Menu Costs and Price Change Distributions: Evidence from Japanese Scanner Data

Yukiko Saito
Fujitsu Research Institute

Tsutomu Watanabe*
Hitotsubashi University

June 14, 2007

Abstract

This paper empirically examines the menu cost hypothesis by analyzing the distribution of price changes using daily scanner data covering almost all items sold at about 200 Japanese supermarkets. We arrive at the following findings. First, as implied by the menu cost hypothesis, small price changes were indeed rare. The price change distribution for goods with sticky prices has a dent at the vicinity of zero inflation. In contrast, no such dent can be observed in the price change distribution for goods with flexible prices. Second, we find that the longer the time that has passed since the last price change, the higher is the probability that a large price change occurs. Combined with the fact that the price adjustment probability is a decreasing function of the price duration, this means that although the price adjustment probability decreases as the price duration increases, once a price adjustment occurs, the magnitude of such an adjustment is large. Third, while the price change distribution is symmetric on a short time scale, it is asymmetric on a long time scale, with the probability of a price decrease being significantly larger than the probability of a price increase. The asymmetry on a long time scale seems to be related to the deflation that the Japanese economy has experienced.

JEL Classification Numbers: E30
Keywords: Menu costs; Price rigidity; Price change distributions; Price duration

*Correspondence: Tsutomu Watanabe, Institute of Economic Research and Research Center for Price Dynamics, Hitotsubashi University, Kunitachi, Tokyo 186-8603, Japan. Phone: 81-42-580-8358, fax: 81-42-580-8333, e-mail: tsutomu.w@srv.cc.hit-u.ac.jp.

1 Introduction

The menu cost hypothesis has various important implications, which have been the subject of numerous empirical investigations. The implications of the menu cost hypothesis can be divided into two major issues: those relating to the probability of the occurrence of a price change and those relating to the distribution of price changes conditional on the occurrence of a change.

An example of the former is that in a high-inflation economy the probability of price adjustments is higher than in a low-inflation economy. In a high-inflation economy, the divergence of actual prices from target prices (i.e., the prices firms would choose if there were no transaction costs involved in price adjustments) rapidly increases over time if prices are not adjusted. We would therefore expect firms to adjust prices frequently in order to avoid large deviations. This implication has been empirically tested, for example, by Lach and Tsiddon (1992) that uses data for Israel.

Another implication is that the probability of a price adjustment increases the longer the period that prices are not adjusted. Under the usual circumstances, where the variance of target prices monotonously increases with time, the probability that firms adjust prices increases if prices have not been adjusted for a long time because the probability that the target price goes out of the inactive range also increases. A number of studies have empirically examined this issue using hazard functions (e.g., Alvarez et al. 2005, Campbell and Eden 2006, Gagnon 2005, Nakamura and Steinsson 2007). However, many of these studies found that the hazard function is downward sloping; in other words, the longer the period in which prices are not adjusted, the lower is the probability that prices are adjusted - a result is in conflict with menu cost theory.

An example of the latter issue - namely, the distribution of price changes conditional on the occurrence of a change - is that the menu cost hypothesis suggests that small prices changes are unlikely to occur. When deviations of actual prices from target prices are small, rather than incurring the transaction costs involved (i.e., menu costs) and implementing small changes in prices, it may be less costly to refrain from any price adjustments in the first place. Thus, based on the menu cost hypothesis, we would expect that only relatively large price changes that are worth the transaction costs incurred are observed in practice. However, research by Kashyap (1995), Carlton (1986), Lach and Tsiddon (2005), Midrigan (2006) and others, examining the distribution of price changes conditional on the occurrence of a change, have found that price changes that are small are in fact by no means rare. On the other hand, it is interesting to note in this context that Kackmeister (2005) reports that while in the nineteenth century small price changes were indeed rare, this is no longer the case when focusing on the more recent period.

Dividing empirical studies on the menu cost hypothesis in this way, i.e., into those related to the probability of price adjustments and those related to the distribution of price changes conditional on the occurrence of a change, we come across an important implication; namely, that the distribution of price changes depends on the past. The aforementioned analyses using hazard functions examine whether price adjustment probability in the current period depend on whether there has been a price adjustment in $t-1$, and $t-2$, and so on; namely, such studies focus on the history dependence of price adjustment probability. Similarly, in studying price change distributions, it should be worthwhile to pay attention to history dependence, i.e., it is important to examine how the occurrence of price adjustments in previous periods affect the price change distribution in the current period. In other words, it is important to examine the relationship between price duration (i.e., the duration between the last price adjustment and any price adjustment in the current period) and the price change distribution in the current period.

On the relationship between price duration and the price change distribution, the menu cost hypothesis is very clear: there is no relationship. Put differently, the menu cost hypothesis posits that the price change distribution is the same, regardless of the length of time since the last price change. In the menu cost hypothesis, the deviation of the actual price from the target price is a state variable, and when the value of this state variable meets a specific condition, price adjustment occurs. Put differently, if we were to gather instances of price adjustments, in each instance, the size of the deviation of the actual price from the target price should be identical. This means that as long as firms base their decision on the magnitude of price adjustments in correspondence to the magnitude of the deviation, identical price change distributions should be observed in each instance. On the other hand, if price adjustments are not state dependent but time dependent, price duration and the price change distribution are not necessarily unrelated because deviations from the target price are not always the same. This suggests that looking at the correlation between price duration and the price change distributions would be useful to distinguish between state- and time-dependent models.

Focusing on the above issue, the purpose of this paper is to empirically examine the implications of the menu cost hypothesis regarding the price change distribution. To this end, we use daily scanner data covering almost all the goods sold in about 200 supermarkets in Japan. The observation period is from 1998 to 2005. The total number of individual price observations in the dataset (consisting of the no. of items $\times$ no. of outlets $\times$ no. of days) is roughly 3 billion. The results of the analysis using this unique dataset can be summarized as follows.

First, we find that, small price changes were indeed rare. We arrive at this conclusion by classifying goods by their average price adjustment probability during the sample period and plotting the price change distribution for each subgroup. The price change distribution of goods with sticky prices has a dent in the center of the distribution (i.e., the part where price changes are small). In contrast, in the price change distribution of goods with flexible prices, such a dent could not be seen. Moreover, we observe that while outside the central part of the price change distribution (i.e., the part where price changes are large) follows a power law distribution, the central part diverges from it.

The second result of our investigation is that there is a stable relation between price duration and the distribution of price changes. Specifically, it was found that the longer the price duration, the deeper becomes the dent in the center of the distribution of price changes. In other words, in the case that a long time has passed since the last price adjustment, this will result in a large price change. On the other hand, when we examine the relationship between price duration and the probability of price adjustment using a hazard function, we arrive at the same result as preceding studies, namely, that a decreasing hazard was observed. Putting these two findings together indicates that although the price adjustment probability declines the longer the price duration, once a price adjustment occurs, the price change will be large as a result. The findings thus suggest the possibility that the longer the price duration, the higher become the menu costs.

The third result of our analysis is that on a short time scale the distribution of price changes is symmetric, but this is not the case on a long time scale. Specifically, defining price changes using the price today and five days earlier and plotting the distribution of price changes, the distribution is almost perfectly symmetric. However, defining price changes using the price today and the price 80 days earlier or more and plotting that distribution, the distribution becomes asymmetric, with the probability of a price decrease being significantly greater than the probability of a price increase. The asymmetry on a long time scale is especially striking since 2000, suggesting that the asymmetry is related to the deflation the Japanese economy has experienced.[1]

[1] Looking at the symmetry and asymmetry on a long time scale for each year, we find that at the beginning of the 1990s, the probability of a price rise was significantly larger than the probability of a price fall. This period represents

The rest of this paper is organized as follows. The next section provides a description of the data used in this paper. Sections 3, 4, and 5 then respectively present our results with regard to the price adjustment probability, the characteristics of the price change distribution and particularly the frequency of price changes of small magnitude, and the analysis on the relationship between price duration and the distribution of price changes. Section 6 concludes the paper.

2 Data Description

The data used in this paper have been collected by Nikkei Digital Media Inc. The frequency of the data is daily and the sample period is from 1988 to 2005. The number of outlets covered in 2005 was 181. Individual items are identified by their JAN (Japanese Article Number) Code and the total number of different items sold in 2005 was approximately @@. The total number of observations for 2005 was about 280 million (defined as the no. of articles $\times$ no. of outlets $\times$ no. of days), while the total for the entire sample period is approximately 2.7 billion observations.

[Insert Tables 1 and 2]

Tables 1 and 2 show the number of outlets and items for each year as well as the turnover in outlets and items in the dataset during the observation period. The number of outlets that are included in the dataset throughout the entire period is 17. The number of items sold by those 17 outlets in 1989 was approximately 230,000 and has subsequently risen steadily, reaching roughly 460,000 in 2004. During this period, tens of thousands of items were newly launched each year, but about the same number of items were also withdrawn. The ratio of the number of newly launched items relative to existing items was about 35 percent, while the withdrawal rate was about 30 percent, indicating that the turnover in items was quite rapid.

3 The Probability of Price Changes

In addition to changes in regular prices, changes in selling prices at outlets also include the effect of temporary sales. In order to remove the effect of such sales, we use the following filter. Let the sales price of a particular product at a particular outlet on day t be represented by $\hat{P}_t$; then we define P_t

the last phase of the asset price bubble and at this time, the CPI were also on an upward trend. The results here thus show that the exact opposite asymmetries occur in period of inflation and periods of deflation.

as :

$$P_t \equiv \max\{\hat{P}_t, \hat{P}_{t-1}, \dots, \hat{P}_{t-k+1}\} \tag{1}$$

That is, P is the largest value of $\hat{P}$ in the last k days. P coincides with the regular price under the assumption that (1) the selling price returns to the regular price on days when there is no sales and (2) there are no sales of a consecutive k days. We set $k = 5$ throughout the paper.

Next, we define the index showing the occurrence of price adjustments as

$$I_t^d \equiv \begin{cases} 1 & \text{if} \quad P_t \neq P_{t-d} \\ 0 & \text{if} \quad P_t = P_{t-d} \end{cases} \tag{2}$$

If a price adjustment occurs between day $t-d$ and day t, then I_t^d becomes 1. On the other hand, if no price adjustment occurs during this period, I_t^d is 0.

[Insert Figure 1]

The upper panel of Figure 1 shows the distribution of the price adjustment probability $\Pr(I_t^d = 1)$ calculated by item and by outlets when $d = 5$. The horizontal axis depict the price adjustment probability divided into 20 bins. The interval farthest to the right shows the $(2^{-1/2}, 1]$ bin, the next one to the left the $(2^{-2/2}, 2^{-1/2}]$ bin, followed by the $(2^{-3/2}, 2^{-2/2}]$ bin, etc. The farther to the left, the stickier the prices. For example, a price adjustment probability of 1/8 means that the probability of its occurrence within a period of d days is 1/8, and with $d = 5$, prices will be adjusted at a frequency of once in 40 days. The vertical axis shows the frequency. It should be noted that the analysis here only consider items with a lifespan of more than 100 days.

To begin with, looking at the distribution by item, we find that although its peak at a price adjustment probability of 1/4, the tails of the distribution are extremely long and the highest frequency can in fact be found in the bin farthest to the left, where the price change probability is less than 1/724. This shows that there are large differences between items in their price adjustment probability. Looking at the distribution by outlet, we also find a wide dispersion of price adjustment probabilities, although this is not as great as in the distribution by item. The distribution when $d = 10$ is shown in the middle panel of Figure 1, which shows that in the distribution by item, the dispersion continues to be large.

