

**Consumption smoothing at retirement:
average and quantile treatment effects in the
regression discontinuity design**

Daniel Burkhard

15-12

November 2015

DISCUSSION PAPERS

Consumption smoothing at retirement: average and quantile treatment effects in the regression discontinuity design*

Daniel Burkhard**

University of Bern

November 12, 2015

Abstract

Standard economic models predict that individuals smooth consumption over the life cycle. In contrast, there exists controversial empirical evidence showing that consumption declines at retirement. This paper investigates whether there is evidence for this so-called *Retirement Consumption Puzzle* in Switzerland. Baseline regression discontinuity estimates of average treatment effects are complemented by quantile treatment effects, where all estimates take the potential endogeneity of retirement into account. The findings suggest that disposable income significantly decreases after retirement, although there is substantial treatment effect heterogeneity. The reduction in income transmits to a negative but considerably less pronounced effect on overall consumption expenditures, indicating that households simultaneously adjust their savings. The results further show that food consumption at home is not or even positively affected by retirement, whereas expenditures in restaurants and hotels significantly decline.

JEL classification: C21, J14, J26

Keywords: Retirement Consumption Puzzle, Consumption Smoothing, Household Expenditure, Regression Discontinuity, Quantile Treatment Effect

*I am grateful to Michael Gerfin, Blaise Melly, Kaspar Wüthrich, and seminar participants at the University of Bern for helpful suggestions and comments. All errors are my own. Financial support from the Swiss National Science Foundation (Grant 100018-135379) is gratefully acknowledged.

**University of Bern, Department of Economics, Schanzeneckstrasse 1, CH-3001 Bern, Switzerland, email: daniel.burkhard@vwi.unibe.ch.

1 Introduction

Standard economic theory has recognized that one of the key motives for savings is consumption smoothing over the life cycle (Browning and Lusardi, 1996). Starting with the early work of Hamermesh (1984), who found that savings of couples in the United States are insufficient to maintain the previous level of consumption, researchers began to investigate whether individual and household consumption are continuous at the transition from the labor force to retirement. For the United States, Haider and Stephens (2007) and Bernheim et al. (2001) showed a drop in consumption at retirement that cannot be accounted for by standard life-cycle models. Similar results featuring the discontinuity at retirement, often entitled the *Retirement Consumption Puzzle*, have been found for the United Kingdom (Banks et al., 1998), for Germany (Schwerdt, 2005), and for Italy (Miniaci et al., 2010). Possible explanations for this non-stable path of consumption include unexpected liquidity shocks, time-inconsistent behavior, or changing preferences.

A second strand of the literature started asking the question whether consumption is measured in the appropriate way since the available data is typically observed on the household level and often focuses on food expenditures. Battistin et al. (2009) and Wakabayashi (2008), using data from Italy and Japan, respectively, also found evidence for the drop in consumption but further reported a reduction in the number of household members when the household head retires. This implies that lower household consumption does not necessarily lead to lower consumption of individuals. Other studies found no reduction when focusing on a broader definition of consumption than food expenditures, e.g., Aguiar and Hurst (2005), Fisher et al. (2008), and Aguila et al. (2011).¹

So far, research on the *Retirement Consumption Puzzle* has focused on average treatment effects (ATE), where *treatment* refers to the state of being retired. There is, however, strong reason to consider the impact of retirement along the distribution of consumption. To study poverty among the elderly, it is of key importance to know whether ATE are a good approximation to the impact on the lower tail of the distribution. To the best of my knowledge, the paper by Fisher and Marchand (2014) is the only exception that estimates quantile treatment effects (QTE) of retirement. However, their approach of considering retirement as exogenous is restrictive. For most countries, workers face incentives to retire before or after the ordinary retirement age and factors like the financial wealth are likely to have an impact on the timing of retirement. In fact, many studies that estimate average effects of retirement follow an *instrumental variable* or *regression*

¹Further criticism stressed the role of expectations about the timing of retirement and expected consumption levels thereafter. Ameriks et al. (2007) showed that the expected reduction tends to be larger than the actual decline at retirement whereas Smith (2006) only found a drop for workers who involuntarily retire. For a general overview, see Hurst (2008).

discontinuity approach that allows for endogenous retirement (see, for instance, Battistin et al., 2009; Stancanelli and Van Soest, 2012).

This paper explores the effect of retirement on disposable income and aggregate consumption expenditures, but also on disaggregated measures of consumption including food expenditures, food consumption in kilograms, spending on restaurants and hotels, transportation expenditures, expenditures related to housing, and healthcare expenditures. The data on elderly households is provided by the Swiss Household Budget Survey (HBS) and the empirical model exploits the discontinuity in the retirement rate induced by the legal retirement age within a fuzzy regression discontinuity design. The resulting estimates of local ATE are complemented by QTE that are estimated using the approach proposed by Frandsen et al. (2012). To the best of my knowledge, the present paper is the first to estimate QTE of endogenous retirement. Moreover, it is the first to study the *Retirement Consumption Puzzle* for Switzerland.

The results show that retirement decreases disposable income, on average, by 2400 Swiss Francs per month. In percentage terms this amounts to a drop of roughly -35% compared to non-retired households. The effect, however, is smaller for a large part of the outcome distribution and gets more negative towards the upper tail, indicating substantial treatment effect heterogeneity. There is a less pronounced effect of approximately -9% on aggregate consumption expenditures, which reveals that households simultaneously adjust their savings. Retirement is found to have no or even weakly positive effects on food consumption at home. There is, however, a negative impact on expenditures in restaurants and hotels. Additionally, my results provide evidence for lower spendings related to housing, whereas healthcare expenditures are unaffected by retirement.

The structure of the paper is as follows: Section 2 describes the institutional background for Switzerland. Section 3 discusses the regression discontinuity estimators applied in the present paper. The data are introduced in section 4. Section 5 addresses the validity of the regression discontinuity design and presents the results for the average and quantile treatment effects of retirement. Conclusions are presented in section 6. All figures and tables are collected in the appendix.

