



Working Paper Series

This paper can be downloaded without charge
from: <http://www.richmondfed.org/publications/>



Richmond • Baltimore • Charlotte

Real Rates and Consumption Smoothing in a Low Interest Rate Environment: The Case of Japan*

Jonathon Lecznar
Boston University[†]

Thomas A. Lubik
Federal Reserve Bank of Richmond[‡]

May 2017

Working Paper No. 17-08

Abstract

We study the dynamics of consumption, the real interest rate and measures of labor input in Japan over the period from 1985-2014. We identify structural breaks in macroeconomic aggregates during the 1990s and associate them with the zero interest rate policy pursued by the Bank of Japan and the surprise increase in the consumption tax rate in April 1997. Formal estimation using the Generalized Methods of Moments shows that the mid-1990s are characterized by breaks in the structural parameters governing household consumption and labor supply decisions. Specifically, following the tax hike and during the low nominal rate period, Japanese households became less risk averse and exhibited a higher degree of habit formation.

JEL CLASSIFICATION: C26; E21; E43

KEY WORDS: Euler equation; GMM; nominal interest rate, labor supply

*The views expressed in this paper are those of the authors and should not be interpreted as those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

[†]Department of Economics, Boston University, Boston, MA 02215. Email: jlecznar@bu.edu.

[‡]Research Department, P.O. Box 27622, Richmond, VA 23261. Email: thomas.lubik@rich.frb.org.

1 Introduction

In this paper, we show that the behavior of aggregate consumption changed considerably in Japan in early 1997. Evidence from raw comovement patterns, structural break tests, and more formal GMM-based estimation on structural Euler-equations for consumption growth all indicate that the behavior of aggregate consumption suffered a break during that time period. Based on the historical record, we can in principle correlate this finding with two dramatic policy actions: the Bank of Japan's (BoJ) implementation of a highly accommodative low-interest policy in mid-to-late 1995 and a 2 percentage point rise in the consumption tax rate to 5% in April 1997.¹ We argue that the results in our paper show fairly conclusively that the tax change led to a break in the aggregate consumption series to the effect that it became more serially correlated afterward. This can be explained in terms of a simple consumption-choice model whereby Japanese households formed stronger habit preferences toward their purchases following the tax increase.

The economy of Japan is a congenial environment to study the behavior of aggregate consumption. The period from the mid-1980s through early 2010s in Japan can be tersely described as a boom, then bust, followed by a long period of primarily stagnation and intermittent deflation. Throughout this period there were marked changes in multiple facets of governmental policy. With regards to monetary policy, the BoJ lowered rates to hitherto historic lows in 1995, only to eventually go further in 1999 by introducing the zero interest rate policy (ZIRP). The BoJ's policy rate has not deviated very far from zero ever since. On the fiscal policy side, numerous rounds of fiscal stimulus were passed beginning in 1992, labor laws on temporary employment were relaxed in 1998, and a tax on consumption was initially introduced in 1989 then subsequently raised in 1997.

We first assess whether key macroeconomic time series exhibit changes in behavior over the period from 1985 through 2014. In particular, we consider measures of consumption, the real interest rate, and the extensive and intensive margins of employment. A simple ocular inspection of the data suggests that they do, as both consumption growth and the real rate of interest appear to begin behaving differently in the mid-1990s. Using a bevy of structural break tests, we identify the second quarter of 1995 as a break in the real interest rate series, which coincides with the onset of a period when the BoJ held the policy rate fixed at 50 basis points. We also find a break in consumption growth in the second quarter of 1997, which coincides with the hike in the consumption tax rate that is often regarded

¹Initially announced in November 1994, the tax rise was associated with an anticipatory spike in consumption followed by a sharp drop, then protracted stagnation.

as the starting point of the lost decade. Moreover, we find evidence of a structural break in the behavior of employment and hours worked that started earlier in the 1990s. The picture that emerges of Japan's economy during the 1990s is one of considerable change in the macroeconomic environment.

Given the dramatic changes in the economic and policy environments during this time, we ask whether the standard consumption Euler-equation is a good and consistent descriptor of consumption growth throughout such an environment. Economic theory suggests that the key explanatory variable for consumption growth is the real rate of interest. A convenient way of thinking about this relationship can be found in the optimal savings decisions of households. More specifically, we consider the canonical consumption Euler-equation arising from constant relative risk aversion (CRRA) preferences with risk aversion parameter σ . It describes how consumption C_t , a (gross) nominal interest rate R_t , and (gross) inflation π_t are related to each other:

$$C_t^{-\sigma} = \beta R_t E_t C_{t+1}^{-\sigma} \frac{1}{\pi_{t+1}}, \quad \sigma > 0. \quad (1)$$

β is a parameter that discounts future consumption, and E_t is a rational expectations operator. This relationship can also be expressed in a more compact form by rewriting it in terms of a log-linear approximation:

$$E_t \Delta \tilde{C}_{t+1} = \frac{1}{\sigma} \left(\tilde{R}_t - E_t \tilde{\pi}_{t+1} \right) = \frac{1}{\sigma} \tilde{r}_t, \quad (2)$$

where tilde ‘ \sim ’ denotes logarithmic deviations from the steady state. The real interest rate, \tilde{r}_t , is defined as the log-difference of the nominal rate and expected inflation.

The Euler-equation (2), and its variants discussed below, provide testable implications for how consumption and real rates comove under the assumption of underlying optimizing behavior.² This relationship also implies that the strength of the responsiveness of consumption growth to changes in the real rate is dictated by the degree of risk aversion.³ However, underlying this time-series relationship is the assumption of structural stability which requires both that σ and the theoretical framework that gave rise to this conjectured relationship remain constant over the period considered. The statistical tests on the Japanese macroeconomic time series described above give us strong reason to believe that the assumption of structural stability is violated during this period.

²A rise in the real rate increases current savings and thus lowers current consumption. Consumption is therefore expected to increase from the current period to the next, which induces positive comovement between the two variables. In the data, this simple relationship is violated since it leaves out additional conditioning variables as we document in the main part of the paper.

³For example, increased risk aversion, signified by an increase in σ , implies that increases in the real rate result in a weaker rise in consumption.

We therefore develop a baseline specification that generalizes the basic relationship conjectured in equation (2) to incorporate external habit formation. Using the Generalized Method of Moments (GMM) to estimate this relationship, we find that a standard specification with habits in consumption fits the data over the full sample reasonably well, with structural parameter estimates in line with previous results in the literature. We cannot, however, accept the null hypothesis of structural stability of the estimating equation, which leads us to consider subsamples. From our estimation, we find evidence suggesting that the 1997 consumption tax rate hike changed the nature of consumer behavior. In particular, following the tax increase, we find that consumers became less risk averse, thus responding more to real rate changes, while exhibiting stronger habit preferences. Additionally, we test alternative specifications incorporating the extensive and intensive margins of employment. We conclude that the inclusion of employment in the utility function is not necessary for capturing consumption dynamics in Japan. This is similar to findings reported by Kiley (2010) who, for the US during 1960Q1-2004Q4, also finds evidence for habit persistence but against nonseparability in consumption and leisure. Overall, we provide evidence supporting the predictive power and structural stability of the habits-based estimating equation, although not in its constituent parameters.

