

JSPS Grants-in-Aid for Scientific Research (S)
Understanding Persistent Deflation in Japan

Working Paper Series

No. 096

June 2017

Why Has Japan Failed to Escape from Deflation?

Kota Watanabe
Tsutomu Watanabe

UTokyo Price Project
702 Faculty of Economics, The University of Tokyo,
7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan
Tel: +81-3-5841-5595
E-mail: watlab@e.u-tokyo.ac.jp
<http://www.price.e.u-tokyo.ac.jp/english/>

Working Papers are a series of manuscripts in their draft form that are shared for discussion and comment purposes only. They are not intended for circulation or distribution, except as indicated by the author. For that reason, Working Papers may not be reproduced or distributed without the expressed consent of the author.

Why Has Japan Failed to Escape from Deflation?

Kota Watanabe*

Tsutomu Watanabe[†]

June 27, 2017

Abstract

Japan has failed to escape from deflation despite extraordinary monetary policy easing over the past four years. Monetary easing undoubtedly stimulated aggregate demand, leading to an improvement in the output gap. However, since the Phillips curve was almost flat, prices hardly reacted. Against this background, the key question is why prices were so sticky. To examine this, we employ sectoral price data for Japan and seven other countries including the United States, and use these to compare the shape of the price change distribution. Our main finding is that Japan differs significantly from the other countries in that the mode of the distribution is very close to zero for Japan, while it is near 2 percent for other countries. This suggests that whereas in the United States and other countries the “default” is for firms to raise prices by about 2 percent each year, in Japan the default is that, as a result of prolonged deflation, firms keep prices unchanged.

JEL Classification Number: E31; E5

Keywords: deflation; price stickiness; Phillips curve; inflation expectations; inflation norms; quantitative easing; menu cost models; sectoral price data

*Canon Institute for Global Studies (CIGS) and University of Tokyo. E-mail: watanabe.kota@canon-igs.org. We would like to thank Kosuke Aoki, Christian Broda, Ippei Fujiwara, Hideo Hayakawa, Kazuo Monma, Roberto Rigobon, Shigenori Shiratsuka, Kazuo Ueda, Hirohide Yamaguchi, Hiroshi Yoshikawa, and David Weinstein for valuable comments on earlier versions of this paper, and Takuto Murase for preparing part of the dataset used in the paper. This research forms part of the project on “Understanding Persistent Deflation in Japan” funded by a JSPS Grant-in-Aid for Scientific Research (No. 24223003).

[†]Graduate School of Economics, University of Tokyo. E-mail: watanabe@e.u-tokyo.ac.jp, Website: <https://sites.google.com/site/twatanabelab/>.

1 Introduction

From the second half of the 1990s onward, Japan suffered a period of prolonged deflation, in which the consumer price index (CPI) declined as a trend. During this period, both the government and the Bank of Japan (BOJ) tried various policies to escape from deflation. For instance, from 1999 to 2000, the BOJ adopted a “zero interest rate policy” in which it lowered the policy interest rate to zero. This was followed by “quantitative easing” from 2001 until 2006. More recently, in January 2013, the BOJ adopted a “price stability target” with the aim of raising the annual rate of increase in the CPI to 2 percent. In April 2013, it announced that it was aiming to achieve the 2 percent inflation target within two years and, in order to achieve this, introduced Quantitative and Qualitative Easing (QQE), which sought to double the amount of base money within two years. Further, in February 2016, the BOJ introduced a “negative interest rate policy,” in which the BOJ applies a negative interest rate of minus 0.1 percent to current accounts held by private banks at the BOJ, followed, in September 2016, by the introduction of “yield curve control,” in which the BOJ conducts JGB operations so as to keep the 10-year JGB yield at zero percent. See Table 1 for an overview of recent policy decisions made by the BOJ.

However, these efforts by the BOJ to escape from deflation have not been successful. That is, while the year-on-year (*y-o-y*) rate of change in the CPI excluding fresh food did initially turn positive, price developments subsequently weakened again. Specifically, after leaving negative territory and returning to zero in May 2013, the *y-o-y* rate of change in the CPI turned positive in June that year and rose to 1.5 percent by April 2014 (excluding the effects of the consumption tax hike in April 2014). This was a direct consequence of the depreciation of the yen in 2012-2015 (when the yen depreciated about 40 percent vis-à-vis the US dollar from 78 yen/dollar in 2012 to 125 yen/dollar in 2015), which was induced by monetary easing by the BOJ. However, since then, inflation has gradually slowed as consumption demand declined and, if the effect of the consumption tax hike is excluded, has fallen to under 1 percent since October 2014. Given this, it looks unlikely that the BOJ’s 2 percent target will be achieved in the foreseeable future.

This paper investigates why it has been so difficult to escape from deflation. To do so, we focus on the fact that Japan’s deflation since the mid-1990s consisted of an extremely mild decline in prices. The largest monthly rate of *y-o-y* decline in the CPI during this period was only around 2 percent, and for the period as a whole the average was slightly less than 1 percent. Therefore, even though it was deflation, it was very mild and did not turn into the kind of severe deflation one might associate with the term; it certainly also did not turn into the kind of deflationary spiral

many policymakers had feared. Of course, the fact that it was possible to avoid severe deflation itself is a good thing. However, from the perspective of escaping from deflation, it is possible that because the decline in prices was “neither here nor there,” the market forces to turn the price decline into an increase were weak. Put differently, the fact that prices did not decline much even during the period of weak aggregate demand means that the Phillips curve was almost flat. It is this flattening of the Phillips curve that made it more difficult to escape from deflation. Given that the flattening of the Phillips curve was at least partially caused by an increase in price stickiness, the key question is why prices became so sticky in the 1990s. Against this background, the aim of this paper is to measure to what extent price stickiness increased over the last two decades and examine why this happened.

