meta-analysis is available at meta-analysis.cz/substitution.

2 Estimates of the Elasticity

To estimate the EIS, researchers often follow Hall (1988) and use the log-linearized consumption Euler equation. That is, they regress consumption growth on the intertemporal price of consumption, the real rate of return:

$$\Delta c_{t+1} = \alpha_i + EIS \cdot r_{i,t+1} + \epsilon_{i,t+1}. \quad (1)$$

Here Δc_{t+1} denotes consumption growth at time $t + 1$, $r_{i,t+1}$ denotes the real return on asset i at time $t + 1$ (for instance the stock market return or treasury bill return), and $\epsilon_{i,t+1}$ denotes the error term. The error term is correlated with $r_{i,t+1}$, and researchers thus use instruments for $r_{i,t+1}$, typically including the values of asset returns and consumption growth known at time t . There are of course many potential modifications to (1), many ways in which it can be

estimated, and many different data that can be used in the estimation; we discuss these issues in detail in Section 3 and control for the context in which researchers obtain their estimates.

The first and crucial step of meta-analysis is the selection of studies that are included. We start with an extensive search in Google Scholar (the search query and the list of studies are available in the online appendix). There are thousands of papers on the topic, so a good search query is needed to identify studies that are likely to contain empirical estimates of the EIS. We adjust our query until it includes most of the well-known empirical papers among the top 50 hits. For the selection of studies we prefer Google Scholar to other databases commonly used in meta-analysis, such as EconLit or Scopus, because Google Scholar provides powerful fulltext search.

The search yields about 1,500 hits in total, but on closer examination we find that papers identified in the bottom half of the search list are unlikely to contain usable empirical estimates of the EIS. We read the abstracts of the first 700 papers to see which can be included in the meta-analysis, and it seems that more than 300 studies contain usable estimates of the EIS. At this point it is clear that to capture the context in which researchers obtain the estimates we have to collect about 30 variables reflecting methodology. Since a typical study (especially a typical working paper) reports many different estimates (using different sets of instrumental variables, for example), we find it unfeasible to include all studies and decide to focus on published studies only and read these studies in detail. An alternative solution is to select just one representative estimate from each study, published or unpublished, and discard the other estimates, but often it is unclear what the preferred estimate would be. We stop the search on January 1, 2013 and identify 169 published studies that provide estimates of the EIS and detailed information on methodology.

Aside from saving us several months of work, the restriction of the sample to published studies has two additional benefits. First, publication status is a simple indicator of quality because published studies are peer-reviewed. Second, published papers are typically better written and typeset, which makes the collection of data easier and reduces the danger of mistakes. But even when we focus solely on published papers, we have to collect about 80,000 data points by hand (the published literature provides 2,735 estimates of the EIS and for each we collect 30 aspects of methodology). Two of the co-authors, therefore, collect the data simultaneously and

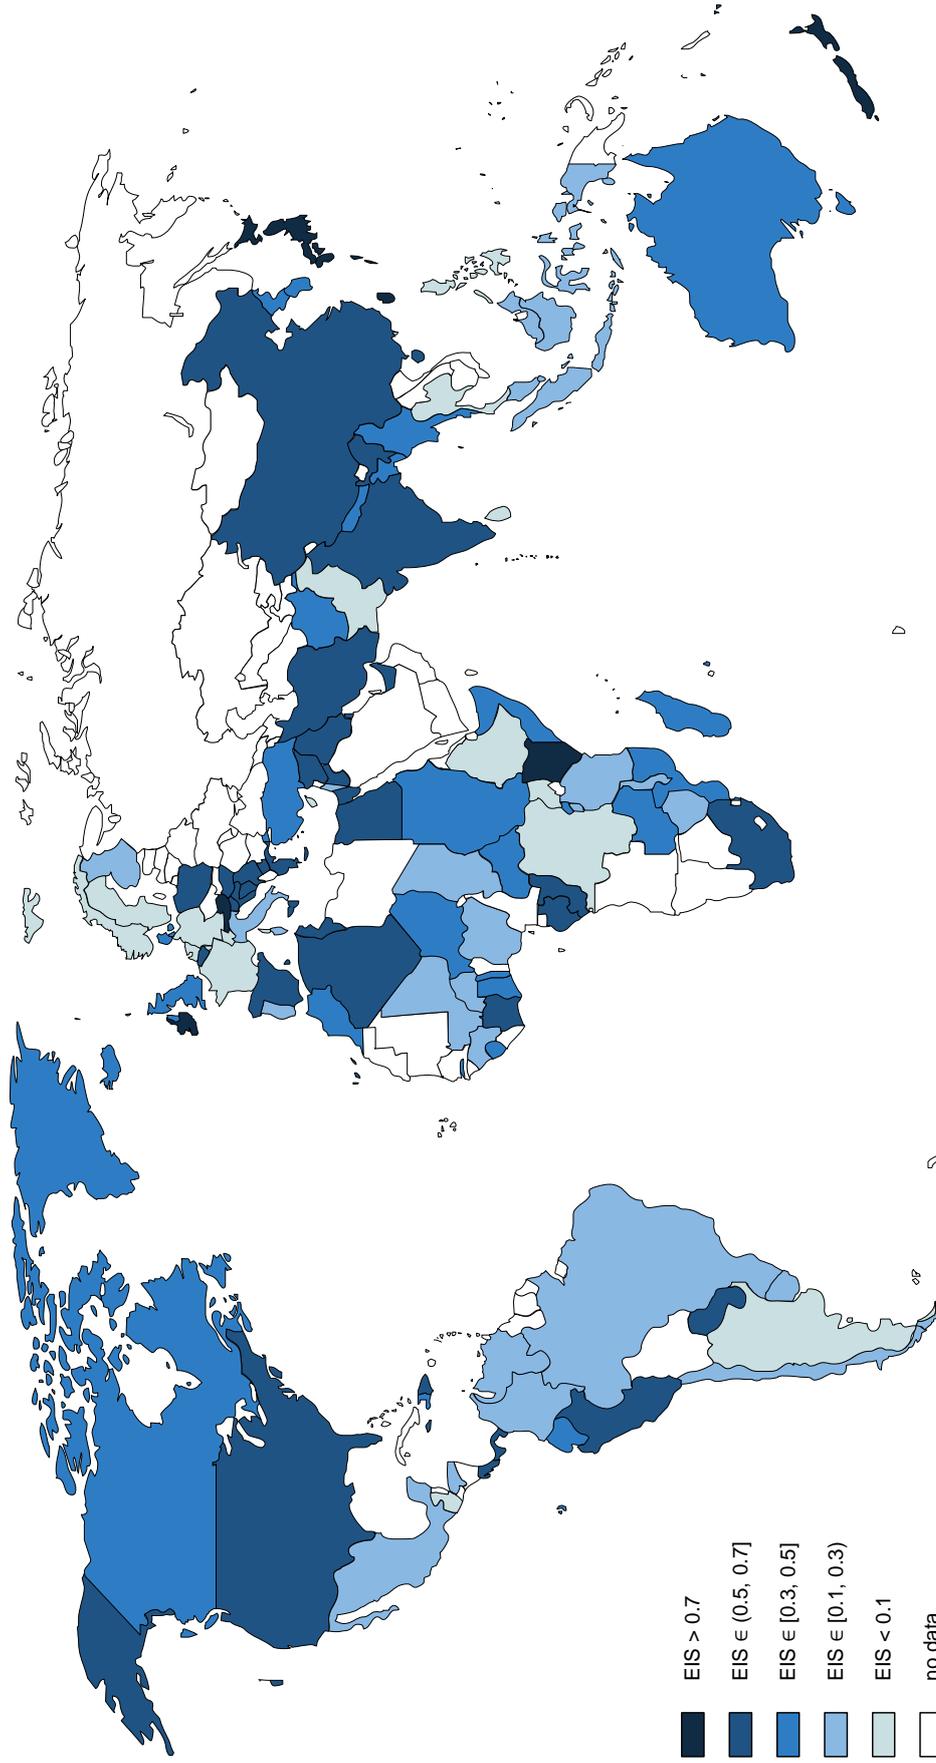
check the resulting data set for errors. The final database used in the paper is available in the online appendix. Judging from the surveys of meta-analyses by Nelson & Kennedy (2009) and Doucouliagos & Stanley (2013) we believe this paper is the largest meta-analysis conducted in economics so far.

Out of the 169 studies included in the meta-analysis, 33 are published in the top five journals in economics, which underlines the importance of the EIS and the amount of research dedicated to its estimation. All studies combined receive on average more than two thousand citations per year in Google Scholar, which indicates that the estimates are heavily used. Our sample includes studies published over three decades: from 1981 to 2012; the median study uses data from 1970 to 1994 and provides 8 estimates of the elasticity. The estimates span 104 different countries, even though about half of all estimates are computed for the US. The mean reported estimate of the EIS is 0.5—for this and all other computations we exclude estimates that are larger than 10 in absolute value (2.5% of the data). Such large estimates seem implausible, but the threshold is arbitrary. In Section 5 we explain that the choice of threshold does not affect our results much. Finally, when each study is given the same weight (as opposed to each estimate being given the same weight), the mean EIS is 0.7. This is close to, for example, the baseline calibration of $2/3$ used by Smets & Wouters (2007).

But the worldwide mean represents a poor guide for the calibration of the EIS in most countries, as Figure 2 illustrates (numerical values for the countries are provided in Table A1 in the Appendix). The estimated EIS differs a lot across countries, typically lying between 0 and 1. Such heterogeneity can make a big difference to the modeled effectiveness of monetary policy, among other things, as we showed in Figure 1. For some countries only a handful of estimates are available, so some of the country averages we report may be quite imprecise and influenced by the estimation method. Nevertheless, for six countries we have more than 50 estimates (the least covered of these countries is Sweden, with 63 estimates reported in 11 studies). Among these countries we find the largest EIS for Japan (0.9), followed by the US (0.6), the UK (0.5), Canada (0.4), Israel (0.2), and Sweden (0.1). The cross-country heterogeneity in the estimated EIS is substantial and calls for an explanation.

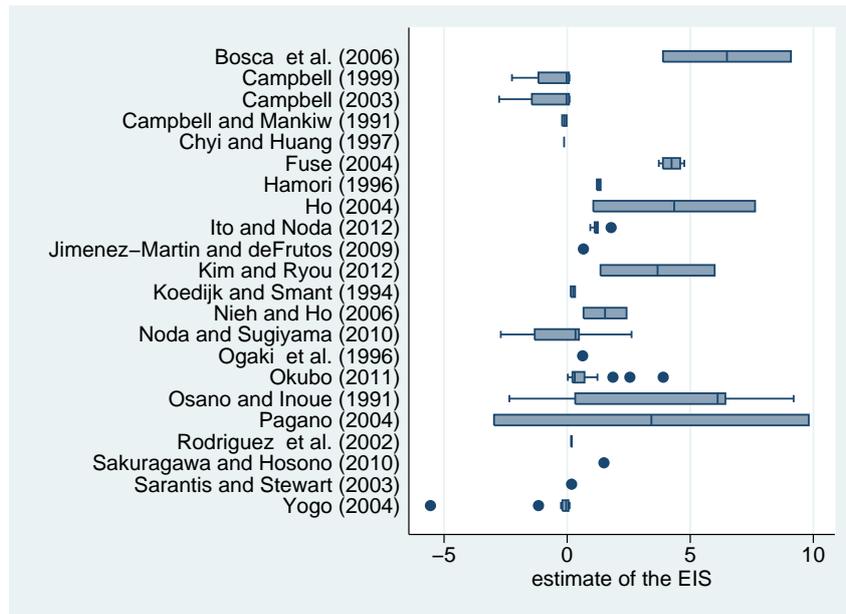
When looking for the sources of cross-country heterogeneity, however, it is also important to take into account that researchers employ different methods to estimate the EIS. Figure 3 shows

Figure 2: Country heterogeneity in the elasticity of intertemporal substitution



Notes: For each country the figure depicts the mean estimate of the EIS reported in the literature; numerical values are provided in Table A1. Estimates larger than 10 in absolute value are excluded.