The plan for our investigation is as follows. In the next section, we provide an overview of the macroeconomic and monetary history of Japan with a focus on the period since the 1980s as a background for the formal empirical analysis. We derive the consumption Euler equation we intend to estimate in Section 3. We discuss the general specification and highlight specific and nested parameterizations. Section 4 contains the main body of results, while section 5 considers some alternative approaches and robustness analyses. We conclude in section 6.

2 A Short Macroeconomic History of Japan over the Lost Decades

In this section, we provide some background on the development of the Japanese economy from the beginning of our sample period in the mid-1980s through the lost decades of the 1990s and 2000s. The first half of our sample was characterized by fast growth and a run-up in asset prices, whereas the latter half of this period saw a dramatic decline of the economy, followed by a sluggish and incomplete recovery that has often been referred to as stagnation. As background for our empirical analysis with its focus on the relationship between aggregate consumption and the real interest rate, we start out with a brief narrative

of the key episodes over this period. We then provide a more detailed statistical analysis with the aim of establishing some key facts.

2.1 A Brief Narrative⁴

The 1980s were a period of relative calm in the Japanese economy as it emerged from the decades-long rebuilding process after the end of World War II. Especially the mid-to-late 1980s were a time of strong growth, as GDP growth rose from 3.3% on average over 1980Q1-1987Q2 to 5.7% in 1987Q3-1990Q2. At the same time, it was a period of aggregate price stability since from 1982-1989 inflation remained well contained between 0-3%. However, primary stock and land price indices rose 300% in 1983-1989. This era has come to be known as the bubble period that laid the foundation for Japan's lost decades. In hindsight, there were many events potentially contributing to this boom-bust cycle. The significant appreciation in the Yen during this period⁵ induced the BoJ to lower the three-month Gensaki rate⁶ from 7% to 3.75% in 1985Q4-1987Q3. Substantial money growth ensued, with broad money growing at greater than 9% annually between 1986-1988 and peaking at 11% in 1989. The higher money growth rate initially stimulated real variables and asset prices, as the Japanese public had come to expect price stability.

Eventually, as inflation picked up – from 1% in 1985Q1-1989Q1 to 2.8% in 1989Q2-1991Q4 – and asset prices reached staggering heights, the BoJ responded by pushing interest rates sharply higher: the Gensaki three-month rate rose from 4.3% to 7.6% between 1989Q2 and 1990Q4. During this same period, a 3% consumption tax, the first of its kind in Japan, was enacted in April 1989. In response to these contractionary policies, real GDP barely grew over 1992Q2-1995Q1, despite the passing of an initial round of fiscal stimulus in 1992. At the same time, asset prices in general began to fall, punctuated by an approximately 60% drop in the Nikkei stock index between 1990 and 1992. Land prices also exhibited a marked decline beginning in 1991 and continuing into the 2000s. The asset price collapse and prevalence of nonperforming loans resulted in a largely insolvent financial sector and the failure of many smaller institutions, primarily between 1992-1995.⁷ Commensurate with these events, broad money growth slowed to a 3.8% annual rate.

⁴This section draws heavily on Hetzel (1999), Ono and Rebick (2003), Ito and Mishkin (2004), and Fortin and Sicsic (2009).

⁵From February 1985 to November 1985 alone, the dollar fell by 20% against the Yen.

⁶The Gensaki rate pertains to bond repurchase agreements. The one- and three-month Gensaki rates were the relevant policy rates at this time, as short-term government bonds were first available in 1986.

⁷In order to stem this tide of failing financial institutions, the Japanese government switched from guaranteeing individual deposits up to ¥10 million to a complete guarantee.

The BoJ responded to the sharp economic downturn by progressively lowering the overnight rate, eventually hitting a hitherto historical low of 0.5% in September 1995. After the BoJ was granted formal independence in April 1998, the rate was further lowered to 25 basis points in September 1998. It was over this period that inflation began its inexorable decline toward deflationary territory. In line with the overall drop in asset prices and money growth, the GDP deflator began falling in 1991 from 2.6% annually to -0.7% in 1995 and -0.6% in 1996. At the same time, GDP growth stabilized at a low but positive level, despite numerous adverse factors: real wages continued to rise in the 1990s, depressing employment growth; and the yen nearly doubled in value relative to the dollar between 1990 and 1995. Coinciding with the economic slowdown in the 1990s, the weakening of the social compact of life-long employment began to occur. This is evidenced by a fall in the percentage of employed workers considered regular-employees from 80% in 1994 to 66% in 2008. Additionally, the likelihood of being employed by the same employer for at least a decade declined between 1992 and 2002 from 63% to 49%.⁸

Arguably the most consequential policy change was initiated in November 1994, when the Diet passed a bill to raise the consumption tax rate from 3% to 5%, effective in April 1997. The anticipated rise in the consumption tax rate contributed positively to a fleeting recovery via the acceleration of big-ticket purchases. However, the decline soon after was sharp: 1997Q2 GDP fell by an annualized rate of 3.9%; consumption growth dropped to an annualized -10% between 1997Q1 and Q2. The effects were also protracted: despite a brief uptick in the CPI to 2.5% annual growth in the middle of 1998, by 2003 it was 3 percentage points below its 1997 level. Nominal GDP fell by 4% between 1997 and 2002. To add to the economic headwinds fomented by the consumption tax rise, 1997 coincided with the expiration of temporary income tax cuts and the onset of the Asian Financial Crisis. While very far from a causal relationship, as numerous other events were occurring contemporaneously, both the 1989 and 1997 consumption tax raises were soon followed by abrupt economic slowdowns.

The BoJ initiated the original near-ZIRP in February 1999, under the promise of maintaining it until deflationary concerns were dispelled. By 1999Q4, house and stock prices neared their early 1980s levels. Although amid deflation, an effectively contractionary monetary policy,⁹ and governmental pressure, the BoJ abandoned the ZIRP in August 2000 by

⁸A revision to the Labor Standard Law in 1998 is often cited as a key contributing factor in this change. The revision increased the maximum length of fixed-term contracts from one to three, then eventually five years.

⁹Hetzel (1999) argues, “[t]he combination of zero, or negative, expected inflation with an equilibrium real rate near zero means that even the low market rates currently observed in Japan are consistent with

raising its target rate to 25 basis points. This decision would come to be seen as a policy mistake, with negative growth subsequently returning. Amid an economic slowdown and continued deflation, the ZIRP was reinstated in March 2001, coupled with the promise of being instituted until the inflation rate remains steadily above zero. In addition to returning to the ZIRP, the BoJ simultaneously instituted a two-fold “quantitative easing policy”: first, it switched the policy target from short-term interest rates to the BoJ’s net current account position; and second, it started a program of purchasing long-term government bonds. The scope of these policies, along with the institution of further measures,¹⁰ continued to expand markedly through March 2003. Despite these efforts, economic growth was muted and deflation remained present through the 2000s. Real GDP per capita rose by only 2.1%, and the GDP deflator fell 10.4% between 2000-2009.