The analysis in this paper makes extensive use of sectoral price data. Japan’s CPI is constructed from items such as margarine and shampoo, and the total number of items is 588. For each of these items, about 570 prices are collected by price collectors of the Statistics Bureau of Japan every month. Therefore, the total number of prices collected each month reaches about 250,000. Roughly speaking, the headline CPI announced each month is the average of this large number of prices, and averages such as this certainly are an important statistical indicator. However, the average is not the only useful indicator. Other statistical indicators - for example, higher-order moments as well as the mode of the distribution of price changes for individual items - also provide a wealth of information on developments in prices. This paper seeks to take advantage of such information by looking at what kind of shape the distribution of price changes takes and how this distribution evolves over time. We supplement this analysis by comparing the shape of the item-level price change distribution for Japan with those for seven other countries including the United States to determine whether and why prices are stickier in Japan than the other countries.

The rest of the paper is organized as follows. Section 2 provides a description of the salient features of Japan’s deflation, in particular the mildness of the price decline and the flattening of the Phillips curve. Section 3 then examines the shape of the distribution of price changes of individual items in the consumer price statistics and how this shape changed over time. Section 4 examines the relationship between the share of items for which the rate of change is close to zero and the overall CPI inflation rate. Section 5 then conducts a cross-country comparison of the price change distribution in order to identify causes behind the prolonged deflation in Japan, while Section 6 provides a summary and discusses some policy implications.

2 Characteristics of Japan's Deflation since the Mid-1990s

Japan's deflation has two important characteristics. First, deflation has lasted for a long time. Prices declined for two decades from the mid-1990s, so that it has clearly been prolonged. Second, however, in terms of the pace of deflation, even when prices were falling relatively fast, they never fell at a rate beyond 2 percent, and the average for the period is only 1 percent. In this sense, deflation was mild.

These two characteristics become even clearer when comparing Japan's experience with the deflation in the United States during the Great Depression. The deflation during this period was severe, with prices falling at an annual rate of more than 8 percent. On the other hand, the deflation during the Great Depression ended in about three years, so the duration was short. The reasons for these differences in the deflation rate and the duration are not easy to identify because the two deflationary episodes concern not only different countries but also different eras. That being said, these differences may stem, at least partially, from differences in the price setting behavior of producers and distributors in the two countries.

As pointed out by Gordon (1981), during the Great Depression in the United States, firms quickly adjusted prices in response to changes in demand and supply. That is, prices were flexible. On the other hand, in Japan price flexibility has fallen in recent years, so that even when demand and supply conditions change, there has been a strong tendency for firms not to change prices immediately. For example, Abe et al. (2008), based on a survey of Japanese manufacturers, report that more than 90 percent of firms responded that they "do not immediately change prices even if demand and supply conditions change." Moreover, in a survey published in the Annual Report on the Japanese Economy and Public Finance 2013, only 21 percent of firms answered that they fully passed on increases in marginal costs to consumers, while the majority indicated that they could not fully pass on such cost increases.¹

A possible reason that Japan's deflation since the mid-1990s was mild is that firms did not see a need to lower prices at a quick pace. For example, it is possible that firms did not experience any change in marginal costs, so that they did not have to change prices. However, it is very likely that firms *did* experience a substantial change in marginal costs. This can be seen by looking at the evolution of the Phillips curve for Japan (Figure 1). In the 1970s and 1980s, there existed a relationship such that when unemployment decreased, inflation increased. That is, an increase in demand would lead to an increase in production and a fall in unemployment; this, in turn, increased marginal production costs, pushing up inflation. However, from the 1990s onward, this relationship

rapidly weakened. The slope of the Phillips curve became much shallower and for the period since 2000 has been close to zero. Looking at the period since 2000 in more detail, even though the unemployment rate fluctuated within the range of 3.9 percent to 5.4 percent, the rate of change in the CPI hovered in a narrow range between -1.4 percent (in 2009) and 1.4 percent (in 2008), and in many years remained close to zero. Even during the global financial crisis, which resulted in a sharp rise in the unemployment rate, the CPI did not fall much.

The flattening of the Phillips curve has important implications with regard to overcoming deflation. Spurring demand through QQE and fiscal stimulus measures should help to reduce unemployment and lead to an improvement in the output gap. Yet, if the Phillips curve is more or less flat, a reduction in the unemployment rate or improvement in the output gap is less likely to result in price increases, making it more difficult to overcome deflation. Watanabe (2013) shows, through a numerical exercise based on the estimated slope of the Phillips curve, that the output gap would have to increase by 5.3 percentage points each year in order to achieve the 2 percent inflation target in two years. This means, assuming that Japan's potential GDP growth rate is 1 percent, it would require real GDP to grow at a rate of 6.3 percent per year, which is clearly unrealistic. In other words, overcoming deflation only by stimulating demand is impossible; instead, it is vital to return the Phillips curve to its original slope or to shift the Phillips curve upward.

3 The Distribution of Item-Level Price Changes

The frequency with which firms adjust prices is an important parameter determining the slope of the Phillips curve. For example, the coefficient associated with the output gap in the Phillips curve derived in a Calvo (1983) setting is positively related with the probability of price adjustment, which is assumed to be exogenously given and constant over time. If the probability of price adjustment is low, changes in economic conditions (as represented by changes in the unemployment rate) will not be factored into prices quickly, so that the slope of the Phillips curve will be shallow.

This suggests that the change in the slope of the Phillips curve in Japan may have been caused by a decline in the frequency of price adjustment. This raises a host of questions such as: Did the frequency of price adjustment actually decline? And if it did decline, for which items is this the case? What are the causes of such a decline in the frequency of price adjustment? And has the adjustment frequency started to increase again since April 2013, when QQE began? In order to find answers to these questions, this section looks at price developments in the individual items making up the CPI. Specifically, the shape of the distribution of price changes of individual items and the evolution of that distribution over time are examined.

3.1 The shape of and changes in the item-level price change distribution

Figure 2 shows the distribution of price changes (monthly y-o-y changes) for each of the 588 items making up the CPI. The vertical axis shows the sum of the CPI weights of the items included in each bin on the horizontal axis. The blue line shows the distribution for March 2014, when the inflation rate was relatively high due to monetary easing, while the black line shows the distribution just before the start of the BOJ's policies to overcome deflation (December 2012). Looking at the blue line, the density is highest in the bins from -0.75 percent to -0.25 percent and from -0.25 percent to +0.25 percent, and these two bins alone make up about 50 percent of the total CPI weight.

