House Price Dispersion in Boom-Bust Cycles: Evidence from Tokyo

Takaaki Ohnishi
Takayuki Mizuno
Tsutomu Watanabe

House Price Dispersion in Boom-Bust Cycles: Evidence from Tokyo

Takaaki Ohnishi[*]   Takayuki Mizuno[†]   Tsutomu Watanabe[‡]

June 29, 2019

Abstract

We investigate the cross-sectional distribution of house prices in the Greater Tokyo Area for the period 1986 to 2009. We find that size-adjusted house prices follow a lognormal distribution except for the period of the housing bubble and its collapse in Tokyo, for which the price distribution has a substantially heavier upper tail than that of a lognormal distribution. We also find that, during the bubble era, sharp price movements were concentrated in particular areas, and this spatial heterogeneity is the source of the fat upper tail. These findings suggest that, during a bubble, prices increase markedly for certain properties but to a much lesser extent for other properties, leading to an increase in price inequality across properties. In other words, the defining property of real estate bubbles is not the rapid price hike itself but an increase in price dispersion. We argue that the shape of cross-sectional house price distributions may contain information useful for the detection of housing bubbles.

JEL Classification Numbers: R10; R30; C46
Keywords: housing bubbles; anomaly detection; lognormal distributions; power-law tail; hedonic models; market segmentation; submarkets

1   Introduction

Property market developments are of increasing importance to practitioners and policy-makers. The financial crises of the past two decades have illustrated just how critical the health of this sector can be for achieving financial stability. For example, the recent financial crisis in the United States in its early stages reared its head in the form of the subprime loan problem. Similarly, the financial crises in Japan and Scandinavia in the 1990s were all triggered by the collapse of bubbles in the real estate market. More recently, the rapid rise in real estate prices - often supported by a strong expansion in bank lending - in a number of emerging market economies has become a concern for policymakers. Given these experiences, it is critically important to analyze the relationship between property markets, finance, and financial crisis.

One of the most urgent tasks in this respect is the development of methods to detect bubbles in property markets, which is key to making reliable forecasts about future financial crises and/or to mitigating them. Consider the case of residential properties. In real estate economics, the fundamental value of a house is determined as the present discounted value of current and future income flows resulting from renting the house to someone (see, for example, Himmelberg et al. 2005). Normally, house prices remain close to their fundamental value due to price arbitrage. However, in some cases, arbitrage forces are not present and prices deviate substantially from their fundamental value. This is what is called a housing bubble (see, for example, Martin and Ventura 2018). To detect housing bubbles defined in this way requires reliable estimates of the fundamental value of house prices, which in turn requires knowing market participants' expectations of rental prices in the coming years. However, such expectations are not observable, so that the fundamental value of house prices is quite difficult to estimate. Thus, it is next to impossible for researchers or policymakers to tell, when they observe a price hike, whether this is the result of a rise in fundamental values or something else.

In this study, we propose an alternative approach to detecting real estate bubbles. Specifically, we propose making use of information on the cross-sectional dispersion of real estate prices. It is often believed that *all* prices rise equally during a bubble period. However, this is not the case. What happens instead is that there are pronounced price rises for certain properties but not for others, so that price inequality across properties increases during a bubble. In other words, the defining characteristic of real estate bubbles

is not the rapid price hike itself but an increase in price dispersion.

Given these considerations, the present study addresses the following empirical issues. First, we examine whether the price distribution is close to a normal distribution, as is often assumed in empirical studies on house price indexes, or whether it has fatter tails than a Gaussian distribution. Second, we investigate how the shape of the price distribution is affected by house attributes, including the size and location of a house. Third, we consider how the shape of the distribution changes over time. In order to examine these three issues, we focus on the housing bubble Japan experienced in the late 1980s and its burst in the early 1990s.

This study is closely related to recent research on the cross-sectional distribution of house prices, including the studies by McMillen (2008), Van Nieuwerburgh and Weill (2010), Määttänen and Terviö (2014), Glaeser et al. (2012), Gyourko et al. (2013), and Blackwell (2018). The main interest of Gyourko et al. (2013), Van Nieuwerburgh and Weill (2010), and Määttänen and Terviö (2014) is the relationship between the house price distribution and the income distribution. For example, Määttänen and Terviö (2014) ask whether the recent increases in income inequality in the United States have had any impact on the distribution of house prices. On the other hand, McMillen (2008) focuses on the change in the house price distribution over time and asks whether the change in the price distribution comes from a change in the distribution of house characteristics such as size, location, age, and so on, or from a change in the implicit prices associated with those characteristics. The focus of our paper is closely related to the issues discussed in these studies but differs from them in some important respects. First, this study is the first attempt to specify the shape of the house price distribution, paying particular attention to the tail part of the distribution. Second, this study examines the effect of a housing bubble on the cross-sectional price distribution. While steep increases in the *mean* of house prices in various countries in recent decades have received substantial attention in the literature, the change in the shape of the cross-sectional price distribution has received much less attention.[1] In this study, we seek to fill this gap.

---

[1]Some of the recent studies, including Ferreira and Gyourko (2012, 2017), Sinai (2013), and Zhang (2018), focus on the higher moments of the house price distribution in the context of boom-bust cycles in housing markets. For example, Ferreira and Gyourko (2012, 2017) highlight the presence of substantial regional heterogeneity in housing prices in the recent US boom-bust cycle, stating that "the recent US housing boom cannot be interpreted as a single, national event, as different markets began to boom across a decade-long period from the mid-1990s to the mid-2000s, some of them multiple times" (Ferreira and Gyourko 2017, p. 1). However, these studies do not examine changes in the shape of the cross-sectional house price distribution over boom-bust cycles.

Another strand of literature closely related to the present study is research on early warning indicators of financial crises, including the studies by Kaminsky and Reinhart (1999), Borio and Lowe (2002), Drehmann and Juselius (2014), and Anundsen et al. (2016). These studies are all based on the idea that financial imbalances, or vulnerabilities, build up over time before crises arise, and that it should be possible to capture the cumulative processes that sow the seeds of the subsequent crisis by carefully monitoring key variables such as the credit to GDP ratio, the debt service ratio, and real estate prices. Specifically, these studies estimate the level of a critical threshold for the deviation of a variable from its trend such that the deviation tends to be above the threshold in crisis periods and below in normal periods. For example, Drehmann and Juselius (2014) employ quarterly time series data for 10 candidate variables from 26 countries and spanning the period 1980-2012 to compare the performance of those variables as early warning indicators for 19 financial crises that occurred during the observation period. They find that the credit to GDP ratio outperforms other measures. The purpose of this study is similar to that of these studies on early warning indicators, but the methodology we employ is quite different. Early warning exercises can be seen as a statistical procedure to detect anomalies in variables like the credit/GDP ratio around crises. Key to this procedure is prior information on crisis dates. However, as often pointed out in previous studies, it is not always easy to obtain precise information on crisis dates, and, more importantly, crises are by their very nature rare events. Given this, it is not necessarily easy to improve the precision of early warning indicators to a sufficiently high level such that policymakers can rely on them to take remedial actions. In this study, we do *not* seek to characterize how key variables behave in abnormal times. Instead, we focus on the behavior of key variables in normal times. Our strategy is to have a clear idea about how those variables behave in normal times and then to detect deviations from such normal behavior - that is, to detect anomalies that may serve as early warning signals.[2]

Our main findings are as follows. First, the cross-sectional distribution of house prices has a fat upper tail, and the tail part is close to a power-law distribution. This is confirmed

---

[2]Anomaly detection techniques can be categorized into supervised and unsupervised detection techniques depending on whether they rely on a training dataset that has labeled instances for normal as well as anomalous classes (see, for example, Chandola et al. 2009). The procedure adopted in the early warning studies can be regarded as *supervised* anomaly detection, as they label some observations of, for example, the credit/GDP ratio as anomalies using information on crisis dates. In contrast, our procedure is essentially *unsupervised* anomaly detection, since we do not label observations as normal or anomalous. See Chandola et al. (2009) for more on the distinction between supervised and unsupervised anomaly detection.

by the goodness-of-fit test proposed by Malevergne et al. (2011). On the other hand, the cross-sectional distribution of house sizes, as measured by the floor space, has an upper tail that is less fat than that of the price distribution and is close to an exponential distribution. These two findings suggest a particular functional form of hedonic regression to identify the size effect. We construct size-adjusted prices by subtracting the house size (multiplied by a positive coefficient) from the log price and find that the size-adjusted price follows a lognormal distribution for most of the observation period. An important exception is the period of the housing bubble and its collapse in 1987-1995, during which the price distribution in each year has a power-law tail even after controlling for the size effect. This suggests that the shape of the size-adjusted price distribution, especially the shape of the tail part, may contain information useful for the detection of housing bubbles.

Second, we conduct additional analyses to investigate in more detail the mechanism responsible for the deviation of the size-adjusted price distribution from a lognormal distribution during the bubble period. Specifically, we divide the area covered by our sample (Greater Tokyo) into small pixels and find that size-adjusted prices almost follow a lognormal distribution *within* each of these pixels even during the bubble period, but the means and variances of the distributions are highly dispersed *across* different pixels. This finding implies that the sharp price hike during the bubble period was concentrated in particular areas, and this spatial heterogeneity is the source of the fat upper tail observed for the bubble period.[3] We interpret this as evidence suggesting that during a bubble period there is an increase in market segmentation.

The rest of the study is organized as follows. Section 2 provides a brief overview of the Japanese housing bubble in the late 1980s. Section 3 then explains the dataset and the empirical strategy we employ. Next, Sections 4 and 5 present our size- and location-adjustments to house prices. Finally, Section 6 concludes the paper.

2   Overview of the Japanese Real Estate Bubble in the Late 1980s

In this section, we provide a brief overview of what happened during the Japanese real estate bubble in the late 1980s and its collapse in the 1990s, as well as how the government

---

[3]In this context, it is interesting to note that Cochrane (2005) argues that an important feature of the tech stock bubble in the late 1990s is that it was concentrated in stocks related to internet business. Cochrane (2005, p. 191) observes that "if there was a 'bubble,' or some behavioral overenthusiasm for stocks, it was concentrated on Nasdaq stocks, and Nasdaq tech and internet stocks in particular."

and the central bank responded.

2.1 The real estate bubble in the late 1980s and its collapse in the 1990s

Figure 1 shows changes in the mean of the cross-sectional house price distribution in the upper panel, the standard deviation in the middle panel, and the transaction volume in the lower panel. The data is compiled from individual listings in a real estate advertisement magazine published on a weekly basis.[4] We see that the mean price exhibits a sharp increase between the beginning of 1987 and the beginning of 1988. Previous studies refer to this as the first phase of the housing bubble in Tokyo. After a short break in 1988, prices started to rise again in 1989 and continued to do so until the fall of 1990. This is the second phase of the housing bubble. Soon after the fall of 1990, prices started to turn down, followed by a slow but persistent decline for more than a decade until prices bottomed out in 2002, when the mean price reached the level before the bubble started in 1987. Prices finally began to rise again in 2003 and continued to rise until registering a sharp decline in 2008 due to the global financial crisis.

Turning to the standard deviation shown in the middle panel, this exhibits a sharp rise during the first phase of the bubble and stayed high during the second phase.[5] Finally, the bottom panel, which shows the transaction volume, indicates that the number of transactions exhibits a sharp increase at the beginning of 1989, exactly when the mean price started to rise, although the transaction volume remained practically unchanged during the first phase of the bubble. Somewhat interestingly, the transaction volume remained at a high level even in 1991 and 1992, when the mean price had already started to decline.

