



BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

Deflation expectations and Japan's lost decade

by Roberto Piazza

June 2015

Number

274



BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional papers)

Deflation expectations and Japan's lost decade

by Roberto Piazza

Number 274 – June 2015

The series Occasional Papers presents studies and documents on issues pertaining to the institutional tasks of the Bank of Italy and the Eurosystem. The Occasional Papers appear alongside the Working Papers series which are specifically aimed at providing original contributions to economic research.

The Occasional Papers include studies conducted within the Bank of Italy, sometimes in cooperation with the Eurosystem or other institutions. The views expressed in the studies are those of the authors and do not involve the responsibility of the institutions to which they belong.

The series is available online at www.bancaditalia.it.

ISSN 1972-6627 (print)

ISSN 1972-6643 (online)

Printed by the Printing and Publishing Division of the Bank of Italy

DEFLATION EXPECTATIONS AND JAPAN'S LOST DECADE

Roberto Piazza*

Abstract

Starting from the early 1990s, GDP in Japan stagnated for about a decade while inflation has been persistently low, at times even negative. This paper provides new stylized facts about the Japanese deflationary process and puts these facts into the context of the literature addressing the origins of the Japanese "lost decade". In order to properly understand the evolution of inflation in Japan and the role of monetary and fiscal policy, a crucial question is whether the deflationary process, a phenomenon specific only to Japan during the period under consideration, was anticipated by agents. The paper suggests a positive answer to the question. In particular, I show that once a "global" inflation forecasting error, common across advanced countries, is removed from Japanese inflation expectations, then the remaining "idiosyncratic" inflation forecasting error is close to zero at various forecasting horizons. This indicates that the Japan-specific deflationary process was fully anticipated by agents, with medium-term (and "idiosyncratic") inflation expectations disanchoring very soon along the deflationary path.

JEL Classification: E2, E5, G21.

Keywords: deflation, lost decade, inflation expectations.

Contents

1. Introduction.....	5
2. Inflation expectations in Japan	6
2.1 Global and idiosyncratic inflation forecasting errors	7
2.2 Robustness: alternative inflation expectations.....	9
3. Explaining the Japanese macroeconomic performance	11
3.1 Lack of aggregate demand.....	14
2.2 Structural slowdown	16
4. Policy lessons and conclusions	18
References	22

* Bank of Italy, Economics, Statistics and Research.

1 Introduction

In recent years the main advanced countries have found themselves grappling with persistently low inflation rates on various occasions. Conventional and unconventional monetary policy instruments have been deployed to unprecedented extents in an effort to push inflation back toward desirable levels. Although low or even negative inflation was a common feature of the gold standard era, since the Great Depression most advanced countries have consistently maintained positive, and at times very high, inflation rates. Today, central banks have positive inflation targets. Therefore, if we restrict our attention to the era of fiat currency, the emergence of an environment with zero or negative inflation would be unprecedented.

A notable exception to the historical inflation experience in the advanced economies is Japan. In Japan inflation has been low or negative for almost two decades now, with the '90s being defined a lost decade, during which GDP growth was particularly weak. The paper presents both new and old evidence about deflation in Japan and the country's unimpressive growth since the early '90s. It examines into two distinct, albeit related, issues. The first is whether inflation expectations timely and correctly embedded the subsequent deflationary spiral in Japan. The second is whether the literature has been able to identify unambiguously the most important causes of the creeping deflation and slow growth in Japan. The answers that this paper provides are the following.

Regarding the first issue, the empirical analysis suggests that the *underlying* deflationary path - which was specific to Japan alone among developed countries - was correctly anticipated by economic agents. Medium-term inflation expectations of professional forecasters, when appropriately calculated, rapidly became disanchored, notwithstanding a series of reductions in the policy rate by the Bank of Japan. One implication of this observation is that the active use of interest rate policies may not be enough to allow the central bank to uniquely anchor medium- to long-run inflation expectations, and the economy might then be exposed to self-fulfilling deflationary spirals (for a theoretical exposition of this point see Benhabib, Schmitt-Grohe, Uribe [2001], Cochrane [2011] and Piazza [2014]).

Regarding the second issue, the literature surveyed in this paper concurs on the fact that Japan's slow growth during the '90s was associated with previous and coincident booms in firms' investment and leverage, followed by a rapid and permanent fall in investment rates

*Bank of Italy, Directorate General for Economics, Statistics and Research, Advanced Economies and International Finance Division.

and a less rapid and more prolonged deleveraging process. However, when we move from the description of stylized facts to the identification of causal links, there appears to be hardly any consensus. Some put the blame for the Japanese deflation and stagnation on a persistent lack of aggregate demand, which was insufficiently contrasted by monetary or fiscal policy stimulus (Krugman [1998], Posen [1998]). Others point out that the slowdown in productivity growth in Japan, possibly caused by an inefficient allocation of credit by the banking sector, is the true underlying factor (Hoshi and Kashyap [2004]). According to this approach and to some additional evidence, the fall in the investment rate throughout the '90s was not an indication that investment was constrained by the ongoing deleveraging process but that, in the new environment of low productivity growth, there was no need for Japan to continue increasing capital at the same rate as before (Hayashi and Prescott [2002]). Overall, as simple calculations show, the plausibility of explanations that rest on a lack of stimulus and those that instead focus on structural factors relies on one single point: how large the increase in the marginal propensity to save in the Japanese private sector was during the '90s.

The rest of the paper proceeds as follows. Section 2 analyses the evolution of inflation expectations in Japan. Section 3 provides an overview of the literature that examines the macroeconomic performance of Japan since the late '80s. Section 4 concludes by drawing some policy lessons.

2 Inflation expectations in Japan

After reaching over 3% in 1990-1991, CPI inflation in Japan began to fall. Once the one-off effects of the 2% VAT increase in the spring of 1997 are removed from actual CPI dynamics, inflation turns out to have fallen steadily for roughly a decade, reaching a minimum of -0.9% in 2002.¹ By 1999, the short-term policy rate had reached zero, and Japan had effectively fallen into a liquidity trap.

