DISCUSSION PAPER SERIES

No. 7945

INCOME DISTRIBUTION AND HOUSING PRICES: AN ASSIGNMENT MODEL APPROACH

Niku Määttänen and Marko Terviö

LABOUR ECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: www.cepr.org/pubs/dps/DP7945.asp

ISSN 0265-8003

INCOME DISTRIBUTION AND HOUSING PRICES: AN ASSIGNMENT MODEL APPROACH

Niku Määttänen, Research Institute of the Finnish Economy (ETLA) and Helsinki Center for Economic Research
Marko Terviö, Aalto University, Helsinki Center for Economic Research and CEPR

Discussion Paper No. 7945
August 2010

Centre for Economic Policy Research
53–56 Gt Sutton St, London EC1V 0DG, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 7945

August 2010

ABSTRACT

Income Distribution and Housing Prices: An Assignment Model Approach*

We present a framework for studying the relation between the distribution of income and the distribution of housing prices that is based on an assignment model of households with heterogeneous incomes and houses of heterogeneous quality. The equilibrium distribution of prices depends on both distributions in a tractable but nontrivial manner. We infer the unobserved distribution of quality from the joint distribution of income and housing prices, and use it to generate counterfactual price distributions under counterfactual income distributions. We apply the model to estimate the impact of recently increased income inequality on the distribution of house prices in 6 US metropolitan regions. We find that the recent increase in income inequality caused average house prices to be lower. The impact of uneven income growth on house prices has been positive only within the top decile.

JEL Classification: D31 and R21
Keywords: assignment models, housing prices and income distribution

Niku Määttänen
ETLA
Lönnrotinkatu 4 B
FIN-00120
Helsinki

Email: niku.maattanen@etla.fi.

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=157639

Marko Terviö
Aalto University
PO Box 1210
00101 Helsinki
Finland

Email: Marko.Tervio@hse.fi

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=163017

* We thank Essi Eerola, Pauli Murto, Otto Toivanen, and Juuso Välimäki for useful suggestions. Määttänen thanks Suomen arvopaperimarkkinoiden edistämissäätiö and Terviö thanks the European Research Council for financial support.

Submitted 19 July 2010

1 Introduction

One question raised by the recent increases in income inequality is whether it has had an impact on the distribution of housing prices. Could it be that ever higher top incomes lead the rich to bid the prices of best locations ever higher? It has been argued that the increase in consumption inequality has been less than the increase in income inequality for essentially this reason. We present an assignment model framework to study this question and, in the process, also generate theoretical insights on the relation of income inequality and house prices.

In our view, from a distributional perspective, a central feature of the housing market is that housing is not a fungible commodity but comes embedded in indivisible and heterogeneous units. The supply of houses is more or less fixed in short and medium run, and in terms of the most desirable locations, also in the long run. Another key feature is that housing is a normal good that takes up a large part of household expenditure, so income effects should not be assumed away. Our modeling approach is based on an assignment model with non-transferable utility. The market consists of a population of households who each own one house and each wish to live in one house. Even though there are no complementarities in the usual sense, this setup results in positive assortative matching: wealthier households end up living in the higher-quality houses.

To focus on the impact of changes in the income distribution we assume that households have the same preferences: the only reason why the wealthy live in better houses is that they can better afford them. We show how the equilibrium distribution of house prices depends on the shapes of both the distribution of house quality and the distribution of income in a tractable although nontrivial manner. With this model, we are equipped to study questions such as the impact of increased income inequality on housing prices.

In our model the distribution of incomes and house quality are exogenous, while the distribution of house prices are endogenous (with the exception of the cheapest or "marginal" house). In general, the joint distribution of houses and income is arbitrary, which results potentially in a lot of trading between households. Equilibrium prices depend on the joint distribution of endowments, not just on the marginal distributions of income and house quality. With an arbitrary initial endowment distribution the equilibrium is quite complicated, and not useful for empirical purposes. However, we interpret the observed house prices as the equilibrium prices that emerge after all trading opportunities have been exploited. Under this "post-trade" assumption we can ask what distribution of unobserved house quality, together with the observed distribution of incomes, would give rise to the observed price distribution as the equilibrium outcome of our model. We find that a suitably parametrized CES utility function allows us to match the change in the price distribution under the assumption that it is caused by the change in incomes and a stable quality distribution. We then use the inferred distribution of house qualities and this preferred utility parametrization to generate counterfactuals that allow us to quantify the impact of the changes in income inequality on house prices.

In our empirical application we use data from all six metropolitan regions that are covered by the American Housing Survey (AHS) both in 1998 and 2007. We consider counterfactual income distributions for 2007 where all incomes grow uniformly since 1998 at the same rate as the actual mean income in the city. (I.e. the shape of the counterfactual distribution is the same as the actual shape in 1998). This counterfactual generates house prices that are on average $0-10\%$ higher, depending on the city. (Due to top coding, all results omit the top 3% of the price distribution). This implies that the increase in inequality has resulted in lower prices on average than would have prevailed under uniform income growth. The impact of uneven income growth on house prices has been positive only within the top decile, with magnitudes of up to 15%.

The reason why uniform income growth would have lead to higher prices at the bottom of the quality distribution is intuitive: if low-income households had higher incomes they would use some of it to bid for low-quality houses. However, in a matching market with positive sorting, any changes in prices spill upwards in the quality distribution. This is because the binding outside opportunity of any (inframarginal) household is that they must want to buy their equilibrium match rather than the next best house. The equilibrium price gradient—the price difference between two "neighboring" houses in the quality distribution—is pinned down by how much the households at the respective part of the income distribution are willing to pay for the quality difference. The price level of any particular house is given by the summation of all price gradients below, plus the price of the marginal house. This is why any local price changes spill upwards in the distribution, and downward pressure on prices from the bottom of the distribution counteracts the local increase in willingness-to-pay among better-off households whose incomes are now higher than under uniform growth. In principle, it is possible for all house prices to go down in response to an increase in inequality, although we don't find this to be the case in any of the six cities in our data.

In the next section we discuss related literature. In Section 3 we present the model, and in Section 4 we show how it can be used for inference and counterfactuals. Our empirical application is presented in Section 5, and Section 6 concludes.

2 Related Literature

Our model is an assignment model augmented with non-transferable utility. Assignment models are models of matching markets with symmetric information and no frictions; for a review see Sattinger (1993). Both sides of the market are assumed to have a continuum of types, so there is no room for "bargaining" as all agents have arbitrarily close competitors. Assignment models typically include an assumption of a complementarity in production, which results in assortative matching and equilibrium prices that depend on the shapes of the type distributions on both sides of the market but in a reasonably tractable way. Assignment models have usually been applied to labor markets, where the productive complementarity is between job types and worker types, as in Sattinger (1979) and Teulings (1995), or between workers themselves in a team production setting, as in Kremer (1993). In our setup there is no complementarity in the usual sense, but equilibrium nevertheless involves assortative matching by wealth and house quality, essentially because housing is a normal good. We don't restrict the shapes of the distributions, and our nonparametric method for inferring the unobserved type distribution and using it to construct counterfactuals is similar to Terviö (2008).

The closest existing literature to our paper is concerned with the dispersion of house prices between cities, while abstracting away from heterogeneity within cities. Van Nieuwerburgh and Weill (2010) study house price dispersion across US cities using a dynamic model, where there is matching by individual ability and regional productivity. Within each city housing is produced with a linear technology, but there is a city-specific resource constraint for the construction of new houses. This causes housing to become relatively more expensive in regions that experience increases in relative productivity. Houses are non-tradable across cities while labor is mobile, so intuitively this result is similar to the Balassa-Samuelson effect in trade theory. In their calibration Van Nieuwerburgh and Weill find that, by assuming a particular increase in the dispersion of ability, they can reasonably well generate the observed increase in wage dispersion and the (larger) increase in the house price dispersion across cities. Gyourko, Mayer, and Sinai (2006) have a related model with two locations and heterogeneous preferences for living in one of two possible cities. One of the cities is assumed to be a more attractive "superstar" city in the sense that it has a binding supply constraint for land. An increase in top incomes results in more bidding for the scarce land, thus leading the price of houses in the superstar city to go up. Moretti (2010) considers a two-city model with two types of labor, where changes in relative housing prices between cities can be affected by productivity (demand for labor) and amenities (location preference). Worker utility is linear but there is heterogeneity by location preference, in equilibrium the marginal worker within each skill group has to be indifferent between cities. Moretti finds that a fifth of the observed increase in college wage premium between 1980 and 2000 was absorbed by higher cost of housing, and that the most plausible cause for this is an increase in demand for high-skill workers in regions that attracted more high-skill workers.

Most of the dynamic macroeconomic models with housing assume that housing is a homogenous malleable good. In any given period, there is then just one unit price for housing. An exception is the property ladder structure that is used in Ortalo-Magne and Rady (2006) and Sanchez-Marcos and Rios-Rull (2008), where there are two types of houses: relatively small "flats" and bigger "houses". For our purposes, such a distribution would be far too coarse. In general, the macro literature focuses on the time series aspects of a general level of housing prices, and abstracts away from the cross-sectional complications of the market. By contrast, we focus on the cross-sectional and distributional aspects of the housing market, and abstract away from the time-series aspect.

One step in our empirical application is that we estimate the elasticity parameter of a constant elasticity of substitution utility function for housing and other consumption. This links our paper to a literature that uses structural models to estimate that parameter. Two very recent papers are Li, Liu and Yau (2009) and Bajari, Chan, Krueger and Miller (2010). These studies estimate the elasticity parameter within a life cycle model using household level data from the

US. However, as far as we know, we are the first to exploit changes in the distribution of housing prices to estimate household preference parameters. This is possible in our model because housing prices are in general a non-linear function of housing quality.

There is a long tradition in explaining heterogeneous land prices in urban economics, going back to Von Thünen (1826) and Alonso (1964). In urban economics models the exogenous heterogeneity of land is due to distance from the center. The focus in urban economics is on explaining how land use is determined in equilibrium, including phenomena such as parcel size and population density. In modern urban economics[1] there are also some models with income effects. Heterogeneity of land is modeled as a transport cost, which is a function of distance from center, and price differences between locations are practically pinned down by the transport cost function.

Models with heterogeneous land have been used in urban economics in connection with endogenous public good provision, in the tradition of Tiebout (1956). Epple and Sieg (1999) estimate preference parameters in a structural model where the equilibrium looks like assortative matching by income and public good quality, although the latter is a choice variable at the level of the community. Glazer, Kanniainen, and Poutvaara (2008) analyze the effects of income redistribution in a setup where heterogeneous land is owned by absentee landlords. They show that the presence of (uniformly distributed) heterogeneity mitigates the impact of tax competition between jurisdictions because taxation that drives some of the rich to emigrate also leads them to vacate high-quality land, allowing the poor to consume better land than before.

Matching models have long been applied to the housing market from a more theoretical perspective, although it is perhaps more accurate to say that housing has often been used in theoretical matching literature as the motivating example of an indivisible good that needs to be "matched" one-to-one with the buyers. The classic reference is Shapley and Scarf (1974), who present a model where houses are bartered by households who are each endowed with and each wish to consume exactly one house. They show that, regardless of the preference orderings by the households, there always exists at least one equilibrium allocation. Miyagawa (2001) extends the model by adding a second, continuous good, i.e., "money." He shows that the core assignment of houses can be implemented with a set of fixed prices for the houses. In

[1] See e.g. Fujita (1989).

Miyagawa's model utility is quasilinear, so there is no potential for income effects. The results obtained in this literature are not directly applicable in our setup, as we have both indivisible and continuous goods and utility is concave in the continuous good. There are many more papers on two-sided assignment, where two distinct classes of agents, "buyers" and "sellers," are matched, but these are less relevant for our setup where anyone who becomes a buyer must also become a seller.[2]

3 Model

We begin by introducing our setup in the context of an arbitrary initial endowment, which here consists of a house of a particular quality and a level of income for every household. We then restrict the possible endowments to "post-trade" allocations, that is, we assume that all mutually beneficial trades have already been made, and the role of equilibrium prices is merely to enforce the no-trade equilibrium. This latter case is not only tractable but also more useful empirically. Our interpretation of cross-sectional data is that, at current prices, each household wishes to live in its current house.

Consider a one-period pure exchange economy, where a unit mass of households consume two goods, housing and a composite good. Preferences are described by a utility function u, same for all households. Houses come in indivisible units of exogenous quality, and utility depends on the quality of the house, denoted by x. Preferences are standard: u is strictly increasing and differentiable, with diminishing marginal rates of substitution. Every household is endowed with and wishes to consume exactly one house. A household's income, denoted by θ, is interpreted as its endowment of the composite good y. There are no informational imperfections, or other frictions besides the indivisibility of houses.

A household endowment (x, θ) is described by a point in $[0,1] \times \mathbb{R}_+$, where in the horizontal dimension $i = F_x(x)$ represents the quantile in the distribution of house quality, and the vertical dimension represents the amount of composite good. As preferences are homogeneous, the same indifference map applies to all households. Figure 1 depicts this economy. An allocation is a joint distribution (of the unit mass of households) over the endowment space. Assume

[2] See, for example, Legros and Newman (2007).

that households are initially distributed smoothly over the endowment space so that there are no atoms and no gaps in the support of either marginal distribution, and both means are finite. The indivisibility of houses means that the resource constraint for x is rather stark: the marginal distribution of x cannot be altered by trading. For the continuous good, only the mean of consumption must match the mean of the endowments ($Ey = E\theta$) which is assumed positive but finite.

Equilibrium consists of a price function p for houses and a matching of households to houses; the composite good is used as the numeraire. Budget constraints are downward sloping curves, because house prices must be increasing in quality (by the monotonicity of u). Figure 1 depicts the budget curve of a household endowed with income $\tilde{\theta}$ and a house of quality $\tilde{x}$, it is defined by $\tilde{\theta} + p(\tilde{x}) = y + p(x)$, where the right side is the cost of consumption. We refer to the left side of the budget constraint as the household's wealth. Wealth is endogenous because it depends on the market value of the house that one is endowed with. However, households are atomistic, so from their point of view wealth and prices are exogenous.

All households with an endowment on the same budget curve trade to the point where the budget curve is tangent to an indifference curve. In terms of Figure 1, the resource constraint requires that the proportion of households with an endowment located below the budget curve that contains $\{i, \theta^*\}$ is equal to i, the proportion of houses that are of quality $x(i)$ or less. In general, the resource constraints lead to rather complicated integral equations, although, by discretizing the house types, the equilibrium can be solved numerically using standard recursive methods. However, we focus on the post-trade allocation, which simplifies the analysis considerably. In particular, we don't need to know whether an arbitrary initial endowment is associated with a unique equilibrium and whether it can be decentralized with a price function. What we need is the following lemma.

