CORE
Metadata, citation and similar papers at core.ac.uk
Provided by Helsingin yliopiston digitaalinen arkisto

Discussion Papers

Income Distribution and Housing Prices: An Assignment Model Approach

Niku Määttänen
ETLA and HECER

and

Marko Terviö
Aalto University and HECER

Discussion Paper No. 314
January 2011

ISSN 1795-0562

HECER – Helsinki Center of Economic Research, P.O. Box 17 (Arkadiankatu 7), FI-00014
University of Helsinki, FINLAND, Tel +358-9-191-28780, Fax +358-9-191-28781,
E-mail info-hecer@helsinki.fi, Internet www.hecer.fi

HECER
Discussion Paper No.314

Income Distribution and Housing Prices: An Assignment Model Approach*

Abstract

We present a framework for studying the relation between the distributions of income and house prices that is based on an assignment model where households are heterogeneous by incomes and houses by quality. Each household owns one house and wishes to live in one house; thus everyone is potentially both a buyer and a seller. The equilibrium distribution of prices depends on both distributions in a tractable but nontrivial manner. We show how the impact of increased inequality on house prices is in principle ambiguous, but can be inferred from data. We estimate the impact of increased income inequality between 1998 and 2007 on the distribution of house prices in 6 US metropolitan regions. We find that the increase in income inequality had a negative impact on average house prices. The impact of uneven income growth on house prices has been positive only within the top decile.

JEL Classification: D31, R21.

Keywords: Assignment models, housing prices, income inequality.

Niku Määttänen

Research Institute
of the Finnish Economy (ETLA)
Lönnrotinkatu 4 B
00120 Helsinki
FINLAND

e-mail: niku.maattanen@etla.fi

Marko Terviö

Department of Economics
Aalto University School of Economics
P.O. Box 21240
FI-00076 Aalto
FINLAND

e-mail: marko.tervio@hse.fi

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

1 Introduction

We present an assignment model of house price determination where houses are heterogeneous by quality and households are heterogeneous by income. Our main purpose is to study the relation of the distributions of income and house prices. A central motivating question is the impact of income inequality on the distribution of housing prices. It has been argued that the increase in consumption inequality has been less than the increase in income inequality because as the rich get richer they bid the prices of best locations ever higher. According to our theoretical results the impact of increased income inequality on top house prices is ambiguous. In our empirical application we estimate that the aggregate impact of recently increased inequality on house prices in the US has been negative, and positive only within the top decile.

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. What we refer to as "houses," for brevity, are really bundles of land and structures (including homes in multi-unit dwellings). The quality of land is inherently heterogeneous because locations differ in their attractiveness due to factors such as distance from the center and view of the sea. The supply of structures is more or less fixed in the short term, although adjustable in the long run. (Qualify of structures can also have a fixed component, due to zoning restrictions or the scarcity value of vintage architecture.) Another key feature of housing is that it takes up a large part of household expenditure, so income effects may be quite significant. Our modeling approach is based on an assignment model with non-transferable utility, which allows for income effects.

We consider a single metropolitan region, where the set of households is fixed. The distributions of income and house quality are exogenous, while the distribution of house prices is endogenous, with the exception of the cheapest or "marginal" house. The market does not consist of initially distinct classes of buyers and sellers but, rather, of a population of households who each own one house and each wish to live in one 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 the equilibrium conditions in our setup would be quite

complicated. However, we focus on the equilibrium prices that emerge after all trading opportunities have been exploited. In empirical terms, we assume that all households prefer to live in their current house. 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 also find that a suitably parametrized CES utility function allows us to match the change in the price distribution under the assumption that it was caused by the change in incomes while house qualities remained unchanged. We then use the inferred distribution of house qualities and our preferred utility parametrization to generate counterfactuals that measure the impact of the changes in income inequality on house prices.

In most assignment models a productive complementarity makes it efficient to match "the best with the best," i.e., total output is maximized by Positive Assortative Matching (PAM). Our setup is a pure exchange economy, but there is another driving force towards PAM, namely the diminishing marginal rate of substitution. The wealthy must live in the most desirable locations and have the highest levels of non-housing consumption, or else there would be unrealized gains from trade. Indeed, in our model, the only reason why the wealthy inhabit the best houses is that they can best afford them. In order to focus on the impact of changes in income distribution we assume homogeneous preferences, so the ordering of houses by market price is also the ordering by quality.

In our empirical application we use data from all six metropolitan regions that were 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 contribution of uneven income growth on house prices has been positive only within the top decile, with magnitudes of up to 12%.

The reason why the counterfactual of 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 relevant part of the income distribution are willing to pay for the quality difference. The price level of any particular house is then given by the summation of all price gradients below, plus the price of the marginal house. Conversely, after an increase in income inequality, 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. In principle, it is possible for all house prices to go down in response to an increase in inequality (but 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 our theoretical results. In Section 4 we show how the model 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 with non-transferable utility. Assignment models are models of matching markets that focus on the combined impact of indivisibilities and two-sided heterogeneity; for a review see Sattinger (1993). All other frictions, such as imperfect information or transaction costs, are assumed away. Both sides of the market are assumed to have a continuum of types, so there is no market power or "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 for constructing 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 competition for scarce land, thus leading the price of houses in the superstar city to go up. In Ortalo-Magne and Prat (2010) household location choice between regions is modeled as part of a larger portfolio problem, where each region has a fixed amount of (infinitely divisible) housing capital. Different regions offer different income processes, so location decisions as well as house prices are affected by hedging considerations.

Moretti (2010) has argued that the recent increase in income inequality in the US overstates the actual increase in consumption inequality, due to changes in house prices between cities. He 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 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 by Ortalo-Magne and Rady (2006) and Rios-Rull and Sanchez-Marcos (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. 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 recent papers are by 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 cross-sectional 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 the classic Von Thünen model, and Alonso (1964). In urban economics models the exogenous heterogeneity of land is due to distance from the urban center. The focus 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 the 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)

[1] See, for example, the textbook by Fujita (1989).

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 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 exists also a large literature on two-sided assignment, where two ex ante distinct classes of agents, "buyers" and "sellers," are matched, but these are further from our setup as we have no such distinction.[2] In our one-sided setup the reservation price of a seller depends on the opportunities available to her as a buyer.

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 level of 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

[2] See, for example, Legros and Newman (2007) and Caplin and Leahy (2010).

mutually beneficial trades have already been made, so the role of equilibrium prices is merely to enforce the no-trade equilibrium. The post-trade case is not only tractable but also empirically useful. 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, differentiable, and quasi-concave. 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) can be 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 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 strictly 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.

[Figure 1 here]

The initial endowment is described by a distribution of households over the consumption space, where house quality is on horizontal and composite good ("money") on vertical axes. Indifference curves are depicted in gray, while the red curve depicts the budget curve of a household endowed with income $\tilde{\theta}$ and a house of quality $\tilde{x}$. The shape of the budget curve depends on the price function p.

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. We don't need to know whether an arbitrary initial endowment is associated with a unique equilibrium. 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, the 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. We prove the existence of an equilibrium allocation in the Appendix under a finite (but arbitrarily large) number of house types. The equilibrium is associated with unique prices, up to a constant that can be interpreted as the opportunity cost of the worst house.

The increasing curve in Figure 2 depicts the equilibrium allocation for one 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.

[Figure 2 here]

Under equilibrium prices, all households must be located on a curve (depicted blue in this example) where they reach the highest possible utility along their budget curve.

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 $p(0) = p_0$ 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 region). 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_0 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 distributions 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 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 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. For brevity, we call this a "regular" equilibrium.

Definition. A *regular equilibrium allocation* is one where there is a strictly monotonic increasing relation between household income and house price.