Okina et al. (2001) identify three characteristics of the Japanese economy during the bubble period. The first characteristic is a rapid and substantial rise in asset prices, including stock and real estate prices. Stock prices exhibited a rapid rise during the initial stage of the bubble; the Nikkei 225 began to accelerate in 1986 and the index hit a peak of 38,915 yen at the end of 1989, when it was 3.1 times higher than at the time of the

---

[4]More details on the dataset used in the paper are provided in the next section.

[5]We also see a secular increase in price dispersion since 1993. We are not quite sure why this is the case, but recent studies, including those by Van Nieuwerburgh and Weill (2010) and Gyourko et al. (2006) find some evidence that the rise in house price dispersion across regions in the United States is related to the change in income distribution across regions. For example, Van Nieuwerburgh and Weill (2010) find that the cross-sectional coefficient of variation of house prices across 330 metropolitan statistical areas in the United States increased from 0.15 in 1975 to 0.53 in 2007. Through a counterfactual simulation, they show that this increase in the dispersion of house prices is accounted for mostly by the increase in income inequality.
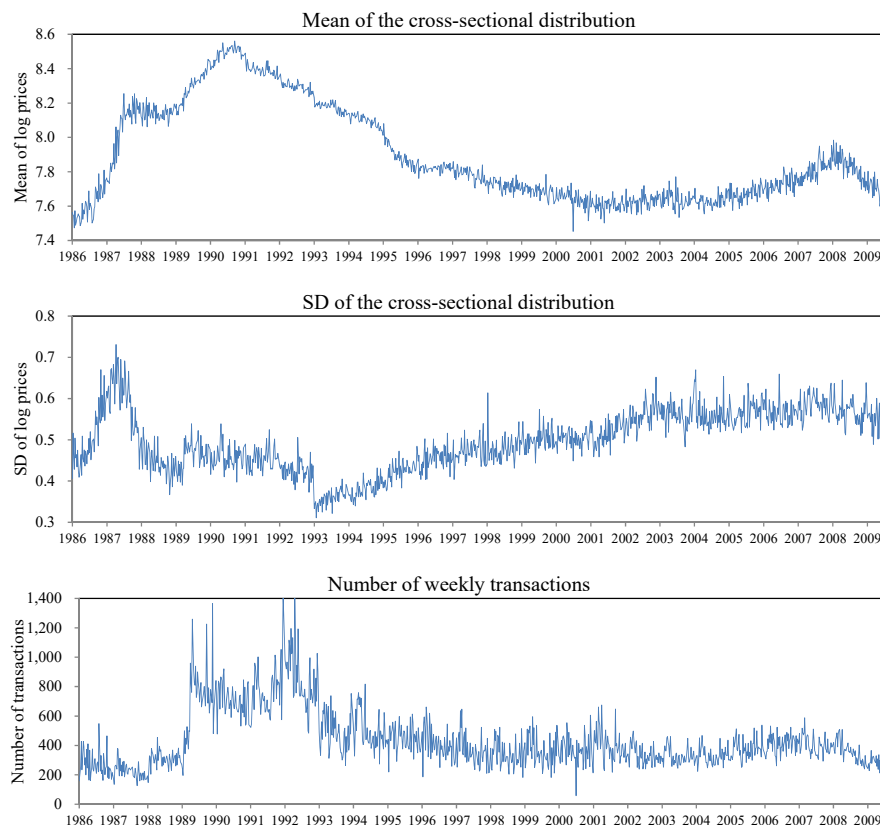
Figure 1: Weekly fluctuations in prices and transaction volume

Plaza Agreement in September 1985. Land prices followed stock prices with a lag of about one year, increasing sharply in 1987 and peaking in late 1990, as mentioned above, with increases spreading from Tokyo to other major cities such as Osaka and Nagoya, and then to other cities of smaller size. The second characteristic of this period is very high economic growth. The business cycle hit a trough in November 1986 and the economy then expanded for 51 months until 1991. Real GDP grew at an average annual rate of 5.5 percent during this period, driven by business fixed investment, housing investment, and expenditure on consumer durables. The third characteristic is a rapid increase in money supply and credit. The annual growth rate of money supply reached more than 10 percent in April-June 1987 as a result of monetary easing by the Bank of Japan as well as financial deregulation. Also, bank borrowing and financing from capital markets substantially increased in 1988 and 1989.

These three characteristics of the Japanese economy during the bubble period were

closely related with each other. In particular, land prices, bank borrowing, and business investment were tightly linked through the credit cycle mechanism highlighted by Kiyotaki and Moore (1997). Shimizu and Watanabe (2013), using information from land registry data, examine empirically whether such a link indeed exists by looking at the market value of land owned by firms and the amount of bank lending to those firms. They show that rapid price rises in the late 1980s raised the value of land as collateral, making it possible for banks to extend larger loans.

2.2 Causes of the real estate bubble in the 1980s

Previous studies on the real estate bubble in the 1980s identify the following as factors behind the emergence and expansion of the bubble: (1) aggressive behavior of financial institutions; (2) financial deregulation; (3) inadequate risk management by financial institutions; (4) the introduction of capital adequacy requirements for banks; (5) protracted monetary easing by the Bank of Japan; (6) taxation and regulation; (7) overconfidence and euphoria; (8) demographic changes; (9) overconcentration of economic functions in Tokyo, and Tokyo becoming an international financial center.[6] Obviously, these factors are not mutually independent. On the demand side, demographic changes, over-concentration, and euphoria are three important factors, creating excess demand in the real estate market. On the other hand, taxation and regulation were critically important on the supply side, because land taxation and land-related regulations made land supply price-inelastic, thereby making it impossible to eliminate excess demand without raising real estate prices. On the monetary side, expansionary monetary policy and banks' loose lending behavior are two important factors that made it possible for firms and individuals to have easy access to liquidity.[7] In the rest of this section, we will focus on the issues associated with demographic changes and taxation.

Demographic changes and housing demand   Mankiw and Weil (1989) argue that demographic changes, such as baby booms and busts, have an impact on housing demand and therefore on housing prices. This directly follows from the Ando-Modigliani life cycle hypothesis; namely, people buy houses during their working career and sell them in old age. A similar idea has recently been proposed by Nishimura (2011), who argues that the real estate bubble in the 1980s and its collapse in the 1990s were caused, at least

---

[6]See Okina et al. (2001) for more details on each of these factors.
[7]See Hoshi and Kashyap (2000) on banks' lending behavior during the bubble and its burst.

partially, by demographic changes. His argument is based on a simple comparison of the inverse dependency ratio (i.e., the ratio of the working-age population between 15 and 64 years of age to the rest of population, who are dependent on the working-age population) and real land prices. The inverse dependency ratio hit a trough in 1980 and exhibited a gradual increase until reaching a peak in 1990, when the children of the post-World War II baby boomers reached working age. The ratio then started to decline again from 1990 onward and has continued to do so in the more than two decades since then. These ups and downs in the inverse dependency ratio more or less coincide with those in real land prices. Nishimura (2011) interprets this as evidence that demographic changes affect land prices through changes in land demand.

Employing international panel data on 22 countries for 1970-2009, Takáts (2012) finds that an increase in the population by 1 percent is associated with an increase in real house prices by 1 percent, and that an increase in the old age dependency ratio by one percent is associated with a decrease in real house prices by 2/3 of a percent. These results imply that, on average, demographic factors raised house prices in advanced economies by around 30 basis point per annum in the past 40 years and will, on average, decrease house prices in advanced economies by around 80 basis points per annum over the next 40 year.

Following Mankiw and Weil (1989), Shimizu and Watanabe (2010) estimate the weighted average of housing demand of each generation, using the population share of each generation as weights, for each prefecture from 1975 to 2008. They find that the change in housing demand and the change in house price are positively correlated, implying that house prices tend to go up in a particular prefecture when housing demand there increases due to demographic reasons. Focusing on the bubble period (i.e., 1985-1990), they find that the cross-sectional correlation is particularly high, although the very sharp price increases in prefectures with large populations, such as Tokyo and Osaka, cannot be accounted for solely by demographic factors.

Land taxation and its reform in the early 1990s   It has been pointed out by many practitioners and researchers that the taxation system in Japan gave land owners an incentive to hold land for speculative purposes, thereby making land supply inelastic to price changes.[8] In this respect, the following two characteristics of the Japanese tax system play

---

[8]In Japan, land is taxed at three stages: inheritance tax, registration and license tax, and/or real estate acquisition tax are applied when land is acquired; fixed property tax, urban planning tax, and/or special land-holding tax are applied while land is held; and corporation tax, income tax, and/or inhabitant tax are

Table 1: Chronology of policy actions related to the real estate bubble in the 1980s and its collapse in the 1990s

1986	Tokyo Metropolitan Government introduces regulations on land transactions
1987	National Land Use Planning Act is amended
	Closer monitoring of land transactions starts
1990	Local Property Tax Act is amended
	Tax enforcement for agricultural land in urban areas is strengthened
1990	National Land Use Planning Act is amended
	Surveillance of resale of properties within two years is introduced
1990	Ministry of Finance issues a directive and guidance requesting financial institutions
	(1) to contain the increase in lending to property-related firms to within the increase
	in total lending, and (2) to report lending to the real estate industry,
	construction industry, and non-banks
1991	Land Value Tax Act is enacted
1992	Land Value Tax is newly introduced
1993	Cooperative Credit Purchasing Co. is established
1995	Tokyo Kyodo Bank (Japanese RCC) is established
1998	Act on Emergency Measures for Early Strengthening of Financial Functions is enacted

a key role. First, tax imposed on the holding of land has been extremely low in Japan when compared with other countries. This is particularly true for agricultural land, and the rate of tax imposed on the holding of agricultural land, even for agricultural land in urban outskirts, has been extremely low. This low cost of holding land, especially agricultural land, results in the inefficient use of land and in land being left unused, and makes holding land more profitable than holding financial and other non-financial assets, promoting speculative land transactions. Second, as for taxes on capital gains, the tax rates were not sufficiently high to prevent speculative transactions; that is, for individuals, an income tax of 20 percent was imposed on capital gains up to 40 million yen and a consolidated tax of 50 percent above this threshold: for companies, apart from ordinary corporation tax, an additional tax of 20 percent was imposed on capital gains on land held for ten years or less.

In the early 1990s, a series of tax reforms was implemented as a part of policy responses to suppress speculative land transactions following the publication of the report on "Basic Issues Regarding Revisions in the Land Tax System" in May 1990 by the Land Tax System Subcommittee. A chronology of tax reforms is presented in Table 1. The tax reforms implemented in the early 1990s consist of three parts. First, in January 1992, property tax rates were raised and a land value tax was newly introduced to make land less attractive as an investment, thereby suppressing speculative land transactions. Second, capital

applied to capital gains when land is transferred.

gains taxation was strengthened. In particular, a punitively high tax rate was imposed on capital gains from transactions in land that has been held only for a short period. For example, in the 1991 reform, a separation tax of 30 percent was additionally introduced for capital gains on transactions in land held by companies for two years or less. Finally, the tax rate associated with the holding of land for agricultural purposes in urban areas was raised to the same level as that for land for residential purposes in order to eliminate the special treatment of land for agricultural purposes. Practices and methods regarding land appraisals for inheritance tax purposes were also amended.

