A nonparametric analysis of the consumption function with a role for financial development

Joshua Duarte, Pedro Bação, João Sousa Andrade*†‡

December 15, 2021

Abstract

In this paper we analyze the effect of financial development on consumption using nonparametric regression methods for panel data. In high-income countries, financial development has a positive impact on consumption. However, the responsiveness of consumption to financial development decreases as the level of financial development rises. In low-income countries, intermediate levels of financial development appear to be associated with lower consumption, while variations in the level of financial development in the tails appear not to affect consumption. The response of consumption to the remaining regressors has a magnitude and sign along the lines of what previous literature on the subject has found. However, the results suggest that these responses are nonlinear, depending not only on the level of the regressor, but also on the level of financial development.

JEL Classification: E20, E44, G20.

Keywords: Consumption, Financial Development, Nonparametric Regression, Local Linear Least Squares.

---

*Duarte: University of Coimbra, joshua.duarte@uc.pt. Bação: CeBER/University of Coimbra, pmab@fe.uc.pt. Andrade: CeBER/University of Coimbra, jasa@fe.uc.pt. Corresponding author: Pedro Bação. Address: University of Coimbra, CeBER, Faculty of Economics, Av Dias da Silva 165, 3004-512 Coimbra, Portugal. Email: pmab@fe.uc.pt. Telephone: +351239790570.

†Acknowledgments: This research has received financial support from the Portuguese Foundation for Science and Technology (FCT) under the fellowship SFRH/BD/111038/2015, with national funds from MCTES and european funds from the European Social Fund.

‡The authors have no conflict of interest to declare.

1 Introduction

The structure of the economies and the way they work has changed markedly over the last decades. One of the major changes of the past decades has been the unprecedented level of financial development achieved. This development leads to the question of whether and how what happens in the financial system affects the rest of the economy. In particular, given the importance of consumption smoothing across time in standard models, consumption is a natural focus for research on the wider effects of financial development. In fact, as Deaton (1992) stresses, most of the models of consumption assumed that it is possible to smooth consumption over several time periods. This implies that there is some form of a financial system in operation. Changes to that financial system that expand or restrict the ability to smooth consumption may therefore affect consumption. This is the primary reason for posing the hypothesis that financial development may impact consumption. The model in Guerrieri and Lorenzoni (2017) shows that a recession originated in the financial system can impact consumption, as well as the interest rate and output, through a credit crunch. Other examples of authors that show how consumption might possibly be linked with changes to the financial system can be found in Bandiera et al. (2000), Carvalho et al. (2012), Martin and Ventura (2011), Dewachter and Wouters (2014) and He and Krishnamurthy (2014).

Related to the previous argument is the possibility that financial development may have not just a direct impact on consumption, but also an indirect effect. The indirect impact would consist of a change in the way consumption responds to its other, traditional, determinants. This has been studied in, e.g., Lee (2013), Fisher et al. (2012) and Estrada et al. (2014).

In this paper we use nonparametric methods to estimate a consumption model and analyze whether and how consumption reacts to financial development. The nonparametric approach has two advantages over a parametric model. First, it allows us to form an idea of what the link between consumption and financial development might look like without imposing strong a priori assumptions about it. Second, it allows us to dodge most of the discussion on what form should the consumption function take, i.e., how should the traditional variables enter it. Despite these advantages, applications of nonparametric methodologies to the analysis of aggregate consumption analysis are relatively scarce. Notable exceptions are Swofford and Whitney (1987), Delgado and Miles (1997), Easaw et al. (2005), Cherchye et al. (2007), Cherchye et al. (2009) and Bruno (2014).

The results we obtain show some evidence of a nonlinear effect of financial development on consumption, with the effect depending on the value of the remaining regressors as well. The effect is negative for low values of the remaining regressors and positive for high values of the remaining regressors.

The paper is organized as follows. Section 2 describes the models and the nonparametric techniques that we used to estimate them. Section 3 provides a brief presentation of our data. Section 4 does the same for the nonparametric regression models. Section 6 concludes the paper.

2 Model and Methodology

2.1 Model

In this section we present the consumption model that we will use throughout the paper to study consumption behavior and its relation to financial development.

Our choice of variables is made in line with has been done in the literature, which gives serious importance to income, wealth and the interest rate as determinants of consumption. Several other regressors have been suggested in the literature over the years, such as wealth and credit, but in accordance to the methodology we use we have to be parsimonious regarding the number of variables we include in the model in order to avoid the *curse of dimensionality* problem that plagues nonparametric estimations. Still, in the robustness checks we contemplate the possibility of including some of these variables in the model. Finally, we add a measure of financial development to the model in order to test our hypothesis.

Regarding notation, despite the fact that we use panel data, in this section and in section 2 we present the formulas concerning the nonparametric models and techniques as if we were using cross-sectional data, i.e., we omit the time subscript, still the notation assumes that the panel dataset includes n cross-section units, with T time observations each, corresponding to a sample of nT observations. This is related to the way in which we deal with the fixed effects, which is by including a discrete variable *Country* in the equations.

The model of interest is a nonparametric version of a consumption function, given by:

$$\log C_i = g(\boldsymbol{x}_i) + u_i \tag{1}$$

$$\boldsymbol{x_i} = \begin{bmatrix}\log Y_i & r_i & \log W_i & FD_i\end{bmatrix} \tag{2}$$

where C is a measure of consumption, Y is a measure of current income, r is the real interest rate, FD is a measure of the level of financial development, W_i is a measure of wealth, $g(\cdot)$ is an unknown smooth function which we wish to estimate and u_i is the error term. Our choice of regressors to include in the consumption function follows some of the standard references like Mankiw (1982) or Ando and Modigliani (1963). We also include a proxy for financial development as a way to assess what role does it have in the consumption function.

We control for fixed effects by including a discrete unordered variable *Country*, which identifies the country to which the observation belongs. This is very similar to the dummy approach in the parametric setting. In the parametric case, allowing for fixed effects is tantamount to including individual dummies in the model. In the nonparametric case, one variable with a different value for each country (e.g., 1 for the first country, 2 for the second country, and so forth) is enough to control for these fixed effects. Thus, the nonparametric model allowing for fixed effects is:

$$\log C_i = g(\boldsymbol{x}_i, Country_i) + u_i \tag{3}$$

In the following sections we will be describing the methodologies used and to this effect we use the model in equation 1 to illustrate the procedures. Our description of the methodologies used here draws heavily on the presentations provided by Li and Racine (2007) and Henderson and Parmeter (2015).

2.2 Nonparametric Regression - Local Linear Least Squares

The logic behind the Local Linear Least Squares (LLLS) method is that we are minimizing the weighted squared distance between the dependent variable and a local linear approximation to the unknown function $g(\cdot)$ we wish to estimate. Taking our equation 1 as an example, expanding the function $g(\boldsymbol{x}_i)$ around $\boldsymbol{x}$ would give:

$$\begin{aligned} \log C_i &= g(\boldsymbol{x}_i) + u_i \\ &\approx g(\boldsymbol{x}) + (\boldsymbol{x}_i - \boldsymbol{x})\beta(\boldsymbol{x}) + u_i \end{aligned} \tag{4}$$

where $\beta(\boldsymbol{x})$ is the gradient vector of length q at point $\boldsymbol{x}$. From here we move on to the minimization problem:

$$\min_{g(\boldsymbol{x}),\beta(\boldsymbol{x})} \sum_{i=1}^{nT} \big[\log C_i - g(\boldsymbol{x}) - (\boldsymbol{x}_i - \boldsymbol{x})\beta(\boldsymbol{x})\big]^2 K_h(\boldsymbol{x}_i, \boldsymbol{x}) \tag{5}$$

where $K_h(\boldsymbol{x_i}, \boldsymbol{x})$ is a product-kernel weighting function. The product-kernel function takes the form:

$$K_h(\boldsymbol{x}_i, \boldsymbol{x}) = \prod_{d=1}^{q} k\Big(\frac{x_{i,d} - x_d}{h_d}\Big) \tag{6}$$

with k being the univariate kernel chosen to smooth each of the regressors, and the index d identifying the regressor in the model.

The solution, $\hat{\delta}(\boldsymbol{x})$, to this minimization problem can be written in matrix notation as:

$$\hat{\delta}(\boldsymbol{x}) = \begin{bmatrix} \hat{g}(\boldsymbol{x}) \\ \hat{\beta}(\boldsymbol{x}) \end{bmatrix} = \left[\boldsymbol{X}'K(\boldsymbol{x})\boldsymbol{X}\right]^{-1}\boldsymbol{X}'K(\boldsymbol{x})\log\boldsymbol{C} \tag{7}$$

In the above equation, $\log\boldsymbol{C}$ is the $nT \times 1$ vector of observations of the dependent variable, $\boldsymbol{X}$ is a $nT \times (q+1)$ matrix with the ith row given by $\boldsymbol{X}_i = [1, (\boldsymbol{x}_i - \boldsymbol{x})]$ and $K(\boldsymbol{x})$ a $nT \times nT$ diagonal matrix with the kernel functions $K_h(\boldsymbol{x}_i, \boldsymbol{x})$.