Figure 3: Method heterogeneity in the elasticity of intertemporal substitution for Japan



Notes: The figure is a box plot of estimates of the EIS corresponding to Japan that are reported in the studies in our sample. Estimates larger than 10 in absolute value are excluded.

how the reported EIS differs across studies even if it is estimated for the same country. For illustration we select Japan, which is the third most often examined country in the literature (after the US and the UK). Dozens of studies estimate the elasticity for the US and the UK and it would be difficult to squeeze them into a box plot, but the conclusion would be the same even for these countries. We see that individual studies report very different estimates and often the within-study distributions of the estimates do not overlap. Therefore, in all the estimations we also control for the methodology employed by the researchers.

3 Why Do the Estimates Differ?

We consider five country characteristics that may influence the reported magnitude of the EIS:

Income Most studies examining heterogeneity in the EIS focus on the role of income. The hypothesis states that poor consumers substitute less consumption intertemporally because their consumption bundle contains a larger share of necessities, which are more difficult to substitute between time periods compared with luxury goods. Moreover, if subsistence requirements represent an important portion of the poor's consumption, the poor have limited discretion for

intertemporal substitution in consumption. This hypothesis has been supported by analyses of micro data (for example, Blundell *et al.*, 1994; Attanasio & Browning, 1995), as well as cross-country data (Atkeson & Ogaki, 1996; Ogaki *et al.*, 1996). We use GDP per capita to capture the differences in income across countries.

Asset market participation We expect households participating in asset markets to be more willing to substitute consumption intertemporally. Exposure to the stock market, for example, may be correlated with households' awareness of the payoffs from intertemporal substitution and, in general, with the forward-looking nature of their consumption. Moreover, Attanasio *et al.* (2002) and Vissing-Jorgensen (2002) argue that consumption Euler equations are not valid for households not participating in the corresponding asset market, and find larger estimates of the EIS for stockholders and bondholders compared with households that do not own these assets. Similarly, Mankiw & Zeldes (1991) find a larger EIS for stockholders than for other households. To capture this country characteristic we use the database of stock market participation developed by Giannetti & Koskinen (2010).

Liquidity constraints Liquidity-constrained households have less opportunities for intertemporal substitution in consumption (Wirjanto, 1995). The resulting consumption of liquidity-constrained households may be linked to income, as it is for the rule-of-thumb consumers of Campbell & Mankiw (1989), and lacks the forward-looking element of the response to the expected real rate of return. Bayoumi (1993), for example, finds that financial deregulation in the UK brought a substantial increase in the proportion of households with a positive EIS. Attanasio (1995) provides a survey of the literature on the effects of liquidity constraints on intertemporal consumption choice. To capture liquidity constraints we use two alternative measures: credit availability defined as the ease of access to loans and reported by the Global Competitiveness Report, and a measure of financial reform reported by the IMF (Abiad *et al.*, 2010).

Asset return Almost all estimations and applications of the EIS assume the elasticity to be constant with respect to the rate of return of the asset in question. In a recent paper, however, Crossley & Low (2011) reject the hypothesis of a constant EIS. To see whether the estimated EIS differs systematically for countries with different returns, we include a measure of the real interest rate defined as the lending rate adjusted for inflation as measured by the GDP deflator.

Culture and institutions The willingness of households to substitute consumption into an uncertain future may be associated with culture and institutions. For example, Porta *et al.* (1998) suggest that institutions have an important influence on financial decisions. It has also been found that trust, or social capital more generally, is an important factor for stock market participation and financial development (Guiso *et al.*, 2004, 2008). Moreover, a large cross-country survey on time discounting and risk preferences (Wang *et al.*, 2011; Rieger *et al.*, 2011) shows the importance of cultural differences. To capture the economic culture of the country we use two measures: the rule of law index (taken from the World Bank Global Governance Indicators), which captures the extent to which people have confidence in the rules of society, and the index of generalized trust in society (Bjoernskov & Meon, 2013).

A detailed description and summary statistics for each variable used in our analysis are reported in Table A2 in the Appendix. A few difficult issues of data collection are worth discussing at this point. First, some variables are not available for all 104 countries in our data set. Data on stock market participation are available for only 28 countries, which we call “core countries” in the analysis, and we also conduct a separate set of regressions without the variable on stock market participation (and, therefore, using almost all countries in the data set). Second, a few estimates of the EIS use data from several countries; for example, the euro area. We keep such estimates in the data set and compute average values of the corresponding country-level characteristics. Third, different studies use data from different time periods to estimate the EIS. Whenever possible, we compute the average of the country characteristic corresponding to the data period. For example, if a study uses data from 1980 to 1994, we use the average value of the real interest rate of that period. This adjustment significantly increases the variation in country-level variables.

We also consider 30 variables reflecting the different aspects of methodology used to estimate the EIS. For ease of exposition we divide these method choices into variables reflecting the definition of the utility function (5 aspects), data characteristics (6 aspects), general design of the analysis (7 aspects), the definition of main variables (4 aspects), estimation characteristics (4 aspects), and publication characteristics (4 aspects).

Utility function An important feature of studies estimating the EIS is whether the elasticity is separated from the coefficient of relative risk aversion. Only about 5% of all the estimates in our sample estimate the parameters separately, usually employing the utility function put forward by Epstein & Zin (1989). Habits in consumption are assumed by 4% of researchers. Some studies assume non-separability between durables and non-durables (4% of estimates), following Ogaki & Reinhart (1998), who argue that assuming separability can produce a downward bias in the estimate of the elasticity. A similar fraction of studies allow for non-separability between private and public consumption, while 5% of studies allow for non-separability between tradable and non-tradable goods.

Data The studies differ greatly in the number of cross-sectional units (usually households or countries) used in the estimation and in the length of the time span of the data. We also include a variable reflecting the average year of the data period to see whether there is a trend in the estimated EIS over time. We include a dummy variable for studies using micro data (about 20% of our data set). Many authors (for example, Attanasio & Weber, 1993) argue that estimating Euler equations on macro data can lead to biased results because of the omission of demographic factors. Moreover, we include dummy variables reflecting the frequency of the data used for the estimation. Most studies use quarterly data (57%); some employ monthly data (10%). Annual data are typically used by micro studies.

Design We include a dummy variable for studies using synthetic cohort data (about 5% of our data set). Most authors assume a time-additive utility function, which results in the EIS being equal to the inverse of the coefficient of relative risk aversion. Some studies focusing on risk preferences regress asset returns on consumption growth and report the inverse of the EIS (almost a third of all the studies in our data set). Nevertheless, Campbell (1999) notes that using the asset return as the response variable may aggravate the problem of weak instruments in estimating the parameter. To see whether this method choice has a systematic effect on the results, we include a dummy variable called *Inverse estimation*.

As we noted earlier, some micro studies on the EIS explore potential heterogeneity in the parameter; they typically estimate the elasticity for different subsets of households. The definition of subsets differs, but researchers usually ask whether richer households or households

participating in asset markets show a larger elasticity of intertemporal substitution. To capture this effect we include a dummy variable *Asset holders*. Next, Campbell & Mankiw (1989), among others, show that because of the time aggregation of consumption the instrument set for asset returns should not contain first lags of variables. But still about 30% of all the estimates are computed using first lags of variables among the instruments.

Gruber (2006) stresses that studies using micro data should include year fixed effects for the identification to come from cross-sectional variation and not from time series variation correlated with consumption. Nevertheless, 3% of the studies in our data set use data from the Panel Study of Income Dynamics but do not include year fixed effects. About a quarter of the studies include income in the estimation to test for excess sensitivity of consumption to current income, and we control for this aspect of methodology as well. We also include the number of demographic controls used in micro studies to explain household-level variation in consumption.

Variable definition Most studies use non-durable consumption as the response variable, but some 20% of the estimates are computed using total consumption. About 6% of studies use food as a proxy for consumption, which according to Attanasio & Weber (1995) can produce biased estimates if food is not separable from other types of consumption. The asset return is typically defined as the interest rate on treasury bills, but almost 20% of studies use the stock market return. Mulligan (2002), however, explains that the rate of return should be measured as the return on a representative unit of capital, and we include a dummy variable for this aspect of methodology.

Estimation We have noted that the log-linearized consumption Euler equation is the favorite framework for estimation of the EIS. But Carroll (2001), for example, criticizes the common practice on the grounds that higher-order terms may be endogenous to omitted variables in the regression resulting from the log-linear Euler equation. Thus we include a dummy variable for studies using the exact Euler equation to see whether log-linearization affects the estimates of the elasticity in a systematic way. Next, the regression parameters are typically estimated using GMM, but a third of studies use two-stage least squares, and 10% of studies disregard endogeneity and employ OLS.

Publication characteristics Some novel methods are employed by only a few studies and their influence on the results cannot be examined in a meaningful way using meta-analysis. For this reason we also include variables reflecting the quality of studies not captured by the method variables introduced above. We include publication year to capture innovations in methodology, the number of citations of the study in Google Scholar, the recursive RePEc impact factor of the journal, and a dummy variable for studies published in the top five general interest journals in economics. The data on citations and impact factors were collected on January 31, 2013.

4 Meta-Regression Analysis

Our intention is to explore whether the country characteristics described in the previous section are associated with the reported EIS, but also to control for the type of methodology used in the studies. That is, we employ the following “meta-regression”:

$$EIS_k = a + \beta \cdot \text{Country variables}_k + \gamma \cdot \text{Method variables}_k + \theta_k. \quad (2)$$

The problem is that there are 30 method variables and it is not clear which ones should be included. We cannot include all of them in an OLS regression because the specification would contain many redundant variables. Some meta-analysts use sequential t -tests to exclude the least significant variables, but such an approach is not statistically valid. In this paper we opt for a technique designed to tackle such regression model uncertainty: Bayesian model averaging (BMA). BMA runs many regressions with different subsets of the explanatory variables on the right-hand side and then constructs a weighted average over these regressions (aside from a robustness check, we always include the country-level variables in all BMA regressions). For applications of BMA in economics, see, for instance, Fernandez *et al.* (2001); Ciccone & Jarocinski (2010); Moral-Benito (2012). Because model uncertainty is inevitable in meta-analysis (it is usually unclear whether some aspects of methodology could influence the results in a systematic way, and the potential aspects are many), BMA has also been frequently used in this field (Moeltner & Woodward, 2009; Irsova & Havranek, 2013; Havranek & Rusnak, 2013).

