

JSPS Grants-in-Aid for Creative Scientific Research
Understanding Inflation Dynamics of the Japanese Economy
Working Paper Series No.44

Real Rigidities: Evidence from an Online Marketplace

Takayuki Mizuno

Makoto Nirei
and
Tsutomu Wanatabe

August 21, 2009

Research Center for Price Dynamics
Institute of Economic Research, Hitotsubashi University
Naka 2-1, Kunitachi-city, Tokyo 186-8603, JAPAN
Tel/Fax: +81-42-580-9138
E-mail: sousei-sec@ier.hit-u.ac.jp
<http://www.ier.hit-u.ac.jp/~ifd/>

Real Rigidities: Evidence from an Online Marketplace

Takayuki Mizuno* Makoto Nirei[†] Tsutomu Watanabe[‡]

August 19, 2009

Abstract

Are prices sticky due to the presence of strategic complementarity in price setting? If so, to what extent? To address these questions, we investigate retailers' price setting behavior, and in particular strategic interaction between retailers, using a unique dataset containing by-the-second records of prices offered by retailers on a major Japanese price comparison website. We focus on fluctuations in the lowest price among retailers, rather than the average price, examining how quickly the lowest price is updated in response to changes in marginal costs. First, we find that, when the lowest price falls rapidly, the *frequency* of changes in the lowest price is high, while the *size* of downward price adjustments remains largely unchanged. Second, we find a positive autocorrelation in the frequency of changes in the lowest price, and that there tends to be a clustering where once a change in the lowest price occurs, such changes occur in succession. In contrast, there is no such autocorrelation in the size of changes in the lowest price. These findings suggest that retailers imitate each other when deciding to adjust (or not to adjust) their prices, and that the extensive margin plays a much more important role than the intensive margin in such strategic complementarity in price setting.

JEL Classification Number: E30

Keywords: price stickiness; real rigidities; strategic complementarities in price setting; time-dependent pricing; state-dependent pricing; adjustment hazard; online markets

*Correspondence: Takayuki Mizuno, Institute of Economic Research, Hitotsubashi University, Kunitachi, Tokyo 186-8603, Japan. E-mail: mizuno@ier.hit-u.ac.jp. This research is a part of the project entitled: Understanding Inflation Dynamics of the Japanese Economy, funded by JSPS Grant-in-Aid for Creative Scientific Research (18GS0101), headed by Tsutomu Watanabe.

[†]Institute of Innovation Research, Hitotsubashi University. E-mail: nirei@iir.hit-u.ac.jp.

[‡]Research Center for Price Dynamics and Institute of Economic Research, Hitotsubashi University. E-mail: tsutomu.w@srv.cc.hit-u.ac.jp.

1 Introduction

Since Bils and Klenow's (2004) seminal study, there has been extensive research on price stickiness using micro price data. One vein of research along these lines concentrates on price adjustment events and examines the frequency with which such events occur. An important finding of such studies is that price adjustment events occur quite frequently. For example, using raw data of the U.S. consumer price index (CPI), Bils and Klenow (2004) report that the median frequency of price adjustments is 4.3 months. Using the same U.S. CPI raw data, Nakamura and Steinsson (2008) report that when sales are excluded, prices are adjusted with a frequency of once every 8 to 11 months. Similar studies focusing on other countries include Dhyne et al. (2006) for the euro area and Higo and Saita (2007) for Japan.

The frequency measure of price changes adopted in these studies is an indicator of price stickiness at the micro price level. When comparing the price stickiness at the micro price level suggested by this frequency measure with the price stickiness observed at the macro level, the estimated micro stickiness is clearly too low. Assuming that the measurements of both micro and macro stickiness are correct, a possible explanation for the discrepancy between the two is the existence of some sort of correlation in price adjustments across firms at the micro level. This idea, labeled "real rigidities" or "strategic complementarity" in price setting, has been advocated by Ball and Romer (1990), Kimball (1995), and others. For example, suppose that price adjustment events occur according to a Poisson process. Firms with an opportunity to adjust their prices may seek to imitate other firms that do not adjust prices, therefore keeping the extent of the price adjustment small. In this case, the intensive margin becomes small through strategic complementarity, while the fraction of adjusters, that is, the extensive margin, remains unchanged. The sum of the two margins becomes smaller, implying that prices become stickier at the macro level.

This means that to understand the extent of macro price stickiness, it is insufficient to only examine the frequency of micro price adjustments; instead, it is also necessary to examine the extent to which strategic complementarity in price setting is present. Against this background, the aim of the present paper is to empirically examine the existence and extent of strategic complementarity in price setting. To do so, what is important is to collect price data of competing firms and retailers, that is, price data for individual firms and retailers that potentially imitate each other. Unfortunately, such data are seldom available. For the compilation of CPI statistics, for example, a particular retailer is chosen to represent a particular region for a particular product and sales price data are then collected from that retailer. This means that it is virtually impossible to obtain price data from competing retailers from CPI raw data.

Another potential source is scanner data, but as far as the authors are aware, there exists no comprehensive dataset comprising prices for competing retailers. For example, Saito and Watanabe (2009) use a dataset consisting of prices collected from about 400 retailers, but those retailers do not necessarily compete in the same area, and the dataset therefore is not useful for examining strategic complementarity.

In order to investigate the price setting behavior of retailers that directly compete with and imitate each other, we use online market data in this paper. Specifically, the data we use are the selling prices offered by retailers on Kakaku.com, a major price comparison site in Japan. Transactions on this site concentrate on consumer electronics such as TV sets, video equipment, digital cameras, personal computers, etc. For example, for the AQUOS LC-32GH2, a liquid crystal television model made by Sharp, prices were provided by about 40 firms. These retailers have registered with Kakaku.com beforehand and have entered a contract stating that they will pay fees to Kakaku.com reflecting the number of customers transferred to their website via Kakaku.com. These 40 retailers monitor at what prices the other retailers offer the same product at any particular moment and based on this adjust the price at which they offer that product themselves. Thus, these retailers can be said to be engaged in moment-by-moment price competition in the virtual market provided by the price comparison site.

The kind of online market provided by Kakaku.com differs from regular (offline) markets on which research on price stickiness so far has concentrated in the following respects. In an online market, consumers can easily (and almost costlessly) see a list of the prices at which each retailer offers a particular product. Even though search costs for prices in online markets are not zero,¹ they are extremely small when compared with search costs in offline markets. This means that it is easy for consumers to find out the lowest price for a product at any given time as well as the price at which they could purchase the product if they went to a particular retailer. From the perspective of the retailer offering the lowest price, this implies that if it raises the price, it will lose out to other retailers and consumer “clicks,” that is, consumer demand, will decrease. As Baye et al. (2009) have shown, since search costs for prices in online markets are low, if a change in relative prices brings about a change in ranking, this will result in a discontinuously large change in demand. Thus, if a retailer falls from first rank to second rank in terms of price offered, it will face a substantial decline in demand; on the other hand, since the retailer already offers the lowest price, the increase in clicks achieved through a further reduction in its price is unlikely to be very large. In this sense, for a retailer offering the lowest price for a particular product, the price elasticity of demand differs greatly for a price increase and a price decrease. For a retailer facing such a kinked demand curve, the

¹See, for example, Ellison and Ellison (2009).

optimal strategy is to leave the price it offers unchanged, even if there is a change in marginal costs, such as procurement prices, at least up to a degree.² In other words, as long as other retailers do not move, the optimal strategy for this retailer is to also not move itself, and in this sense, strategic complementarity in price setting arises. Such strategic complementarity is likely to be stronger in online markets, where the costs involved in the propagation of price information to other firms in the industry and consumers are low, than in offline markets.

In this paper, following Caballero and Engel (2007), it is assumed that each retailer has a target price for a product, and a retailer, who is allowed to adjust its price, will set it in accordance with this target price. We will assume that this target price depends on marginal costs, as usual, but it also depends on the prices offered by other retailers. This is how strategic complementarity in price setting is incorporated. Using this setting, we are able to see how a price cut by retailer A, which is driven by a decline in marginal costs, induces a price cut by retailer B, which in turn induces further responses by retailers C, D, E, and so on. What we would like to know is how quickly (or how slowly) prices respond to a decline in marginal costs during such a chain reaction of price cuts. To do so, in this paper, we will focus on the price offered by adjusters, thereby isolating the effect of strategic complementarity on the adjustment speed from other effects. In this respect, it is important to note that it is the lowest price (i.e., the lowest price among all prices offered by retailers in a particular period) rather than the average of adjusters' prices that has the largest impact on retailers' pricing behavior in the next and subsequent periods. In this sense, the lowest price is a key variable in understanding strategic complementarity in price setting. Given this understanding, we will pay special attention to fluctuations in the lowest price.