Was Japan's deflationary spiral anticipated, or did deflation catch agents by surprise? Modern macroeconomic theory, which emphasizes the role of agents' expectations about the formation of prices, suggests that this is an important question for understanding the Japanese experience. In addition, inflation expectations permeate the debate on the role and effectiveness of monetary policy. Most monetary theories emphasize the role of the central bank in anchoring medium-term expectations to a level deemed desirable. Within this framework, therefore, explanations of the mechanics behind a deflationary spiral depend very much on establishing whether the phenomenon was an unexpected shock or whether it was matched by a shift in medium-term inflation expectations. In the case of Japan, if the prolonged reduction in inflation was unanticipated, in the sense that inflation expectations did not move initially but instead remained stable and positive for a relatively long period of time, then the process cannot have been caused, at least initially, by a disanchoring of expectations. On the contrary, if the deflation was fully anticipated, then Japanese policymakers wished to maintain positive inflation but may have lacked the credibility or the tools to achieve their goal. In the

¹The series for inflation and inflation expectations used in this paper (except for Figure 3) are net of the effects of the 1997 VAT increase. For the inflation series, I assume that the pass-through to prices was completed within four quarters of the tax hike in 1997Q2. Correspondingly, I distribute the total 2% effect on prices pro-rata between 1997 and 1998. For inflation expectations, I assume that the tax hike was fully anticipated by forecasters starting with the Consensus Forecast survey of October 1996, and that forecasters distributed the total 2% effect on prices according to the pro-rata rule described above.

latter case, the evolution of Japanese inflation could be interpreted as caused by self-fulfilling deflationary expectations, which Japanese policymakers were unable to reverse.

The discussion above highlights the importance of carrying out an analysis of the behaviour of inflation expectations to understand the origins of the Japanese deflation. There are, however, limited sources of information to assess the evolution of inflation expectations in Japan during the late '80s and the '90s. I consider two surveys. The first is the Consensus Forecast survey among professional forecasters. These data have three main advantages: expectations are quantitative, they are measured at different time horizons, and they can be compared across countries. However, there are also some disadvantages, namely that the first observation on longer term inflation expectations from the half-yearly Consensus Forecast survey has been available only since 1990, and that these expectations reflect the views of only a few forecasters, rather than those of the general public.² To address these shortcomings, I compare Consensus Forecast data with results from the Consumer Confidence survey, conducted among Japanese households by the Bank of Japan, which includes a question on expected inflation. The advantage of this survey is its long time span (the dataset I use begins in 1971) and its reliance on a large sample of households, but inflation expectations are measured only on a qualitative scale (increase/stable/decrease) and only at a one-year forecasting horizon. I transform the qualitative information of the Consumer Confidence survey into quantitative expectations using the procedure developed by Carlson and Parkin (1975) and applied to Japan by Ueda (2010)³ I find that during the '90s one-year-ahead inflation expectations from Consensus Forecast and the Consumer Survey are highly correlated and on average close to each other, a fact that I take as a sign of robustness in favour of the conclusion I draw from the Consensus dataset alone.

2.1 Global and idiosyncratic inflation forecasting errors

The analysis of inflation expectations is based on the following equation

$$\pi_t = \pi_{t|t-1}^e + \varepsilon_t^g + \varepsilon_t^i$$

where π_t is the realized consumer price inflation between time $t - 1$ and time t , and $\pi_{t|t-1}^e$ is the conditional expectation agents formed for that variable at time $t - 1$. Moreover, the *total* forecasting error $\pi_t - \pi_{t|t-1}^e$ is split between a component due to a *global* unanticipated inflation shock ε_t^g , stemming for instance from shocks to commodity prices or from global business cycles, and to an *idiosyncratic* shock ε_t^i , specific only to the Japanese economy.

Global and idiosyncratic average forecasting errors can be constructed (in logarithmic form) for forecasting horizons n of more than one year in the following way

$$\frac{1}{n} \sum_{j=0}^{n-1} \pi_{t-j} = \frac{1}{n} \sum_{j=0}^{n-1} \pi_{t-j|t-n}^e + \bar{\varepsilon}_t^g + \bar{\varepsilon}_t^i \quad (1)$$

where for $n \geq 2$ the shocks $\bar{\varepsilon}_t^g$ and $\bar{\varepsilon}_t^i$ are the average global and idiosyncratic forecasting errors over the n forecasting periods. Note that the first term on the left-hand side of the above equation is the average realized inflation rate between time $t - n$ and t , while the first

²See De Pooter, Robitaille, Walker, Zdinak (2014).

³I thank professor Ueda for kindly sharing his dataset and calculations.




Figure 1: Realized (solid line) and expected (dashed line) inflation. Expected inflation is calculated as actual inflation minus the idiosyncratic forecasting error. Dates in the graph refer to the end of the forecasting period. Sources: Thomson Reuters Datastream, Statistics Bureau of Japan, Consensus Forecast (October survey).

term on the right-hand side is the average of the inflation expectations formed at time $t - n$ for the inflation rates between time $t - n$ and time t .

For any given forecasting horizon $n \geq 1$ the corresponding time t global forecasting error $\bar{\varepsilon}_t^g$ is estimated to be the average of the *total* forecasting errors $\frac{1}{n} \sum_{j=0}^{n-1} \pi_{t-j} - \frac{1}{n} \sum_{j=0}^{n-1} \pi_{t-j|t-n}^e$ calculated for Germany, the UK, and the US. This procedure implicitly assumes that the idiosyncratic forecasting errors specific to the three countries are eliminated when taking cross-country averages and so only the global shock remains. The sample period runs from $t = 1990$ to $t = 2006$, during which there was no *permanent* deflationary process in Germany, the UK, or the US, implying that the global shocks $\bar{\varepsilon}_t^g$ should not contain any permanent, unexpected, global deflationary process. Thus, if Japan did indeed experience a permanent deflationary process that was largely unexpected it cannot be embedded in the $\bar{\varepsilon}_t^g$ shock, but should rather show up as a sequence of consistently negative values for the estimated *idiosyncratic* forecast error $\bar{\varepsilon}_t^i$.

At any time t the solid line in Figure 1 depicts, for forecast horizons $n = 1$ and $n = 3$, the average realized inflation rates between time t and time $t - n$, i. e. the left-hand side of equation (1). The difference at t between the solid and the dashed line is by construction equal to the estimated *idiosyncratic* forecasting error ε_t^i , obtained by plugging the estimated global forecasting error ε_t^g into (1). Figure 1 suggests a stark conclusion: once inflation forecasting errors common to the other main advanced countries are filtered out, agents had quite accurately anticipated the occurrence of a Japan-specific deflationary process. Something peculiar happened in Japan during the '90s that did not happen in other advanced countries: the *underlying* large and persistent reduction in inflation was correctly anticipated by economic agents.