3 Data and Empirical Strategy

3.1 Data

We use a unique dataset that we compiled from individual listings in a widely circulated real estate advertisement magazine published on a weekly basis by Recruit Co., Ltd., one of the largest vendors of residential lettings information in Japan. The dataset covers the Greater Tokyo Area for the period 1986 to 2009, thus including the bubble period in the late 1980s and its collapse in the first half of the 1990s. It contains 724,416 listings for condominiums and 1,602,918 listings for single family houses.[9] In this paper we use data only for condominiums. The Greater Tokyo Area covered in the dataset includes the 23 special wards of Tokyo, other areas making up the Tokyo Metropolis, as well as adjacent cities and suburbs. The dataset covers more than 95 percent of all transactions in the 23 special wards (i.e., central Tokyo), while the coverage for the other areas is somewhat more limited. This dataset has been used in a series of studies, including Shimizu et al. (2010), which compares hedonic and repeat-sales measures in terms of their performance.

3.2 Empirical strategy

A widely used approach to deal with product heterogeneity in terms of quality is hedonic analysis, which has been employed in a number of studies to analyze real estate prices. The core idea of hedonic analysis is that the value of a product is the sum of the values of individual product characteristics. Rosen (1974) argues that a house can be regarded as a

---

[9]The dataset contains full information about the evolution of the posted price for a housing unit from the week when it was first listed until the final week when it was removed because of a successful transaction. In this study, we only use the price in the final week, since this can be safely regarded as sufficiently close to the contract price. The number of listings shown in the text does not include those prices listed before the final week.

set of attributes and that different houses consist of different sets of attributes. Most importantly, such a set of attributes essentially cannot be unbundled, so that each household buys one house and obtains a flow of housing services from it.[10] Shimizu et al. (2010) start their analysis by assuming that the value of a house is the sum of the values of attributes such as its floor space, its age, the commuting time to the nearest station, and so on, and run hedonic regressions using these attributes as independent variables.

This idea has important implications regarding the shape of the cross-sectional distribution of house prices. To show this, let us start by assuming that the price of house i at a particular point in time, which is denoted by P_i,[11] is the sum of K components:

$$P_i = F(X_{i1}, X_{i2}, \ldots, X_{ik}, \ldots, X_{iK}) \tag{1}$$

where P_i and X_{ik} are both random variables and $X_{i1}, \ldots, X_{iK}$ are assumed to be independent from each other. Furthermore, we assume a multiplicative functional form such that

$$P_i = \prod_{k=1}^{K} X_{ik}. \tag{2}$$

Taking the logarithm of both sides of this equation yields

$$\ln P_i = \sum_{k=1}^{K} x_{ik}, \tag{3}$$

where $x_{ik} \equiv \ln X_{ik}$. This equation appears frequently in hedonic analyses of house prices. It simply states that the price of a house is equal to the sum of K random variables.

Given this setting, the central limit theorem tells us that the sum of these random variables converges to a normal distribution if the number of attributes, K, goes to infinity. Let us denote the variance of x_{ik} by s_k^2 and define the average variance $\bar{s}_K^2$ as

$$\bar{s}_K^2 \equiv \frac{1}{K} \left(s_1^2 + s_2^2 + \cdots + s_K^2 \right).$$

Then, according to the Lindeberg-Feller central limit theorem, the sum of random variables $\sum_{k=1}^{K} x_{ik}$ converges to a normal distribution as K goes to infinity, if the average variance

---

[10]A similar idea has been adopted in recent studies applying assignment models to the housing markets including Landvoigt et al. (2015), Piazzesi and Schneider (2016), Määttänen and Terviö (2014), and Rios-Rull and Sanchez-Marcos (2008). In these models, equilibrium prices are determined so that households with high (low) demand for housing services are assigned to high (low) quality houses. These studies show that the distribution of equilibrium prices across houses is determined by the distribution of quality across houses and the distribution of characteristics across households. Rosen (1974) and Landvoigt et al. (2015) show that the assumption of no costless unbundling implies limits to arbitrage in housing markets.

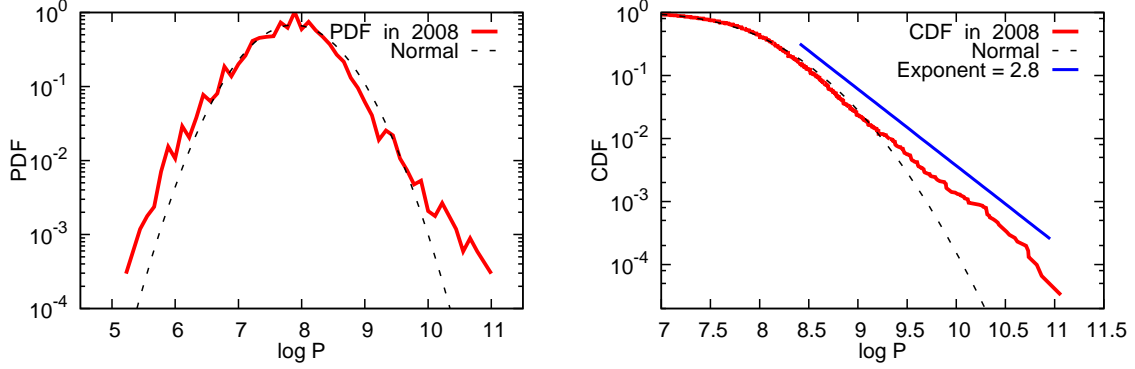[11]Note that the subscript for time is dropped here to simplify the exposition.

Figure 2: House price distribution in 2008

$\bar{s}_K^2$ converges to a finite constant (namely, $\lim_{K\to\infty} \bar{s}_K^2 = \bar{s}^2$) and the following condition is satisfied:

$$\lim_{K\to\infty} \frac{\max_{k\leq K}\{s_k\}}{K\bar{s}_K} = 0. \tag{4}$$

In other words, the theorem states that the sum of random variables will tend to be normally distributed as far as condition (4), which is referred to as the Lindeberg condition, is satisfied. A notable feature of this result is that it does *not* require that the variables in the sum come from the same underlying distribution. Instead, the theorem requires only that no single term dominates the average variance, as stated in (4). Put differently, the Lindeberg condition states that none of the random variables is dominantly large relative to their sum.[12] A famous textbook example of the central limit theorem is the distribution of persons' height. The height distribution of, say, mature men of a certain age can be considered normal, because height can be seen as the sum of many small and independent effects. Similarly, the log price of houses will be normally distributed if house prices are determined as the sum of many small and independent effects.

The above argument suggests that the lognormal distribution can be seen as a benchmark for the cross-sectional distribution of house prices. However, some previous studies on house price distributions find that the actual distributions have fatter tails than a lognormal distribution. For example, McMillen (2008), using data on single family houses in Chicago for 1995, shows that the kernel density estimates for the log price are asymmetric, with a much fatter lower tail. Against this background, we examine the extent to which the house price distribution deviates from a lognormal distribution using our observations

---

[12]For more on this theorem, see Feller (1968). Greene (2003) provides a compact description of various versions of the central limit theorem including this one.

for 2008. The results are presented in Figure 2, where the left panel shows the probability density function (PDF), with the horizontal axis representing the yen price in logarithm and the vertical axis representing the corresponding density, also in logarithm. The empirical distribution is shown by the red line, while the lognormal distribution with the same mean and standard deviation as in the data is shown by the black dotted line. The figure indicates that the tails of the empirical distribution are fatter than those of the lognormal distribution. In particular, the upper tail of the empirical distribution is much fatter than that of the lognormal distribution. To examine the differences in the upper tail more closely, we accumulate the densities from the right (upper) tail to produce the cumulative distribution function (CDF), which is shown in the right panel. In this panel, the value on the vertical axis corresponding to the value of 9.2 on the horizontal axis, for example, is 0.01, meaning that the fraction of houses whose prices are equal to or higher than that price level is 1 percent. We now see more clearly that the upper tail of the empirical distribution is fatter than that of the lognormal distribution. For example, the fraction of housing units whose price deviates from the mean by more than 3σ is about 1.47 percent, while the corresponding number for the lognormal distribution is only 0.26 percent.

What causes the empirical distribution to deviate from the benchmark (i.e., the lognormal distribution)? This is the main topic we address in this paper. Our hypothesis is that some of the factors that determine house prices are dominantly volatile, so that the Lindeberg condition is violated. Denoting these dominant factors by vector Z_i, the house price distribution, $\Pr(P_i = p)$, can be decomposed as follows:

$$\Pr(P_i = p) = \sum_z \Pr(P_i = p \mid Z_i = z)\Pr(Z_i = z). \tag{5}$$

Note that the house price distribution conditional on Z_i, namely $\Pr(P_i = p \mid Z_i = z)$, should be a lognormal distribution, since the dominant factors are now fully controlled for. This means that the right-hand side of equation (5) is a weighted sum of lognormals, with the weights being given by $\Pr(Z_i = z)$. We know that the sum of lognormals with different means and variances is no longer a lognormal (see, for example, Feller 1968), and our hypothesis is that this is why the house price distribution deviates from the benchmark. Given this hypothesis, we proceed as follows in the remainder of the study: we first specify the dominant factors and then eliminate them, thereby constructing prices that are adjusted for these factors; finally, we examine whether these adjusted prices follow a lognormal distribution.
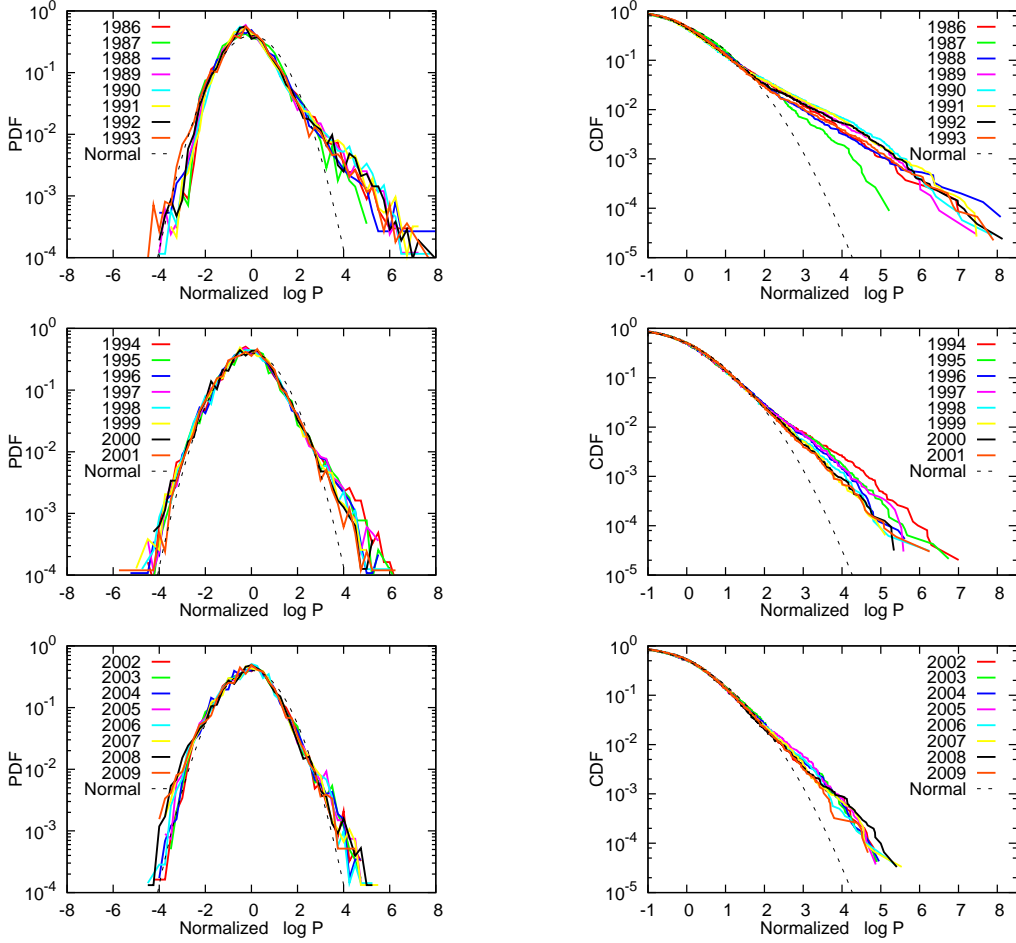
Figure 3: House price distributions by year

Diewert et al. (2011) show that the most important characteristics determining the price of a property are the land area of the property, the livable floor space area of the structure, and the location of the property. Similarly, previous studies on house prices in Japan, including Shimizu et al. (2010), find that the two key determinants of the price of a housing unit (especially in the case of condominiums) are its floor space and its location. These empirical evidence suggests that the size and the location of a property are key candidates for the Z variables. We will identify and eliminate the size effect in the next section and the location effect in Section 5.