This, in our model, is equivalent to the case where income and wealth are perfectly rank correlated. As explained earlier, the equilibrium in our setup has to satisfy perfect rank correlation between wealth and house price, which is a weaker requirement.[4] The purpose of this simplification is to make sure that the analytics of the no-trade equilibrium can be used even under changes in income distribution, as order-preserving changes in incomes are then guaranteed to keep the ranking of households by wealth unchanged. A change in the ordering by wealth would generate trading and would thus not fall within the scope of the no-trade case.

[4]In our data the monotonic relation between income and house prices emerges very naturally as a "side-effect" of kernel smoothing the data under the minimal assumptions needed to make wealth monotonic in house price.

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 a regular 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

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

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

$$\begin{aligned} \hat{p}(i) - p(i) &= p_0 + \int_0^i \frac{u_x(x(j), \hat{\theta}(j))}{u_y(x(j), \hat{\theta}(j))} x'(j)\, \mathrm{d}j - \left(p_0 + \int_0^i \frac{u_x(x(j), \theta(j))}{u_y(x(j), \theta(j))} x'(j)\, \mathrm{d}j \right) \\ &= \int_0^i \left(\frac{u_x(x(j), \hat{\theta}(j))}{u_y(x(j), \hat{\theta}(j))} - \frac{u_x(x(j), \theta(j))}{u_y(x(j), \theta(j))} \right) x'(j)\, \mathrm{d}j. \end{aligned} \tag{4}$$

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 (4) 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 (4) 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. This would happen, for instance, if the quality differences between

best houses are relatively small so that the rich are less inclined to use income growth to bid up the price of the next better house.

The impact of a simple increase in income levels is characterized by the next Proposition.

Proposition 5 *Suppose the endowments form a regular 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* i*.*

The reasoning is similar as in the Proof of Proposition 4, but even simpler because the new price gradient is greater than the original gradient at all quantiles.

As an immediate corollary, an increase in income levels will also increase the variance of house prices. A simple extreme case is where all incomes catch up with the highest income level. Proposition 4 tells us that such a complete elimination of income inequality would increase both the levels and the dispersion of house prices. Intuitively, all households would then be competing for the best houses and the price difference between a low quality and a high quality house must become relatively large to make some household willing to hold the low quality house.

The exogenous bottom price $p(0)$ is held constant under comparative statics, so the price distribution we are really referring to is for quantiles $(0, 1]$. A change in the lowest price, p_0, would cause all prices to change by that same amount. As the model does not include the possibility to move out of the market, the level of the constant is irrelevant for the attractiveness of any trade: it increases both the buying and selling prices, and thus washes out of all transactions.

3.3 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." This is similar in spirit to two-sided matching in that, by construction, sellers' reservation prices can be considered exogenous to the problem. The absentee landlord assumption can be introduced to our model by assuming that the revenue from house sales goes to atomistic outside agents who are not buyers in this market and have no market power as sellers.

Consider the household at quantile i of the income distribution. Again, by Lemma 1, equilibrium must involve positive assortative matching by wealth, which now 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.

$$i = \arg\max_{j \in [0,1]} u\left(x\left(j\right), \theta\left(i\right) - p\left(j\right)\right) \quad \text{for all } i \in [0,1] \tag{5}$$

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:

$$p'\left(i\right) = \frac{u_x\left(x\left(i\right), \theta\left(i\right) - p\left(i\right)\right)}{u_y\left(x\left(i\right), \theta\left(i\right) - p\left(i\right)\right)} x'\left(i\right). \tag{6}$$

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

If all houses were owned by an absentee monopolist, then the self-selection constraint (6) 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 lowest price $p(m)$ would then have to be pinned down by some outside opportunity for the buyers (e.g. an exogenous utility level from living in another region).

3.4 The case with CES

For the empirical application we assume CES utility,

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

with Cobb-Douglas utility defined in the usual fashion at $\rho = 0$. If wealth is just equal to income θ (as in the absentee landlord setup above), then the equilibrium price gradient (6) is

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

Under the post-trade assumption, where wealth is equal to the sum of income and the equilibrium value of one's house, $\theta(i) + p(i)$, the prices cancel out in the right-hand side and this can

be solved as

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

When p and θ are observed, then x can be solved for under a given elasticity parameter ρ. The other preference parameter, α, is absorbed by the units of x and can then be normalized at, say, one half.

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 insights. Here we present an example with a closed-form solution.

Assume that both income and house quality follow Pareto distributions, so that $\theta \tilde{}$Pareto(θ_0, η) and $x \tilde{}$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 $\rho < 0$. (This expenditure share is really the share of housing out of total wealth, which includes lifetime permanent income.) 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 extreme example provides the simplest demonstration of 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 social 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, subject to budget constraint.

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

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 Inferring quality and computing counterfactuals

In the empirical application we assume that 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. Under this assumption we can infer the unobserved distribution of house qualities from the observed relation between household income and house prices. That is, for given distributions p and θ, and for a given utility function u, we can infer the distribution of x. (We discuss 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.

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 the lowest price is always taken to be the lowest price in the data.

We now 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 household 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 reality, but it emerges very naturally in our data with some smoothing (details will be explained later).

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, and we obtain a discrete equivalent of the price gradient.[5]

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

Equivalently, solving for x_h yields the inference formula

$$x_h = \left(x_{h-1}^\rho + \frac{1-\alpha}{\alpha} \left[(\theta_h + p_h - p_{h-1})^\rho - (\theta_h)^\rho \right] \right)^{\frac{1}{\rho}}. \tag{21}$$

Note that the value of $\alpha \in (0,1)$ is without consequence when x is unobserved, because changing α is equivalent to changing the units of x. We set $\alpha = 1/2$. (The levels of x are not themselves useful, and this normalization is without loss of generality for the counterfactuals that we are interested in). Denoting $\hat{x} = x^\rho$, (21) can now be solved as

$$\hat{x}_h = \hat{x}_0 + \sum_{j=1}^{h} \left[(\theta_j + p_j - p_{j-1})^\rho - (\theta_j)^\rho \right], \tag{22}$$

which includes an undefined constant of integration $\hat{x}_0$, i.e. the quality of the worst occupied house.

Our main empirical interest is in constructing prices under counterfactual income distributions. Conceptually this is easiest to understand as combining inferred $\hat{x}_h$, observed p_0, and posited counterfactual $\hat{\theta}_h$ in the discrete price formula (20), to obtain counterfactual prices $\hat{p}_h$. In the CES case the steps can be combined to yield the counterfactual prices directly as

$$\hat{p}_h = \hat{p}_{h-1} + \left((\theta_h + p_h - p_{h-1})^\rho - (\theta_h)^\rho + \left(\hat{\theta}_h \right)^\rho \right)^{\frac{1}{\rho}} - \hat{\theta}_h. \tag{23}$$

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, or "city"): Baltimore, Boston, Houston, Minneapolis, Tampa, and Washington. The choice of MAs was strictly determined by the availability of data. The most populous MAs are

[5] In the discrete setup there is 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.

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 because it is done at city-level. 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).

In the AHS metro survey data, the house prices have been censored separately at each MA at the 97% percentile. Thus we have to exclude 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 bottom 5% of houses in each MA, and so the 5th percentile price will be the lowest (and thus exogenous) house price in our analysis. All of our results refer to this restricted sample, except where otherwise mentioned.

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.

We also need to set an interest rate r to make the units of yearly income compatible with the house price: income θ is measured as annual income divided by the interest rate. (In effect, we assume that households face an infinite time-horizon, and expect no changes in the future.) We fix the interest rate at $r = 0.05$. Changing r over a reasonable range ($2 - 8\%$) makes little difference to our inferred elasticity parameter ρ or the results of the counterfactual experiments that we present below.

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] Rank correlation between income and house price is between 0.38 (Boston) and 0.52 (Houston) 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.

[Figures 3 and 4 here]

Actual and kernel smoothed household income plotted against house value quantile by metropolitan area and year. Gray dots depict excluded or top-coded 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. The only apparently exceptional case is Houston where lowest incomes have grown faster than the mean.