The main findings of this paper can be summarized as follows. First, looking at time series of the lowest price, one notable characteristic is that when the lowest price falls rapidly, the frequency of changes in the lowest price is high. That is to say, at a time when the lowest price declines rapidly, each retailer scrambles to be the one offering the lowest price, and the outcome is that the lowest price is revised frequently. On the other hand, though, the size of the adjustment when the lowest price is updated does not change much even when the lowest price falls rapidly. This result suggests that the price setting of each retailer is not characterized by time dependent pricing but by state dependent pricing; at the same time, it suggests that changes in the lowest price are affected not by changes in the intensive margin (the size of price declines) but by changes in the extensive margin (the frequency of price declines).

²That the existence of this kind of kinked demand curve may give rise to price stickiness was already argued by Negishi (1979) in the context of an oligopolistic competition model. Moreover, joining time dependent pricing à la Calvo (1983) and a smoothed-off kinked demand curve, Kimball (1995) shows that strategic complementarity arises and macro price rigidity increases.

The second finding is that the longer it is since the lowest price has been adjusted, the lower tends to be the probability that the next price adjustment occurs. Put differently, there tends to be a clustering, so that once an adjustment of the lowest price occurs, such adjustments occur in succession. On the other hand, when the lowest price is not adjusted for a long time, the probability of a price adjustment event is low. These patterns indicate that retailers imitate each other when deciding to adjust their prices and also when deciding not to do so. Both types of behavior are an indication of strategic complementarity where each retailer's target price depends on the prices offered by other retailers.

Third and finally, we find a positive autocorrelation in the frequency of changes in the lowest price, which is consistent with the presence of clustering in changes in the lowest price. The frequency of changes in the lowest price is correlated with the frequency of such changes as much as a week earlier. In contrast, there is no such autocorrelation in the size of changes in the lowest price. This again suggests that the extensive margin plays a much more important role than the intensive margin in strategic complementarity in price setting.

The rest of the paper is organized as follows. Section 2 provides a description of the Kakaku.com dataset used in this paper and discusses some of the characteristics of these data. In Section 3, we then examine the characteristics of price changes using autocorrelations and hazard functions. Next, in Section 4, based on Caballero and Engel (2007), we construct a model of state-dependent pricing with strategic complementarities, conduct a simulation analysis, and examine whether the model can replicate the facts found in Section 3. Section 5 concludes the paper. In the Appendix, we solve the dynamic optimization problem encountered by retailers in determining their optimal strategy in a setting where the demand they face is determined by their price ranking.

2 Data Description

2.1 Overview

The data we use in this paper are from Kakaku.com, a major Japanese price comparison website. The website is operated by Kakaku.com, Inc., and at present about 1,300 retailers use the site as part of their sales activities. A wide range of products is offered through the website, but the most important are consumer electronics and personal computers, and if items with a different barcode are counted as separate products, about 300,000 products are offered. The number of monthly users is about 12 million.

By visiting the website, users of Kakaku.com can obtain information on the characteristics of a product they are interested in, find a list of retailers offering that product, and the price

at which each retailer offers that product. In addition, users can also obtain information on the characteristics of retailers such as whether they charge delivery fees (and if so, how much), whether they accept credit card payment, whether they accept cash on delivery, the address of the distribution center, whether they have an online shop, and retailer ratings by customers who have used the retailer. Consumers visiting Kakaku.com use this information to choose a retailer from which to purchase that product and then click a button on the Kakaku.com website saying “Go to retailer.” They are then transferred to the website of the selling retailer, where they go through the retailer’s sales procedure and finally purchase the product. Each retailer pays fees to Kakaku.com corresponding to the number of customers sent from the Kakaku.com website to that of the retailer, which is how Kakaku.com, Inc. generates its income. Kakaku.com, Inc. does not collect any fees directly from consumers.

There are various types of price comparison websites and some focus on gathering, on their own, prices advertised on internet websites and posting a list of these. However, different from this type of price comparison websites, the special characteristic of Kakaku.com is that Kakaku.com, Inc. and each retailer enter a contract (on the payment of fees depending on transferred customers, etc.) before any prices are listed. Therefore, retailers are well-informed about what other retailers have registered with Kakaku.com. Moreover, based on information sent to them by Kakaku.com, Inc., retailers check three or four times a day, or more frequently, the prices offered by other retailers, the overall rank of their own price, whether the number of customers transferred to their own site is large or small, and, if necessary, adjust their own price. Each retailer conducts such monitoring activity with regard to all the products it offers.

The dataset used in this paper consists of the records, with a time stamp up to the second, of all prices offered by each retailer (for a total of around 4 million records) and the history of customer clicks on the “Go to retailer” button (around 24 million records) for all products offered during the 731 days from November 1st, 2006 to October 31st, 2008. At which price retailers offered a product may be considered to be an indication of the supply-side situation of the product, while which shop consumers clicked on represents the demand side.³

This is not the first study to use prices from price comparison sites. However, few studies have used a dataset that includes information on how the price at which retailers offer a product changes over time, and which retailer consumers click on. The dataset closest to the one in this paper is the one used by Baye et al. (2009) from a price comparison site in the United Kingdom. However, their dataset consists of daily aggregated data and does not allow observations on

³Of course it is important to note that just because the “Go to retailer” button was clicked on, this does not necessarily mean that this ultimately resulted in a purchase. However, looking at the correlation between the number of customers referred from Kakaku.com and information on the number of actual sales obtained from scanner data for several of the retailers registered with Kakaku.com, it can be confirmed that this is extremely high. This result shows that the number of clicks is a sufficiently useful proxy for the actual number of sales.

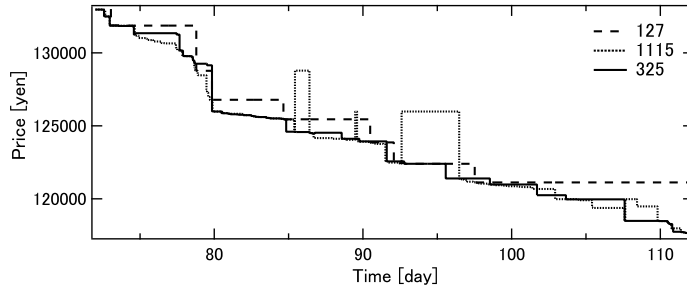


Figure 1: Illustration of prices offered by retailers

competition between retailers for the lowest price on an hour-by-hour, minute-by-minute, and, in some cases, even second-by-second basis.

Figure 1 shows an example of the fluctuations in the prices at which three retailers competing for the lowest price were offering a liquid crystal TV made by Sharp (AQUOS LC-32GH2). The figure illustrates the following. First, there is a strong downward trend in the price for this product, and within a period of about 50 days, the price fell from 130,000 yen to a little below 120,000 yen. However, what also becomes clear is that the retailers did not continuously lower prices day-by-day. Rather, after maintaining a particular price level for several days or weeks, they lowered the price discontinuously by several hundred or several thousand yen. Then, soon thereafter, they again maintained the same price level. This occurred repeatedly. The observation that price adjustment events are infrequent and that when they occur, price adjustments are discontinuous is in line with patterns shown in previous studies using CPI raw data or scanner data.⁴ Second, the three retailers do not adjust prices at the same time but instead each adjust their price at a different time and to a different extent. Although the prices offered by the three retailers overall show the same trend, a closer look reveals that the price gap between the retailers fluctuates, and that competition is fierce, with first one retailer and then the other taking the lead.

2.2 Frequency of price adjustments

Studies since Bills and Klenow (2004) using micro price data typically employ the frequency of price adjustments as a measure for price stickiness. Figure 2 shows the distribution of the price duration (that is, the interval from one price change until the next price change) for the AQUOS LC-32GH2 liquid crystal TV. The data we are using here are for the ten months from

⁴Studies using CPI raw data or scanner data report that sales (that is, temporary price drops) occur frequently. However, at least on the basis of Figure 1, it appears that the retailers on Kakaku.com do not conduct sales. In the dataset used in our paper, sales cannot be observed as frequently as in the case of offline retailers.