Lemma 1 *In equilibrium there is positive assortative matching (PAM) by household wealth and house quality.*

That is, in equilibrium, the ranking of households by wealth and by house quality must be the same. The proof is in the Appendix. In short, diminishing marginal rate of substitution guarantees PAM: of any two households, the wealthier must live in the better house, or else the two could engage in mutually profitable trade. The twist here is that the ordering by wealth is not known beforehand, because the value of the house is endogenous. So, despite PAM, the equilibrium allocation is not obvious and depends on the shape of the joint distribution of $\{x, \theta\}$. The benefit of Lemma 1 is that it guarantees that the equilibrium allocation is essentially one-dimensional, so we can index both households and houses by the house quality quantile i.

Lemma 2 *In equilibrium all households* $i \in [0, 1]$ *are located on a curve* $\{x(i), y^*(i)\}$ *in the endowment space that is continuous almost everywhere. If there are jumps they are upwards.*

This follows directly from Lemma 1: as wealth and therefore utility are increasing in i, downward jumps in y^* (as a function of i) are ruled out. Similarly, allocations supported over any thick region in endowment space would violate PAM. Only upward jumps in y^* are not ruled out, but there can only be a countable number of them or else y^* would not stay finite. Hence y^* is continuous almost everywhere. However, y^* does not have to be increasing; indeed, in Section 3.5 we will construct a (somewhat contrived) example where y^* is strictly decreasing.

The increasing curve in Figure 2 depicts the equilibrium allocations for a particular example. Households below the curve are the net suppliers of quality: they are endowed with a relatively high quality house and trade down in order to increase their consumption of the composite good. Households endowed with a house of quality $x(i)$ and income level $\theta(i) = y^*(i)$ do not trade. Assuming a full support $[x_0, x_1] \times [\theta_0, \theta_1]$ for the distribution of endowments, the end points of the equilibrium curve are necessarily $\{x_0, \theta_0\}$ and $\{x_1, \theta_1\}$, the endowments of the unambiguously poorest and richest households in this economy, who have either nothing to offer or gain in exchange.

We have now characterized what the allocation must look like after all trading opportunities have been exhausted. From now on we will restrict our analysis to this post-trade world. In a pure exchange economy, the post-trade allocation can be interpreted as "just another" endowment. For notational convenience we will be referring to this "endowment" of the composite good as θ.

3.1 Equilibrium price gradient

Suppose that all trading opportunities have been exhausted, so that the current allocation is an equilibrium allocation. Let's denote by $\theta(i)$ the observed allocation of composite good for owners of houses of quality $x(i)$. Now, by definition, equilibrium prices p must result in every household preferring to live in its own house, so that

$$i = \arg\max_{j \in [0,1]} u\left(x\left(j\right), \theta\left(i\right) + p(i) - p\left(j\right)\right) \tag{1}$$

holds for all $i \in [0,1]$. Since households are atomistic, they take p as given. When the associated first-order condition, $u_x x' - u_y p' = 0$, is evaluated at the optimal choice $j = i$ the prices cancel out inside the utility function. (That this optimum is global is guaranteed by Lemma 1.) Solving for p' we obtain an equation for equilibrium prices:

$$p'\left(i\right) = \frac{u_x\left(x\left(i\right), \theta\left(i\right)\right)}{u_y\left(x\left(i\right), \theta\left(i\right)\right)} x'\left(i\right). \tag{2}$$

This price gradient is *the* key equation of our model. Combined with the exogenous boundary condition $p\left(0\right) = p_0$ it can be solved for the equilibrium price function p. The boundary condition can be interpreted as the opportunity cost for the lowest-quality house, or as the reservation price for the poorest household stemming from some exogenous outside opportunity (such as moving to another housing market). The continuity of u and x implies that p is continuous.[3]

The intuition behind the price gradient (2) is that the price difference between any neighboring houses in the quality order depends only on how much the relevant households—at that particular quantile of the wealth distribution—are willing to pay for that particular quality difference. This depends on their marginal rate of substitution between house quality and other goods, which in general depends on the level of wealth. The price level at quantile i is the sum of the outside price $p\left(0\right)$ and the integral over all price gradients (2) below i. This is our next proposition.

Proposition 3 *Suppose θ is an equilibrium allocation. The equilibrium price function is then unique up to an additive constant p_0 and given by*

$$p\left(i\right) = p_0 + \int_0^i \frac{u_x\left(x\left(j\right), \theta\left(j\right)\right)}{u_y\left(x\left(j\right), \theta\left(j\right)\right)} x'\left(j\right) dj. \tag{3}$$

[3] If θ has a discontinuity, as is allowed by Lemma 2, then p has a kink.

Note that the equilibrium price at any quantile i depends on the distribution of housing quality and income at all quantiles below i. Hence changes at any part of the price distribution spill upwards but not downwards. Loosely speaking, in terms of a discrete setup, this asymmetry in the direction of price spillovers can be understood by considering the problems faced by the richest and poorest households. If the richest household were to get even richer this would have no implication on prices, as it would not make the second richest household willing to pay more for the best house. By contrast, were the poorest household to increase its income slightly (but so that it still remained the poorest), this would increase its willingness to pay for the second worst house, thus increasing the second poorest household's wealth and opportunity cost of living in its house. This, in turn, will increase the second poorest household's willingness to pay for the third worst house, and so on, causing the local price increase at the bottom keep spilling upwards in the distribution.

3.2 Absentee landlords: a digression

In urban economics models with heterogeneous land it is standard to assume that all land is initially owned by competitive outside sellers or "absentee landlords."[4] In our model the absentee landlord assumption can be introduced by assuming that the revenue from house sales goes to outside agents who are not buyers in this market. Consider the household at quantile i of the income distribution. Again, by Lemma 1, equilibrium must involve positive assortative matching by wealth, which here consists only of income θ, and house quality. Thus p must result in every household buying a house of the same quality rank as is their rank in the wealth distribution.

$$i = \arg\max_{j \in [0,1]} u\left(x\left(j\right), \theta\left(i\right) - p\left(j\right)\right) \quad \text{for all } i \in [0,1] \tag{4}$$

Here the price of the house actually chosen is not part of household wealth and so it does not cancel out of the price gradient. As a result, equilibrium prices are now defined as a nonlinear ordinary differential equation:

$$p'\left(i\right) = \frac{u_x\left(x\left(i\right), \theta\left(i\right) - p\left(i\right)\right)}{u_y\left(x\left(i\right), \theta\left(i\right) - p\left(i\right)\right)} x'\left(i\right). \tag{5}$$

[4] See, for example, Fujita (1989).

Combined with a boundary condition this can still be solved for the equilibrium price function p. The boundary condition must satisfy $p_0 \leq w(0)$, or else the poorest household cannot afford to live anywhere.

If all the houses were owned by an absentee monopolist, then the self-selection constraint (5) would still have to hold, but the seller could restrict the quantity sold. This would require the monopolist to be able to credibly commit to not selling the lowest quality houses up until some quantile $m \in (0,1)$. The outside opportunity of the buyers would then pin down the lowest price $p(m)$.

3.3 Comparative statics

In this section we analyze the comparative statics of equilibrium prices with respect to changes in income distribution. Here we assume that the economy begins in an equilibrium where there is a strictly positive relation between income θ and house quality. The purpose of this simplification is to streamline the discussion, as the changes in price distribution cannot cause the ranking of households by wealth to change.

It is worth noting that, in the light of our model, the claim that an increase in income inequality must lead to an increase in the prices of best houses is incorrect.

Proposition 4 *Suppose that the endowments form an equilibrium allocation, and that the income distribution experiences a mean-preserving and order-preserving spread, where incomes decrease below quantile $h \in (0,1)$ and increase above h. Then housing prices will either i) decrease everywhere or ii) decrease everywhere except at quantiles $(h', 1]$, where $h' > h$.*

Proof. Denote the new distributions by hats. By definition, the new income distribution satisfies

$$\hat{\theta}(i) < \theta(i) \qquad \text{for } i \in [0,h),$$
$$\hat{\theta}(i) > \theta(i) \qquad \text{for } i \in (h,1],$$
$$\int_0^1 \left(\hat{\theta}(j) - \theta(j)\right) \mathrm{d}j = 0.$$

Applying (3), the change in prices at any $i \in [0,1]$ is

$$\begin{aligned}\hat{p}(i) - p(i) &= p_0 + \int_0^i \frac{u_x(x(j), \hat{\theta}(j))}{u_y(x(j), \hat{\theta}(j))} x'(j)\,\mathrm{d}j - \left(p_0 + \int_0^i \frac{u_x(x(j), \theta(j))}{u_y(x(j), \theta(j))} x'(j)\,\mathrm{d}j \right) \\ &= \int_0^i \left(\frac{u_x(x(j), \hat{\theta}(j))}{u_y(x(j), \hat{\theta}(j))} - \frac{u_x(x(j), \theta(j))}{u_y(x(j), \theta(j))} \right) x'(j)\,\mathrm{d}j. \qquad (6)\end{aligned}$$

The inverse of the marginal rate of substitution between y and x, $u_x(x,\theta)/u_y(x,\theta)$, is increasing in θ, and $x' > 0$, so the integrand in (6) is negative at all j where $\hat{\theta}(j) < \theta(j)$, i.e., for $j < h$. Similarly, the integrand is positive for $j > h$. The definite integral in (6) must therefore be strictly negative at $i = h$, where it reaches its minimum, and increasing above h. If, at some $h' > h$ the definite integral reaches zero then it will be positive at all $i > h'$, but it might not reach zero before $i = 1$, in which case the price change is negative at all $i \in (0,1]$. ■

Intuitively, if income is redistributed from poor to rich, this will increase the local price gradient (2) at the top quantiles, as the willingness-to-pay for extra quality goes up for the rich. But, for the same reason, the price gradient at bottom quantiles goes down. Above h, the change in price gradient is positive, but the negative spillover from below will dominate until some $h' > h$, and it can be that the cumulative impact of positive gradients is not enough to overtake the negative impact. It is therefore possible for all house prices to go down in response to an increase in inequality.

The impact of a simple increase in income levels is perhaps less surprising.

Proposition 5 *Suppose the endowments form an equilibrium allocation. If all incomes rise in an order-preserving manner, then housing prices will increase at all quantiles* $i \in (0,1]$ *and this increase is increasing in the quantile* i.

The reasoning is similar as in the Proof of Proposition 4, but even simpler because the new price gradient is greater than the old gradient at all quantiles. As an immediate corollary, an increase in income levels will also increase the variance of house prices.

3.4 The case with CES

For the empirical application we assume CES utility,

$$u(x,y) = (\alpha x^{\rho} + (1-\alpha) y^{\rho})^{\frac{1}{\rho}}, \quad \text{where } \rho < 1 \text{ and } \alpha \in (0,1), \tag{7}$$

with Cobb-Douglas utility defined in the usual fashion at $\rho = 0$. If wealth is exogenous (as in the absentee landlord setup), then the equilibrium price gradient (5) is

$$p'(i) = \frac{\alpha}{1-\alpha} \left(\frac{w(i) - p(i)}{x(i)} \right)^{1-\rho} x'(i). \tag{8}$$

Assuming the endowments form an equilibrium allocation, so that $w(i) = \theta(i) + p(i)$, this can be solved as

$$p(i) = p_0 + \frac{\alpha}{1-\alpha} \int_0^i \left(\frac{\theta(s)}{x(s)} \right)^{1-\rho} x'(s)\, \mathrm{d}s. \tag{9}$$

Example: Pareto distributions Equilibrium prices have a closed-form solution in our model only under specific assumptions. Our empirical applications do not require such a solution, but they are useful for theoretical extensions. Here we present an example with a closed-form solution.

Assume that both income and house quality follow Pareto distributions, so that θ˜Pareto(θ_0, η) and x˜Pareto(x_0, γ), where $w_0 > p_0$ and $\eta, \gamma > 2$. The associated quantile functions are $\theta(i) = \theta_0 (1-i)^{-\frac{1}{\eta}}$ and $x(i) = x_0 (1-i)^{-\frac{1}{\gamma}}$. If utility takes the CES form (7) then, under the post-trade interpretation, the equilibrium prices (9) can be solved in closed form:

$$p(i) = p_0 + \frac{\alpha}{1-\alpha} \frac{\xi}{\gamma} \theta_0^{1-\rho} x_0^{\rho} \left((1-i)^{-\frac{1}{\xi}} - 1 \right) \tag{10}$$

where $\xi \equiv \frac{\eta}{1-\rho(1-\eta/\gamma)}$. This means that prices are distributed according to a Generalized Pareto Distribution. The expenditure share of housing $a(i) \equiv p(i) / (p(i) + \theta(i))$ can be shown to have a limiting value at 1 if $\rho > 0$ and at 0 if $\rho < 0$. In the knife-edge Cobb-Douglas case it is

$$a(1) = \frac{\alpha\eta}{\alpha\eta + (1-\alpha)\gamma}. \tag{11}$$

The expenditure share a is then everywhere increasing if $a(1) > a_0 \equiv p_0 / (p_0 + \theta_0)$ (decreasing if $a(1) < a_0$). If housing at the extensive margin can be created at constant marginal cost then the poorest household faces in effect linear prices and it is reasonable to assume that $a(0) = \alpha$.

In this case the expenditure share of housing is strictly increasing in income if and only if $\eta > \gamma$, i.e., when the variance of wealth is lower than the variance of house quality.

This example illustrates how one cannot expect the expenditure shares to be constant across income levels, even if utility function takes the Cobb-Douglas form. The expenditure share of housing is not directly given by preferences because the prices faced by the consumers are nonlinear. The standard CES result that expenditure shares are independent of income is based on all goods being fungible, so that there is essentially just one type of housing.

Example: Degenerate wealth distribution Suppose all households have the same wealth level $\bar{w}$ and preferences are Cobb-Douglas. Now prices must make every household indifferent between every housing unit. Then (8) reduces to

$$p(i) = \bar{w} - (\bar{w} - p_0) \left(\frac{x_0}{x(i)} \right)^{\frac{\alpha}{1-\alpha}} \tag{12}$$

where $\bar{w} > p_0$ must be assumed. This (admittedly contrived) example provides the simplest demonstration for why the expenditure share of housing cannot be expected to be constant even if preferences are the same for all – some households simply must end up with the lower quality houses and for this they must be compensated with higher consumption of other goods.