	Permanent income		Income		House value		Rank correlation	
	1998	2007	1998	2007	1998	2007	1998	2007
Baltimore	0.16	0.17	0.37	0.40	0.27	0.28	0.38	0.48
Boston	0.13	0.17	0.36	0.44	0.25	0.23	0.48	0.46
Houston	0.18	0.20	0.39	0.41	0.30	0.27	0.53	0.52
Minneapolis	0.14	0.16	0.33	0.36	0.23	0.23	0.50	0.49
Tampa	0.19	0.23	0.41	0.52	0.29	0.31	0.52	0.48
Washington	0.14	0.17	0.33	0.39	0.25	0.23	0.50	0.49

Table 1. Summary statistics for inequality (Gini coefficient) of permanent income, income, and house value, and for the Spearman rank correlation between income and house value. "Permanent income" is income kernel smoothed relative to house price.

[Figure 5 here]

Distribution of permanent income relative to mean by city and year.

[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 assortative matching by wealth and house price.

In reality the value of the house consists of the value of land and the value of structures. It is, of course, possible to adjust the quality of the structures to some extent. We abstract away from adjustable quality and concentrate on relatively short term variation in our empirical application. However, even in the long term, according to Davis and Heathcote (2007), most growth and volatility in U.S. house prices has been due to variation in land values.[7]

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 while restricting it to be the same at 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 (22). 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 $\hat{p}^{07}$. We set the price of the lowest quality house (5th percentile in the full sample) 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 apply 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. Notice how assuming Cobb-Douglas preferences ($\rho = 0$) would result in systematic underprediction of the prices. 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 $|\log(p^{07}(i)) - \log(\hat{p}^{07}(i|\rho))|$. Thus we are trying to match the entire 2007 price distribution

[7] Davis and Palumbo (2008) estimate the share of land in house values for detached single-family homes; in our sample cities it ranges between 31% (Houston) and 76% (Boston) in 2004.

with just two parameters, the other being the bottom price p_0 which we match exactly by construction. 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.

[Figure 6 here]

Actual and predicted price distributions under various values of ρ in two example cities.

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 common for all MAs. The common preferences case results in elasticity of 0.50, which 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 income elasticity in this 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, $\log(p^{07}/\hat{p}^{07})$, by quantile of house value. Our simple model seems to provide a reasonable fit in most cities. The exceptions are Tampa and Washington. According to the model, the prices of the best houses should have increased much more there than they actually did. It may be that the supply of top quality housing is relatively elastic in these two cities, whereas we assume a fixed quality distribution. At least for Tampa, it seems plausible that there are relatively attractive undeveloped locations, while in a city like Boston the best locations are mostly already either built up or protected from development. Elastic housing supply is also an example of an omitted factor that would bias our estimate of the elasticity parameter. However, as we show below, the results of our main counterfactual experiment are, in the end, relatively insensitive to the value of the elasticity parameter.

	N_{98}	N_{07}	$1/(1-\rho)$	MAPE
Baltimore	2306	1211	0.42	4.6
Boston	1960	1075	0.63	2.8
Houston	1963	1105	0.89	2.3
Minneapolis	2741	1524	0.48	2.6
Tampa	2106	1255	0.49	8.3
Washington	2356	1315	0.32	9.1
All	13432	7485	0.50	12.2

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. "All" refers to the case where preference parameter (ρ) is restricted to be the same in all cities.

[Figure 7 here]

Errors in predicted 2007 house prices by metropolitan area. Solid blue line refers to the baseline case, where preferences are estimated separately in each MA, and dashed red line to the case where preferences are restricted to be the same everywhere.

5.3 Income elasticity of housing expenditure

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).

In order to illustrate these features, let us define income elasticity of housing demand at quantile i using a 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. Housing expenditure changes from $p(i)$ to $p(j)$, where

$$j = \arg \max_{s \in [0,1]} (u(x(s), (1 + \triangle)\theta(i) + p(i) - p(s)) \tag{24}$$

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

In equilibrium, this elasticity can be determined just from current prices and incomes. Since there is positive assortative matching by housing quality and wealth, we can find the new housing expenditure $p(j)$ simply by finding where the household will be located in the wealth distribution after its income is changed. Thus j from (24) solves $(1 + \triangle)\theta(i) + p(i) = \theta(j) + p(j)$. The change in wealth position is independent of the utility function (as long as it exhibits the diminishing MRS required for positive sorting). In the end, this result stems from the assumption that houses are indivisible and have fixed qualities.

Figure 8 shows the 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.[8] We can calculate this elasticity only up to the point where the hypothetical income increase would lift the households outside our data range, which at the top is the top-coding threshold. Figure 8 reveals how income elasticity varies quite substantially over the distribution. Intuitively, how much more one would spend on housing following an increase in income depends on the increase in housing quality that would be available. In this model, the housing quality that one extra dollar can buy varies over the distribution, as it depends on the shape of the quality distribution and on the incomes of competing buyers.

[Figure 8 here]

Estimated income elasticity of housing expenditure by quantile in the distribution of wealth.

Given that there is assortative matching, the shape of the utility function matters only when there are aggregate changes. The counterfactual embedded in the individual household's demand elasticity holds prices constant because one household has a vanishingly small impact on

[8]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 $p(i)$.

prices. By contrast, if all household incomes go up by $\triangle$ percent then prices react; by how much depends on preferences. (Now no one will actually move, because everyone's rank in wealth distribution stays the same.) By substituting into (9), we see that under CES utility all price differences ($p - p_0$) blow up by a factor of $(1 + \triangle)^{1-\rho}$.

The observation that the elasticity of demand varies 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 can depend on the associated changes in the distribution of income. This is also illustrated by the following counterfactual experiments.

5.4 Counterfactuals

The empirical question we set out to answer was how changes in income distribution influence housing prices. We now apply our methodology to a specific data set to compute the impact of increased income inequality between 1998 to 2007 on housing prices in six US metropolitan areas. Specifically, given the inferred values for ρ (see Table 2) and the quality distributions x^{98} in each MA, we compute the predicted price distributions in 2007 under a counterfactual income distribution that has the same shape as 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. When interpreting these results it is important to keep in mind that we keep all other things, including the supply of housing, fixed.

Results are shown in Figure 9. It displays the relative difference between counterfactual and empirical housing prices. Our benchmark case is the one where the elasticity parameter is estimated separately for each MA. In addition, we consider the case where the elasticity parameter is restricted to be the same for all MAs.

Consider first the benchmark case. The impact of increased income inequality is qualitatively similar in all but one MA: it has lowered housing prices until about 80-90th percentile

and increased prices in the upper tail.[9] The exception is Houston, where changes near the bottom of the income distribution work 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 there as in other MAs.

By and large, the results are similar in the case where the elasticity parameter is the same for all MAs. In particular, the point at which the price effect turns from negative to positive is almost the same in both specifications. Also the increase in the price of even the best houses (those at the 97th percentile) is always moderate. 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. As explained in Section 3.2, this negative price impact at the bottom spills upwards in the quality distribution. This 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.

[Figure 9 here]

The impact of increased income inequality on housing prices.

The average price changes are shown in Table 3. The first two columns display the mean relative effect of the change in the shape of the income distribution on house prices for the MA-specific and common preference specifications. The third column displays the absolute price change in the benchmark case. The mean impact of change in income inequality in the benchmark specification varies from -10.1% in Tampa to 0.2% in Houston.

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. The last two columns in Table 3 display the mean effects of this radical counterfactual. According to the model, overall income inequality lowers housing prices by $19 - 52\%$ relative to the counterfactual of total income equality.

[9] Recall that our income measure is an estimated permanent income.