2.2 Data, Preliminary Results, and Some Stylized Facts

We now want to establish some stylized facts to inform the empirical analysis to be conducted later. We focus on the period shortly before the asset price run-up in the mid-1980s through the Great Recession and its aftermath. To this end, we collect quarterly data from 1985Q3 through 2013Q4, published by the Statistics Division of the Cabinet Office of Japan and available via the Haver database. All quantity variables are normalized by total population and are seasonally adjusted. We compute annualized growth rates as 400 times the quarter-over-quarter log-difference. We follow Kiley (2010) and measure consumption as nondurable goods and services. The series is converted into real values using the consumer price index (CPI) with 2010 as the base year. We compute the real interest rate as the difference between a short-term nominal interest rate and a measure of expected inflation. For the former, we choose the uncollateralized overnight call rate, which is the BoJ’s policy rate. We measure it as the effective, end of period, annual rate. Expected inflation is approximated by the annualized growth rate in the CPI between the subsequent and the current quarters. Our maintained assumption is that the realized one-period-ahead inflation rate is a good proxy for its one-step-ahead forecast. Similarly, current inflation is computed as the annualized growth rate in the CPI between the current and previous quarter. We use two measures of labor supply, namely total employment from the Japanese Labor Force Survey, which captures the extensive margin of labor adjustment. Alterna-

contractionary monetary policy.”

¹⁰In the fall of 2002, the BoJ began buying stocks from banks. Additionally, over the course of 2003, the BoJ adds bank bills and commercial paper, along with asset-backed securities and commercial papers to its portfolio. The goal of these policies was to remove risky assets from bank’s balance sheets.

tively, an intensive measure is given by aggregate weekly hours worked in nonagricultural industries.

Figure 1 illustrates the primary relationship we investigate in our empirical exercise. We plot the growth rate of nondurables consumption, ΔC , against the nominal, $i_{Nominal}$, and the real rate of interest, i_{Real} , where each series is constructed as described above. The graph conveys the impression that there are three distinct episodes of post-1985 Japanese macroeconomic history. From the start of our sample through the rise and collapse of Japanese asset prices and the accompanying recession up until the mid-1990s, consumption growth is volatile, with highs of close to 15% almost matched by lows of close to -10%. Over the course of this period the nominal rate declines from 8.5% in the early 1990s to a level of just above zero in 1998. This trend behavior of the nominal rate is matched by the real rate, although the latter appears more volatile.¹¹ The policy rate hits zero in early 1999, which coincides with the second episode we can identify in Figure 1. From then on, consumption growth is less volatile and remains at a lower level, as does the real rate of interest. Since the nominal rate is at the zero lower bound, any movement in the real rate is thus driven by changes in expected inflation in the way we constructed the real rate. The picture changes again with the onset of the Great Recession when consumption and real rate *volatility* rise again, whereas the nominal rate remains at zero. This sequence of episodes indicates that the relationship between consumption growth and the real rate may have undergone changes that are related to the zero lower bound on nominal interest rates. This is one of the questions that we take up in our paper.

In order to establish a baseline for the changes in these relationships, we compute simple correlations that are reported in Table 1. We split the sample in 1997Q2, which we visually identify as a likely break date. Using more sophisticated statistical methods, we confirm below that this date is, in fact, consistent with a break in the consumption series. Over the full sample period, consumption and the real rate are positively correlated as measured by a correlation coefficient of 0.39. This correlation declines between the two subsamples from a value of 0.38 to 0.25. We provide further evidence of the changing nature of this relationship in Figure 2, where we report five-year rolling window correlations between consumption growth and the real rate. While the correlation is positive over the full sample and the subsamples, the size of the correlation varies in line with the three episodes we

¹¹It is, of course, a central empirical question which direction the cause-and-effect relationship runs. Does the policy rate follow the real rate down in the worldwide decline of interest rates? Or is policy such that it is accommodative and working through an expected inflation channel. More discussion and some recent evidence is provided by Laubach and Williams (2015) and Lubik and Matthes (2015a).

identified in the previous paragraph. The first period exhibits a correlation of around 0.35, which at the onset of the zero interest rate policy rises to well above 0.5. The correlation comes down sharply in 2008 when the rolling window starts to include data points from the Great Recession. Throughout this period the correlation remains below 0.3 and is thus lower than the correlation at the beginning of our sample.¹²

We also look at the behavior of measures of labor input over the sample period. As we will show below, economic theory allows us to link the behavior of consumption growth and the real rate to changes in employment via the intertemporal Euler-equation. Focusing only on the consumption-real rate relationship may run danger to an omitted variable bias in how consumption growth is determined. For a first assessment of the potential importance of labor, we plot the growth rate of total employment, ΔN , our extensive margin, against consumption growth and the real rate in Figure 3. Employment is noticeably less volatile than the other two series. Moreover, the contemporaneous correlation coefficients in Table 1 suggest that employment growth is only weakly correlated with consumption, if at all, and only mildly stronger with the real rate. Noticeably, the correlation is strongest in the second half of the sample, albeit negative with respect to both consumption and the interest rate. While the latter may not be surprising since higher rates tend to be contractionary and thereby reduce employment, the former fact may be unexpected. We assess this finding more formally when we estimate a theoretical relationship between these variables below. Finally, Figure 4 depicts five-year rolling window correlations. The relationship between consumption and employment does not appear to change as markedly as that with the real rate. Toward the end of the sample, the relationship turns decidedly negative, while there is a period in the mid-1990s whereby this relationship is noticeably positive.

In the next step, we assess the possibility of breaks in the time series of interest more formally. Table 2 reports results from various structural break tests on the series for non-durable goods and services consumption, the real interest rate, and two labor market variables: total employment and average hours worked. Overall, the results confirm what the more casual eyeballing tests above suggested. We find robust evidence of structural breaks in all variables throughout the 1990s and around the time of the onset of the BoJ's zero interest rate policy. There are, however, some interesting differences among the series.

The tests clearly identify 1997Q2 as the break period in our consumption series. The

¹²Lubik and Matthes (2015b) highlight the importance of modeling time variation explicitly in aggregate time series. They advocate the use of time-varying parameter VARs with stochastic volatility to delineate different sources of time variation and apparent breaks in data. Applying this methodology to our question at hand goes beyond the scope of the paper but is a topic for future research.

date of this break aligns ominously with the April 1, 1997, increase in the consumption tax, indicating the possibility that a policy change coincided with or induced the break. Turning to the specific test results, the sequential Bai and Perron (2003) test for the number of breaks in a series indicates a single break over the full sample period. The onset of the Great Recession, on the other hand, seems not to line up with a break, as none of the tests indicate a break around the 2007-2008 period.¹³ We take this as supportive of our focus on changes in the policy environment, be it a shift in the BoJ's policy stance or changes in consumption-relevant tax rates as a driver of changes in the macroeconomic environment. Continuing with the evidence for consumption, the Andrews (1993) test for a single unknown breakpoint also picks 1997Q2. To assess the robustness of this finding, we performed standard Chow-tests for a range of known break dates around this period. Again, 1997Q2 emerges as a break date. At the same time, there is some uncertainty over the exact break date, as we can reject the null hypothesis of no break in the period 1996Q2-1997Q4 for typical significance levels.