3.5 Planner's problem: another digression

Suppose a planner decided on the allocation, with the objective of maximizing average utility. (Equivalently, suppose the households were allowed to agree on the allocation behind a veil of ignorance.) The fixed distribution of house quality forces the planner to impose the unequal distribution of x on the households. The planner's problem consists of dividing the total endowment of the composite good between the households according to some positive function y.

$$\max_{y \geq 0} \int_0^1 u(x(i), y(i))\, \mathrm{d}i \text{ st. } \int_0^1 y(i)\, di = \bar{\theta}. \tag{13}$$

Pointwise maximization reveals that it is optimal to equate the marginal utility of the composite good across households

$$u_y(x(i), y^*(i)) = \lambda \quad \text{for all } i, \tag{14}$$

where the constant $\lambda > 0$ is determined by the resource constraint. By differentiation of (14) we see that

$$dy^*/dx = -u_{xy}/u_{yy}, \tag{15}$$

so if $u_{xy} < 0$ then those who are given better houses are given less money. In such a case the planner's allocation would also be an example of an equilibrium allocation where there is a strictly negative relation between non-housing consumption and house quality.

Example. CES-utility.

For CES utility $-u_{xy}/u_{yy} = y/x$, so, by (15), the solution must be of the form

$$y^* = kx. \tag{16}$$

The constant $k > 0$ is determined by the resource constraint:

$$k = \bar{\theta} / \int_0^1 x(i)\mathrm{d}i. \tag{17}$$

Under CES utility the goods are complementary in the sense that, behind a veil of ignorance, an individual would prefer to allocate more money for the states of the world where she has a better house (this already follows from $u_{xy} > 0$).

The planner's solution under CES preferences is intuitively unappealing, because CES exhibits risk neutrality with respect to wealth. In a setup with aggregation over states of the world a more general utility function is needed.

Example. CRRA-CES utility is defined as $v(x,y) = \frac{1}{\psi}u(x,y)^{\psi}$, where u is the ordinary CES-utility (7) and $\psi \leq 1$ captures relative risk aversion. The expression (15) is now, after some simplification,

$$dy^*/dx = \frac{\alpha x^{\rho-1} y^*}{\alpha(1-\rho)x^{\rho} + (1-\alpha)(1-\psi)\left(y^*\right)^{\rho}} \left(\psi - \rho\right). \tag{18}$$

Thus $dy^*/dx > 0$ if and only if $\psi > \rho$. Under CRRA-CES utility there is tension between the complementarity of consumption between the two goods, which is decreasing in ρ, and between the risk aversion, which is decreasing in ψ. The complementarity drives the planner to allocate more non-housing consumption to households who get better houses, while risk aversion drives to the opposite direction. When complementarity dominates ($\rho < \psi$), as is necessarily the case under standard CES ($\psi = 1$) the planner will allocate more money to the lucky recipients of the better houses.

Naturally, risk aversion over wealth would make no difference in our main model as there is no uncertainty. Any positive monotone transformation of u cancels out of individual optimization and market equilibrium conditions, leaving (1) and (2) unaffected.

4 Inference about the quality distribution

In the empirical application we assume that the observed prices correspond to the equilibrium prices that emerge after all trading opportunities have been exhausted. We think this is a reasonable interpretation of cross-sectional data because, at any point in time, only a small fraction of households trade houses. The model can then be used to infer the unobserved distribution of x from the observed relation between household income and house prices. For given distributions p and θ, and for a given utility function u, we can infer the distribution of x up to a constant. (We return to the inference of preferences in the next section.) This can be done by treating x as the unknown in the differential equation (2), while normalizing the boundary condition $x(0) = x_0$ at any positive value. The levels of x are not themselves interesting (they only have meaning in the context of the model), but they are needed to generate a number of interesting counterfactuals, for which the normalization of x_0 will be without loss of generality.

For our empirical application we discretize the model and assume CES utility. In each market, we have data on prices (house values) $p_0 < p_1 < \cdots < p_N$ and the associated incomes θ_h. The minimal requirement for the data to be consistent with the model is that observed wealth $p + \theta$ must be in strictly positive relation with p (Lemma 1). This relation is, of course, not perfect in micro data, but it emerges very naturally in our data with some smoothing (details will be explained later).

We know that equilibrium must involve positive assortative matching. Intuitively, the basis for inference is the "incentive compatibility" condition that makes household h want to buy its equilibrium match—which is house h—instead of any other house h'. For CES utility, these conditions are

$$\left(\alpha x_h^\rho + (1-\alpha)\left(\theta_h\right)^\rho\right)^{\frac{1}{\rho}} \geq \left(\alpha x_{h'}^\rho + (1-\alpha)\left(\theta_h + p_h - p_{h'}\right)^\rho\right)^{\frac{1}{\rho}}, \text{ for all } h, h' \in \{0, \ldots, N\}. \tag{19}$$

Thanks to PAM, we can ignore this constraint for all other household pairs except for those who are "neighbors" in the rank by house quality, ($h' = h - 1$). For convenience, we assume that these constraints hold as an equality.[5] The associated equality can then be solved for x_h to yield the inference formula

$$x_h = \left(x_{h-1}^\rho + \frac{1-\alpha}{\alpha}\left[\left(\theta_h + p_h - p_{h-1}\right)^\rho - \left(\theta_h\right)^\rho\right]\right)^{\frac{1}{\rho}}. \tag{20}$$

Denoting $\hat{x} = x^\rho$ this can be solved for

$$\hat{x}_h = \hat{x}_0 + \frac{1-\alpha}{\alpha}\sum_{j=1}^{h}\left[\left(\theta_j + p_j - p_{j-1}\right)^\rho - \left(\theta_j\right)^\rho\right] \tag{21}$$

which includes an undefined constant of integration $\hat{x}_0$, i.e. the quality of the worst occupied house. Note that when we infer x the value of $\alpha \in (0,1)$ is without consequence for all observable variables. This is because, in the utility function, changing α is equivalent to changing the units of x. Since x is not observed but only inferred within the model, this merely changes the scale of measurement for x, and has no impact on the monetary counterfactuals that we are interested in.

The purpose of inferring the distribution of an abstract quality unit is that it allows us to construct interesting counterfactuals. Having inferred the distribution of x based on the observed distributions of θ and p, we can then posit a counterfactual income distribution $\hat{\theta}$ and generate the implied counterfactual distribution of house prices, by combining $\hat{\theta}$ and x in the (discrete equivalent of) equilibrium price relation (9). Note, however, that as p_0 is exogenous, our model only explains the differences in prices relative to the marginal unit of housing, $p - p_0$. In the counterfactuals that follow the lowest price is always taken to be the lowest price in the data.

[5]There is, in the discrete setup, a match-specific rent that the "neighbors" could bargain over. Here we in effect assume that "sellers" have all bargaining power. Making the opposite assumption makes virtually no difference to the empirical results.

5 Empirical application

5.1 Data and smoothing

We use income and price data from the American Housing Survey (AHS) for six metropolitan areas (MA): Baltimore, Boston, Houston, Minneapolis, Tampa, and Washington. The choice of MAs was strictly determined by the availability of data. The most populous MAs are covered only in national-level data (the National Surveys), where house prices are top coded at a common national threshold. This results in a disproportionate censoring of observations in the largest MAs. Furthermore, in those MAs where the national data is not decimated by top coding it suffers from small sample size, which unfortunately renders the national surveys useless for our purposes. Fortunately, certain mid-size cities are covered in separate "Metro surveys," which have been conducted at irregular intervals, but are less affected by the top-coding problem. Our sample includes all MAs from which there is metro survey data both in 2007 (the most recent year) and 1998 (the second most recent year for the MAs surveyed in 2007).

Our income measure is total disposable income, including taxes and transfers, during the last year. House price is based on the survey question where respondents were asked to estimate the current market value of their house. We consider only homeowners, which amounts to assuming that rental housing forms an entirely separate market.

In the AHS metro survey data, the house prices have been censored separately at each MA at the 97% percentile. Thus we have to drop the top 3% from our analysis. Furthermore, there are apparently significant data quality issues at the bottom of the price distributions, with many house prices observed in the range of a few hundred or thousands of dollars. For this reason we drop the observations with the lowest house prices; in the benchmark case we drop the bottom 5%. In the model the lowest price is always exogenous; so here the 5th percentile price will be exogenous to our analysis. We also illustrate how some of the results change if we instead drop the bottom 15%.

We observe the joint distribution of income θ and house price p in each of the six MAs for both 1998 and 2007. To be consistent with the equilibrium of our model, the levels of observed wealth $w = \theta + p$ should be perfectly rank correlated with house value p (recall Lemma 1). To achieve this, we first reduce the relation of income and house price into a curve, by using kernel regression to estimate $\theta(i)$ as $E[\theta|F_p(p) = i]$, where F_p is the empirical CDF of p. Figures 3 and 4 display the actual and kernel smoothed incomes relative to the quantile by house price. There is a strong positive relationship between house value and income.[6] We denote the distributions for years $t \in \{98, 07\}$ by $\{\theta^t, p^t\}$, where θ^t is smoothed and p^t is raw data.

Table 1 displays Gini coefficients for our estimated permanent income (i.e. smoothed income), annual income, and house value. All MAs feature an increase in income inequality from 1998 to 2007. Housing values show no such systematic pattern. Figure 5 displays the distributions of smoothed income relative to its mean in 1998 and 2007 on log scale. The only apparently exceptional case is Houston where the lowest incomes have grown faster than the mean.

Gini coefficients	Permanent income		Income		House value	
	1998	2007	1998	2007	1998	2007
Baltimore	0.16	0.17	0.37	0.40	0.27	0.28
Boston	0.13	0.17	0.36	0.44	0.25	0.23
Houston	0.18	0.20	0.39	0.41	0.30	0.27
Minneapolis	0.14	0.16	0.33	0.36	0.23	0.23
Tampa	0.19	0.23	0.41	0.52	0.29	0.31
Washington	0.14	0.17	0.33	0.39	0.25	0.23

Table 1. Patterns of inequality.

5.2 Inferring preferences

We now impose the CES-utility function specified in (7) and infer the elasticity parameter ρ. We infer ρ both separately for each MA and restricting it to be the same across all MAs. Given ρ and $\{\theta^{98}, p^{98}\}$, we first infer the quality distribution in 1998, denoted by x^{98}, using the inference formula in (21). We also need to set an interest rate r to make the units of yearly income compatible with the house price: θ is measured as annual income divided by the interest rate.

[6]We use the Epanechnikov kernel and a bandwidth of 9%, except in Tampa where a bandwidth of 11% is required for the smoothed data to conform to PAM by wealth and house price.

We fix the interest rate at $r = 0.05$.[7] Changing r over a reasonable range ($2 - 8\%$) makes little difference to our inferred elasticity parameter ρ or the results from the counterfactual experiments that we present below.[8]

Having inferred x^{98} from 1998 data, we then use the equilibrium price formula—the discretized equivalent of (9)—to predict the 2007 housing price distribution given actual observed wealth in 2007 ($\theta^{07} + p^{07}$) and x^{98}. In other words, assuming that the quality distribution x is fixed, we ask what would be the predicted price distribution in 2007 given the observed 2007 income distribution. We denote this prediction by $\hat{p}^{07}$. The price of the lowest quality house (5th percentile in the benchmark case) is set equal to the actual value, again because the model does not explain the absolute price level but rather the difference over the lowest quality house. We also have to implicitly assume that the ranking of households by income has not changed, as we are applying a static model in both years.

We compare the model's predicted price distribution to the empirical 2007 distribution for a range of values for ρ, where each comparison involves the entire procedure described above. Figure 6 illustrates this for the cases of Boston and Tampa. It shows the observed 2007 housing price distribution as well as the predicted housing price distributions under various values for ρ. Our preferred elasticity parameter is the one where the mismatch between the empirical and predicted 2007 housing price distribution is the smallest. Formally, we pick ρ in order to minimize the mean of absolute percentage errors (MAPE), i.e. the mean value of $|\log(p^{07}(i)) - \log(\hat{p}^{07}(i|\rho))|$. In the case where ρ is restricted to be the same in all MAs, it is chosen by minimizing the unweighted average of MAPE across the six MAs.

Table 2 shows for each MA the number of observations by year, the inferred value for the elasticity of substitution, and the minimized value of MAPE. The last row corresponds to the case where the elasticity parameter is restricted to be the same in all MAs, the result is 0.50.

[7] Dividing the annual income by the annual interest rate is to assume that households face an infinite time-horizon. Alternatively and equivalently, we could multiply the house price with the annual interest rate. None of the results we present would change.

[8] This is true even if we assume that the interest rate is different in 1998 and in 2007. Intuitively, this is because inference is based on the distributions of income and housing prices, not just their average levels. By contrast, measuring the proper interest rate is crucial if one tries to estimate demand elasticities by using time variation in some aggregate house price index.

This is in the middle of the MA-specific elasticities which are also relatively close to each other, ranging from 0.32 to 0.89. (We discuss the meaning of "elasticity" in our context below.) The measure of fit varies substantially over different MAs. The fit is clearly better in Boston, Houston, and Minneapolis than in Tampa or Washington. Figure 7 shows the relative price error, $\log(p^{07}(i)) - \log(\hat{p}^{07}(i))$, by quantiles. The solid line corresponds to the benchmark case, where the cheapest 5% of houses are dropped, and the dashed line to the case where the lowest 15% are dropped. Apparently there is something going on in the upper tail of the distribution in Tampa and Washington that is not captured by our model. According to the model, the prices of the best houses should have increased much more there than they actually did. However, in most cities, our simple model seems to provide a reasonable fit.

	N_{98}	N_{07}	$1/(1-\rho)$	MAPE
Baltimore	2306	1211	0.42	4.30
Boston	1960	1075	0.63	2.79
Houston	1963	1105	0.89	2.29
Minneapolis	2741	1524	0.48	2.61
Tampa	2106	1255	0.49	8.14
Washington	2356	1315	0.32	9.13
Common ρ	13432	7485	0.50	12.18

Table 2. Number of observations by year, estimated elasticity of substitution, and the minimized mean absolute percentage error (MAPE) for the predicted price in 2007.

5.3 On income and price elasticities

The inferred elasticities in Table 2 are in line with many studies that use household data (see e.g. Li et al., 2009, and the references therein). However, as always in structural estimation, the interpretation of the parameters depends on the specifics of the model. Our estimates are not directly comparable with those obtained in studies that do not take into account the friction arising from the indivisibility of houses. In our setup, the effective income and price elasticities of housing demand are not determined solely by assumptions about preferences, because prices are a nonlinear function of quality. This nonlinearity also implies that income and price elasticities of housing expenditure will vary across income levels (despite CES utility).

We define the income elasticity of housing demand at quantile i using the counterfactual increase in housing expenditure that would result if a single household were to alone experience a change in income. Following a $\triangle$ percent change in income at quantile i, the household will reoptimize its consumption. The housing expenditure of a household at quantile i changes from $p(i)$ to $p(j)$, where

$$j = \arg\max_{s\in[0,1]} (u(x(i), (1+\triangle)\theta(i) + p(i) - p(s)) \tag{22}$$

is the household's new quantile in the distribution of wealth and housing quality, given that everyone else's incomes stay the same.