Interestingly, the degree of precision with which agents forecast the Japan-specific deflationary path is even greater when we move from relatively short-term (one year) inflation expectations to medium term expectations (three years). This can be more clearly seen in Figure 2, which plots the idiosyncratic and global components of the forecasting errors. At the one year horizon the idiosyncratic error has mean equal to -0.05% and standard deviation of 0.55%, while at




Figure 2: Global (solid line) and idiosyncratic (dashed line) errors. Sources: Thomson Reuters Datastream, Statistics Bureau of Japan, Consensus Forecast (October survey).

the three-year horizon⁴ the mean is -0.01% and the standard deviation is 0.35% . These results indicate that medium-term expectations became disanchored from a largely positive level fairly early in the deflationary path. This conclusion is important because it opens up the possibility that the disanchoring of the medium-term inflation expectations in Japan may not have been just the consequence of the ongoing deflation but instead its cause. For reasons on which we do not speculate here, economic agents ended up coordinating their expectations around the view that medium-term inflation was going to fall in Japan. This sunspot disanchoring of expectations, in turn, may have led actual inflation to fall.

While, as seen, the *idiosyncratic* forecasting error was on average very close to zero, the same is not true of the *total* forecasting error. Over the period 1992-2004 the total forecasting error was persistently and significantly negative, a result due entirely to the behaviour of the *global* forecasting error. Thus, over the period under consideration, agents persistently overestimated future inflation in the main advanced countries especially at the three-year horizon, as seen in the decomposition of the total forecasting error provided in Figure 2.

2.2 Robustness: alternative inflation expectations

The scarcity of long time series on inflation expectations makes it difficult to ascertain whether the results of the previous section are robust to the use of alternative measures of expectations.⁵ To partially address this issue I compare one-year inflation expectations taken from yjr half-yearly Consensus Forecast survey with those that can be obtained from the Consumer Confidence survey of the Bank of Japan. I find that the two measures are highly correlated, a conclusion that provides at least partial proof of robustness to my results.⁶

The Consumer Confidence survey provides a long time series of answers to a question asking consumers whether in the next year prices would go up, remain unchanged, or go down. I quantify these qualitative responses by introducing a variation on the methodology developed by Carlson and Parkin (1975) and applied by Ueda (2010), whose original dataset

⁴At the three-year horizon, the available observations are reduced from 17 to 15.

⁵For an overview of the role and measurement of inflation expectations see Visco(1984).

⁶The conclusions of the paper are also robust to using the April Consensus Forecast survey instead of the October survey.

I employ. In a nutshell, in the original procedure of Carlson and Parkin (1975) agents are assumed to believe that inflation between time $t - 1$ and time t is a normal random variable $\tilde{\pi}_t$ with mean $\pi_{t|t-1}^e$ and standard deviation $\sigma_{t|t-1}$. The goal of the procedure is to use the qualitative information from the survey to ‘pin down’ the moments of the normal distribution between $t - 1$ and t , thus obtaining the expected value of inflation $\pi_{t|t-1}^e$. The shape of the distribution is ‘pinned down’ by making two assumptions. The first is that there exists a time-invariant parameter $\delta > 0$ such that the probability of $\tilde{\pi}_t$ falling below the threshold $-\delta$ equals the share of individuals indicating that they expect prices to fall, and the probability of $\tilde{\pi}_t$ rising above the threshold δ equals the share of individuals that indicating that they expect prices to increase. Clearly, the probability of $\tilde{\pi}_t$ falling between $-\delta$ and δ must be equal to the share of respondents that indicating that prices will not change. Thus, a natural interpretation of δ is that it provides the *precision* with which respondents make explicit that they expect inflation to be different from zero. Having used the shares of respondents to pin down two percentiles of the normal distribution, it is possible to show that the time varying mean can now be written as a function $\pi_{t|t-1}^e(\delta)$ of the unknown parameter δ . In turn, δ is obtained through an estimation procedure based on the assumption that expectations are unbiased on average over the sample period. Using our usual notation, this assumption implies setting

$$\sum_{t=1}^T [\pi_t - \pi_{t|t-1}^e(\delta)] = 0 \quad (2)$$

where $t = 1, \dots, T$ indicates the *entire* sample over which the transformation from qualitative to quantitative expectations is performed. Having obtained an estimate $\hat{\delta}$ by plugging into (2) the actual inflation rates π_t , one can then construct for every t the quantitative value $\pi_{t|t-1}^e = \pi_{t|t-1}^e(\hat{\delta})$ of the one year ahead inflation expectation.

The procedure outlined above, while useful for obtaining quantitative inflation expectations, has an important drawback for the purpose of our empirical investigation. It requires us to assume, through (2), that inflation expectations are unbiased, which is exactly the property that I tried to establish in the previous section. To make this point clear, assume that we apply the procedure above to the sample period $t = 1990$ to $t = 2006$. Then, by construction, we are bound to obtain a series of inflation expectations whose average forecasting error during the deflationary path is zero. To overcome this problem, I propose to follow an alternative route, which I lay out in two steps.

For the first step I continue to assume that δ is constant over time, but I modify the estimation procedure in (2) using inflation rates and responses to the Consumer Confidence survey for the quarterly five-year window 1980Q1-1989Q4.⁷ The estimated $\hat{\delta}$ is then used, as in an out-of-sample procedure, to calculate inflation expectations $\pi_{t|t-1}^e(\hat{\delta})$ for the period 1990-2006. The left-hand panel in Figure 3 compares the result from this modified Carlson-Parkin procedure to the Consensus Forecast expectations. Two aspects deserve attention. First, the correlation between the two series is very high (87%), providing an indication that variations in Consensus expectations are coherent with variations in expectations obtained from other sources. The second is that in the second part of the sample, which corresponds to the period in which annual inflation had already stabilized around very low values (Figure 1, left

⁷While δ is estimated using quarterly data, Figure 3 reports only the inflation expectations for the fourth quarter of each year. Selecting the Consumer Confidence inflation expectation formed in the fourth quarter of each year is consistent with the decision to use the October survey from Consensus Forecast.