Based on the finding of large differences between items in the probability of price adjustments, we classify items into ten subgroups. Concretely, we label items with a price adjustment probability belonging to $(1/2, 1]$ as $G0$. Similarly, those with a probability belonging to $(1/4, 1/2]$ as $G1$, those with a probability belonging to $(1/8, 1/4]$ as $G2$, etc. The tenth subgroup $G9$ comprises items with a probability of less than $1/512$. $G0$ is the subgroup in which menu costs are the smallest and prices are the most flexible. Progressing to $G1$, $G2$, etc., menu costs become increasingly larger and prices increasingly stickier.

Finally, the bottom panel of Figure 1 depicts the price adjustment probability for each year. It shows that although the price adjustment probability was stable from 1988 onward throughout the 1990s, a large increase can be observed from 2000. However, it is possible that this trend is affected by the fact that this was a period in which there was a substantial turnover in outlets (i.e., a large number of entries and exits). Therefore, in order to remove this effect, we calculate the price adjustment probability based only on the 17 outlets that existed throughout the entire observation period. The result is the same as in the case when all outlets are included, i.e., a large increase in the probability can be observed.

The fact that the price adjustment probability changes in each year shows that the stochastic process for prices changes over time, implying that it would be dangerous to examine the price adjustment probability and the associated price change distributions for the entire sample period, without paying a particular attention to such a non-stationarity. In the next two sections, therefore, we concentrate in our analysis on the period from 1988 to 2002 unless otherwise mentioned, when the price adjustment probability was relatively stable.

4 Unconditional Price Change Distributions

In this and the next sections, we will look at price change distributions. This section examines *unconditional* price change distributions, while the next section investigates *conditional* price change distributions in the sense that they are distributions conditional on the history of price change events.

Specifically, let us define gross inflation rates from $t - d$ to t as

$$\Pi_t^d \equiv P_t / P_{t-d}. \tag{3}$$

The unconditional distributions we investigate in this section is $f(\Pi_t^d \mid I_t^d = 1)$, while the conditional distributions we will look at in the next section is $f(\Pi_t^d \mid I_t^d = 1, I_{t-d}^d = I_{t-2d}^d = \cdots = I_{t-nd}^d = 0,\ I_{t-(n+1)d}^d = 1)$.

4.1 Small price changes

The main interest of this section is whether and to what extent there exist small price changes in price change distributions. Simple versions of menu cost models imply that there should not exist small price changes. Based on this understanding, several researchers, including Kashyap (1995), Carlton (1986), Midrigan (2006), and Lach and Tsiddon (2006), look for small price changes in price change distributions for the United States and Israel. They all find that the density associated with small price changes is not zero, regarding this as evidence against the menu cost hypothesis.

However, the mere existence of non-zero density for small price changes is not inconsistent with the menu cost hypothesis, simply because any small price changes could occur if menu costs are sufficiently small. For example, let us think about the model proposed by Caballero and Engel (2006), where the menu cost for a product is assumed to change stochastically over time. In this model, an event of price change could occur even if the price is not far from the target level, as long as the realized value for the menu cost is sufficiently small. Specifically, the adjustment hazard assumed in their model implies that the probability of a price change increases with the deviation of the price from the target level, but only gradually (i.e., not discontinuously). Given this model, it would not be surprising even if we observe non-zero density for small price changes.

As another possibility for non-zero density for small price changes, suppose that most firms adopt state dependent pricing, while a limited number of firms adopt time dependent pricing in the sense of Calvo (1983). If this is the case, it is possible that we observe non-zero density for small price changes, but it would be too early to reject the menu cost hypothesis based on that observation.

The above considerations suggest that an appropriate criterion is not whether the density for small price changes is exactly equal to zero or not, but whether it is sufficiently small or not. In what follows, we first specify and estimate a probability density function, PDF, that governs large price changes, and then extrapolate it to small price changes so as to compare between the observed and predicted density for small price changes. The menu cost hypothesis implies an overprediction in the sense that the predicted density is greater than the observed one.

[Insert Figure 2]

Figure 2 shows $f(\Pi_t^5 \mid I_t^5 = 1)$ for the 10 product subgroups. The sample period is 1988 to 2002, and the observations from all outlets are included. The horizontal axis represents 20 bins for the gross inflation rate, consisting of $(0, 2^{-9/20})$, $[2^{-9/20}, 2^{-8/20})$, $[2^{-8/20}, 2^{-7/20})$, $[2^{-7/20}, 2^{-6/20})$, $\ldots$, $[2^{-1/20}, 2^{0/20})$, $(2^{0/20}, 2^{1/20}]$, $\ldots$, $(2^{8/20}, 2^{9/20}]$, $(2^{9/20}, \infty)$.

The distribution for the subgroup $G0$, i.e., the subgroup with the most flexible prices, shows that the densities for two bins at the center of the distribution, $[0.96, 1)$ and $(1, 1.03]$, are higher than those for the other bins.[2] That is, the densities for small price changes are smaller than those for large price changes. We see a similar regularity for the subgroup $G1$. However, we cannot see such a single modal distribution for the other subgroups with prices being stickier than $G1$. For example, the distribution for the subgroup $G2$ has a dent at its center in that the densities associated with $[0.96, 1)$ and $(1, 1.03]$ are slightly smaller than the densities for the bins surrounding these two. We can see such a dent at the center of the distributions for $G3$, $G4$, and $G5$ as well. The fact that small price changes are likely for goods with flexible prices, but less likely for goods with sticky prices is consistent with the menu cost hypothesis.

4.2 Robustness

To check the robustness of this finding, we conduct the same exercise as in Figure 2, but using different samples. First, we extract items whose prices are above 200 yen in order to see whether a dent at the center of a distribution is created not by menu costs, but by monetary indivisibility. Figure 3 clearly shows that this is not the case.

[Insert Figure 3]

Next, we change the sample period from 1988-2002 to 1988-2005. As we saw in the previous section, the probability of price adjustments substantially increased in and after 2003. We are thus curious about how such heterogeneity in terms of the price change probability across years would affect price change distributions. Figure 4 shows a regularity that is quite different from what we see in Figure 2; namely, densities associated with the two bins, $[0.96, 1)$ and $(1, 1.03]$, are now greater than others for the subgroups ranging from $G0$ to $G7$. In other words, small price changes are not rare any more for these subgroups. By scrutinizing the data, we can confirm that this is a direct consequence of heterogeneity of the price change probability across years. For example, those items which are classified into $G3$ have a similar distribution as in Figure 2 for the period of 1988-2002, but they exhibit a single modal distribution for the period of 2003-2005, during which not only the probability of price changes

[2]The bins $[0.96, 1)$ and $(1, 1.03]$ correspond to $[2^{-1/20}, 2^{0/20})$ and $(2^{0/20}, 2^{1/20}]$, respectively.

increased, but also the likelihood of small price changes increased substantially for those items. An interpretation is that menu costs for these $G3$ items were very low in 2003-2005, leading to an increase in the price change probability as well as an increase in the likelihood of small price changes. The result of this exercise shows that heterogeneity across time, not to mention heterogeneity across products, could create a serious bias about estimates of price change distributions.

[Insert Figure 4]

4.3 Distortions of price change distributions at the vicinity of zero inflation

Figure 5 shows cumulative density functions, CDFs, of price changes for the 10 product subgroups. Figure 5.1 shows CDFs for the case of price decreases, i.e., $\Pi < 1$, where the vertical axis depicts $\log \Pr(\Pi^5 < c \mid I^5 = 1)$ with $c \leq 1$, and the horizontal axis represents the value of c. Similarly, Figure 5.2 shows CDFs for the case of price increases, $\Pi > 1$, and the vertical axis represents $\log \Pr(\Pi^5 > c \mid I^5 = 1)$, where $c \geq 1$.

[Insert Figure 5]

We can see in Figure 5.1 that every point of a CDF is on a straight line except two or three points from the right, which are close to $\Pi = 1$. Given that the horizontal axis is expressed in logarithm, and the vertical axis represents the log of CDF, this fact implies that the PDF and the CDF governing large price changes are given by the form of

$$f(\Pi \mid I = 1) \propto \Pi^{-\alpha}; \tag{4}$$

$$F(\Pi \mid I = 1) \propto \Pi^{-(\alpha-1)}, \tag{5}$$

where α is a positive parameter. A distribution with these form of PDF and CDF is referred to as power law distribution or Pareto distribution, where α is the power law exponent.[3] One of the most important characteristics of this form of distribution is that it has a very long tail. Using the U.S. scanner data, Midrigan (2006) find that a price change distribution has tails fatter than those of a normal distribution, and that density at the vicinity of zero inflation is greater than those of a normal distribution. Our finding is perfectly consistent with the first one, although it is in sharp contrast with the second one.

[3] See, for example, @@@ (@@@@) for more on power law distributions.

A careful examination of Figure 5.1 reveals that seven points from the left, namely points corresponding to $c = 0.73, 0.75, 0.78,\ 0.81, 0.84, 0.87, 0.90$, are on a straight line, but the remaining three points, namely those which correspond to $c = 0.93, 0.96, 1.00$, are below the straight line. This suggests a methodology to quantitatively evaluate distortions at the central part of a price change distribution: we first fit a straight line for the seven points by an ordinary least squares regression, then extrapolate the line to the remaining three points, and finally obtain prediction errors, which is defined as predicted minus actual values, as a measure for distortions of a distribution at the vicinity of zero inflation. The red lines shown in Figure 5.1, which are obtained in this way, indicate that actual values for the three points from the right tend to be smaller than the predicted ones, suggesting that small price changes are less likely to occur as compared with large price changes. The same method is applied to the case of price increases, and the result is presented in Figure 5.2.

[Insert Table 3]

Table 3 presents estimates for distortions at the vicinity of zero inflation. For example, the number at the upper left corner, -0.003, represents how much the predicted value deviates from the actual one in terms of the third point from the right, namely $\log \Pr(\Pi^5 < 0.93 \mid I^5 = 1)$, for the product subgroup $G0$. We can read from this table that predictions errors tend to be larger for $G1$ than for $G0$, and those for $G2$ are larger than those for $G3$, and so on, implying that distortions at the central part of a distribution become greater for those products with stickier prices.

However, this relationship between the degree of distortions and price stickiness is not a monotonic one. That is, prediction errors tend to decrease with price stickiness for the product subgroups $G4$, $G5$, and so on, and finally those errors become very close to zero or even below zero for the products with very sticky prices, $G8$ and $G9$, in which a price change occurs every 1280 days or more. To know the reason behind this non-monotonic relationship, we go back to Figure 5.1, to find that the slope of an estimated line, namely the power law exponent, tends to become smaller with price stickiness. This tendency is particularly clear for $G6$, $G7$, and $G8$. This implies that the variance of a distribution, which governs large price changes, becomes larger with price stickiness. On the other hand, the three points corresponding to $c = 0.93, 0.96, 1.00$ are located almost at the same place for these subgroups, and consequently, predictions errors for them approach to zero. This result about products with very sticky prices is obviously not consistent with the menu cost hypothesis, suggesting that reasons behind price stickiness for those products might be different from reasons for the products whose prices are not very sticky.

[Insert Figure 6]

Figure 6 illustrates the difference between an estimated and observed price change distributions. The upper panel of Figure 6 shows an estimated distribution for the subgroup $G3$, which has the largest prediction errors as shown in Table 3. The upper panel is obtained based on the assumption that the PDF governing small price changes is identical to the one that governs large price changes. Specifically, the estimated straight lines in Figure 5 indicate that the PDF governing large price changes for the product subgroup $G3$ is given by

$$f(\Pi \mid I = 1) \propto \begin{cases} \Pi^{-0.149} & \text{for} \quad \Pi < 1 \\ \Pi^{-0.151} & \text{for} \quad \Pi > 1 \end{cases} \tag{6}$$

A simple comparison between the upper and lower panels of Figure 6 clearly indicates that the actual PDF has a dent at the vicinity of zero inflation (namely, $\Pi = 1$).

4.4 Symmetry or asymmetry of price change distributions

So far we have set d at $d = 5$. But one may wonder how price change distributions would change when we choose different values for d. Note that changing the value of d is equivalent to changing time scale when we look at the time-series data. In particular, we are interested in how symmetry or asymmetry would be changed when we look at the data on a longer time scale.