The rate of change of the CPI excluding fresh food for March 2014 was +1.3 percent, and the figure clearly shows that this increase does not reflect a uniform increase in prices of 1.3 percent; rather, there is substantial heterogeneity across items in the rate of price change. Comparing the distribution for March 2014 with that for December 2012 shows that the former has a fatter upper tail (for items whose prices increased) and a thinner lower tail (for items whose prices decreased) than the latter, and it is this that is responsible for the rise in the CPI inflation rate from -0.2 percent in December 2012 to +1.3 percent in March 2014. On the other hand, the shape of the central part of the distribution (representing items whose prices remained unchanged) is almost the same.

In order to examine when the distribution started to take this shape with the peak in the vicinity of zero, Figure 3 plots the share over time of items for which the rate of price change was close to zero. Specifically, four different definitions for "close to zero" are used: ± 0.5 percent, ± 0.3 percent, ± 0.1 percent, and 0 percent. Then, for each year, items for which the rate of change was close to zero are identified and the CPI weights of those items are summed up, and it is this value which is shown on the vertical axis in the figure. Starting with the purple line, which defines "close to zero" as ± 0.5 percent, the share of items whose rate of price change was close to zero was only 10-20 percent during the period of high inflation in the 1970s. This indicates that during this inflationary period, almost all items registered price changes every year and only for a limited number of items was the rate of change close to zero, presumably due to circumstances specific to those items. During the 1980s, as overall CPI inflation decelerated, the share of such items increased to a level of about 20-25 percent. Note that the sharp drop in 1989 seen in the figure reflects the effect of the introduction of the consumption tax, as a result of which prices inclusive of tax increased. Similarly, in 1997, the increase in the consumption tax rate resulted in another sharp drop in the share.

However, it is from 1995 that the most conspicuous change in the share of such items can be

observed: with the exception of 1997, the share rose rapidly until 1999. In that year, the share of items whose rate of price change was close to zero reached 55 percent, and it has remained at that high level, with some fluctuations, ever since.

3.2 Item-level price change distributions before and after QQE

Next, let us examine if and how items whose price remained unchanged contributed to inflation under the BOJ's monetary easing since spring 2013. Table 2 shows the transition probabilities from 2012 to 2014. Items are divided into those whose price in March 2012 had risen on a y-o-y basis, those whose price had remained unchanged, and those whose price had fallen, and the table shows for each of these groups how prices subsequently developed. Specifically, taking the topmost row as an example, the table is read as follows: among items whose price had fallen in 2012, 23 percent also registered a price decline in 2014, while for 30 percent of these items, prices in 2014 remained unchanged, and 48 percent of items that had registered a fall in 2012 actually saw a price increase in 2014. Note that for the construction of this transition matrix, price changes between -0.75 percent and +0.75 percent are defined as "zero" price changes, while anything below is defined as a decrease and anything above as an increase.

The results in Table 2 show that an extremely high share - 79 percent - of items that registered no price change in 2012 remained in that category in 2014. On the other hand, the share of items that transitioned from no price change in 2012 to an increase in 2014 is only 16 percent. In other words, items that registered no price change in 2012 made hardly any contribution to the increase in CPI inflation, suggesting that the BOJ's inflation targeting policy and QQE had very little effect on these items. On the other hand, among items that registered a price increase in 2012, 64 percent also registered an increase in 2014 and contributed to raising the inflation rate. Moreover, among items that registered a price decrease in 2012, 48 percent saw a price increase instead, and this also contributed to pushing up inflation. Thus, items with flexible prices contributed to raising the inflation rate, while items with sticky prices tended to register no price change and were the major obstacle to increases in the CPI inflation rate.

4 Why Have Prices Become Stickier?

4.1 Exogenous vs. endogenous changes in price stickiness

There are two possible reasons as to why prices in Japan have become stickier since the mid-1990s. The first is a structural change in the economy, such as a change in the competitive environment that firms operate in, leading firms to change their price setting behavior and resulting in the

increase in price stickiness. Previous studies point to a variety of factors that may have resulted in structural changes in the competitive environment. For instance, the second half of the 1990s is a period when the rise of new firms in emerging economies such as China gathered pace, intensifying global competition. This may have created a situation in which firms were unable to pass on any increases in marginal costs to customers.

An alternative to this is that the increase in price stickiness since the mid-1990s may have been caused endogenously rather than exogenously. Ball and Mankiw (1994) argue based on a menu cost model that price stickiness can change endogenously depending on the level of trend inflation. That is, when trend inflation is high, the profits forgone for a firm by not adjusting prices will be considerable. Such a firm would fall behind if it alone does not raise prices while its rivals do. Because the profits forgone due to not adjusting prices are large, firms will choose to adjust prices despite incurring menu costs. As a result, prices are flexible and the Phillips curve is steeper. In contrast, when trend inflation is close to zero, as has been the case in Japan since the mid-1990s, the profit forgone due to not adjusting prices is smaller than the menu costs, so that firms will put off adjusting prices. Prices are stickier and the slope of the Phillips curve is smaller. In this way, changes in trend inflation lead to endogenous changes in price stickiness. See Levin and Yun (2007) and Bakhshi et al. (2007) for more on this issue.

Note that whether the increase in price stickiness arose exogenously or endogenously has quite different implications. If the increase in price stickiness is due to exogenous structural changes, prices should continue to be stickier and it seems unlikely that the slope of the Phillips curve will return to its previous level. In contrast, if the increase in price stickiness emerges endogenously with a decline in trend inflation, then price stickiness should decline and the slope of the Phillips curve should return to where it was once trend inflation picks up.

4.2 Relationship between the share of items with no price changes and the rate of inflation

To determine whether the increase in price stickiness occurred exogenously or endogenously, this subsection examines the relationship between price stickiness and trend inflation. Specifically, we check whether the share of items whose prices are not adjusted is inversely correlated to the rate of inflation.