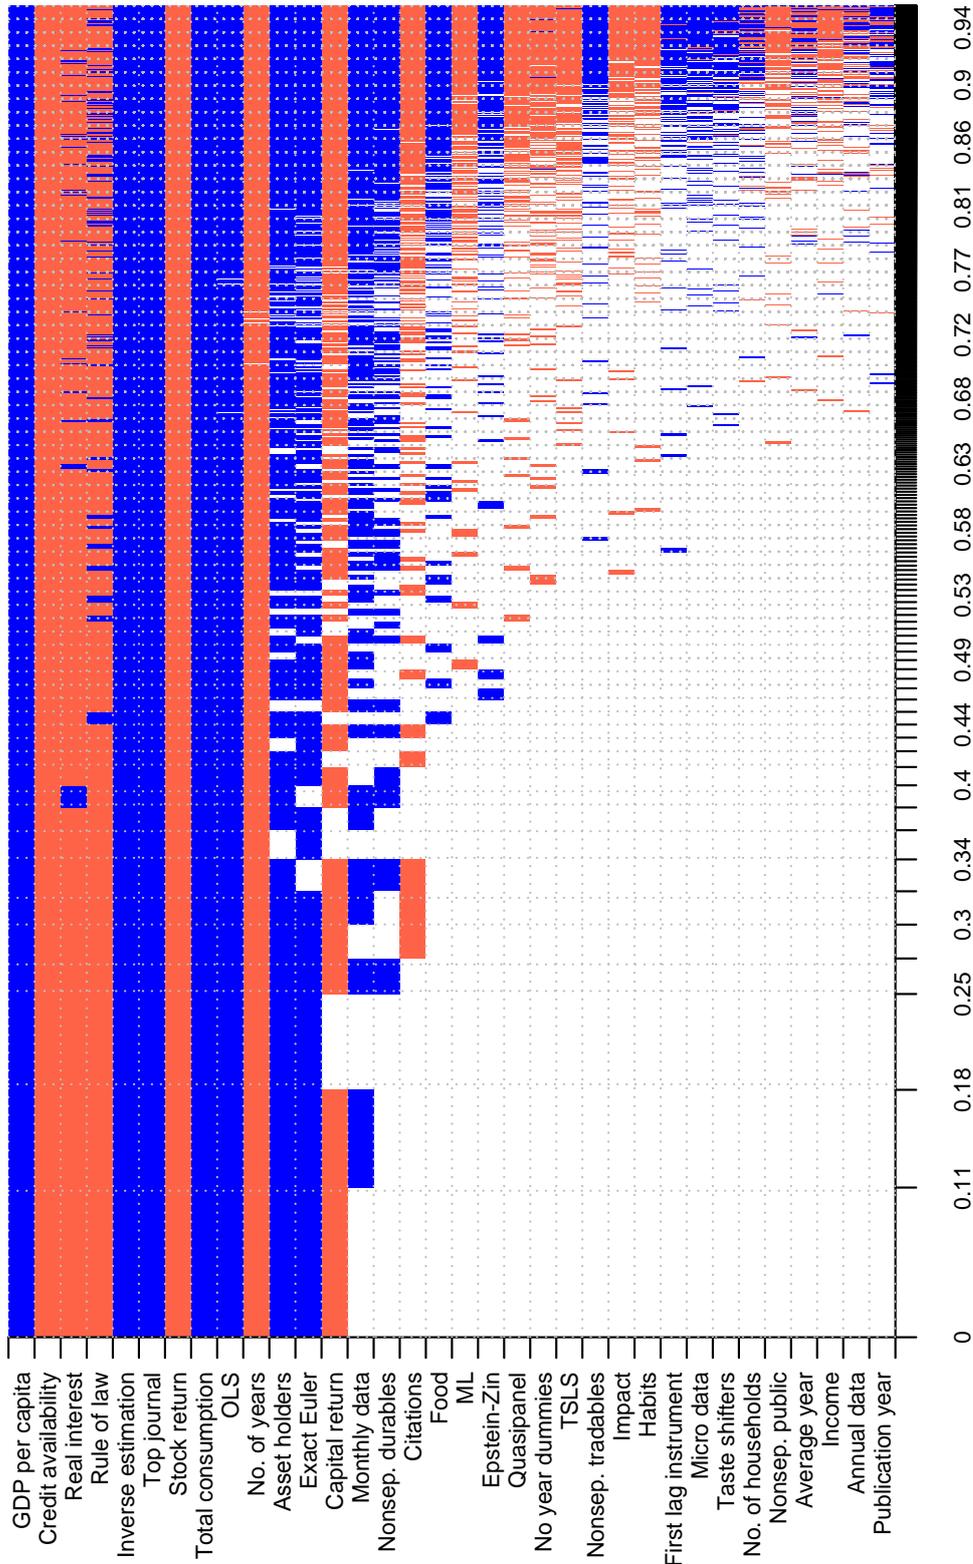
Bayesian model averaging is described in detail by Feldkircher & Zeugner (2009), for instance, and here we only give intuition for the technical terms needed for the evaluation of the

results. The weights used in the BMA estimation are called *posterior model probabilities* and capture how well individual regressions fit the data—thus the weights are analogous to adjusted R-squared or information criteria used in frequentist econometrics. For each variable the sum of the posterior probabilities of models in which the variable is included indicates the so-called *posterior inclusion probability*, which is analogous to statistical significance. If the posterior inclusion probability of a variable is close to one, almost all models that are effective in explaining the variance in the reported EIS include that variable. BMA provides us with a large number of regressions, and from these we can compute for each variable the *posterior coefficient distribution*. The posterior coefficient distribution gives us the posterior mean (analogous to the estimate of a regression coefficient) and posterior standard deviation (analogous to the standard error of an estimated regression parameter).

Because we have 30 method variables, there are 2^{30} potential regressions with different combinations of the method variables. To compute all these regressions would take several weeks, so we opt for the Metropolis-Hasting algorithm, a Markov chain Monte Carlo method. The Metropolis-Hastings algorithm walks through the most important part of the model mass—the models with high posterior model probabilities. For all BMA estimations we use one million burn-ins and two million iterations to ensure a good degree of convergence. We employ the beta-binomial prior advocated by Ley & Steel (2009): the prior model probabilities are the same for all possible model sizes. We set the Zellner’s g prior following Fernandez *et al.* (2001). These priors are quite conservative and reflect the fact that we know little about the true model size and parameter signs. In the next section, however, we check if our results are robust to a different choice of priors. All of the computations are performed using the R package `bms` available at `bms.zeugner.eu`. Codes for all our estimations are available in the online appendix.

In our first BMA estimation we do not include *stock market participation*, which is available for only 28 countries, and use data for as many countries as possible. The estimation is illustrated in Figure 4. The columns in the figure denote individual models; the variables are sorted by posterior inclusion probability in descending order. A blue cell (darker in grayscale) implies that the variable is included and its estimated sign is positive. A red color (lighter in grayscale) implies that the variable is included and the estimated sign is negative. Blank cells imply that the corresponding variable is not included in the model. Only the 5,000 models with

Figure 4: Model inclusion, all countries



Notes: Response variable: estimate of the elasticity of intertemporal substitution. Columns denote individual models; variables are sorted by posterior inclusion probability in descending order. Blue color (darker in grayscale) = the variable is included and the estimated sign is positive. Red color (lighter in grayscale) = the variable is included and the estimated sign is negative. No color = the variable is not included in the model. The horizontal axis measures cumulative posterior model probabilities. Only the 5,000 models with the highest posterior model probabilities are shown.

Table 1: Explaining the differences in the estimates of the EIS, all countries

Response variable:	Bayesian model averaging			Frequentist check (OLS)		
	Estimate of the EIS	Post. mean	Post. std. dev.	PIP	Coef.	Std. er.
<i>Country characteristics</i>						
GDP per capita	0.134	0.074	1.000	0.126	0.084	0.138
Credit availability	-0.037	0.059	1.000	-0.033	0.055	0.553
Real interest	-0.005	0.007	1.000	-0.003	0.006	0.635
Rule of law	-0.020	0.092	1.000	-0.019	0.074	0.800
<i>Utility</i>						
Epstein-Zin	0.018	0.074	0.069			
Habits	-0.004	0.032	0.021			
Nonsep. durables	0.122	0.199	0.309			
Nonsep. public	-0.001	0.019	0.012			
Nonsep. tradables	0.006	0.043	0.027			
<i>Data</i>						
No. of households	0.000	0.003	0.012			
No. of years	-0.201	0.055	0.982	-0.196	0.048	0.000
Average year	0.015	0.940	0.012			
Micro data	0.002	0.026	0.017			
Annual data	0.000	0.008	0.010			
Monthly data	0.160	0.167	0.531	0.263	0.090	0.004
<i>Design</i>						
Quasipanel	-0.015	0.068	0.059			
Inverse estimation	0.530	0.067	1.000	0.512	0.137	0.000
Asset holders	0.349	0.181	0.849	0.421	0.089	0.000
First lag instrument	0.002	0.015	0.021			
No year dummies	-0.027	0.131	0.054			
Income	0.000	0.008	0.011			
Taste shifters	0.001	0.011	0.015			
<i>Variable definition</i>						
Total consumption	0.373	0.085	0.997	0.379	0.102	0.000
Food	0.051	0.147	0.141			
Stock return	-0.344	0.077	0.999	-0.385	0.163	0.021
Capital return	-0.207	0.148	0.723	-0.288	0.077	0.000
<i>Estimation</i>						
Exact Euler	0.219	0.131	0.792	0.283	0.244	0.250
ML	-0.023	0.084	0.085			
TOLS	-0.006	0.035	0.043			
OLS	0.420	0.111	0.984	0.440	0.119	0.000
<i>Publication</i>						
Publication year	0.018	0.843	0.010			
Citations	-0.018	0.032	0.268			
Top journal	0.482	0.085	1.000	0.442	0.074	0.000
Impact	-0.001	0.005	0.025			
Constant	-0.579	NA	1.000	-0.330	0.874	0.706
Observations	2,526			2,526		

Notes: EIS = elasticity of intertemporal substitution. PIP = posterior inclusion probability. Country characteristics are always included in all models of the BMA. In the frequentist check we only include method characteristics with PIP > 0.5. Standard errors in the frequentist check are clustered at the country level. More details on the BMA estimation are available in Table A3 and Figure A1.

the highest posterior model probabilities are shown, but we can see that they capture almost all of the cumulative model probabilities.

The best models in terms of posterior probabilities are depicted on the left. The very best one includes only 9 out of the 30 method variables at our disposal; the variables included are *inverse estimation*, *top journal*, *stock return*, *total consumption*, *OLS*, *no. of years*, *asset holders*, *exact Euler*, and *capital return*. *Monthly data* is not included in the best model, but it belongs to most of the other good models, and has a posterior inclusion probability larger than 0.5. All other method variables have posterior inclusion probabilities below 0.5, which indicates that they do not matter much for the magnitude of the estimated elasticity. Concerning the country-level variables (which are included in all models), we can see that *GDP per capita* and *credit availability* have the same estimated influence on the EIS no matter what method variables are included. In contrast, the estimated signs for *real interest* and *rule of law* are unstable and depend on the specification of the model.

The numerical results of the BMA estimation are summarized in Table 1. For each variable we report the estimated posterior mean for the regression parameter and the corresponding posterior standard deviation together with the posterior inclusion probability (for country-level variables the posterior inclusion probability is one by definition). In the right-hand part of the table we report the results of the frequentist check of our BMA estimation; that is, we also run a simple OLS. In the OLS we only include variables that proved to be relatively important in the BMA exercise (those with posterior inclusion probabilities above 0.5) and cluster the standard errors at the country level. We can see that the results of the frequentist check are very similar to the BMA results. Diagnostics of the BMA estimation are available in Table A3 and Figure A1 in the Appendix.