[Insert Figure 7]

Figure 7 presents the result of this exercise: we plot $\Pr(\Pi^d > 1 \mid I^d = 1)$ and $\Pr(\Pi^d < 1 \mid I^d = 1)$ for various values for d ($d = 5, 10, 20, \ldots, 1280$). First, we can see that the two probabilities are both equal to 0.5 for $d = 5$ and $d = 10$, indicating that price change distributions are symmetric on a short time scale. This is consistent with the fact that most of the PDFs in Figure 2 are symmetric. Progressing to longer time scales, however, the probability of a price decrease, $\Pr(\Pi^d < 1 \mid I^d = 1)$, monotonically increases: it reaches above 0.6 when $d = 1280$, indicating a substantial asymmetry.

Golosov and Lucas (2006) propose a model about firms' pricing decisions in an environment with both idiosyncratic and aggregate productivity disturbances. Probably it would be safe to say that we observe the effects of idiosyncratic disturbances on price change distributions when we choose small values for d, such as 5 and 10 days, while we observe those of aggregate disturbances on price change distributions when we choose large values for d, such as 640 and 1280 days. If this is the case, the results presented in Figure 7 show that idiosyncratic disturbances themselves are symmetric, which is consistent with the assumption of symmetric idiosyncratic shocks adopted by Golosov and Lucas (2006), while aggregate disturbances, including monetary policy shocks, are strongly asymmetric.

[Insert Figure 8]

Asymmetric distributions on a long time scale themselves might not be so surprising, but one may wonder why the probability of a price decrease (not a price increase) becomes higher on a longer time scale. This might be related to the fact that our sample period overlaps at least partially with the period when the Japanese economy experienced deflation. To investigate further on this possibility, we plot in Figure 8 the two probabilities, $\Pr(\Pi^d > 1 \mid I^d = 1)$ and $\Pr(\Pi^d < 1 \mid I^d = 1)$, for each year of our sample period, 1988-2005.

First, we see that the probability of a price decrease was higher than 0.5 from 2000 onward until 2005: we observe particularly strong asymmetry in 2004 and 2005. If we recall that the Japanese rate of inflation, measured either by the CPI or by the GDP deflator, was below zero during the period of 2000-2005, we may be allowed to interpret this asymmetry as reflecting deflation during this period. The second finding from Figure 8 is that there was another asymmetry at the beginning of the 1990s, in the sense that the probability of a price increase (rather than the probability of a price decrease) significantly exceeded 0.5. This period was a final stage of the asset price bubble in which the CPI inflation rate eventually started to rise fairly sharply. The observed asymmetry might have arisen from such an inflationary pressure in the Japanese economy.

5 Conditional Price Change Distributions

The probability of price adjustments and price change distributions could both potentially depend on past events. Among various types of history dependence, researchers have been interested in the dependence of the price adjustment probability upon how long the time has passed since the last price change. For example, Campbell and Eden (2006) finds from the US scanner data that the price adjustment probability is inversely correlated with price duration, i.e., a decreasing hazard function.

Similar results have been reported by Alvalez et al. (2005) for Eurpean countries.[4]

An important thing to note here is that price duration could be correlated not only with the price adjustment probability, but also with the price change distribution conditional on the occurence of a price change. The latter correlation is the main interest of this section. To our knowledge, there is no serious studies on the latter correation as far as prices of goods and services are concerned.[5]

The menu cost hypothesis has clear implications about these two kinds of history dependence. First, the hazard function should be upward sloping. Under the assumption that the variance of the target price monotonically increases with time, the deviation of the actual price from the target price becomes larger as the time elapses since the last price adjustment. Thus the probability of price adjustment is an increasing function of price duration.

Second, the price change distribution should be independent of price duration. According to a simple *Ss* rule, price adjustment occurs only when the deviation from the target price reaches *s*, and the new price is set at *S*, therefore the price change always equals to *S*-*s*.[6] There is no mechanism in such a simplified version of menu cost models that would yield a correlation between the price change distribution and price duration. More generally, the deviation of the actual price from the target price is a state variable in menu cost models, and when the value of this state variable meets a specific condition, price adjustment occurs. Put differently, if we were to gather instances of price adjustments, in each instance, the size of the deviation of the actual price from the target price should be identical. As long as firms base their decision on the magnitude of price adjustments in correspondence to the magnitude of the deviation, identical price change distributions should be observed, irrespective of how long the time have passed since the last price change.

5.1 Duration and the price change probability

We start by looking at the relationship between price duration and the probability of price adjustments. We denote price duration as n, where $n \geq 0$, and calculate the probability defined by

$$\Pr[I_t^d = 1 \mid I_{t-d}^d = I_{t-2d}^d = \cdots = I_{t-nd}^d = 0,\ I_{t-(n+1)d}^d = 1]. \tag{7}$$

[4]These studies deal with prices for those goods and services that are typically included in consumer price indexes, while other studies, such as Engle and Russel (1998), investigate history dependence for asset prices. For example, Zhang et al. (2001) estimate a hazard function for the IBM stock price, and find that it is not even monotonic, but is of an inverted U shape.

[5]However, there are several studies that investigate the relationship between asset price change distributions and price duration. See, for example, Russel and Engle (1998) that studies such a relationship using IBM stock prices.

[6]See, for example, Sheshinski and Weiss (1977).

for each product subgroup. Figure 9 presents the results: Figure 9.1 for the case of $d = 5$ and Figure 9.2 for the case of $d = 20$. The horizontal axis depicts nd, which express price duration by the number of days since the last price adjustment.[7] As we can see from Figure 9.1, there is a clear negative correlation between price duration and the price change probability for the product subgroup $G0$. Even for other subgroups, the price change probability is inversely correlated with price duration, except that there is a small spike at 20 days for these subgroups. A similar result is obtained in Figure 9.2. These results are consistent with the existing studies on the shape of hazard functions.

[Insert Figure 9]

5.2 Duration and the price change distribution

The observed negative correlation between price duration and the price adjustment probability is obviously inconsistent with the menu cost hypothesis, but it has an even more unpleasant implication about the dynamic evolution of prices. That is, under the assumptions that (1) the variance of target prices increases with time, and that (2) price change distributions are independent of price duration, a decreasing hazard implies that the deviations of actual prices from target prices monotonically increases with time. If we believe that firms never allow actual prices to substantially deviate from target prices, then we have to throw away either of the two assumptions. One possibility is that price change distributions are correlated with price duration in that large price changes become more likely to occur as the time passes since the last price adjustment.

[Insert Figure 10]

To examine this possibility, we first calculate

$$\log \Pr[\Pi_t^d < c \mid I_{t-d}^d = I_{t-2d}^d = \cdots = I_{t-nd}^d = 0,\ I_{t-(n+1)d}^d = 1] \qquad \text{for} \quad c \leq 1 \tag{8}$$

$$\log \Pr[\Pi_t^d > c \mid I_{t-d}^d = I_{t-2d}^d = \cdots = I_{t-nd}^d = 0,\ I_{t-(n+1)d}^d = 1] \qquad \text{for} \quad c \geq 1 \tag{9}$$

for each product subgroup, with d being set at $d = 5$.[8] And then we estimate distortions of price change distributions at the vicinity of zero inflation using the same method as in Section 4.3. Table 4 presents estimates for such distortions. If we look at the upper panel, which shows estimated distortions for

[7] Note that we calculate the price change probability even for $n = 0$, i.e., $\Pr[I_t^d = 1 \mid I_{t-d}^d = 1]$. In this sense Figure 9 differs from the usual hazard function which starts at $n = 1$.

[8] Figures 10.1 and 10.2 show CDFs defined by (8) and (9) for product subgroup $G4$.

$\Pi < 1$, we see almost no correlation between the estimated distortions and price duration for the subgroup $G0$ and $G1$. In other words, price change distributions are independent of price duration in terms of relative frequency of small and large price changes, which is consistent with the menu cost hypothesis. However, if we look at the columns for $G2$, $G3$, and $G4$ at the upper panel, we see a positive correlation between the estimated distortions and price duration. This tendency is more clearly seen at the lower panel, which shows results for $\Pi > 1$: for example, the estimated distortion for $G4$ increases from 0.051 at the duration of 10 days to 0.095 at the duration of 40 days, and finally to 0.228 at the duration of 70 days. In other words, as far as positive inflation is concerned, small (positive) price changes become less likely to occur relative to large (positive) price changes as the time elapses since the last price adjustment.

[Insert Table 4]

Together with the observed negative correlation between the price adjustment probability and price duration (i.e., a decreasing hazard function), this result implies that although the price adjustment probability declines the longer the price duration, once a price adjustment occurs, the price change is likely to be large. This could be interpreted as indicating that firms raise prices by a substantial amount in order to eliminate large deviations from target prices, which is an unavoidable consequence of the long absence of price adjustments. Moreover, it should be noted that there is a clear asymmetry between a price increase and decrease; namely, a large price increase is likely to occur after the long absence of price adjustments, while the same thing is not true for a large price decrease. These empirical results suggest that menu costs increase as the time passes since the last price adjustment.

6 Conclusion

This paper has empirically examined the menu cost hypothesis by analyzing the distribution of price changes using daily scanner data covering almost all items sold at about 200 Japanese supermarkets. We have arrived at the following findings. First, as implied by the menu cost hypothesis, small price changes were indeed rare. The price change distribution for goods with sticky prices has a dent in the center. In contrast, no such dent can be observed in the price change distribution for goods with flexible prices. Second, we find that the longer the time that has passed since the last price change, the higher is the probability that a large price change occurs. Combined with the fact that the price adjustment probability is a decreasing function of the price duration, this means that although the price adjustment probability decreases as the price duration increases, once a price adjustment occurs, the magnitude of such an adjustment is large. Third, while the price change distribution is symmetric on a short time scale, it is asymmetric on a long time scale, with the probability of a price decrease being significantly larger than the probability of a price increase. The asymmetry on a long time scale seems to be related to the deflation that the Japanese economy has experienced.

References

[1] Alvarez, L. J., P. Burriel, and I. Hernando (2005). "Do Decreasing Hazard Functions for Price Changes Make Sense?," Working Paper No. 461, European Central Bank.

[2] Bils, M., and P. J. Klenow (2004). "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy* 112(5), 947-985.

[3] Caballero, Ricardo J., and Eduardo M.R.A. Engel (2006). "Price Stickiness in *Ss* Models: Basic Properties," MIT.

[4] Calvo, Guillermo A. (1983). "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics* 12, 383-98.

[5] Campbell, J. R., and B. Eden (2006). "Rigid Prices: Evidence from U.S. Scanner Data," Vanderbilt University.

[6] Carlton, D. W. (1986). "The Rigidity of Prices," *American Economic Review* 76(4), 637-658.

[7] Dotsey, M., R. King, and A. Wolman (1999). "State-Dependent Pricing and the Genaral Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics* 114(2), 655-690.

[8] Engle, Robert F., and Jeffrey R.Russell (1998). "Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data," *Econometrica* 66, 1127-1162.

[9] Gagnon, E. (2005). "Price Setting Under Low and High Inflation: Evidence from Mexico," Northwestern University.

[10] Golosov, M., and R. E. Lucas (2006). "Menu Costs and Phillips Curves," MIT.

[11] Kackmeister, A. (2005). "Yesterday's Bad Times are Today's Good Old Times: Retail Price Changes in the 1890's were Smaller, Less Frequent, and More Permanent," Finance and Economics Discussion Series, Federal Reserve Board.

[12] Kashyap, A. K. (1995). "Sticky Prices: New Evidence from Retail Catalogs," *Quarterly Journal of Economics* 110(1), 245-274.

[13] Klenow, P. J., and O. Kryvtsov (2005). "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation," Stanford University.