To examine whether this is the case, the shares of items whose prices increased, fell, or remained unchanged are calculated for each month from January 1971 to March 2014. In the panels of Figure 4, these shares are depicted on the vertical axis and monthly observations are plotted against the

inflation rate, which is shown on the horizontal axis. It should be noted that “unchanged prices” here are defined as price changes that are strictly zero, while any positive price change is defined as an increase and any negative price change as a decrease. Also, note that observations for April 1989-March 1990 and April 1997-March 1998, when the inflation rate was respectively affected by the introduction of and hike in the consumption tax, are excluded from the sample.

Starting with the panel depicting the relationship between the share of items whose prices rose and the inflation rate, this indicates that in the period 1971-1994, which is represented by the blue circles, a higher inflation rate was associated with a higher share of items whose prices increased. The same relationship can be observed for the period 1995-2014, represented by the red squares, but the blue circles and the red squares display a break somewhere around 3 percent inflation. From 1995 onward, the share of increasing price items declines more quickly with the rate of inflation.

Turning to the share of items whose prices fell, the figure indicates that in the period 1971-1994 this tends to fall when inflation rises. However, similar to the share of increasing price items, the share of decreasing price items displays a break in the neighborhood of 3 percent and rapidly increases when inflation falls to zero.

A simple linear regression shows that in the period 1971-1994, a 1 percentage point fall in the inflation rate led to a 1.3 percentage point decline in the share of items whose price increased and a 0.9 percentage point rise in the share of items whose price fell, indicating that the decline in the former was greater than the increase in the latter. Consequently, a 1 percentage point decline in the inflation rate raised the share of items whose price remained unchanged by 0.4 percentage points. The fact that the blue circles in Figure 4(c) are downward sloping provides a graphic representation of this relationship. In other words, price stickiness as measured by the share of unchanged items increases as the inflation rate declines from a positive value to zero. This finding suggests that the observed increase in price stickiness is, at least partially, due to the endogenous mechanism described by Ball and Mankiw (1994).²

Figure 4(c) also shows that price stickiness decreases as the inflation rate falls from zero into negative territory.³ A simple regression shows that the share of items whose price remained unchanged increases by 0.008 percentage points for a 1 percentage point deviation in the rate of inflation from zero in a negative direction. Interestingly, the corresponding figure is greater when the rate of inflation deviates in a positive direction (0.003), suggesting that the extent to which prices become less sticky is greater when the inflation rate moves from zero into negative territory than when it moves into positive territory. Note that this is the opposite of downward price rigidity.

5 Cross-country Comparison of Price Change Distributions

In this section, we conduct a cross-country comparison to investigate more on causes of the increase in price stickiness in Japan. Specifically, we compare item-level price change distributions for major industrial countries including Japan and the United States, with a particular focus on price stickiness in these countries.

5.1 Item-level price change distributions as of March 2014

Let us begin by comparing the item-level price change distribution for Japan as of March 2014 with those for the United States, Canada, and the United Kingdom (Figure 5). Starting with the United States, the peak of the distribution lies between 2 and 3 percent. Although there is also a small peak around zero, the peak around 2-3 percent is much higher. The shape of the distribution is quite different from that for Japan. It seems that in the United States price increases of around 2-3 percent are the default, meaning that unless there are special circumstances, US firms raise their prices every year in the range of 2-3 percent. On the other hand, in Japan the default is to leave prices unchanged, and this likely gives rise to the difference in the distributions.

Looking at the results for the other two countries, the peaks of the distributions for both Canada and the United Kingdom are in the range of 1-2 percent. In other words, for all three countries – the United States, Canada, and the United Kingdom – the peak of the distribution is in the vicinity of 2 percent. These countries have an inflation target of 2 percent and the peak of the distribution and the level of the inflation target more or less coincide with each other. We will come back to this issue later.

5.2 The fraction of items with a near-zero price change conditional on the mean of the item-level price change distribution

As we saw in the previous section, the shape of the item-level price change distribution differs depending on the level of trend inflation. Therefore, a simple comparison of price change distributions in different countries at a particular point in time may be misleading, since the rate of inflation in different countries differs at any particular point in time. In what follows, we instead focus on the distribution of item-level price changes in period t conditional on the mean of the distribution in the same period (i.e., the rate of inflation in period t). We compare such “conditional” distributions for eight countries, namely, Japan, the United States, Canada, Germany, France, Switzerland, Italy, and the United Kingdom.⁴

The first thing we do is to calculate, for each country, the fraction of items with a near-zero

price change for a given overall inflation rate. A “near-zero” price change here is defined as a price change within ± 0.25 percent. The results are shown in Figure 6, which is read as follows. The value on the vertical axis corresponding to, say, 2 percent on the horizontal axis represents the fraction of items with a near-zero price change in those months in which the mean of the distribution – i.e., the overall inflation rate – was in the vicinity of 2 percent. As can be seen, a common feature across countries is that the fraction of items with near-zero inflation tends to decrease with the rate of inflation. In other words, prices tend to be stickier the closer the rate of inflation is to zero.

However, a number of important differences across countries can be observed. First, the fraction of items with near-zero inflation is significantly higher in Japan than the other countries, irrespective of the rate of inflation shown on the horizontal axis. In other words, the difference in price stickiness between Japan and the other countries cannot be accounted for by differences in trend inflation. Note that this result may be due to differences in terms of the data granularity. The number of items in the Japanese item-level price data is 588, which is comparable to the number of items in the data for Germany and the United Kingdom but much larger than for the other countries (see footnote 4 for the number of items in the other countries). The granularity of the Japanese data is finer in this sense, so that it is possible that there are more zeros in the Japanese data. To check whether this is the reason for the higher fraction of items with near-zero inflation, we construct a new dataset in which we reduce the number of items from 588 to about 300 through aggregation and repeat the same exercise as in Figure 6. However, we find no significant difference in the result.

Second, the fraction of items with near-zero inflation for Japan takes the highest value when the inflation rate is somewhere around -2 percent and tends to become smaller as the rate of inflation goes deeper into negative territory, which is again consistent with the prediction by Ball and Mankiw (1994). The same tendency can be seen for the other countries, but the location of the peak is different from that of Japan. For the United States and the United Kingdom, the peak is negative but not that far from zero. For Canada and Switzerland, the peak is located in positive territory. Meanwhile, for the other three countries (Germany, France, and Italy), no peak is observed, since there are not many observations with a negative inflation rate. However, looking at the curves for these countries, it appears that any peak likely would *not* be located in positive territory - at least in the case of Germany and France.