2 Institutional background

The Swiss pension system is classified in three complementary institutions that are called pillars. Besides the compulsory social insurance for the entire country (first pillar) and mandatory occupational pensions (second pillar), individuals can voluntarily save for

retirement by contributing to a private pension plan (third pillar).² While the first pillar only provides a relatively small monthly transfer, the goal of the second pillar is to maintain the previous standard of living. For all three pillars, the legal retirement age is 65 for men, and 64 for women. While the retirement age of men has been stable since 1948, it has been revised several times for women. Most recently, it has been increased from 63 to 64 in 2005.

The legal retirement age, however, is not strictly enforced. It is allowed to take early retirement or to postpone retirement by at most five years. The first pillar allows for at most two years of early retirement, which is financed by a reduction of the annual payments of 13.6%. If retirement is taken one year before the legal retirement age, the annual payments are 6.8% lower. People who postpone retirement receive payments that increase by approximately five percentage points annually compared to the ordinary pension. For the second pillar, the earliest possible retirement age is 58. Early retirement also leads to lower benefits, but the extent of the reduction differs by insurer. Furthermore, a woman (man) who retires at age 58 does not receive payments from the first-pillar social insurance before she (he) is aged 62 (63) and, therefore, becomes eligible for first-pillar retirement.

The pension benefits of the first and the second pillar both depend, to some extent, on the amount that has been contributed during the work-life. Both systems are financed by payroll deductions that are under federal regulation, but the structures of the two systems differ. The first pillar consist of a pay-as-you-go system that induces substantial redistribution from high to low income groups since the benefits are capped at a relatively low level. Additionally, the payment to married couples is limited to 150% of the maximum pension of an individual. In contrast, the second pillar as well as the third pillar are fully-funded pension systems. Therefore, the benefits are solely determined by the amount of individual contributions and paid on the individual level.

3 Empirical model

Let Y_i denote an outcome variable, e.g., consumption expenditures, for household i . Under the assumption of exogenous retirement, it would be straightforward to regress Y_i on a binary retirement indicator R_i and some vector of covariates. This assumption, however, is not likely to hold. Although there exists a legal retirement age, it is not binding and a considerable fraction decides to retire at a different age. Unobserved factors as financial assets are likely to have direct effects on the outcomes of interest as

²There is a financial incentive for private savings since contributions to the third pillar reduce income taxes.

well as the timing of retirement. A convenient way to handle the endogeneity problem is to exploit the discontinuity in the retirement rate induced by the legal retirement age within a regression discontinuity (RD) design. The RD approach was introduced in the evaluation literature by Thistlethwaite and Campbell (1960) and recently summarized by Imbens and Lemieux (2008) as well as Lee and Lemieux (2010). The conditions to identify the QTE are adapted from Frandsen et al. (2012).

3.1 Regression discontinuity design

Figure 1 plots the fraction of retirees against age, where the variable age is normalized such that the value zero corresponds to the legal retirement age. At the legal retirement age, the retirement rate discontinuously increases by approximately 30 percentage points. To formalize this discontinuity, let the running variable X_i be the age of an individual minus the legal retirement age. Additionally, let x_0 denote the point where $X_i = 0$. Then the probability of retirement can be expressed as

$$\Pr(R_i = 1|X_i) = \begin{cases} g_1(X_i) & \text{if } X_i \geq x_0 \\ g_0(X_i) & \text{if } X_i < x_0 \end{cases}, \text{ where } g_1(x_0) \neq g_0(x_0). \quad (1)$$

The jump in the probability of retirement at x_0 is embedded in assumption 1.

Assumption 1. $\lim_{X_i \downarrow x_0} \Pr(R_i = 1|X_i = x_0) > \lim_{X_i \uparrow x_0} \Pr(R_i = 1|X_i = x_0)$.

It is further useful to assume that the response to exceeding the threshold x_0 is monotone (assumption 2), which rules out observations changing their status from retired to non-retired at the legal retirement age.

Assumption 2. $\lim_{x \rightarrow x_0} \Pr(R_i^1 \geq R_i^0|X_i = x) = 1$, where $R_i^0 = \lim_{X_i \uparrow x_0} R_i(X_i)$ and $R_i^1 = \lim_{X_i \downarrow x_0} R_i(X_i)$.

Using the functions $g_1(X_i)$ and $g_0(X_i)$, we can rewrite equation (1) as

$$\mathbb{E}[R_i|X_i] = \Pr(R_i = 1|X_i) = g_0(X_i) + [g_1(X_i) - g_0(X_i)]T_i, \quad (2)$$

where

$$T_i = \mathbb{1}(X_i \geq x_0).$$

One needs to construct suitable estimates for $g_1(X_i)$ and $g_0(X_i)$, locally at x_0 , since the

difference between the two can be used to estimate the *first-stage* effect of T_i on R_i :

$$\rho_{fs} = \lim_{X_i \downarrow x_0} \mathbb{E}[R_i|X_i = x_0] - \lim_{X_i \uparrow x_0} \mathbb{E}[R_i|X_i = x_0]. \quad (3)$$

Note that there are no observations precisely at x_0 with $T_i = 0$, but we may approximate that point using $g_0(X_i)$ with X_i arbitrarily close to x_0 .

Similar to the retirement rate, there may be a discontinuity in the dependent variable Y_i at x_0 . If all factors that have an impact on the outcome, except for the retirement rate, stay constant or change in a continuous way, then a discontinuity in the dependent variable can be attributed to the discontinuous increase in the retirement rate.³ This condition is formally imposed by assumption 3, where F_{Y^1} and F_{Y^0} are the distribution functions of the potential outcomes with treatment (Y^1) and without treatment (Y^0). Note that assumption 3 is stronger than required for identification of the ATE but essential for the QTE.

Assumption 3. $F_{Y^r|R^0, R^1, X}(y|r^0, r^1, x)$ is continuous in x at x_0 , for $r^0, r^1 \in \{0, 1\}$. $\mathbb{E}[R^z|X = x]$ is continuous at x_0 , for $z \in \{0, 1\}$.

Finally, assumption 4 requires that observations close to x_0 exist.