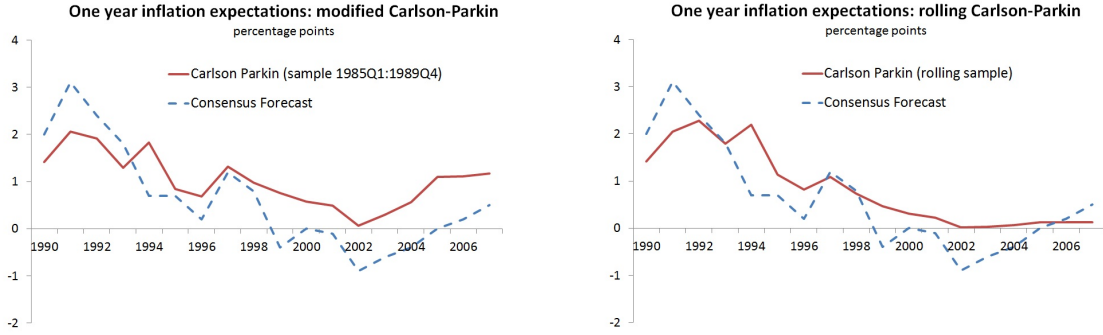


Figure 3: Inflation expectations from the modified Carlson-Parkin methods (solid line) and Consensus Forecast (dashed line) errors. These series are not corrected for the VAT tax increase. Source: data for the calculations are from Ueda (2010).

panel), there is a persistent level gap between Consensus and the modified Carlson-Parkin expectations.

One reasonable cause of the opening of this gap is that, as Japan progressively moved from a largely positive inflation regime to a regime with zero or slightly negative inflation, consumers started to communicate with greater *precision* whether they indeed expected inflation to be exactly at zero or if they expected it to have even small positive or negative variations. If this were the case, then the precision threshold δ would not be constant but instead would vary over time. To check for this possibility I introduce a second modification to the Carlson-Parkin method, which allows for the construction of inflation expectations $\pi_{t|t-1}^e(\delta_{t-1})$ based on a time varying parameter δ_{t-1} estimated using only past information. Specifically, I estimate δ_{t-1} by applying the condition (2) to the five years rolling window period ending in the year $t - 1$.⁸ The estimated $\hat{\delta}_{t-1}$ is then used to derive the expectation $\pi_{t|t-1}^e(\hat{\delta}_{t-1})$ at time $t - 1$. The results of this procedure are displayed in the right-hand panel of Figure 3. The main conclusion is that the persistent gap between the two series in the later part of the sample is now reduced. The average distance between Consensus and Consumer Confidence inflation expectations is now 0.25% (from 0.38% when the fixed estimation sample is used), while the correlation remains high at 86%. Hence, as the two series are highly correlated and very close to each other on average, using either of them makes no significant difference when calculating the idiosyncratic forecast.

3 Explaining Japan’s macroeconomic performance

In this section I revisit what we know and what is still largely a ground for debate on the causes of the Japanese deflation and GDP growth slowdown in the '90s. In doing so I also point out the implications, for the literature, of the findings on inflation expectations that were presented in the previous section.

⁸I also impose that δ must be equal to the maximum between the value obtained from the estimation using (2) and an exogenous lower bound which I set equal to 0.1%. The reason for this choice is that for the period 2002-2006 the value $\hat{\delta}$ estimated through (2) would be slightly negative.




Figure 4: Sources: Bank of Japan Flow of Funds and Cabinet Office.

I begin with what we know. The Japanese deflation and low growth (for which some qualifications are in order: see below) were preceded by a boom in investment and in firms' leverage, followed by a prolonged deleveraging and a sharp reduction in investment rates (Figure 4). Non-residential investment increased from 12% of GDP in 1983 to 18% in 1991. Over the same period the leverage of non-financial corporations increased from 142% of GDP to 175%. After 1991, the investment rate fell quite sharply, never to regain its previous peak. In the run-up to the cyclical peak, the debt accumulation and leverage of non-financial firms followed closely the pattern of investment. Differences in the evolution of the two variables instead emerged on the downward part of the cycle. In particular, leverage remained high for a prolonged period of time, up to 1998. Moreover, compared with just three years of declining investment rates, the deleveraging phase was slower and longer (about five years). After 2003, firms' leverage remained stable. It is widely accepted that the aforementioned investment and leverage cycles account for the strong GDP expansion as seen in Figure 6, in which the solid line indicates Japanese per capita GDP relative to the US.

To sum up, it is widely accepted that explaining the evolution of investment and firms' debt is the key to understanding what happened to Japan after 1991. It is in providing this explanation that the divergence of opinion and lack of consensus among economists enter. At the cost of an over-simplifying, I categorize the literature into two branches. One posits that the unimpressive growth performance in Japan was due to a 'lack of aggregate demand' (negative output gap); the other emphasizes that the marked GDP slowdown is explained simply by a reduction in productivity growth (and thus, by a slowdown in potential GDP).

The findings in the previous section have consequences for these two explanations of the Japanese growth slowdown. Here, the main conclusion to be drawn is the following. Once expected real interest rates are constructed using our measure of 'idiosyncratic' expected inflation (which, as we know, is very similar to realized inflation), then the 'lack of aggregate demand' approach requires the assumption of a much larger increase, during the '90s, in Japan's marginal propensity to save than is implied by a structural approach. To spell out this point I use a simple, overarching framework composed of two equations

$$\frac{\hat{C}_{t+1}}{\hat{C}_t} = \frac{1 + r_t^e}{1 + r^n} \quad (3)$$

$$1 + r^n \equiv \frac{1 + g}{\beta} \quad (4)$$

where r^n is the *natural real interest rate*, β is the subjective discount factor, \hat{C} is the deviation of consumption from a potential trend growth g , and r_t^e is the *expected* real (market) interest rate between time t and time $t + 1$.

Condition (3) is the usual Euler equation giving the first order condition for inter-temporal maximization of a household that is a net saver (and also for a household which is a net borrower, in the absence of binding borrowing constraints). Equation (4) defines the ‘natural real interest rate’ r^n , which is the real interest rate that would prevail in the market in the absence of deviation of actual output from its ‘potential’ level. Iterating forwards (3) and assuming that, as time goes to infinity, consumption deviations are zero, we obtain a version of the Euler equation that is slightly more convenient to use for our purposes,

$$\log \hat{C}_t = - \sum_{j=0}^{\infty} [\log(1 + r_{t+j}^e) - \log(1 + r^n)] \quad (5)$$