This uncertainty over the break date is echoed in our results for the real interest rate.¹⁴ We find strong evidence for the existence of a single break in 1995Q2. As before, a simple Chow test indicates that we can reject the null of no break for a wide range of break dates around this time. We note that the second quarter of 1995 coincides with the date when the call rate settled on 50 basis points for an extended period after coming down substantially in the wake of the collapse of the asset price bubble. This date also coincides with the period when the correlation between consumption growth and the real rate changes substantially.

A key hypothesis we investigate in our paper is whether the behavior of consumption growth is partially explained by the behavior of employment due to nonseparabilities in the utility function. We consider total employment as a measure of the extensive margin of labor input and average hours worked for the intensive margin. While the Bai-Perron test and the Andrews test for an unknown break both point toward 1992Q2 for the total employment series, a break in average hours worked can be rejected.¹⁵ This suggests that the

¹³We should note, however, that the length of the Great Recession subsample is short enough to raise small-sample concerns for these break tests, especially since the onset of the Great Recession is close to the 15% trimming of the overall sample as recommended in the literature. Nevertheless, we will take a separate look at the Great Recession period in our robustness section.

¹⁴We choose to focus on the real rate since it is the key variable for understanding consumption growth. Moreover, the fact that the nominal rate was subject to the zero lower bound can be considered as independent evidence of a break in the nominal rate as the economy changes its underlying dynamics in this case. The question thus remains whether a commensurate break in the behavior of expected inflation offsets the break in the nominal rate or not.

¹⁵The Bai-Perron test finds evidence for one break in 1990Q1 that is just about significant at the 5% level and close to the start-point of the 15% trimming range. Given the evidence from the other tests we discount

economic upheaval in the 1990s that culminated in the ZIRP and a lost decade started with a structural break in the behavior of the extensive margin of employment. The obvious corollary is that the Japanese model of lifetime employment suffered its demise with a downward adjustment in employment growth.

We can now summarize our preliminary empirical findings as follows. We find substantial evidence of a structural break in the behavior of several aggregate time series in the 1990s. The behavior of employment, particularly along the extensive margin, changed in the early 1990s. Interestingly, this timing roughly coincides with the loosening of the “job for life” model that is considered to have begun in the mid-1990s (see Fortin and Sicsic, 2009). This was followed by a break in the behavior of the real rate around 1995 when the BoJ began a policy of very low interest rates reaching zero in 1998. Lastly, the consumption growth series experienced a structural break in 1997Q2. This date is ominous because of the change in the consumption tax from 3% to 5% on April 1, 1997, which is widely credited as the starting point for deflation in Japan and pushing the economy into a long recession. The next step in our study is to analyze the behavior of consumption growth and its determinants in light of the consumption Euler-equation.

3 A Consumption Euler-Equation: Theory and Empirics

The key theoretical building block for our analysis is the consumption Euler-equation that is derived from a household’s utility-maximization problem. Assuming risk aversion, a household and its members desire to smooth consumption over time. This can be accomplished, for instance, by holding and investing in interest-bearing assets, such as nominal bonds. These assets deliver payoffs to sustain consumption when other sources of income decline; they provide a vehicle for savings and transfer income over time when there is a temporary windfall. The optimal intertemporal consumption choice depends on the effective real rate of return of the asset portfolio. As is well known, the generic optimality condition for such an optimization problem is:

$$\lambda_t = \beta E_t \lambda_{t+1} R_t / \pi_{t+1}, \quad (3)$$

where λ_t is the marginal utility of wealth, R_t is the (gross) nominal return, and $\pi_t = P_t / P_{t-1}$ is the (gross) inflation rate. P_t is an aggregate price index. $0 < \beta < 1$ is the household’s discount factor.

However, the marginal utility of wealth λ_t is generally unobservable. In order to derive this finding,

testable implications from this relationship, we need to link it to observable variables. Since λ_t is also the Lagrange-multiplier on the household's budget constraint, we can connect it to the marginal utility of consumption: $\lambda_t = MU_t(\cdot)$. Depending on the specification of the utility function, namely its parametric form and the type of its arguments, we can then estimate the resulting Euler-equation using limited-information methods. We follow Kiley (2010) in choosing a broad set of specifications for the utility function.

Our first specification, which we use to establish a baseline for the parameter estimates, allows for habit formation in consumption. We assume that a household's utility depends on current consumption as well as the previous period's consumption level. Formally, we can capture this by the utility function $U(\mathcal{C}_t) = \frac{\mathcal{C}_t^{1-\sigma}}{1-\sigma}$, where $\mathcal{C}_t = C_t - \theta C_{t-1}$ is effective consumption under habit formation, and $0 \leq \theta < 1$ is the habit parameter. Furthermore, we allow for curvature in the utility function, where $\sigma > 0$ is the intertemporal substitution elasticity. Assuming that agents have external habits,¹⁶ that is, that they do not take into account that today's consumption choice affects tomorrow's habit stock, the optimality condition is:

$$\lambda_t = (C_t - \theta C_{t-1})^{-\sigma}. \quad (4)$$

Substituting into the generic Euler-equation and computing a log-linear approximation results in:

$$E_t \Delta \tilde{\mathcal{C}}_{t+1} = \theta \Delta \tilde{\mathcal{C}}_t + \frac{1-\theta}{\sigma} \left(\tilde{R}_t - E_t \tilde{\pi}_{t+1} \right). \quad (5)$$

We note that for $\theta = 0$, the expression reduces to the standard case without habits as in equation (2). Habit formation simply redistributes the consumption adjustment mechanism away from rapid interest rate movements toward slower intrinsic consumption movements. Conditional on the current level of consumption, increases in the real rate imply higher expected consumption growth. However, depending on the underlying factors, the relationship could turn on its head such that real rate increases are associated with lower expected consumption growth, which in turn would require lower current consumption growth.

The second specification we consider allows for an additional variable in marginal utility, namely labor input N_t . We assume the preference formulation, $U(\mathcal{C}_t, 1 - N_t) = \frac{1}{1-\sigma} \mathcal{C}_t^{1-\sigma} v(1 - N_t)$, where $\mathcal{C}_t = C_t - \theta C_{t-1}$ is effective consumption and $v(\cdot)$ is utility derived from leisure $1 - N_t$. This specification implies the Euler-equation:

$$E_t \Delta \tilde{\mathcal{C}}_{t+1} = \theta \Delta \tilde{\mathcal{C}}_t + \eta \frac{1-\theta}{\sigma} \frac{N}{1-N} E_t \Delta \tilde{N}_{t+1} + \frac{1-\theta}{\sigma} \left(\tilde{R}_t - E_t \tilde{\pi}_{t+1} \right), \quad (6)$$

¹⁶With internal habits, agents do internalize this feedback effect. This makes the analytics more cumbersome, since it introduces additional leads and lags in the consumption Euler-equation.

where $\eta = -v'(\cdot)(1 - N)/v(\cdot) > 0$ is the labor supply elasticity and N is the steady-state value of the labor input. When $\theta = 0$, the specification reduces to the non-habits case, but it still allows for expected employment growth to enter the Euler-equation. We note that this specification does not alter how current consumption growth and the real rate affect expected consumption, that is, the respective coefficients on these terms remain the same. In that sense, expected employment growth simply enters as an additional regressor. However, the coefficients on these variables are connected via cross-coefficient restrictions imposed by theory (and the specific form of the utility function). This specific functional form allows us to separate identification of the parameters of interest. The habit parameter θ is identified off the first term in (6), while the substitution elasticity σ can be identified off the last term from the movements of the real rate. Given a calibrated level of long-run employment, we can then identify the labor supply elasticity η from the movements of the labor variable.