[14] Lach, S., and D. Tsiddon (2005). "Small Price Changes and Menu Costs," The Hebrew University of Jerusalem

[15] Midrigan, V. (2005). "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations," Ohio State University.

[16] Nakamura, E., and J. Steinsson (2007). "Five Facts about Prices: A Reevaluation of Menu Cost Models," Harvard University.

[17] Russell, Jeffrey R., and Robert F. Engle (1998). "Econometric Analysis of Discrete-Valued Irregularly-Spaced Financial Transactions Data Using a New Autoregressive Conditional Multinomial Model," UCSD Economics Discussion Paper 98-10.

[18] Sheshinski, E., and Y. Weiss (1977). "Inflation and Costs of Price Adjustment," *Review of Economic Studies* 44(2), 287-303.

[19] Zhang, Michael Yuanjie, Jeffrey R. Russell, and Ruey S. Tsay (2001). "A Nonlinear Autoregressive Conditional Duration Model with Applications to Financial Transaction Data," *Journal of Econometrics* 104, 179-207.

Table 1: Number of Outlets, Products, and Observations

	No. of outlets	Entry	Exit	No. of products	No. of observations
1988	29	-	-	380232	22361602
1989	45	16	0	550567	37701574
1990	49	5	1	628840	44520684
1991	53	4	0	643970	48946978
1992	59	9	3	720541	53849713
1993	65	6	0	795821	59106207
1994	101	38	2	1339484	85605788
1995	110	23	14	1753299	113546038
1996	130	22	2	2118326	142485548
1997	146	20	4	2640191	162906492
1998	168	26	4	3346375	206299100
1999	168	4	4	3524924	216163437
2000	185	21	4	4169871	228330310
2001	185	2	2	4480215	260976696
2002	186	13	12	4952856	267585172
2003	185	3	4	4606987	230512281
2004	172	17	30	5058395	265404717
2005	181	15	6	5046844	279849443

Table 2: Turnover of Products in the 17 Outlets

	Number of products in the 17 outlets	Entry	Exit	Entry rate	Exit rate
1989	231942	80107	73882	0.345	0.319
1990	242605	84545	78180	0.348	0.322
1991	231488	67063	71559	0.290	0.309
1992	221828	61899	63364	0.279	0.286
1993	222170	63706	64353	0.287	0.290
1994	260540	102723	72713	0.394	0.279
1995	313019	125192	91897	0.400	0.294
1996	343012	121890	101909	0.355	0.297
1997	380912	139809	117468	0.367	0.308
1998	409046	145602	133590	0.356	0.327
1999	422121	146665	142888	0.347	0.339
2000	440935	161702	150383	0.367	0.341
2001	439873	149321	149506	0.339	0.340
2002	450879	160512	158003	0.356	0.350
2003	425856	132980	147450	0.312	0.346
2004	466156	187750	167677	0.403	0.360

Table 3: Prediction Errors by Product Subgroup

Predicted minus Actual	G0	G1	G2	G3	G4	G5	G6	G7	G8	G9
$\log \Pr(\Pi^5 < 0.93 \mid I^5 = 1)$	-0.003	0.052	0.068	0.072	0.061	0.039	0.021	0.011	0.003	0.017
$\log \Pr(\Pi^5 < 0.96 \mid I^5 = 1)$	-0.018	0.114	0.158	0.165	0.145	0.104	0.063	0.023	-0.043	-0.083
$\log \Pr(\Pi^5 < 1.00 \mid I^5 = 1)$	-0.109	0.170	0.256	0.271	0.242	0.179	0.115	0.056	-0.028	-0.063

Predicted minus Actual	G0	G1	G2	G3	G4	G5	G6	G7	G8	G9
$\log \Pr(\Pi^5 > 1.07 \mid I^5 = 1)$	-0.004	0.053	0.068	0.075	0.065	0.039	0.011	-0.006	-0.027	0.009
$\log \Pr(\Pi^5 > 1.03 \mid I^5 = 1)$	-0.020	0.116	0.161	0.172	0.158	0.117	0.074	0.042	-0.006	0.007
$\log \Pr(\Pi^5 > 1.00 \mid I^5 = 1)$	-0.113	0.170	0.260	0.284	0.269	0.217	0.165	0.140	0.104	0.103

Table 4: Prediction Errors by Price Duration

Predicted minus actual values for

$$\log \Pr[\Pi_t^d < 0.93 \mid I_{t-d}^d = \cdots = I_{t-nd}^d = 0,\ I_{t-(n+1)d}^d = 1]$$

Duration [days]	G0	G1	G2	G3	G4
10	0.043	0.061	0.069	0.066	0.029
15	0.046	0.064	0.071	0.075	0.039
20	0.047	0.067	0.075	0.081	0.049
25	0.048	0.068	0.074	0.090	0.080
30	0.045	0.065	0.067	0.080	0.062
35	0.048	0.066	0.074	0.081	0.061
40	0.045	0.068	0.077	0.090	0.047
45	0.043	0.070	0.080	0.095	0.057
50	0.043	0.071	0.085	0.098	0.075
55	0.038	0.067	0.081	0.105	0.074
60	0.037	0.065	0.076	0.093	0.045
65	0.039	0.064	0.079	0.097	0.060
70	0.032	0.065	0.081	0.091	0.075

Predicted minus actual values for

$$\log \Pr[\Pi_t^d > 1.07 \mid I_{t-d}^d = \cdots = I_{t-nd}^d = 0,\ I_{t-(n+1)d}^d = 1]$$

Duration [days]	G0	G1	G2	G3	G4
10	0.055	0.064	0.066	0.065	0.051
15	0.060	0.079	0.086	0.092	0.087
20	0.061	0.080	0.091	0.100	0.085
25	0.061	0.073	0.079	0.081	0.037
30	0.064	0.078	0.087	0.089	0.037
35	0.065	0.080	0.096	0.101	0.057
40	0.068	0.085	0.104	0.124	0.095
45	0.072	0.090	0.116	0.149	0.128
50	0.077	0.087	0.114	0.139	0.096
55	0.074	0.082	0.103	0.122	0.083
60	0.077	0.083	0.101	0.135	0.132
65	0.083	0.087	0.112	0.163	0.167
70	0.088	0.091	0.115	0.193	0.228