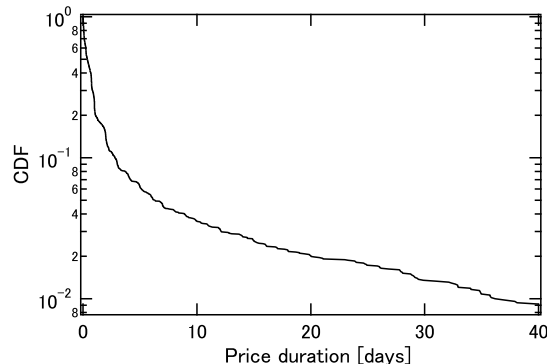


Figure 2: Cumulative distribution function of the price duration

November 2006 to June 2007. The number of retailers providing a price for this product during this period is not fixed, but on average there are 40 retailers. We use all the price spells for these retailers. Note that the period from when a retailer begins to offer this product until it first changes the price is not regarded as a price spell. Moreover, we similarly exclude the period immediately before a retailer stops offering a product.

In the figure, the horizontal axis shows the price duration, while the vertical axis shows the corresponding value of the cumulative distribution function (CDF). For example, the value on the vertical axis corresponding to a price duration of 10 days is 0.04, indicating that the share of price spells of 10 days or more in the total is 4 percent. The average price duration is 1.93 days, and the median is 0.34 days. In other words, on average, the probability that a price adjustment event occurs on any given day is 0.518 ($=1/1.93$), although the probability of price adjustment within 24 hours since the last price change is 0.750. If price adjustment events occur according to a Poisson process, price duration follows an exponential distribution. Because the vertical axis here is shown in logarithmic scale, if the price duration follows an exponential distribution, the measured CDF should form a straight line.⁵ However, as can be seen in the figure, the CDF takes a convex shape, suggesting that the price duration does not follow an exponential distribution.

In previous research using micro price data, it is repeatedly reported that the hazard function is downward sloping, that is, the longer the time since the last price adjustment, the lower is the probability of an adjustment event occurring. Figure 2 shows that the same is the case here. However, previous studies measure hazard functions using pooled data for several

⁵Denote price duration by y and its PDF and CDF by $f(y)$ and $F(y)$, respectively. Then, a hazard function $h(y)$ is related to the CDF of y as follows: $h(y) = -\frac{d}{dy} \log[1 - F(y)] = \frac{f(y)}{1 - F(y)}$. If the price duration obeys a Poisson process, $h(y)$ is constant, so that the derivative of $\log[1 - F(y)]$ with respect to y is constant as well, implying that the CDF forms a straight line in a log plot.

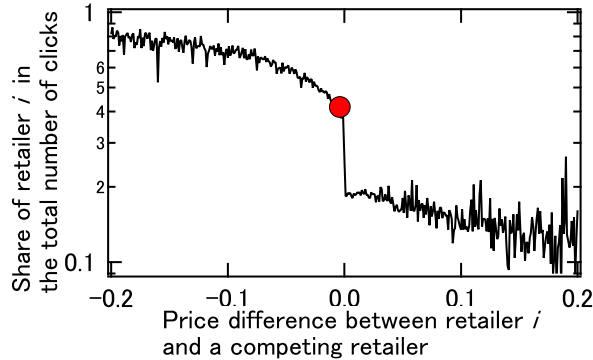


Figure 3: The relationship between price rank and the share in the number of clicks

products, and it has been pointed out that it is possible that the downward sloping hazard function may be the result of heterogeneity in price adjustment probability across different products. In contrast, the price spells used in Figure 2 is collected only from one product, so that this kind of problem of product heterogeneity does not arise, although we cannot rule out the possibility that there exists non-trivial heterogeneity in price adjustment probabilities across different retailers, and this may give rise to the convex shape.

2.3 Price ranks

In the online market provided by Kakaku.com, the price rank has an important role. To illustrate this, Figure 3 shows how the price rank for each retailer affects the probability that it will be clicked on by consumers for the AQUOS LC-32GH2 liquid crystal TV. Specifically, it shows, for the case that retailer i is ranked first or second, how price differences with competing retailers affect the number of clicks. On the horizontal axis, the figure depicts the price difference between retailer i and competing retailers for each point in time. For example, a value of -0.1 indicates that the price offered by retailer i , ranking first, is 10 percent lower than that of the second-ranked retailer; a value of 0.1 indicates that the price offered by retailer i in the second rank is 10 percent higher than that of the first-ranked retailer. The vertical axis shows the share of retailer i in the total number of clicks.

Let us assume that retailer i offers the lowest price and is in the position indicated by the red circle. If it raises the price even only a little, it will be overtaken by a rival retailer and its number of clicks will decline. According to the figure, as a result of the price increase, the number of clicks would fall discontinuously. In fact, through only a small price increase, the number of clicks would almost halve.⁶ On the other hand, if the retailer i , from the position

⁶While the figure shows the competitive relationship between the first and the second rank, the discontinuous

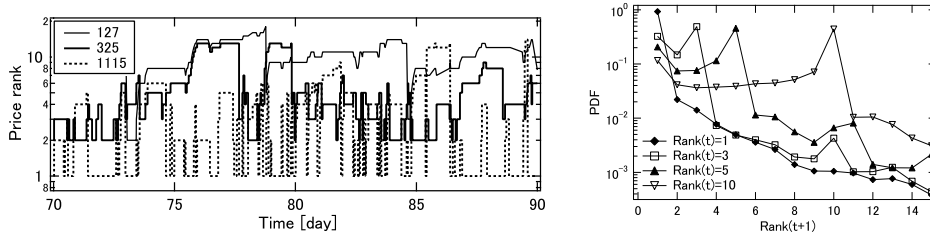


Figure 4: Transition in price rank for each retailer

indicated by the red circle, lowers the price, the number of clicks it receives will increase because the price advantage vis-à-vis the second-ranked retailer will expand further, but that increase in the number of clicks will not be that great. Thus, the elasticity of demand for an increase and a decrease in price differs and in this sense the demand curve is kinked. Therefore, for retailer i , even if marginal costs change, as long as that change is relatively small, the optimal behavior is not to move from the position indicated by the red circle. Negishi (1979) and Kimball (1995) have shown that in a situation where the demand curve is kinked and rivals do not move, the optimal response is not to move oneself, and that this strategic complementarity in price setting, in turn, raises macro price stickiness. Figure 3 suggests that it is possible that in online markets, where search costs are low, this mechanism works even more strongly.

Another aspect with regard to price rank concerns heterogeneity between retailers. If, for example, all retailers were homogenous and the demand curve kinked as shown in Figure 3, then all retailers would offer the same price and therefore all would occupy the first rank. However, this is not what occurs in practice. The left part of Figure 4 shows how the rank of the three retailers depicted in Figure 1 moves over time. This shows that retailer no. 1115 stayed in first rank for a long time and even when it fell into second place, it soon returned to first place. In contrast, retailer no. 127 stayed around 10th rank for most of the time. Thus, there is a tendency for top-ranking retailers to always remain in the top ranks, for middle-ranking retailers to always remain in that vicinity, and for low-ranking retailers to be always positioned in lower ranks. In order to show this point more clearly, the right part of Figure 4 examines what rank 1st, 3rd, 5th, and 10th-ranking retailers just before a price adjustment held immediately after adjusting their price. For instance, for retailers that held 5th rank before adjusting their prices, the probability that they ranked 5th, 4th, or 3rd after lowering their prices was high. A similar tendency can also be found for other ranks. These two figures

change in the number of clicks seen here can also be observed for the 2nd and 3rd rank, the 3rd and 4th rank, etc. A similar discontinuity in the number of clicks is found by Baye et al. (2009) using data from a British price comparison site.

show that each retailer has a particular rank that it targets, but this target rank differs across retailers.⁷

The fact that each retailer has a different target rank has an important implication for how the empirical analysis should be conducted using this dataset. For example, retailers that aim for the first rank will probably imitate other retailers aiming for the first rank when it comes to deciding whether to adjust or not adjust prices. On the other hand, such retailers would be very unlikely to imitate retailers targeting, say, the fifth rank. This implies that if we pool all retailers, with their different target ranks, it is highly likely that we will fail to detect the presence of strategic complementarity. To avoid this, we need to collect observations for a set of retailers that are homogeneous in terms of their target rank. The strategy we will take in this paper is to focus on the pricing behavior of those retailers who aim for the first rank, thereby removing heterogeneity in terms of target rank. Given this strategy, we are particularly interested in the lowest price, since this price is an important determinant (maybe the most important determinant) of the pricing behavior of those retailers who aim for the first rank. In this sense, the lowest price is a key variable in understanding strategic complementarity in price setting. In the remainder of the paper, we will therefore mainly investigate various properties of the lowest price.