Figure 8 shows the effective income elasticity of housing expenditure in each MA. For each quantile i, the elasticity is computed as the midpoint arc elasticity around $\theta^{07}(i)$ with a 10% income change.[9] Elasticity is drawn only up to the point where the income increase shifts households above the top-coding threshold. The Figure reveals how income elasticity varies quite substantially over the distribution.[10] Intuitively, how much more one is willing to spend on housing following an increase in income depends on the increase in housing quality that would be available. In this model, the increase in housing quality that one extra dollar would buy varies over the distribution as it depends also on the incomes of competing buyers.

Similarly, we can also consider the price elasticity of housing expenditures by assuming that housing prices change proportionally over the whole distribution. The resulting change in housing expenditure at quantile i is computed by letting the household at i reoptimize their consumption under the new prices

$$j = \arg\max_{s\in[0,1]} (u(x(i), \theta(i) + (1+\triangle)(p(i) - p(s))). \tag{23}$$

[9]Using a smaller Δ results in otherwise similar but more "erratic" elasticity curves. This is because the distribution of prices has not been smoothed, so small changes in i can result in large changes in $p(i)$.

[10]The income elasticity goes up to 12 for the first quantiles in Tampa, where the price gradient is extremely steep at the bottom of the distribution (where we suspect data quality issues).

The proportional change in housing prices has the same effect on relative prices as a proportional change in incomes. The resulting price elasticities (not shown) roughly mirror the income elasticities (they have the opposite sign). However, the definition of price elasticity is problematic in this context where prices are endogenous, as the "counterfactual" change in prices in (23) is not consistent with the optimizing behavior of other households in the model.

Finally, we can also derive the "housing wealth effect", or the elasticity of non-housing consumption with respect to housing prices using (23). The change in housing prices changes the consumption of a household at quantile i from $\theta(i)$ to $\theta(i) + (1 + \triangle)(p(i) - p(j))$. The resulting elasticities (not shown) are positive but on average less than 0.2. An increase in housing prices increases household wealth and decreases the relative price of non-housing consumption. These effects work to increase non-housing consumption. However, part of the increase in household wealth is spent on housing. Indeed, given the relatively low estimated elasticity of substitution parameters (see Table 2), an increase in housing prices makes households increase their expenditure share for housing. This effect limits the increase in non-housing consumption.

The observation that these elasticities vary substantially over the distribution suggest that estimating household preferences without taking into account the indivisibility and heterogeneity of houses is problematic. For instance, the effects of a given change in aggregate income on aggregate housing demand may depend heavily on the associated changes in the distribution of income. This is also illustrated by the following counterfactual experiments.

5.4 Counterfactuals

The key question we set out to ask was how changes in income distribution influence housing prices. We now apply our methodology to answer this question for our specific data set and compute the impact of increased income inequality between 1998 to 2007 on housing prices. Specifically, given the inferred values for ρ (see Table 2) and the quality distributions x^{98} in each MA, we compute the predicted equilibrium price distributions in 2007 under a counterfactual income distribution that has the same shape as it did in 1998 but the same mean as in 2007. (The counterfactual income distribution is also assumed to preserve the ranking of households by income.) We then compare this counterfactual price distribution with the fitted price distribution that is obtained by plugging in the actual 2007 income distribution; the difference between the two is the impact of increased inequality.

Results are shown in Figure 9. It displays the relative difference in housing prices that follows the counterfactual income distribution and the empirical one. The effect is qualitatively similar in all but one MA: the increase in income inequality has lowered housing prices until about 80-90th percentile and increased prices in the upper tail, the average impact being negative. The exception is Houston, where the change near the bottom of the income distribution works to increase the prices of the lowest quality houses. This reflects the fact, visible in Figure 5, that in Houston, unlike in the other MAs, the lowest incomes have increased relative to the mean income, even though the Gini coefficient has increased as in other MAs.

The first two columns in Table 3 display the mean effect of the change in income distribution both in relative terms and in absolute terms. The third column corresponds to the case where ρ is restricted to be the same in all MAs. The mean effect is negative in all MAs except for Houston where it is slightly above zero. Intuitively, the increase in income inequality results in lower incomes at the bottom of the distribution relative to the counterfactual of uniform income growth. This works to lower the prices of the lowest quality houses. However, as explained in Section 3.3, this negative price impact at the bottom spills upwards in the quality distribution. This negative effect counteracts the local increase in willingness-to-pay among the lower part of the better-off households who in actuality saw their incomes rise faster than the mean.

More speculatively, we also consider the impact of the overall income inequality on housing prices by computing the equilibrium prices (again, given the same estimated quality distribution x^{98}) while assuming that all households have the mean income observed in 2007. Figure 10 displays the effect on the entire distribution of housing prices and the last two columns in Table 3 display the mean effects. According to the model, overall income inequality lowers housing prices by $19 - 52\%$ relative to the counterfactual of absolute income equality.

	Change (%)	Change ($)	Change (%), common ρ	Total (%)	Total ($)
Baltimore	-3.9	-14,300	-3.6	-43	-262,000
Boston	-4.3	-22,600	-5.9	-20	-122,000
Houston	+0.2	+3,100	+0.6	-19	-34,700
Minneapolis	-2.4	-7,090	-2.5	-26	-100,000
Tampa	-9.6	-25,800	-9.9	-52	-260,000
Washington	-6.7	-37,600	-4.0	-47	-464,000

Table 3. Change: The impact of the change in the shape of the income distribution on house prices between 2007 and 1998, relative to what house prices would have been under uniform income growth. Total: The impact of income inequality on house prices in 2007, relative to what house prices would be under equal income.

6 Conclusion

We have presented a new framework for studying the relationship between income distribution and housing price distribution. The key element of the model is that houses are heterogenous and indivisible, and that housing is a normal good. The equilibrium is tractable but yields a number of nontrivial implications. House prices, even at the top, can in principal go either way in response to an increase in income inequality.

The equilibrium conditions of our model allow us to estimate the housing quality distribution without imposing restrictions on the shapes of the distributions. We make a number of further assumptions in order to interpret the cross-sectional data on income and house prices in the context of our model. As our main empirical application we analyze housing prices in six US metropolitan regions in 1998 and 2007. We obtain a theory-driven estimate for the impact of increased income inequality by means of constructing counterfactuals. Specifically, we asked how the housing price distribution would differ from the actual price distribution if the income distribution in 2007 were different, namely if only the mean had changed but inequality had not increased since 1998.

Our model lends itself to several other applications in future work. Most directly, it is a natural framework for studying how various housing subsidy schemes impact housing prices. The intuition developed in the context of the price gradient already suggests that housing subsidies targeted for the poor will not merely be capitalized into prices of low-quality housing but will spill necessarily upwards in the quality ladder. A serious analysis of this issue will require us to extend the model to incorporate non-owner-occupied housing. Another related question is the estimation of the impact of rent control and the repeal thereof, which would result in a simultaneous supply and demand shock in a local housing market. Again, the impact on the entire distribution of housing prices is likely to be nontrivial. Further development to incorporate intertemporal and multi-region perspectives also offers some promising directions.

7 Appendix

Proof of Lemma 1, Positive assortative matching (PAM) by wealth and house quality.

Define the reservation price $\Delta(x'|x,y)$ by

$$u(x,y) = u(x', y - \Delta), \tag{24}$$

so that Δ is the maximum price that a household with endowment $(x,y) \in \mathbb{R}^2_+$ is willing to pay to switch into a house of type x'. Thus $\{x', y + \Delta(x'|x,y)\}$ traces an indifference curve that goes through the endowment (x,y); it is strictly decreasing, and Δ changes sign at $x' = x$. The reservation price of switching into a worse house is negative.

Consider two households, $h = 1, 2$, endowed with houses $x_1 < x_2$. There is trade between them if and only if their reservation prices for the trade sum up to something positive, i.e., if

$$S(x_1, y_1, x_2, y_2) := \Delta(x_2|x_1, y_1) + \Delta(x_1|x_2, y_2) > 0. \tag{25}$$

The wealthier household has, by definition, a larger budget set, so PAM by wealth and house quality is equivalent with PAM by utility and house quality. We want to show that if $x_1 < x_2$ and $u(x_1, y_1) > u(x_2, y_2)$ then there must be trade, as this rules out any violations of PAM in equilibrium.

Consider a point $\{x_2, y_2'\}$ where $y_2' = y_1 - \Delta(x_2|x_1, y_1)$. This is along the indifference curve going through (x_1, y_1) and vertical in relation to (x_2, y_2). Swapping positions between (x_2, y_2')

and (x_1, y_1) does not change utility so if $y_2 = y_2'$ then $\Delta(x_1|x_2, y_2') + \Delta(x_2|x_1, y_1) = 0$. If $\Delta(x_1|x_2, y)$ is strictly decreasing in y then there will be trade if and only if $y_2 < y_2'$, which is equivalent to $u(x_1, y_1) > u(x_2, y_2)$. Differentiating (24) we obtain

$$\frac{d\Delta}{dy} = 1 - \frac{u_y\left(x_2, y\right)}{u_y\left(x_1, y - \Delta\right)}. \tag{26}$$

We know that u_y is positive and decreasing for a fixed x, but now the comparison is at two different levels of x, yet on the same indifference curve. Use $\tilde{y}(x)$ to denote the indifference curve. Thus (26) is negative if

$$\frac{d}{dx}\left(u_y\left(x, \tilde{y}(x)\right)\right) = u_{xy} + u_{yy}\left(\frac{d\tilde{y}(x)}{dx}\right) = u_{xy} - u_{yy}\left(\frac{u_x}{u_y}\right) \geq 0 \tag{27}$$

But this is just the condition for MRS_{yx} to be decreasing in y, i.e. for MRS_{xy} to be increasing in y. Notice that, while $u_{xy} \geq 0$ is a sufficient condition for PAM, it is also stronger than necessary.

References

ALONSO, WILLIAM (1964). *Location and Land Use.* Harvard University Press, Cambridge, MA.

BAJARI, PATRICK; CHAN, PHOEBE; KRUEGER, DIRK AND DANIEL MILLER (2010). A Dynamic Model of Housing Demand: Estimation and Policy Implications. *NBER working paper* 15955.

BUITER, WILLIAM (2008). "Housing Wealth Isn't Wealth." *NBER Working Paper*, 14204.

EPPLE, DENNIS AND HOLGER SIEG (1999). "Estimating Equilibrium Models of Local Jurisdictions." *Journal of Political Economy*, 107(4), pp. 645–681.

FUJITA, MASAHISA (1989). *Urban Economic Theory: Land Use and City Size.* Cambridge University Press, Cambridge, UK.

GLAZER, AMIHAI; VESA KANNIAINEN AND PANU POUTVAARA (2008). "Income Taxes, Property Values, and Migration." *Journal of Public Economics*, 92, pp. 915–923.

GYOURKO, JOSEPH; CHRISTOPHER MAYER AND TODD SINAI (2006). "Superstar Cities." *Working Paper.*

LEGROS, PATRICK AND ANDREW F NEWMAN (2007). "Beauty Is a Beast, Frog Is a Prince: Assortative Matching with Nontransferabilities." *Econometrica*, 75(4), pp. 1073-1102.

LI, WENLI; HAIYOUNG LIU AND RUI YAO (2009). "Housing over Time and over the Life Cycle: A Structural Estimation." *Federal Reserve Bank of Philadelphia Working Paper*.

MIYAGAWA, EIICHI (2001). "House allocation with transfers." *Journal of Economic Theory*, 100(2), pp. 329–355.

MORETTI, ENRICO (2010). "Real Wage Inequality." *Working paper.*

ORTALO-MAGNÉ, FRANCOIS AND SVEN RADY (2006). "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints." *Review of Economic Studies*, 73, pp. 459-85.

RÍOS-RULL, JOSÉ-VICTOR AND VIRGINIA SANCHEZ-MARCOS (2008): "An Aggregate Economy with Different House Sizes." *Journal of The European Economic Association*, P&P, 6(2-3), pp. 705-14.

SATTINGER, MICHAEL (1979). "Differential Rents and the Distribution of Earnings." *Oxford*

Figure 1.

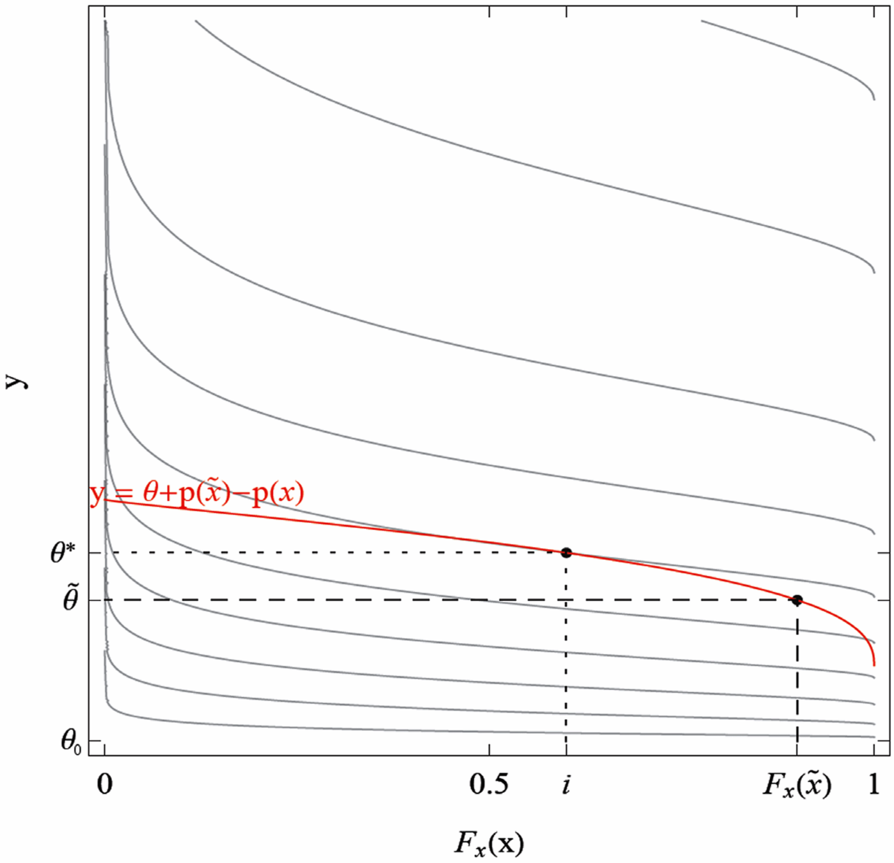

Figure 2.