When using equation (5), a crucial issue is how to calculate the expected real interest rate r_t^e . In the absence of reliable measures for ex-ante inflation expectations, the expected real interest rate is often proxied by the *ex-post* real interest rate, calculated as the difference between the ex ante nominal interest rate and ex post inflation. Instead, when inflation expectations are available, the expected real rate is calculated as the difference between the ex ante nominal rate and the *ex ante* inflation expectation. Interestingly, in our case, if we adjust inflation expectations by netting out the global forecasting error $\bar{\varepsilon}_t^g$, then the expected real rate r_t^e calculated using ex post inflation is virtually identical to the one calculated using the ex ante adjusted inflation expectations. The value of the ex ante expected real rate r_t^e resulting from the adjustment procedure is depicted in Figure 5⁹. This ‘adjusted’ measure of the ex-ante real interest rate is an interesting benchmark, for the following reasons. First, it gives the counterfactual evolution of the ex ante real rate in Japan had agents made no error in estimating the *global* component of inflation. Second, as we have shown that agents made only small errors in estimating the *idiosyncratic* component of inflation (so that $\pi_t + \bar{\varepsilon}_t^i \approx \pi_t$), then our adjusted real rate is also very close to the ex ante real rate calculated under the assumption of perfect foresight. Finally, how does our adjusted ex ante real rate compare with the unadjusted ex ante real rate, i.e. the rate calculated using the original Consensus inflation expectations, which include also the *global* forecasting error? The answer is straightforward once we look at the right-hand panel of Figure 2, which shows that, in our sample, the *global* forecasting error is always negative, apart from the first and last year, when it is positive but relatively close to zero. It follows then that the adjusted real rate is always larger than the unadjusted one, except at the ends of the sample, when the two rates are close to each other.

3.1 Lack of aggregate demand

According to one interpretation, since the bust of the investment cycle in the late ’80s, Japan has been suffering from a prolonged negative output gap. The boom years of the period 1986-1991 were a combination of two underlying trends, a catching up of Japan with the

⁹Using our notation, the adjustment procedure is then equivalent to using the value $\pi_t + \bar{\varepsilon}_t^i$ as our measure of ex-ante expected inflation for the period t .



Figure 5: For a given year t , the real rate is calculated as the difference between the three-year nominal rate on government bonds, prevailing in the market three years before t , and the average inflation expected over the three years leading up to t . Expected inflation is calculated including only the idiosyncratic error component. Sources: Thomson Reuters Datastream.

technological frontier (i.e. the US), and a bubble economy driven by over-investment in low-yield inefficient projects. Because of the bubble component, part of the high Japanese growth in the period 1986-1991 was indeed unsustainable, and a downward correction was bound to occur. In this sense, a period of below-trend GDP growth, potentially reflected in a reduction in Japanese per capita GDP relative to the US, is not surprising. What is surprising, instead, is that the Japanese growth was so low for such a prolonged time that even the genuine catching-up of the '80s was eventually undone. After 1998, Japanese per capita GDP remained, relative to the US, at the same stable level as in the period 1980-1985. Some authors consider that a complete undoing, due to purely structural reasons, of the catching-up phase is so unreasonable that we should posit instead that Japan had been operating below potential since at least 1993. In turn, the negative output gap is interpreted (or, rather, is defined) as a lack of aggregate demand, which is assumed to be the ultimate cause of the Japanese deflation. The policy prescription to close the output gap, it is argued, is thus public stimulus.

But what kind of stimulus? Here the authors are divided. Posen (1998), for instance, argues that with a zero nominal interest rate monetary stimulus becomes ineffective, and the only policy option that can work is a strong fiscal stimulus. According to Posen, «Fiscal policy works, when it is tried», even when the fiscal impulse has a *one-time* nature, as the case of the 1995 stimulus in Japan would prove. The effectiveness of that stimulus, it is argued, is not surprising as «the Great Depression offered many important lessons for macroeconomic policy. One of the least controversial is that there is a role for government revenues and expenditures to automatically move in the opposite direction of the business cycle» (page 30).

In his influential paper assessing the Japanese situation, Krugman (1998) offers quite a different view: «My point is that the Great Depression, which is the usual, indeed perhaps

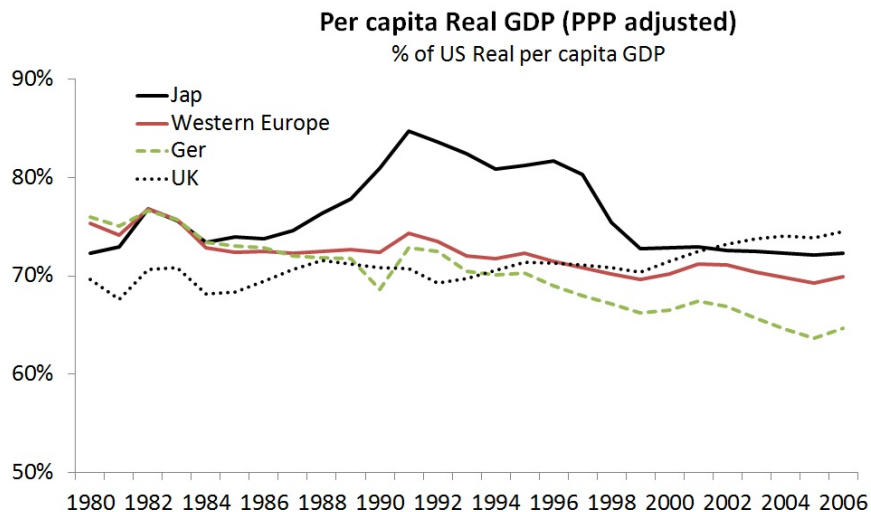


Figure 6: Source: Maddison (2008), *Historical Statistics of the World Economy: 1-2008 AD.*

the sole, motivating view that a *one time* fiscal expansion can produce a sustained recovery, does not appear to fit the story line well» (page 1960, italics added). Rather than focusing on a fiscal stimulus, Krugman (1998) calls for a monetary expansion. In Krugman’s view, the negative output gap in Japan was due to a lack of demand created by the discrepancy between a negative ‘natural real interest rate’ and the observed positive real rate (stemming from a non negative nominal rate and zero or negative inflation). A permanent monetary expansion, aimed at raising inflation well above the price stability target, would have been the preferred policy to follow, in order to increase inflation enough to push the real interest rate negative levels. Even at the zero lower bound, he argues, monetary policy can raise inflation. All the central bank has to do is to «commit to be irresponsible», i.e. commit to give up its price stability goal by permanently increasing the money stock¹⁰.