We estimate the Euler-equations detailed above using GMM. It is well known that a GMM approach is quite sensitive to the instrument set being used and the method utilized to compute the weighting matrix in small samples, especially with respect to the heteroskedasticity and autocorrelation (HAC) robust estimation of the variance-covariance matrix. Moreover, we are mindful of weak instrument problems, which we address separately in our robustness section. In order to maintain consistency across the different specifications, we use as our baseline method the Newey-West HAC estimator with fixed bandwidth. The estimator is evaluated in a feasible manner by iterating to convergence. Experimenting with different specifications, we found that this baseline method provided overall quite satisfactory results. We will point out deviations from this baseline using different methods where appropriate.¹⁷

4 Empirical Results

We present the key results in two steps. We first estimate a baseline specification that omits the role of additional covariates in the Euler-equation. Based on the initial results, we then augment this version by the inclusion of employment as prescribed in the previous section.

¹⁷Hall (2005) has an extensive discussion of the care that needs to be taken when interpreting the results from different empirical GMM methods when the underlying theoretical model is misspecified.

4.1 A Baseline Consumption Euler-Equation

We first estimate a specification that only includes habits in consumption, namely equation (5), which establishes a baseline for the extended version of the Euler-equation. This is a standard specification in the empirical consumption-based asset pricing literature and in macroeconomic models, which has proved to deliver reasonable performance. The key aspect is that habit formation introduces a lagged term in the estimating equation, which is designed to capture the serial correlation in consumption growth data. We estimate two structural parameters, the intertemporal substitution elasticity σ and the habit parameter θ . The results are reported in Table 3.

We first consider estimates from a baseline instrument set that includes the second through fourth lags of nondurables and services expenditures growth, CPI inflation and the overnight call rate. The results from the full sample estimation, from 1986Q4 - 2014Q4, are representative of the literature and are in line with our prior expectations. The intertemporal substitution elasticity σ is estimated at 1.44 with a standard error of 0.26, while the estimate of the habit parameter θ at 0.17 is not significantly different from zero. The J-statistic for a test of the overidentifying restrictions indicates that the moment conditions are valid at a p-value of 0.58, which is par for the course in consumption Euler-estimations. Alternative empirical specifications, including more lags and alternative weighting matrix estimators result in estimates of these parameters in the same ballpark. In particular, a statistically significant substitution elasticity ranging between 1 and 2 and a small habit coefficient that is often not statistically distinguishable from zero. However, closer inspection of the empirical results reveals what may have been apparent from the discussion of the raw data above. The behavior of Japanese aggregate time series has changed in the mid-to-late 1990s as the BoJ entered its period of pursuing the ZIRP. When we plot the residuals from the baseline consumption-Euler specification (see Figure 5) we note a sizeable drop in their volatility around 1997 and an increase in volatility, albeit less pronounced, around 2007, the start of the Great Recession. This clearly suggest that the baseline specification misses out on key aspects of the data.

In the next step, we estimate the consumption Euler-equation over sub-periods. As our analysis of the raw data has shown, a break in the real interest rate likely occurred in 1995Q2. This coincided with the nominal rate reaching a level of 50 basis points for an extended period, before being lowered further. At the same time, we also identified a break in consumption growth in 1997Q2, the timing of which corresponded with a hike in the consumption tax. Visual inspection of the graph of residuals in Figure 5 shows that

the break seems closer to the middle of 1997. Formally, we can assess whether the sample period has experienced a structural break in terms of the consumption Euler-equation by performing the Andrews-Fair Wald test for the null hypothesis of structural stability.¹⁸ The sup Wald-statistic (not reported, but available on request) points toward a structural break in 1996Q4 at a p-value of 0.42, whereby the p-value for 1997Q1 is 0.40. Given our evidence from the raw data, we therefore decide to split the sample into two, starting the second period from 1997Q1 onward.¹⁹

The results from the subsample estimation using the same instrument set as before are reported in Table 3. In the first sample period, the substitution elasticity rises to 2, whereas the habit parameter estimate is negative and insignificant.²⁰ The p-value of this specification is 0.65. The results for the second subsample are quite different, however. With a p-value of 0.70, σ now falls to a (still significant) 0.42, whereas the estimate of θ is 0.52 with a standard error of 0.12. We obtain the same pattern for all empirical specifications and weighting matrix choices. Specifically, the full sample parameter values are an “average” of the subsample estimates, whereby Japanese consumers became less risk averse but allowed for more habits in consumption at the turn of the ZIRP period.

We can also interpret these findings in terms of the properties of the data. As the Euler-equation (5) indicates, the introduction of habits adds lagged consumption growth to the specification. The habit parameter θ is therefore identified from the degree of persistence in consumption. In the first subsample, this parameter is indistinguishable from zero, while in the second sample period it rises to around 0.5. In other words, the ZIRP and the consumption tax hike made consumption growth more persistent when compared to the preceding period.²¹ The substitution elasticity σ is then identified from the responsiveness of consumption growth to real rate movements given θ . The elasticity $(1 - \theta)/\sigma$ in the first

¹⁸We include the Great Recession period as part of the second sample since the zero-lower-bound issue was present during that time span as well. Moreover, the shorter sample period for this episode raises concerns about the power of these tests.

¹⁹Following the recommendation of Hall and Sen (1999) to identify the source of instability we also perform their O-test. Since the Wald test indicates structural instability, this can stem from parameter instability or instability in the instrument set. We find that the O-test statistic is highly insignificant with values at 0.99 for a wide range of possible break dates. This suggests a broader source of instability than just changes in the parameters. One likely candidate is explicit stochastic volatility, which also seems indicated by the behavior of the residuals in the baseline regression. However, analyzing this aspect further goes beyond the scope of our paper.

²⁰The results remain the same when we fix $\theta = 0$ for this period.

²¹There is a potential fallacy here in that it is well known that it is difficult to disentangle intrinsic (via habits) from extrinsic (via exogenous shocks) sources of persistence in rational expectations models (see Nason and Smith, 2008, for further discussion and examples). That is, the source of increased persistence in consumption growth is not an increase in habit formation but via a more persistent real rate. The use of a structural model and the embedded cross-coefficient restrictions only guards partially against this.

period is 0.55, then rises to 1.15 in the ZIRP period. Less volatile real rate movements – which stem almost exclusively from changes in expected inflation – now have a larger effect on consumption growth.

The findings from the baseline structural estimation therefore lend support to the view that the rise in the consumption tax rate changed the nature of Japanese consumer behavior, in that both attitudes toward risk and assessment of relative consumption choices were affected. An alternative interpretation, however, is that the baseline model is misspecified due to the omission of an explanatory variable. That is, what appears as a structural break in the estimating equation simply reflects the changing nature of an omitted variable. We now assess this hypothesis by turning to an alternative specification of the consumption Euler-equation.