This can also be easily extended to include a discrete variable, namely the unordered variable *Country*. In this case the minimization problem is:

$$\min_{\delta(\boldsymbol{x}^c)} \Big[\log\boldsymbol{C} - \boldsymbol{X}\delta(\boldsymbol{x}^c)\Big]' W(\boldsymbol{x}^c)\Big[\log\boldsymbol{C} - \boldsymbol{X}\delta(\boldsymbol{x}^c)\Big] \tag{8}$$

where $\boldsymbol{x}^c = (\boldsymbol{x}, Country)$ and the matrix $\boldsymbol{X}$ is the same as in the case without discrete regressors. $W(\boldsymbol{x}^c)$ is an $nT \times nT$ diagonal matrix containing the product kernels $W(\boldsymbol{x}_i^c, \boldsymbol{x}^c)$ given by:

$$W(\boldsymbol{x}_i^c, \boldsymbol{x}^c) = K_h(\boldsymbol{x}_i, \boldsymbol{x})L_\lambda(Country_i, Country) \tag{9}$$

This product kernel—suggested by Li and Racine (2003)—admits both the case of the continuous regressors and of the unordered discrete variable *Country*. To smooth the variable *Country* we use the kernel suggested by Aitchison and Aitken (1976), which is given by:

$$L_\lambda(Country_i, Country) = \begin{cases} 1-\lambda & \text{if } Country_i = Country \\ \frac{\lambda}{P-1} & \text{if } Country_i \neq Country \end{cases} \tag{10}$$

where λ is the bandwidth for the discrete unordered variable and P is the number of different values the discrete unordered variable can take. From this problem we obtain the solution:

$$\hat{\delta}\big(\boldsymbol{x}^c\big) = \left[\boldsymbol{X}'W(\boldsymbol{x}^c)\boldsymbol{X}\right]^{-1}\boldsymbol{X}'W(\boldsymbol{x}^c)\log\boldsymbol{C} \tag{11}$$

2.3 Bandwidth Selection - AIC Cross Validation

We select the bandwidth using the methodology of Hurvich et al. (1998). This method is based on the Akaike Information Criterion (AIC, Akaike, 1998). It is a version of the AIC, modified so as to be more appropriate for model selection in the nonparametric setting. The formula for it is:

$$
\begin{aligned}
AIC_C(h) &= \ln(\hat{\sigma}^2) + \frac{1 + \frac{tr(\boldsymbol{H})}{nT}}{1 - \frac{tr(\boldsymbol{H})+2}{nT}} \\
\hat{\sigma}^2 &= \frac{1}{nT}\sum_{i=1}^{nT}\big[\log C_i - \hat{g}(\boldsymbol{x}_i)\big]^2
\end{aligned}
\tag{12}
$$

where the matrix $\boldsymbol{H}$ is given by:

$$
\boldsymbol{H} = \begin{bmatrix}
\frac{K(\boldsymbol{x}_1,\boldsymbol{x}_1)}{\sum_{l=1}^{nT} K(\boldsymbol{x}_1,\boldsymbol{x}_l)} & \frac{K(\boldsymbol{x}_1,\boldsymbol{x}_2)}{\sum_{l=1}^{nT} K(\boldsymbol{x}_1,\boldsymbol{x}_l)} & \cdots & \frac{K(\boldsymbol{x}_1,\boldsymbol{x}_n)}{\sum_{l=1}^{nT} K(\boldsymbol{x}_1,\boldsymbol{x}_l)} \\
\frac{K(\boldsymbol{x}_2,\boldsymbol{x}_1)}{\sum_{l=1}^{nT} K(\boldsymbol{x}_2,\boldsymbol{x}_l)} & \frac{K(\boldsymbol{x}_2,\boldsymbol{x}_2)}{\sum_{l=1}^{nT} K(\boldsymbol{x}_2,\boldsymbol{x}_l)} & \cdots & \frac{K(\boldsymbol{x}_2,\boldsymbol{x}_n)}{\sum_{l=1}^{nT} K(\boldsymbol{x}_2,\boldsymbol{x}_l)} \\
\vdots & \vdots & \ddots & \vdots \\
\frac{K(\boldsymbol{x}_n,\boldsymbol{x}_1)}{\sum_{l=1}^{nT} K(\boldsymbol{x}_n,\boldsymbol{x}_l)} & \frac{K(\boldsymbol{x}_n,\boldsymbol{x}_2)}{\sum_{l=1}^{nT} K(\boldsymbol{x}_n,\boldsymbol{x}_l)} & \cdots & \frac{K(\boldsymbol{x}_n,\boldsymbol{x}_n)}{\sum_{l=1}^{nT} K(\boldsymbol{x}_n,\boldsymbol{x}_l)}
\end{bmatrix}
\tag{13}
$$

2.4 Model Specification Test - Full Parametric against Full Nonparametric

Naturally, when employing a nonparametric approach, one is interested in knowing whether the additional effort is worth while. In other words, one is interested in testing whether the nonparametric approach represents an improvement over a parametric model. A specification test, of the null hypothesis of a full parametric specification against a full nonparametric specification, is described in Hsiao et al. (2007). This test has the form of a conditional moment test— similar to that in Zheng (1996)—where the null hypothesis is:

$$
H_0: \quad \mathrm{P}\big[\,\mathrm{E}(\log C_i|\boldsymbol{x}_i) = m(\boldsymbol{x}_i, \beta)\big] = 1 \tag{14}
$$

where $m(\cdot)$ is a known function, β is a vector of unknown parameters and $\boldsymbol{x}_i$ is a vector containing the values of the regressors corresponding to observation i. We want the null of the test to be the full parametric linear specification, so we set $m(\boldsymbol{x}_i, \beta) = \boldsymbol{x}_i'\beta$, i.e., the null we are interested in is:

$$
H_0: \quad \mathrm{P}\big[\,\mathrm{E}(\log C_i|\boldsymbol{x}_i) = \boldsymbol{x}_i'\beta\big] = 1 \tag{15}
$$

The test statistic is constructed through the sample analogue of:

$$
I \stackrel{\text{def}}{=} \mathrm{E}[u_i\,\mathrm{E}(u_i|\boldsymbol{x}_i)f(\boldsymbol{x}_i)] \tag{16}
$$

where $u_i = \log C_i - \boldsymbol{x}_i'\beta$. The sample analogue of this moment condition is:

$$\hat{I}_n = \frac{1}{nT}\sum_{i=1}^{nT}\hat{u}_i\hat{\mathrm{E}}_{-i}(u_i|\boldsymbol{x}_i)\hat{f}_{-i}(\boldsymbol{x}_i) = \frac{1}{(nT)^2}\sum_{i=1}^{nT}\sum_{\substack{j=1\\ j\neq i}}^{nT}\hat{u}_i\hat{u}_j\boldsymbol{W}(\boldsymbol{x}_i,\boldsymbol{x}_j) \qquad (17)$$

with the subscript $-i$ indicating that it refers to the leave-one-out estimator obtained when observation i is omitted. Additionally, $\boldsymbol{W}(\boldsymbol{x}_i,\boldsymbol{x}_j)$ is a generalized product kernel which allows for the presence of both continuous and discrete variables. The test statistic can then be normalized as follows:

$$\hat{J}_n \stackrel{\text{def}}{=} nT|\boldsymbol{h}|^{1/2}\hat{I}_n/\sqrt{\hat{\Omega}}$$
$$\hat{\Omega} = \frac{2|\boldsymbol{h}|}{(nT)^2}\sum_{i=1}^{nT}\sum_{\substack{j=1\\ j\neq i}}^{nT}\hat{u}_i^2\hat{u}_j^2\boldsymbol{W}^2(\boldsymbol{x}_i,\boldsymbol{x}_j) \qquad (18)$$

where $|\boldsymbol{h}|$ is the product of the bandwidths obtained via a cross-validation procedure for the nonparametric regression of $\log C$ on $\boldsymbol{x}$. The distribution for the test statistic can be obtained via two-point wild bootstrap.

2.5 Model Tests - Variable Significance Tests

Often, a paramount issue of interest after estimating a model is whether the explanatory variables are statistically significant. In the case of standard linear parametric models, the usual t-statistics will do the job. In nonparametric models we need to use special procedures to obtain similar tests. We will present the significance tests for nonparametric models in this section. The null hypothesis for the tests in this section is that the variables under test do not affect the dependent variable. We will continue to use the model given by equation 1 to illustrate the procedures.