	Change (%)		Change ($k)	Total (%)	Total ($k)
	benchmark	common ρ	benchmark	benchmark	benchmark
Baltimore	-4.4	-3.6	-16.0	-43	-270
Boston	-4.5	-5.9	-23.3	-20	-123
Houston	+0.2	+0.6	+2.7	-19	-35
Minneapolis	-2.6	-2.5	-7.7	-26	-102
Tampa	-10.1	-9.9	-27.6	-52	-266
Washington	-6.6	-4.0	-36.7	-47	-460

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 the income distribution and the housing price distribution. In the model houses are heterogeneous and indivisible. The key element is that all houses are owned by the households, and, due to concave utility, their reservation prices as sellers depend on the opportunities available to them as buyers, which in turn depend on their incomes and on the prices of other houses. Thus, our model provides a framework for analyzing how income differences get capitalized into house prices. The equilibrium is tractable under the assumption that households prefer to live in their current house.

The equilibrium can be understood intuitively by considering the price gradient, which is, loosely, the price difference between "neighboring" houses in the quality distribution. The price gradient expresses how much households that inhabit a particular part of the quality distribution in equilibrium are willing to pay for the quality difference over next best house. This depends on their marginal rate of substitution between house quality and other goods, which in general depends on their level of wealth. The price level at any quantile in the distribution is the sum over all price gradients below.

The natural comparative static in our setup is order-preserving changes in income distribution, as these have an impact on prices without generating trading. The model yields a number of theoretical implications about the relation of income and house price distributions. Any increase in income levels will increase both the level and dispersion of house prices, but an increase in income inequality will decrease house prices except (possibly) in a segment adjacent to the top. House prices at the top can go either way in response to an increase in income inequality, depending on the details of both supply and demand sides of the market.

The equilibrium conditions enabled us to estimate the housing quality distribution without imposing restrictions on the shapes of the distributions. As our main empirical application we analyzed housing prices in six US metropolitan regions. We obtained a theory-driven estimate for the impact of recent increases in income inequality on house prices. Specifically, we asked how the 2007 price distribution would differ from the actual price distribution if income of every household would have grown at the actual mean rate since 1998. We found that the impact of increased inequality on prices has been modest but negative on average, and positive only at the top decile. This is because the cumulative impact of reductions in the price gradient at the bottom households, whose income growth did not keep up with the mean growth, dominates the positive effects almost all the way to the top of the distribution.

Our model opens the possibility for other applications. Most directly, it is a natural framework for studying how various housing and income subsidy schemes impact housing prices. The intuition of the price gradient makes it clear why housing subsidies targeted for the poor will not merely be capitalized into prices of low-quality housing but will spill upwards in the quality ladder. A serious empirical analysis of this issue will require the inclusion of non-owner-occupied housing.

The counterfactuals presented in this paper took the form of order-preserving changes in the exogenous income distribution. This was crucial for being able to apply the formulae derived under the assumption of no-trade equilibrium. The lack of trading also meant that price changes do not have welfare effects–they are merely changes in paper wealth. In order to have welfare effects, price changes have to generate trading. One interesting and challenging topic for further study is the welfare effects of regulations, for instance transaction taxes, which can be expected to distort the matching of houses and households. A particularly policy-relevant issue is the

impact of repealing rent control, which would result in a simultaneous supply and demand shock. Again, the impact on the entire distribution of housing prices is likely to be nontrivial. Our characterization of the (no-trade) equilibrium should be helpful even in applications that involve trading, although these will require more involved numerical methods.

7 Appendix

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

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

$$u(x,y) = u(x', y-\Delta), \tag{25}$$

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

$$S(x_1,y_1,x_2,y_2) := \Delta(x_2|x_1,y_1) + \Delta(x_1|x_2,y_2) > 0. \tag{26}$$

The wealthier household has 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 (25) we obtain

$$\frac{d\Delta}{dy} = 1 - \frac{u_y(x_2,y)}{u_y(x_1,y-\Delta)}. \tag{27}$$

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 (27) is negative if

$$\frac{d}{dx}\left(u_y\left(x,\tilde{y}(x)\right)\right) = u_{xy} + u_{yy}\left(\frac{d\tilde{y}(x)}{dx}\right) = u_{xy} - u_{yy}\left(\frac{u_x}{u_y}\right) \geq 0 \tag{28}$$

This is just the condition for diminishing marginal rate of substitution MRS_{yx} to be decreasing in y (i.e., for MRS_{xy} to be increasing in y), which holds for any quasi-concave utility function.

Proof: Existence of equilibrium Let's now discretize the house types, so that $0 < x_0 < \cdots < x_N < \infty$. For brevity, we will refer to households endowed with a type-k house as households of type k. Denote the mass of type-k households with income equal or lower than θ by $F_k(\theta)$. We assume that all of these conditional income distributions are continuous with full support $[\underline{\theta}, \bar{\theta}]$, where $0 < \underline{\theta} < \bar{\theta} < \infty$. The mass of type-$k$ households is $m_k > 0$, so $F_k\left(\bar{\theta}\right) = m_k$. Recall that we have normalized the total mass at $\sum m_k = 1$.

With discretized house types we require an additional assumption:

$$u(x_{k-1}, \bar{\theta}) > u(x_k, \underline{\theta}) \text{ for all } k = 1, \ldots, N. \tag{29}$$

This means that distribution of autarky utility has overlap between neighboring household types. Together with Lemma 1, this implies that there must be trade between them, which in turn will guarantee that price increments are uniquely determined.

Lemma 1 shows that equilibrium utility is increasing in the quality of the house consumed. The full support of incomes and the overlap of autarky utilities imply that the consumption levels y of households that consume a house of type k must be distributed with full support on some interval $[\underline{y}_k, \bar{y}^k] \in [\underline{\theta}, \bar{\theta}]$. The unequivocally poorest and richest households do not have anyone to trade with, so $\underline{y}_0 = \underline{\theta}$ and $\bar{y}_N = \bar{\theta}$.

The wealthiest household that consumes a type$-k$ house must be indifferent between trading up to the next house type. This indifference condition defines the price increments $\Delta_k = p_k - p_{k-1}$ as functions of upper bounds $\bar{y}_k$,

$$\gamma\left(\bar{y}_k | x_k, x_{k+1}\right) = \left\{\Delta_{k+1} : u\left(x_k, \bar{y}_k\right) = u\left(x_{k+1}, \bar{y}_k - \Delta_{k+1}\right)\right\} \text{ for all } k = 0, \ldots, N-1. \tag{30}$$

Note that γ is single-valued for all $\bar{y}_k \in [\underline{\theta}, \bar{\theta}]$, with image in $(0, \bar{y}_k)$.

The same indifference conditions can also be described in terms of the lower bounds of the consumption intervals, with $u(x_k, \underline{y}_{k+1} + \Delta_{k+1}) = u(x_{k+1}, \underline{y}_{k+1})$, so

$$\underline{y}_{k+1} = \bar{y}_k - \Delta_{k+1} = \bar{y}_k - \gamma\left(\bar{y}_k | x_k, x_{k+1}\right). \tag{31}$$

Thus the lower bounds as well as the price increments are uniquely determined by the upper bounds, so the equilibrium allocation can be described in terms of the $\bar{y}_k$ alone, where $k = 0, \ldots, N-1$.[10]