4 Size-adjustment to House Prices

4.1 Distribution of unadjusted house prices

Figure 3 presents the PDF and CDF of the cross-sectional price distribution for each year from 1986 to 2009. To make the price distributions in different years comparable, we normalize the log prices in year t by subtracting the mean in year t (i.e., the mean of log prices in year t) and dividing by the standard deviation in year t (i.e., the standard deviation of log prices in year t). The lognormal lines in the figure represent the CDF of a standard lognormal distribution. Note that the CDFs are constructed in the same way as in Figure 2, that is, the value on the vertical axis corresponding to a price level is the sum of the densities *above* that price level.

The first thing we see from the figure is that, as in Figure 2, the PDFs and the CDFs show fatter upper tails than a lognormal distribution. More importantly, we see that the deviation from a lognormal distribution tends to be larger for the late 1980s and early 1990s. Specifically, the PDFs in these years are substantially skewed to the right, indicating that during the bubble period house prices did not rise by the same percentage for every housing unit; instead, price increases were concentrated in particular housing units, so that relative prices across houses changed significantly.
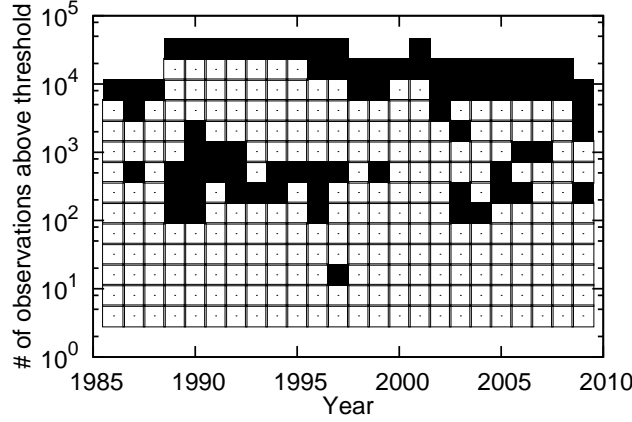
The CDFs in this figure provide more detailed information regarding the shape of the price distributions. We see that the CDF for each year forms an almost straight line in this log-log graph, implying that the house price distribution is well approximated by a power law distribution (or a Pareto distribution) at least in the tail part, the PDF and CDF of which are given by

$$\Pr(P_{it} = p) = \frac{\zeta_t m_t^{\zeta_t}}{p^{\zeta_t+1}}; \quad \Pr(P_{it} \geq p) = \left(\frac{m_t}{p}\right)^{\zeta_t}; \quad p > m_t > 0 \tag{6}$$

where P_{it} denotes the price of house i in period t, and ζ_t and m_t are time-variant positive parameters.[13] The shape of a power law distribution is mainly determined by the parameter ζ_t, which is referred to as the exponent of the power law distribution. Smaller values for ζ_t imply fatter tails. Note that the CDF given in (6) implies that

$$\ln \Pr(P_{it} \geq p) = -\zeta_t \ln p + \zeta_t \ln m_t.$$

---

[13]See Gabaix (2009) for an extensive survey of empirical and theoretical studies on power laws in various economic contexts such as income and wealth, the size of cities and firms, and stock market returns.

■: The null (power law) is rejected at the 1 percent significance level.
□: The null is not rejected at the 1 percent significance level.

Figure 4: Power law distribution versus lognormal distribution

In other words, the log of the cumulative probability should be linearly related to the log price, and the slope of a linear line between the two variables should be equal to $-\zeta_t$. The CDFs in Figure 3 suggest the presence of such a linear relationship between the log price and the log of the cumulative probability. We see from the CDF in Figure 2 that ζ_{2008} is about 2.8. Similarly, we find from the corresponding figures for the other years (which are not shown due to space limitations) that ζ for those years also takes a value of around 3.[14]

As a goodness-of-fit test, we employ the test proposed by Malevergne et al. (2011). Specifically, we test the null hypothesis that, beyond some threshold u, the upper tail of the house price distribution is characterized by a power law distribution,

$$\Pr(P = p; \ \alpha) = \alpha \cdot \frac{u^{\alpha}}{p^{\alpha+1}} \cdot 1_{p \geq u}$$

against the alternative that the upper tail follows a lognormal beyond the same threshold,

---

[14]Note that we cannot obtain estimates for ζ_t from Figure 3. The CDFs in Figure 3 are for normalized prices, which are defined by $[P_{it} \exp(-\mu_t)]^{1/\sigma_t}$, where μ_t and σ_t are the mean and the standard deviation in year t. Therefore, the slope of each CDF in Figure 3 is given by $\sigma_t \zeta_t$ (rather than ζ_t), if the original price follows the power law distribution given by (6). Many examples of power law distributions have been provided. For example, the net worth of Americans follows a power law distribution with an exponent of 1.1; the frequency of the use of words follows such a distribution with an exponent of 1.2; the population of US cities has an exponent of 1.3; the number of hits on websites has an exponent of 1.4; the magnitude of earthquakes has an exponent of 2.8; and price movements in financial markets have an exponent of 3 (or lower). The exponents for the house price distributions estimated here are greater than most of these figures, implying that the tail parts of the house price distributions are less fat than those in the other examples of power law distributions.

i.e.,

$$\Pr(P = p; \ \alpha, \beta) = \left[\sqrt{\frac{\pi}{\beta}} \exp\left(\frac{\alpha^2}{4\beta}\right) \left(1 - \Phi\left(\frac{\alpha}{\sqrt{2\beta}}\right)\right) \right]^{-1} \frac{1}{p} \exp\left(-\alpha \ln \frac{p}{u} - \beta \ln^2 \frac{p}{u}\right) \cdot 1_{p \geq u},$$

where $\Phi(\cdot)$ represents the CDF of a standard normal distribution. Note that this is equivalent to testing the null that the upper tail of the *log* price follows an exponential distribution against the alternative that it follows a normal distribution. For this transformed test, Del Castillo and Puig (1999) have shown that the clipped empirical coefficient of variation $\hat{c} \equiv \min(1, c)$ provides the uniformly most powerful unbiased test, where c is the empirical coefficient of variation.

The result of our goodness-of-fit test is presented in Figure 4, where the horizontal axis represents the year and the vertical axis represents the number of observations above the threshold u. For example, 10^3 on the vertical axis means that the threshold u is set such that the number of observations above u is 10^3. A black square indicates that the null is rejected at the 1 percent significance level for a particular year-threshold combination, while a white square indicates that the null is not rejected at the same significance level. The figure shows that a power law distribution provides a good approximation for the 500 most expensive houses, while a lognormal distribution provides a better approximation for the set of less expensive houses.

4.2   Distribution of house sizes

Previous studies on wealth (or income) distributions across households have typically found that those distributions are characterized by fat upper tails, and that they follow a power law distribution. Given that houses form an important part of households' wealth, it may be not that surprising that we detect a similar pattern in the house price distribution. However, the result that house prices follow a power law distribution is not consistent with the argument based on the central limit theorem. Why do house prices follow a power law distribution rather than a lognormal distribution? As a first step to address this question, we decompose the house price distribution as follows:

$$\Pr(P_{it} = p) = \sum_s \Pr(P_{it} = p \mid S_i = s) \Pr(S_i = s), \tag{7}$$

where S_i represents the size of housing unit i, which is measured by the floor space of that unit. The term $\Pr(S_i = s)$ represents the distribution of house sizes, while the term $\sum \Pr(P_{it} = p \mid S_i = s)$ represents the distribution of house prices conditional on house
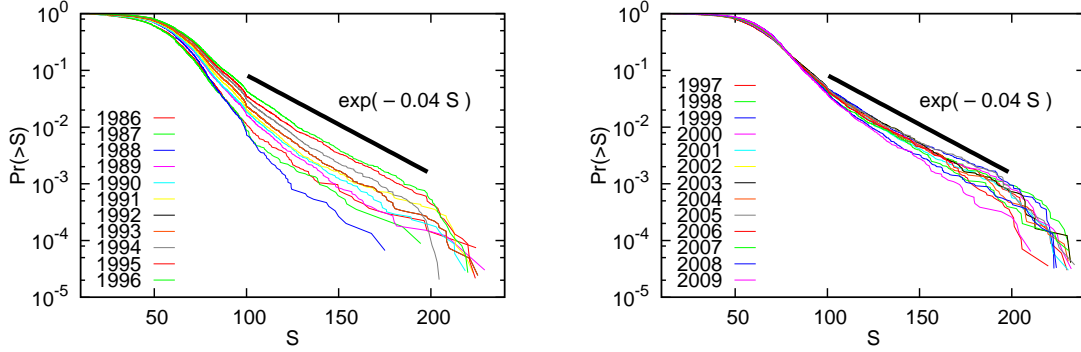
Figure 5: Cumulative house size distributions

size. An important thing to note is that even if each of these conditional distributions is lognormal, the weighted sum of lognormals with different mean and variance is not a lognormal distribution. This is a potential source of the power law tails that we observe in our house price data.

We start by examining the term $\Pr(S_i = s)$ in equation (7). Figure 5 presents the CDFs of house sizes for each year, with the floor space, measured in square meters, on the horizontal axis and the log of the CDF on the vertical axis. We see that the CDF for each year is close to a straight line in this semi-log graph, implying that the size distribution can be approximated by an exponential distribution whose PDF and CDF are given by

$$\Pr(S_i = s) = \lambda_t \exp\left(-\lambda_t s\right); \quad \Pr(S_i \geq s) = \exp\left(-\lambda_t s\right); \quad \lambda_t > 0. \tag{8}$$

Note that the CDF shown above implies that

$$\ln \Pr(S_i \geq s) = -\lambda_t s,$$

so that the log of the CDF depends linearly on house size. This is what we see in Figure 5. The slope of the CDF line, namely the value of λ, is almost identical for the different years and is somewhere around 0.04.