3 Stylized Facts Regarding Fluctuations in the Lowest Price

3.1 Intensive versus extensive margins

In this section, we look at various time-series properties of the lowest price as a first step to examining the pricing behavior of retailers who aim for the first rank. The first issue we will investigate is whether and how the presence of strategic complementarity in price setting affects the frequency of changes in the lowest price and the size of price adjustment at the time of changes in the lowest price. In Kimball's (1995) setting, where price adjustment events are assumed to occur according to a Poisson process, the presence of strategic complementarity can affect pricing behavior, if at all, only through a change in the size of price adjustments. In his model, firms that are allowed to adjust their price imitate those who are not allowed to do so. Therefore, the adjusters choose to conduct a smaller price adjustment, and this leads to higher

⁷To examine why retailers target different ranks is beyond the scope of this paper. However, according to interviews with retailers conducted by the authors, the fact that retailers' target rank differs is related to the inventory turnover ratio. For example, for retailers whose scale of operation is small, the amount of capital, including borrowing from banks, is typically also small, and they need to raise the inventory turnover ratio. For this reason they aim at the lowest price. On the other hand, retailers that have relatively abundant capital and that do not need to turnover inventory that quickly choose to set a relatively high price and sell the product at a more measured pace over a longer period of time.

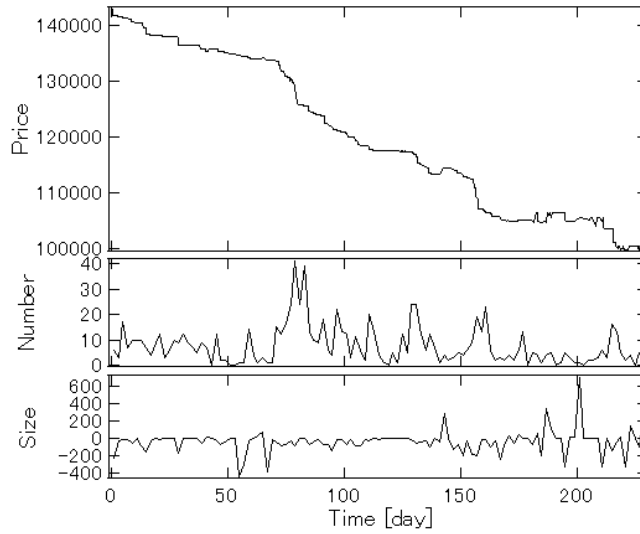


Figure 5: Fluctuations in the lowest price

price stickiness at the macro level. On the other hand, Ball and Romer (1990) investigate the role of strategic complementarity in an economy with menu costs, in which not only the size but also the timing of price adjustments is endogenously determined. They show that, provided that other firms do not adjust their prices, it would be optimal for a firm not to adjust its price even when marginal costs change, as long as this change in marginal costs is sufficiently small. This implies a positive correlation between (the absence of) price adjustments by one firm and (the absence of) price adjustments by a rival firm.

The upper panel of Figure 5 shows fluctuations in the lowest price of the AQUOS LC-32GH2 liquid crystal TV over a sample period of 230 days starting from November 1, 2006. The lowest price was updated 880 times during this sample period, implying that it was updated about 3.8 times a day. As seen in the figure, the lowest price was on a downward trend throughout the sample period, but a closer inspection of the figure reveals that there were three phases during which the decline in the lowest price accelerated: between the 70th day and the 80th day; between the 120th day and the 130th day; and between the 150th day and the 160th day. The rapid decline in the lowest price during these three phases was probably associated with a substantial decline in marginal costs, such as a change in procurement prices.

An important question concerns whether the rapid decline in the lowest price came from increases in the frequency of changes in the lowest price (the extensive margin) or changes in the size of downward price adjustments (the intensive margin). The middle and lower panels of the figure show the two margins, respectively. Specifically, the middle panel presents the number

Table 1: Correlations between the price changes and the intensive and extensive margins

Price changes	Correlations between the price change and:	
	Extensive margin	Intensive margin
Above 1,000 yen	0.64	0.08
500-1,000 yen	0.39	-0.01
100-500 yen	0.19	0.34
1-100 yen	0.20	0.43

Note: “Price” here means “lowest price.” The correlation between price changes and the extensive margin is calculated as the coefficient of correlation between price changes during two-day intervals and the frequency of price changes during those intervals when those price changes fall within the range indicated in the left-hand column. Similarly, the correlation between price changes and the intensive margin is calculated as the coefficient of correlation between price changes during two-day intervals and the average size of all adjustments during those intervals when those price changes fall within the range indicated in the left-hand column.

of adjustments in the lowest price for each two-day period, while the lower panel shows the average size of adjustments in the lowest price for each two-day period. As can be clearly seen in the middle panel, the number of adjustments increased significantly during the three phases. For example, the number of adjustments reached as high as 40 during the first phase, which is more than four times the regular level. In contrast, the lower panel shows no clear sign of an increase in the average size of adjustments during the three phases.

To investigate the contribution of the extensive and the intensive margin in more detail, we calculate the coefficient of correlation between the extensive margin and the price change for each two-day period (i.e., the difference between the lowest price at the first second of a two-day period and at the final second of the period), as well as the coefficient of correlation between the intensive margin and the price change for each two-day period. The result is presented in Table 1, showing that the extensive margin is highly correlated with the price change in periods when the price decline exceeded 1,000 yen. In contrast, there is no significant correlation with the intensive margin in these periods. These results are consistent with what we saw in Figure 5. The same can be confirmed for periods when the price declines were in the range of 500-1,000 yen. On the other hand, in periods when the price decline was smaller (i.e., below 500 yen), the coefficient of correlation for the intensive margin is now significant and even higher than the one for the extensive margin, implying that the intensive margin contributes more to price changes than the extensive margins during those quiet periods without large price fluctuations.

The above results indicate that the frequency of price adjustments changes over time depending on the environment surrounding retailers, thus clearly rejecting the idea that price

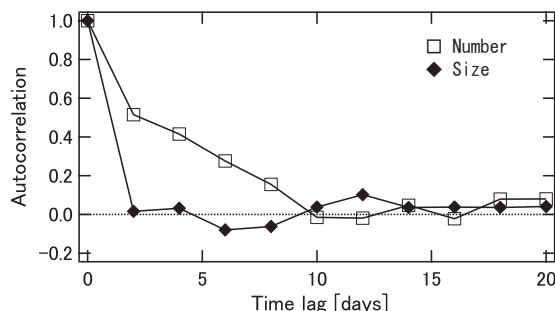


Figure 6: Autocorrelation functions of the intensive and the extensive margin

adjustment events occur according to a Poisson process. The results also indicate that the extensive margin plays a more important role than the intensive margin on days with large price changes, suggesting that large price changes result from fierce competition among retailers, with first one retailer and then the other taking the lead.

3.2 History dependence in intensive and extensive margins

When the lowest price is updated, at least one retailer (and sometimes several retailers) changes its price to a *new* level. This new price should be uncorrelated with past values of the lowest price, if there is neither strategic complementarity in price setting nor persistence in marginal costs. In this case, the size of adjustment of the lowest price, or the intensive margin, should be uncorrelated with its past values. In other words, the intensive margin should be unpredictable. Similarly, the extensive margin should be unpredictable as well. Put differently, the intensive and extensive margins can be correlated with their past values only when there is strategic complementarity in price setting or persistence in marginal costs.

Figure 6 shows the estimated autocorrelations for both the intensive and the extensive margin. The extensive margin has a high positive correlation of above 0.5 with its value two days before, and this autocorrelation gradually decays until it reaches zero vis-à-vis the extensive margin ten days before. In contrast, the intensive margin has no significant correlation with its past values. To understand why we arrive at such different results for the intensive and extensive margins, suppose that the positive autocorrelation for the extensive margin comes not from strategic complementarity but from persistence in marginal costs. In this case, each retailer would set its price equal to marginal costs when it decides to adjust its price. Then the intensive margin should be correlated with its past values through marginal costs, which is clearly inconsistent with what we observed in the data. An alternative possibility is that there exists strategic complementarity but there is no strong persistence in marginal costs.

Here one may wonder why strategic complementarity creates a positive autocorrelation for the extensive margin, but not for the intensive margin. One possible interpretation is that each retailer imitates its rivals in the sense that it adjusts its price immediately after the rivals adjust their prices, but it does not go so far as to mimic the size of rivals' price adjustments. For example, suppose that a retailer has decided to adjust its price following a rival's price cut by 1,000 yen. However, when the retailer proceeds to decide the size of the adjustment, it is no longer influenced by the rival and instead chooses a price adjustment, say by 500 yen, taking its own unique circumstances into consideration.