Note that if the mode of the conditional distribution is located somewhere around zero irrespective of the rate of inflation, we would expect the fraction of items with zero inflation to monotonically increase as the rate of inflation comes down from positive to zero, reaching its maximum when the inflation rate is zero. However, the result that this does not happen in some of the

countries suggests that the mode of the conditional distribution in those countries deviates from zero depending on the inflation rate.

5.3 The mode of the item-level price change distribution conditional on the mean of the distribution

As we saw in Figure 5, the mode of the item-level price change distribution as of March 2014 was in the vicinity of zero for Japan, while it was above zero for the United States, Canada, and the United Kingdom. To investigate this in more details, we now show in Figure 7 the mode of the *conditional* distribution for each of the eight countries.⁵ We see that for Japan the mode is much lower than for the other countries irrespective of the rate of inflation. For example, when the inflation rate, which is shown on the horizontal axis, is 3 percent, the mode is about 1 percent for Japan, while it is around 3 percent for the United States.

A feature that is common to all countries is that the mode and the mean of the item-level price change distribution do not necessarily coincide. Specifically, the mean tends to be greater than the mode when the mean is high, indicating that the distribution is skewed to the right. Such positive skewness arises when prices rise for a particular set of sectors (for example, due to a energy price hike or currency depreciation) but remain unchanged for the other sectors, thereby resulting in relative price changes. In contrast, when the mean is low, the mean tends to be smaller than the mode, indicating that the distribution is skewed to the left. The presence of such a positive correlation between the mean and the skewness of the cross-sectional price change distribution has been discussed by Ball and Mankiw (1995) among others.⁶

It is important to note that, for Japan, the mean and the mode coincide when the inflation rate is in the vicinity of zero, indicating that the threshold associated with symmetry/asymmetry of the distribution is somewhere around zero. However, this is not true for the other countries. For example, the threshold is slightly above 3 percent for the United States, much higher than for Japan. This means that the inflation rate can reach 3 percent in the United States without any changes in relative prices across sectors. However, for Japan, 3 percent inflation can be achieved only with a positive supply shock. To the extent that supply shocks are inherently short-lived, this implies that 3 percent inflation would not be sustainable for Japan. It may be the case that inflation expectations are well anchored somewhere near 2-3 percent for countries like the United States, but they are anchored at zero percent for Japan, so that Japan has to rely on positive supply shocks, such as yen depreciation, to achieve positive inflation.

5.4 The dispersion of the item-level price change distribution conditional on the mean of the distribution

Finally, we examine how the dispersion of the distribution, which is referred to as relative price variability (RPV), is related to the rate of inflation. Previous studies on RPV empirically show that RPV tends to increase with the rate of inflation. Sheshinski and Weiss (1977) and Weiss (1993) provide an explanation of these empirical results based on menu cost models. Specifically, Weiss (1993) argues that, at moderate inflation rates, an increase in inflation raises the relative size of each price increase; in addition, the degree of synchronization in firms' price adjustments is rather low, so that RPV rises with inflation. However, Weiss (1993) also argues that, at very high rates of inflation, most firms raise prices within the same period (so that price adjustments are synchronized) and at more similar rates, so that RPV *declines* as the rate of inflation increases.

Figure 8 shows the relationship between RPV, which is measured as the difference between the 10th and 90th percentiles of a distribution, and the inflation rate for the eight countries. We find the presence of a U-shaped relationship in all countries except Italy and Germany, for which the number of observations with negative inflation rates is quite limited.⁷ More interestingly, the bottom of the U-shape differs across countries. It is slightly below zero for Japan, but positive for the other countries. For example, the bottom is located at 1 percent for the United States, 1.5 percent for Canada, 0.5 percent for Switzerland, 1.5 percent for France, and 1.5 percent for the United Kingdom.

The theoretical argument by Weiss (1993) implies that RPV takes a minimum value when price changes are synchronized to the utmost extent. Suppose, for example, that price indexation is adopted in an economy and the rate of indexation is exogenously fixed at $\bar{\pi}$. In this case, price adjustments are synchronized across almost all firms, and RPV is minimized when the inflation rate is equal to $\bar{\pi}$.⁸ This suggests that the observed difference in the bottom of the U-shape between Japan and the other countries stems from the difference in $\bar{\pi}$ across countries. Note that the literal definition of $\bar{\pi}$ is the rate of price change associated with price indexation. However, a more realistic interpretation is that there is an implicit understanding among sellers and buyers that prices will be revised each year at a fixed rate, which is given by $\bar{\pi}$. In this sense, $\bar{\pi}$ may be interpreted as the reference level of inflation shared by firms.⁹ Figure 8 indicates that this reference inflation level is somewhere around 1 to 2 percent in all countries other than Japan, where it is in the vicinity of zero.

6 Summary and Policy Implications

The main findings of this paper can be summarized as follows. First, for the majority of the 588 items constituting the CPI, making up about 50 percent of the CPI in terms of weight, the y-o-y rate of price change was close to zero. In this sense, price stickiness was high. This situation started during the onset of deflation in the second half of the 1990s and continued even after the rate of inflation turned positive in 2013 and 2014 due to monetary easing by the BOJ.

Second, we show that the fraction of items with a near-zero price change increased as the overall inflation rate fell and approached zero, which is consistent with the argument that as inflation approaches zero, prices become stickier due to menu costs.

Third, the cross-country comparison of item-level price change distributions showed that Japan differs significantly from other countries in that the mode of the distribution is very close to zero for Japan, while it is near 2 percent for the other countries examined, including the United States. Importantly, this result remains unchanged even if we control for differences in trend inflation across countries, implying that the observed increase in price stickiness in Japan since the second half of the 1990s cannot be fully accounted for by the menu cost-type argument.