Assumption 4. $F_X(x)$ is differentiable at x_0 and $\lim_{x \rightarrow x_0} f_X(x) > 0$.

For the ATE, it is useful to model Y_i by two separate functions $f_1(X_i)$ and $f_0(X_i)$ on either side of the threshold x_0 :

$$\mathbb{E}[Y_i|X_i] = \begin{cases} f_1(X_i) & \text{if } X_i \geq x_0 \\ f_0(X_i) & \text{if } X_i < x_0 \end{cases} \quad (4)$$

The difference in the mean outcomes between the two functions at x_0 is the *reduced-form* effect of T_i on Y_i . It is defined as

$$\rho_{rf} = \lim_{X_i \downarrow x_0} \mathbb{E}[Y_i|X_i = x_0] - \lim_{X_i \uparrow x_0} \mathbb{E}[Y_i|X_i = x_0]. \quad (5)$$

The *reduced form* divided by the *first stage* identifies the average causal effect ρ of retirement R_i on the outcome variable Y_i ,

$$\rho = \frac{\rho_{rf}}{\rho_{fs}} = \frac{\lim_{X_i \downarrow x_0} \mathbb{E}[Y_i|X_i = x_0] - \lim_{X_i \uparrow x_0} \mathbb{E}[Y_i|X_i = x_0]}{\lim_{X_i \downarrow x_0} \mathbb{E}[R_i|X_i = x_0] - \lim_{X_i \uparrow x_0} \mathbb{E}[R_i|X_i = x_0]}. \quad (6)$$

³Section 5.1 addresses the continuity of other factors that potentially have an impact on Y_i .

The functions f_1 , f_0 , g_1 , and g_0 can be estimated nonparametrically. Under parametric assumptions, the retirement effect ρ is estimated by two-stage least squares (2SLS).

The causal effect ρ is, without assuming treatment effects homogeneity, an average treatment effect that is local in two dimensions. It is the average treatment effect for compliers, i.e., individuals that retire because they reach the legal retirement age, identified locally at x_0 .

3.2 Average treatment effects

3.2.1 Parametric estimation

In general, the ATE can be estimated by parametric models with *age*-polynomials of (low) order p to approximate the f - and g -functions. To cover the *first-stage* effect, the parametrizations $g_0(X_i) = \gamma_0 + \sum_{j=1}^p \gamma_j X_i^j$ and $g_1(X_i) = \gamma_0 + \pi T_i + \sum_{j=1}^p \gamma_{Tj} T_i X_i^j$ are applied. The difference between the two functions at x_0 , i.e., at $X_i = 0$, is the parametric *first-stage* effect π . Adding the analogous parametrizations for $f_1(X_i)$ and $f_0(X_i)$, we can write the *first-stage* and the *reduced-form* equations as

$$R_i = \gamma_0 + \pi T_i + \sum_{j=1}^p \gamma_j X_i^j + \sum_{j=1}^p \gamma_{Tj} T_i X_i^j + u_i \quad (7)$$

and

$$Y_i = \alpha_0 + \lambda T_i + \sum_{j=1}^p \alpha_j X_i^j + \sum_{j=1}^p \alpha_{Tj} T_i X_i^j + v_i. \quad (8)$$

Note that this specification allows for different intercepts and curvatures on either side of the cutoff x_0 . In principle, one could estimate equations (7) and (8) separately by ordinary least squares and calculate the parametric point estimate of the causal effect in a second step by the fraction $\hat{\rho}_p = \hat{\lambda}/\hat{\pi}$. This estimator is numerically equivalent to a 2SLS regression of the dependent variable Y_i on constant, R_i , $\sum_{j=1}^p X_i^j$, and $\sum_{j=1}^p T_i X_i^j$, where T_i itself is used as the instrument for the endogenous R_i . The estimator based on 2SLS is more convenient than two separate regressions because the 2SLS standard errors are appropriate.⁴

Gelman and Imbens (2014) show that polynomials of higher order (second, third, or higher) are likely to have poor properties (undesirable weights, sensitivity of the estimate to the order of the polynomial, and poor inference). Therefore, this paper applies linear splines on both sides of the cutoff and additionally considers nonparametric estimates

⁴This is not immediately the case for the manual calculation using $\hat{\rho}_p = \hat{\lambda}/\hat{\pi}$.

obtained by local linear regression that are not subject to these concerns.

3.2.2 Potential rounding errors

The identification strategy of this paper exploits the discontinuity in the retirement rate that occurs at the legal retirement age. Ideally, we would like to know the age with high precision, for instance in months or even in days. In the available data, however, age is measured in years which can be seen as the *true* age rounded down to the nearest integer. Dong (2015) argues that this potentially induces a rounding error in the estimate of the retirement effect but shows that the bias can be corrected under a relatively mild assumption. For the present paper, it is sufficient to assume that birthdays are uniformly distributed over the year. The bias-corrected estimate for functions with first-order polynomials is $\hat{\rho}_{bc} = (\hat{\lambda} - \frac{1}{2}\hat{\alpha}_{T1})/(\hat{\pi} - \frac{1}{2}\hat{\gamma}_{T1})$, where $\hat{\lambda}$, $\hat{\alpha}$, $\hat{\pi}$, and $\hat{\gamma}$ are the estimates of the corresponding parameters in equations (7) and (8).⁵ The resulting $\hat{\rho}_{bc}$ cannot be estimated by 2SLS and is a nonlinear combination of the estimates of two different regressions. Therefore, the coefficients of the two regressions are combined by seemingly unrelated estimation to perform inference.

3.2.3 Nonparametric estimation

The estimates based on parametric functions are supplemented by a nonparametric analysis of the retirement effect. Nonparametric estimation of the functions f_1 , f_0 , g_1 , and g_0 identifies the reduced form (ρ_{rf}) as well as first stage (ρ_{fs}) effect and thereby allows us to construct a local Wald estimate of the causal effect, as in equation (6), that does not rely on parametric assumptions.