2.5.1 Racine (1997) test for continuous variables

The first test we use is the one proposed by Racine (1997). The starting point for this test is the idea that if a variable is not relevant in explaining the dependent variable, then its marginal effect should be zero over its domain. Letting $\boldsymbol{X}$ be the vector of regressors and $\boldsymbol{X}_{(j)}$ the set of j regressors whose significance we are testing, the null hypothesis of this test can be expressed as:

$$H_0: \quad \mathbf{d} = \frac{\partial\,\mathrm{E}(\log C|\boldsymbol{X})}{\partial\boldsymbol{X}_{(j)}} = \mathbf{0}_j \quad \text{for all } \boldsymbol{x} \qquad (19)$$

Using an aggregate L_2 norm measure, the null hypothesis can be expressed as:

$$H_0: \quad \lambda = \mathrm{E}\left(\boldsymbol{\iota}'\boldsymbol{\Delta}\right) = 0 \tag{20}$$

where $\boldsymbol{\iota}$ is a vector of ones of length j and $\boldsymbol{\Delta}$ is a vector whose elements are the squares of the corresponding elements of $\mathbf{d}$ (the vector of derivatives with respect to the variables under test). The test statistic is constructed using a sample analogue of λ:

$$\hat{\lambda} = (nT)^{-1} \sum_{i=1}^{nT} \sum_{j} \left[\hat{\beta}_j(\boldsymbol{x}_i)\right]^2 \tag{21}$$

where $\hat{\boldsymbol{\beta}}(\boldsymbol{x}_i)$ is the estimated gradient vector at point $\boldsymbol{x}_i$, and the second summation sums the gradients corresponding to the j variables which significance we are testing. This test statistic is then pivotized in two ways. The first is by dividing the pointwise gradient estimates by the asymptotic approximations to their standard errors (SE), thus obtaining:

$$\hat{\lambda} = (nT)^{-1} \sum_{i=1}^{nT} \sum_{j} \left[\frac{\hat{\beta}_j(\boldsymbol{x}_i)}{SE(\hat{\beta}_h(\boldsymbol{x}_i))}\right]^2 \tag{22}$$

The second way in which the test is pivotized is by taking the $\hat{\lambda}$ from equation 22 and dividing this test statistic by an estimate of its standard error obtained via nested resampling. The pivotal test statistic is then:

$$\hat{t} = \frac{\hat{\lambda}}{SE(\hat{\lambda})} \tag{23}$$

The distribution of the test statistic is then obtained via bootstrap.

2.5.2 Racine et al. (2006) test for Discrete Variables

Now we describe the testing procedure proposed by Racine et al. (2006). Now the starting point is that if a discrete variable has no influence on the dependent variable, then the value it takes makes no difference to the value of the dependent variable. In describing this test we will be using the model given by equation 3 instead of of equation 1 as the former includes a discrete variable, *Country*. The null for this test can be stated as:

$$H_0: \quad \mathrm{E}(\log C|\boldsymbol{x}, Country) = \mathrm{E}(\log C|\boldsymbol{x}) \text{ almost everywhere} \tag{24}$$

where $\boldsymbol{x}$ is the vector of continuous regressors. Considering the way equation 3 is written, the null for this test can also be written as:

$$H_0: \quad g(\boldsymbol{x}_i, Country_i = l) = g(\boldsymbol{x}_i, Country_i = 1) \text{ almost everywhere} \\ \text{for } l = 2, \ldots, n \tag{25}$$

This suggests that the test be based on:

$$I = \sum_{l=2}^{n} \mathrm{E}\left([g(\boldsymbol{x}, \boldsymbol{Country} = l) - g(\boldsymbol{x}, \boldsymbol{Country} = 1)]^2 \right) \tag{26}$$

In fact, I is always non-negative and $I = 0$ if and only if H_0 is true. One can use the nonparametric estimate of $g(\cdot)$ to compute the test statistic:

$$\hat{I} = (nT)^{-1} \sum_{i=1}^{nT} \sum_{l=2}^{n} [\hat{g}(\boldsymbol{x}_i, Country_i = l) - \hat{g}(\boldsymbol{x}_i, Country_i = 1)]^2 \tag{27}$$

The distribution of the test statistic is then obtained using bootstrap.

3 Data

Our dataset is a panel of yearly observations for 46 countries from 2000 to 2014. For consumption we use total household consumption expenditures, for disposable income we use net disposable income and for the interest rate the real interest rate. Tables 1 to 3 report information about the coverage of the dataset, descriptive statistics and correlations.

The data for household consumption expenditures in constant 2010 prices in local currency was obtained from the National Accounts Main Aggregates Database (United Nations Statistics Division).

We retrieved the data for net disposable income in current local currency from AMECO for Belgium, Denmark, Germany, Ireland, Greece, Spain, France, Italy, Luxembourg, Netherlands, Austria, Portugal, Finland, Sweden, United Kingdom, Norway, Switzerland, United States, Japan, Canada, Mexico, South Korea, Australia, Bulgaria, Latvia, Lithuania, Poland, Czech Republic, Estonia, Hungary, Romania and Slovakia; from OECD for South Africa, New Zealand, Chile and Russia; and for all the other countries we obtained the data from UNdata (http://data.un.org/).

The data for each country's total wealth in current USD comes from Credit Suisse's Global Wealth Databooks.

Data on the interbank interest rates was collected from FRED for South Africa, Russia, New Zealand, Norway, Mexico, Israel and Iceland; from the

Table 1: List of countries

Armenia	Australia	Austria
Belgium	Bulgaria	Canada
Chile	Cyprus	Czech Republic
Denmark	Estonia	Finland
France	Germany	Greece
Hungary	Iceland	Ireland
Israel	Italy	Jamaica
Japan	Kuwait	Latvia
Lithuania	Luxembourg	Mexico
Netherlands	New Zealand	Norway
Poland	Portugal	Republic of Korea
Romania	Russian Federation	Slovakia
Slovenia	South Africa	Spain
Sweden	Switzerland	Thailand
Ukraine	United Kingdom	United States
Venezuela		

Table 2: Descriptive Statistics

Variables	N	mean	sd	min	max
r	690	0.0404	0.0381	-0.02	0.345
$\log C$	690	14.01	0.833	11.29	15.22
$\log Yd$	690	14.40	0.893	11.61	15.91
$\log W$	690	10.796	1.345	7.012	13.033
FD	690	0.583	0.23	0.098	1

Table 3: Cross-correlation table

	r	$\log C$	$\log Yd$	$\log W$	FD
r	1				
$\log C$	-0.487	1			
$\log Yd$	-0.487	0.982	1		
$\log W$	-0.430	0.956	0.946	1	
FD	-0.431	0.820	0.799	0.831	1

IMF's International Financial Statistics for Kuwait, South Korea, Thailand, Venezuela, Chile, Jamaica, Armenia, Australia, Canada and Switzerland; and from AMECO for all the other countries.

To obtain the real interest rate we employ the usual formula:

$$\text{real rate} = \frac{1 + \text{nominal rate}}{1 + \text{inflation rate}} - 1 \tag{28}$$

We use data on the consumer price index obtained from the World Bank for all countries except Chile and Venezuela (for which we resorted to FRED), to compute the inflation rate.

To obtain the net disposable income and total wealth in real terms, we calculated a consumption deflator using data on consumption expenditure in current prices and consumption expenditure at constant 2010 prices. To convert in USD we use the exchange rates from the National Accounts Main Aggregates Database.

We then compute per capita variables for consumption, disposable income and total wealth by dividing the real, USD measured, versions of each of these variables by each country's population. The data on population is retrieved from the National Accounts Main Aggregates Database for all countries.

The broad index of financial development which we use is the final aggregation level of a series of subindices computed in Svirydzenka (2016). The procedure starts by applying principal component analysis to a set of financial system data in order to obtain a measure of efficiency, a measure of access and a measure of depth for both the financial markets and the financial institutions of a given country. In the intermediate step, the measures of efficiency, access and depth are aggregated for financial markets and for financial institutions. This produces two indicators, one of the development level of the financial markets and the other of the development level of financial institutions. The last step aggregates these two indices in order to obtain the broad-based index of financial development, which reflects the development of the financial system of a country.

A last remark should be done regarding the dataset used in this paper. This paper is the result of a chapter from a PhD thesis for which the public funding is disclosed in the acknowledgments and from which a few other papers have also resulted some of which share the same dataset as the present paper, and as such a very similar data description section.

4 Results

In this section we report the results from the nonparametric regressions. We also report in appendix the results from a fully parametric model with the same regressors. We used a local linear regression technique, with a second-order Gaussian kernel for the continuous variables and the Aitchison and Aitken (1976) kernel for the unordered discrete variable. The bandwidths are chosen according to the Kullback-Leibler cross-validation criterion proposed by Hurvich et al. (1998). The results in this section were obtained using the resources available from the *np* package for R (Hayfield and Racine, 2008) and from the companion website for Henderson and Parmeter (2015).