4.2 An Euler-Equation with Employment

Intertemporal consumer choice implies that the change in the marginal utility of consumption is driven by the real interest rate. In the simple model, where utility depends on consumption only, this translates into a direct relationship between consumption growth and the real rate. However, the macroeconomic literature abounds with alternative specifications for consumer utility that allow for additional arguments interacting with consumption choice. The margin we consider is the labor-leisure trade-off as derived in Section 3. We therefore estimate equation (6) with GMM as an alternative to the benchmark specification. This leaves us with an additional parameter to estimate, namely the labor supply elasticity $\eta \geq 0$. Moreover, the derivation of equation (6) shows that the coefficient on expected employment growth also depends on the steady-state value of employment. We set $N = 2/3$ in all of our estimation exercises based on the sample mean in the total employment series when normalized by population. One aspect of the exercise that we focus on is to what extent the additional variable can capture the unexplained residuality in the baseline, and more specifically, whether there are still apparent breaks in the residual series. In this section, we use total employment, that is, the extensive margin of labor adjustment, as our observable variable for labor input. We present results for the intensive margin, namely hours worked, as a robustness check in the following section. The estimation results are reported in Table 3.

The estimates of the structural parameters go qualitatively in the same direction as those for the specification with only habits. For the full sample period from 1986Q4 - 2014Q4, the substitution elasticity $\sigma = 0.90$, which is halfway between the two subsample estimates

of 1.71 and 0.35. We also note that the subsample estimates are very close to those of the habits-only specification, whereas the full sample estimates differ significantly. The habit parameter θ rises from zero to 0.53 over the subsample, which is the same estimate as in the baseline. Interestingly, the full sample estimate is $\theta = 0.53$ and therefore identical to the second subsample estimate, albeit with a higher standard error. These results confirm our previous findings: the tax hike apparently instigated a break in household preferences. As in the benchmark specification, a higher degree of habit formation may pick up stronger serial correlation in consumption growth.

The specification with employment allows us to estimate the aggregate labor supply elasticity η in addition to the other preference parameters. For the full sample, we find that $\eta = -0.31$ (with a standard error of 0.24), which is an inadmissible value given the specification of utility. However, the P-value of this specification is considerably higher than that of the baseline specification, which suggests that the estimation algorithm attempts to compensate for underlying behavior in the time series that a theoretically consistent specification over the full sample cannot fully accommodate. When we restrict estimates of η to be non-negative, the estimated value is zero (reported in Table 3). For the two subsamples the estimates of η are 0.81 and 0.11, respectively, whereby the latter is not statistically significant at conventional values. Moreover, the p-values are essentially the same as in the habits-only specification. When we look at the time series of the estimated residuals from this specification (not reported, but available upon request), we find a more or less identical pattern as in the benchmark. Prior to the break in 1997Q1, the residuals are considerably more volatile than in the second half of the sample, but the overall degree of variation in the residuals appears very similar to those in the benchmark without employment in the regression.

Finally, we also perform structural break tests on the GMM estimating equation (results not reported). The Hall-Sen test for structural instability finds strong evidence for overall instability throughout the middle of the full sample period (which is the period we are focusing on), similar to our benchmark findings. We also cannot reject the null hypothesis of stability based on the Andrews-Fair Wald test. However, the highest p-value of 0.31 is reached in 1996Q2, which is almost a year earlier than in the benchmark and not obviously related to any policy decisions that may have led to this break. This finding could be explained with reference to the breaks in the behavior of employment that we found in the raw data. As Table 2 shows, there is a likely break in the employment series in 1992Q2. Given the excess volatility of the residuals in the first half of the sample, the effects of a

break in one series may have taken time to affect other series as well. Incidentally, this break date is also close to the break in the real rate series in 1995Q2. However, we are discounting this finding to the extent that we have to restrict the estimate of the labor supply estimate in order to get valid results. We therefore conclude that allowing for substitutability between consumption and labor, as measured by total employment, is not necessary for capturing consumption dynamics in Japan. The dominant factor appears to be the break in the consumption series.

5 Robustness and Further Empirical Results

We consider a few additional empirical exercises to further substantiate our findings for the benchmark specification. First, we re-estimate the Euler-equation with labor input using data on hours worked to assess whether changes in the behavior of the intensive margin are important. Second, we look at the issue of the robustness of the estimates given that GMM often has to contend with problems of weak or invalid instruments. In the third exercise, we consider the behavior of our variables of interest before and after the Great Recession which arguably is another period of potential structural change.

5.1 The Intensive Margin of Labor Adjustment

The specification of the Euler-equation (6) includes a term for labor input but is in principle silent on what the variable N measures. In the benchmark, we used total employment as the observable series. An arguably more relevant series is the number of hours worked, which is a broader measure of labor input since it also captures the intensive margin. In order to assess the robustness of our benchmark results, we therefore redo the previous analysis with this alternative labor supply measure. The standard break tests conducted on the raw data in Table 2 suggest that we can date one structural break in 1990Q1. However, the evidence is less statistically robust than for the other series, especially since the break date is close to the start of the effective sample period. In economic terms, it may very well be that Japanese employers adapted to a break in total employment caused, for instance, by changes in employment or retirement law, by adjusting on the intensive margin. As a result, the path of overall labor input would remain largely unaffected.

When we reestimate the Euler-equation (6) with the alternative series we find that the results hew closely to those for the benchmark (see Table 4). Overall, using hours worked data results in a more elastic labor supply, presumably on account of higher variations in the hours margin. As before, the second subsample implies a much higher estimate for the habit

parameter, which picks up the higher degree of serial correlation in consumption growth following the sales-tax hike in early 1997. Moreover, the residuals from this regression depict the same pattern as evident before. We therefore conclude that our initial findings are robust to the use of alternative labor supply data, specifically with respect to the importance of habit persistence after the tax hike and the relative unimportance of the labor-leisure trade-off in explaining consumption growth.

5.2 Weak Instruments

A general concern in GMM estimation is that the instruments may be weak in the sense that they are not correlated strongly enough with the endogenous variable or that the correlation patterns among the instruments are such that the parameter estimates and their standard errors are not reliable. In this case any hypothesis tests based on a specific instrument should be regarded with caution. To assess this possibility for our benchmark specification we therefore conduct two sets of weak instrument diagnostics. First, we perform the Cragg and Donald (1993) test, which is the multivariate analog of a standard F-test.²² We perform the test for our baseline instrument set, which includes lags of consumption growth, CPI inflation, and the call rate, but also for variations of the instrument set in terms of additional variables and combinations of various lags. We apply this test for the models with habits only and also for the extended specification using employment data. We find across the board that we cannot reject the null hypothesis of the presence of weak instruments. The values of the test statistics never quite reach 2 in the full sample and in the subsamples, whereas the critical values from Stock and Yogo (2005) are around 10 for a significance level of 10%. This leaves us with the impression that the results above need to be interpreted with caution as they possibly reflect distortions to inference from weak instruments. However, since the results of the various specification point in the same direction, these concerns should also not be overinterpreted.