The functions on both sides of the cutoff x_0 are estimated by local linear regression using the triangular kernel. This procedure shrinks the weight of the observations that are away from the cutoff and thereby reduces their impact on the estimates of the discontinuities. The bandwidth is chosen optimally using the data-driven approach by Imbens and Kalyanaraman (2012), which is specific to the RD framework. The sensitivity of the estimates to the bandwidth choice is evaluated by reporting the effects that are found when using bandwidths of half and doubled size.

3.3 Quantile treatment effects

As the impact of retirement is potentially different in different parts of the outcome distributions, the ATE are supplemented with quantile treatment effects to examine such

⁵Dong (2015) provides a general expression for the bias correction with polynomial functions of higher order.

treatment effect heterogeneity. In particular, the QTE characterize the impact on the entire outcome distributions and thereby allow for a more complete evaluation of the effects on consumption of elderly households.

I apply the methodology proposed by Frandsen et al. (2012). Their procedure is suitable to nonparametrically estimate unconditional QTE within the RD framework. The estimand is a local QTE, defined as

$$\delta(\tau) = Q_{Y^1|C,X=x_0}(\tau) - Q_{Y^0|C,X=x_0}(\tau), \quad (9)$$

where $Q_{Y|C,X=x_0}(\tau)$ denotes the τ quantile of Y for compliers, locally at $X = x_0$. The QTE is an effect on the distribution that does not represent the effect on a specific individual or household without further assumptions.⁶ The estimator exploits the fact that quantiles of interest can be expressed in terms of the baseline distribution function by $Q_{Y|C,X=x_0}(\tau) = \inf\{u : F_{Y|C,X=x_0}(u) \geq \tau\}$. Using this equality, equation (9) reads as:

$$\delta(\tau) = \inf\{u : F_{Y^1|C,X=x_0}(u) \geq \tau\} - \inf\{u : F_{Y^0|C,X=x_0}(u) \geq \tau\}. \quad (10)$$

Frandsen et al. (2012) provide the expressions for $F_{Y^1|C,X=x_0}(y)$ and $F_{Y^0|C,X=x_0}(y)$, i.e., the distribution functions of the potential outcomes for compliers at x_0 . In particular, they show that

$$F_{Y^1|C,X=x_0}(y) = \frac{\lim_{X_i \downarrow x_0} \mathbb{E}[\mathbb{1}(Y_i \leq y)R_i|X_i = x_0] - \lim_{X_i \uparrow x_0} \mathbb{E}[\mathbb{1}(Y_i \leq y)R_i|X_i = x_0]}{\lim_{X_i \downarrow x_0} \mathbb{E}[R_i|X_i = x_0] - \lim_{X_i \uparrow x_0} \mathbb{E}[R_i|X_i = x_0]}$$

$$F_{Y^0|C,X=x_0}(y) = \frac{\lim_{X_i \downarrow x_0} \mathbb{E}[\mathbb{1}(Y_i \leq y)(1 - R_i)|X_i = x_0] - \lim_{X_i \uparrow x_0} \mathbb{E}[\mathbb{1}(Y_i \leq y)(1 - R_i)|X_i = x_0]}{\lim_{X_i \downarrow x_0} \mathbb{E}[1 - R_i|X_i = x_0] - \lim_{X_i \uparrow x_0} \mathbb{E}[1 - R_i|X_i = x_0]}$$

and point out that the sample analogs, $\hat{F}_{Y^1|C,X=x_0}(y)$ and $\hat{F}_{Y^0|C,X=x_0}(y)$, can be estimated by local linear weighted two-stage least squares.

Analogous to the nonparametric estimates of the ATE, the triangular kernel and the bandwidth suggested by Imbens and Kalyanaraman (2012) are applied to specify the weights for estimating the QTE.⁷ The QTE are estimated over a grid of $\{0.1, 0.15, \dots,$

⁶In particular, this involves a rank invariance assumption.

⁷Similar to Yu and Jones (1998), Frandsen et al. (2012) suggest to specify the bandwidths used for estimation of the QTE in terms of a reference bandwidth h_{mean} that is optimal for a mean (e.g., Imbens and Kalyanaraman, 2012). This approach, however, makes the estimation of the QTE computationally more demanding because the bandwidths differ by quantile and the distribution functions $F_{Y^1|C,X=x_0}(y)$ and $F_{Y^0|C,X=x_0}(y)$ therefore have to be estimated several times. In practice, the optimal bandwidths for estimating the QTE are similar to those for a mean. At the median, the optimal bandwidth is $h_{\tau=0.5} \approx$

0.9} and the standard errors are obtained by the bootstrap.

4 Data

4.1 Data source and sample selection

This paper explores pooled cross-section data of the years 2006 to 2011 from the HBS, conducted by the Swiss Federal Statistical Office. In this dataset, the unit of observation is the household. For key personal attributes, however, additional individual data is available. This includes, most importantly, the age and the labor market status of each household member. Besides data on individuals, the HBS provides comprehensive income and expenditure data. The sample used to estimate the causal effects of retirement is restricted to households where the household head's age is inside the interval of ± 15 years around the legal retirement age.⁸ Note that the household head is defined as the person in the household with the highest income and can therefore be male or female. Households are considered to be retired when the household head is retired and to be non-retired when the household head is active in the labor market. In the period under study, the legal retirement age was 65 for male workers, whereas it was 64 for females. Therefore, the sample includes households with male heads aged 50 to 80 and households with female heads that are between 49 and 79 years old.

Additionally, I exclude households living with individuals that are not part of the same family, very large households with more than seven members, and couples with an age difference of more than 20 years. These restrictions protect against potential distortions arising from economic activities of individuals belonging to other families and large age differences. After sample selection, there are 8524 households in the data, of which 4444 are couples, 2590 are individuals living alone, and 1490 are households with three or more members.

4.2 Descriptive statistics

Table 1 provides summary statistics for all outcome variables, the running variable, the retirement indicator, and supplementary covariates. The mean of overall consumption expenditures of non-retired households is 5816 Swiss Francs (CHF) per month. Retired households only spend 4296 CHF on average, where the difference in means is significant

$1.09 \times h_{mean}$, whereas the bandwidths at the 10th and the 90th percentiles are $h_{\tau=0.1,0.9} \approx 1.24 \times h_{mean}$.