4.1 Model Specification Tests

We start with the specification tests of the full parametric formulation against the full nonparametric alternative described in section 2.4. We performed the test on the following parametric models:

$$\log C_{i,t} = \beta_0 + \boldsymbol{x}'_{i,t}\beta + u_{i,t} \quad (a)$$

$$\log C_{i,t} = \beta_0 + \boldsymbol{x}'_{i,t}\beta + \sum_{i=2}^{n} \gamma_i D_i + u_{i,t} \quad (b)$$

where D_i is a dummy variable for country i, taking the value 1 if the observation belongs to that country and zero otherwise. The vector $\boldsymbol{x}_{i,t}$ includes all the regressors in equations 1 and 2. We also performed the specification test on the two models but without including financial development:

$$\log C_{i,t} = \alpha_0 + \alpha_1 \log C_{i,t-1} + \alpha_2 \log Y_{i,t} + \alpha_3 r_{i,t} + u_{i,t} \quad (c)$$

$$\log C_{i,t} = \alpha_0 + \alpha_1 \log C_{i,t-1} + \alpha_2 \log Y_{i,t} + \alpha_3 r_{i,t} + \sum_{i=2}^{n} \gamma_i D_i + u_{i,t} \quad (d)$$

The results reported below use the kernels and cross-validation procedures that we described previously. The results of the tests presented in table 4 show a strong rejection of all these parametric specifications in favor of a nonparametric approach.

Table 4: Specification Test: Parametric Vs Nonparametric

	$\hat{J}_n$	p-value $\hat{J}_n$	$\hat{I}_n$	p-value $\hat{I}_n$
(a)	24.182	$<$0.01	0.015	$<$0.01
(b)	3.548	$<$0.01	0.000	$<$0.01
(c)	9.592	$<$0.01	0.002	$<$0.01
(d)	4.069	$<$0.01	0.000	$<$0.01

4.2 Bandwidths

Next we report on the bandwidths we used in estimating our model. Additionally we report a rule-of-thumb upper bound for the bandwidths. For continuous variables, this upper bounds is equal to two times the standard deviation of the variable. This follows the suggestion from Hall et al. (2007). Hall et al. (2007) argue that a cross-validation method will select very large smoothing values (larger than a few standard deviations) for variables for that enter linearly in a model estimated by LLLS. For the discrete variable the upper bound is given by $(P-1)/P$. In this case, if the variable hits the upper bound this means that the variable is irrelevant in the model estimated by LLLS.

The bandwidths we obtained for the models based on equation 3, to be estimated, are reported in table 5, along with the upper bounds. We see that for wealth, country and financial development, the estimated bandwidths are below the upper bounds, providing evidence in favor of the hypothesis that these variables influence consumption in a nonlinear way. As for income and the interest rate the estimated bandwidth is above the upper bound in the local linear model which suggests that this variable enters the model in a way very close to linear.

Table 5: Estimated Bandwidths

	Bandwidth	Upper Bound
$\log Yd$	23840.92	1.786
r	7480.951	0.076
$\log W$	0.729	2.689
FD	0.096	0.460
$Country$	0.064	0.978

4.3 Variable Significance

Next we present the results of the nonparametric variable significance tests. Table 6 shows the p-values of these tests for each of the variables.

The null of this test is interpreted in the same way as the null of the traditional parametric significance test, meaning that a rejection of the null indicates that the variable has a statistically significant impact on the dependent variable. The results of the tests indicate that the variables included in the regression are all relevant in explaining consumption behavior despite wealth only having statistical significance at the 5% level.

Table 6: Significance Test

	Pivotal	Non-Pivotal
r	<0.01***	<0.01***
$\log Yd$	<0.01***	<0.01***
$\log W$	<0.01***	0.048**
FD	<0.01***	<0.01***
$Country$	<0.01***	<0.01***

4.4 Model Interpretation

We now take the results from the previous sections together with the plots in appendices B and C and attempt to interpret the overall results provided by our nonparametric model of consumption. We added vertical dotted lines to the plots that indicate the quartiles of the variable in the x-axis. We also added rugs to the plots showing the sample distribution of the variable in the axis near which they are represented. In the rug plot each tick indicates one observation meaning that the heavier the shade on the rug, the denser the distribution is around that area.

We produced two kinds of plots from our estimates: partial regression plots and gradient plots. The partial regression plot shows the estimated consumption function for a set of values of a given regressor, keeping all the other regressors at a specific value. This allows us to observe how consumption behaves in response to each of its regressors individually. Gradient plots show how the derivative of the estimated function with respect to a specific variable varies with that variable, while holding the other regressors constant. This means that the gradient plots show how the marginal effects of each regressor on consumption vary with that regressor. Note that across most of the plots we keep the interest rate fixed at the median value. The reason for this is to avoid the ambiguity from discussion on whether high values of the interest rate are associated with high values or low values of income and wealth as this would be outside the scope of our study.

From the plots we can easily observe that within our sample income has a linear and positive effect on consumption. This effect changes slightly with the value of the remaining regressors as evidenced by the slight change in the slope of the partial regression plot and the different values of the gradient.

The effect of the interest rate as shown in the plots is linear on the value of the interest rate itself while changing slightly with the value of other regressors. The confidence intervals in the gradient plots do include zero for the combinations of values we chose for the remaining regressors. While this may seem to contradict the significance test, we must note that the significance test takes into account the full spectrum of values the regressors might take.

The effect of wealth on the other hand shows a nonlinear behavior on both the value of wealth itself and the other variables. As we can see in the gradient plots the sign of this effect starts of negative for a combination of low values of wealth and the remaining regressors and becomes positive with higher values of the other variables.

Finally, the effect of financial development on consumption is also of the nonlinear kind. In this case, the sign of the effect seems to be determined essentially by the remaining regressors with it being negative for combinations of low income and low wealth and being positive otherwise. The effect however is decreasing with the value of financial development up to a relatively high threshold value after which the effect becomes stronger with increasing financial development.

As an additional exercise to study the effects of financial development over consumption we redo the partial regression plots and the gradient plots for the remaining regressors as functions of each variable, with the other regressors set at their means, except for financial development which we set at its first (Q1), second (Q2), third (Q3) and fourth (Q4) quartiles (Q4 is the maximum of the index of financial development), corresponding to each of the four panels in the before mentioned plots.

This analysis shows us that most of the effect of financial development is contained within its own gradient. The case where we see the most difference across all plots is when financial development goes from its Q1 to Q2 which appears to to trigger either a change in sign and behavior of the effect (as we can see in the gradient of wealth) or a change in the magnitude of the effect (as seen in the other regressors).

In sum, while we find evidence against the fully parametric specification in the statistical tests, our results do show that for some of the traditional explanatory variables of the consumption function, the link is very close to linear in the value of the variables themselves while leaving some space for potential interaction effects between these variables as seen when the effect changes for different combinations of regressor values. Financial innovation within our samples comes to play but mostly with a direct effect self-contained within its own gradient.

5 Robustness Checks

We also perform a series of brief robustness checks on our results to observe how sensitive they are to some of the methodological and data decisions that we took. The aspects we control for are:

1. Kernel - One of the methodological decisions that might have an impact over the results is the choice of kernel. The one we used, the Gaussian kernel, typically fares well in applied studies. Nonetheless, in this section we experiment with a different kernel for our estimations, the Epanechnikov kernel which is known to have better performance in terms of terms of efficiency. There are no noticeable differences in the pivotal version of the significance tests but for the non-pivotal version, financial development and wealth now fail to reject the null. The estimated bandwidths only show differences in the case of the interest rate which now is below the upper bound. As for the gradients and the partial regression plots while giving the same information they are now noisier.

2. Cross-Validation Method - The bandwidth selection procedure can have a deep impact on the results of a nonparametric estimation. In doing this we use a data-driven method, cross validation, specifically the one described in section 2.3. We experimented redoing the analysis but this time selecting the bandwidths by least squares cross validation. In the significance tests, financial development now fails to reject the null of both kinds of tests. The estimated bandwidths also have only minor differences, the exception being the one for the interest rate which is now below the upper bound. Regarding the plot analysis, there are some differences in the partials for financial development and wealth while keeping basically the same conclusions. As for the gradients for wealth and interest rate behave differently while having the same qualitative results.

3. Regression Technique - Our choice of methodology in terms of regression was the local linear least squares. Among the nonparametric local polynomial regressors, the two most widely used are the local linear and the local constant estimatos. Our preference for the local linear is associated with it having been found to have several advantages over the local constant methodology in applications such as ours. Still, we assess how different our results would be had we used the local constant estimation technique. First, the significance test for financial development now rejects the null in both the pivotal pivotal version

of the test. The bandwiths now are all below the upper bounds. The partial regression plots become less smooth while still giving similar results. The gradients on the other hand are significantly noisier.

What we observe from our robustness check analysis is that our results from the previous sections are fairly resilient to some of the pertinent changes one can introduce into the model. Specifically, our positive effect of financial development over consumption in situations where the other regressors are at high values is present in all the estimations.

6 Conclusion

The essential role consumption plays in macroeconomics makes it important that one understands what are the variables at play in the determination of its behavior. More so when economies have been deeply shaken by the last financial crisis. Most of the theories of consumption depend to some degree upon the possibility of smoothing consumption expenditures across the consumer's lifetime. This smoothing is mostly done through the financial system. Given the financial system's evolution in the past few decades, and the importance it appears to have gained in the economic structure, it is relevant to ask what impact may these changes have on consumption behavior.