5.3 The Great Recession

In the final exercise, we look specifically at the behavior of aggregate consumption during the Great Recession. A priori, we might expect that the onset of the Great Recession in 2008-2009 could possibly change the behavior of the consumption equation. This seems not to be the case. When we estimate the Euler-equation in its various forms over the Great Recession sample we do not find significant differences from the second subsample period

²²The Cragg-Donald test cannot be performed on a nonlinear equation. We therefore replace all composite nonlinear parameters with new coefficients, following the same procedure as Kiley (2010).

estimates, which began in 1997Q2. This is also evident when we look at the pattern of residuals in Figure 5, where the recession is noticeable but not to the same degree as in the first subsample. More formally, we conduct our usual set of break tests, which find support for a break in either 2009Q1 or 2009Q4 depending on the specific benchmark model, but only at low levels of significance with p-values around 0.30. We can therefore conclude that the findings from our benchmark specifications are robust.

6 Conclusion

We show in this paper that the behavior of aggregate consumption in Japan changed considerably in early 1997. Evidence from raw comovement patterns, structural break tests, and more formal GMM-based estimation on structural Euler-equations for consumption growth indicates that the behavior of aggregate consumption suffered a break during that time period. We can in principle correlate this finding with two dramatic policy actions. First, the BoJ implemented a highly accommodative low-interest policy in mid-to-late 1995, which was accompanied by deflation and a strong appreciation of the yen. In fact, the data show a break in the behavior of the real rate of interest during that period. Second, an ill-timed move by the Japanese government to raise consumption taxes in April 1997 resulted in an anticipatory spike in consumption. This was followed by a sharp drop in economic activity and protracted stagnation. This date directly coincides with a break in the consumption series. We argue that the results in our paper show fairly conclusively that the tax change led to a break in the aggregate consumption series to the effect that it became more serially correlated afterward. This can be explained in terms of a simple consumption-choice model whereby Japanese households formed stronger habit preferences toward their purchases following the tax increase.

A byproduct of our analysis is to show that a consumption Euler-equation with habits in preferences can fit the consumption behavior well in terms of its relationship with the real rate of interest. Despite formal evidence of a break in the real rate series and the fact that the BoJ's monetary policy was subject to the zero lower bound - which could have arguably interfered with the consumption-real rate relationship - we do not find any indication that this is the case during the ZIRP from the late 1990s through the 2000s. Finally, we do not find any evidence that the labor-leisure trade-off plays a significant role in explaining consumption growth beyond the real-rate channel. Our results make arguably a strong case that tax policy is at the core of the economic slump that Japan suffered in the 1990s and further out. Yet, this insight is particularly timely because of the enactment of two recent

consumption tax rate increases: in April 2014 and 2017, the rate rose from 5% to 8% and then 10%, respectively. The imposition of these additional tax hikes, upon the backdrop of our results, argues for further investigation into the relationship between tax policy and consumer behavior in the aggregate.

Our work can be extended in several additional directions. First, there are some remaining concerns as to the validity of the results because of the low power of structural break tests and the presence of weak instruments. Second, the analysis should be broadened to consider alternative specifications of the Euler-equation, especially in regards to preferences. Third, and assuming that the results hold true, the analysis should be expanded to include other intertemporal relationships such as asset pricing or investment equations. Lastly, this analysis could also be utilized to inform models that explicitly model structural breaks as an equilibrium phenomenon.

References

- [1] Andrews, Donald W.K. (1993): “Tests for Parameter Instability and Structural Change with Unknown Change Point”. *Econometrica*, 61, 821-856.
- [2] Bai, Jushan, and Pierre Perron (2003): “Computation and Analysis of Multiple Structural Change Models”. *Journal of Applied Econometrics*, 18(1), 1-22.
- [3] Cragg, John G., and Stephen G. Donald (1993): “Testing Identifiability and Specification in Instrumental Variable Models”. *Econometric Theory*, 9, 222–240.
- [4] Fortin, Aurélien, and Michaël Sicsic (2009): “Japan’s changing labour market and how its affecting its growth models ”. *Trésor-Economics*, 65, 1–8.
- [5] Hall, Alastair R. (2005): *Generalized Method of Moments*. Oxford University Press. Oxford.
- [6] Hall, Alastair R. and Amit Sen (1999): “Structural Stability Testing in Models Estimated by Generalized Methods of Moments”. *Journal of Business and Economic Statistics*, 17, 335-348.
- [7] Hetzel, Robert L. (1999): “Japanese Monetary Policy: A Quantity Theory Perspective”. *Federal Reserve Bank of Richmond Economic Quarterly*, 85(1), 1-25.
- [8] Ito, Takatoshi and Frederic S. Mishkin (2004): “Two Decades of Japanese Monetary Policy and the Deflation Problem ”. *NBER Working Paper*, 10878.

- [9] Kiley, Michael T. (2010): “Habit Persistence, Nonseparability between Consumption and Leisure, or Rule-of-Thumb Consumers: Which Accounts for the Predictability of Consumption Growth?” *The Review of Economics and Statistics*, 92(3), 679-683.
- [10] Laubach, Thomas, and John C. Williams (2015): “Measuring the Natural Rate of Interest Redux”. *Federal Reserve Bank of San Francisco Working Paper* 2015-16.
- [11] Lubik, Thomas A., and Christian Matthes (2015a): “Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches”. *Federal Reserve Bank of Richmond Economic Brief*, 15-10.
- [12] Lubik, Thomas A., and Christian Matthes (2015b): “Time-Varying Parameter Vector Autoregressions: Specification, Estimation, and an Application”. *Federal Reserve Bank of Richmond Economic Quarterly*, 101(4), 323-352.
- [13] Nason, James M. and Gregor W. Smith (2008): “The Phillips Curve: Lessons from Single-Equation Estimation”. *Federal Reserve Bank of Richmond Economic Quarterly*, 94, 361-395.
- [14] Ono, Hiroshi, and Marcus E. Rebick (2003): “Constraints on the Level of Efficient Use of Labor in Japan ”. *NBER Working Paper*, 9484.
- [15] Stock, James, and Motohiro Yogo (2005): “Testing for Weak Instruments in Linear IV Regression”. In: James Stock and D.W.K. Andrews (Eds.): *Identification and Inference for Econometric Models: Essays in Honor of Thomas J. Rothenberg*. Cambridge University Press, Cambridge, MA.

Table 1: Sample Correlations

Correlation Coefficients: 1985Q3 - 1997Q1

	Consumption	Employment	Real Rate
Consumption	-	-0.03	0.38
Employment	-0.03	-	0.08
Real Rate	0.38	0.08	-

Correlation Coefficients: 1997Q2 - 2013Q4

	Consumption	Employment	Real Rate
Consumption	-	-0.14	0.25
Employment	-0.14	-	-0.165
Real Rate	0.25	-0.165	-

Correlation Coefficients: 1985Q3 - 2013Q4

	Consumption	Employment	Real Rate
Consumption	-	-0.04	0.39
Employment	-0.04	-	-0.26
Real Rate	0.39	-0.26	-

Notes: We report contemporaneous sample correlations for the 1985Q3 - 2013Q4 period and two subsamples. Consumption is measured as the growth rate of nondurables goods and services per capita; the real rate is measured as the overnight call rate less the one-step-ahead CPI inflation rate.