The fact that the extensive margin has a positive autocorrelation up to ten days has important implications. To illustrate this, suppose that marginal costs fall from m to m' ($m' < m$) and stay there permanently. Think about a pure Calvo-type setting in which a Calvo event occurs at the probability of 0.75 per day (this is taken from Section 2.2). In this case, the first Calvo event arises for some retailer 20 minutes after the fall in marginal costs,⁸ so that the lowest price is updated by $m' - m$ at that time, and stays there permanently. Importantly, the change in the lowest price occurs only once. But this does not mean that all retailers update their prices to m' in 20 minutes; in fact, it takes 1.33 days ($= 1/0.75$) before every retailer completes a downward price adjustment. From this simple calculation, we learn the following. First, the data tell us that it takes ten days before the adjustment in the lowest price is complete, which is much longer than 20 minutes. This implies that the downward adjustment of the lowest price by $m' - m$ is partitioned into small pieces. In fact, the size of changes in the lowest price that was most frequently observed during the sample period of 230 days was 10 yen, accounting for 31.5 percent of all price adjustments, followed by 1 yen with a share of 24.7 percent in the total. On the other hand, the share of larger size adjustments, like 100 yen or 1,000 yen, is less than one percent. This is consistent with the presence of strategic complementarity in price setting. The second implication is related to the time required to complete the adjustment for all retailers. The estimate of 1.33 days is clearly wrong, because, given that, even in the case of the lowest price, it takes ten days for the adjustment to be completed, it is highly likely that it takes far more than ten days for the adjustment to be completed for all retailers. This simply means that it is quite misleading to estimate the adjustment speed from the observed frequency of price adjustments.

To investigate the autocorrelation of the extensive margin in more detail, we estimate something similar to the hazard function. Suppose that retailer i changes its price at a particular point in time, and that the total number of price adjustments by other retailers, including price adjustments that do not update the lowest price, after that point in time is n . We then calculate

⁸The number of retailers is assumed to be 100.

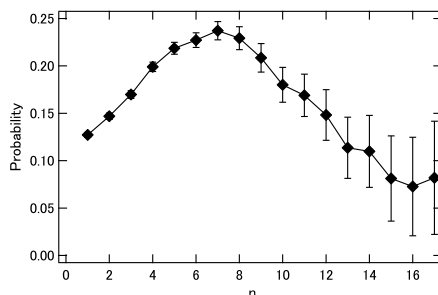


Figure 7: Probability of price adjustment conditional on the number of adjustments by rivals

the probability of a price adjustment by retailer i conditional on the occurrence of the n price adjustments by the other retailers. Without strategic complementarity in price setting, retailer i would not be affected by the n price adjustments implemented by the other retailers, so that the probability of a price adjustment by retailer i does not depend on n at all. However, with strategic complementarity, a larger n means that the price offered by retailer i is obsolete, and is far away from the current lowest price. Therefore, retailer i has a stronger incentive to regain the first rank by cutting its price, so that the conditional probability increases with n . Note that the hazard function defined in this way has a different shape depending on whether strategic complementarity is or is not present.

Figure 7 shows the estimation result. The horizontal axis represents the value of n ; for example, $n = 4$ means that four price adjustments by other retailers occur after the last price adjustment by retailer i , and the corresponding value on the vertical axis represents the conditional probability of a price adjustment by retailer i . This figure resembles those frequently used in previous studies, such as Alvarez et al. (2005) among others, except that the horizontal axis represents not the elapsed time but the number of price adjustments by other retailers. As often pointed out in those studies, when estimating hazard functions, it is important to use data for a homogenous set of retailers. In particular, we need to remove the heterogeneity among retailers in terms of their target rank that we saw in Section 2.3. To do so, we focus only on retailers that are positioned at the 5th rank or above immediately after the last price adjustment. As seen in the figure, the conditional probability of price adjustment increases with n at least until $n = 7$. This result clearly rejects the null hypothesis of no strategic complementarity in price setting. Instead, the result indicates that retailer i 's decision is significantly influenced by the pricing behavior of other retailers. Note that the conditional probability starts to decline from $n = 8$ onward, which looks quite similar to the downward-sloping hazard functions repeatedly reported in previous studies on price stickiness.

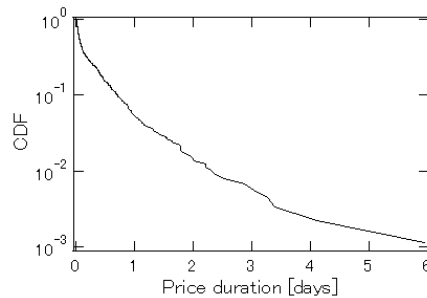


Figure 8: Distribution of price durations in the lowest price

The declining conditional probability could be interpreted as reflecting the fact that there still remains some heterogeneity among retailers in terms of the probability of price adjustments: prices offered by retailers whose probability of price adjustment is low tend to survive longer, thus generating the downward-sloping part of the estimated hazard function.

Finally, the fact that the extensive margin shows a positive autocorrelation implies that the lengths of price spells, that is, the length of the period between two consecutive changes in the lowest price, are correlated with each other. To see this, we show the distribution of price durations in Figure 8, where the horizontal axis represents price duration while the vertical axis represents the corresponding value of the cumulative distribution function. As we explained before, if the change in the lowest price occurs according to a Poisson process, price duration should follow an exponential distribution and the CDF therefore should form a straight line in a log plot. However, the CDF takes a convex shape, implying that the longer it is since the lowest price has been adjusted, the lower tends to be the probability that the next price adjustment occurs. Put differently, there tends to be a clustering, so that once an adjustment of the lowest price occurs, such adjustments occur in succession. On the other hand, when the lowest price is not adjusted for a long time, the probability of a price adjustment event is low.

4 State-Dependent Pricing with Strategic Complementarities

In this section, we will introduce the generalized Ss model proposed by Caballero and Engel (2007) to describe state-dependent pricing with strategic complementarities, and conduct numerical simulations with calibrated parameters to see whether we can replicate the facts found in the last section.

4.1 Caballero and Engel’s generalized Ss model

Consider a setting in which each retailer reviews its price and, if necessary, adjusts it. Suppose that the opportunity of a price review arrives according to a Poisson process with a probability of $1 - \theta$. When this opportunity arrives for a retailer, it compares its current price with what Caballero and Engel (2007) refer to as the target price. The retailer changes its price if the discrepancy between the two is sufficiently large, and sets a new price equal to the target price.

A key assumption in this setting concerns how the target price is determined. Here we assume that there are two types of retailers. The first type of retailers pays no attention to the prices offered by rival retailers. The target price of these retailers is assumed to be equal to marginal costs, which is denoted by m_t . On the other hand, the second type of retailers considers the prices offered by rivals in deciding their target prices. Specifically, their target price is equal to a weighted average of the average value of the prices offered by those who changed their prices in the previous period (“adjusters”) and the average value of the prices offered by those who did not change their prices in the previous period (“non-adjusters”), which are denoted by P_{t-1}^A and P_{t-1}^{NA} , respectively. The fraction of the first type in the total number of adjusters in each period is assumed to be constant and given by $1 - \alpha$. Finally, we assume that a typical retailer is a combination of the two types with the weights $1 - \alpha$ and α . Then the target price of this retailer, which is denoted by P_{it}^* , is given by

$$P_{it}^* = (1 - \alpha)m_t + \alpha [\omega P_{t-1}^A + (1 - \omega)P_{t-1}^{NA}] \quad (1)$$

where ω is a parameter between 0 and 1, and P_{it}^* , m_t , P_{t-1}^A , and P_{t-1}^{NA} are all in logarithm.

The probability that retailer i changes its price conditional on that it is allowed to review its price is denoted by Λ and assumed to depend on x_{it} , which is defined by $x_{it} \equiv P_{it-1} - P_{it}^*$. The function $\Lambda(x_{it})$ is what Caballero and Engel (1993a) refer to as the “adjustment hazard function.” This is a useful tool to discriminate between state-dependent and time-dependent pricing. If the probability of price adjustment depends upon a state variable, x , this indicates state-dependent pricing, and if not, this indicates time-dependent pricing. We make two assumptions about the shape of the adjustment hazard function. First, the probability of price adjustment becomes higher as the actual price deviates more, positively or negatively, from the target level, so that $\Lambda'(x) > 0$ for $x > 0$ and $\Lambda'(x) < 0$ for $x < 0$. Caballero and Engel (1993b) call this the increasing hazard property. Second, the adjustment hazard function is assumed to be symmetric.