Demand for type-k houses is the sum of demands from each household type. Consider type$-j$ households endowed with income θ. They will consume a type-k house if their wealth is in the same range as of those type$-k$ households who consume their endowment: $p_j + \theta \in [p_k + \underline{y}_k, p_k + \bar{y}_k]$. The bounding inequalities for these intervals can be written as

$$\begin{aligned} \theta &\leq \bar{y}_k + p_k - p_j \\ \theta &\geq \underline{y}_k + p_k - p_j = \bar{y}_{k-1} + p_{k-1} - p_j \end{aligned} \tag{32}$$

for $k = 0, \ldots, N-1$. (Recall that $\bar{y}_N = \bar{\theta}$.) Total demand for type$-k$ houses is

$$Q_k\left(\bar{y}\right) = \sum_{j=0}^{N} \left[F_j\left(\bar{y}_k + p_k - p_j\right) - F_j\left(\bar{y}_{k-1} + p_{k-1} - p_j\right)\right] \tag{33}$$

$$\begin{aligned} &= \sum_{j=0}^{k-1} \left[F_j\left(\bar{y}_k + \left(\Delta_{j+1} + \cdots + \Delta_k\right)\right) - F_j\left(\bar{y}_{k-1} + \left(\Delta_{j+1} + \cdots + \Delta_{k-1}\right)\right)\right] \\ &\quad + F_k\left(\bar{y}_k\right) - F_j\left(\bar{y}_{k-1} - \Delta_k\right) \\ &\quad + \sum_{j=k+1}^{N} \left[F_j\left(\bar{y}_k - \left(\Delta_{k+1} + \cdots + \Delta_j\right)\right) - F_j\left(\bar{y}_{k-1} - \left(\Delta_k + \cdots + \Delta_j\right)\right)\right], \end{aligned} \tag{34}$$

where Δ_k are functions of $\bar{y}$ as defined by (30) and (31). Define the excess demand functions as

$$Z_k\left(\bar{y}\right) = \max\left\{0, Q_k\left(\bar{y}\right)\right\} - m_k, \tag{35}$$

where the infeasible negative demands allowed by (33) are removed. The nonnegativity constraint is binding if the upper bound $\bar{y}_k$ is below the lower bound $\underline{y}_k$ implied by $\bar{y}_{k-1}$.

[10] Equivalently, the allocation can be described in terms of the lower bounds $\underline{y}_k$ or prices Δ_k, where $k = 1, \ldots, N$. This choice of independent variables is a matter of convenience.

Each household has zero net demand for housing units, because not trading amounts to demanding your own house. Thus $Z_N(\bar{y}) = -\sum_{k=0}^{N-1} Z_k(\bar{y})$ and market equilibrium can be defined as

$$Z_k(\bar{y}) = 0 \text{ for } k = 0, \ldots, N-1, \tag{36}$$

or, in vector form, $Z(\bar{y}) = 0$. This is a system of N equations in N unknowns.

Consider the best response of an imaginary "player" who attempts to minimize the absolute value of excess demand for type$-k$ houses with the sole instrument of choosing the upper bound $\bar{y}_k$. (The top house type, where $\bar{y}_N = \bar{\theta}$ by construction, does not have a "player" representing it.) Use $\bar{y}_{-k}$ to denote the vector $\bar{y}$ without the k:th element. The best response function is

$$a_k(\bar{y}_{-k}) = \arg\min_{a \in [\underline{\theta}, \bar{\theta}]} \left| Z_k(a|\bar{y}_{-k}) \right|. \tag{37}$$

It is straightforward to show that Z_k is continuous and increasing in a, but there may be a flat region. This happens when a is below the lower bound $\underline{y}_k$ determined by $\bar{y}_{-k}$. Then demand for $k-$type houses is zero, and $Z_k(a|\bar{y}_{-k}) = -m_k$. However, the highest possible lower bound is strictly below $\bar{\theta}$, by (31) it is $\bar{\theta} - \gamma\left(\bar{\theta}|x_{k-1}, x_k\right)$, so Z_k must be strictly increasing in some interval $[\theta, \bar{\theta}]$. There are two possibilities

$$a_k(\bar{y}_{-k}) = \begin{cases} \bar{\theta}, & Z_k\left(\bar{\theta}|\bar{y}_{-k}\right) \leq 0 \\ \theta, & Z_k(\theta|\bar{y}_{-k}) = 0 \end{cases}, \tag{38}$$

where $\theta \in (\underline{\theta}, \bar{\theta})$. Thus a_k is single-valued but may not be able to eliminate excess supply for k. Hence $a(\bar{y}) = \left(a_0(\bar{y}_{-0}), \ldots, a_{N-1}\left(\bar{y}_{-(N-1)}\right)\right)$ defines a continuous vector-valued function from $[\underline{\theta}, \bar{\theta}]^N$ to itself. Thus, by Brouwer Fixed Point theorem, there exists a fixed point $y^* = a(y^*)$.

Next we prove, by contradiction, that a fixed point must be a market equilibrium. Suppose that $y^* = a(y^*)$ but $Z_k(y^*) < 0$ for some k, so $y_k^* = \bar{\theta}$. Now consider the highest such k, so that $Z_h(y^*) = 0$ for all higher types $h \in \{k+1, \ldots, N-1\}$. Since $\underline{y}_{k+1} = \bar{\theta} - \gamma\left(\bar{\theta}|x_k, x_{k+1}\right) > \underline{\theta}$ (due to overlap of autarky utilities) at least some $k+1$-types are trading down, while no one among types k or lower is trading up to houses above k. Hence $Z_h > 0$ for at least one $h < k$, a contradiction with (38). Intuitively, all demand for houses $k+1$ and higher would have to come from households endowed with those houses. Since some of them are trading down, there must be excess demand for some $h > k$.

Thus any fixed point of Z indeed defines an equilibrium allocation of our model. Finally, since γ is strictly increasing, the equilibrium allocation is associated with unique price increments Δ_k, so associated equilibrium prices p_k are unique up to an additive constant. (However, we have not shown that there is a unique equilibrium allocation).

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.

CAPLIN, ANDREW AND JOHN LEAHY (2010). "A Graph Theoretic Approach to Markets for Indivisible Goods." *NBER Working Paper 16284.*

DAVIS, MORRIS AND JONATHAN HEATHCOTE (2007). "The Price and Quantity of Residential Land in the United States." *Journal of Monetary Economics*, 54, pp. 2595–2620.

DAVIS, MORRIS AND MICHAEL PALUMBO (2008). "The Price of Residentical Land in Large US Cities." *Journal of Urban Economics*, 63, pp. 352–384.

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." *NBER Working Paper 12355.*

KREMER, MICHAEL (1993): "The O-Ring Theory of Economic Development." *Quarterly Journal of Economics*, 108, pp. 551–575.

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

ORTALO-MAGNÉ, FRANÇOIS AND ANDREA PRAT (2010). "Spatial Asset Pricing: A First Step." *Working paper.*

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 Economic Papers*, 31, pp. 60-71.

SATTINGER, MICHAEL (1993). "Assignment Models of the Distribution of Earnings." *Journal of Economic Literature*, 31, pp. 831-880.

SHAPLEY, LLOYD AND HERBERT SCARF (1974). "On Cores and Indivisibility." *Journal of Mathematical Economics*, 1(1), pp. 23–27.

TERVIÖ, MARKO (2008). "The Difference that CEOs Make: An Assignment Model Approach." *American Economic Review* 98(3), pp. 642–668.

TEULINGS, COEN N (1995). "The Wage Distribution in a Model of the Assignment of Skills to Jobs." *Journal of Political Economy*, 102, pp. 280–315."

TIEBOUT, CHARLES M (1956). A Pure Theory of Local Expenditure." *Journal of Political Economy*, 64, pp. 416–424.

VAN NIEUWERBURGH, STIJN AND PIERRE-OLIVIER WEILL (2010). "Why Has House Price Dispersion Gone Up?" *Review of Economic Studies*, 77(4), pp.1567–1606.

Figure 1.