A commitment to being irresponsible is the wrong way to go, according to Koo (2009). In his view, an irresponsible central bank risks undermining the credibility of a currency, with consequences that are disastrous compared with those of a mild deflation. Moreover, Koo (2009) rejects Krugman’s assertion that the exact identification of the actual causes of the excessively low aggregate demand is of secondary importance, as a macroeconomic stimulus is likely to cure the problem in any case. Instead, Koo argues that it is crucial to understand the roots of the deflation, which he identifies as the deleveraging process of Japanese firms. Japan’s slow growth is seen as having been the consequence of a balance sheet recession. The attempt by firms to pay off their pre-existing and excessive debt, built up during the bubble years, is what caused the fall in aggregate demand and in the price level in Japan. Fiscal expansion should have been used to counterbalance this fall in private demand. The argument of a deleveraging-driven recession was later formalized by Krugman and Eggertsson (2012)

¹⁰In an open economy, Svensson (2013) argues that an increase in the inflation rate, and a consequent lowering of real interest rates, can be achieved through a central bank’s commitment to devalue the currency.

and by Eggertsson and Mehrotra N. (2014).¹¹ As Koo points out, the important result of having identified the cause of the Japanese deflation and low growth is that we can predict that they will end when the deleveraging process is over and leverage has thus stabilized (see Figure 4).

I conclude this section by discussing the consequences, for demand-side explanations, of the finding that the deflation in Japan was anticipated. Recall that explanations based on lack of demand argue that the output gap in Japan was persistently negative over the period considered. In terms of equation (5) this translates into assuming negative values for $\log \hat{C}_t$. To obtain this result we need to assume that there was a sequence of positive deviations $r_t^e - r^n$ between the *expected* market and the natural real interest rate. Because, even after removing the global inflation shock, real rates were in Japan were still *expected* to fall (Figure 5), to obtain positive values for $r_t^e - r^n$ the natural rate r^n must have decreased faster than r_t^e .¹² But, as we have seen, though, explanations based on lack of demand tend to downplay the presence of a significant slowdown in potential GDP growth in Japan. It follows that a significant fall in g cannot be the reason for the required rapid fall in r^n , which must then be due to a sharp rise in β , i.e. a large exogenous increase in agents' desire to save.¹³

3.2 Structural slowdown

Part of the literature dismisses output gap based explanations of the Japanese growth slow-down, and focuses on structural aspects. To introduce this point let us go back to Figure 6, in which I also plot the evolution of per capita output, relative to the US, of Western European countries, with Germany and the UK as two polar sub-cases. Following the partial European catch-up with the US during the post-war period, after 1980 the relative per capita output of the European economies has remained roughly stable, at just over 75% of the US level. Quite strikingly, over the long-run a similar narrative of a partial and interrupted catching-up with the technological frontier appears to hold for Japan as well. The only difference between Western Europe and Japan was that Japan had an above trend investment driven growth in the years 1987-1991, followed by a compensating below trend growth in 1992-1998.

Hayashi and Prescott (2002) show that Japan's growth slowdown could be attributed to a slowdown in the growth rate of Japanese productivity and to a change in labor market regulations. The evolution of inflation is not given any particular weight as a cause or as a symptom of the stagnation, and is thus left out of the analysis. According to Hayashi and Prescott (2002), the investment boom of the years 1987-1991 was immediately followed by a progressive fall in the return on capital, which is a possible indication that, at least from an ex-post perspective, the previous capital build-up was excessive. Also, a legislated reduction in the length of the working week from 48 hours to 40 hours, introduced gradually between 1988 and 1998 provided a one-off effect causing a temporary reduction of the growth rate of potential output to below trend values. According to this overall line of reasoning, the bust in the investment rate in Japan should not be surprising, in fact «Assuming that TFP growth remains low, Japan cannot rely on capital deepening for growth in per-working-age-person

¹¹These papers show how an exogenously imposed deleveraging process can give rise to a negative 'natural real interest rate' situation.

¹²As noted previously, if the global inflation shock is not removed, and the 'unadjusted' expected real rate is used, we conclude that expected real rates has fallen even faster in Japan

¹³As seen, if exogenous borrowing constraints suddenly become binding, then some agents are exogenously forced into saving, replicating the effects on the model of an exogenous increase in β .

output as it did in the past, as the Japanese capital stock is nearly at its steady state value» (page 220).

Hayashi and Prescott (2002) also present an empirical analysis indicating that after 1991-1997 there was no credit crunch in Japan, and thus firms' investment was not constrained by the availability of credit during that period. These conclusions are coherent with the findings of Motonishi and Yoshikawa (1999) and are strongly endorsed by Koo (2009). Authors generally agree that if a credit crunch did occur, it was relatively short-lived and took place amid the 1997 recession.

The structuralist explanation of Hayashi and Prescott (2002) does not provide a detailed causal justification for the the observed slowdown in productivity growth in Japan. Caballero, Hoshi and Kashyap (2008), and Hoshi and Kashyap (2004) go deeper into this issue arguing that the productivity slowdown was caused by banks' lending practices. According to the authors, after the asset and investment bubble burst, and with the ensuing banking crisis, Japanese banks adopted a policy of debt evergreening. This policy kept insolvent firms afloat and prevented more efficient firms from gaining market shares and replacing the old and inefficient ones. The drag on the economy's aggregate productivity was substantial and was aggravated by the reluctance of policy-makers to force banks to clean up their balance sheets, putting an end to their zombie-lending practice. The delay with which the Japanese authorities dealt with the stock of non-performing loans in the banking sector may also explain why firms' leverage remained high until 1998, well beyond the conclusion of the investment cycle (Figure 4). For a detailed account of policymakers' actions amid the Japanese banking crisis see Nakaso (2001).

Recapitulating, the evolution of the expected real interest rate r_t^e , constructed from the 'idiosyncratic' component of expected inflation, has its own specific story to tell when we look at the performance of Japanese growth through the lens of equation (5). The main conclusion is that a structural explanation requires much smaller variations in the propensity to save (proxied by β) than a lack-of-demand approach. The reasons are essentially two. First, for supporters of a structural explanation, there was no significant output gap in Japan, meaning that $\log \hat{C}_t = 0$. Hence, for (5) to hold, it is now enough for the natural rate r^n to fall exactly as fast (and not faster, as required if we postulate a negative output gap) as r_t^e . Second, structural explanations point out that the rate g of potential growth decreased significantly in Japan. It follows that the (small) required fall in r^n can to a large extent be explained by a fall in g , rather than by an increase in β .

To see this last point more clearly, consider the following observations. From Figure 5 note that over the sample period considered in this paper the expected real rate r_t^e in Japan fell slowly by 3 percentage points overall.¹⁴ Moreover, from Figure 6 we can see that the fall in the relative Japanese per capita GDP since 1992-1993 was about 5 percentage points. If we were to use a pure *demand side* approach, then we could postulate that, during that same period, Japan experienced an output gap of about -5%. Using these observations, let us perform the following back-of-the-envelope exercises. Consider an economy in which at time t the expected real interest rate r_t^e falls instantly by 3 percentage points and stays low for a period of five years.¹⁵ Postulate, moreover, that at time t the economy suffers from an output gap equal to

¹⁴Incidentally, this would also be the overall fall in the expected real rate if the 'unadjusted' ex ante real rate were used.