Figure 1: The Price Adjustment Probability by Product, Outlet, and Year

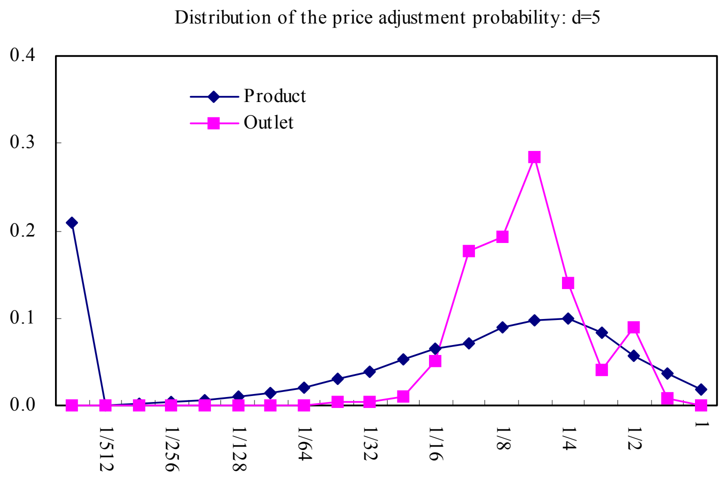

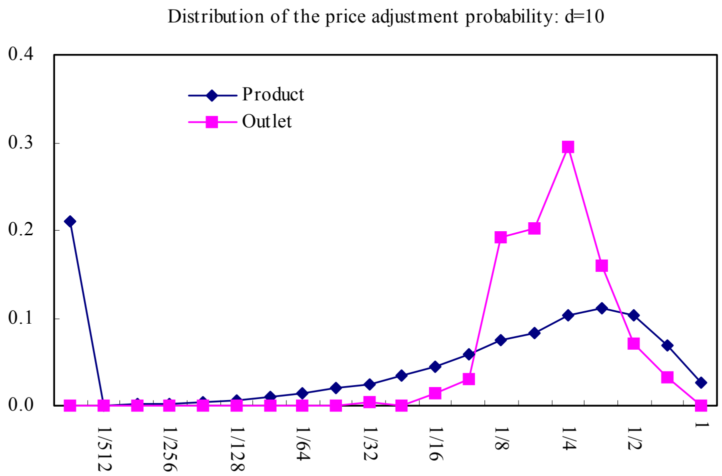

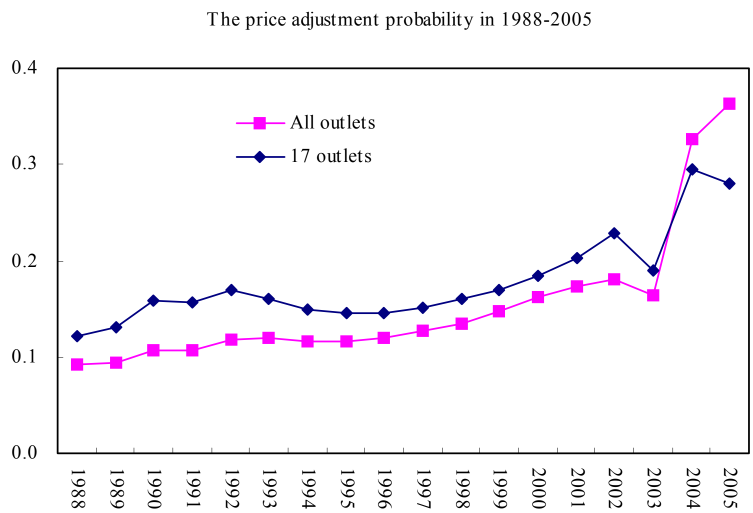

Figure 2: Price Change Distributions in 1988-2002

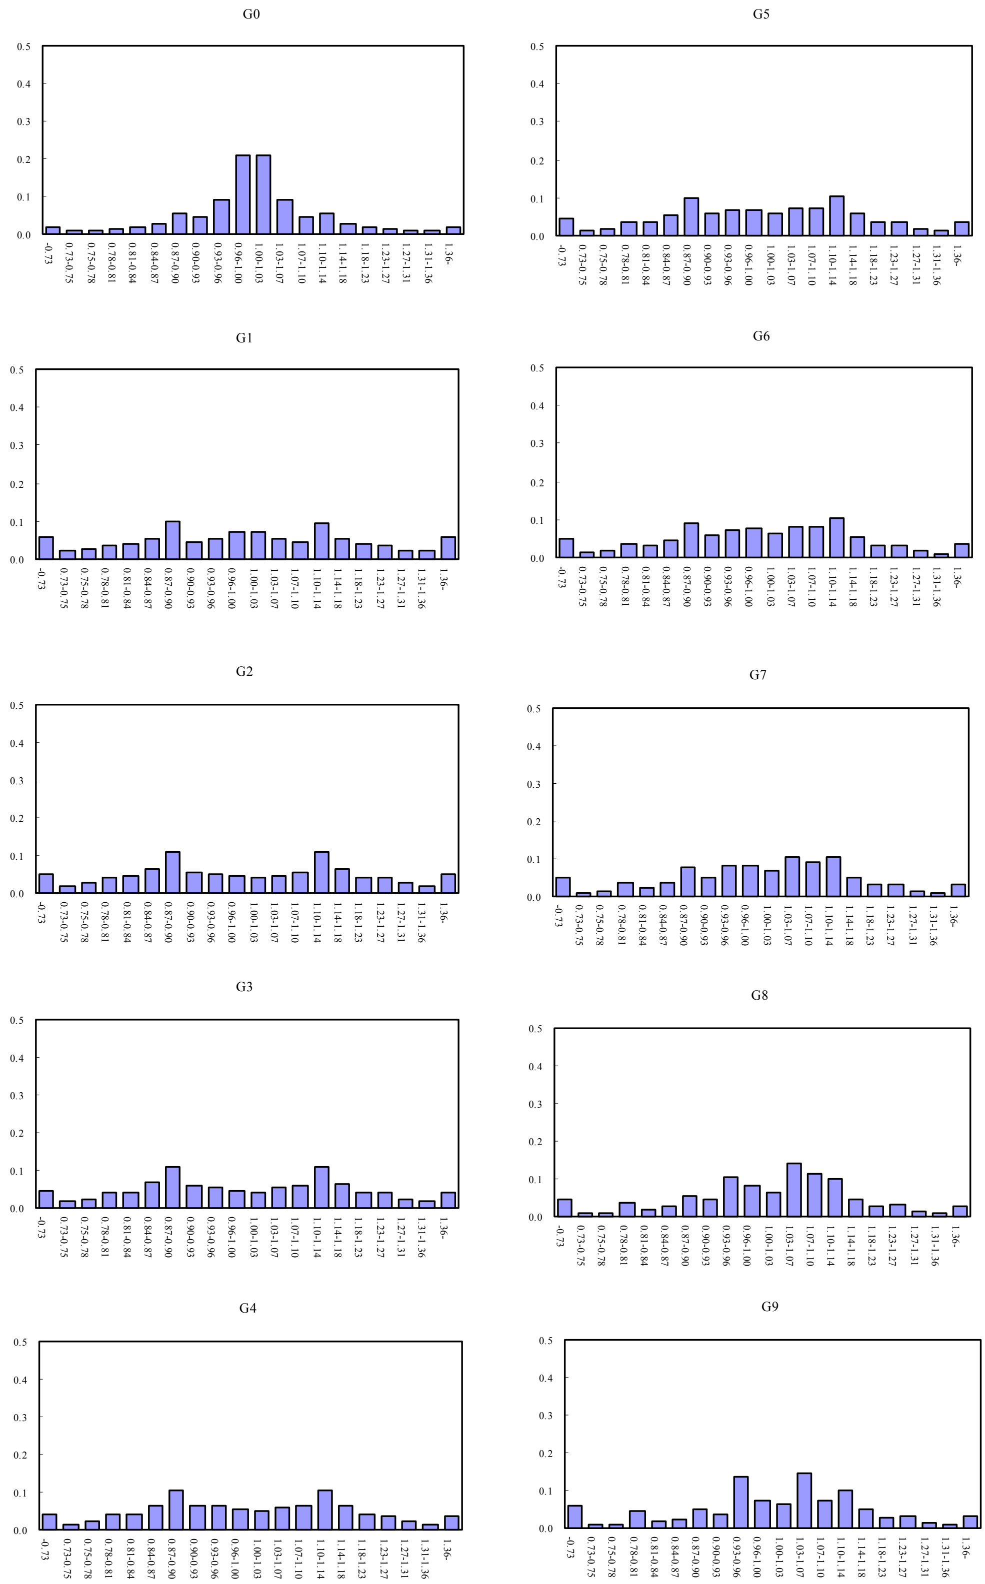

Figure 3: Price Change Distributions for Products with Prices being Above 200 Yen