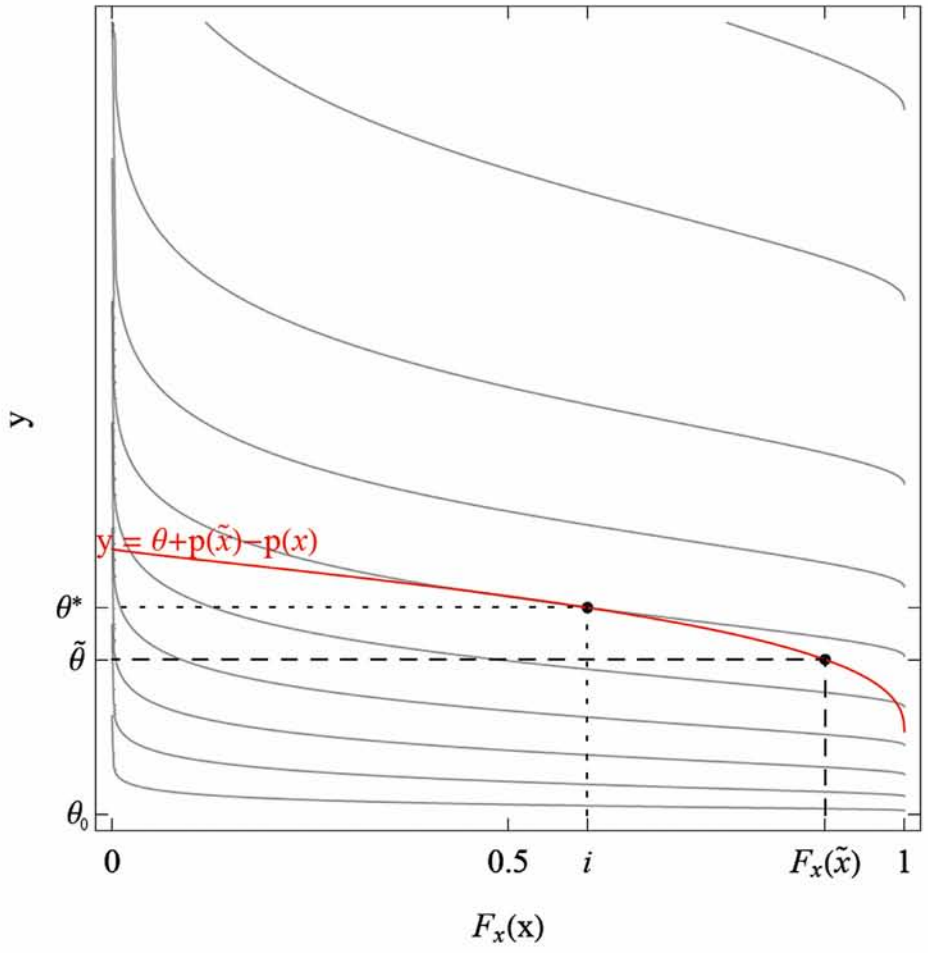

Figure 2.

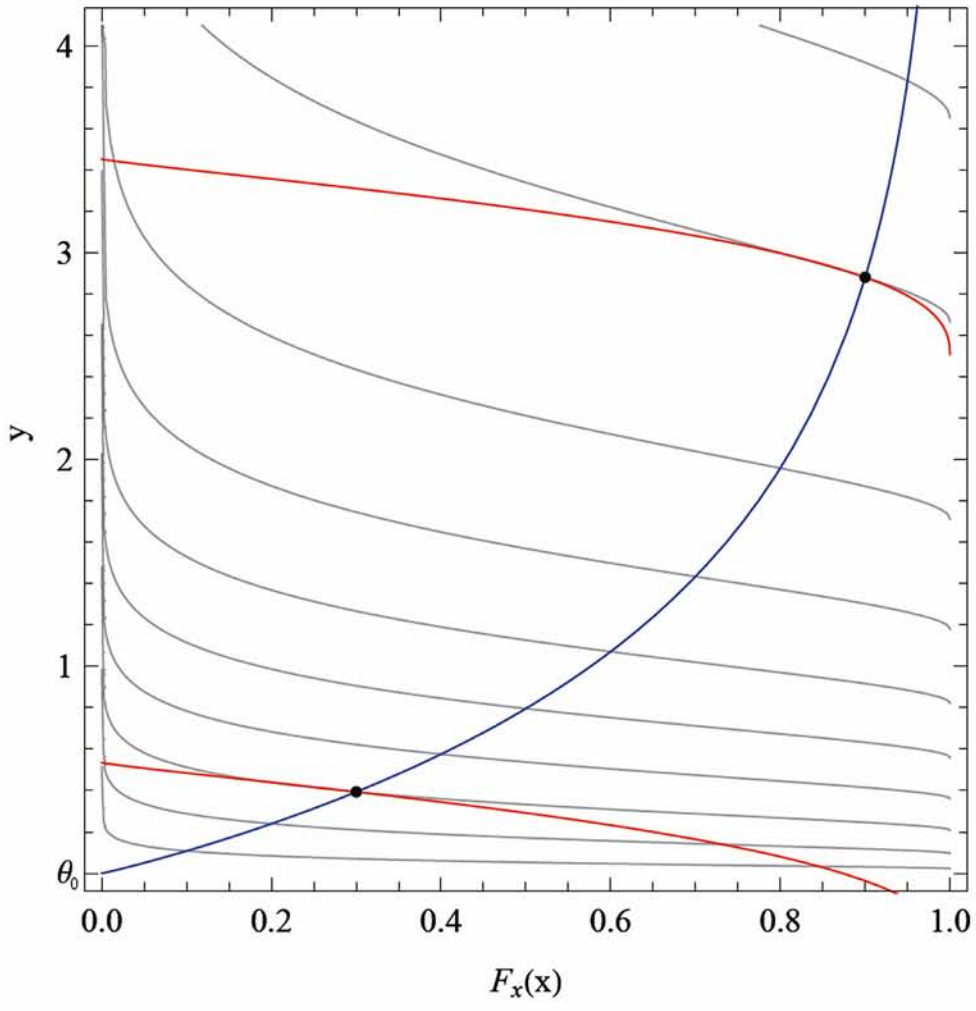

Figure 3.