Given the above setting, we are able to see how the aggregated price level responds to shocks in marginal costs. Denoting the response of the price of retailer i in period t to a marginal cost

shock in period t by $\Delta P_{it}(\Delta m_t, x)$ and its aggregated counterpart by $\Delta P_t(\Delta m_t)$, we have

$$\begin{aligned}\Delta P_t(\Delta m_t) &= \int \Delta P_{it}(\Delta m_t, x) h(x) dx \\ &= (1 - \theta) \int [x + (1 - \alpha)\Delta m_t] \Lambda [x + (1 - \alpha)\Delta m_t] h(x) dx\end{aligned}\quad (2)$$

where $h(x)$ is the cross-section distribution (ergodic distribution) of the state variable x . Differentiating this equation with respect to Δm_t and evaluating at $\Delta m_t = 0$ yields

$$\lim_{\Delta m_t \rightarrow 0} \frac{\Delta P_t}{\Delta m_t} = (1 - \alpha)(\mathcal{A} + \mathcal{E}) \quad (3)$$

where \mathcal{A} and \mathcal{E} are defined by

$$\mathcal{A} \equiv (1 - \theta) \int \Lambda(x) h(x) dx; \quad \mathcal{E} \equiv (1 - \theta) \int x \Lambda'(x) h(x) dx \quad (4)$$

The expression on the left-hand side of (3) is Caballero and Engel's (2007) measure of price flexibility, which is basically the impulse response function. Their measure of price flexibility is determined by three factors. The first is the term $\int \Lambda(x) h(x) dx$ in (4), which represents the frequency of price adjustments conditional on the arrival of price review event. As is easily seen in (3), a higher frequency of price adjustments leads to greater price flexibility in terms of the impulse response function. The second factor is the term \mathcal{E} in (3), which represents the extent of state dependence and takes a positive value under the assumption of the increasing hazard property. The third factor is α , which represents the fraction of retailers that imitate rival retailers. With a larger value of α , fewer retailers pay attention to marginal costs and therefore the aggregated price level becomes less sensitive to changes in marginal costs.

Next, we would like to examine some of the implications for history dependence in the average of adjusters' prices. First, using equation (1), we can calculate the autocorrelations in the size of adjustments of the average of adjusters' price under the assumption that m_t follows a random walk and that $\omega = 1$. If retailer i is allowed to review its price in period t and decides to change it, it will set the new price at the target level given by (1), so that we have $P_t^A = P_{it}^*$. Substituting this into (1), we have

$$P_t^A = (1 - \alpha) \sum_{k=0}^{\infty} \alpha^k m_{t-k},$$

indicating that P_t^A depends on the past values of marginal costs unless α is equal to zero, although P_t^A is the average of *new* prices set in period t . This history dependence stems from the backward-looking behavior of the second type of retailers who imitate the average of the prices offered by adjusters in the previous period. The above equation implies that the coefficient of correlation between $P_t^A - P_{t-1}^A$ and $P_{t-j}^A - P_{t-j-1}^A$ is given by α^j , implying

that autocorrelations in the size of adjustments in the average of adjusters' prices is entirely determined by the parameter α in this special case.

Second, we calculate how a decline in P_{t-1}^A affects P_t^A . Denoting a change in P_{t-1}^A by ΔP_{t-1}^A , and the resulting change in P_t^A by ΔP_t^A , we have

$$\Delta P_t^A(\Delta P_{t-1}^A) = \frac{\int (x + \alpha\omega\Delta P_{t-1}^A) \Lambda(x + \alpha\omega\Delta P_{t-1}^A) h(x) dx}{\int \Lambda(z + \alpha\omega\Delta P_{t-1}^A) h(z) dz} \quad (5)$$

As before, we differentiate this with respect to ΔP_{t-1}^A and evaluate at $\Delta P_{t-1}^A = 0$ to get

$$\lim_{\Delta P_{t-1}^A \rightarrow 0} \frac{\Delta P_t^A}{\Delta P_{t-1}^A} = \alpha\omega \left(1 + \frac{\mathcal{E}}{\mathcal{A}} \right). \quad (6)$$

In the case of time dependent pricing, $\Lambda'(x) = 0$ and thus $\mathcal{E} = 0$, so that the right-hand side of (6) is equal to $\alpha\omega$. In this case, a decline in P_{t-1}^A leads to a decline in P_t^A only through a decrease in the size of price adjustments (i.e., the intensive margin effect). However, in the general case where $\Lambda'(x)$ is not zero, a decline in P_{t-1}^A leads to an increase in x for each retailer, thereby increasing the adjustment hazard for retailers that were with $x > 0$ before the shock to P_{t-1}^A occurs (and that therefore sought to lower their price) and decreasing it for retailers that were with $x < 0$ before the shock occurs (and that therefore wanted to raise their price). Put differently, more retailers decrease prices and fewer retailers increase prices, thereby leading to a negative response of P_t^A . Note that the negative impact on P_t^A is reinforced by the presence of this extensive margin effect. More importantly, the decline in P_t^A , in turn, leads to a decline in P_{t+1}^A through the same intensive and extensive margin effects. This is how the presence of strategic complementarity creates positive autocorrelations both in the size and the frequency of adjustments in the adjusters' price.

4.2 Numerical simulations

In this subsection we simulate the model with calibrated parameters, which are set as follows. We assume that the number of retailers is ten, and that the length of each time step is 3.6 minutes, or 40 steps=1 day. As for the probability of price review, we assume that each retailer reviews its price, on average, 0.4 times a day, which implies $1 - \theta = 0.01$. This is based on the interviews we conducted with major players in the Kakaku.com market.

As for marginal costs m_t , we assume that the decline in marginal costs occurs according to a Poisson process, and that the probability of the decline is 0.002 per time step, implying that a decline in marginal costs occurs, on average, 30 times a year. The size of the decline in marginal costs is assumed to be 1 percent of the price level. Combining these two, marginal costs m_t and therefore the price level decline, on average, by 30 percent per year, which is consistent with what is observed in the data.

Following Caballero and Engel (2006), the adjustment hazard function is assumed to be of quadratic form, and can be represented by:

$$\Lambda(x) = \begin{cases} 1 & \text{if } x < -\lambda_0 \\ (x/\lambda_0)^2 & \text{if } -\lambda_0 \leq x \leq \lambda_0 \\ 1 & \text{if } \lambda_0 < x \end{cases} \quad (7)$$

where λ_0 is a positive parameter. The value of λ_0 is assumed to be 0.005, implying that the probability of adjustment is equal to unity if the deviation of the current price from the target level is above 5 percent, while it is less than unity if the deviation is less than 5 percent. Finally, the parameters appearing in equation (1) are assumed to be $\alpha = 0.9$ and $\omega = 0$. The value of α is based on the empirical autocorrelation in the frequency of changes in the lowest price.⁹

Before discussing the simulation results, let us describe how the model will behave using a simple example. Suppose that marginal costs decline from m to m' ($m' < m$), and that this decline is sufficiently large. For a retailer that is given the opportunity to review its price immediately after this event, the target price is given by $P^* = (1 - \alpha)m' + \alpha m$. For the second retailer that is given the opportunity to review its price, the target price is $P^* = (1 - \alpha^2)m' + \alpha^2 m$. Similarly, for the k -th retailer that is allowed to review its price, the target price is given by $P^* = (1 - \alpha^k)m' + \alpha^k m$. Since the target price coincides with P^A , P^A for the k -th retailer is $P^A = (1 - \alpha^k)m' + \alpha^k m$. As can be seen from this equation, P^A monotonically declines with k , implying that the price adjustment by a retailer at a particular point of time always results in updating the lowest price at that time.¹⁰

The simulation results are presented in Figure 9. The upper panel shows part of the time-series of the lowest price, as well as the frequency of adjustments in the lowest price and the size of adjustments in the lowest price. The variables in the figure are defined in exactly the same way as those in Figure 5. The figure shows that the lowest price is on a downward trend, but sometimes exhibits a rapid decline, which occurs immediately after a decline in marginal costs. Decomposing the decline in the lowest price into the frequency and the size of adjustments, we see that both the frequency and the size of adjustments significantly increase at times when the lowest price declines rapidly. What is going on here is that the decline in marginal costs leads to a decline in the target price, as shown in equation (1), which, in turn, decreases the target price for other retailers. The intensive and the extensive margin increase during such a chain reaction of price cuts.