¹⁵For simplicity, I am condensing the slow 3 percentage points fall in r_t^e , which in our sample took place over a period of about fifteen years into a sudden fall lasting five years.

$\hat{C}_t = -5\%$ and that at the end of the five years the output gap is closed (i.e. $\hat{C}_{t+j} = 0$ for $j \geq 5$, which by (5) implies $r_{t+1}^e = r_{t+j}^n$ for $j \geq 5$). From equation (5) we then have, approximately, that $r_{t+j}^e - r_{t+j}^n = 1\%$ over the five year period which, together with the 3 percentage point fall in r^e , implies that the natural real rate r^n must have fallen by 4 percentage points during the period under consideration. As we assumed that potential growth g remains unchanged, then the 4 percentage points decrease in r^n must be completely absorbed by a rise in β . Let us now redo the exercise using a *structural approach*. In this case, as the assumed output gap is zero, we just need to posit that the overall 3% fall in r_t^e must be matched by an overall 3% fall in r^n . If, moreover, we assume that potential growth falls suddenly, for instance, by 2 percentage points, then the sudden increase in the propensity to save β must now account only for a remaining 1% reduction in r^n , rather than the much larger 4% calculated under the *demand side* approach. It is worth noting that OECD data indicate that the saving rate¹⁶ of Japanese households has *fell* in the period 1990-2006. Together with the absence of any significant deleveraging in the household sector (Figure 4), these observations cast doubts on the possibility that a large increase in the propensity to save of households is at the root of an underperformance of GDP growth in Japan driven by lack of demand.

4 Policy lessons and conclusions

The bursting of the Japanese housing and stock market bubble in the early '90s and the roots of the subsequent banking crisis, deflation and low growth, have been the subject of much debate. Policymakers have at times been harshly criticized for “not doing enough” to avert what to some looks like a disastrous outcome and a large cumulative loss of output. Some advanced countries, especially in Europe, are now experiencing similar situations, with a prolonged period of low growth and very low inflation. Over the past few years, private sector deleveraging and banking difficulties have been the most important problems policymakers have struggled with. In this context, the Japanese experience is brought once more to the fore in the search for useful policy lessons. This note highlights two lessons that can be useful to inform policy discussions in the current context. The first is about the nature of the Japanese deflation, while the second is about the identification of policy responses.

Over the period 1991-2006 short- and medium-term inflation expectations consistently overshoot realized inflation. Hence, ex post real interest rates in Japan were from 50 to 100 basis points higher than lenders and borrowers had expected ex-ante. As the years 1991-2006 also coincided with a period of significant corporate deleveraging in Japan, the higher ex-post real rates may have contributed to exert significant pressure on debtors, with possible consequences for the dynamics of aggregate demand and the price level. One natural question then arises: to what extent did disinflation surprise the market? The answer provided by the paper is a qualified ‘no’. During the years 1991-2006, the underperformance of realized inflation was a global phenomenon that occurred in all The main advanced countries. Among them, however, only in Japan did deflation take root, while in all the other advanced countries inflation remained on average close to 2%. It is therefore fair to argue that there must be something special, or ‘idiosyncratic’, about the Japanese deflation. To further investigate this point the paper looks again at inflation expectations in Japan. This time, however, the global forecasting error is removed, so that only idiosyncratic factors are left. Once the effects of

¹⁶Total household net savings as a share of GDP.

global factors are eliminated, short- and medium-term inflation expectations indicate that the Japanese deflation during 1991-2006 was fully and accurately anticipated. Underlying inflation expectations got disanchored very soon along the deflationary path from their previous 2% level, notwithstanding the progressive reduction to zero of the Bank of Japan's policy rate. This early disanchoring of inflation expectations has had interesting consequences. It suggests, in particular, that the Japanese deflation may have had a self-fulfilling nature.

From a policy-making perspective, self-fulfilling deflations have quite different properties from deflations induced by shocks to the natural real interest rate (as in Eggertsson and Woodford [2004]). Self-fulfilling deflations cannot be stopped by aggressive reductions in the policy rate or by forward guidance on its future evolution. Moreover, in self-fulfilling deflations government spending does not have any appreciable effect on GDP, as along self-filling deflation paths fiscal multipliers should be close to zero (Mertens and Ravn [2012] and Piazza [2015]). Thus, the first main lesson to be drawn from the Japanese experience is that self-fulfilling deflations may indeed occur and that their prompt identification matters in terms of choosing the appropriate policy response. The policy response must necessarily be non-conventional, in the sense that it must go beyond the standard use of the short-term interest rate policy by the central bank.¹⁷

The second lesson comes from looking beyond the evolution of inflation and considering some additional issues as well, such as the growth performance of Japan and the macroeconomic effects of leverage cycles and of banking crises. Here, for policymakers seeking an uncontroversial historical case that could inform their formulation of current policies, the lesson to be drawn is quite dismal. After over twenty years of debate, the literature is still very much divided on how to interpret the Japanese case and the role that has been played (and, going forward, should be played) by policymakers. The split that exists in the literature is first of all between demand-side and structuralist interpretations.

Structuralist interpretations point out that the Japanese low growth can be satisfactorily explained by a shift in the dynamics of aggregate productivity, due for instance to the inability of the banking sector to allocate credit efficiently. On the contrary, demand-side explanations point to the presence in Japan of a persistent negative output gap, which calls for active intervention by policymakers. This strand of literature is, in turn, internally divided when it comes to specifying what stimulus policies are needed, with some authors favouring fiscal deficits and others monetary expansions. Overall, whether the fiscal or the monetary expansion option is ultimately embraced, supporters of stimulus policies agree that in Japan policymakers did too little: larger deficits or far larger monetary expansions were needed.

The divide, on the Japanese case, between supporters of demand-side explanations, and between them and structuralists, runs deep into the foundations of macroeconomic theory, and has a parallel in the different views on the causes of the Great Depression. Supporters of demand-side explanations argue that an excessive desire to save by private agents in Japan, following the initial burst of the bubble, should have been offset by higher government spending or lower real interest rates. The paper points out that, for this 'lack of aggregate demand' approach to hold, the increase in the private marginal propensity to save in Japan during the '90s would have had to be large, much larger than implied by a structural approach. Empirically, there is no evidence of an increased desire to save by the Japanese household sector since the '90s.