The fact that house sizes follow an exponential distribution implies that the tails of the size distribution are less fat than those of the price distribution. For example, for 2008, the fraction of housing units whose size deviates from the mean by more than 3σ is only 0.94 percent, while the corresponding number for the price distribution is 1.47 percent.[15]

---

[15]To see why the tails of the house size distribution are less fat than the tails of the price distribution, consider a simple example in which household A has 100 times as much wealth as household B, so that A spends 100 times as much money on a house as B. What does A's house look like? Does it have a bathroom

4.3 Size-adjusted prices

We now turn to the relationship between the price of a house and its size, which is represented by the conditional probability $\Pr(P_{it} = p \mid S_i = s)$ in equation (7). We propose a hedonic model that is consistent with the fact that house prices and sizes follow, respectively, a power law distribution with an exponent of ζ_t and an exponential distribution with an exponent of λ_t. That is, the log prices are determined by

$$\ln P_{it} \sim \left(\frac{\lambda_t}{\zeta_t} \right) S_i + \epsilon_{it}, \tag{9}$$

where ϵ_{it} is a normally distributed random variable, which, as we saw in Section 3.2, can be interpreted as the sum of many small and independent factors. To obtain equation (9), we first note that the PDF of the exponential distribution given in (8) implies that $(\lambda_t/\zeta_t)S_i$ follows an exponential distribution with an exponent of ζ_t if S_i itself is an exponential distribution with an exponent of λ_t. Next, we can show that the sum of the random variable that follows an exponential distribution and the random variable that follows a normal distribution is well approximated by the exponential distribution when the sum takes large values (because of the much fatter tails of an exponential distribution). Combining the two, the right-hand side of (9) is well approximated by an exponential distribution with an exponent of ζ_t when the sum of the two terms on the right-hand side takes large values. On the other hand, the fact that P_{it} follows a power law distribution with an exponent of ζ_t implies that $\ln P_{it}$ follows an exponential distribution with an exponent of ζ_t. Thus, we can confirm that each side of equation (9) follows an identical distribution with an identical exponent.[16]

To empirically test the hedonic model given by (9), we first examine for a linear relationship between the log price of houses and their size. The upper panels of Figure 6 show the floor space on the horizontal axis and the median of the log price corresponding to that size on the vertical axis. These panels indicate that there exists a stable linear rela-

---

that is 100 times larger than the one in B's house? Alternatively, does it have 100 bathrooms? Needless to say, neither is true, because even a person of A's wealth would have little use for such a gigantic bathroom (or so many bathrooms). Instead, it is more likely that the size of A's house (and therefore the size and/or number of its bathroom) is only, say, 10 times greater and, consequently, the unit area price of A's house is 10 times higher than B's. This is similar to the finding by Bils and Klenow (2001) that richer households not only consumer more goods but also higher quality and therefore more expensive goods.

[16]The price-size relationship described by equation (9) provides an answer to the question regarding the choice of functional form for hedonic price equations, which has been extensively discussed in previous studies such as Cropper et al. (1988). The novelty of our approach is that we derive this functional form not from economic theory but from the statistical fact that house prices and sizes follow a power law and an exponential distribution, respectively.
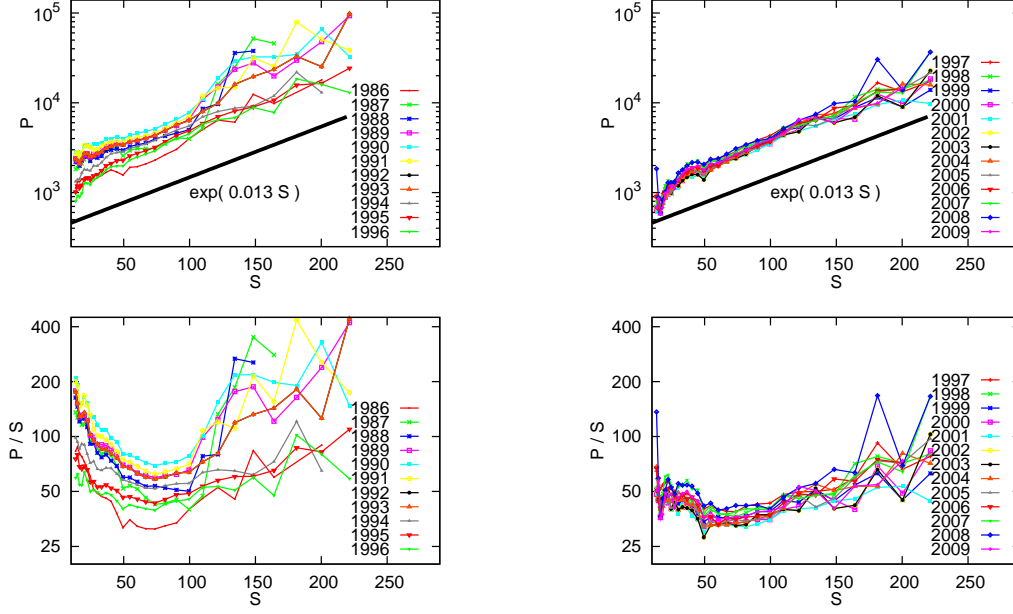
Figure 6: Relationship between house size and price

tionship between the two variables. Furthermore, equation (9) implies that the price per square meter, $P/S = [\exp(\lambda/\zeta)S + \text{positive constant}]/S$, decreases with S when S is small and increases with S when S is sufficiently large, so that there should exist a U-shaped relationship between the price per square meter and house size. The lower panels of Figure 6, in which the vertical axis now measures P/S, confirms this prediction.

Second, we run an ordinary least squares regression of the form

$$\ln P_{it} = a_t S_i + b_t + \eta_{it} \tag{10}$$

to see whether the disturbance term η_{it} is indeed normally distributed as assumed in (9). The regression results are presented in Figures 7 and 8. Figure 7 shows the estimates of a and b for each year. The estimate of a is almost identical across years and is around 0.013, implying that an increase in house size by a square meter leads to a 1.3 percent increase in the house price. More importantly, the estimate of a is very close to the value predicted by (9). That is, the value of ζ is around 3, as we saw in Section 4.1, and the value of λ is about 0.04, as we saw in Section 4.2, so that the coefficient on S_i, namely λ/ζ, should be somewhere around 0.013 ($= 0.04/3$). This is quite close to the point estimate of a for each year.[17] Turning to the estimate of b, this exhibits substantial fluctuations: it increases by

---

[17]Note that the price per square meter, $\exp(aS + b)/S$ takes its minimum value when S is equal to
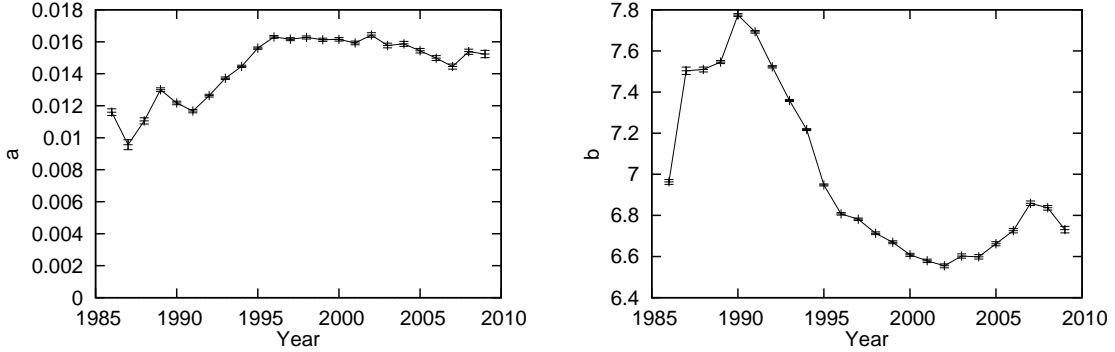
21

Figure 7: Price-size regressions

more than 20 percent per year from 1986 to 1990 and then declines by 10 percent per year from 1990 to 2002.

Figure 8 shows the CDFs of size-adjusted prices, which are defined as

$$\tilde{P}_{it} \equiv \left[P_{it} \exp(-\hat{a}_t S_i - \hat{b}_t) \right]^{1/\hat{\sigma}_t}, \tag{11}$$

where $\hat{a}_t$ and $\hat{b}_t$ are the estimates of a_t and b_t, and $\hat{\sigma}_t$ is the estimate for the standard deviation of η_{it}. Note that the hedonic model given by (9) implies that $\tilde{P}_{it}$ should be a lognormal distribution. The CDFs of the size adjusted prices are shown in the three panels on the right-hand side of Figure 8, while the price distributions *without* size adjustments from Figure 3 are replicated on the left-hand side. Comparing these two sets of CDFs, we see that the CDFs of the size-adjusted prices are much closer to the CDF of a lognormal distribution. More specifically, the CDFs for 2002 to 2009, which are shown in the lower right panel, are almost identical to the CDF of a lognormal distribution. The same applies to the CDFs for 1994 to 2001, which are shown in the middle right panel. However, the CDFs for 1986 to 1993, which are presented in the upper right panel, are still far from the CDF of a lognormal distribution, although they are slightly closer to it than the CDFs of the non-adjusted prices.

We apply the same methodology to land price data for Japan.[18] An advantage of the land price data is that it covers the period 1974-2008, which is much longer than the period

---

$1/a$. Given the estimate of a, this implies that the price per square meter takes its minimum value when $S = 1/0.013 \approx 75$, which is consistent with what we see in the lower two panels of Figure 6.