Concerning method variables, our results suggest that the type of utility function does not affect the reported estimates of the EIS in a systematic way. On the other hand, we find that certain aspects of the data are important, namely, that studies using longer time series report smaller estimates of the elasticity and that monthly frequency of data is associated with larger estimates. Both these effects, however, are rather small. An important aspect of study design is whether the EIS is estimated directly in a regression with consumption growth as the response variable or if the inverse of the EIS is estimated in a regression where asset return is on the

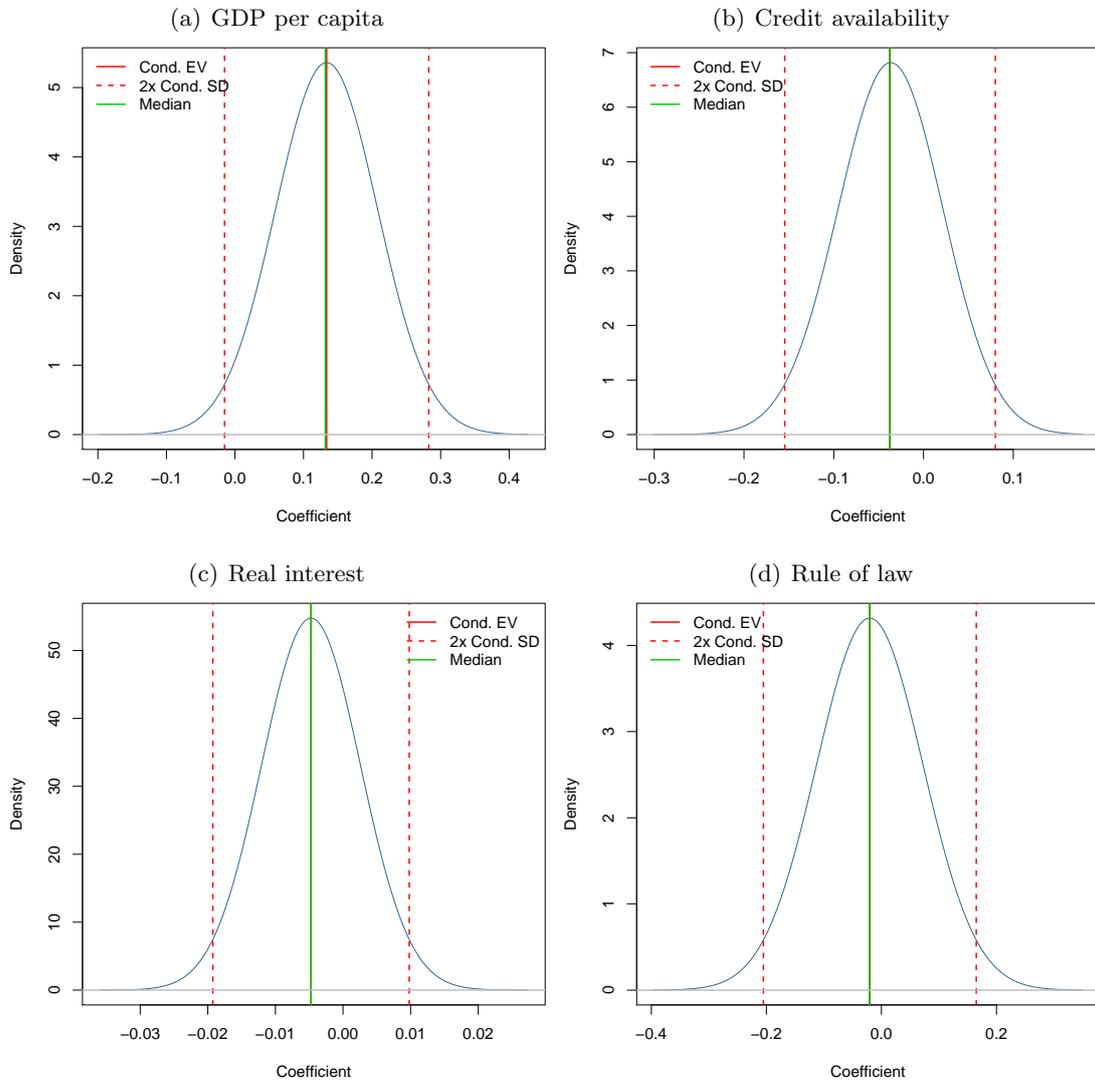
left-hand side. In the latter case the implied elasticity tends to be larger on average by 0.5, which is a significant difference considering that the mean of all the reported estimates is 0.5 and the practical relevance of such changes of the EIS is large, as illustrated in Figure 1.

When the elasticity of intertemporal substitution is estimated for a sub-sample of rich households or stockholders, the estimate tends to be substantially larger as well: by 0.35. Thus poor households and non-asset holders seem to display a significantly smaller EIS, which is in line with Mankiw & Zeldes (1991), Blundell *et al.* (1994), and Vissing-Jorgensen (2002), among others. The definitions of the two main variables in the consumption Euler equations—consumption and asset return—are important as well. When total consumption is used instead of non-durable consumption, the study is likely to find a larger EIS. Also, the use of bond returns as the measure of asset returns, in contrast to the use of stock returns or returns on a unit of capital, is associated with a larger reported EIS.

Studies that estimate the exact consumption Euler equation (that is, studies that do not use log-linear approximation) usually report a larger elasticity. Failure to acknowledge endogeneity when regressing consumption growth on asset returns results in substantial overestimation of the EIS: by about 0.4. Finally, our results also indicate that studies published in the top five general interest journals in economics tend to report estimates of the EIS larger by 0.5 compared with studies published in other journals. The difference may reflect aspects of quality that are not captured by the other variables we collected. Papers published in top journals often present novel methodology, and method aspects that have only been used by a few studies are difficult to examine in a meta-analysis framework.

The country-level variables, which are the main focus of our paper, are included in all the regressions, so for these variables the posterior inclusion probabilities reported in Table 1 are not informative. Instead we need to look at the posterior distribution of the regression coefficients reported in Figure 5. From the figure we can see that the estimated regression parameters for *credit availability*, *real interest*, and *rule of law* are close to zero. The dashed lines denote values that lie two standard deviations from the mean of the estimated regression parameter; therefore, they can be interpreted as analogous to 95% confidence intervals in frequentist econometrics. Even for *GDP per capita* the interval includes zero, but only marginally, which is analogous to borderline statistical significance at the 5% level. The frequentist check of BMA reported

Figure 5: Posterior coefficient distributions for country characteristics



Notes: The figure depicts the densities of the regression parameters encountered in different regressions (with different subsets of control variables on the right-hand side). For example, the regression coefficient for *GDP per capita* is positive in almost all models, irrespective of the control variables included. The most common value of the coefficient is approximately 0.13. On the other hand, the coefficient for *Rule of law* is negative in one half of the models and positive in the other half, depending on which control variables are included. The most common value is 0.

in Figure 5 shows statistical significance at the 10% level (and p-values larger than 0.5 for the other three country-level variables). We conclude that there seems to be a positive association between income and the elasticity of intertemporal substitution; the economic significance of this association is examined at the end of this section.

As a next step we add the variable *stock market participation* to the model, which reduces the number of countries to 28—the ones for which information on stock market participation is available—and we label them “core countries.” We are especially interested in the effect the new variable has on the estimated EIS, but we also examine the robustness of our results compared with the case where data for all countries were included. Even though this new BMA estimation includes far fewer countries, it only loses about 270 observations, because most studies estimate the EIS using data from the core countries.

The results of the BMA estimation with *stock market participation* are reported in Table 2; more details and diagnostics are available in Table A4 and Figure A2 in the Appendix. Concerning method characteristics, there are several changes compared with the estimation using all countries. First, it matters for the reported EIS whether the assumed utility function allows for non-separabilities between durable and non-durable consumption goods: allowing for non-separabilities is associated with larger estimated elasticities. Nevertheless, the variable has a posterior inclusion probability of only 0.54 and is not statistically significant in the frequentist check. Second, the posterior inclusion probability of the variable *exact Euler* drops to 0.29, so it seems to be less important when only the core countries are considered. Third, our results for the core countries suggest that highly cited studies report smaller estimates of the elasticity. But again, the corresponding variable has a posterior inclusion probability of only 0.6, and it is not significant in the frequentist check. Moreover, the posterior inclusion probability for this variable decreases sharply below 0.5 when we exclude the most cited study, Hall (1988), who reports small estimates.

Concerning the country-level variables, in the new BMA estimation we find a smaller posterior mean for the coefficient corresponding to *GDP per capita*; the variable also loses statistical significance in the frequentist check (nevertheless, the decrease in the posterior mean may reflect the positive correlation between *GDP per capita* and *stock market participation* of 0.54). The results concerning the remaining three variables do not change much, and the variables

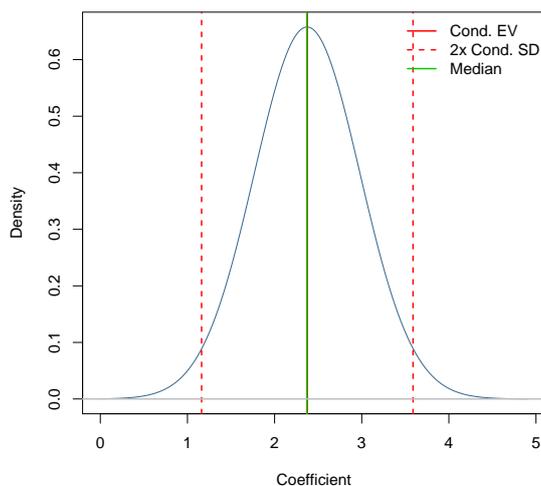
Table 2: Explaining the differences in the estimates of the EIS, core countries

Response variable:	Bayesian model averaging			Frequentist check (OLS)		
Estimate of the EIS	Post. mean	Post. std. dev.	PIP	Coef.	Std. er.	p-value
<i>Country characteristics</i>						
Stock market partic.	2.376	0.607	1.000	2.221	0.542	0.000
GDP per capita	0.080	0.137	1.000	0.116	0.138	0.405
Credit availability	-0.008	0.094	1.000	-0.003	0.122	0.982
Real interest	0.005	0.022	1.000	0.010	0.024	0.680
Rule of law	-0.283	0.193	1.000	-0.296	0.206	0.163
<i>Utility</i>						
Epstein-Zin	0.036	0.110	0.115			
Habits	-0.004	0.034	0.019			
Nonsep. durables	0.240	0.244	0.540	0.471	0.276	0.100
Nonsep. public	0.000	0.015	0.009			
Nonsep. tradables	0.004	0.042	0.016			
<i>Data</i>						
No. of households	-0.001	0.005	0.022			
No. of years	-0.248	0.059	0.996	-0.226	0.059	0.001
Average year	-0.025	0.860	0.010			
Micro data	-0.001	0.022	0.015			
Annual data	0.001	0.012	0.012			
Monthly data	0.141	0.166	0.506	0.326	0.054	0.000
<i>Design</i>						
Quasipanel	-0.107	0.191	0.273			
Inverse estimation	0.575	0.073	1.000	0.598	0.097	0.000
Asset holders	0.210	0.208	0.558	0.372	0.143	0.015
First lag instrument	0.002	0.019	0.022			
No year dummies	-0.007	0.066	0.021			
Income	-0.001	0.012	0.012			
Taste shifters	0.000	0.008	0.010			
<i>Variable definition</i>						
Total consumption	0.416	0.103	0.993	0.409	0.142	0.008
Food	0.016	0.080	0.057			
Stock return	-0.322	0.097	0.974	-0.358	0.158	0.032
Capital return	-0.224	0.164	0.714	-0.331	0.051	0.000
<i>Estimation</i>						
Exact Euler	0.067	0.114	0.287			
ML	-0.022	0.082	0.086			
TSLS	-0.002	0.021	0.022			
OLS	0.394	0.136	0.957	0.385	0.181	0.044
<i>Publication</i>						
Publication year	-0.074	1.288	0.012			
Citations	-0.052	0.048	0.595	-0.089	0.055	0.117
Top journal	0.529	0.104	1.000	0.567	0.103	0.000
Impact	0.000	0.004	0.016			
Constant	0.892	NA	1.000	-0.220	1.427	0.878
Observations	2,254			2,254		