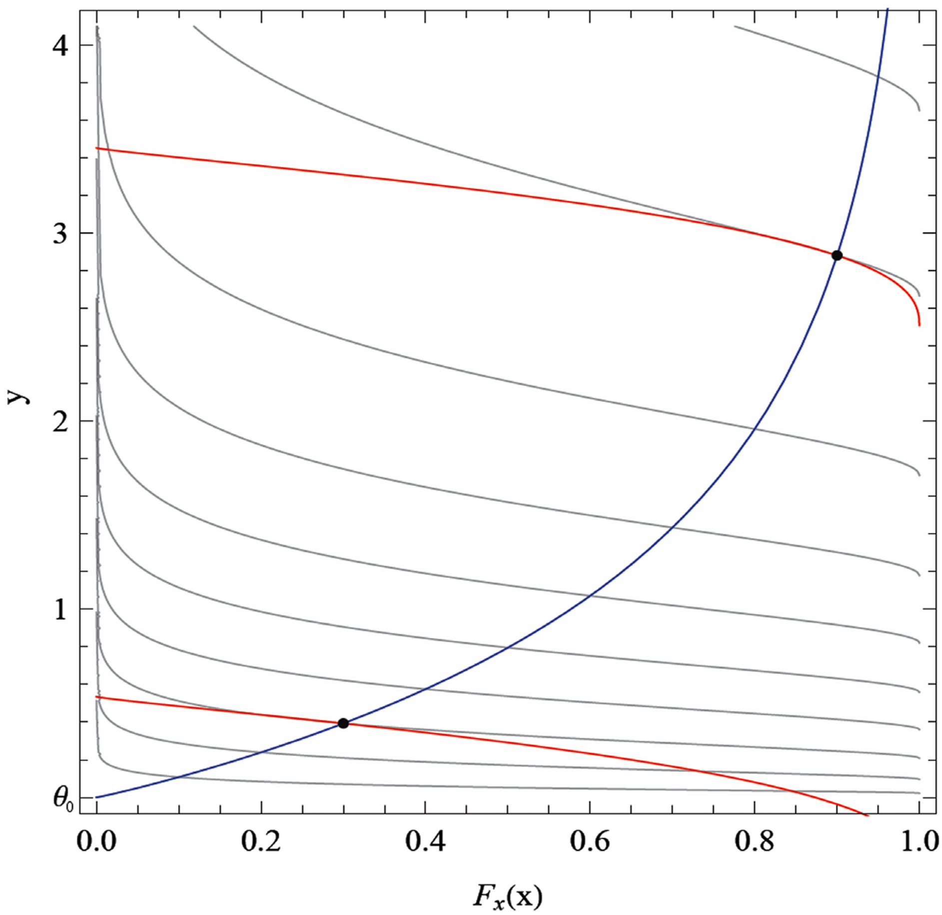

Figure 1.

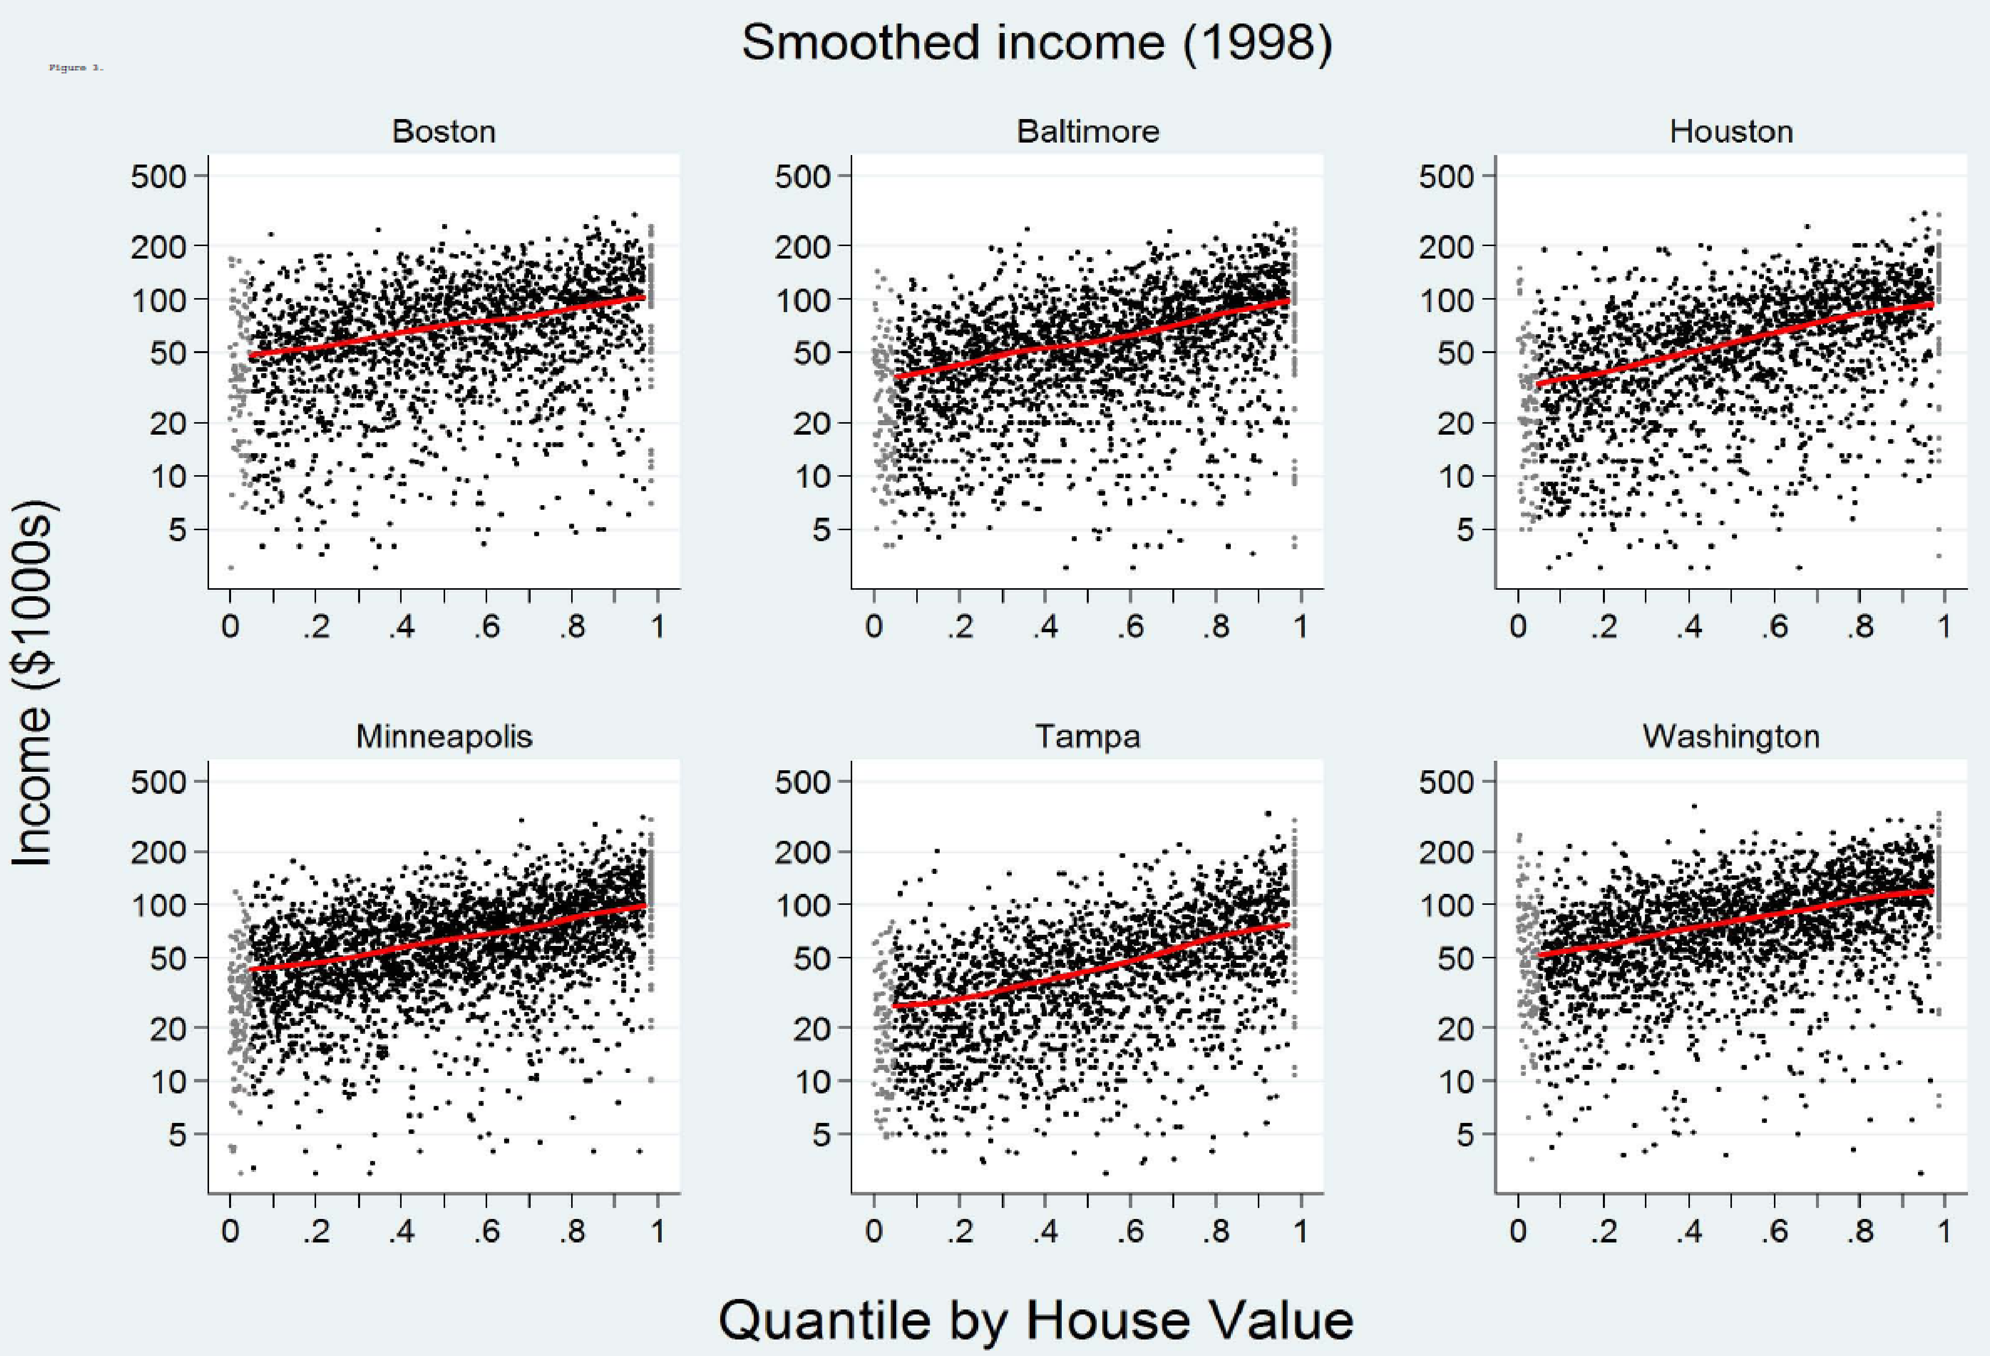

Smoothed income (1998)
Boston
Baltimore
Houston
Minneapolis
Tampa
Washington
Income ($1000s)
500
200
100
50
20
10
5
0
.2
.4
.6
.8
1
Quantile by House Value

Figure 4.

Smoothed income (2007)
Boston
Baltimore
Houston
Minneapolis
Tampa
Washington
Income ($1000s)
2000
1000
500
200
100
50
20
10
5
0
.2
.4
.6
.8
1
Quantile by house value

Figure 5.