Table 2: Structural Break Tests

Break Tests	Consumption	Real Rate	Employment	Hours
Date	1997Q2	1995Q2	1992Q2	1990Q1
<hr/> Bai-Perron <hr/>				
0 vs. 1 break	16.95*	68.37**	57.72**	8.87*
1 vs. 2 break	5.19	0.74	5.19	0.48
<hr/> Andrews <hr/>				
Max LR	7.78 (0.07)	124.10 (0.00)	41.24 (0.00)	2.03 (0.78)
Exp LR	2.32 (0.04)	58.15 (0.00)	16.98 (0.00)	0.31 (0.63)
Ave LR	3.94 (0.02)	48.89 (0.00)	16.35 (0.00)	0.55 (0.60)
<hr/> Chow <hr/>				
	7.78 (0.01)		41.24 (0.00)	2.03 (0.16)

Notes: We report structural break tests for the sample period 1985Q3-2014Q1. For unknown break dates, we trim the data equally by 15%. The effective sample period is thus 1990Q1-2009Q4. Consumption is measured as the growth rate of nondurables goods and services per capita; the real rate is measured as the overnight call rate less the one-step-ahead CPI inflation rate; employment and hours worked are in growth rates. The table reports F-Statistics and, where appropriate, p-values. Bai and Perron (2003) sequentially tests the null hypothesis of L against L+1 breaks. We choose $L^{\max}=5$, and trim the sample by 15%. The 5% critical value is 8.58. Andrews (1993) tests the null hypothesis of no break against the general alternative. We report several versions of the likelihood-ratio test. The Chow-test tests the null hypothesis of no break at a given date.

Table 3: GMM Estimation: Benchmark

Habit Specification								
	J-Stat	P-Value	σ	S.E.	θ	S.E.		
Full Sample 1986Q4-2014Q4	9.46	0.58	1.44	0.26	0.17	0.16		
Sub-Sample 1 1986Q4-1997Q1	8.65	0.65	2.00	0.43	0.00	0.11		
Sub-Sample 2 1997Q2-2014Q4	8.20	0.70	0.42	0.17	0.52	0.12		
Habits & Labor Specification								
	J-Stat	P-Value	σ	S.E.	θ	S.E.	η	S.E.
Full Sample 1986Q4-2014Q4	6.80	0.74	0.90	0.33	0.53	0.22	0.00	0.24
Sub-Sample 1 1986Q4-1997Q1	7.44	0.68	1.71	0.36	0.00	0.10	0.81	0.41
Sub-Sample 2 1997Q2-2014Q4	7.15	0.71	0.35	0.13	0.53	0.12	0.11	0.13

Table 4: GMM Estimation: Robustness

Alternative Labor Data: Hours Worked								
	J-Stat	P-Value	σ	S.E.	θ	S.E.	η	S.E.
Full Sample 1986Q4-2014Q4	2.41	0.99	0.99	0.15	0.61	0.07	0.00	0.06
Sub-Sample 1 1986Q4-1997Q1	8.91	0.54	1.65	0.35	0.00	0.14	0.31	0.17
Sub-Sample 2 1997Q2-2014Q4	4.41	0.93	0.49	0.38	0.62	0.22	0.90	0.06




Figure 1: Consumption Growth and Interest Rates

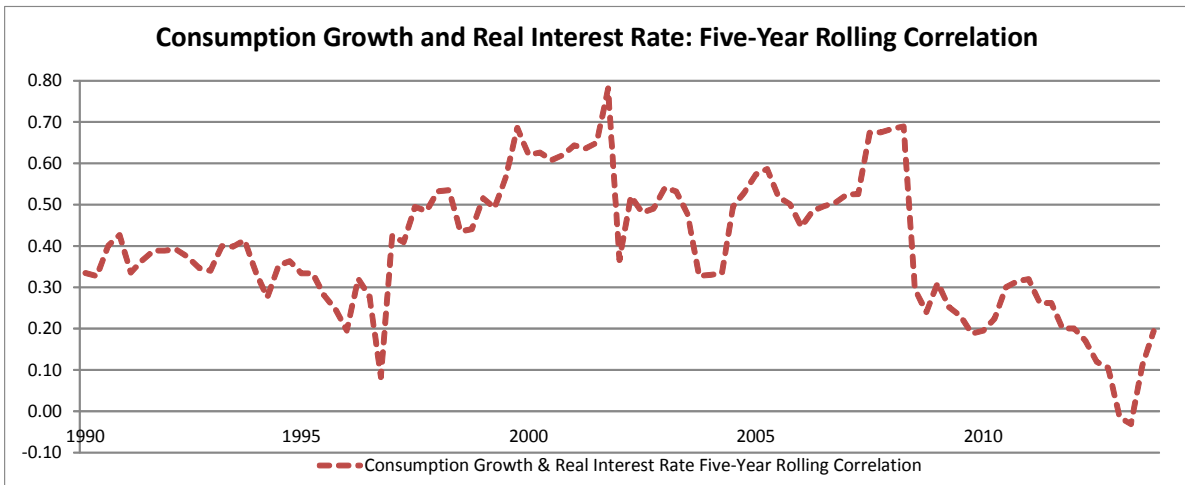


Figure 2: 5-Year Rolling Window Correlations




Figure 3: Employment and Consumption Growth

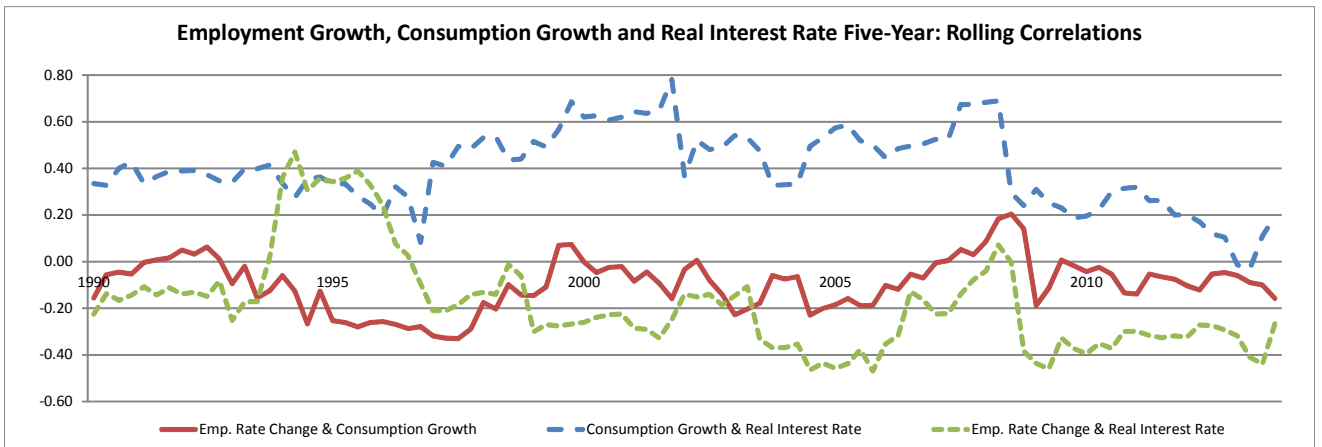


Figure 4: 5-Year Rolling Window Correlations




Figure 5: GMM Residuals: Baseline Specification