Notes: EIS = elasticity of intertemporal substitution. PIP = posterior inclusion probability. Country characteristics are always included in all models of the BMA. In the frequentist check we only include method characteristics with PIP > 0.5. Standard errors in the frequentist check are clustered at the country level. More details on the BMA estimation are available in Table A4 and Figure A2.

still appear to be quite unimportant. In contrast, the newly included *stock market participation* is positively associated with the estimated elasticities, as we can see from Figure 6. The regression parameter for this variable is positive in virtually all regressions in which the variable is included. Also, in the frequentist check the variable is highly statistically significant, with a p-value below 0.001. Our results thus suggest that households in countries with high stock market participation tend to be more willing or able to substitute consumption intertemporally.

Figure 6: Posterior coefficient distribution for *stock market participation*



Notes: The figure depicts the densities of the regression parameters encountered in different regressions (with different subsets of control variables on the right-hand side).

But is the effect of *stock market participation* economically important? The estimated posterior mean for the regression coefficient corresponding to the variable is 2.4, so that an increase in stock market participation of 10 percentage points is associated with an increase in the EIS of 0.24; an important difference according to the simulation shown in Figure 1. In Table 3 we compute what happens to the estimated elasticity if the value of a country-level characteristic changes from its sample minimum to its sample maximum (“maximum effect”) and if the value increases by one standard deviation (“standard-deviation effect”). For variables *GDP per capita*, *credit availability*, *real interest*, and *rule of law*, we prefer to use the coefficients from the BMA estimation with all countries; for the variable *stock market participation* we have to use the value from the estimation with the core countries only. Out of the five country-level variables, *stock market participation* has the largest effect, followed by *GDP per capita*. The other variables do not seem to matter much. The maximum effect of changes in *stock market*

participation is a whopping 0.93; the standard-deviation effect is 0.14, which can also make a difference to the results of structural models, as shown in Figure 1.

Table 3: The economic significance of differences in country characteristics

Variable	Maximum effect	Std. dev. effect
Stock market partic.	0.931	0.141
GDP per capita	0.683	0.088
Credit availability	-0.119	-0.020
Real interest	-0.265	-0.019
Rule of law	-0.087	-0.012

Notes: The table depicts the predicted effects of increases in the variables on the EIS estimates based on the BMA results (the specification with core countries for *stock market participation*; the specification with all countries for the other variables). Maximum effect = an increase from sample minimum to sample maximum. Std. dev. effect = a one-standard-deviation increase.

5 Robustness Checks

In this section we evaluate the robustness of our findings by employing different variants of the BMA specification with the core countries—that is, including the variable *Stock market participation*. First, we run a BMA estimation in which country-level variables are treated in the same way as method variables; in other words, different models may or may not include country-level variables, in contrast to the previous analysis, in which country-level variables were included in all models. Table 4 provides the results (here we do not report results for variables with posterior inclusion probability below 0.5), and more details and diagnostics are available in Table A5 and Figure A3 in the Appendix.

In this estimation the posterior inclusion probabilities for country-level variables are not necessarily 1, and indeed the probabilities for all variables except *stock market participation* are lower than 0.5, which means that these variables do not help us explain the variation in the reported elasticities once the characteristics of methodology are taken into account. In contrast, the posterior inclusion probability of *Stock market participation* is 0.92, which would be characterized as “substantial” in the guidelines for the interpretation of the posterior inclusion probability by Eicher *et al.* (2011). Moreover, in the frequentist check the variable is statistically significant at the 1% level.

The regression parameter for *stock market participation* estimated by BMA is now lower than

in the previous case, but still implies an important effect on the estimated EIS: an increase in stock market participation of 10 percentage points is associated with an increase in the estimated elasticity of 0.18. Concerning the method variables, the results of the robustness check are similar to the baseline case, where the country-level variables are included in all models, but a few differences emerge. First, the data frequency does not seem to be important for the estimated EIS when country and method variables are treated in the same way. Second, the results suggest that estimating the exact Euler equation, instead of the log-linearized version, tends to deliver larger elasticities—we reported the same finding for the BMA estimation with all countries (that is, excluding *stock market participation*). Third, according to this robustness check the number of study citations is not associated with the magnitude of the reported elasticity.

Table 4: Robustness check: no fixed variables

Response variable:	Bayesian model averaging			Frequentist check (OLS)		
	Estimate of the EIS	Post. mean	Post. std. dev.	PIP	Coef.	Std. er.
Stock market partic.	1.775	0.736	0.917	2.128	0.613	0.002
GDP per capita	0.000	0.010	0.008	0.060	0.166	0.721
Credit availability	-0.002	0.016	0.021	0.040	0.129	0.760
Real interest	0.000	0.002	0.008	-0.004	0.026	0.879
Rule of law	-0.013	0.062	0.053	-0.290	0.238	0.234
Inverse estimation	0.563	0.078	1.000	0.535	0.146	0.001
Top journal	0.502	0.103	1.000	0.418	0.074	0.000
Total consumption	0.449	0.095	0.999	0.439	0.101	0.000
No. of years	-0.255	0.056	0.999	-0.232	0.050	0.000
Stock return	-0.340	0.088	0.990	-0.341	0.139	0.022
OLS	0.438	0.120	0.986	0.521	0.148	0.002
Capital return	-0.231	0.160	0.735	-0.282	0.054	0.000
Asset holders	0.277	0.210	0.694	0.404	0.115	0.002
Exact Euler	0.138	0.144	0.522	0.283	0.226	0.221
Constant	0.746	NA	1.000	0.105	1.634	0.950
Observations	2,254			2,254		

Notes: PIP = posterior inclusion probability. Country characteristics and method variables are treated in the same way in the BMA estimation. Results for method characteristics with PIP < 0.5 are not reported. Standard errors in the frequentist check are clustered at the country level. More details on the BMA estimation are available in Table A5 and Figure A3.

The second robustness check involves different priors for the BMA estimation. Now we use the priors that are advocated by Eicher *et al.* (2011) because they typically perform well in forecasting exercises: the unit information g-prior (the prior provides the same amount of information as one observation) and the uniform model prior (each model has the same probability). As we have noted, BMA runs many regressions with different combinations of the explanatory variables on the right-hand side and not all of the variables have to be included. It

follows that models of size 15—the number of explanatory variables divided by two—are most common. If each model has the same probability, with the uniform model prior we implicitly impose the prior that the “true” model explaining the differences in the reported elasticities has 15 explanatory variables, which is apparent from Figure A4 in the Appendix. That is why for the baseline estimation we prefer the random model prior, which gives each model *size* the same prior probability and reflects the fact that we know little *ex ante* about how many variables should be included in the model. The results of the robustness check are reported in Table 5 and for both country-level and method variables they are virtually identical to the baseline case.

Table 5: Robustness check: priors according to Eicher *et al.* (2011)

Response variable: Estimate of the EIS	Bayesian model averaging			Frequentist check (OLS)		
	Post. mean	Post. std. dev.	PIP	Coef.	Std. er.	p-value
Stock market partic.	2.328	0.598	1.000	2.221	0.542	0.000
GDP per capita	0.082	0.137	1.000	0.116	0.138	0.405
Credit availability	-0.018	0.095	1.000	-0.003	0.122	0.982
Real interest	0.007	0.022	1.000	0.010	0.024	0.680
Rule of law	-0.258	0.192	1.000	-0.296	0.206	0.163
Inverse estimation	0.594	0.070	1.000	0.598	0.097	0.000
Top journal	0.554	0.101	1.000	0.567	0.103	0.000
Stock return	-0.345	0.081	0.998	-0.358	0.158	0.032
Total consumption	0.416	0.098	0.998	0.409	0.142	0.008
No. of years	-0.247	0.059	0.998	-0.226	0.059	0.001
OLS	0.383	0.127	0.969	0.385	0.181	0.044
Capital return	-0.305	0.128	0.921	-0.331	0.051	0.000
Asset holders	0.294	0.192	0.771	0.372	0.143	0.015
Citations	-0.067	0.045	0.762	-0.089	0.055	0.117
Nonsep. durables	0.331	0.231	0.738	0.471	0.276	0.100
Monthly data	0.193	0.165	0.641	0.326	0.054	0.000
Constant	1.199	NA	1.000	-0.220	1.427	0.878
Observations	2,254			2,254		

Notes: PIP = posterior inclusion probability. In this specification we employ the priors suggested by Eicher *et al.* (2011), who recommend using the uniform model prior (each model has the same prior probability) and the unit information prior (the prior provides the same amount of information as one observation). Results for method characteristics with PIP < 0.5 are not reported. Standard errors in the frequentist check are clustered at the country level. More details on the BMA estimation are available in Table A6 and Figure A4.

Finally, in the third robustness check we use different proxies for liquidity constraints and institutions. Instead of the measure of credit availability reported in the Global Competitiveness Report we now employ the measure of financial reform published by the IMF; instead of perceptions of the rule of law in society we employ the measure of generalized trust developed by Bjoernskov & Meon (2013). The result concerning *stock market participation* holds: the variable is positively and strongly associated with the elasticity of intertemporal substitution. The other variables are less important, even though *GDP per capita* and *Financial reform* yield

Table 6: Robustness check: alternative proxies for liquidity constraints and institutions

Response variable:	Bayesian model averaging			Frequentist check (OLS)		
	Estimate of the EIS	Post. mean	Post. std. dev.	PIP	Coef.	Std. er.
Stock market partic.	2.399	0.609	1.000	2.342	0.848	0.011
GDP per capita	0.137	0.142	1.000	0.198	0.114	0.095
Financial reform	-0.692	0.307	1.000	-0.777	0.394	0.060
Real interest	0.025	0.023	1.000	0.023	0.032	0.493
Trust	-0.006	0.005	1.000	-0.005	0.004	0.257
Inverse estimation	0.577	0.075	1.000	0.627	0.103	0.000
Top journal	0.543	0.104	1.000	0.602	0.114	0.000
Total consumption	0.423	0.100	0.996	0.416	0.147	0.009
No. of years	-0.236	0.061	0.991	-0.228	0.058	0.001
OLS	0.412	0.126	0.976	0.443	0.189	0.028
Stock return	-0.303	0.101	0.961	-0.299	0.136	0.037
Asset holders	0.299	0.211	0.728	0.406	0.130	0.005
Citations	-0.063	0.049	0.682	-0.093	0.057	0.119
Capital return	-0.182	0.168	0.596	-0.265	0.061	0.000
Nonsep. durables	0.257	0.247	0.570	0.465	0.273	0.101
Constant	-0.440	NA	1.000	-0.797	1.093	0.473
Observations	2,254			2,254		

Notes: PIP = posterior inclusion probability. In this specification we replace *Credit availability* with *Financial reform* and *Rule of law* with *Trust*. Results for method characteristics with PIP < 0.5 are not reported. Standard errors in the frequentist check are clustered at the country level. More details on the BMA estimation are available in Table A7 and Figure A5.

statistical significance at the 10% level in the frequentist check of the BMA estimation. Concerning the method variables, the results are close to the baseline case, with the exception of data frequency, which seems to be unimportant here, similarly to the first robustness check and the BMA estimation with all countries.