In order to find evidence of links between financial development and consumption we used nonparametric regression techniques on consumption, a set of its typical regressors and a measure of financial development. The nonparametric approach allows us to search for nonobvious and nonlinear links that may exist between financial development and consumption, without having to specify a functional form for that link. Our results point to the existence of a nonlinear effect of financial development on consumption, an effect that varies with the value of the other determinants of consumption. In addition, our results provide evidence that while wealth has a nonlinear effect as well, the income and the interest rate's impact on consumption which are linear but with a possible interaction with the other variables in the consumption function.

Our findings help emphasize something that has become clearer since the most recent financial crisis. What happens within the financial sector does make its way through the rest of the economy which is noticeable from the effect we estimated and that the effect of changes to the financial system is dependent upon specific aspects of each given country.

References

Aitchison, J. and Aitken, C. G. (1976). Multivariate binary discrimination by the kernel method. *Biometrika*, 63(3):413–420.

Akaike, H. (1998). Information theory and an extension of the maximum likelihood principle. In Parzen, E., Tanabe, K., and Kitagawa, G., editors, *Selected Papers of Hirotugu Akaike*, pages 199–213. Springer New York, New York, NY.

Ando, A. and Modigliani, F. (1963). The "life cycle" hypothesis of saving: Aggregate implications and tests. *The American Economic Review*, 53(1):55–84.

Bandiera, O., Caprio, G., Honohan, P., and Schiantarelli, F. (2000). Does financial reform raise or reduce saving? *The Review of Economics and Statistics*, 82(2):239–263.

Bruno, G. (2014). Consumer confidence and consumption forecast: a non-parametric approach. *Empirica*, 41(1):37–52.

Carvalho, V. M., Martin, A., and Ventura, J. (2012). Understanding bubbly episodes. *The American Economic Review*, 102(3):95–100.

Cherchye, L., De Rock, B., and Vermeulen, F. (2007). The collective model of household consumption: a nonparametric characterization. *Econometrica*, 75(2):553–574.

Cherchye, L., De Rock, B., and Vermeulen, F. (2009). Opening the black box of intrahousehold decision making: Theory and nonparametric empirical tests of general collective consumption models. *Journal of Political Economy*, 117(6):1074–1104.

Deaton, A. (1992). *Understanding consumption*. Oxford University Press.

Delgado, M. A. and Miles, D. (1997). Household characteristics and consumption behaviour: A nonparametric approach. *Empirical Economics*, 22(3):409–429.

Dewachter, H. and Wouters, R. (2014). Endogenous risk in a dsge model with capital-constrained financial intermediaries. *Journal of Economic Dynamics and control*, 43:241–268.

Easaw, J. Z., Garratt, D., and Heravi, S. M. (2005). Does consumer sentiment accurately forecast uk household consumption? are there any comparisons to be made with the us? *Journal of Macroeconomics*, 27(3):517–532.

Estrada, Á., Garrote, D., Valdeolivas, E., and Vallés, J. (2014). Household debt and uncertainty: Private consumption after the great recession. Working Papers 1415, Banco de España.

Fisher, L. A., Huh, H.-s., and Otto, G. (2012). Structural cointegrated models of us consumption and wealth. *Journal of Macroeconomics*, 34(4):1111–1124.

Friedman, M. (1957). The permanent income hypothesis. In *A theory of the consumption function*, pages 20–37. Princeton University Press.

Guerrieri, V. and Lorenzoni, G. (2017). Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics*, 132(3):1427–1467.

Hall, P., Li, Q., and Racine, J. S. (2007). Nonparametric estimation of regression functions in the presence of irrelevant regressors. *The Review of Economics and Statistics*, 89(4):784–789.

Hayfield, T. and Racine, J. S. (2008). Nonparametric econometrics: The np package. *Journal of Statistical Software*, 27(5).

He, Z. and Krishnamurthy, A. (2014). A macroeconomic framework for quantifying systemic risk. Technical report, National Bureau of Economic Research.

Henderson, D. J. and Parmeter, C. F. (2015). *Applied nonparametric econometrics*. Cambridge University Press.

Hsiao, C., Li, Q., and Racine, J. S. (2007). A consistent model specification test with mixed discrete and continuous data. *Journal of Econometrics*, 140(2):802–826.

Hurvich, C. M., Simonoff, J. S., and Tsai, C.-L. (1998). Smoothing parameter selection in nonparametric regression using an improved akaike information criterion. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(2):271–293.

Lee, J. (2013). Consumption, financial wealth and labor income in Korea. *Japan and the World Economy*, 25:59–67.

Li, Q. and Racine, J. (2003). Nonparametric estimation of distributions with categorical and continuous data. *Journal of Multivariate Analysis*, 86(2):266–292.

Li, Q. and Racine, J. S. (2007). *Nonparametric econometrics: theory and practice*. Princeton University Press.

Mankiw, N. (1982). Hall's consumption hypothesis and durable goods. *Journal of Monetary Economics*, 10(3):417–425.

Martin, A. and Ventura, J. (2011). Theoretical notes on bubbles and the current crisis. *IMF Economic Review*, 59(1):6–40.

Racine, J. (1997). Consistent significance testing for nonparametric regression. *Journal of Business & Economic Statistics*, 15(3):369–378.

Racine, J. S., Hart, J., and Li, Q. (2006). Testing the significance of categorical predictor variables in nonparametric regression models. *Econometric Reviews*, 25(4):523–544.

Svirydzenka, K. (2016). Introducing a new broad-based index of financial development. IMF Working Papers 16/5, International Monetary Fund.

Swofford, J. L. and Whitney, G. A. (1987). Nonparametric tests of utility maximization and weak separability for consumption, leisure and money. *The Review of Economics and Statistics*, 69(3):458–464.

Zheng, J. X. (1996). A consistent test of functional form via nonparametric estimation techniques. *Journal of Econometrics*, 75(2):263–289.

A Parametric Model Results

This appendix presents the result from the full parametric version of our baseline model. The specification is as follows:

$$\log C_{i,t} = \mu_i + \alpha_1 r_{i,t} + \alpha_2 \log Yd_{i,t} + \alpha_3 \log W_{i,t} + \alpha_4 FD_{i,t} + u_{i,t} \tag{29}$$

Table 7: Parametric Model Estimations

	r	$\log Yd$	$\log W$	FD
Coefficients	-0.155	0.807***	-0.025	0.445***
	(0.143)	(0.077)	(0.024)	(0.156)
R^2	0.888			

robust standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

B Partial Regression Plots - LLLS

Figure 1: Partial Regression Plot, Financial Development on Consumption, country fixed effects LLLS

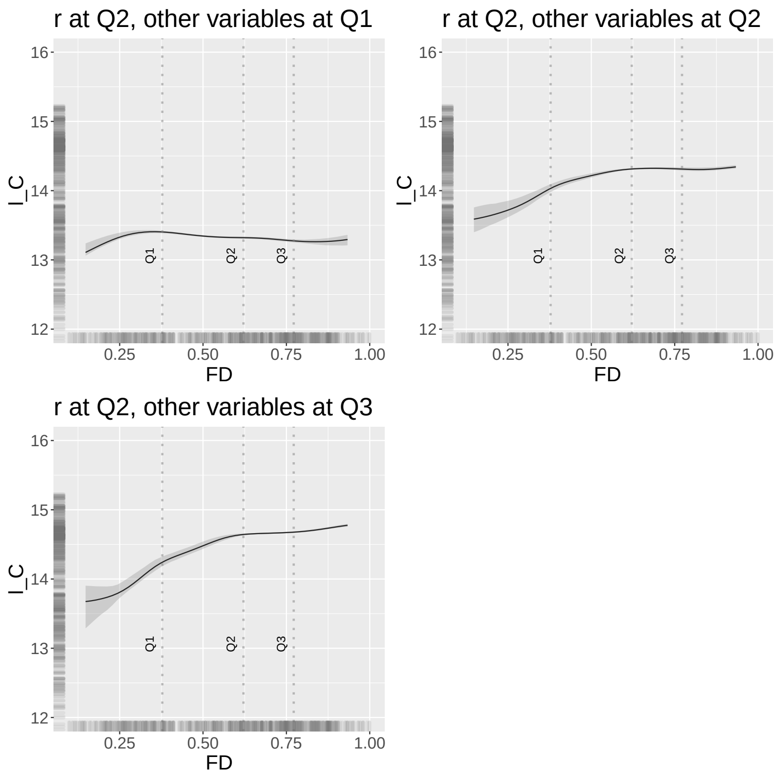

Figure 2: Partial Regression Plot, Wealth on Consumption, country fixed effects LLLS