Our empirical results indicate that inflation expectations are anchored somewhere near 2 percent for the United States and the other countries, but they are anchored at zero percent for Japan. Put differently, whereas in the United States and the other countries the “default” is for firms to raise prices by about 2 percent each year, in Japan the default is that, as a result of prolonged deflation, firms keep prices unchanged. Such pricing behavior in Japan has contributed significantly to the flattening of the Phillips curve, thereby making it difficult to escape from deflation. One of the most important lessons to be learned from the Japanese experience during the Abenomics period is that it is extremely difficult to eliminate the deflationary mindset and the pricing behavior based on this mindset once it has become deeply ingrained in society.

Notes

¹According to the survey in the Annual Report on the Japanese Economy and Public Finance 2013, 23 percent of firms answered they passed on about half, 12 percent answered that they passed on less than half, and 26 percent answered that they could not pass on any cost increases at all.

²Similar results are reported for other countries. See, for example, Gagnon (2009).

³Note that all of the studies to date using micro price data deal with periods of positive inflation and there are no examples of studies on the relationship between the inflation rate when this is zero and price rigidity. Gagnon (2009) also mainly focuses on a period of positive inflation, although in 2001 and 2002 the prices of fresh vegetables and fruits in Mexico fell due to weather factors and the y-o-y rate of change in the CPI for goods turned negative. It appears that during this period, symmetry held in that price stickiness tended to decrease the more the inflation rate in terms of the CPI for goods fell below zero (see Figure 4 in Gagnon (2009)). However, this bout of deflation lasted only for a few months, so that it is not possible to tell whether this is a robust tendency.

⁴The number of items and the observation period for each country are as follows: for Japan, the number of items is 588 and the observation period is 1970 to 2016; for the United States, the number of items is 182 and the observation period is 1970 to 2016; for Canada, the number of items is 170 and the observation period is 1985 to 2016; for Germany, the number of items is 577 and the observation period is 1991 to 2016; for France, the number of items is 262 and the observation period is 1990 to 2016; for Switzerland, the number of items is 268 and the observation period is 1982 to 2016; for Italy, the number of items is 215 and the observation period is 1996 to 2010, and finally, for the United Kingdom, the number of items is 687 and the observation period is 1997-2016. The number of items for each country may not be the same throughout the observation period. Note that the analysis in this and later subsections is based on unweighted data (e.g., unweighted rather than weighted mean of the item-level price change distribution).

⁵See Watanabe and Watanabe (2017) for details on the estimation method of the mode of the item-level price change distribution. See Cavallo and Rigobon (2011) for more on the estimation of the mode of a price change distribution.

⁶Ball and Mankiw (1994) argue that if the distribution of sector-specific shocks, such as oil price hikes, is skewed to the right, firms in a particular sector that experiences a large shock have an incentive to raise their prices, while firms in a sector that experiences only a very small shock do not want to change their prices, since this would incur menu costs. Therefore, even if the mean of the sector shock distribution is zero, the mean of the inflation rates in individual sectors could be positive, creating a positive correlation between the mean and skewness. More recently, Choi (2010) has argued that the degree of price stickiness differs across sectors, so that even common shocks (i.e., shocks common to all sectors) may have different impacts on prices across sectors, thereby creating skewness in sectoral price change distributions.

⁷Choi (2010) finds similar a U-shaped relationship between inflation and RPV for the United States and Japan.

⁸See Watanabe and Watanabe (2017) for a more detailed discussion of this issue .

⁹Okun (1981) and Schultze (1981), among others, argue that firms have implicit contracts with their customers, which constitute social norms. An interpretation of our result is that inflation norms differ between Japan and the other countries, possibly reflecting the experience of prolonged deflation in Japan.

References

- [1] Abe, N., Tonogi, A., and Watanabe, T. (2008). “Who Adjusts Prices: Manufacturers or Retailers?” *Economic Review*, 59(4), 305-316 (in Japanese).
- [2] Bakhshi, H., Khan, H., and Rudolf, B. (2007). “The Phillips Curve under State-dependent Pricing.” *Journal of Monetary Economics*, 54(8), 2321-2345.
- [3] Ball, L., and Mankiw, N. G. (1994). “Asymmetric Price Adjustment and Economic Fluctuations.” *Economic Journal*, 104(423), 247-261.
- [4] Ball, L., and Mankiw, N. G. (1995). “Relative-price Changes as Aggregate Supply Shocks.” *The Quarterly Journal of Economics*, 110(1), 161-193.
- [5] Calvo, Guillermo A. (1983). “Staggered Prices in a Utility-maximizing Framework.” *Journal of Monetary Economics* 12(3), 383-398.
- [6] Cavallo, A., and Rigobon, R. (2011). “The Distribution of the Size of Price Changes.” NBER Working Paper No. 16760.

-
- [7] Choi, C. Y. (2010). "Reconsidering the Relationship between Inflation and Relative Price Variability." *Journal of Money, Credit and Banking* 42(5), 769-798.
- [8] Gagnon, E. (2009). "Price Setting during Low and High Inflation: Evidence from Mexico." *The Quarterly Journal of Economics*, 124(3), 1221-1263.
- [9] Gordon, Robert J. (1981). "Output Fluctuations and Gradual Price Adjustment." *Journal of Economic Literature* 19(2), 493-530.
- [10] Levin, A., and Yun, T. (2007). "Reconsidering the Natural Rate Hypothesis in a New Keynesian Framework." *Journal of Monetary Economics*, 54(5), 1344-1365.
- [11] Okun, A. M. (1981). *Prices and Quantities: A Macroeconomic Analysis*. Washington, D.C.: Brookings Institution.
- [12] Schultze, C. L. (1981). "Some Macro Foundations for Micro Theory." *Brookings Papers on Economic Activity*, 1981(2), 521-592.
- [13] Sheshinski, E., and Weiss, Y. (1977). "Inflation and Costs of Price Adjustment." *The Review of Economic Studies*, 44(2), 287-303.
- [14] Watanabe, K., and Watanabe, T. (2017). "Price Rigidity at Near-Zero Inflation Rates: Evidence from Japan." *CARF Working Paper Series*, CARF-F-408, March 2017.
- [15] Watanabe, T. (2013). "What's Required to Stop Deflation?" *Japan Spotlight*, May/June 2013, Japan Economic Foundation, 30-35.
- [16] Weiss, Y. (1993). "Inflation and Price Adjustment: A Survey of Findings from Micro-Data." In E. Sheshinski and Y. Weiss, eds., *Optimal Pricing, Inflation, and the Cost of Price Adjustment*. Cambridge, Mass.: MIT Press.