⁸The approach of this paper is to use a sufficiently large window combined with flexible parametric models and nonparametric estimators that choose the bandwidth based on some optimization criterion. Note that it is not uncommon to specify a smaller age window of ± 10 years around pension eligibility. For recent examples see, e.g., Battistin et al. (2009) or Stancanelli and Van Soest (2012).

on the 1% level. Most of the disaggregated consumption expenditures, including food consumed at home, spending on restaurants and hotels, transportation (i.e., gasoline, public transport, etc.), and housing (rent, mortgage interest, etc.) are also significantly lower on the 1% level for retired households. Moreover, the average amount of food consumed at home (in kilograms) is lower for retired than for non-retired households, but healthcare expenditures (excluding health insurance premiums) are higher for retired households. Both of these differences are significant on the 1% level. Table 1 further shows that disposable income, defined as gross income minus taxes and compulsory health insurance and social insurance premiums, is on average 2959 CHF lower (−40%) for retired households.⁹

The differences in means, however, cannot be taken as evidence for effects of retirement as retired household heads are on average distinctively older. The difference in mean age between the two groups is thirteen years. As expected, most retired household heads have reached the legal retirement age (86%), whereas only 11% of the non-retired have crossed this threshold. The fraction of female household heads is eleven percentage points larger for retired (41%) than for non-retired households (30%). Retirees live, on average, more often in rented apartments, have less cohabitants in their household, and are married with lower probability.¹⁰

5 Empirical analysis

5.1 Validity of the RD design

Although the validity of the RD design is not testable, this section provides feasible checks for the credibility of its identification assumptions. If no factors having an impact on the outcomes of interest, except for the retirement rate, change discontinuously at the legal retirement age, consistent estimation of the RD effect does not require additional covariates. In contrast, if other factors change simultaneously, the estimated effect is unlikely to be caused by a change in the retirement rate only. It is therefore of great importance to explore the behavior of available covariates along the running variable. Figure 3 plots the fraction of female household heads, the fraction of tenants, the fraction of married household heads, and the average number of household members against age. The covariate averages by age are complemented by the nonparametric local linear estimates of the *reduced-form* effects induced by the legal retirement age. None of these *reduced-form* effects is significantly different from zero (p -values of 0.318, 0.543, 0.350, 0.405). Even

⁹Gross income consists of earned income as well as pension benefits and social insurance benefits.

¹⁰Widows and widowers are considered as non-married.

though it does not guarantee the absence of unobserved confounders, this result provides supportive evidence for the validity of the RD design.

A related check is to examine the density of the running variable *age* around the threshold. A jump in the density at x_0 would be problematic since it provides some evidence for sorting around the threshold and challenges the appropriateness of the RD design (Lee and Lemieux, 2010). Figure 2 shows the density of the running variable and indicates that it is indeed continuous. The formal test suggested by McCrary (2008), which does not reject the null hypothesis of continuity, corroborates this finding.¹¹ This result provides evidence for local random assignment and supports the credibility of the RD design.

5.2 Average treatment effects

The estimates of the local ATE of retirement are shown in table 2. The parametric 2SLS estimates using linear splines are reported in column 1. The results reported in column 2 additionally include dummies indicating whether the household head is female, home owner or a renter, married, and six dummies for the number of individuals living in the household (households with two members are the reference group). Based on the approach proposed by Dong (2015), columns 3 and 4 report the bias-corrected coefficients of the estimates in columns 1 and 2. For the nonparametric estimates, the label *opt. bw.* refers to the optimal bandwidth chosen using the algorithm by Imbens and Kalyanaraman (2012). The *reduced-form* effects of these nonparametric estimates are graphically illustrated in figure 4. The results reported in the last two columns of table 2 are estimated using bandwidths of half and doubled size. All estimates are computed using the HBS sampling weights.

The estimates of the ATE indicate that disposable income decreases by approximately 2400 CHF per month at retirement. This substantial decline is significant on the 1% level for all specifications. Compared to the average disposable income of households with non-retired heads aged 60 to 70 (see table 1), this amounts to a reduction of -35% . Concerning overall consumption expenditures, the parametric estimates indicate a significant decline by 400 to 600 CHF at retirement, which is equivalent to a reduction of -7 to -11% compared to non-retired households aged 60 to 70. Based on nonparametric estimation, the coefficients of retirement are of similar magnitude. They are, however, not significantly different from zero.

The remainder of table 2 addresses six more specific measures of consumption. Re-

¹¹The McCrary (2008) test was conducted using the bin width of 1 for age and the default bandwidth calculation of the *DCdensity* Stata package (bandwidth = 4.85). The discontinuity estimate (log difference in height) is 0.0035 with a standard error of 0.0793.

tirement does not have significant effects on average expenditures on food consumed at home. This finding is the equivalent for all parametric and nonparametric model specifications. Besides food expenditures, food consumption at home measured in kilograms is included as an additional dependent variable. Interestingly, the estimates of some models indicate that retirement has significantly positive effects on this outcome. In particular, these estimates point to an increase of approximately 12% compared to the average of non-retired households with heads aged 60 to 70. Stated differently, these results provide some evidence that households buy cheaper comestibles after retirement, since the amount of food slightly increases but expenditures stay constant.

Expenditures in restaurants and hotels significantly decrease by 170 CHF at retirement, corresponding to roughly -32% compared to the average of non-retired household with heads aged 60 to 70. This reduction may be triggered by lower work-related expenses since retirees are likely to spend more time at home. The effect on transportation expenditures is also negative and lies within the interval of -100 to -150 CHF per month in all models. For all but one specification, however, the standard errors are relatively large and the estimates are not significantly different from zero. Retirement significantly decreases expenditures on housing. The point estimates point to a reduction of approximately 170 CHF per month (-12% compared to non-retired households), which indicates that some households reduce their rent or mortgage interest payments by moving to cheaper apartments. Average healthcare expenditures are unaffected by retirement.