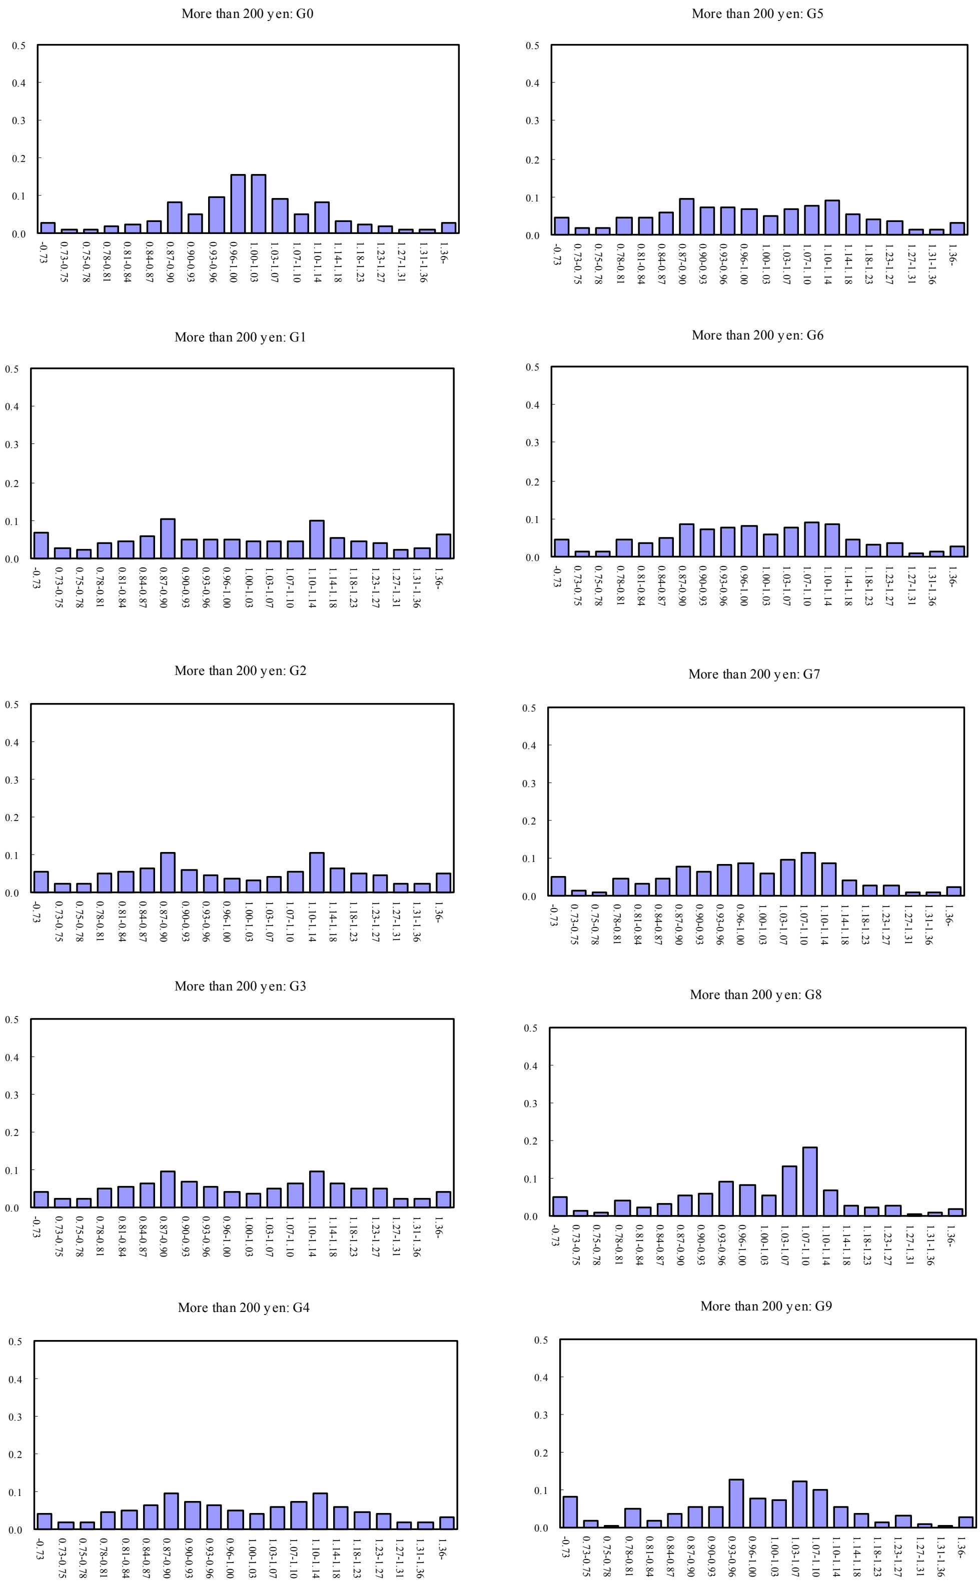

Figure 4: Price Change Distributions in 1988-2005

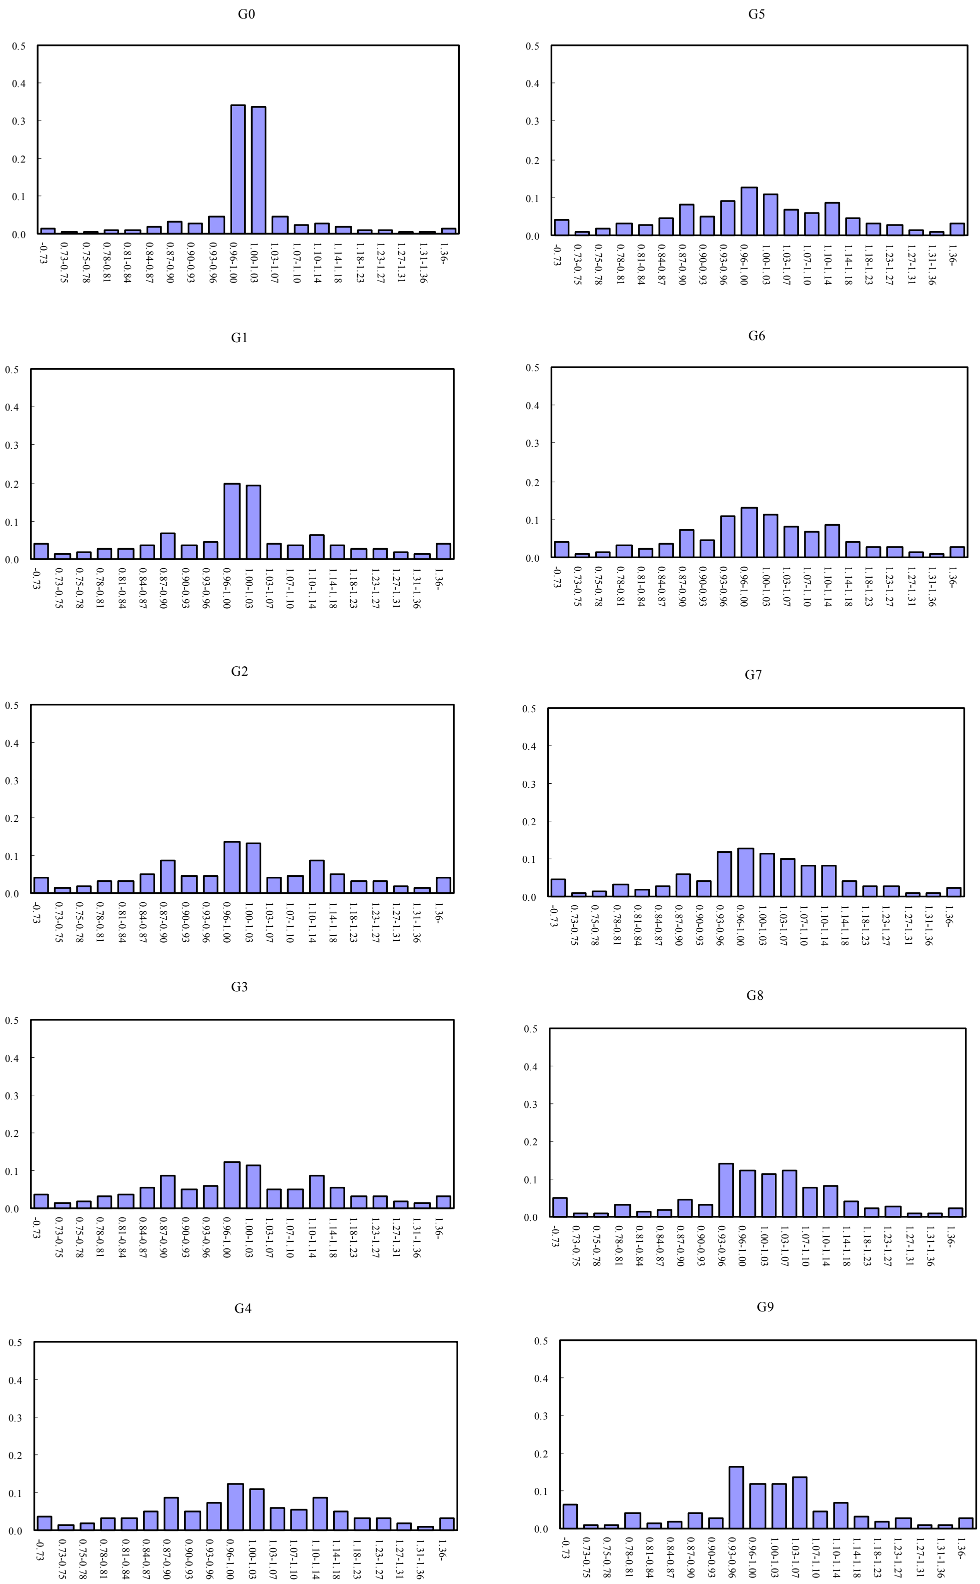

Figure 5.1: Cumulative Distribution Functions for Π<1

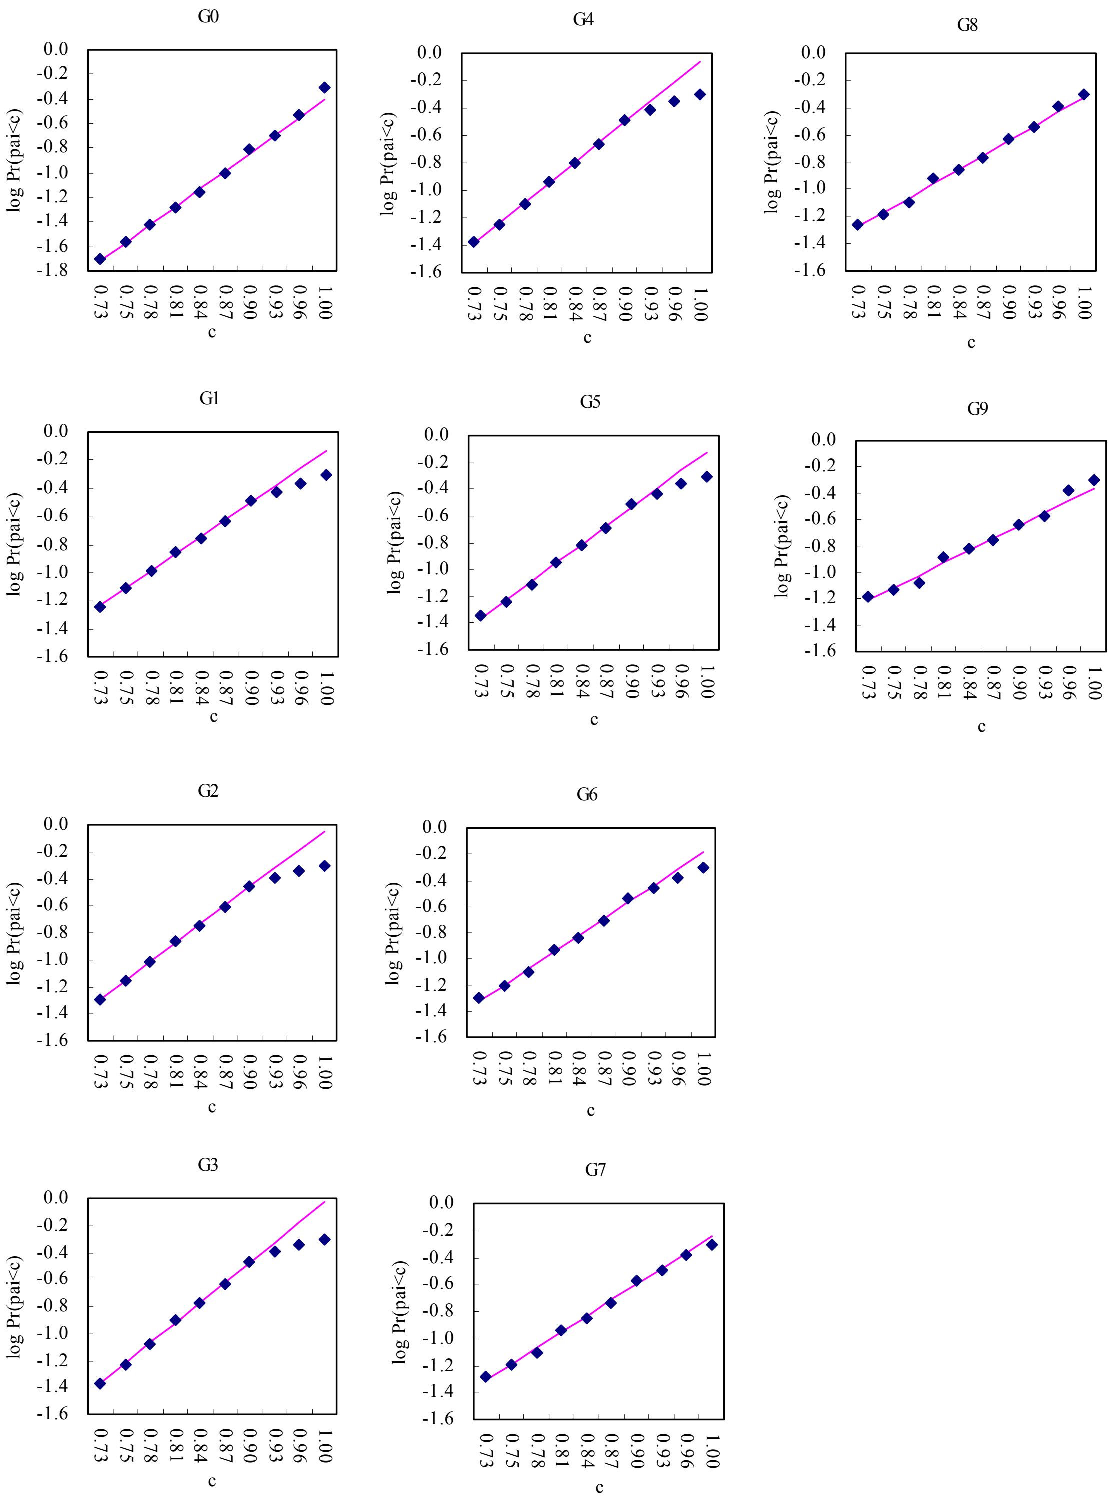

Figure 5.2: Cumulative Distribution Functions for Π>1

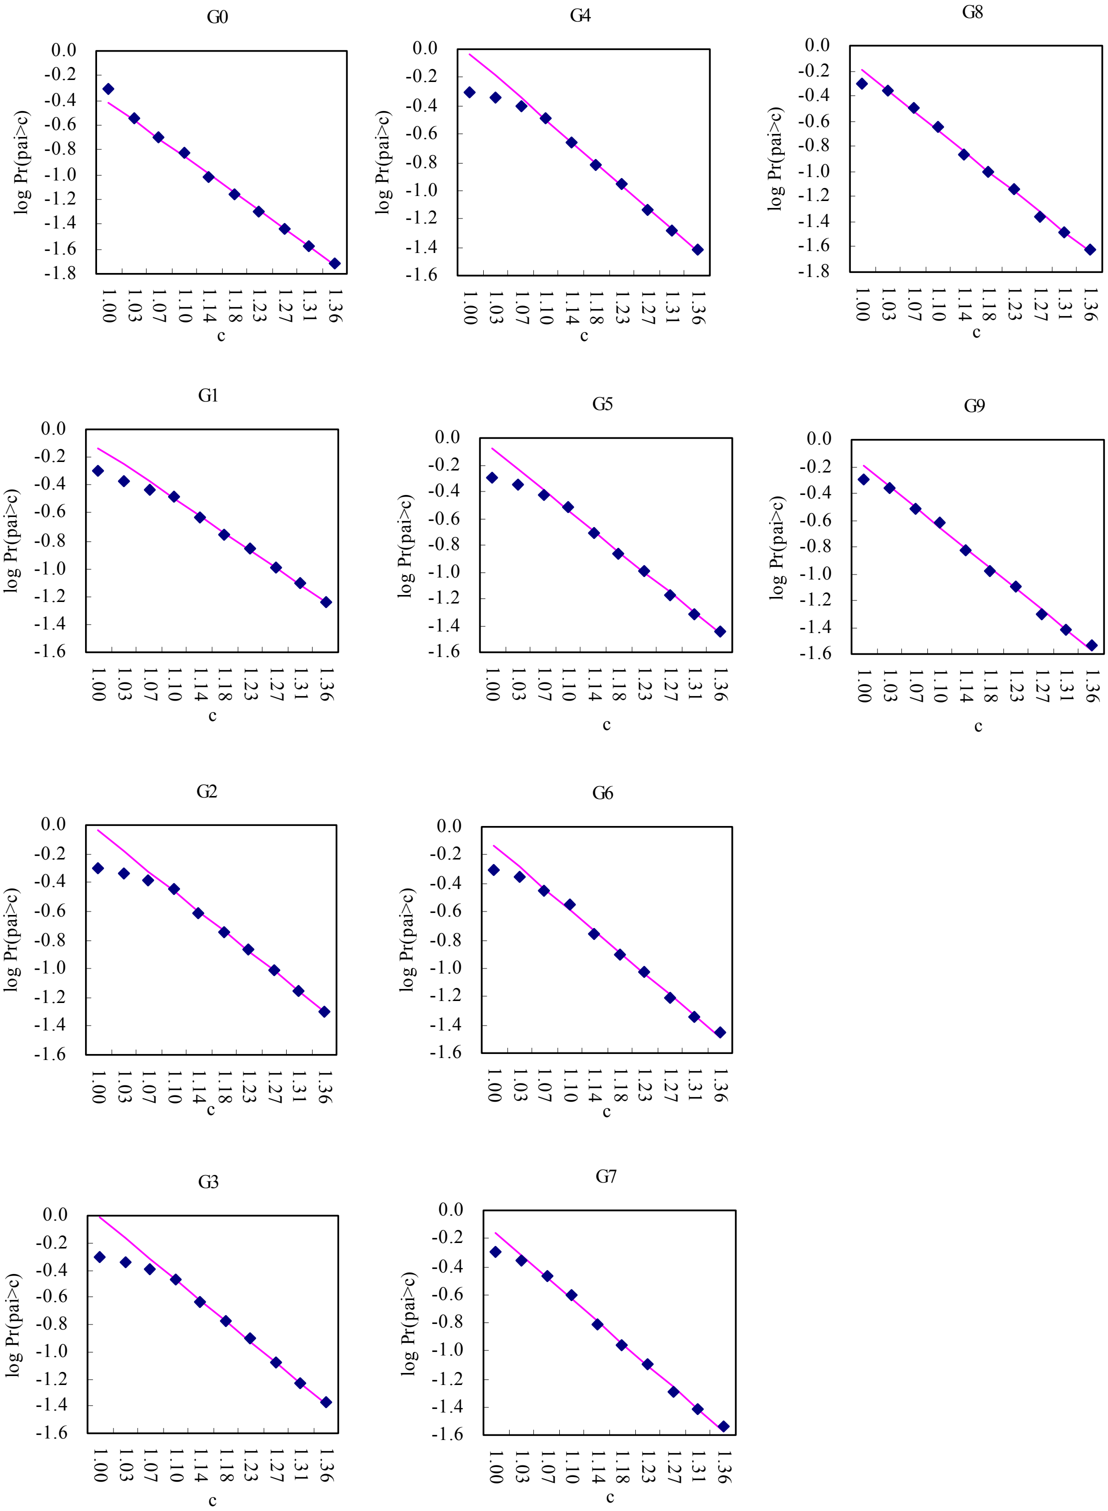