[18]This dataset contains land prices for about 30,000 sites across Japan. The price of each site is assessed by the Land Appraisal Committee once a year and made available to the public. The dataset contains information about various attributes of a particular piece of land, including its address, size, shape, types of land use (residential, commercial, or industrial), etc. In this paper, we focus on land for residential use that is smaller than 300 square meters and located in the Greater Tokyo Area. The number of observations ranges from 2,000 to 5,000 per year.
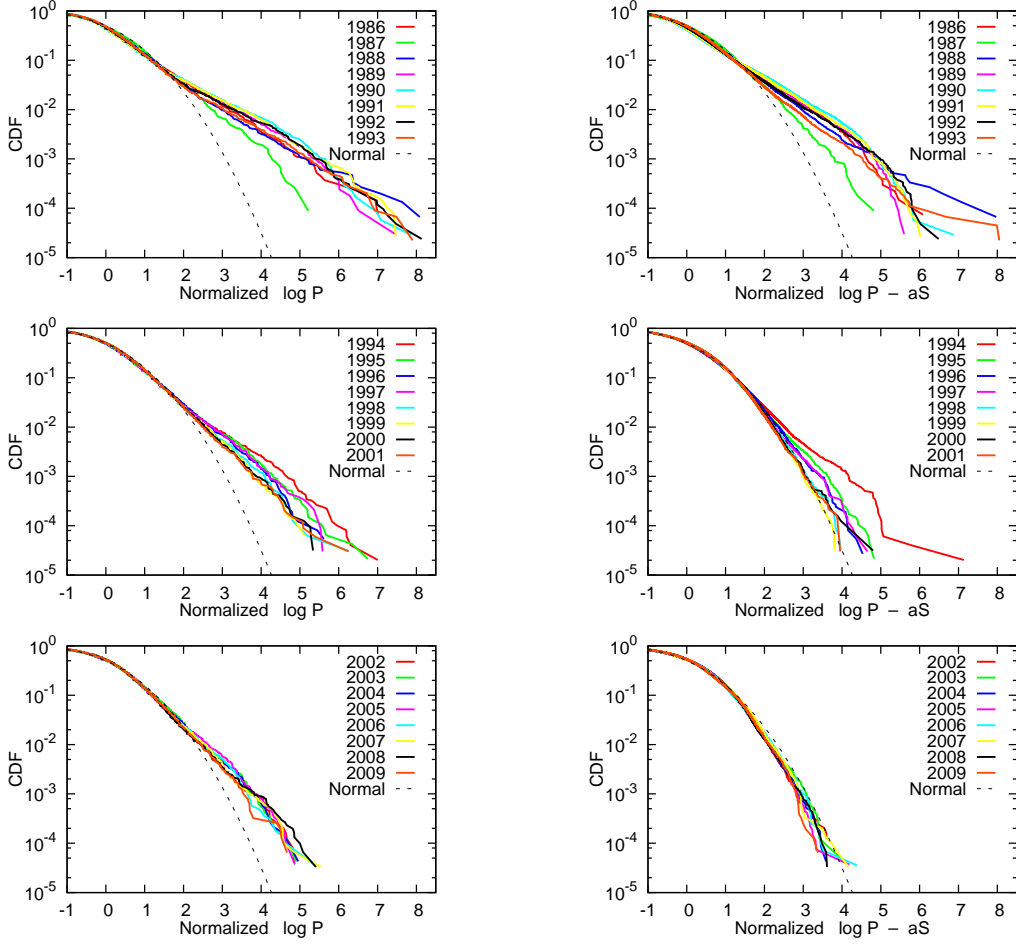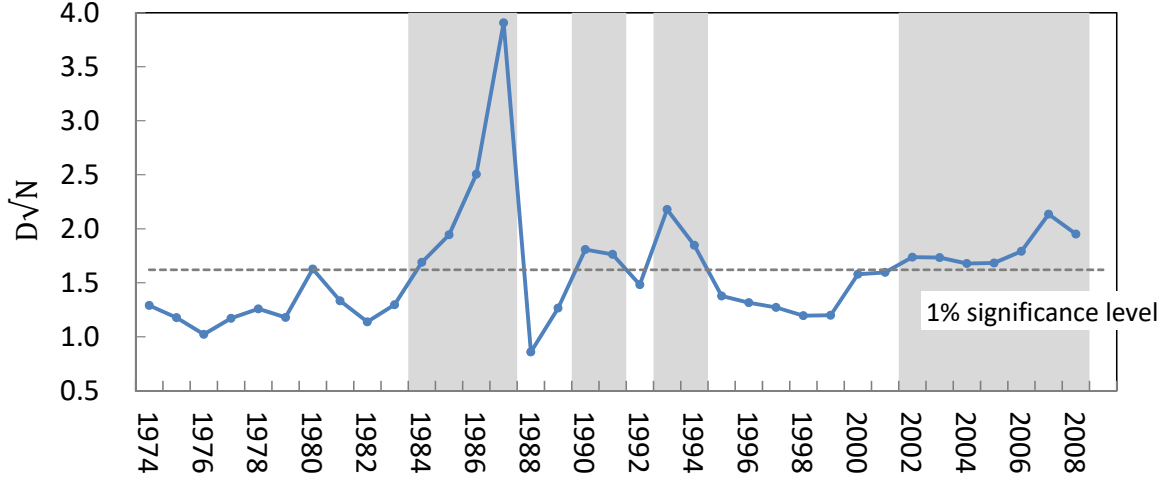
Figure 8: Cumulative distributions of size-adjusted house prices

covered by the house price data used in the previous subsections. Using this dataset, we quantitatively evaluate the degree to which the CDF of size-adjusted prices deviates from the CDF of the normal distribution. We denote this measure of deviation as $D\sqrt{N}$, where D is the maximum difference between the empirical CDF and the Gaussian CDF, and N is the number of observations. This is the statistic used in the Kolmogorov-Smirnov test, which is a nonparametric test to determine if a sample comes from a population with a specific distribution (the normal distribution in our case). Figure 9 shows the result. The null hypothesis that price observations are drawn from the normal distribution is rejected at the 1 percent significance level if $D\sqrt{N} > 1.62$. We see a statistically significant deviation from the normal distribution in 1984-1987, 1990-1991, 1993-1994, and 2002-2008. The period 1984-87 corresponds to the first phase of the bubble, which was explained in Section 2, while the periods 1990-91 and 1993-94 correspond to the second phase. More

23

Note: The shaded areas indicate that the null that price observations are normally distributed is rejected at the 1 percent significance level.

Figure 9: Deviation from the normal distribution of the land price data

interestingly, we see a statistically significant deviation in 2002-2006, when the Bank of Japan implemented quantitative easing to escape from deflation, as well as in 2007-2008, when the Bank of provided liquidity and additional monetary easing in response to the onset of the global financial crisis.

5 Location Adjustment to House Prices

The analysis in the previous section suggested that size-adjusted prices followed a lognormal distribution at least during calm periods without large price fluctuations. This is consistent with the idea that, as stated in (7), the power law tails of the original prices stem from the mixture of lognormal distributions with different mean and variance. At the same time, the analysis in the previous section showed that the fat tails of the price distribution remained largely unchanged during the bubble period (i.e., the late 1980s and the first half of the 1990s) even after controlling for the size effect. This suggests that there still remains some mixing of lognormal distributions.

In this section, we test the hypothesis that the power law tails of the size-adjusted price distribution during the bubble period arise due to the fact that the price distribution consists of a mixture of different lognormal distributions corresponding to different regions.
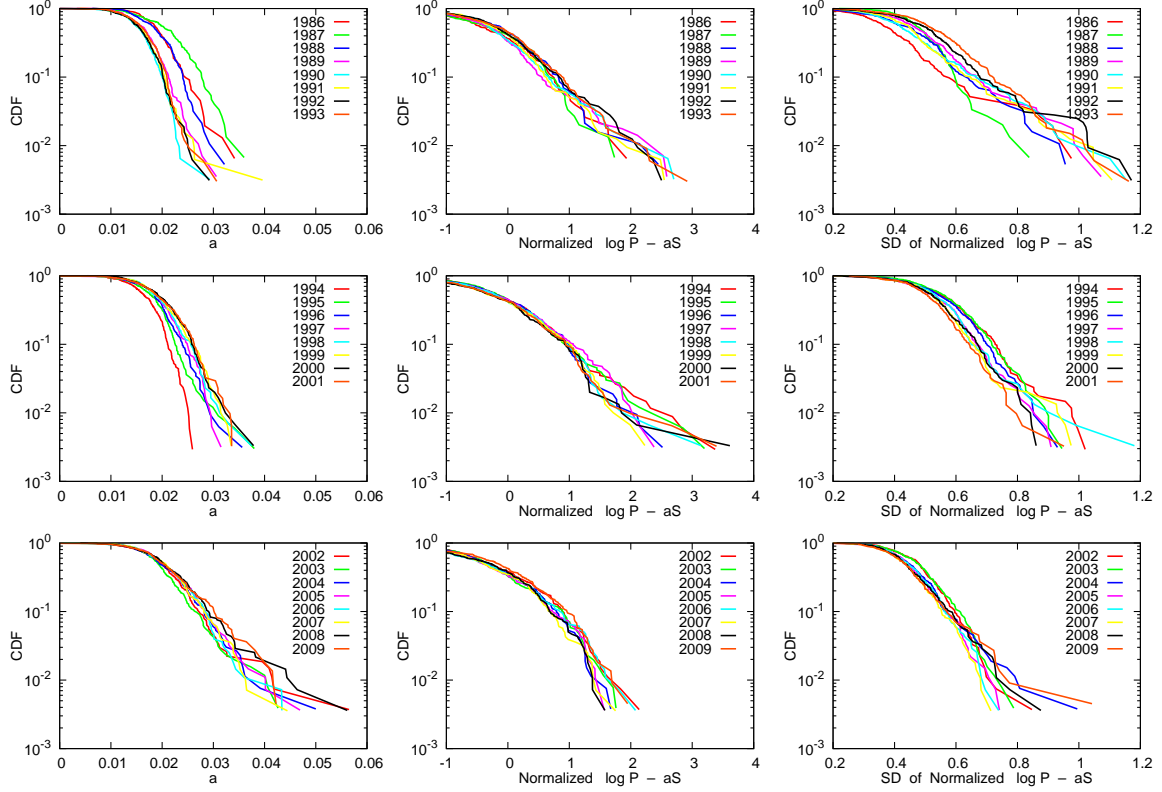
Figure 10: Dispersion of a_r, b_r and σ_r across pixels

To do so, we start by decomposing the size-adjusted price distribution into the sum of conditional distributions:

$$\Pr\left(\tilde{P}_{i_r t} = p\right) = \sum_{\theta} \Pr\left(\tilde{P}_{i_r t} = p \mid \theta_{rt} = \theta\right) \Pr\left(\theta_{rt} = \theta\right), \tag{12}$$

where $\tilde{P}_{i_r t}$ denotes the size-adjusted price of a house located in region r, which is defined by $\tilde{P}_{i_r t} \equiv P_{i_r t} \exp(-a_{rt} S_{i_r} - b_{rt})$. The vector of parameters θ_{rt} is defined as

$$\theta_{rt} \equiv (a_{rt}, b_{rt}, \sigma_{rt}), \tag{13}$$

where the parameters a_{rt}, b_{rt}, and σ_{rt} are the coefficient on the house size variable, the constant term, and the standard deviation of the disturbance term in equation (10), but it is assumed in this section that they could differ depending on the location. The location effect is fully controlled for in the conditional distributions $\Pr(\tilde{P}_{i_r t} = p \mid \theta_{rt} = \theta)$, so that they should be lognormal. According to equation (12), the distribution of $\tilde{P}_{i_r t}$ is a mixture of these lognormal distributions, each of which is for a different region.

25

We first examine the distribution of θ_{rt} across different regions. Specifically, we divide the Greater Tokyo Area into pixels of 0.033 degrees latitude and 0.033 degrees longitude or roughly 3.3 by 3.3 km.[19] Then, using size-adjusted prices within a pixel, we run a regression of the form

$$\ln P_{i_r t} = a_{rt} S_{i_r} + b_{rt} + \eta_{i_r t} \tag{14}$$

for each combination of r and t and obtain $\hat{\theta}_{rt} \equiv (\hat{a}_{rt}, \hat{b}_{rt}, \hat{\sigma}_{rt})$. The regression results are presented in Figure 10.[20] The three panels on the left show the CDFs of $\hat{a}_{rt}$, while the panels in the middle and on the right respectively show the CDFs of $\hat{b}_{rt}$ and $\hat{\sigma}_{rt}$. The CDFs of $\hat{a}_{rt}$ indicate that a is less dispersed across pixels during the period of the bubble and its collapse (1987-1995) than in the other years. On the other hand, the CDFs of $\hat{b}_{rt}$ and $\hat{\sigma}_{rt}$ show that these parameters are more highly dispersed during the same period, implying that the sharp price hike during the bubble period was concentrated in particular pixels. Put differently, the housing market was segmented during this period.

Next, we investigate whether the conditional distributions are close to a lognormal distribution. Using the estimates of θ_{rt} obtained from the regression, we calculate the size-adjusted prices for each pixel, which are given by:

$$\tilde{P}_{i_r t} \equiv \left[P_{i_r t} \exp(-\hat{a}_{rt} S_{i_r} - \hat{b}_{rt}) \right]^{1/\hat{\sigma}_{rt}} . \tag{15}$$

The estimated CDFs of $\tilde{P}_{i_r t}$ are presented in Figure 11 for the years 1986, 1990, 1994, 1998, 2002, and 2006. Note that each of the six panels contains four different lines, each of which corresponds to a different pixel size, namely 4.190 by 4.190 degrees, 0.524 by 0.524 degrees, 0.263 by 0.263 degrees, and 0.033 by 0.033 degrees. The results for 1998, 2002, and 2006 indicate that the CDFs are very close to a lognormal distribution, irrespective of pixel size. This is not very surprising given that, as we saw in the previous section, the CDFs in these years were already close to a lognormal distribution before controlling for the location effect. For the period of the bubble and its collapse, we see more interesting results: for 1986, 1990, and 1994, the estimated CDF tends to be closer to a lognormal distribution the smaller the pixel size.[21]

---

[19]Note that one degree is approximately 100 km.