Except for healthcare expenditures, figure 4 reveals that consumption expenditures and disposable income exhibit clear negative trends along age. In contrast, healthcare expenditures slightly increase with age. Moreover, the bias correction proposed by Dong (2015) has negligible impact on the estimates of the causal effects.

5.3 Quantile treatment effects

This section explores the impact of retirement along the distribution of the dependent variables. Figure 5 shows the point estimates and 95% confidence intervals of the QTE that are constructed using 500 bootstrap replications. All estimates are computed using the HBS sampling weights.

The QTE of retirement on disposable income are significantly negative and stay within the interval of -1000 to -2000 CHF per month for the lower part of the distribution. Starting at the 55th percentile, the estimates get more negative and reach -4000 CHF at the ninth decile, pointing to substantial effect heterogeneity. The decrease towards the end of the distribution explains the average effect of approximately -2400 CHF. In contrast, the QTE for overall consumption expenditures are not significantly different from zero and the points estimates are relatively homogeneous along the entire outcome

distribution.

The estimated QTE for food expenditures are, as the corresponding ATE, not significantly different from zero. Their point estimates, however, increase along the distribution. For the alternative measure of food consumption, food in kilograms, the point estimates follow a similar pattern and are not significantly different from zero in the lower part of the distribution. In the upper part, however, they increase further and become significantly positive at certain percentiles.

Panel (e) of figure 5 illustrates that the impact of retirement on expenditures in restaurants and hotels is significantly negative along the entire distribution. Up to the median, the point estimates are relatively stable within the interval of -50 to -100 CHF per month. In the upper part of the distribution, however, the estimate decreases and reaches -300 CHF around the third quartile. The effects on transportation expenditures are significantly negative at the 15th, 20th, 25th, and the 30th percentiles, but not significantly different from zero for other parts of the distribution. Nevertheless, the point estimates for this outcome substantially decrease along the distribution. They are, however, estimated with less precision. Panel (g) of figure 5 shows that the negative impact of retirement on housing expenditures is driven by significant effects in the upper part of the outcome distribution, whereas the lower part of the distribution is unaffected. Similar to the ATE, the estimated QTE for healthcare expenditures are not significantly different from zero along the entire distribution.

The findings of this section provide valuable insights that could not be uncovered by an analysis that is restricted to ATE. First, the variation in the estimates along the outcome distributions suggests that there is considerable heterogeneity in the causal effects of retirement. Second, the QTE indicate that the sizes of the ATE on most outcomes are determined by large effects in the upper part of the distributions, whereas the impacts at the lower part are often smaller or even insignificant.

6 Conclusion

Using comprehensive income and expenditure data from the HBS, this paper investigates whether there is evidence for the *Retirement Consumption Puzzle* in Switzerland. The findings show that disposable income significantly decreases after retirement, but this effect only weakly transmits to lower consumption expenditures. The decrease in consumption is of remarkably smaller magnitude than the reduction in income and for some specifications not significantly different from zero. Retirement is found to have zero impact on food expenditures, whereas the amount of food in kilograms is positively affected for some specifications, indicating that households buy somewhat cheaper comestibles.

Additionally, the negative impact of retirement on expenditures in restaurants and hotels provides evidence for substitution towards home production of food. Moreover, the significantly negative impact on housing expenditures indicates that some households move to cheaper apartments. Supplementary estimates of QTE show that substantial parts of the ATE are driven by large effects at the upper quantiles of the distributions and thereby point to treatment effect heterogeneity.

Overall, my findings do not support the *Retirement Consumption Puzzle*. First, the reduction in average overall consumption expenditures is substantially lower than the decrease in disposable income, indicating that households simultaneously reduce their savings. Second, there is no effect on food expenditures but weak evidence for a positive impact on the amount of food, which provides some evidence that retirees buy more but cheaper comestibles. They have more leisure time that can be used to shop for bargains and are likely to allocate more time to preparing meals. Finally, the most relevant insights from a policy perspective are the relatively small effects of retirement at the lower part of the outcome distributions. Thus, my results are fairly consistent with consumption smoothing over the life cycle.

There are some limitations to the paper and the data at hand. This study explores pooled cross-section data where the age of each individual is measured in years. In the context of the RD design, however, it would be desirable to observe the age with higher precision, for instance in months. Moreover, the HBS does not include time-use data which would be necessary to examine precisely the substitution from market-produced goods to home-produced goods. A further analysis of the channels of substitution using additional data is certainly worth pursuing in future research. Aside from these limitations, this paper provides evidence against the *Retirement Consumption Puzzle* and shows that the reduction in income at retirement transmits to a considerably less pronounced effect on consumption.

References

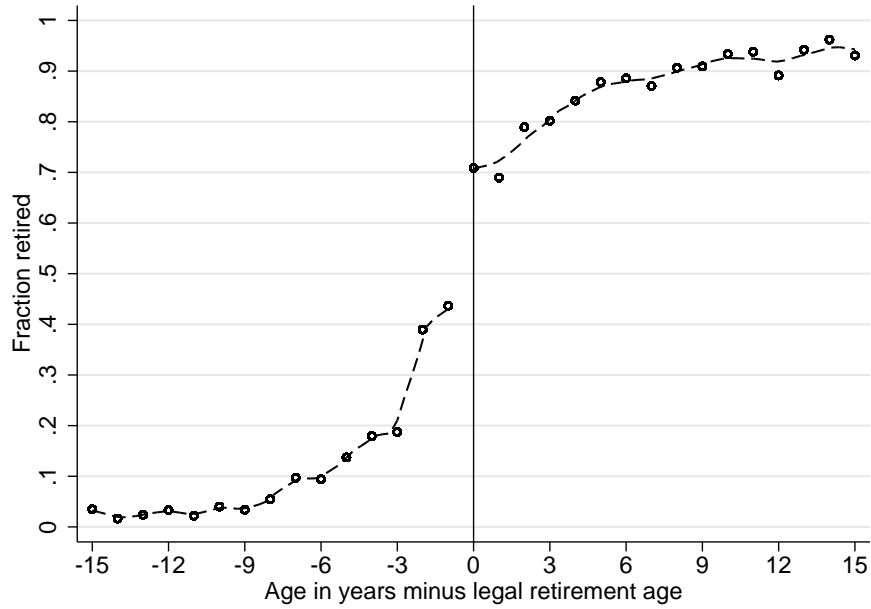
- Aguiar, M., Hurst, E., 2005. Consumption versus expenditure. *Journal of Political Economy* 113 (5), pp. 919–948.
- Aguila, E., Attanasio, O., Meghir, C., 2011. Changes in consumption at retirement: Evidence from panel data. *The Review of Economics and Statistics* 93 (3), pp. 1094–1099.
- Ameriks, J., Caplin, A., Leahy, J., 2007. Retirement consumption: Insights from a survey. *The Review of Economics and Statistics* 89 (2), pp. 265–274.