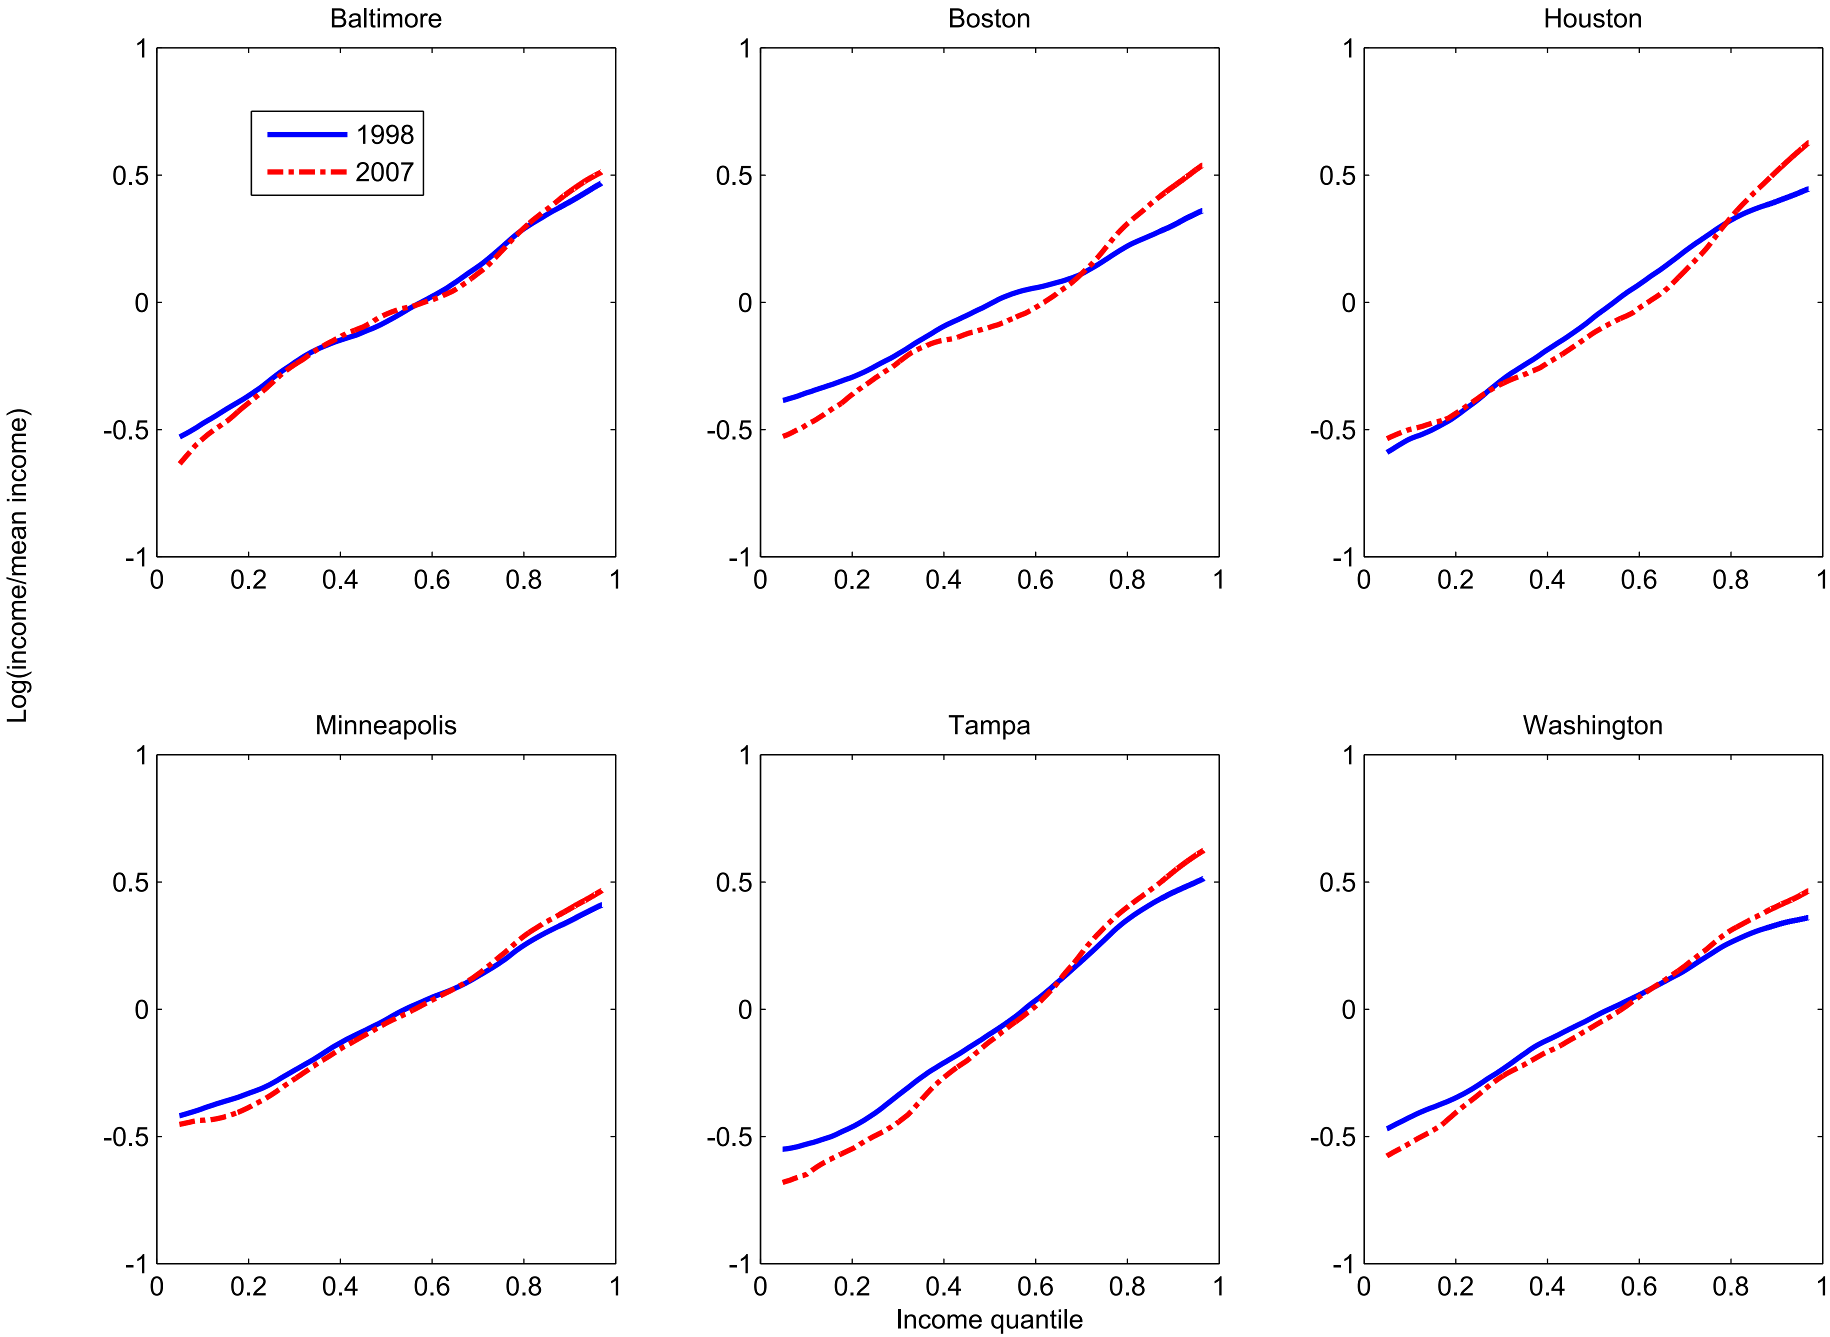

Baltimore
Boston
Houston
Minneapolis
Tampa
Washington
1998
2007
Log(income/mean income)
Income quantile
1
0.5
0
-0.5
-1
0
0.2
0.4
0.6
0.8
1

Figure 6.

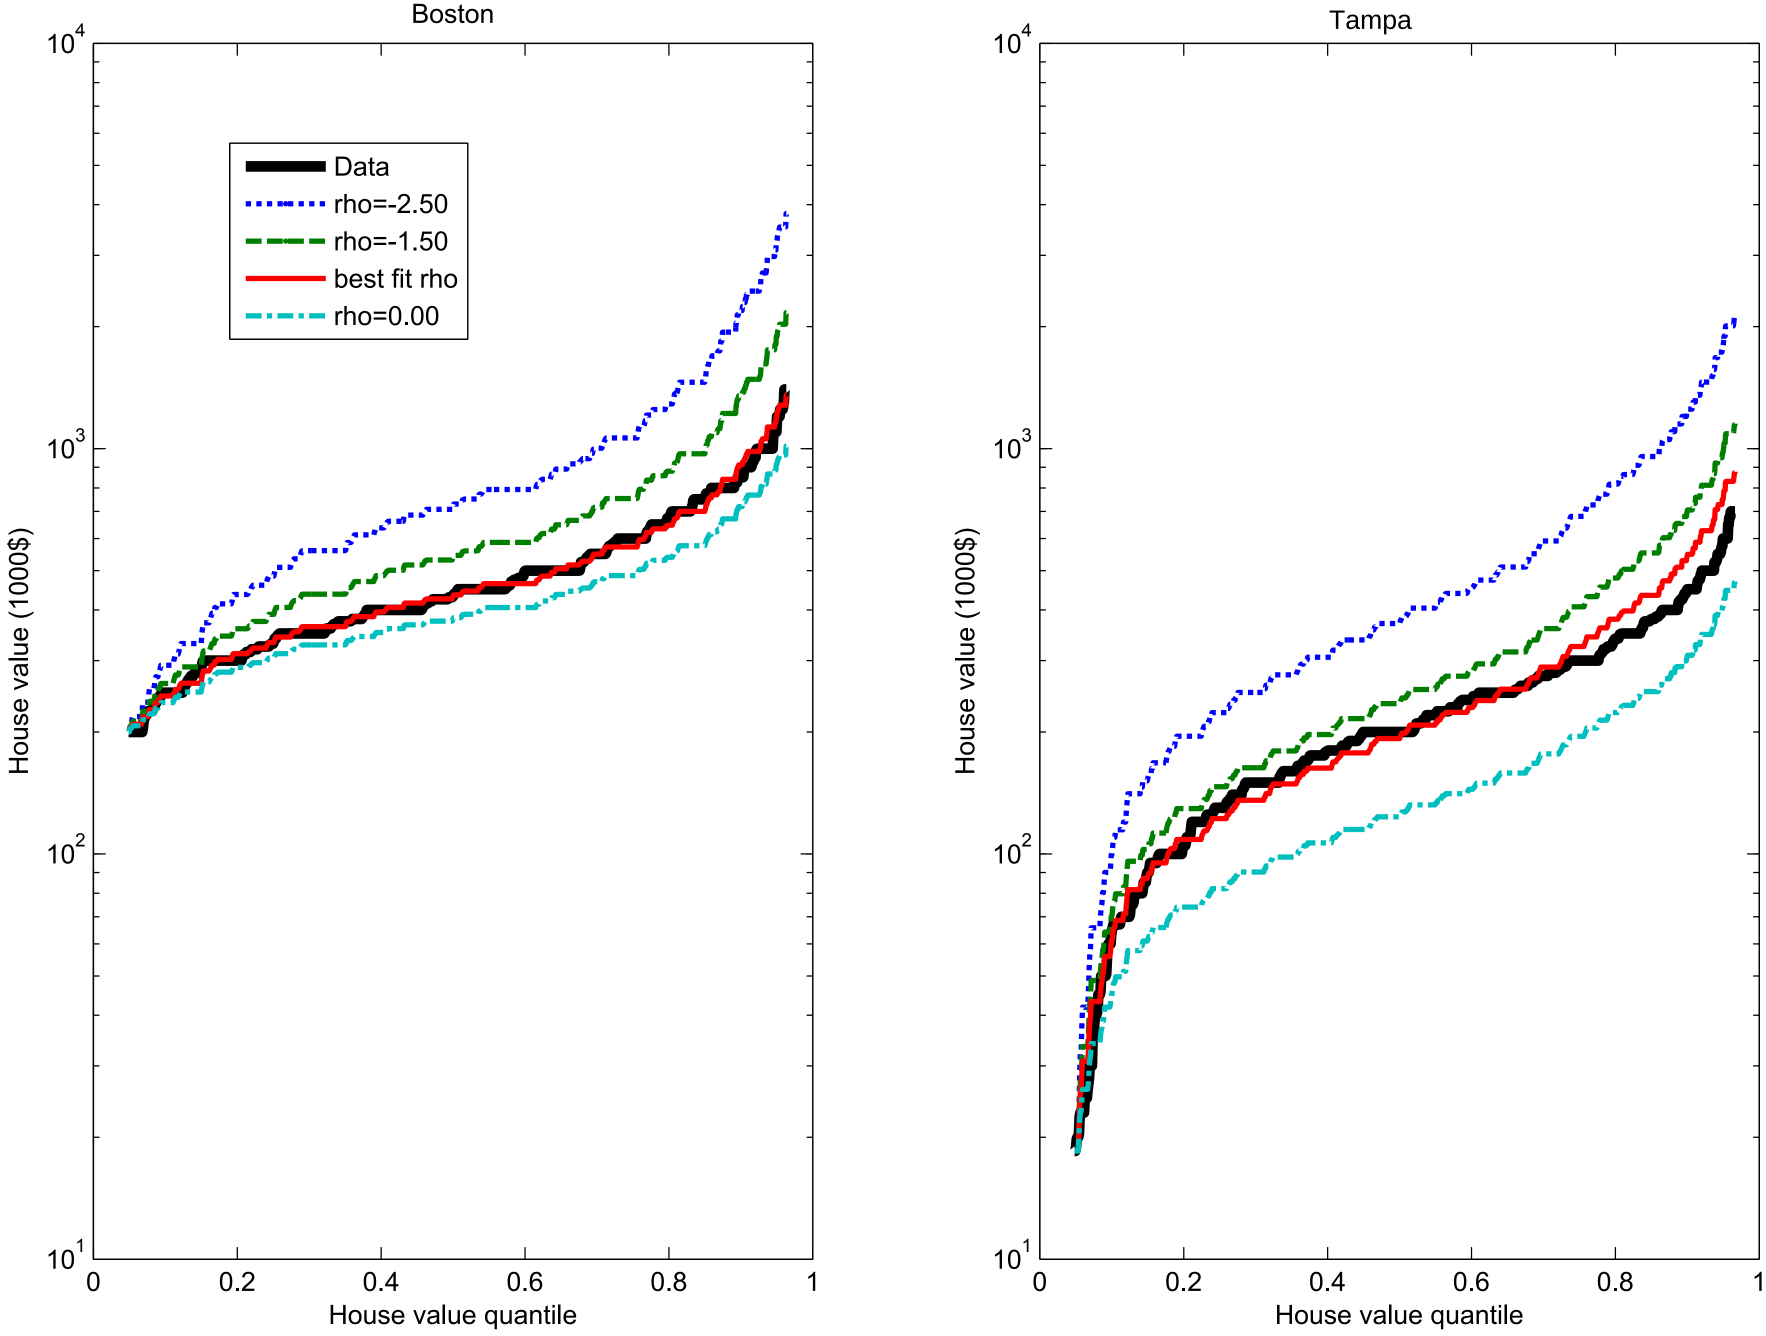
Boston
Tampa
Data
rho=-2.50
rho=-1.50
best fit rho
rho=0.00
House value (1000$)
House value quantile
10^4
10^3
10^2
10^1
0
0.2
0.4
0.6
0.8
1

Figure 7.