Figure 6: Estimated and Actual Probability Density Functions for G3

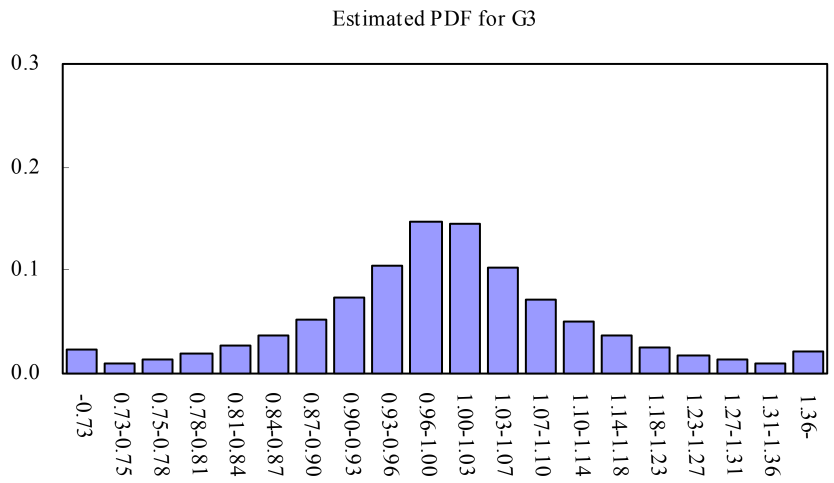

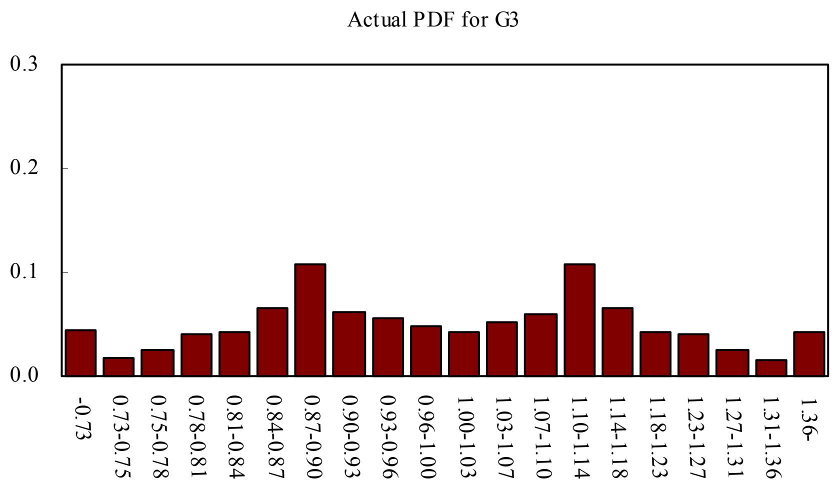

Figure 7: Asymmetry in Price Change Distributions for Different Time Scales

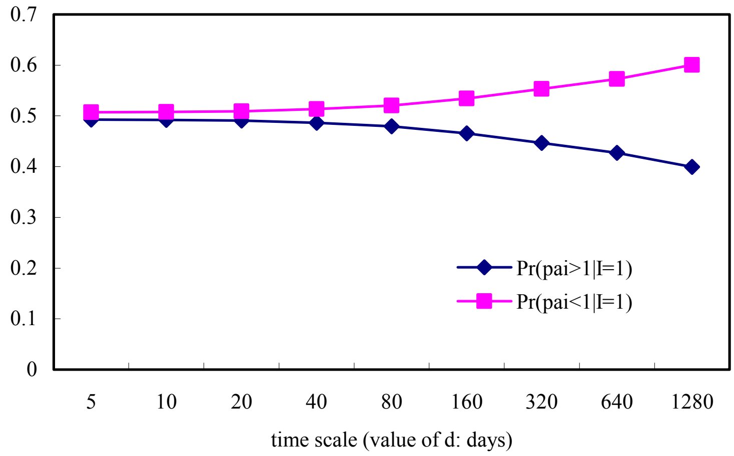

Figure 8: Asymmetry in Price Change Distributions in 1988-2005

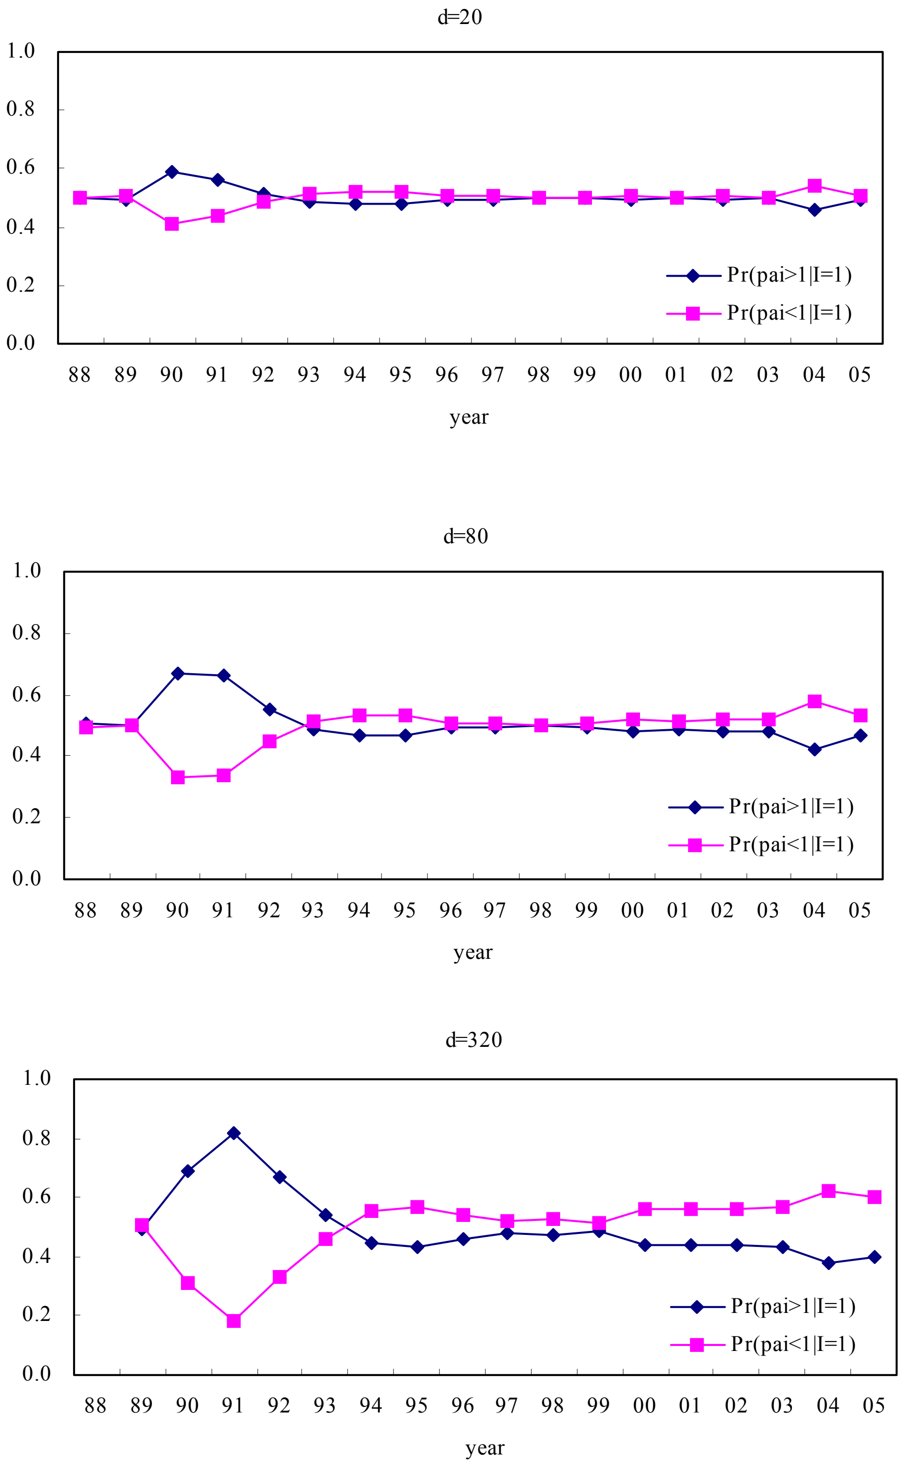