[20]In conducting these regressions, we use only pixels with more than twenty transactions in a year. The number of pixels used in the regressions is about 300 for each year.

[21]It should be noted that the estimated CDF does not fully converge to a lognormal even in the case of the smallest pixels. It may be the case that the CDF becomes much closer still to a lognormal distribution if we were able to reduce the pixel size even further. Unfortunately, we cannot do so because of the limited number of observations.
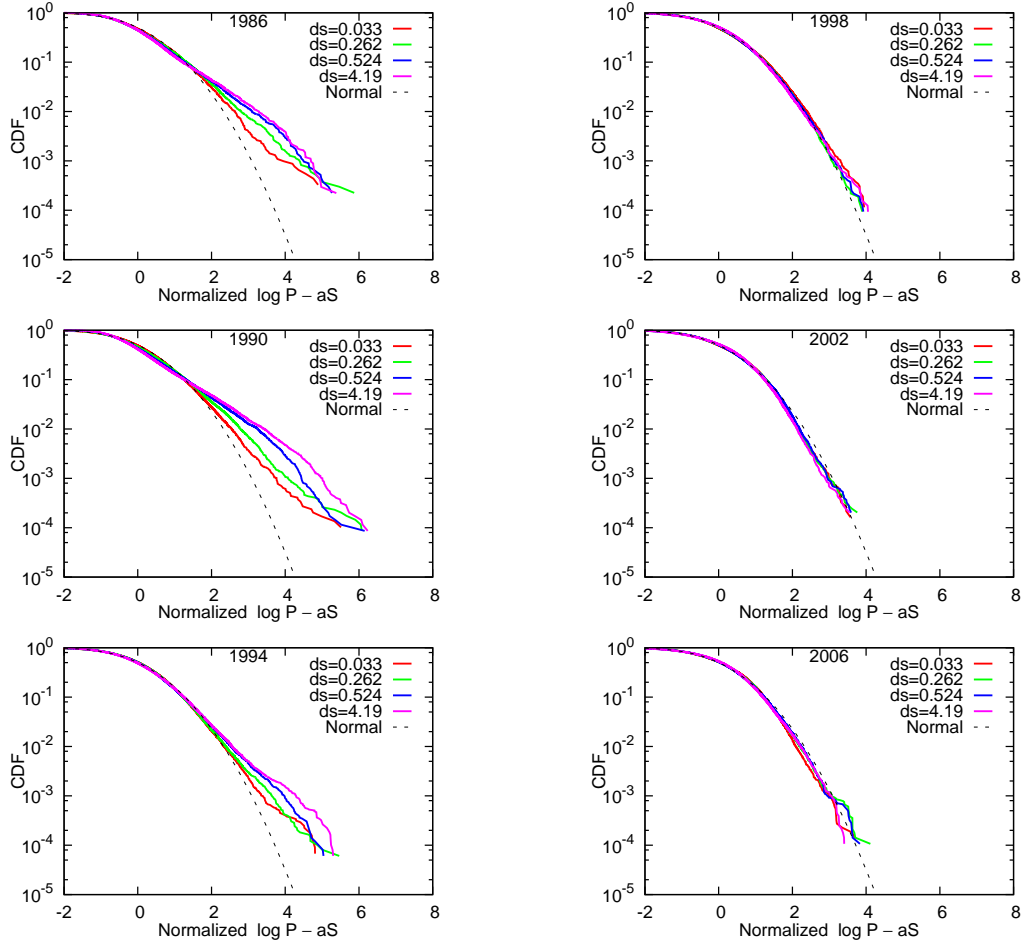
Figure 11: Cumulative distributions of size-adjusted house prices for different pixel sizes

In sum, the analysis in this section shows that the distribution of size-adjusted prices *within* a pixel is fairly close to a lognormal distribution even for the period of the bubble and its collapse, but its mean and standard deviation are highly dispersed *across* different pixels. As a result, the sum of these lognormals turns out to be far from a lognormal distribution during this period. In other words, heterogeneity across pixels in terms of the mean and standard deviation is the source of the fat upper tail of the size-adjusted price distribution during the period of the bubble and its collapse.

The results shown above suggest the presence of submarkets; that is, the prices of houses belonging to a particular submarket are almost identical in the sense that they are drawn from the same lognormal distribution, but the prices of houses belonging to different submarkets differ as they are from different distributions. More importantly, the size of submarkets tends to be smaller in a bubble period than in normal times, implying
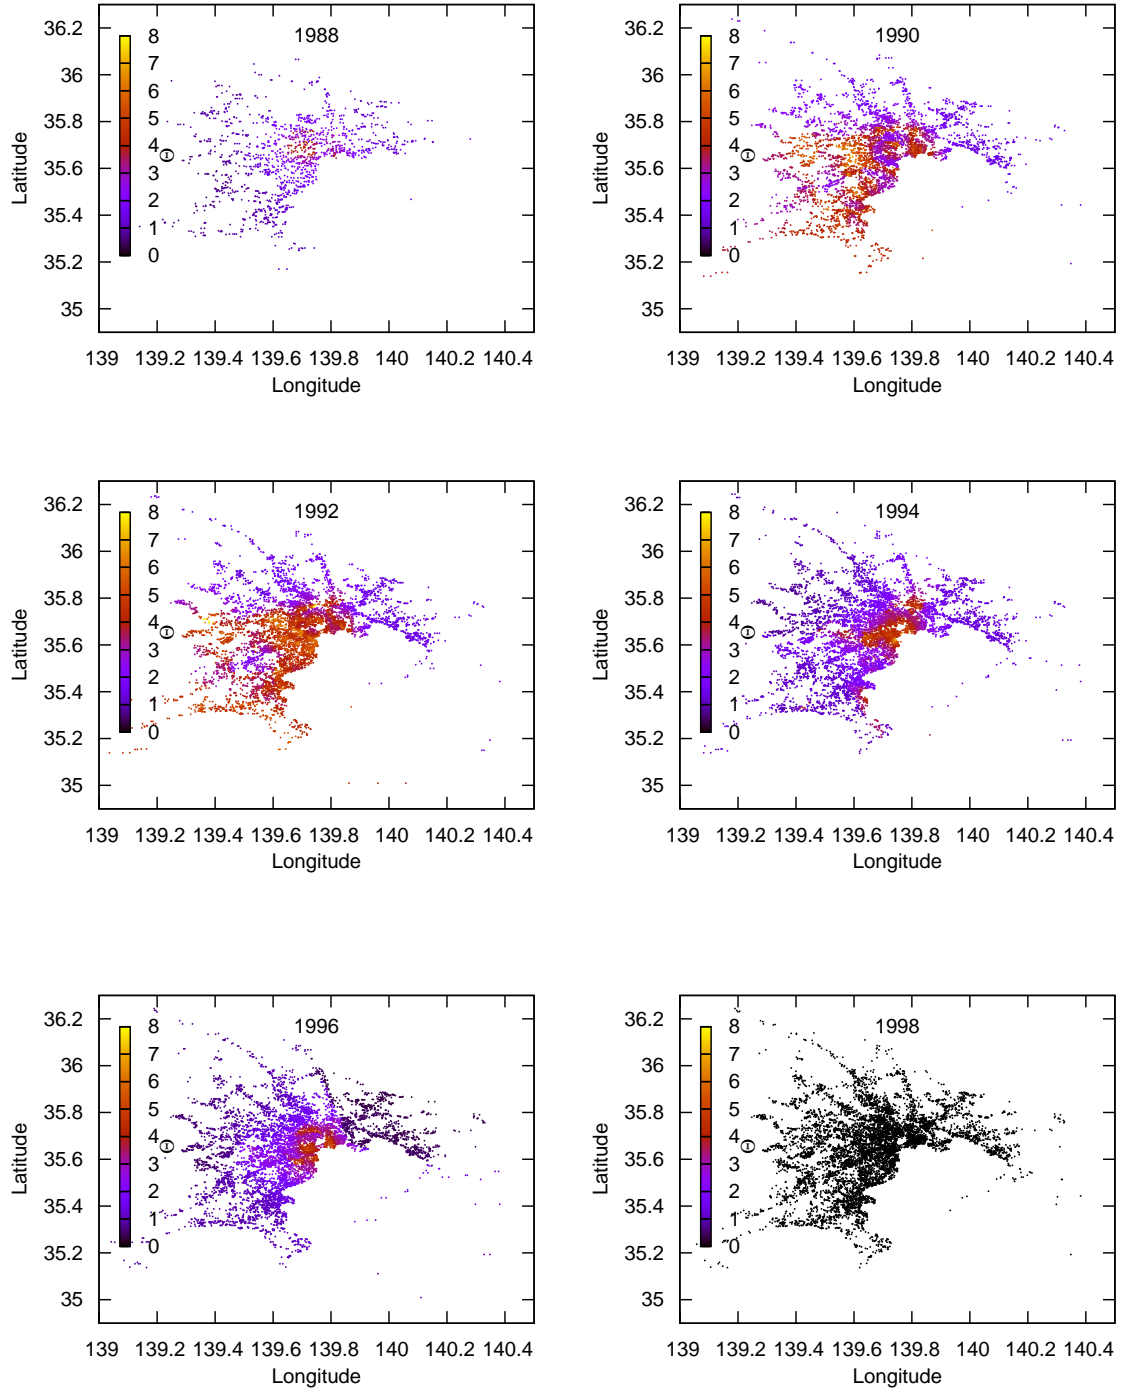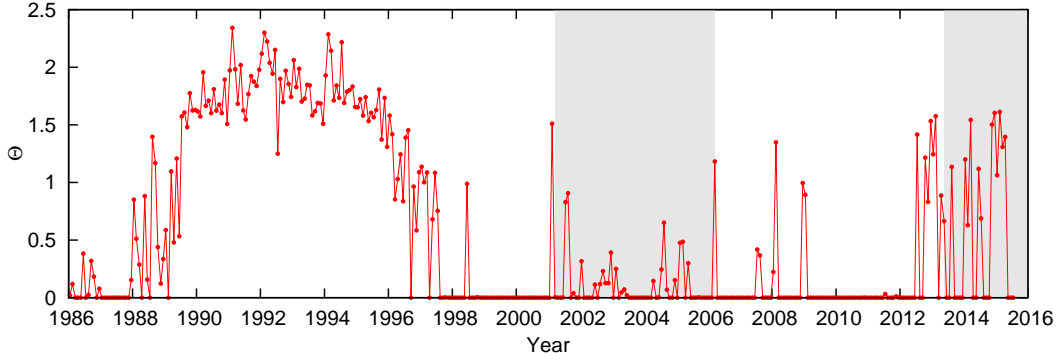
Figure 12: Estimate of the size of submarkets

Note: The first and second shaded areas, respectively, indicate the period of quantitative easing from March 2001 to March 2006 and the period of quantitative and qualitative easing since April 2013.

Figure 13: Evolution of the average size of submarkets

that market segmentation is more substantial in a bubble period.