- Banks, J., Blundell, R., Tanner, S., 1998. Is there a retirement-savings puzzle? *The American Economic Review* 88 (4), pp. 769–788.
- Battistin, E., Brugiavini, A., Rettore, E., Weber, G., 2009. The retirement consumption puzzle: Evidence from a regression discontinuity approach. *The American Economic Review* 99 (5), pp. 2209–2226.
- Bernheim, B. D., Skinner, J., Weinberg, S., 2001. What accounts for the variation in retirement wealth among U.S. households? *The American Economic Review* 91 (4), pp. 832–857.
- Browning, M., Lusardi, A., 1996. Household saving: Micro theories and micro facts. *Journal of Economic Literature* 34 (4), pp. 1797–1855.
- Dong, Y., 2015. Regression discontinuity applications with rounding errors in the running variable. *Journal of Applied Econometrics* 30 (3), 422–446.
- Fisher, J., Marchand, J., 2014. Does the retirement consumption puzzle differ across the distribution? *The Journal of Economic Inequality* 12 (2), 279–296.
- Fisher, J. D., Johnson, D. S., Marchand, J., Smeeding, T. M., Torrey, B. B., 2008. The retirement consumption conundrum: Evidence from a consumption survey. *Economics Letters* 99 (3), 482 – 485.
- Frandsen, B. R., Frölich, M., Melly, B., 2012. Quantile treatment effects in the regression discontinuity design. *Journal of Econometrics* 168 (2), 382–395.
- Gelman, A., Imbens, G., 2014. Why high-order polynomials should not be used in regression discontinuity designs. Working Paper 20405, National Bureau of Economic Research.
- Haider, S. J., Stephens, M., 2007. Is there a retirement-consumption puzzle? Evidence using subjective retirement expectations. *The Review of Economics and Statistics* 89 (2), pp. 247–264.
- Hamermesh, D. S., 1984. Consumption during retirement: The missing link in the life cycle. *The Review of Economics and Statistics* 66 (1), pp. 1–7.
- Hurst, E., 2008. The retirement of a consumption puzzle. Working Paper 13789, National Bureau of Economic Research.
- Imbens, G., Kalyanaraman, K., 2012. Optimal bandwidth choice for the regression discontinuity estimator. *The Review of Economic Studies* 79 (3), pp. 933–959.

- Imbens, G. W., Lemieux, T., 2008. Regression discontinuity designs: A guide to practice. *Journal of Econometrics* 142 (2), 615 – 635.
- Lee, D. S., Lemieux, T., 2010. Regression discontinuity designs in economics. *Journal of Economic Literature* 48 (2), 281 – 355.
- McCrary, J., 2008. Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics* 142 (2), 698 – 714.
- Miniaci, R., Monfardini, C., Weber, G., 2010. How does consumption change upon retirement? *Empirical Economics* 38 (2), 257–280.
- Schwerdt, G., 2005. Why does consumption fall at retirement? Evidence from Germany. *Economics Letters* 89 (3), 300 – 305.
- Smith, S., 2006. The retirement-consumption puzzle and involuntary early retirement: Evidence from the British household panel survey. *The Economic Journal* 116 (510), pp. C130–C148.
- Stancanelli, E., Van Soest, A., 2012. Retirement and home production: A regression discontinuity approach. *The American Economic Review* 102 (3), pp. 600–605.
- Thistlethwaite, D. L., Campbell, D. T., 1960. Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational Psychology* 51 (6), 309 – 317.
- Wakabayashi, M., 2008. The retirement consumption puzzle in Japan. *Journal of Population Economics* 21 (4), pp. 983–1005.
- Yu, K., Jones, M. C., 1998. Local linear quantile regression. *Journal of the American Statistical Association* 93 (441), pp. 228–237.

A Figures

Figure 1: Retirement by age



The dots denote means by age. The dashed lines are local polynomial fits.