Table 1: Monetary Policy Decisions Made by the Bank of Japan in 2013-2016

Jan 22, 2013	Joint statement of the Government and BOJ on overcoming deflation and achieving sustainable economic growth was released. BOJ set an inflation target of 2 percent.
Apr 4, 2013	BOJ introduced Quantitative and Qualitative Monetary Easing (QQE). BOJ decided to double the monetary base to achieve the 2 percent target with a time horizon of about two years.
Oct 31, 2014	BOJ expanded QQE by increasing its annual net purchases of JGBs from 50 trillion yen to 80 trillion yen.
Jan 29, 2016	BOJ decided to apply a negative interest rate of minus 0.1 percent to part of BOJ current account balances.
Sep 21, 2016	BOJ decided to employ yield curve control with the target level of 10-year JGB yields set to zero.

Table 2: Transition Probability Matrix from December 2012 to March 2014

		Item Level Inflation in Mar. 2014		
		Negative	Zero	Positive
Item Level Inflation in Dec. 2012	Negative	23%	30%	48%
	Zero	4%	79%	16%
	Positive	16%	20%	64%
Unconditional Probability		11%	57%	32%

Note: “Zero” here is defined as a price change between ± 0.75 percent.

Figure 1: Phillips curve, 1971-2016




Figure 2: Distribution of price changes across items

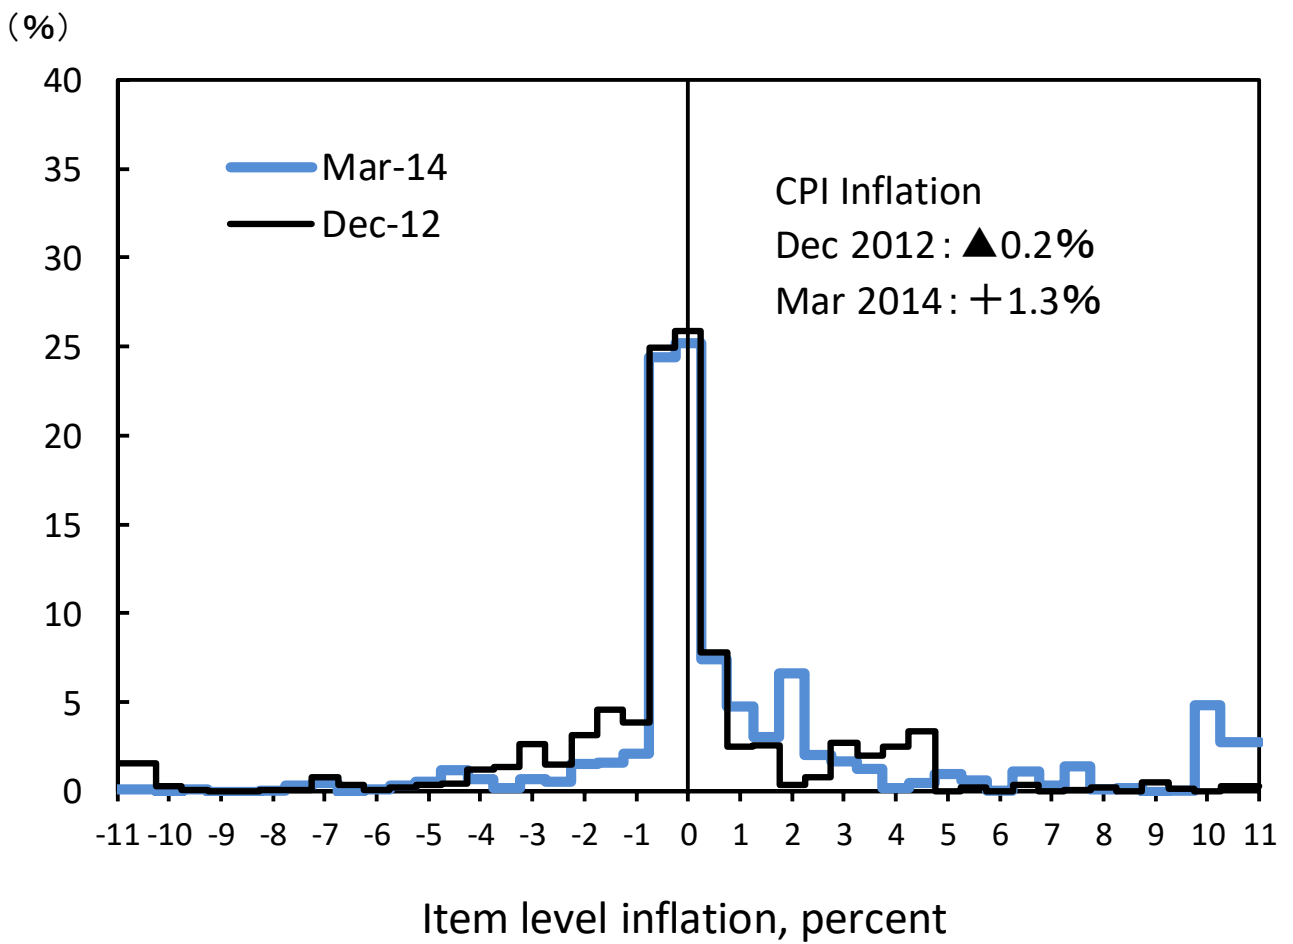


Figure 3: Fraction of items with a near-zero inflation rate




Figure 4: Fraction of items with price increase, price decrease, and no price change