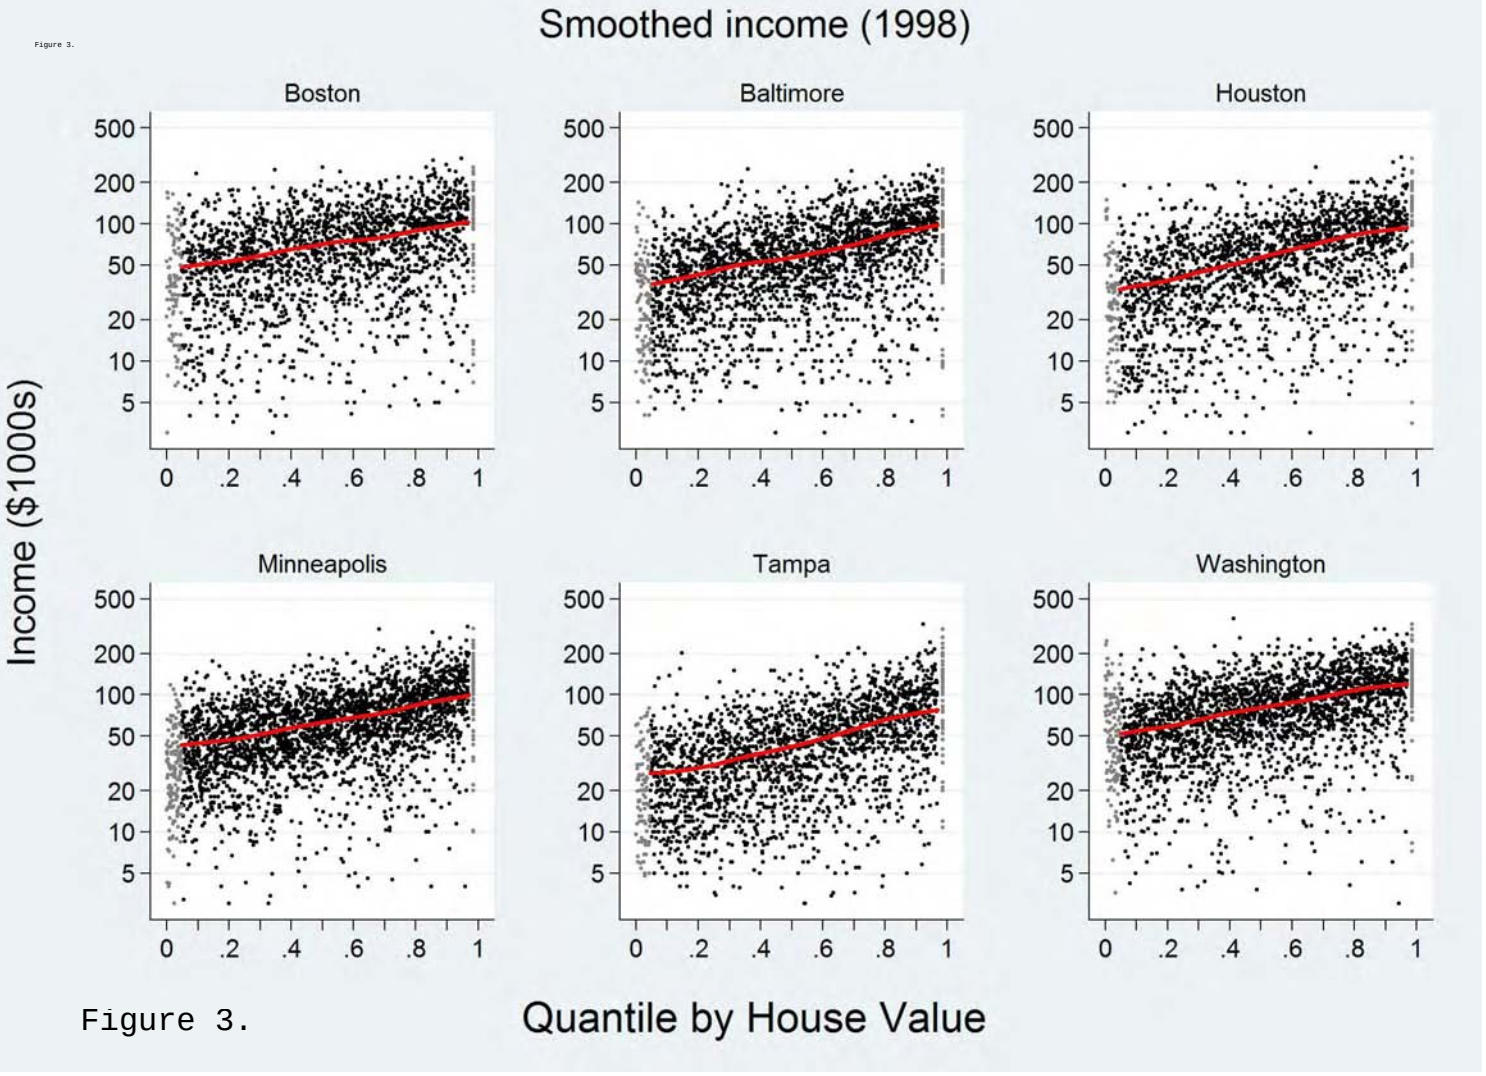

Figure 3.

Figure 4.

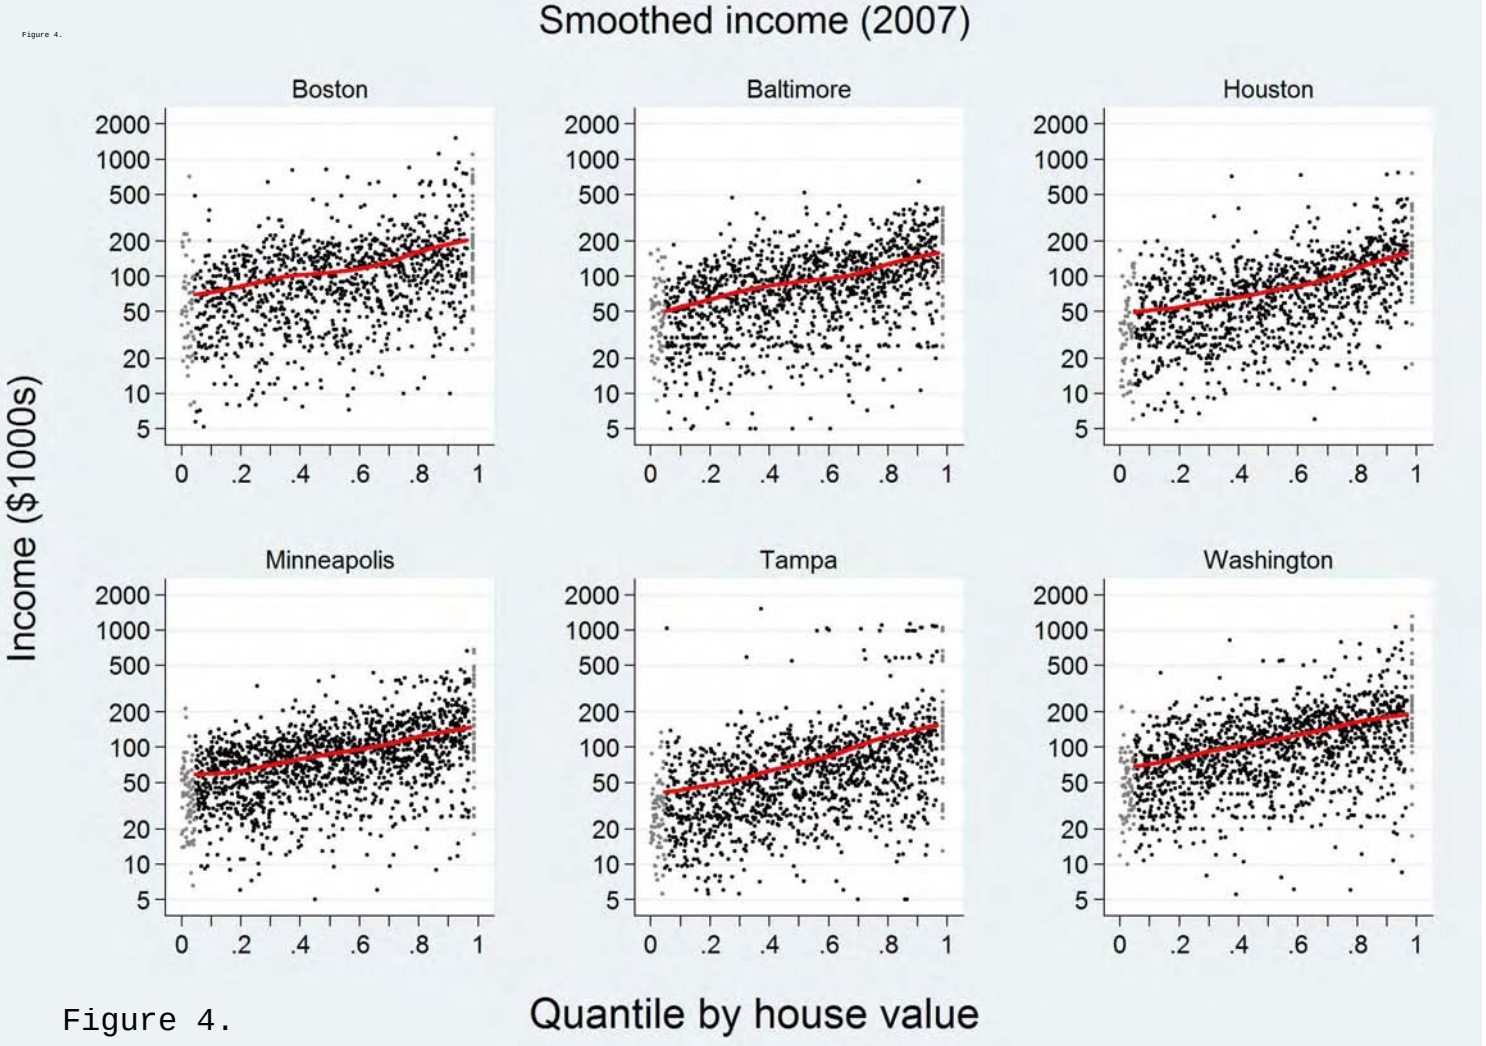

Figure 4.

Figure 5.