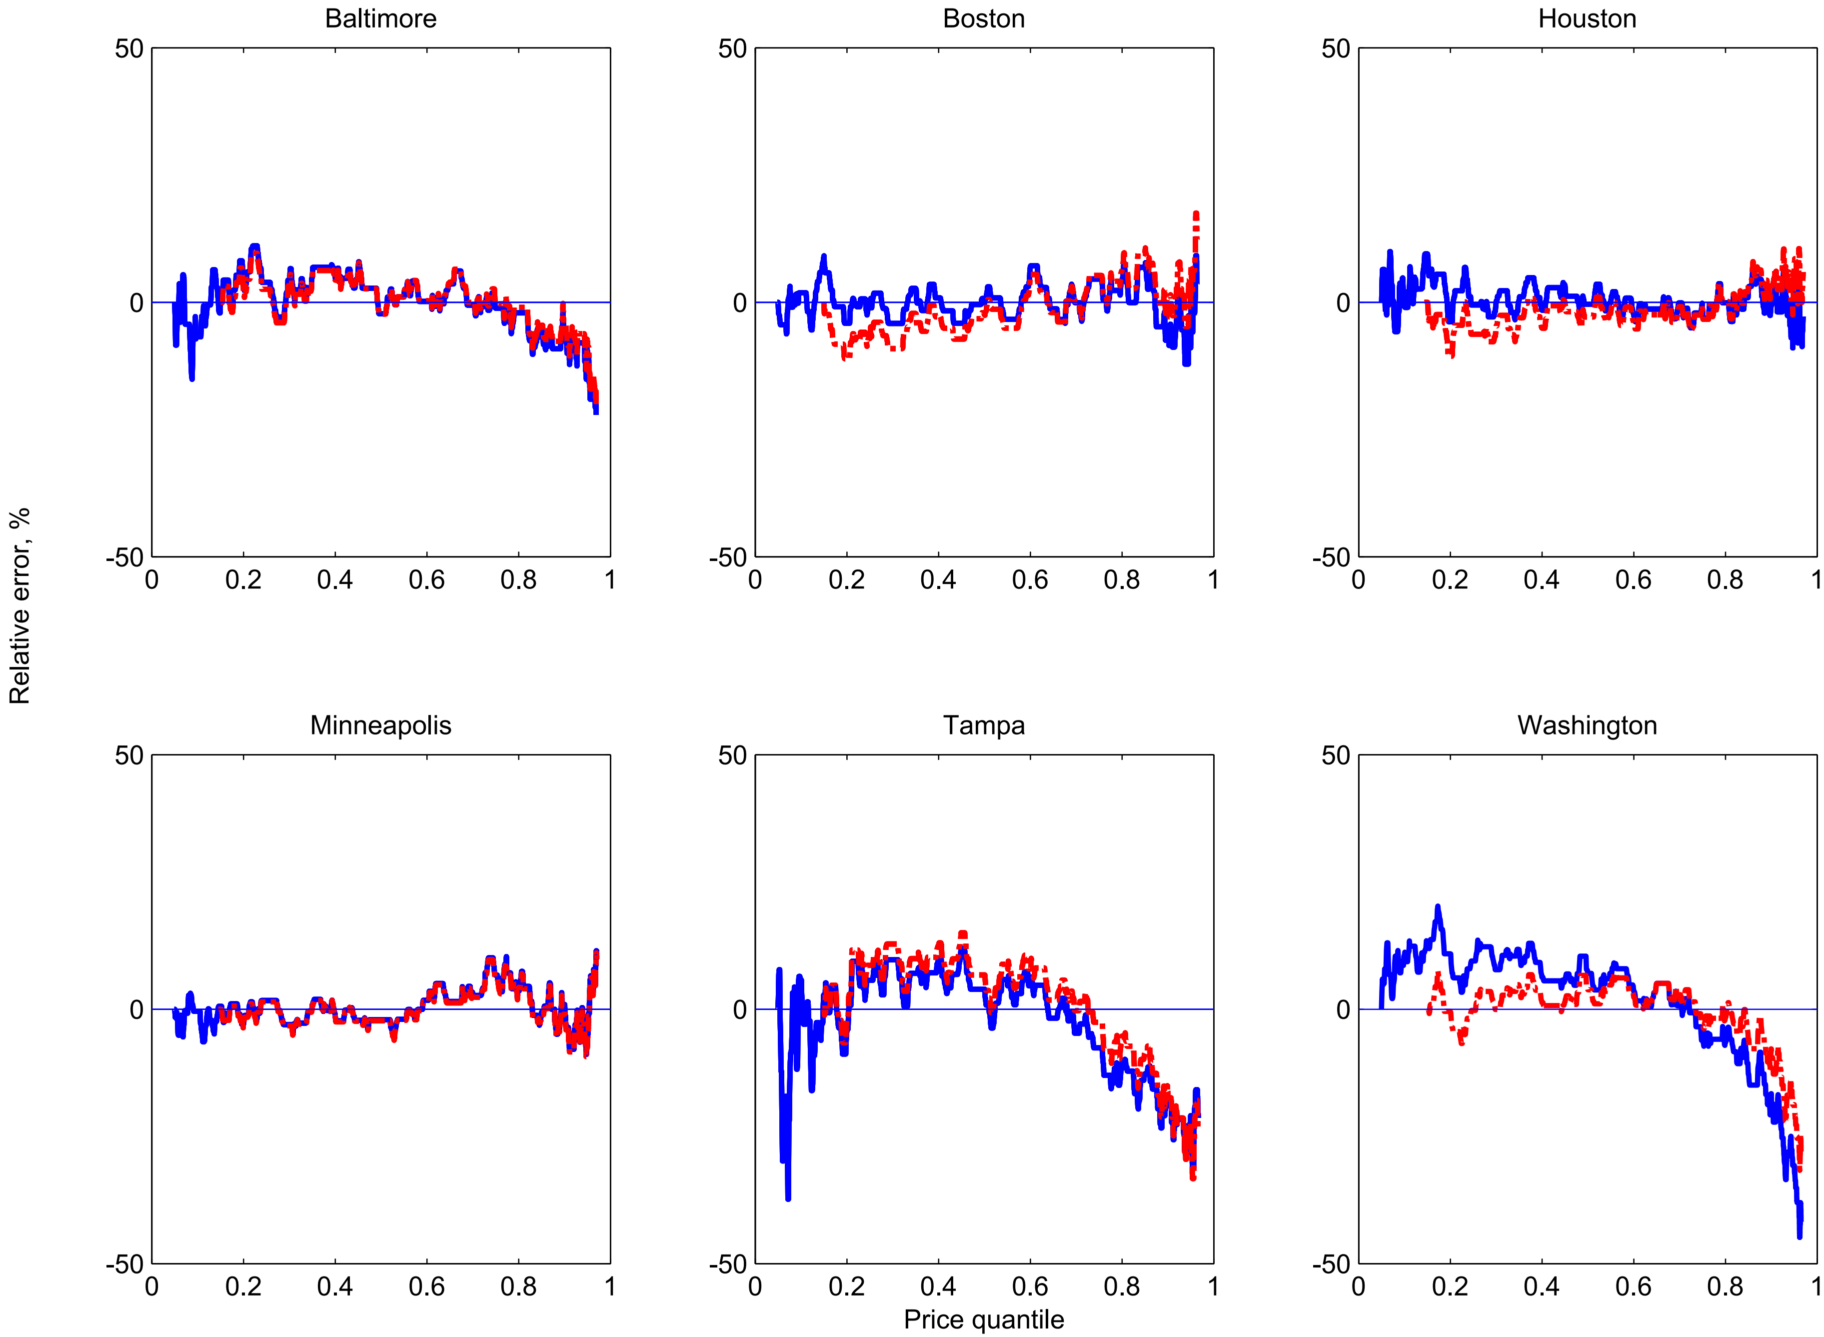
Baltimore
Boston
Houston
Minneapolis
Tampa
Washington
50
0
-50
0
0.2
0.4
0.6
0.8
1
Relative error, %
Price quantile

Figure 8.

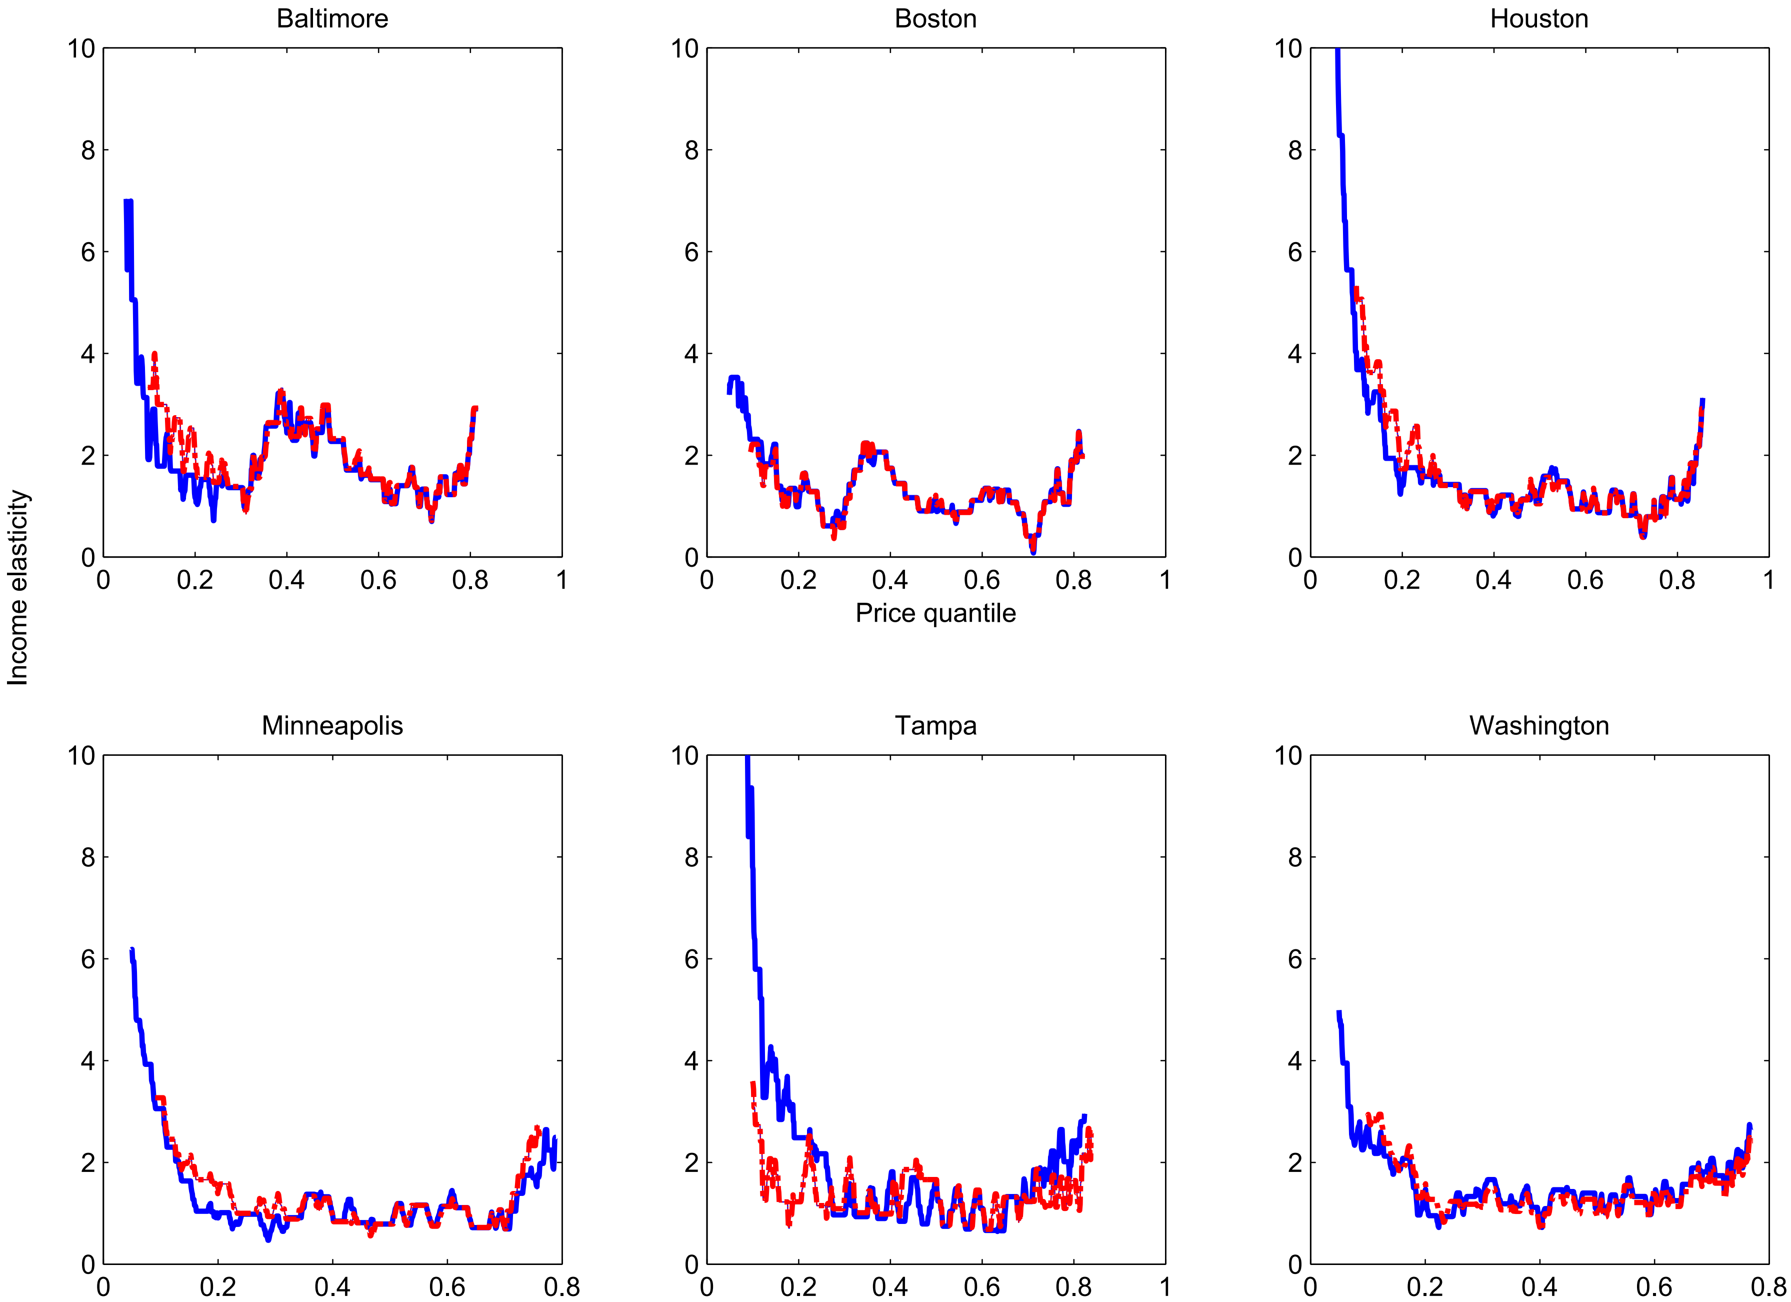
Baltimore
Boston
Houston
Minneapolis
Tampa
Washington
Income elasticity
Price quantile
10
8
6
4
2
0
0
0.2
0.4
0.6
0.8
1

Figure 9.