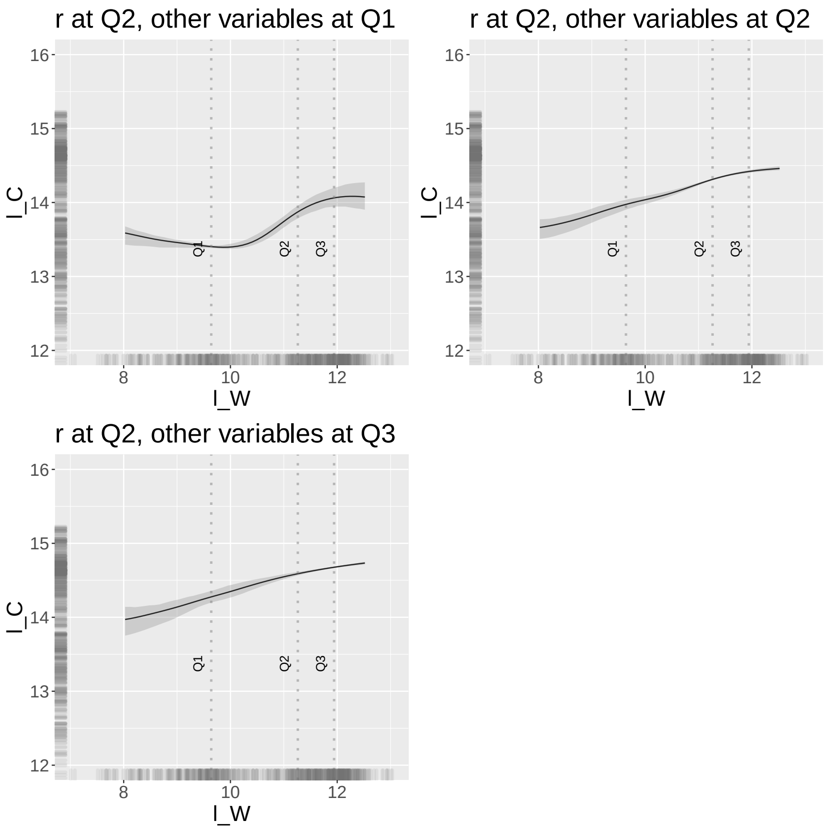

Figure 3: Partial Regression Plot, Income on Consumption, country fixed effects LLLS

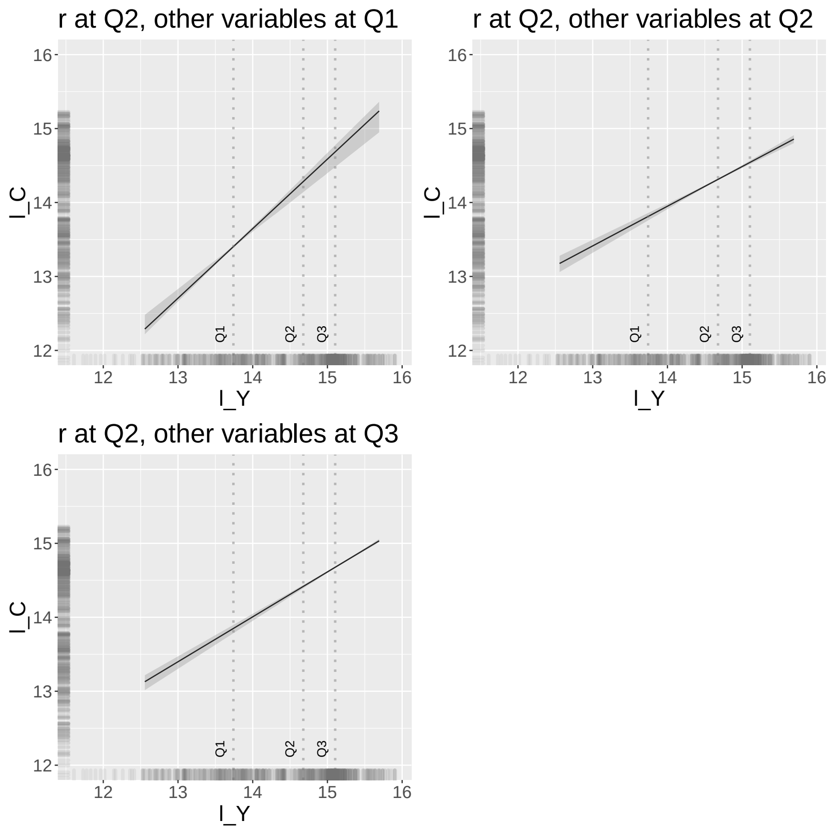

Figure 4: Partial Regression Plot, Interest rate on Consumption, country fixed effects LLLS

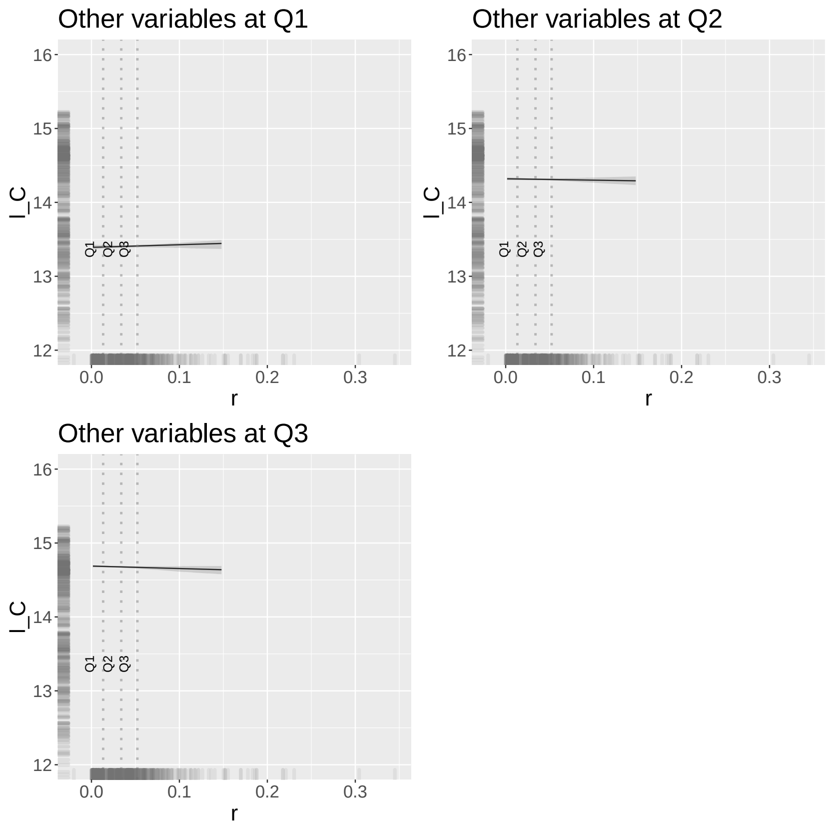

C Gradient Plots - LLLS

Figure 5: Gradient Plot, Financial Development on Consumption, country fixed effects LLLS