As we have noted, for all analyses in the paper we exclude estimates of the EIS larger than 10 in absolute value. It is necessary to exclude outliers because the inverse method of estimation used by some researchers can yield implausible estimates of the elasticity—even larger than 100 in absolute value. Because with the asset return on the left-hand side the researcher estimates the inverse of the EIS (the coefficient of relative risk aversion under the typical power utility), imprecise estimation may yield a coefficient close to zero and imply that the EIS is close to infinity. The threshold of 10 is arbitrary, but we get very similar results with the threshold set to 1, 5, 20, and 100. Moreover, the results are also similar when we include all estimates of the EIS and employ the robust estimator developed by Verardi & Croux (2009) for the frequentist check. As far as we know, a variant of robust estimation is not yet available for the BMA framework.

6 Concluding Remarks

We present a quantitative survey of estimates of the elasticity of intertemporal substitution in what we believe is the largest meta-analysis conducted in economics. We collect 2,735 estimates from 169 published studies and find that the mean elasticity is 0.5, but that the estimates vary greatly across countries and methods. We use Bayesian model averaging to explore country-level heterogeneity while controlling for 30 variables that reflect different techniques used in the estimation of the elasticity. We find that households in countries with higher income per capita and higher stock market participation show larger values of the EIS. Thus, using a unique cross-country data set we corroborate the micro-level findings of Blundell *et al.* (1994) and Attanasio & Browning (1995), who report a larger elasticity for richer households, and Mankiw & Zeldes (1991) and Vissing-Jorgensen (2002), who find a larger EIS for asset holders than for other households. Our results also suggest that researchers obtain systematically larger estimates of the EIS when they estimate the parameter using a sub-sample of rich households or asset holders.

Rich households substitute consumption across time periods more easily because necessities, which are difficult to substitute intertemporally, constitute a smaller fraction of their consumption bundle in comparison with poor households. Moreover, the opportunities for intertemporal substitution for households in developing countries may be restricted by subsistence requirements (Ogaki *et al.*, 1996). Concerning asset holders, Vissing-Jorgensen (2002) points out that the consumption Euler equation need not be valid for households that do not participate in asset markets, leading to estimates of the EIS close to zero. Another possible explanation is that exposure to financial markets, especially the stock market, may make households more forward-looking and willing to substitute consumption in response to changes in expected asset returns.

Several aspects of methodology affect the reported elasticities in a systematic way. For example, the definition of the utility function is important, especially whether researchers allow for non-separabilities between durable and non-durable consumption goods. The size of the data set matters for the estimated elasticities as well. Further, when researchers use asset returns as the response variable and estimate the inverse of the EIS, the implied elasticity tends to be substantially larger—on average by about 0.5 compared to the case where consumption growth

is used as the response variable. The definition of consumption growth (total consumption, non-durables, or food expenditure) and asset return (bond, stock, or capital return) is also important. Ignoring the presence of endogeneity typically leads to overestimation of the elasticity. Finally, the top five general interest journals in economics tend to publish substantially larger estimates than other journals, which may reflect unobserved aspects of study quality.

An important issue that we do not discuss in this paper is publication selection bias. Several commentators have suggested that in empirical economics statistically insignificant results tend to be underreported and that the resulting mean estimate observed in the literature may be biased (DeLong & Lang, 1992; Card & Krueger, 1995; Ashenfelter & Greenstone, 2004; Stanley, 2005). We analyze publication selection bias in the EIS literature in a companion paper, Havranek (2013), and believe that while such bias can affect the mean reported elasticity, it is not related to country-level heterogeneity in the EIS.

References

- ABIAD, A., E. DETRAGIACHE, & T. TRESSEL (2010): "A New Database of Financial Reforms." *IMF Staff Papers* **57(2)**: pp. 281–302.
- AI, H. (2010): "Information Quality and Long-Run Risk: Asset Pricing Implications." *Journal of Finance* **65(4)**: pp. 1333–1367.
- ASHENFELTER, O. & M. GREENSTONE (2004): "Estimating the Value of a Statistical Life: The Importance of Omitted Variables and Publication Bias." *American Economic Review* **94(2)**: pp. 454–460.
- ASHENFELTER, O., C. HARMON, & H. OOSTERBEEK (1999): "A Review of Estimates of the Schooling/Earnings Relationship, with Tests for Publication Bias." *Labour Economics* **6(4)**: pp. 453–470.
- ATKESON, A. & M. OGAKI (1996): "Wealth-varying intertemporal elasticities of substitution: Evidence from panel and aggregate data." *Journal of Monetary Economics* **38(3)**: pp. 507–534.
- ATTANASIO, O. P. (1995): "The intertemporal allocation of consumption: theory and evidence." *Carnegie-Rochester Conference Series on Public Policy* **42(1)**: pp. 39–56.
- ATTANASIO, O. P., J. BANKS, & S. TANNER (2002): "Asset Holding and Consumption Volatility." *Journal of Political Economy* **110(4)**: pp. 771–792.
- ATTANASIO, O. P. & M. BROWNING (1995): "Consumption over the Life Cycle and over the Business Cycle." *American Economic Review* **85(5)**: pp. 1118–37.
- ATTANASIO, O. P. & G. WEBER (1993): "Consumption Growth, the Interest Rate and Aggregation." *Review of Economic Studies* **60(3)**: pp. 631–49.
- ATTANASIO, O. P. & G. WEBER (1995): "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey." *Journal of Political Economy* **103(6)**: pp. 1121–57.
- BARRO, R. J. (2009): "Rare Disasters, Asset Prices, and Welfare Costs." *American Economic Review*

- 99(1)**: pp. 243–64.
- BAYOUMI, T. A. (1993): “Financial Deregulation and Consumption in the United Kingdom.” *The Review of Economics and Statistics* **75(3)**: pp. 536–39.
- BJOERNSKOV, C. & P.-G. MEON (2013): “Is trust the missing root of institutions, education, and development?” *Public Choice* (**forthcoming**).
- BLUNDELL, R., M. BROWNING, & C. MEGHIR (1994): “Consumer Demand and the Life-Cycle Allocation of Household Expenditures.” *Review of Economic Studies* **61(1)**: pp. 57–80.
- BROWNING, M. & A. LUSARDI (1996): “Household Saving: Micro Theories and Micro Facts.” *Journal of Economic Literature* **34(4)**: pp. 1797–1855.
- CAMPBELL, J. Y. (1999): “Asset Prices, Consumption, and the Business Cycle.” In J. B. TAYLOR & M. WOODFORD (editors), “Handbook of Macroeconomics,” volume 1 of *Handbook of Macroeconomics*, chapter 19, pp. 1231–1303. Elsevier.
- CAMPBELL, J. Y. & N. G. MANKIW (1989): “Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence.” In “NBER Macroeconomics Annual 1989, Volume 4,” NBER Chapters, pp. 185–246. National Bureau of Economic Research, Inc.
- CARD, D. & A. B. KRUEGER (1995): “Time-Series Minimum-Wage Studies: A Meta-Analysis.” *American Economic Review* **85(2)**: pp. 238–43.
- CARROLL, C. D. (2001): “Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation).” *Advances in Macroeconomics* **1**: p. 6.
- CHARI, V. V., P. J. KEHOE, & E. R. MCGRATTAN (2002): “Can Sticky Price Models Generate Volatile and Persistent Real Exchange Rates?” *Review of Economic Studies* **69(3)**: pp. 533–63.
- CHETTY, R., A. GUREN, D. MANOLI, & A. WEBER (2011): “Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins.” *American Economic Review* **101(3)**: pp. 471–75.
- CICCONE, A. & M. JAROCINSKI (2010): “Determinants of Economic Growth: Will Data Tell?” *American Economic Journal: Macroeconomics* **2(4)**: pp. 222–46.
- COLACITO, R. & M. M. CROCE (2011): “Risks for the Long Run and the Real Exchange Rate.” *Journal of Political Economy* **119(1)**: pp. 153–181.
- CROSSLEY, T. F. & H. W. LOW (2011): “Is the elasticity of intertemporal substitution constant?” *Journal of the European Economic Association* **9(1)**: pp. 87–105.
- DELONG, J. B. & K. LANG (1992): “Are All Economic Hypotheses False?” *Journal of Political Economy* **100(6)**: pp. 1257–72.
- DISDIER, A.-C. & K. HEAD (2008): “The Puzzling Persistence of the Distance Effect on Bilateral Trade.” *The Review of Economics and Statistics* **90(1)**: pp. 37–48.
- DOUCOULIAGOS, H. & T. D. STANLEY (2013): “Are All Economic Facts Greatly Exaggerated? Theory Competition and Selectivity.” *Journal of Economic Surveys* **27(2)**: pp. 316–339.
- EICHER, T. S., C. PAPAGEORGIOU, & A. E. RAFTERY (2011): “Default Priors and Predictive Performance in Bayesian Model Averaging, with Application to Growth Determinants.” *Journal of Applied Econometrics* **26(1)**: pp. 30–55.
- EPSTEIN, L. G. & S. E. ZIN (1989): “Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework.” *Econometrica* **57(4)**: pp. 937–69.