Comparing this simulation result with the empirical result shown in Figure 5, we see that the model successfully replicates the tendency observed in the data that the extensive margin

⁹To test the robustness of our results, we also conducted simulations with slightly smaller values of α ($\alpha = 0.8$ and 0.7). The results are similar to the one obtained in the baseline case.

¹⁰A sketch of the model that gives rise to such ranking competition is provided in the Appendix.

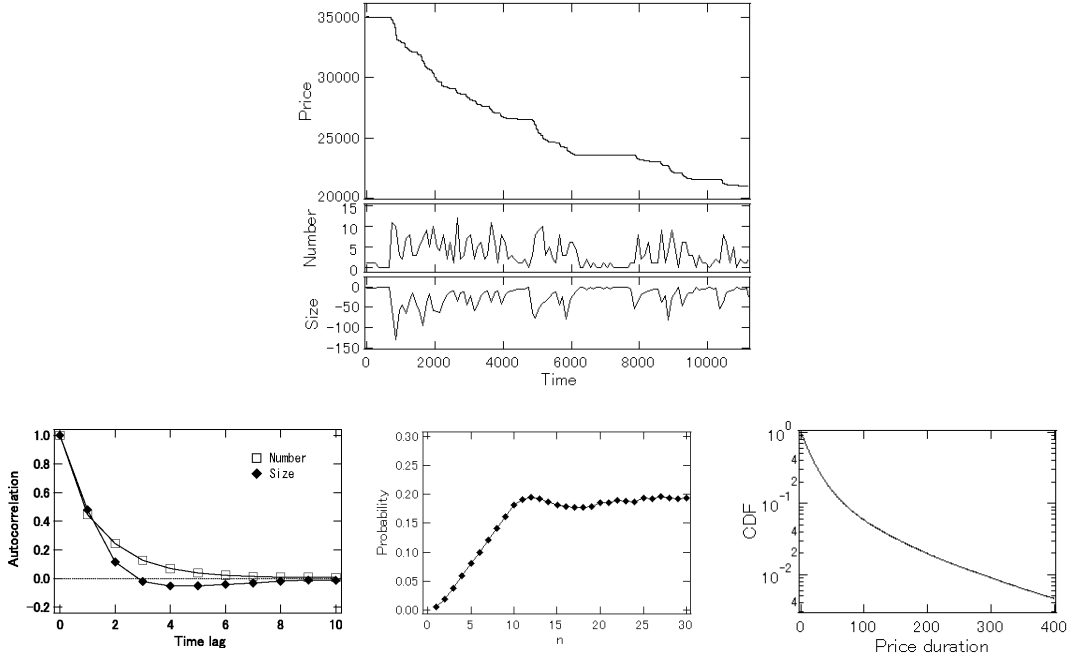


Figure 9: Simulation results

increases when the lowest price declines rapidly. However, the simulation result suggests that the intensive margin increases as well when the lowest price declines rapidly, which is clearly inconsistent with the findings from the data. In Caballero and Engel’s generalized S_s model, the state variable x determines the size of adjustments, but it also determines the likelihood of adjustments through the adjustment hazard function $\Lambda(x)$. This feature is clearly seen in the close synchronization of the intensive and the extensive margin in the simulation result. The empirical result in Figure 5 suggests that the size of adjustments is determined in a way different from the one described by the Caballero-Engel model.

Next, autocorrelations for the size and the frequency of changes in the lowest price are shown in the lower left panel of Figure 9. The result clearly shows that there is a positive autocorrelation both in the intensive and the extensive margin, although the autocorrelation is slightly larger and longer for the extensive margin than for the intensive margin. Note that such persistent movements in the intensive and the extensive margin stem from strategic interactions among retailers due to the presence of strategic complementarity described by equation (1). Comparing this result with the empirical result presented in Figure 6, we see an important difference for the intensive margin in that the empirical autocorrelations are almost zero, suggesting again that the size of adjustments is not determined as described in the Caballero-Engel model.

The lower middle panel shows the probability of a price adjustment by retailer i conditional on the occurrence of the n price adjustments by other retailers. Comparing this with the empirical result presented in Figure 7, the model successfully replicates the pattern that the conditional probability increases with n , although the simulation result does not show the property that the hazard function is at least partially downward sloping, as we saw in the data. Finally, the lower right panel shows the CDF of the price duration in terms of the lowest price (i.e., the length of the period between two consecutive changes in the lowest price). The CDF does not form a straight line but takes a convex shape, as we observed in Figure 8.

5 Conclusion

Are prices sticky due to the presence of strategic complementarity in price setting? If so, to what extent? To address these questions, we investigated retailers' price setting behavior using a unique dataset containing by-the-second records of prices offered by retailers on a major Japanese price comparison website. We focused on fluctuations in the lowest price among retailers, rather than the average price, examining how quickly (or how slowly) the lowest price is updated in response to changes in marginal costs.

We found that when the lowest price falls rapidly, the *frequency* of changes in the lowest price is high, while the *size* of downward price adjustments remains largely unchanged. We also found a positive autocorrelation in the frequency of changes in the lowest price, and that there tends to be a clustering where once a change in the lowest price occurs, such changes occur in succession. In contrast, there is no such autocorrelation in the size of changes in the lowest price. These findings suggest that retailers imitate each other when deciding to adjust (or not to adjust) their prices, and that the extensive margin plays a much more important role than the intensive margin in such strategic complementarity in price setting.

We conducted simulation analysis with calibrated parameters to see whether Caballero and Engel's generalized Ss model can replicate these facts. We successfully reproduced various features of the extensive margin, including its persistence, while we failed to reproduce some features of the intensive margin, including the absence of persistence.

References

- [1] Alvarez, L. J., P. Burriel, and I. Hernando (2005). "Do Decreasing Hazard Functions for Price Changes Make Sense?," *Working Paper* 461, European Central Bank.
- [2] Ball, L., and D. Romer (1990), "Real Rigidities and the Non-Neutrality of Money," *Review of Economic Studies* 57(2), 183-204.

- [3] Basu, S. (2005), “Comment on Implications of State-Dependent Pricing for Dynamic Macroeconomic Models,” *Journal of Monetary Economics* 52(1), 243-247.
- [4] Baye, M. R., J. R. J. Gatti, P. Kattuman, and J. Morgan (2009), “Clicks, Discontinuities, and Firm Demand Online,” May 2009. <http://www.nash-equilibrium.com/baye/Shoppers.pdf>.
- [5] Bills, M., and P. J. Klenow (2004), “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy* 112(5), 947-985.
- [6] Caballero, R. J., and E. M.R.A. Engel (1993a), “Microeconomic Rigidities and Aggregate Price Dynamics,” *European Economic Review* 37(4), 697-717.
- [7] Caballero, R. J., and E. M.R.A. Engel (1993b), “Microeconomic Adjustment Hazards and Aggregate Dynamics,” *Quarterly Journal of Economics* 108(2), 359-383.
- [8] Caballero, R. J., and E. M.R.A. Engel (2006), “Price Stickiness in Ss Models: Basic Properties,” MIT and Yale.
- [9] Caballero, R. J., and E. M.R.A. Engel (2007), “Price Stickiness in Ss Models: New Interpretations of Old Results,” *Journal of Monetary Economics* 54(Supplement 1), 100-121.
- [10] Calvo, G. (1983), “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics* 12(3), 383-398.
- [11] Dhyne, E., L. J. Alvarez, H. L. Bihan, G. Veronese, D. Dias, J. Hoffmann, N. Jonker, P. Lunnemann, F. Rumler, and J. Vilmunen (2006), “Price Setting in the Euro Area and the United States: Some Facts From Individual Consumer Price Data,” *Journal of Economic Perspectives* 20(2), 171-192.
- [12] Dotsey, M., and R. G. King (2005), “Implications of State-Dependent Pricing for Dynamic Macroeconomic Models,” *Journal of Monetary Economics* 52(1), 213-242.
- [13] Ellison, G., and S. F. Ellison (2009), “Search, Obfuscation, and Price Elasticities on the Internet,” *Econometrica*, forthcoming.
- [14] Golosov, M., and R. E. Lucas (2007). “Menu Costs and Phillips Curves,” *Journal of Political Economy* 115(2), 171-199.
- [15] Higo, M., and Y. Saita (2007), “Price Setting in Japan: Evidence from CPI Micro Data,” *Bank of Japan Working Paper Series 07-E-20*, Bank of Japan.
- [16] Kimball, M. (1995), “The Quantitative Analytics of the Basic Neomonetarist Model,” *Journal of Money, Credit and Banking* 27(4), 1241-1277.