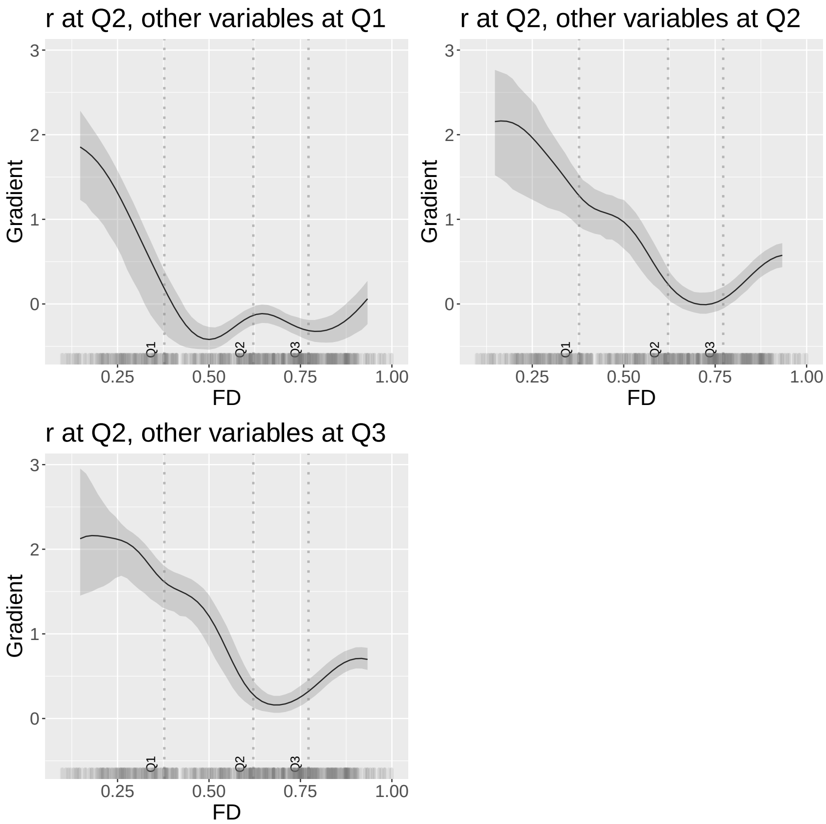

Figure 6: Gradient Plot, Wealth on Consumption, country fixed effects LLLS

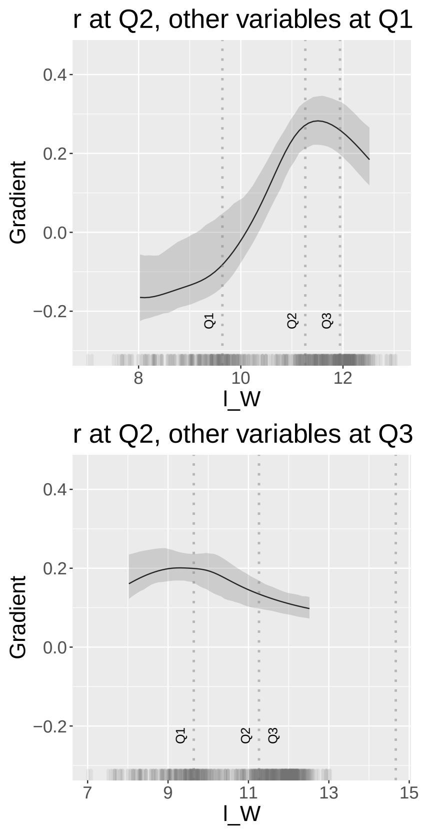

Figure 7: Gradient Plot, Income on Consumption, country fixed effects LLLS

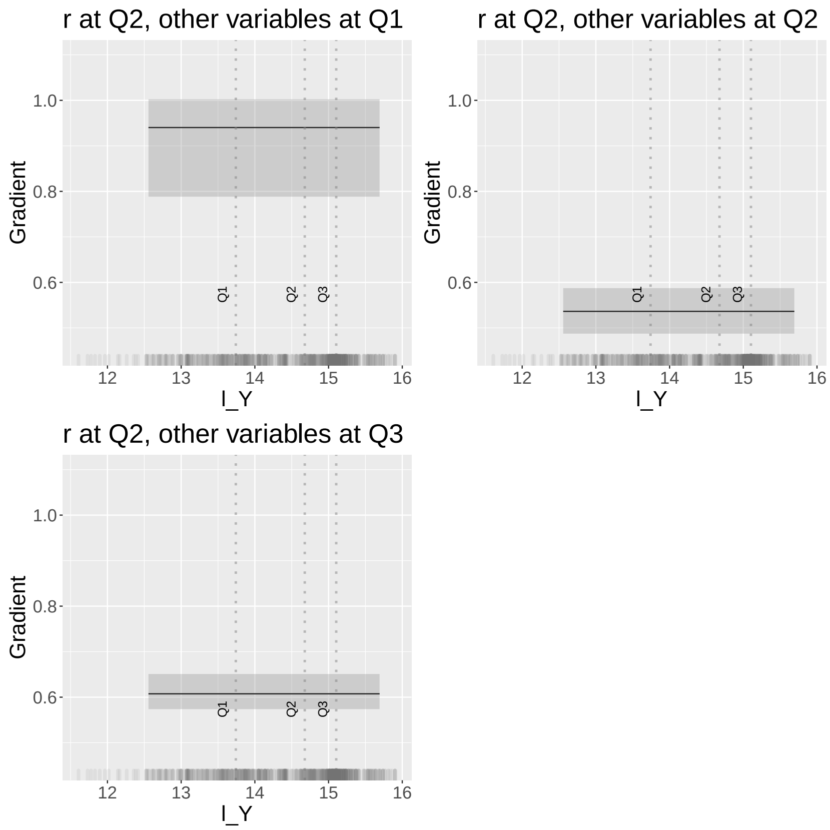

Figure 8: Gradient Plot, Interest rate on Consumption, country fixed effects LLLS

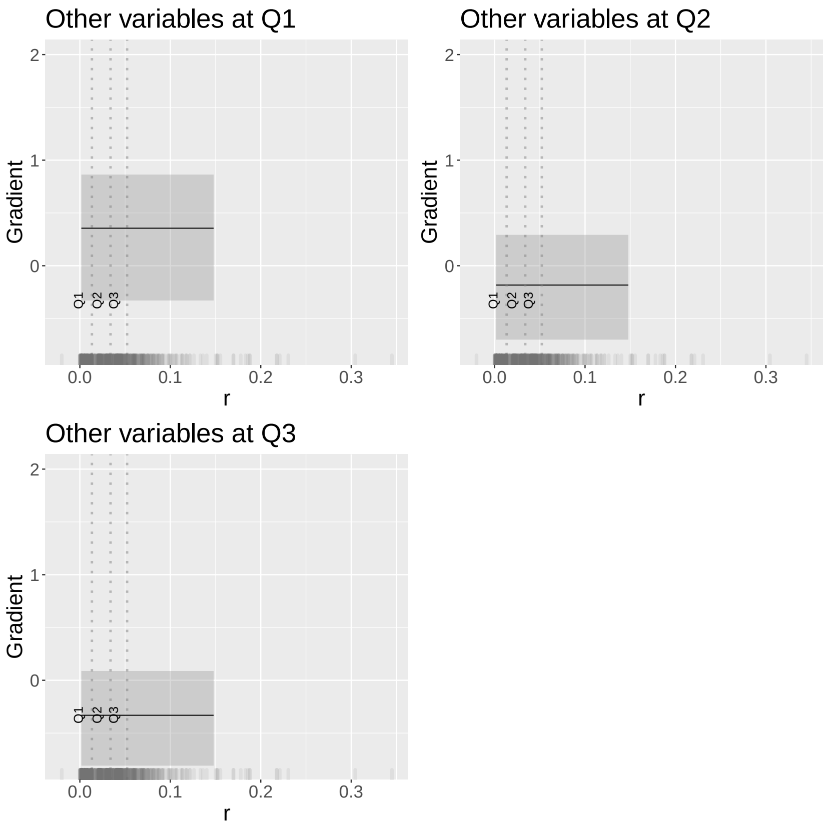

D Partial Regression Plots, by Financial Development Levels

Figure 9: Partial Regression Plot for different values of FD, Wealth, country fixed effects LLLS

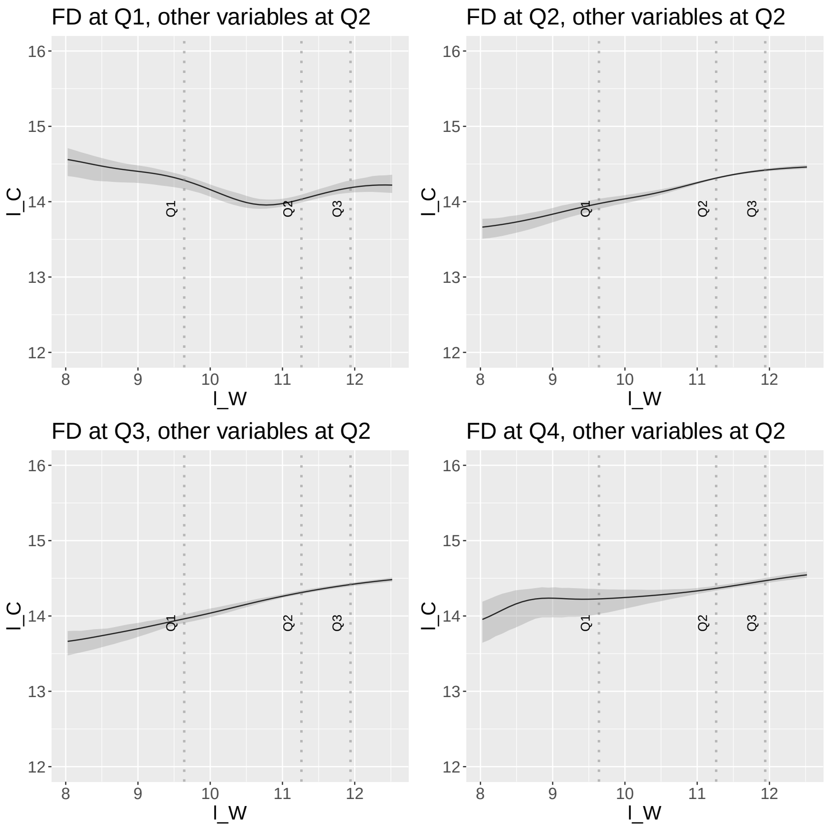

Figure 10: Partial Regression Plot for different values of FD, Income, country fixed effects LLLS