¹⁷Discussing the role and effectiveness of non-conventional policies is beyond the scope of this note.

In the end, the uncontroversial points on which the literature appears to agree are relatively few and are the following. The real estate and stock market bubble burst after a period of very high investment, which was at least in part channeled towards inefficient projects. Over this period, firms' leverage increased substantially. The bust was then marked by a permanent and sizeable fall in the investment rate, coupled with an increase in firms' propensity to save, which in turn led to a long-lasting deleveraging process in the corporate sector. For a long time after the bust, and at least until 1997, firms do not appear to have been constrained in their investment decisions. Thus, the initial fall in the investment rates was not due to lack of access to bank credit or to a credit crunch. Still, the role of the banking channel is almost universally called into question, although specific interpretations differ. Some emphasize the initial laxity in lending policies that contributed to create the bubble phase. Others point to the slow restructuring of banks and the attendant effects on aggregate productivity growth.

After exploring the history of the Japanese deflation between 1990 to 2006 it seems natural to add, as a conclusion, some considerations on the more recent past. We can ask ourselves how the findings of the paper help us to interpret *Abenomics*, the strategy for economic growth based on *three arrows* (monetary, fiscal and structural) inaugurated by the prime minister Shinzo Abe in late 2012. The answer depends very much on our variable of interest. If we focus on GDP, thanks to its mix of demand-side (fiscal and monetary expansion) and supply-side (structural reforms) policies, *Abenomics* should lead to an expansion in real GDP, irrespective of which between the 'lack of demand' or 'structural slowdown' approach is the correct one. For instance, if growth in Japan is held back by structural problems, then demand side policies would be counterproductive (e.g. by raising the public debt). Still, the shortcomings stemming from misguided demand-side 'arrows' could in principle be more than compensated by the beneficial effects of structural policies that increase potential growth (and thus, indirectly, also improve the sustainability of the public finances). This positive conclusion regarding the effect of *Abenomics* on real GDP is of course valid only as long as *all* the 'arrows' are indeed shot.¹⁸ If we move on to consider the long-term impact of *Abenomics* on inflation, then the analysis in the paper suggests that we should distinguish between two situations. If the low inflation is caused by the economy operating below potential, then a mix of temporary fiscal and monetary expansions should be enough to raise the inflation rate (and close the output gap). If, instead, the low inflation is due to a self-fulfilling equilibrium in which inflation expectations become disanchored from the level desired by policy-makers, then *temporary* demand-side policies are not able to raise the inflation rate *permanently*. In this case, a permanently positive inflation can be achieved only through a 'sunspot' shock (i.e. a purely expectational shock, unrelated to the fundamentals of the economy) that suddenly resets and coordinates long-run inflation expectations back to the desired level. *Abenomics*, with its widespread impact on the media and on key volatile financial variables such as the stock market and the exchange rate, may be the needed, albeit fragile, 'sunspot' (or 'shock-and-awe' policy) capable of anchoring inflation expectations to the announced 2% level.

¹⁸At the present time, relatively little appears to have been done on the structural side of the *Abenomics* strategy.

5 References

- Benhabib J., Schmitt-Grohe S., Uribe M. (2001), 'The Perils of Taylor Rules', *Journal of Economic Theory*, 96, 40-69.
- Caballero R. J., Hoshi T., Kashyap A.K. (2008), 'Zombie Lending and Depressed Restructuring in Japan', *American Economic Review*, 98, 1943-77.
- Carlson J.A., Parkin M. (1975), 'Inflation expectations', *Economica*, 42, 123-138.
- Cochrane J. (2011), 'Determinacy and Identification with Taylor Rules', *Journal of Political Economy*, 119, 565-615.
- Cochrane J. (2013), 'The New-Keynesian Liquidity Trap', mimeo, October 24.
- De Pooter M., Robitaille P., Walker I., Zdinak M. (2014), 'Are Long-Term Inflation Expectations Well Anchored in Brazil, Chile and Mexico?', Discussion Paper, Board of Governors of the Federal Reserve System.
- Eggertsson G, Mehrotray N. (2014), 'A Model of Secular Stagnation', mimeo, Brown University.
- Eggertsson G. B., Woodford M. (2003), 'The Zero Bound on Interest Rates and Optimal Monetary Policy', *Brookings Papers on Economic Activity*, 139-211.
- Eggertsson G. B., Woodford M. (2004), 'Policy Options in a Liquidity Trap', *American Economic Review*, 94, 76-79.
- Hoshi T., Kashyap A.K. (2004), 'Japan's Financial Crisis and Economic Stagnation', *Journal of Economic Perspectives*, 18, pages 3-26.
- Krugman P. (1998), 'It's Baaack: Japan's Slump and the Return of the Liquidity Trap', *Brookings Papers on Economic Activity*, 2, 137-205.
- Koo R. (2009), *The Holy Grail of Macroeconomics*, John Wiley & Sons.
- Krugman and Eggertsson (2012), 'Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach', *The Quarterly Journal of Economics*, 127, 1469-1513.
- Hayashi F., Prescott E.C. (2002), 'The 1990s in Japan: A Lost Decade', *Review of Economic Dynamics*, 5, 206-235.
- Maddison A. (2008), 'Historical Statistics of the World Economy: 1-2008 AD'.
- Mertens K., and Ravn M.O. (2012), 'Fiscal Policy In An Expectations Driven Liquidity Trap', mimeo, University College London.

Motonishi T., Yoshikawa H. (1999), 'Causes of the long stagnation during the 1990s: financial or real?', *NBER Working Paper* 7351.

Nakaso H. (2001), 'The financial crisis in Japan during the 1990s: how the Bank of Japan responded and the lessons learnt', *BIS paper* 6.

Piazza R. (2014), 'Self-fulfilling deflations', mimeo, Bank of Italy.

Posen A. S. (1998), *Restoring Japan's Economic Growth*, Institute For International Economics.

Svensson L. E. O. (2003), 'Escaping From A Liquidity Trap And Deflation: The Foolproof Way And Others', *Journal of Economic Perspectives*, 17, 145-166.

Ueda K. (2010), 'Determinants of Households' Inflation Expectations in Japan and the United States', *Journal of the Japanese and International Economies*, 24, 503-518, 2010.

Visco, I. (1984), *Price Expectations in Rising Inflation*, Amsterdam: North-Holland.