- FELDKIRCHER, M. & S. ZEUGNER (2009): “Benchmark Priors Revisited: On Adaptive Shrinkage and the Supermodel Effect in Bayesian Model Averaging.” *IMF Working Papers 09/202*, International Monetary Fund.
- FERNANDEZ, C., E. LEY, & M. F. J. STEEL (2001): “Benchmark priors for Bayesian model averaging.” *Journal of Econometrics* **100(2)**: pp. 381–427.
- GIANNETTI, M. & Y. KOSKINEN (2010): “Investor Protection, Equity Returns, and Financial Globalization.” *Journal of Financial and Quantitative Analysis* **45(01)**: pp. 135–168.
- GRUBER, J. (2006): “A Tax-Based Estimate of the Elasticity of Intertemporal Substitution.” *NBER Working Papers 11945*, National Bureau of Economic Research, Inc.
- GUISO, L., P. SAPIENZA, & L. ZINGALES (2004): “The Role of Social Capital in Financial Development.” *American Economic Review* **94(3)**: pp. 526–556.
- GUISO, L., P. SAPIENZA, & L. ZINGALES (2008): “Trusting the Stock Market.” *Journal of Finance* **63(6)**: pp. 2557–2600.
- HALL, R. E. (1988): “Intertemporal Substitution in Consumption.” *Journal of Political Economy* **96(2)**: pp. 339–57.
- HAVRANEK, T. (2013): “Publication Bias in Measuring Intertemporal Substitution.” *Working paper*, Czech National Bank and Charles University in Prague. Available at meta-analysis.cz/eis.
- HAVRANEK, T. & M. RUSNAK (2013): “Transmission Lags of Monetary Policy: A Meta-Analysis.” *International Journal of Central Banking* (**forthcoming**).
- HOUSE, C. L. & M. D. SHAPIRO (2006): “Phased-In Tax Cuts and Economic Activity.” *American Economic Review* **96(5)**: pp. 1835–1849.
- IRSOVA, Z. & T. HAVRANEK (2013): “Determinants of Horizontal Spillovers from FDI: Evidence from a Large Meta-Analysis.” *World Development* **42(1)**: pp. 1–15.
- JIN, K. (2012): “Industrial Structure and Capital Flows.” *American Economic Review* **102(5)**: pp. 2111–46.
- LEY, E. & M. F. STEEL (2009): “On the effect of prior assumptions in Bayesian model averaging with applications to growth regressions.” *Journal of Applied Econometrics* **24(4)**: pp. 651–674.
- MANKIW, N. G. & S. P. ZELDES (1991): “The Consumption of Stockholders and Non-Stockholders.” *NBER Working Papers 3402*, National Bureau of Economic Research, Inc.
- MOELTNER, K. & R. WOODWARD (2009): “Meta-Functional Benefit Transfer for Wetland Valuation: Making the Most of Small Samples.” *Environmental & Resource Economics* **42(1)**: pp. 89–108.
- MORAL-BENITO, E. (2012): “Determinants of Economic Growth: A Bayesian Panel Data Approach.” *The Review of Economics and Statistics* **94(2)**: pp. 566–579.
- MULLIGAN, C. B. (2002): “Capital, Interest, and Aggregate Intertemporal Substitution.” *NBER Working Papers 9373*, National Bureau of Economic Research, Inc.
- NELSON, J. & P. KENNEDY (2009): “The Use (and Abuse) of Meta-Analysis in Environmental and Natural Resource Economics: An Assessment.” *Environmental and Resource Economics* **42**: pp. 345–377.
- OGAKI, M., J. D. OSTRY, & C. M. REINHART (1996): “Saving Behavior in Low- and Middle-Income Developing Countries: A Comparison.” *IMF Staff Papers* **43(1)**: pp. 38–71.
- OGAKI, M. & C. M. REINHART (1998): “Measuring Intertemporal Substitution: The Role of Durable

- Goods.” *Journal of Political Economy* **106(5)**: pp. 1078–1098.
- PIAZZESI, M., M. SCHNEIDER, & S. TUZEL (2007): “Housing, Consumption and Asset Pricing.” *Journal of Financial Economics* **83(3)**: pp. 531–569.
- PORTA, R. L., F. L. DE SILANES, A. SHLEIFER, & R. W. VISHNY (1998): “Law and Finance.” *Journal of Political Economy* **106(6)**: pp. 1113–1155.
- RAFTERY, A. E., D. MADIGAN, & J. A. HOETING (1997): “Bayesian Model Averaging for Linear Regression Models.” *Journal of the American Statistical Association* **92**: pp. 179–191.
- RIEGER, M. O., M. WANG, & T. HENS (2011): “Prospect Theory around the World.” *Discussion Papers 2011/19*, Department of Finance and Management Science, Norwegian School of Economics.
- RUDEBUSCH, G. D. & E. T. SWANSON (2012): “The Bond Premium in a DSGE Model with Long-Run Real and Nominal Risks.” *American Economic Journal: Macroeconomics* **4(1)**: pp. 105–43.
- SMETS, F. & R. WOUTERS (2007): “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach.” *American Economic Review* **97(3)**: pp. 586–606.
- STANLEY, T. D. (2001): “Wheat from Chaff: Meta-Analysis as Quantitative Literature Review.” *Journal of Economic Perspectives* **15(3)**: pp. 131–150.
- STANLEY, T. D. (2005): “Beyond Publication Bias.” *Journal of Economic Surveys* **19(3)**: pp. 309–345.
- TRABANDT, M. & H. UHLIG (2011): “The Laffer Curve Revisited.” *Journal of Monetary Economics* **58(4)**: pp. 305–327.
- VERARDI, V. & C. CROUX (2009): “Robust regression in Stata.” *Stata Journal* **9(3)**: pp. 439–453.
- VISSING-JORGENSEN, A. (2002): “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution.” *Journal of Political Economy* **110(4)**: pp. 825–853.
- WANG, M., M. O. RIEGER, & T. HENS (2011): “How Time Preferences Differ: Evidence from 45 Countries.” *Discussion Papers 2011/18*, Department of Finance and Management Science, Norwegian School of Economics.
- WIELAND, V., T. CWIK, G. J. MÜLLER, S. SCHMIDT, & M. WOLTERS (2012): “A new comparative approach to macroeconomic modeling and policy analysis.” *Journal of Economic Behavior & Organization* **83(3)**: pp. 523–541.
- WIRJANTO, T. S. (1995): “Aggregate Consumption Behaviour and Liquidity Constraints: The Canadian Evidence.” *Canadian Journal of Economics* **28(4b)**: pp. 1135–52.

A Summary Statistics

Table A1: Meta-analyses of the elasticity of intertemporal substitution for individual countries

Country	Mean EIS	Std. err. of the mean	No. of estimates
Argentina	-0.171	0.221	12
Australia	0.362	0.160	32
Austria	3.149	1.876	6
Belgium	0.677	0.390	10
Brazil	0.107	0.093	19
Burma	0.439	0.042	4
Canada	0.389	0.110	91
Chile	0.137	0.077	7
China	0.530	0.234	5
Colombia	0.158	0.078	8
Denmark	0.488	0.588	7
Finland	0.185	0.320	46
France	-0.034	0.153	44
Germany	0.080	0.163	39
Greece	0.561	0.291	18
Hong Kong	0.099	0.017	33
Iceland	0.352	0.367	4
India	0.515	0.090	5
Indonesia	0.102	0.160	8
Ireland	1.739	0.778	7
Israel	0.235	0.033	65
Italy	0.290	0.162	33
Japan	0.893	0.243	109
Kenya	1.228	0.481	7
Korea	0.423	0.219	32
Malaysia	0.173	0.161	11
Mexico	0.158	0.053	12
Netherlands	0.027	0.221	31
New Zealand	2.206	0.269	4
Norway	-0.386	0.583	4
Pakistan	0.100	0.203	6
Philippines	-0.026	0.111	9
Portugal	0.152	0.258	7
Singapore	0.120	0.131	7
Spain	0.504	0.107	44
Sri Lanka	0.033	0.159	8
Sweden	0.065	0.126	63
Switzerland	-0.434	0.201	31
Taiwan	1.549	1.421	7
Thailand	0.081	0.064	9
Turkey	0.314	0.133	12
UK	0.487	0.070	251
Uruguay	0.117	0.124	5
US	0.594	0.036	1429
Venezuela	0.157	0.093	6

Notes: The table shows mean estimates of the EIS in countries for which at least 4 estimates are reported in the literature. Estimates larger than 10 in absolute value are excluded.

Table A2: Description and summary statistics of regression variables

Variable	Description	Mean	Std. dev.
EIS	Estimate of the elasticity of intertemporal substitution (response variable).	0.492	1.298
Country characteristics			
Stock market partic.	The fraction of households participating in the domestic stock market (source: Giannetti & Koskinen, 2010).	0.246	0.059
GDP per capita	Gross domestic product per capita at purchasing-power-adjusted 2005 dollars (source: Penn World Tables).	9.804	0.658
Credit availability	The ease of access to loans (source: The Global Competitiveness Report, www.weforum.org).	3.523	0.547
Financial reform	The IMF's financial reform index (source: Abiad <i>et al.</i> , 2010).	0.691	0.197
Real interest	The lending interest rate adjusted for inflation as measured by the GDP deflator (source: World Development Indicators).	4.448	3.954
Rule of law	The extent to which agents have confidence in the rules of society, and in particular the quality of contract enforcement (source: World Bank Global Governance Indicators).	1.404	0.611
Trust	Perceptions of general trust in society (source: Bjoernskov & Meon, 2013).	39.09	9.543
Method characteristics			
<i>Utility</i>			
Epstein-Zin	=1 if the estimation differentiates between the EIS and the coefficient of relative risk aversion.	0.053	0.224
Habits	=1 if habits in consumption are assumed.	0.040	0.196
Nonsep. durables	=1 if the model allows for nonseparability between durables and nondurables.	0.041	0.199
Nonsep. public	=1 if the model allows for nonseparability between private and public consumption.	0.044	0.206
Nonsep. tradables	=1 if the model allows for nonseparability between tradables and nontradables.	0.046	0.210
<i>Data</i>			
No. of households	The logarithm of the number of cross-sectional units used in the estimation (households, cohorts, countries).	1.103	2.384
No. of years	The logarithm of the number of years of the data period used in the estimation.	3.184	0.570
Average year	The logarithm of the average year of the data period.	7.590	0.006
Micro data	=1 if the coefficient comes from a micro-level estimation.	0.187	0.390
Annual data	=1 if the data frequency is annual.	0.328	0.469
Monthly data	=1 if the data frequency is monthly.	0.097	0.296
<i>Design</i>			
Quasipanel	=1 if quasipanel (synthetic cohort) data are used.	0.053	0.224
Inverse estimation	=1 if the rate of return is the response variable in the estimation.	0.317	0.465
Asset holders	=1 if the estimate is related to the rich or asset holders.	0.054	0.226
First lag instrument	=1 if the first lags of variables are included among the instruments.	0.305	0.460
No year dummies	=1 if year dummies are omitted in micro studies using the Panel Study of Income Dynamics.	0.030	0.171
Income	=1 if income is included in the specification.	0.241	0.428
Taste shifters	The logarithm of the number of controls for taste shifters.	0.117	0.452
<i>Variable definition</i>			
Total consumption	=1 if total consumption is used in the estimation.	0.203	0.402
Food	=1 if food is used as a proxy for nondurables.	0.059	0.235
Stock return	=1 if the rate of return is measured as the stock return.	0.189	0.392
Capital return	=1 if the rate of return is measured as the return on capital.	0.113	0.317

Continued on next page

Table A2: Description and summary statistics of regression variables (continued)

Variable	Description	Mean	Std. dev.
<i>Estimation</i>			
Exact Euler	=1 if the exact Euler equation is estimated.	0.238	0.426
ML	=1 if maximum likelihood methods are used for the estimation.	0.049	0.216
TSLS	=1 if two-stage least squares are used for the estimation.	0.338	0.473
OLS	=1 if ordinary least squares are used for the estimation.	0.104	0.306
<i>Publication</i>			
Publication year	The logarithm of the year of publication of the study.	7.601	0.004
Citations	The logarithm of the number of per-year citations of the study in Google Scholar.	2.024	1.256
Top journal	=1 if the study was published in one of the top five journals in economics.	0.207	0.405
Impact	The recursive RePEc impact factor of the outlet.	1.089	1.535

Notes: Method characteristics are collected from published studies estimating the elasticity of intertemporal substitution. The list of studies is available in the online appendix at meta-analysis.cz/substitution.