Figure 2: Density of the running variable *age*






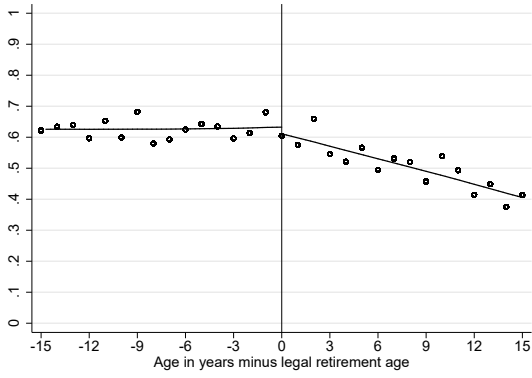
Figure 3: Control variables




(a) fraction of females



(b) fraction of tenants




(c) fraction married




(d) average number of household members


Figure 4: Local linear estimates of the *reduced-form* effects




(a) disposable income



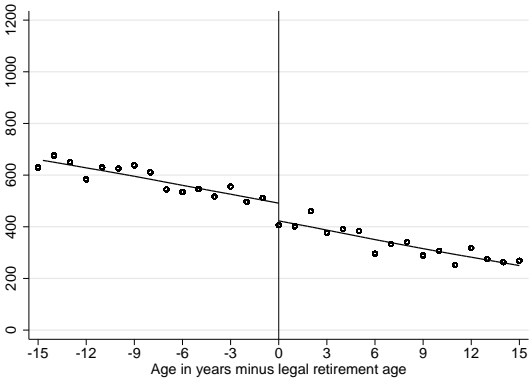
(b) consumption expenditures




(c) food expenditures



(d) food in kilograms




(e) expenditures in restaurants and hotels




(f) transportation expenditures

Figure 4 (Continued): Local linear estimates of the *reduced-form* effects




(g) housing expenditures




(h) healthcare expenditures


Figure 5: Quantile treatment effects




(a) disposable income



(b) consumption expenditures




(c) food expenditures




(d) food in kilograms


Figure 5 (Continued): Quantile treatment effects




(e) expenditures in restaurants and hotels



(f) transportation expenditures



(g) housing expenditures



(h) healthcare expenditures

B Tables

Table 1: Descriptive statistics

	all obs.	non-retired	retired	Aged 60-70	60-70, nonret.	60-70, ret.
Dispos. inc. (CHF)	6120.61 (3517.57)	7418.80 (3674.80)	4459.61 (2452.17)	5584.89 (3136.29)	6787.68 (3539.68)	4624.32 (2369.06)
Consumption (CHF)	5149.31 (2648.68)	5816.02 (2731.31)	4296.28 (2269.76)	5016.11 (2474.47)	5446.34 (2624.49)	4672.52 (2291.20)
Food (CHF)	625.93 (373.76)	680.44 (399.80)	556.18 (324.49)	613.42 (349.22)	636.13 (349.39)	595.28 (348.12)
Food (kg)	67.69 (44.72)	71.92 (48.17)	62.28 (39.22)	66.02 (43.38)	67.14 (46.58)	65.13 (40.64)
Restaurants (CHF)	479.04 (481.18)	589.44 (522.87)	337.79 (377.46)	451.68 (450.22)	533.53 (490.36)	386.32 (403.78)
Transport. (CHF)	646.05 (816.94)	790.29 (850.21)	461.50 (732.04)	609.34 (817.65)	689.23 (827.02)	545.53 (804.65)
Housing (CHF)	1384.11 (737.76)	1504.48 (753.92)	1230.10 (686.54)	1346.36 (730.71)	1453.18 (764.83)	1261.04 (690.71)
Healthcare (CHF)	330.43 (498.29)	299.30 (438.45)	370.25 (563.25)	340.61 (495.93)	314.18 (419.20)	361.71 (548.77)
Age	63.00 (8.74)	57.30 (5.91)	70.31 (5.87)	64.82 (3.20)	62.91 (2.66)	66.35 (2.75)
Retired	0.44 (0.50)	0.00 (0.00)	1.00 (0.00)	0.56 (0.50)	0.00 (0.00)	1.00 (0.00)
Age \geq legal ret. age	0.44 (0.50)	0.11 (0.31)	0.86 (0.35)	0.55 (0.50)	0.26 (0.44)	0.78 (0.42)
Female	0.35 (0.48)	0.30 (0.46)	0.41 (0.49)	0.34 (0.47)	0.30 (0.46)	0.36 (0.48)
Tenant	0.49 (0.50)	0.46 (0.50)	0.51 (0.50)	0.45 (0.50)	0.44 (0.50)	0.47 (0.50)
Married	0.58 (0.49)	0.63 (0.48)	0.51 (0.50)	0.59 (0.49)	0.64 (0.48)	0.56 (0.50)
Size of household	1.91 (0.96)	2.18 (1.10)	1.57 (0.57)	1.71 (0.65)	1.83 (0.73)	1.61 (0.57)
# observations	8524	4845	3679	3275	1363	1912

Means per month (s.d. in parentheses). The category *Restaurants* includes expenditures on hotels. obs.: observations, ret.: retired, nonret.: non-retired, CHF: Swiss Francs. The statistics are weighted by the HBS sampling weights.

Table 2: Local average treatment effects of retirement

Outcome	Linear parametric 2SLS estimates				Nonparametric estimates		
	without covariates	with covariates	bias corr. w/o covar.	bias corr. with covar.	optimal b.w.	$\frac{1}{2} \times$ b.w.	$2 \times$ b.w.
Disposable income	−2269.8*** (328.5)	−2527.6*** (297.9)	−2362.3*** (322.2)	−2564.8*** (292.3)	−2335.0*** (388.5)	−2915.5*** (929.4)	−2289.7*** (343.8)
Consump- tion	−442.2* (261.4)	−600.8** (241.8)	−427.4* (257.6)	−552.5** (238.2)	−427.2 (289.9)	−638.9 (510.7)	−435.4 (271.2)
Food Ex- penditures	27.0 (37.1)	−2.5 (31.7)	26.1 (36.5)	4.7 (31.3)	12.6 (41.2)	−84.6 (80.3)	21.7 (38.3)
Food in kilograms	9.0* (4.7)	4.9 (4.0)	8.8* (4.6)	5.7 (3.9)	6.8 (5.4)	−5.2 (12.0)	8.2* (4.9)
Restaurants and Hotels	−167.4*** (48.2)	−190.4*** (46.8)	−165.5*** (47.2)	−186.9*** (45.8)	−171.8*** (53.7)	−215.4** (95.7)	−169.1*** (50.0)
Transpor- tation	−110.3 (85.2)	−140.5* (82.9)	−106.9 (83.9)	−132.9 (81.6)	−103.9 (93.8)	−151.5 (156.8)	−107.4 (88.2)
Housing	−171.4** (77.7)	−159.9** (76.2)	−175.9** (76.9)	−156.0** (75.4)	−166.6* (87.4)	−238.5 (162.8)	−169.4** (81.0)
Health- care	69.8 (51.1)	67.6 (50.2)	71.6 (51.4)	68.9 (50.5)	75.9 (54.6)	88.6 (65.1)	72.5 (52.5)

Coefficients estimated using a sample of 8524 observations. Heteroskedasticity-robust standard errors are reported in parentheses. corr.: corrected, w/o: without, covar.: covariates, b.w.: bandwidth. Significance levels: * 10%, ** 5%, *** 1%. The sample was weighted by the HBS sampling weights in the estimation.