As a final exercise, we estimate the size of submarkets in the following way. First, we collect the $n-1$ houses that are closest to a house i that was sold in month t. Second, we calculate the size-adjusted price for each of the n houses (i.e., house i and the closest $n-1$ houses) and conduct Geary's test to see whether the log of size-adjusted prices follows a normal distribution. We start this from a small value of n and repeat the exercise until we find n_{it}^* such that the null is rejected for values greater than n_{it}^*. Finally, we define the measure of the size of a submarket, Θ_{it}, as $\Theta_{it} \equiv -\log(n_{it}^*/J_t)$, where J_t is the number of houses that were sold in month t. Note that Θ_{it} is equal to zero if all houses in the Greater Tokyo Area belong to a single market and that the larger Θ_{it} is, the smaller is the size of a submarket. We calculate Θ_{it} for all houses that were sold in each month of the observation period.

Figure 12 shows the estimated value of Θ_{it} for each house sold in 1988, 1990, 1992, 1994, 1996, and 1998. The value of Θ_{it} was close to zero for most of the houses sold in 1988, which means that the size of submarkets was basically very large in that year. However, Θ_{it} exhibited a substantial increase in 1990, especially for houses located in the central part of Tokyo, suggesting that submarkets were of much smaller size than in 1988. This continued until 1998, when the value of Θ_{it} returned to zero for almost all houses sold in that year. Figure 13 shows how the average value of Θ_{it} across i, i.e., $J_t^{-1}\sum_i \Theta_{it}$, fluctuates over time. It turns out that the size of submarkets tended to be small in the bubble period, which is

29

not that surprising given the results shown in Figures 10 and 11. More interestingly, the size of submarkets also tended to be small in the period from March 2001 to March 2006, when the Bank of Japan pursued quantitative easing, and has also been small since slightly before April 2013, when the Bank of Japan introduced quantitative and qualitative easing.

6 Conclusion

In this paper, we found that the cross-sectional distribution of house prices in the Greater Tokyo Area has a fat upper tail and that the tail part is close to a power law distribution. On the other hand, the cross-sectional distribution of house sizes measured in terms of floor space has less fat tails than the price distribution and is close to an exponential distribution. We proposed a hedonic model consistent with these findings and, using data for Greater Tokyo, confirmed that size-adjusted prices follow a lognormal distribution except for the period of the asset bubble and its collapse, for which the price distribution remains asymmetric and skewed to the right even after controlling for the size effect. As for the period of the bubble and its collapse, we found evidence that the sharp price movements were concentrated in particular areas, and that this spatial heterogeneity is the source of the fat upper tail.

The analysis in this paper shows that the cross-sectional distribution of size-adjusted prices is very close to a lognormal distribution for normal times but deviates substantially from a lognormal for the bubble period. This suggests that the shape of the size-adjusted price distribution, especially the shape of the tail part, may contain information useful for the detection of housing bubbles. That is, the presence of a bubble can be safely ruled out if recent price observations are found to follow a lognormal distribution. On the other hand, if there are many outliers, especially near the upper tail, this may indicate the presence of a bubble, since such price observations are very unlikely to occur if prices follow a lognormal distribution. This method of identifying bubbles is quite different from conventional approaches based on aggregate measures of housing prices, which are estimated either by hedonic or repeat-sales regressions, and therefore should be a useful tool to supplement existing methods.

References

[1] Anundsen, A. K., K. Gerdrup, F. Hansen and K. Kragh–Sorensen (2016) "Bubbles and Crises: The Role of House Prices and Credit", *Journal of Applied Econometrics*, Vol. 31, No. 7, pp. 1291-1311.

[2] Bils, M. and Klenow, P. J. (2001) "Quantifying Quality Growth", *American Economic Review*, Vol. 91, No. 4, pp. 1006-1030.

[3] Blackwell, C. (2018) "Power Laws in Real Estate Prices? Some Evidence", *The Quarterly Review of Economics and Finance*, Vol. 69, pp. 90-98.

[4] Borio, C. and P. Lowe (2002) "Assessing the Risk of Banking Crises", *BIS Quarterly Review*, Vol. 7, No. 1, pp. 43-54.

[5] Chandola, V., A. Banerjee and V. Kumar (2009) "Anomaly Detection: A Survey", *ACM Computing Surveys*, Vol. 41, No. 3, pp. 15.

[6] Cochrane, J. H. (2005) "Stocks as Money: Convenience Yield and the Tech-Stock Bubble", in W. C. Hunter, G. G. Kaufman, and M. Pomerleano, eds., *Asset Price Bubbles: The Implications for Monetary, Regulatory, and International Policies*, MIT Press, pp. 175-203.

[7] Cropper, M. L., L. B. Deck and K. E. McConnell (1988) "On the Choice of Functional Form for Hedonic Price Functions" *Review of Economics and Statistics*, Vol. 70, No. 4, pp. 668-675.

[8] Del Castillo, J., and P. Puig (1999) "The Best Test of Exponentiality Against Singly Truncated Normal Alternatives", *Journal of the American Statistical Association*, Vol. 94, pp. 529-532.

[9] Diewert, W. E., J. de Haan and R. Hendriks (2011) "The Decomposition of a House Price Index into Land and Structures Components: A Hedonic Regression Approach", *The Valuation Journal*, Vol. 6, No. 1, pp. 58-105.

[10] Drehmann, M. and M. Juselius (2014) "Evaluating Early Warning Indicators of Banking Crises: Satisfying Policy Requirements", *International Journal of Forecasting*, Vol. 30, No. 3, pp. 759-780.

[11] Feller, W. (1968) *An Introduction to Probability Theory and Its Applications.* Vol. 1, Third Edition, New York: Wiley.

[12] Ferreira, F. and J. Gyourko (2012) "Heterogeneity in Neighborhood-level Price Growth in the United States, 1993-2009", *American Economic Review*, Vol. 102, No. 3, pp. 134-40.

[13] Ferreira, F. and J. Gyourko (2017) "Anatomy of the Beginning of the Housing Boom: US Neighborhoods and Metropolitan Areas, 1993-2009", November 2017.

[14] Gabaix, X. (2009) "Power Laws in Economics and Finance", *Annual Review of Economics*, Vol. 1, No. 1, pp. 255-294.

[15] Glaeser, E. L., J. D. Gottlieb and K. Tobio (2012) "Housing Booms and City Centers", *American Economic Review*, Vol. 102, No. 3, pp. 127-33.

[16] Greene, W. (2003) *Econometric Analysis*, Fifth Edition, Upper Saddle River, NJ: Prentice Hall.

[17] Gyourko, J., C. Mayer and T. Sinai (2013) "Superstar Cities", *American Economic Journal: Economic Policy*, Vol. 5, No. 4, pp. 167-99.

[18] Himmelberg, C., C. Mayer and T. Sinai (2005) "Assessing High House Prices: Bubbles, Fundamentals, and Misperceptions", *Journal of Economic Perspectives*, Vol. 19, pp. 67-92.

[19] Hoshi, T. and A. Kashyap (2000) "The Japanese Banking Crisis: Where Did It Come from and How Will It End?", *NBER Macroeconomics Annual 1999*, Vol. 14. MIT, 2000, pp. 129-212.

[20] Kaminsky, G. L. and C. M. Reinhart (1999) "The Twin Crises: The Causes of Banking and Balance-of-Payments Problems", *American Economic Review*, Vol. 89, No. 3, pp. 473-500.

[21] Kiyotaki, N. and J. Moore (1997) "Credit Cycles", *Journal of Political Economy*, Vol. 105, No. 2, pp. 211-248.

[22] Landvoigt, T., M. Piazzesi and M. Schneider (2015) "The Housing Market (s) of San Diego", *American Economic Review*, Vol. 105, No. 4, pp. 1371-1407.

[23] Määttänen, N. and M. Terviö (2014) "Income Distribution and Housing Prices: An Assignment Model Approach", *Journal of Economic Theory*, Vol. 151, pp. 381-410.

[24] Malevergne, Y., V. Pisarenko and D. Sornette (2011) "Testing the Pareto Against the Lognormal Distributions with the Uniformly Most Powerful Unbiased Test Applied to the Distribution of Cities", *Physical Review E*, Vol. 83, No. 3, 036111.

[25] Mankiw, N. G. and D. N. Weil (1989) "The Baby Boom, the Baby Bust, and the Housing Market", *Regional Science and Urban Economics*, Vol. 19, No. 2, pp. 235-258.

[26] Martin, A. and J. Ventura (2018) "The Macroeconomics of Rational Bubbles: A User's Guide", *Annual Review of Economics*, Vol. 10, pp. 505-539.

[27] McMillen, D. P. (2008) "Changes in the Distribution of House Prices Over Time: Structural Characteristics, Neighborhood, or Coefficients?", *Journal of Urban Economics*, Vol. 64, No. 3, pp. 573-589.

[28] Nishimura, K. (2011) "Population Ageing, Macroeconomic Crisis and Policy Challenges", Prepared for the Panel "The General Theory and the Policy Reponses to Macroeconomic Crisis" at the 75th Anniversary Conference of Keynes' General Theory, University of Cambridge, June 19-21, 2011.

[29] Okina, K., M. Shirakawa and S. Shiratsuka (2001) "The Asset Price Bubble and Monetary Policy: Japan's Experience in the Late 1980s and the Lessons", *Monetary and Economic Studies* (Special Edition), Vol. 19, No. 2, pp. 395-450.

[30] Piazzesi, M. and M. Schneider (2016) "Housing and Macroeconomics", in J. B. Taylor and H. Uhlig, eds., *Handbook of Macroeconomics*, Vol. 2, Elsevier, pp. 1547-1640.

[31] Rios-Rull, J. V. and V. Sanchez-Marcos (2008) "An Aggregate Economy with Different Size Houses", *Journal of the European Economic Association*, Vol. 6, No. 2-3, pp. 705-714.

[32] Rosen, S (1974) "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition", *Journal of Political Economy*, Vol. 82, No. 1, pp. 34-55.

[33] Shimizu, C., K. G. Nishimura and T. Watanabe (2010) "Housing Prices in Tokyo: A Comparison of Hedonic and Repeat-Sales Measures", *Journal of Economics and Statistics*, Vol. 230, No. 6, pp. 792-813.

[34] Shimizu, C. and T. Watanabe (2010) "Housing Bubbles in Japan and the United States", *Public Policy Review*, Vol. 6, No. 3, pp. 431-472.

[35] Shimizu, C. and T. Watanabe (2013) "Sticky Housing Rents and Monetary Policy: The Japanese Experience", Paper presented at the KDI Workshop on "Real Estate Driven Systemic Risk: Country Cases and Their Policy Implications", December 13-14, 2012, Seoul.

[36] Sinai, T. (2013), "House Price Moments in Boom-Bust Cycles", in E. L. Glaeser and T. Sinai, eds., *Housing and the Financial Crisis*, Chicago: University of Chicago Press, pp. 19-68.

[37] Takáts, E. (2012) "Aging and House Prices", *Journal of Housing Economics*, Vol. 21, No. 2, pp. 131-141.

[38] Van Nieuwerburgh, S. and P. O. Weill (2010) "Why Has House Price Dispersion Gone Up?", *Review of Economic Studies*, Vol. 77, No. 4, pp. 1567-1606.

[39] Zhang, L. and Y. Yi (2018) "What Contributes to the Rising House Prices in Beijing? A Decomposition Approach", *Journal of Housing Economics*, Vol. 41, pp. 72-84.