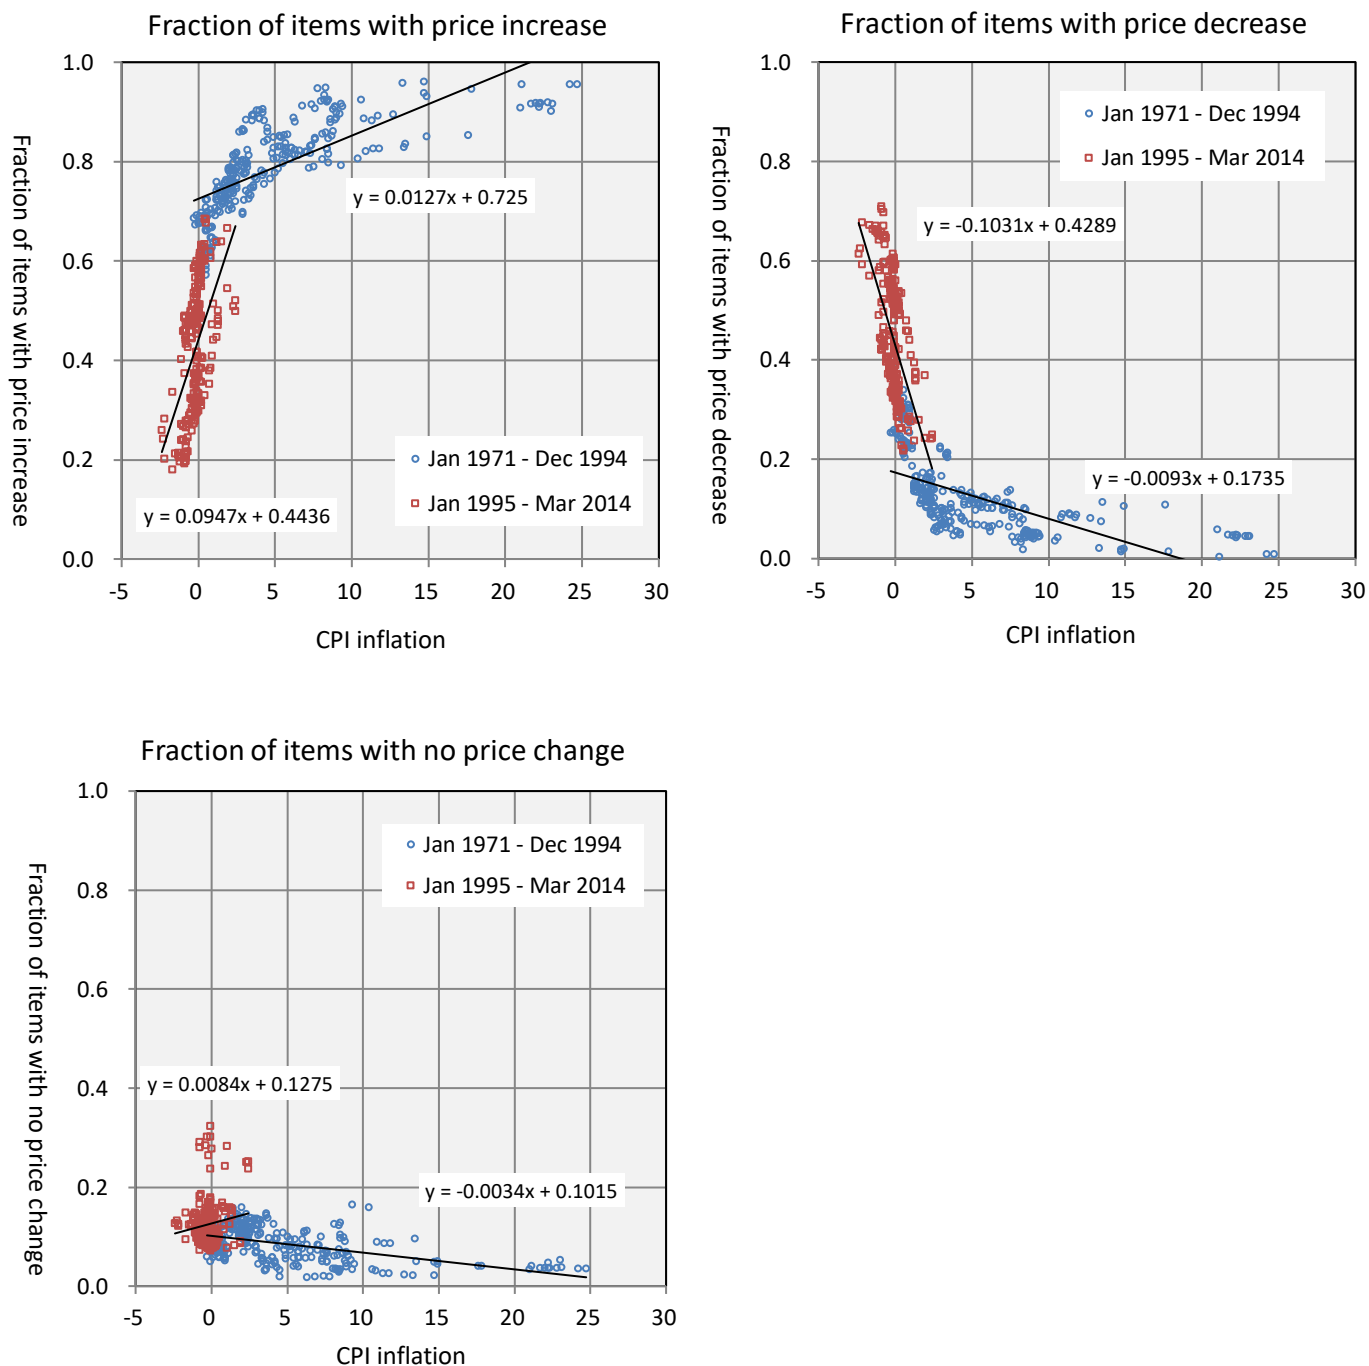


Figure 5: Price change distributions for Japan, US, Canada, and UK, March 2014




Figure 6: Fraction of items with no price change conditional on the inflation rate




Figure 7: Mode of the price change distribution conditional on the inflation rate




Figure 8: Relative price variability conditional on the inflation rate