Figure 9.1: Duration and the Price Adjustment Probability, d=5

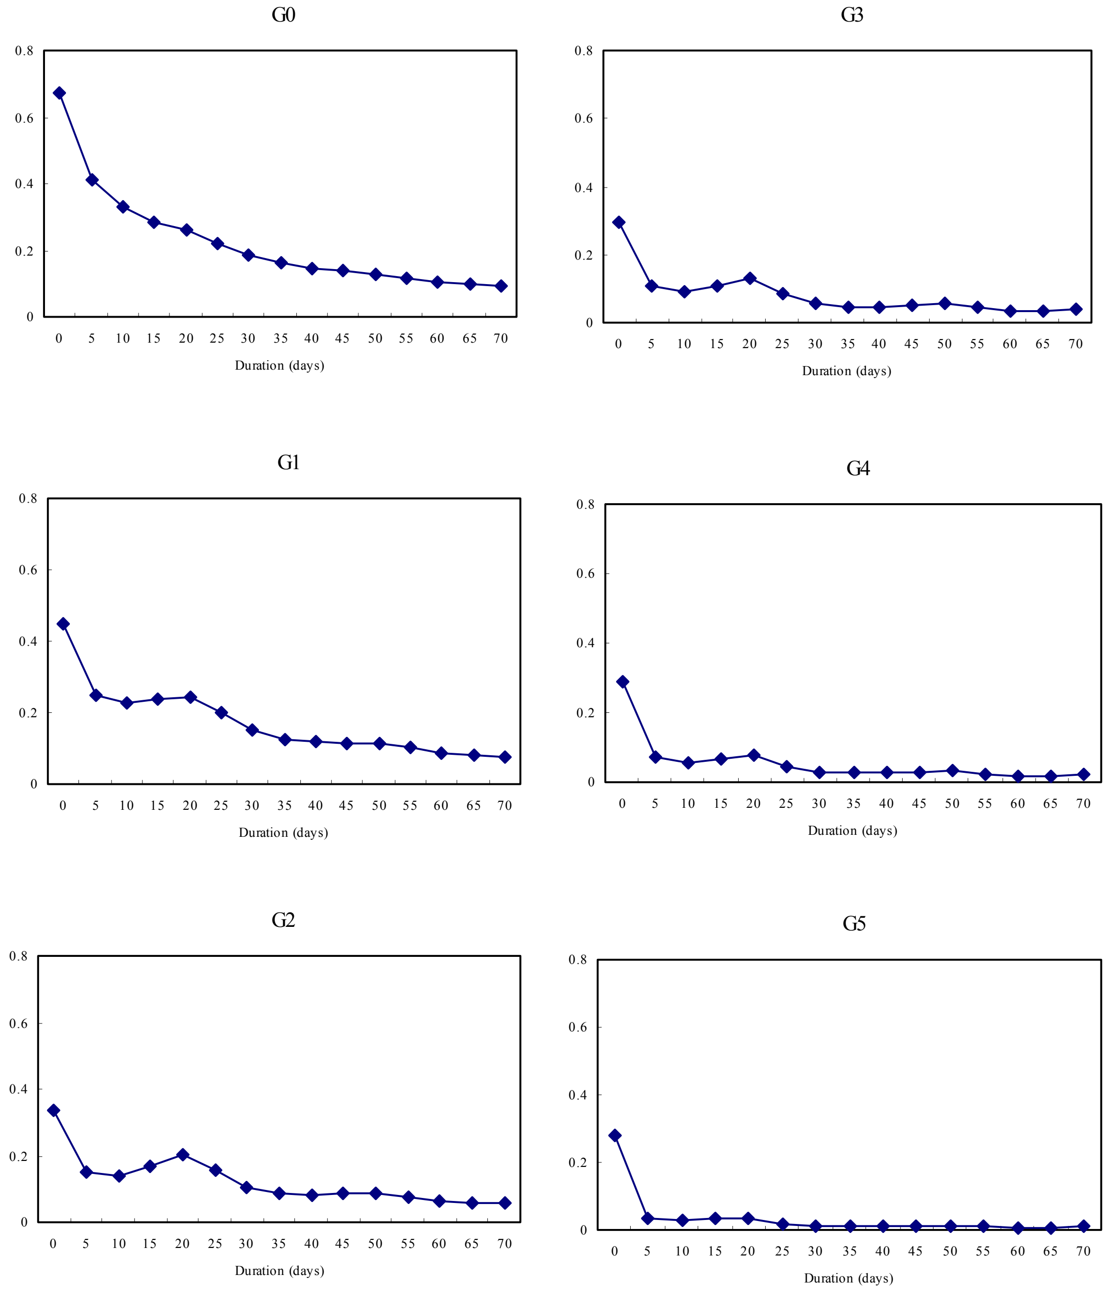

Figure 9.1: Duration and the Price Adjustment Probability, d=10

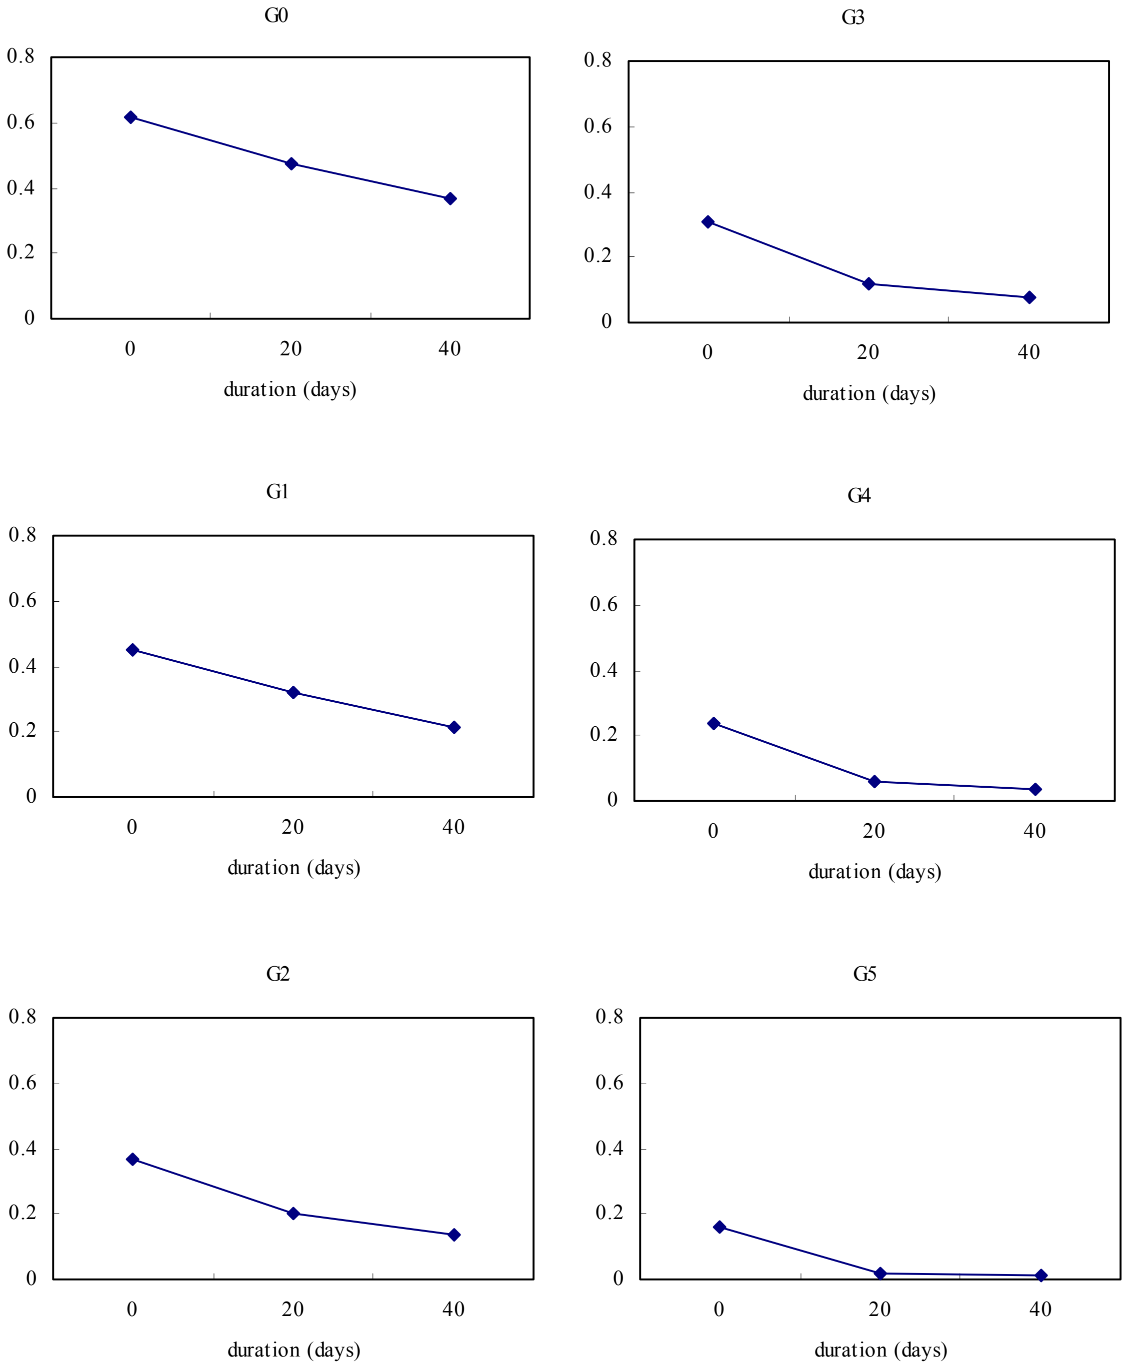

Figure 10.1: CDFs for Product Subgroup G4, Π<1

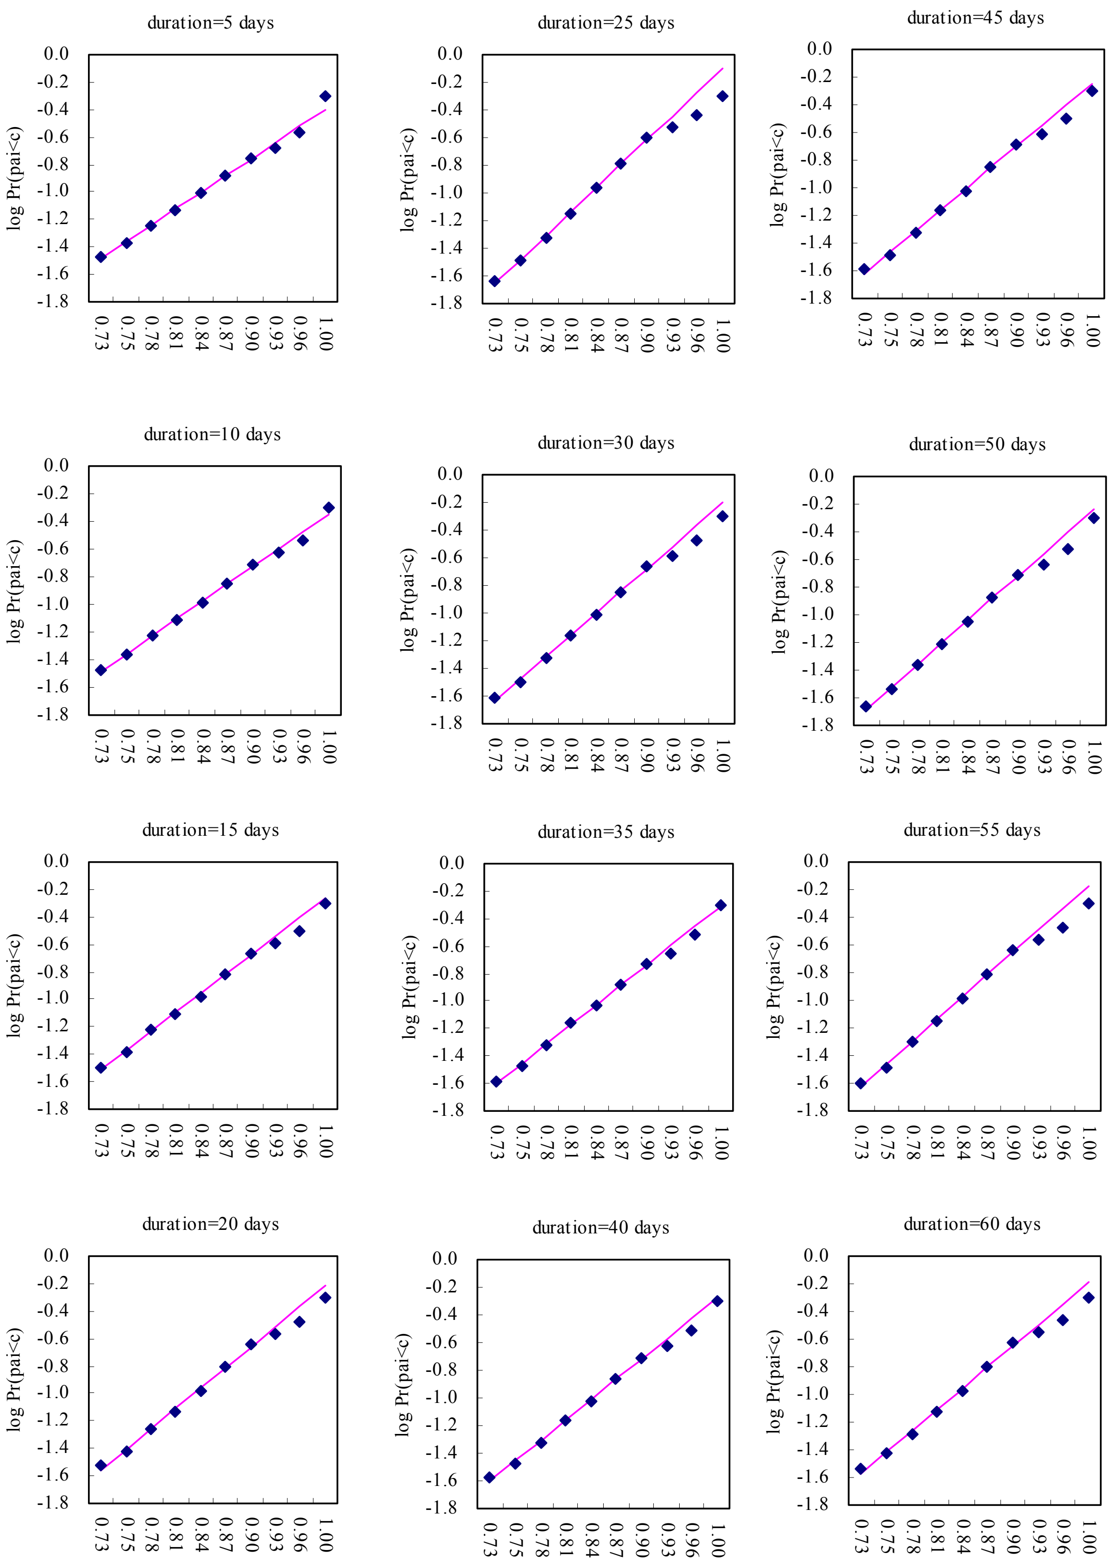

Figure 10.2: CDFs for Product Subgroup G4, Π>1