B Diagnostics of BMA

Table A3: Summary of BMA estimation, all countries

<i>Mean no. regressors</i> 14.1707	<i>Draws</i> $2 \cdot 10^6$	<i>Burn-ins</i> $1 \cdot 10^6$	<i>Time</i> 8.14355 minutes
<i>No. models visited</i> 377,919	<i>Modelspace</i> $1.7 \cdot 10^{10}$	<i>Visited</i> 0.0022%	<i>Topmodels</i> 96%
<i>Corr PMP</i> 0.9999	<i>No. Obs.</i> 2,526	<i>Model Prior</i> random	<i>g-Prior</i> BRIC
<i>Shrinkage-Stats</i> Av= 0.9996			

Notes: The “random” model prior refers to the beta-binomial prior advocated by Ley & Steel (2009): prior model probabilities are the same for all possible model sizes. We set the Zellner’s g prior following Fernandez *et al.* (2001).

Figure A1: Model size and convergence, BMA with all countries

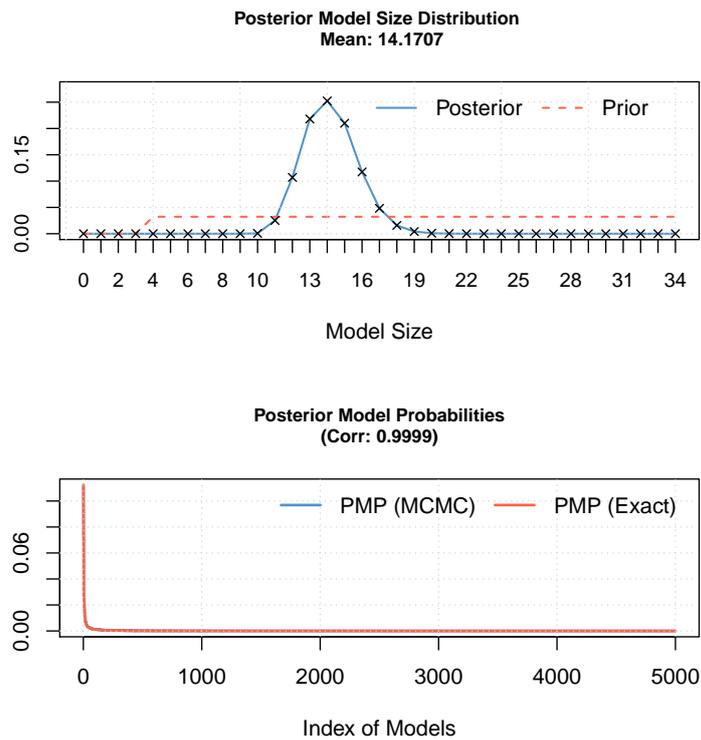


Table A4: Summary of BMA estimation, core countries

<i>Mean no. regressors</i>	<i>Draws</i>	<i>Burn-ins</i>	<i>Time</i>
14.9218	$2 \cdot 10^6$	$1 \cdot 10^6$	8.464817 minutes
<i>No. models visited</i>	<i>Modelspace</i>	<i>Visited</i>	<i>Topmodels</i>
478,214	$3.4 \cdot 10^{10}$	0.0014%	94%
<i>Corr PMP</i>	<i>No. Obs.</i>	<i>Model Prior</i>	<i>g-Prior</i>
0.9996	2,254	random	BRIC
<i>Shrinkage-Stats</i>			
Av= 0.9996			

Notes: The “random” model prior refers to the beta-binomial prior advocated by Ley & Steel (2009): prior model probabilities are the same for all possible model sizes. We set the Zellner’s g prior following Fernandez *et al.* (2001).

Figure A2: Model size and convergence, BMA with core countries

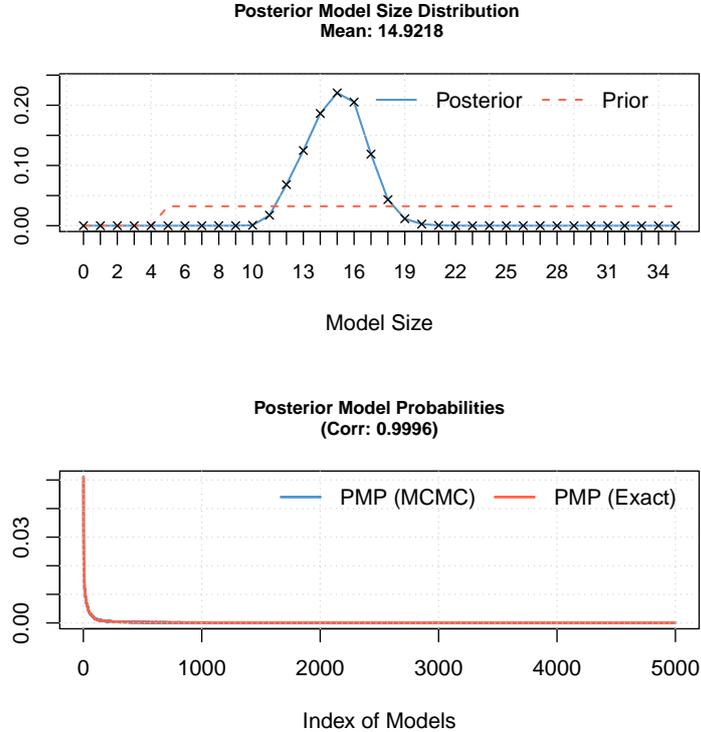


Table A5: Summary of BMA estimation, no fixed variables

<i>Mean no. regressors</i> 10.9643	<i>Draws</i> $2 \cdot 10^6$	<i>Burn-ins</i> $1 \cdot 10^6$	<i>Time</i> 7.003633 minutes
<i>No. models visited</i> 387,615	<i>Modelspace</i> $3.4 \cdot 10^{10}$	<i>Visited</i> 0.0011%	<i>Topmodels</i> 92%
<i>Corr PMP</i> 0.9995	<i>No. Obs.</i> 2,254	<i>Model Prior</i> random	<i>g-Prior</i> BRIC
<i>Shrinkage-Stats</i> Av= 0.9996			

Notes: The “random” model prior refers to the beta-binomial prior advocated by Ley & Steel (2009): prior model probabilities are the same for all possible model sizes. We set the Zellner’s g prior following Fernandez *et al.* (2001).

Figure A3: Model size and convergence, BMA with no fixed variables

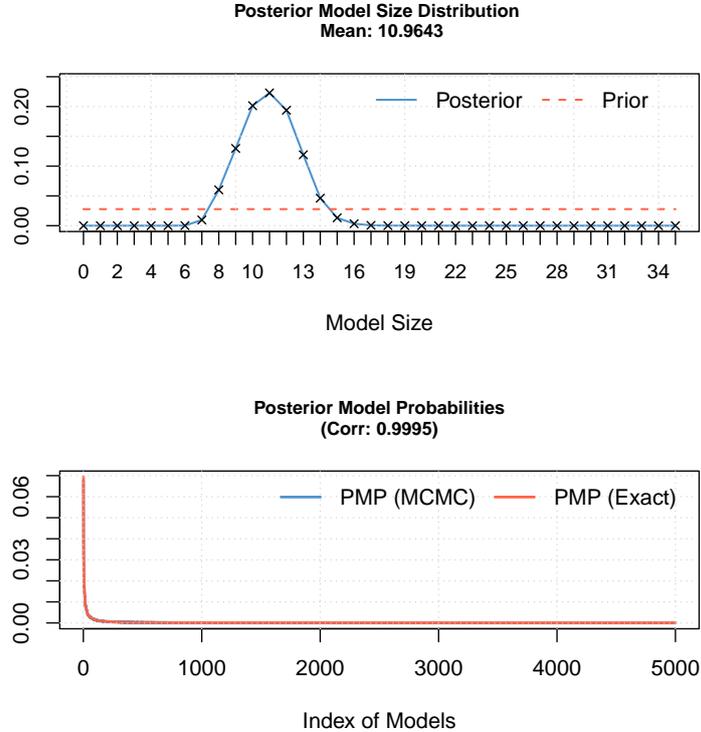


Table A6: Summary of BMA estimation, priors according to Eicher *et al.* (2011)

<i>Mean no. regressors</i> 16.3370	<i>Draws</i> $2 \cdot 10^6$	<i>Burn-ins</i> $1 \cdot 10^6$	<i>Time</i> 8.44965 minutes
<i>No. models visited</i> 497,193	<i>Modelspace</i> $3.4 \cdot 10^{10}$	<i>Visited</i> 0.0014%	<i>Topmodels</i> 90%
<i>Corr PMP</i> 0.9994	<i>No. Obs.</i> 2,254	<i>Model Prior</i> uniform	<i>g-Prior</i> UIP
<i>Shrinkage-Stats</i> Av= 0.9996			

Notes: In this specification we employ the priors suggested by Eicher *et al.* (2011), who recommend using the uniform model prior (each model has the same prior probability) and the unit information prior (the prior provides the same amount of information as one observation).

Figure A4: Model size and convergence, BMA with priors according to Eicher *et al.* (2011)

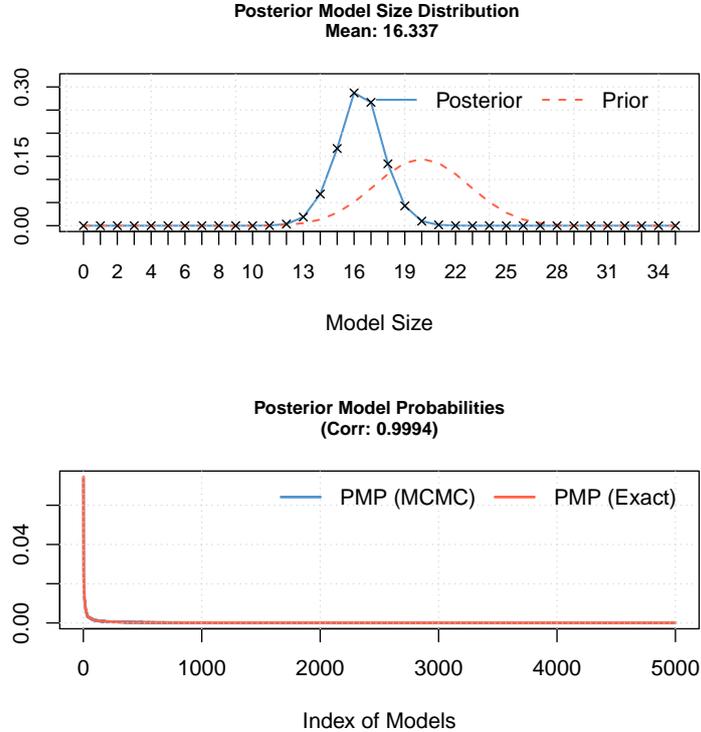


Table A7: Summary of BMA estimation, alternative proxies

<i>Mean no. regressors</i> 14.9921	<i>Draws</i> $2 \cdot 10^6$	<i>Burn-ins</i> $1 \cdot 10^6$	<i>Time</i> 8.557683 minutes
<i>No. models visited</i> 443,396	<i>Modelspace</i> $3.4 \cdot 10^{10}$	<i>Visited</i> 0.0013%	<i>Topmodels</i> 95%
<i>Corr PMP</i> 0.9993	<i>No. Obs.</i> 2,254	<i>Model Prior</i> random	<i>g-Prior</i> BRIC
<i>Shrinkage-Stats</i> Av= 0.9996			

Notes: The “random” model prior refers to the beta-binomial prior advocated by Ley & Steel (2009): prior model probabilities are the same for all possible model sizes. We set the Zellner’s g prior following Fernandez *et al.* (2001).

Figure A5: Model size and convergence, BMA with alternative proxies