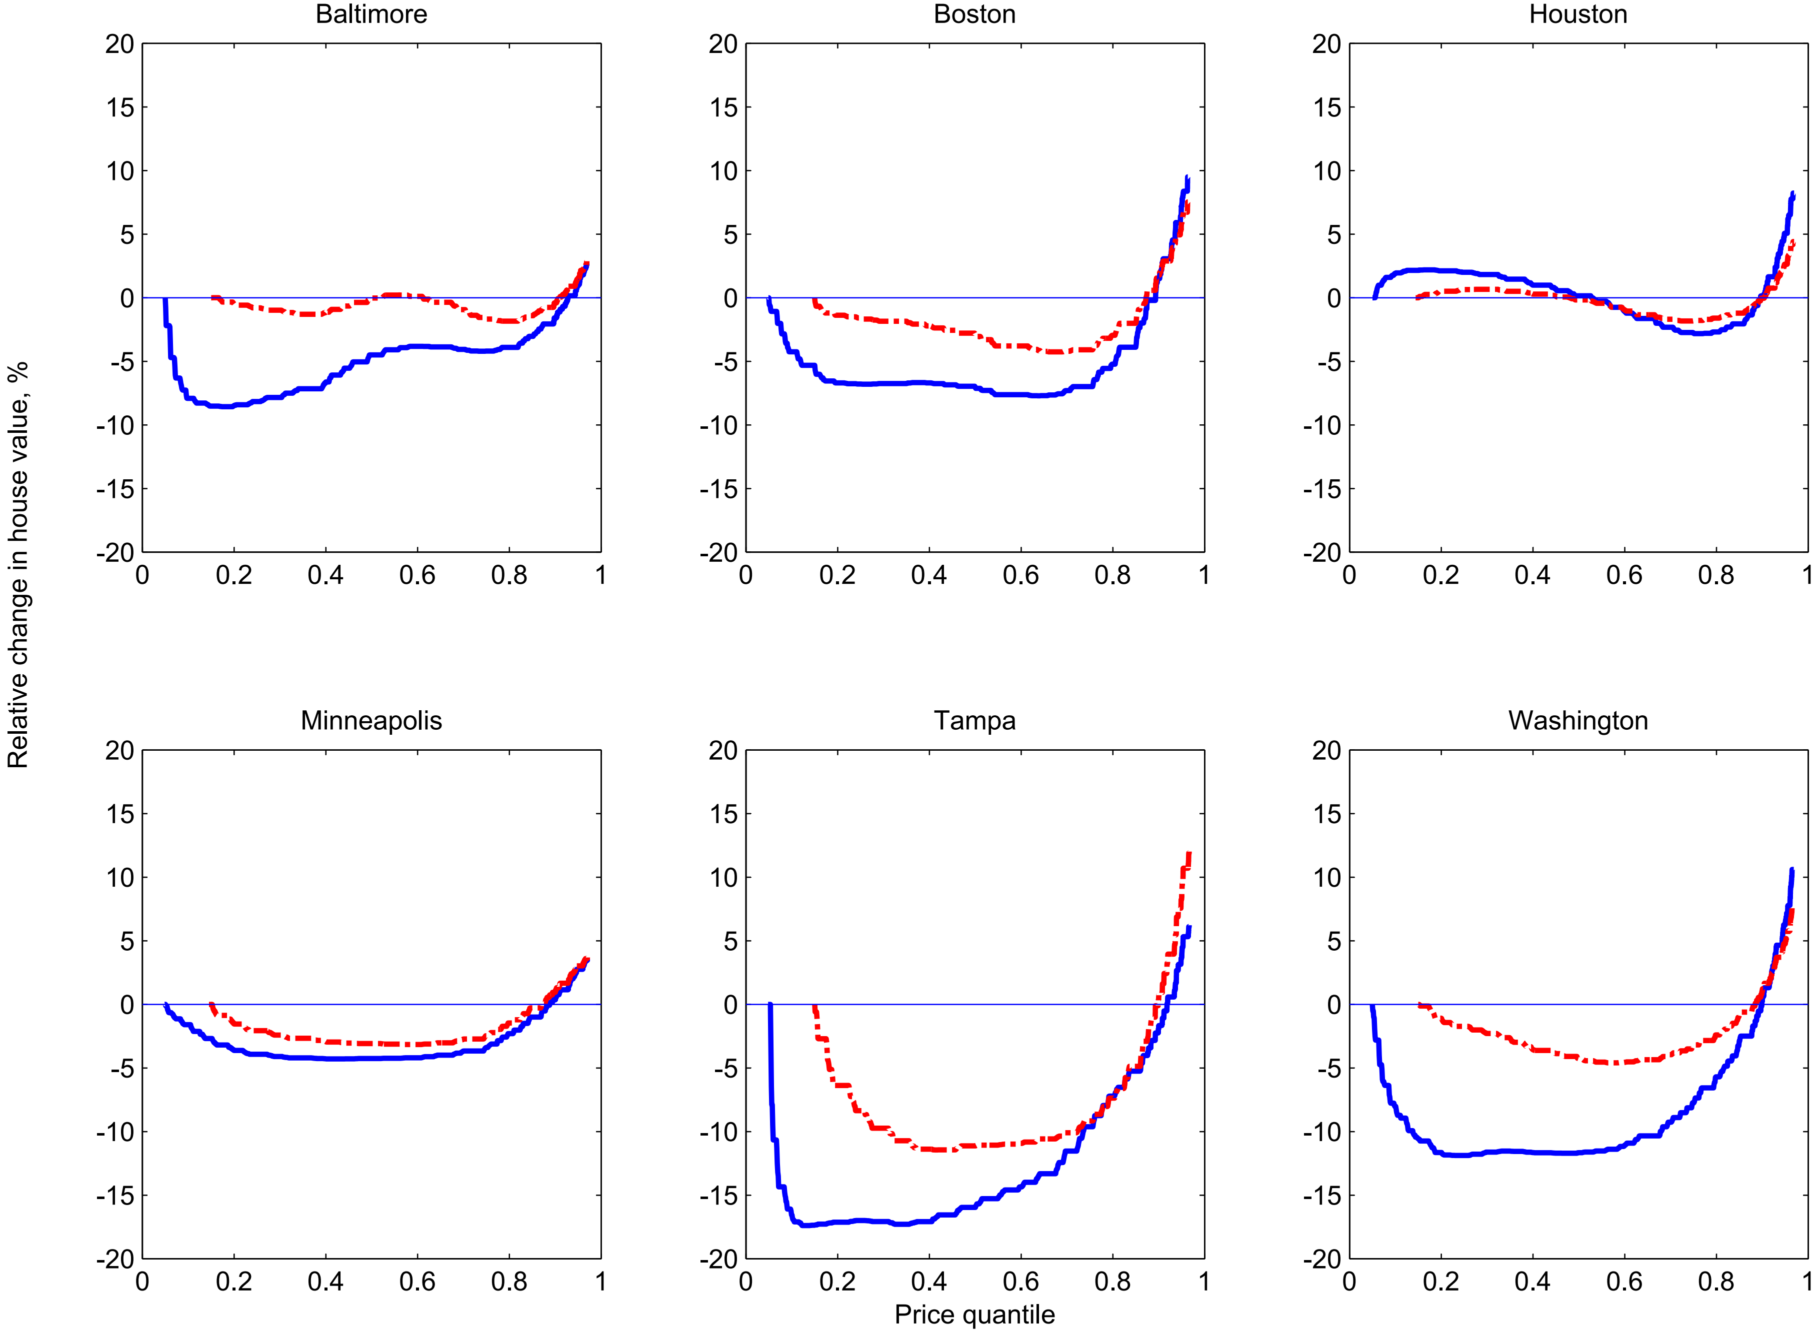

Baltimore
Boston
Houston
Minneapolis
Tampa
Washington
Relative change in house value, %
Price quantile
20
15
10
5
0
-5
-10
-15
-20
0
0.2
0.4
0.6
0.8
1

Figure 10.

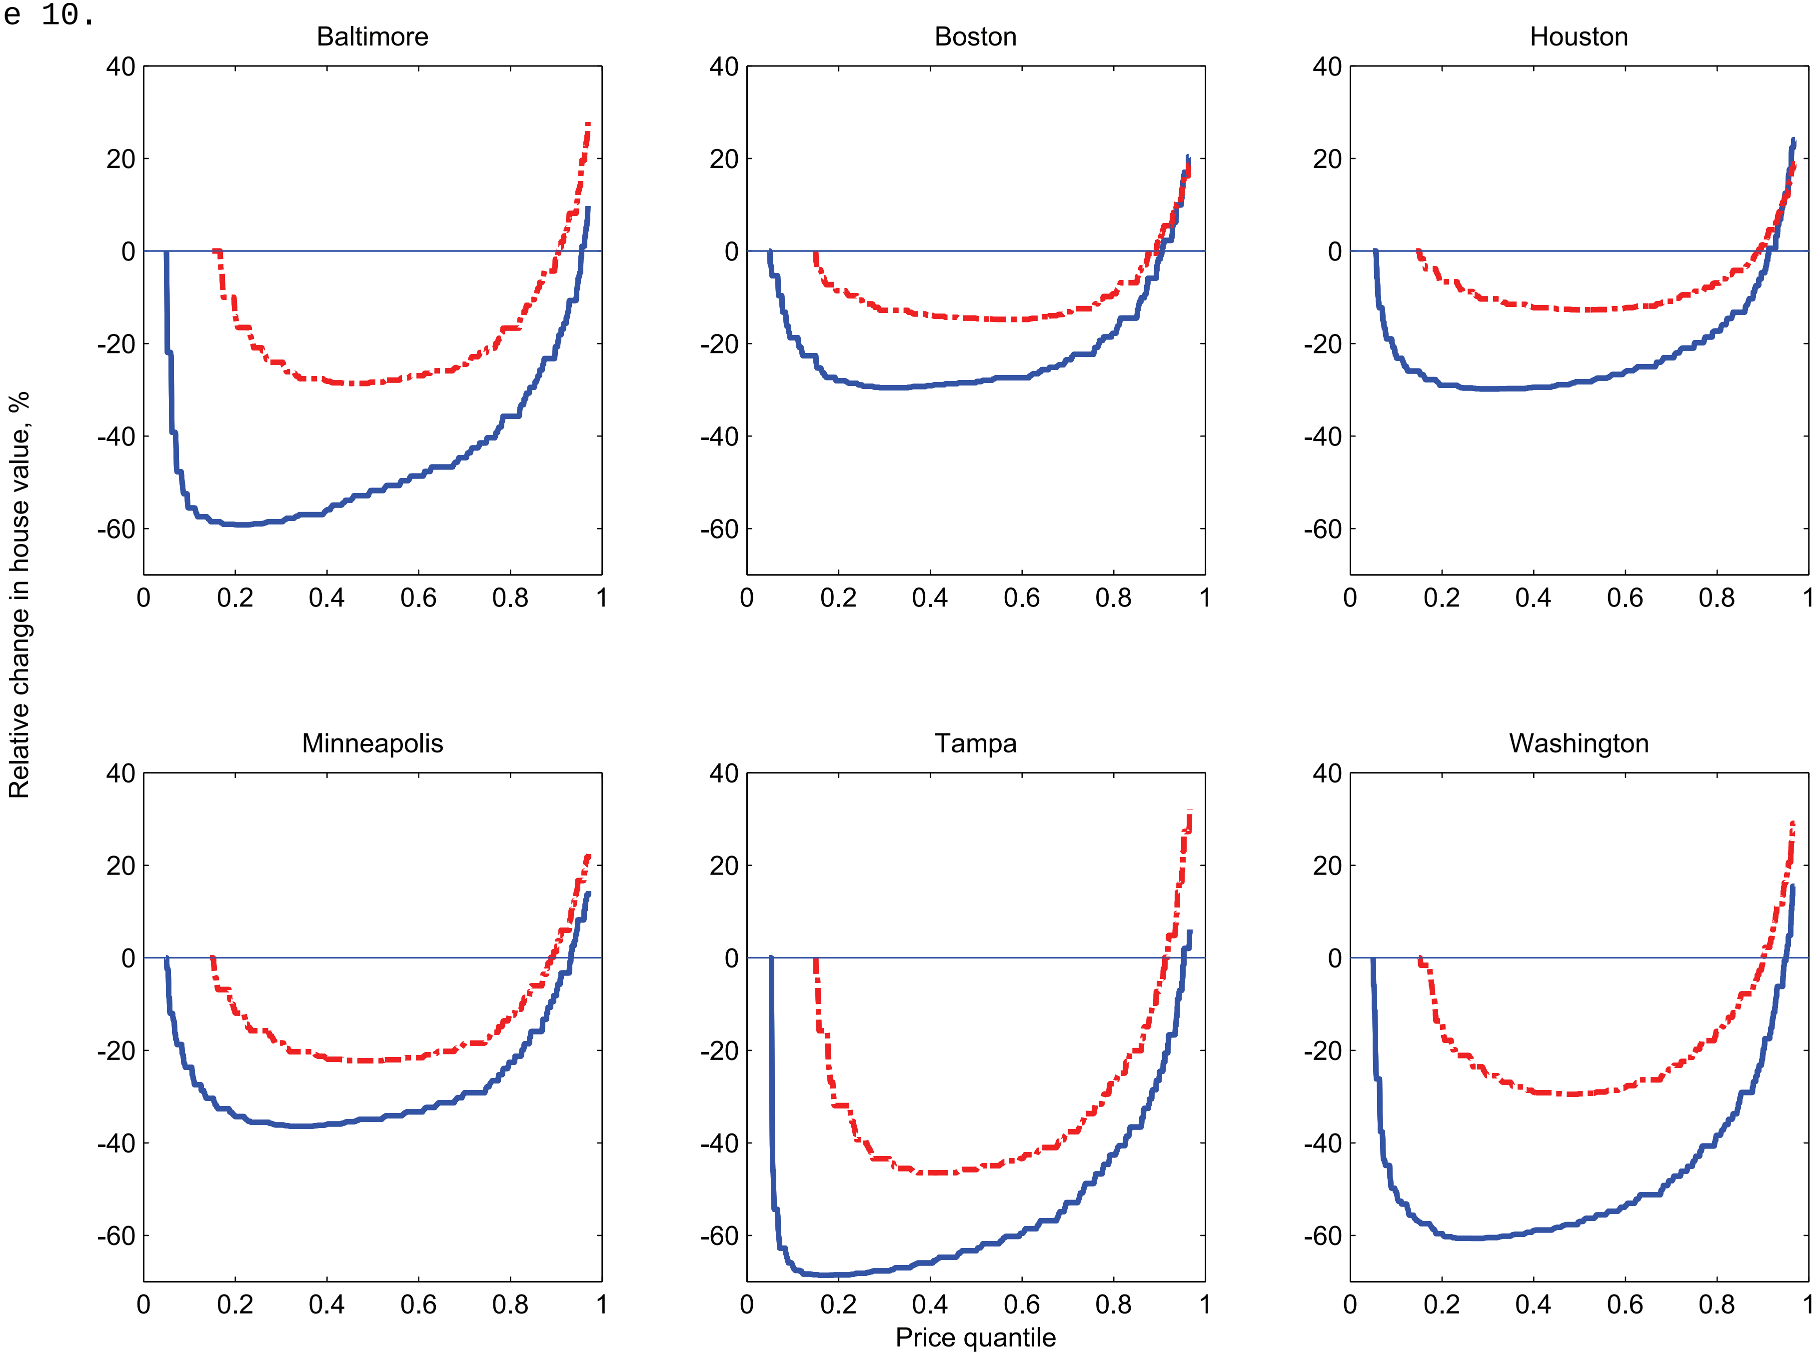

Baltimore
Boston
Houston
Minneapolis
Tampa
Washington
Relative change in house value, %
Price quantile
40
20
0
-20
-40
-60
0
0.2
0.4
0.6
0.8
1