- [17] Klenow, P. J., and O. Kryvtsov (2008), “State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?” *Quarterly Journal of Economics* 123(3), 863-904.
- [18] Nakamura, E., and J. Steinsson (2008), “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics* 123(4), 1415-1464.
- [19] Negishi, T. (1979), *Microeconomic Foundations of Keynesian Macroeconomics*, New York: North-Holland.
- [20] Saito, Y., and T. Watanabe (2009), “Menu Costs and Price Change Distributions: Evidence from Japanese Scanner Data,” Research Center for Price Dynamics Working Paper Series No.17.
- [21] Woodford, M. (2009), “Information-Constrained State-Dependent Pricing,” Columbia University, June 25, 2009.

A A model of ranking competition

Consider N firms $i = 1, 2, \dots, N$ that compete in ranking. Let $r_{i,t}$ denote the rank of firm i in period t . We consider the case that firms compete in a narrow range around a competitive price and obtain a vanishingly small fixed margin for a unit of the good sold at the price. Thus, a profit maximizing firm aims to maximize the number of units it sells. Firms' posted prices are ranked and shown to customers. Firms can take the first ranking position by infinitesimally cutting the posted price. The demand faced by each firm is determined by its ranking. We assume that demand is given by a linear function, $a - br_{i,t}$. Thus, firm i always takes the first position by choosing $r_{i,t} = 1$ whenever it takes action. Let β denote the discount factor.

We assume that there are two kinds of friction in the firm's action. First, as in Calvo pricing, opportunities for firm i to act arrive only stochastically with Poisson rate λ . Such friction occurs when each firm has only a limited number of staff who monitor the market. Second, firm i incurs fixed cost c_i when changing $x_{i,t}$. This fixed cost c_i corresponds to the amount of time staff spend posting the new price.

The rank $r_{i,t}$ changes to first place whenever i takes action. The rank $r_{i,t}$ shifts to $r_{i,t} + 1$ whenever another firm takes action and takes over first place. We first consider the case where new entrants arrive with Poisson rate μ and take the first position. Let $v(r)$ denote the value function of a firm at rank r right after the moment when the rank adjustment event has arrived. Let $W(r)$ denote the value function at a moment when the firm cannot adjust. $V(r)$ denotes the conditional expected value of the firm in the next period when its ranking in this period is r . The expectation is taken over two uncertainties: the arrival of an adjustment event and the arrival of an entrant. The Bellman equation for the firm's dynamic optimization problem is written as follows:

$$v(r) = \max_{\{\text{Adjust, Inact}\}} \{A, W(r)\} \quad (\text{A.1})$$

$$A = a - b - c_i + \beta V(1) \quad (\text{A.2})$$

$$W(r) = a - br + \beta V(r) \quad (\text{A.3})$$

$$V(r) = (1 - \mu)((1 - \lambda)W(r) + \lambda v(r)) + \mu((1 - \lambda)W(r + 1) + \lambda v(r + 1)) \quad (\text{A.4})$$

Within the interior of the inaction region (i.e., within the region of r such that both r and $r + 1$ are in the inaction region), a linear value function solves the Bellman equation as follows:

$$v(r) = W(r) \quad (\text{A.5})$$

$$W(r) = \frac{a - br - b\mu\beta/(1 - \beta)}{1 - \beta} \quad (\text{A.6})$$

$$V(r) = \frac{a - br - b\mu/(1 - \beta)}{1 - \beta} \quad (\text{A.7})$$

Then,

$$A = \frac{a - b - b\mu\beta/(1 - \beta)}{1 - \beta} - c_i \quad (\text{A.8})$$

Then, the inaction region is written as $1 \leq r < r^*$ where the adjustment point r^* is determined at the point where the two payoffs in (A.1) are equalized as $A = W(r^*)$. Note that $W(r)$ for the adjustment region $r \geq r^*$ must satisfy:

$$W(r) = a - br + \beta((1 - \mu)((1 - \lambda)W(r) + \lambda A) + \mu((1 - \lambda)W(r + 1) + \lambda A)) \quad (\text{A.9})$$

Solving this forward, we obtain

$$W(r) = \frac{1 + \beta\lambda A - rb}{1 - \beta(1 - \lambda)} - \frac{b\beta\mu(1 - \lambda)}{(1 - \beta(1 - \lambda))^2} \quad (\text{A.10})$$

for the adjustment region. Using this to evaluate $A = W(r^*)$, we obtain:

$$r_i^* = 1 + \frac{(1 - \beta)c_i}{b} + \frac{\mu\beta}{(\beta + (1 - \beta)/\lambda)(1 - \beta)} \quad (\text{A.11})$$

Hence, the optimal policy for firm i is to take the first rank when its current rank is below the minimum level r_i^* , and not to adjust its rank when $r_{i,t} < r_i^*$. The optimal policy has an (S,s) pattern where $r_{i,t}$ increases during inaction and is reset to 1 when $r_{i,t}$ reaches r_i^* . The region $1 \leq r_{i,t} < r_i^*$ is called the inaction region, and the region $r_{i,t} \geq r_i^*$ is called the adjustment region.

The threshold is easily interpreted. r_i^* is increasing in c_i , i.e., a firm allows its rank to decline further if the adjustment cost is high. r_i^* is decreasing in b because the marginal benefit of a high ranking is greater when b is larger. r_i^* is increasing in μ because the firm's ranking declines faster when μ is high. Finally, for a sufficiently small λ , r_i^* is decreasing in λ . This is because a high rate of adjustment events allows firms to check their ranking frequently and thus allows them to wait longer.

So far we are concerned with a single firm's behavior. Now we turn to characterizing the equilibrium at which μ is consistent with the behavior of all firms. We can characterize two polar cases at which μ is independent of the state of the agents: when all firms are content with the current ranking and thus $\mu = 0$, and when all firms are located in the adjustment region and thus take action when the Calvo event hits.

First we consider the case $\mu = 0$. This can happen when c_i is distributed in such a way that the vector (r_i^*) coincides with vector $(1, 2, \dots, N)$ after a proper rearrangement of the index i . This is the case when the industry has an "order" in which each player's rank corresponds to its costs. In this case, no adjustment occurs once $r_{i,t} = r_i^*$ is achieved for all i under $\mu = 0$, and $\mu = 0$ is consistent with the fact that no firm will adjust in future.

We can extend this case to where a new but transient entrant takes the first rank with very small probability θ . Then, the ranking stays in the proper order in the previous sense (with the policy function derived for the case $\mu = \theta$) for an extended period of time until the new entrant disturbs the order.

Next, we consider the case where there is no new entrant. The rank $r_{i,t}$ is shifted to $r_{i,t} + 1$ only if another firm at a lower rank moves and takes the first position. The actions of firms at

a higher rank than i do not affect i 's rank. Then, the probability for $r_{i,t}$ to be shifted to $r_{i,t} + 1$ is the sum of probabilities for each of $N - r_{i,t}$ firms to take action and take the first position.

Consider a polar case where almost all firms are located in the adjustment region. For example, starting from the stationary case considered above, suppose that the last ranking firm takes the first position due to some exogenous factors such as the need to liquidate its assets. Then all other firms move down one rank ($r_i^* + 1$) and want to adjust their rank whenever an adjustment event arrives. In this case, firm i rationally expects that i moves down one rank with probability $\mu = \lambda(N - r_{i,t})$. Note that in this case the probability μ depends on the firm's rank $r_{i,t}$. Nonetheless, we can still show that a linear function solves the Bellman equation within the interior of the inaction region:

$$v(r) = W(r) = \frac{1}{1 - \beta} \left(a - \frac{\beta\lambda Nb}{1 - \beta + \beta\lambda} \right) - \frac{br}{1 - \beta + \beta\lambda} \quad (\text{A.12})$$

The slope of v is flatter than in the previous case. This is because lower-ranking firms experience a slower erosion of their position under the current specification of μ .

Under the current setup, we can obtain the threshold r_i^* . We consider the case where the liquidation of a firm disrupts the ranking order and all the firms except the liquidated firm are located at $r_{i,t} = r_i^* + 1$. Then, each firm adjusts its ranking whenever an adjustment event arises, and the adjustment events arise with the rate λ . This is consistent with the presumption that $\mu = \lambda(N - r_{i,t} - 1)$.