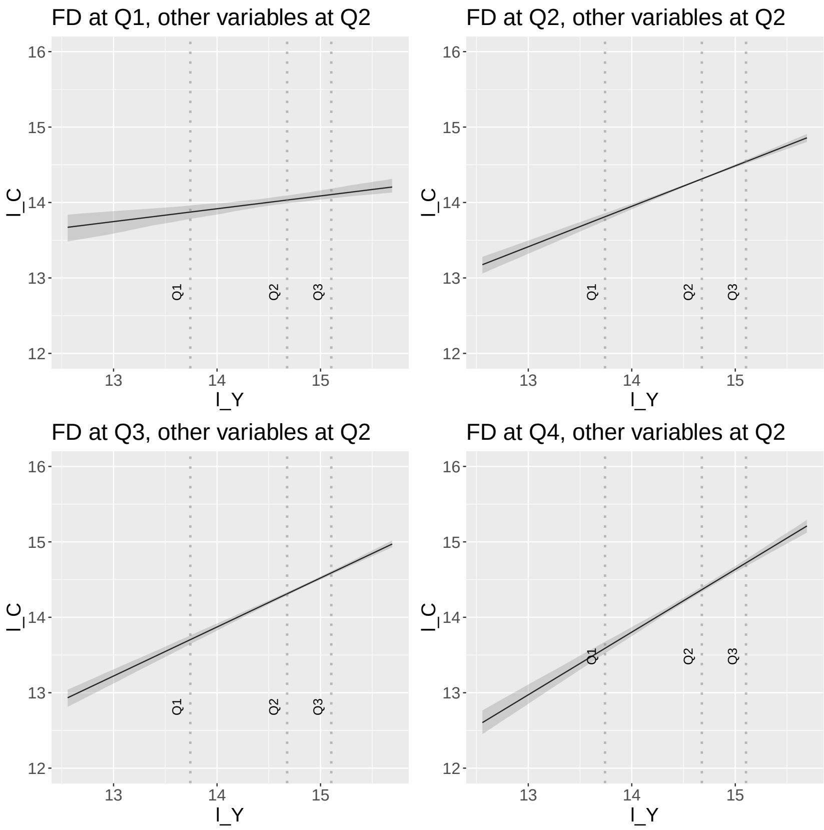

Figure 11: Partial Regression Plot for different values of FD, Interest rate, country fixed effects LLLS

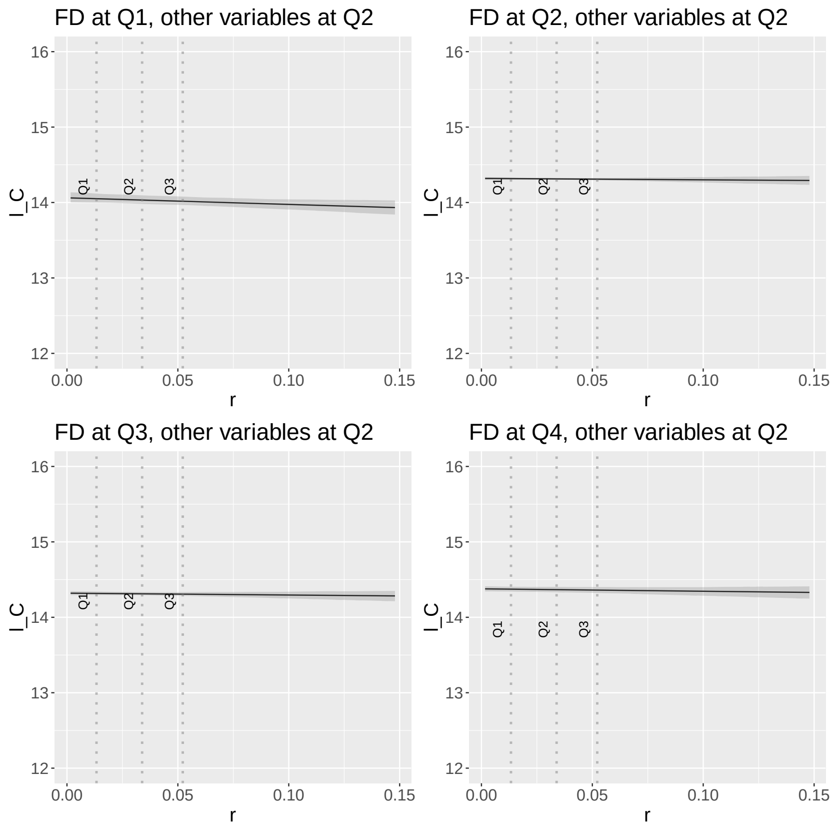

E Gradient Plots, by Financial Development Levels

Figure 12: Gradient Plot for different values of FD, Wealth, country fixed effects LLLS

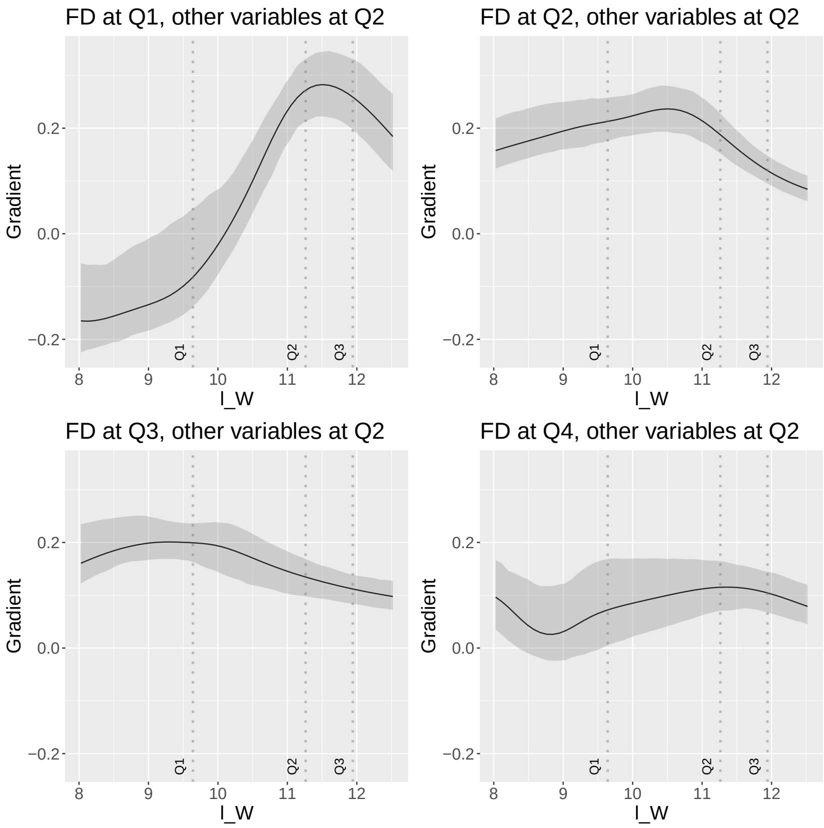

Figure 13: Gradient Plot for different values of FD, Income, country fixed effects LLLS

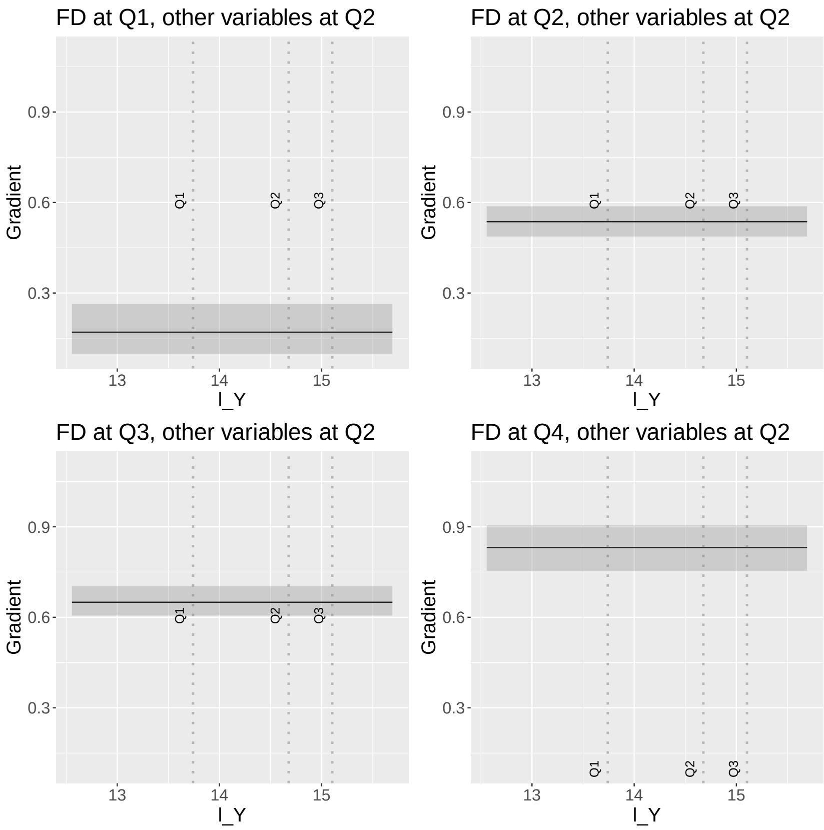

Figure 14: Gradient Plot for different values of FD, Interest rate, country fixed effects LLLS