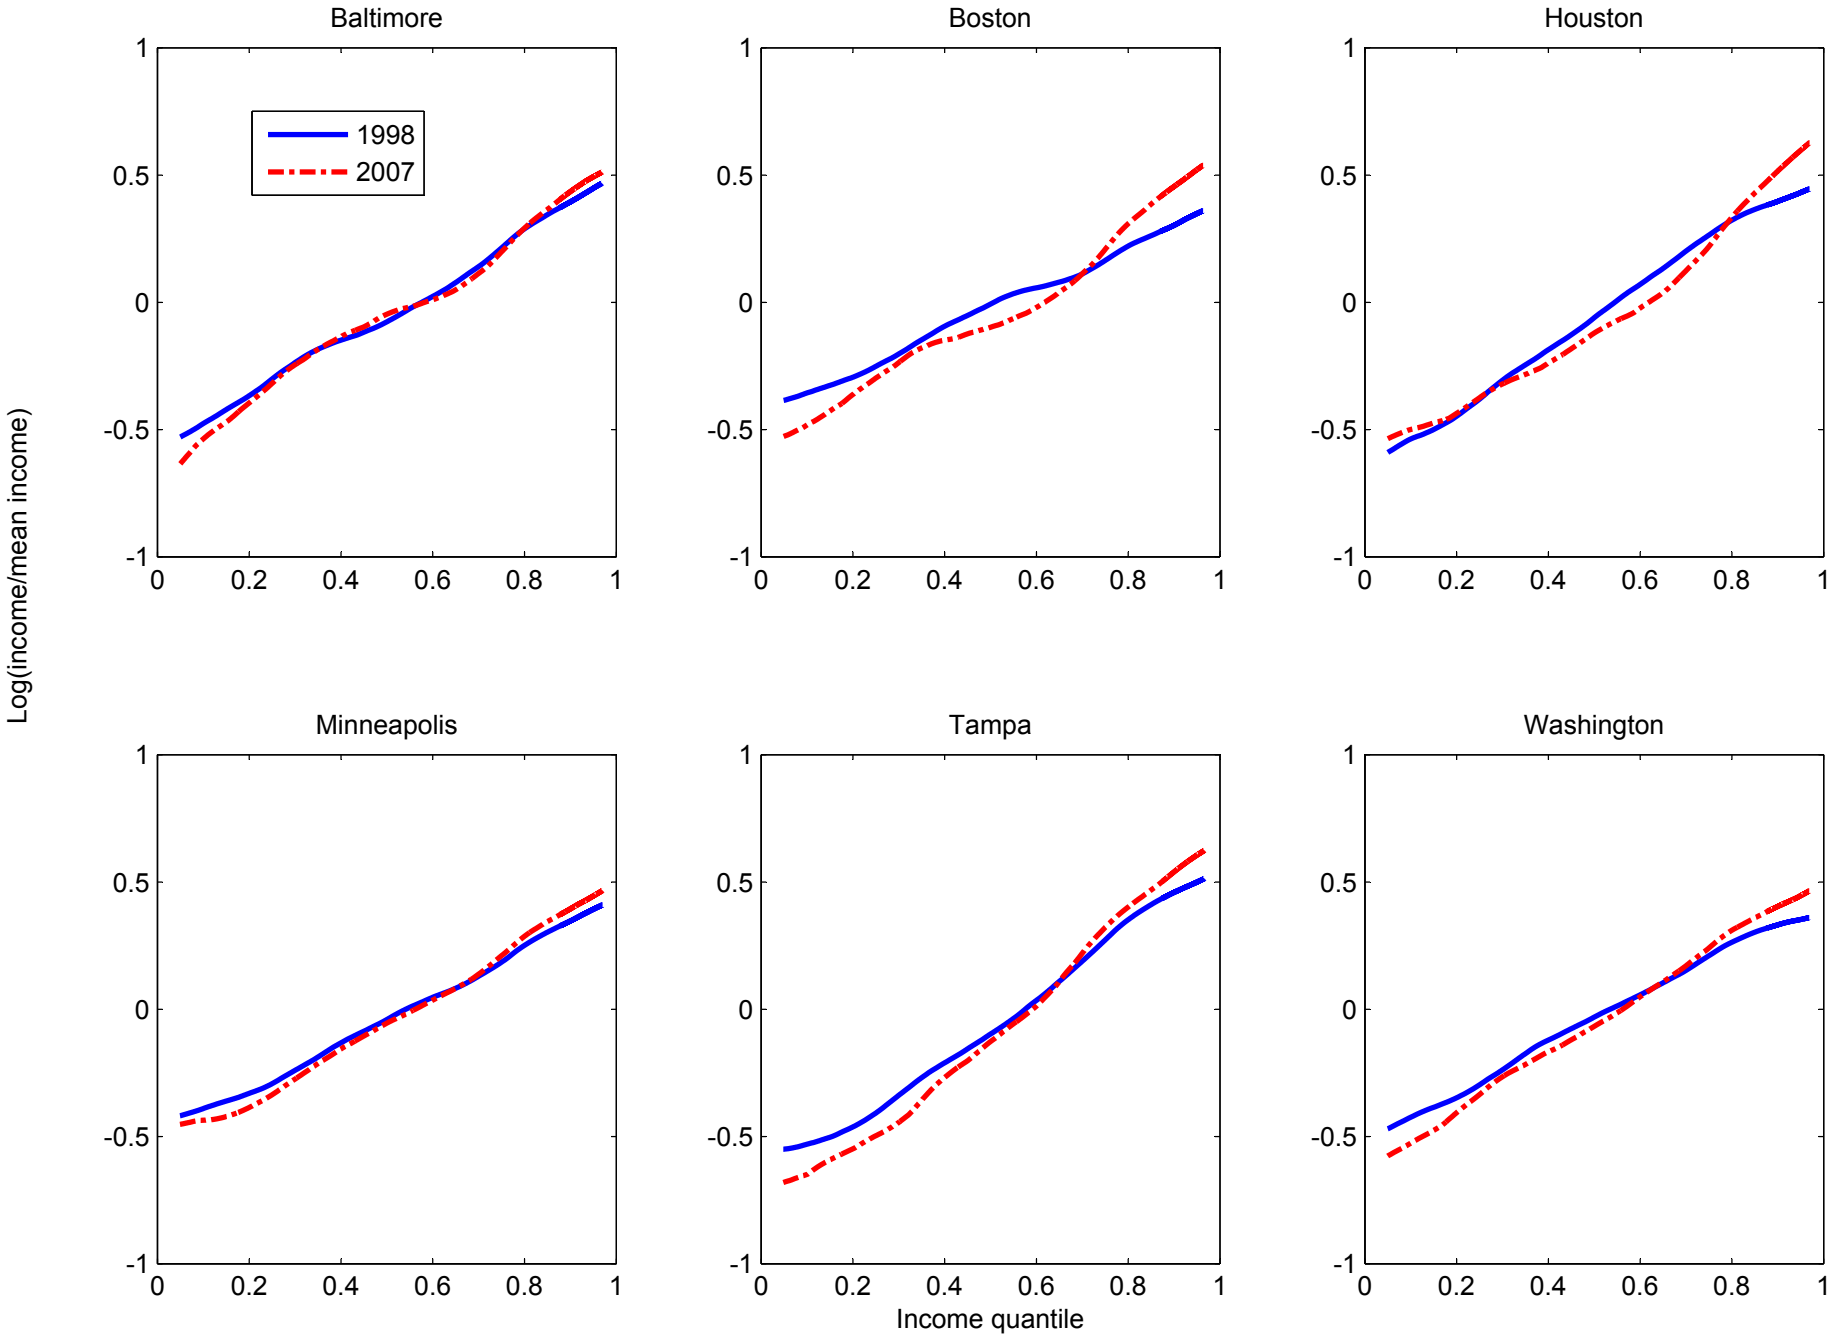
Baltimore
Boston
Houston
Minneapolis
Tampa
Washington
1998
2007
Log(income/mean income)
Income quantile
1
0.5
0
-0.5
-1
0
0.2
0.4
0.6
0.8
1

Figure 6.

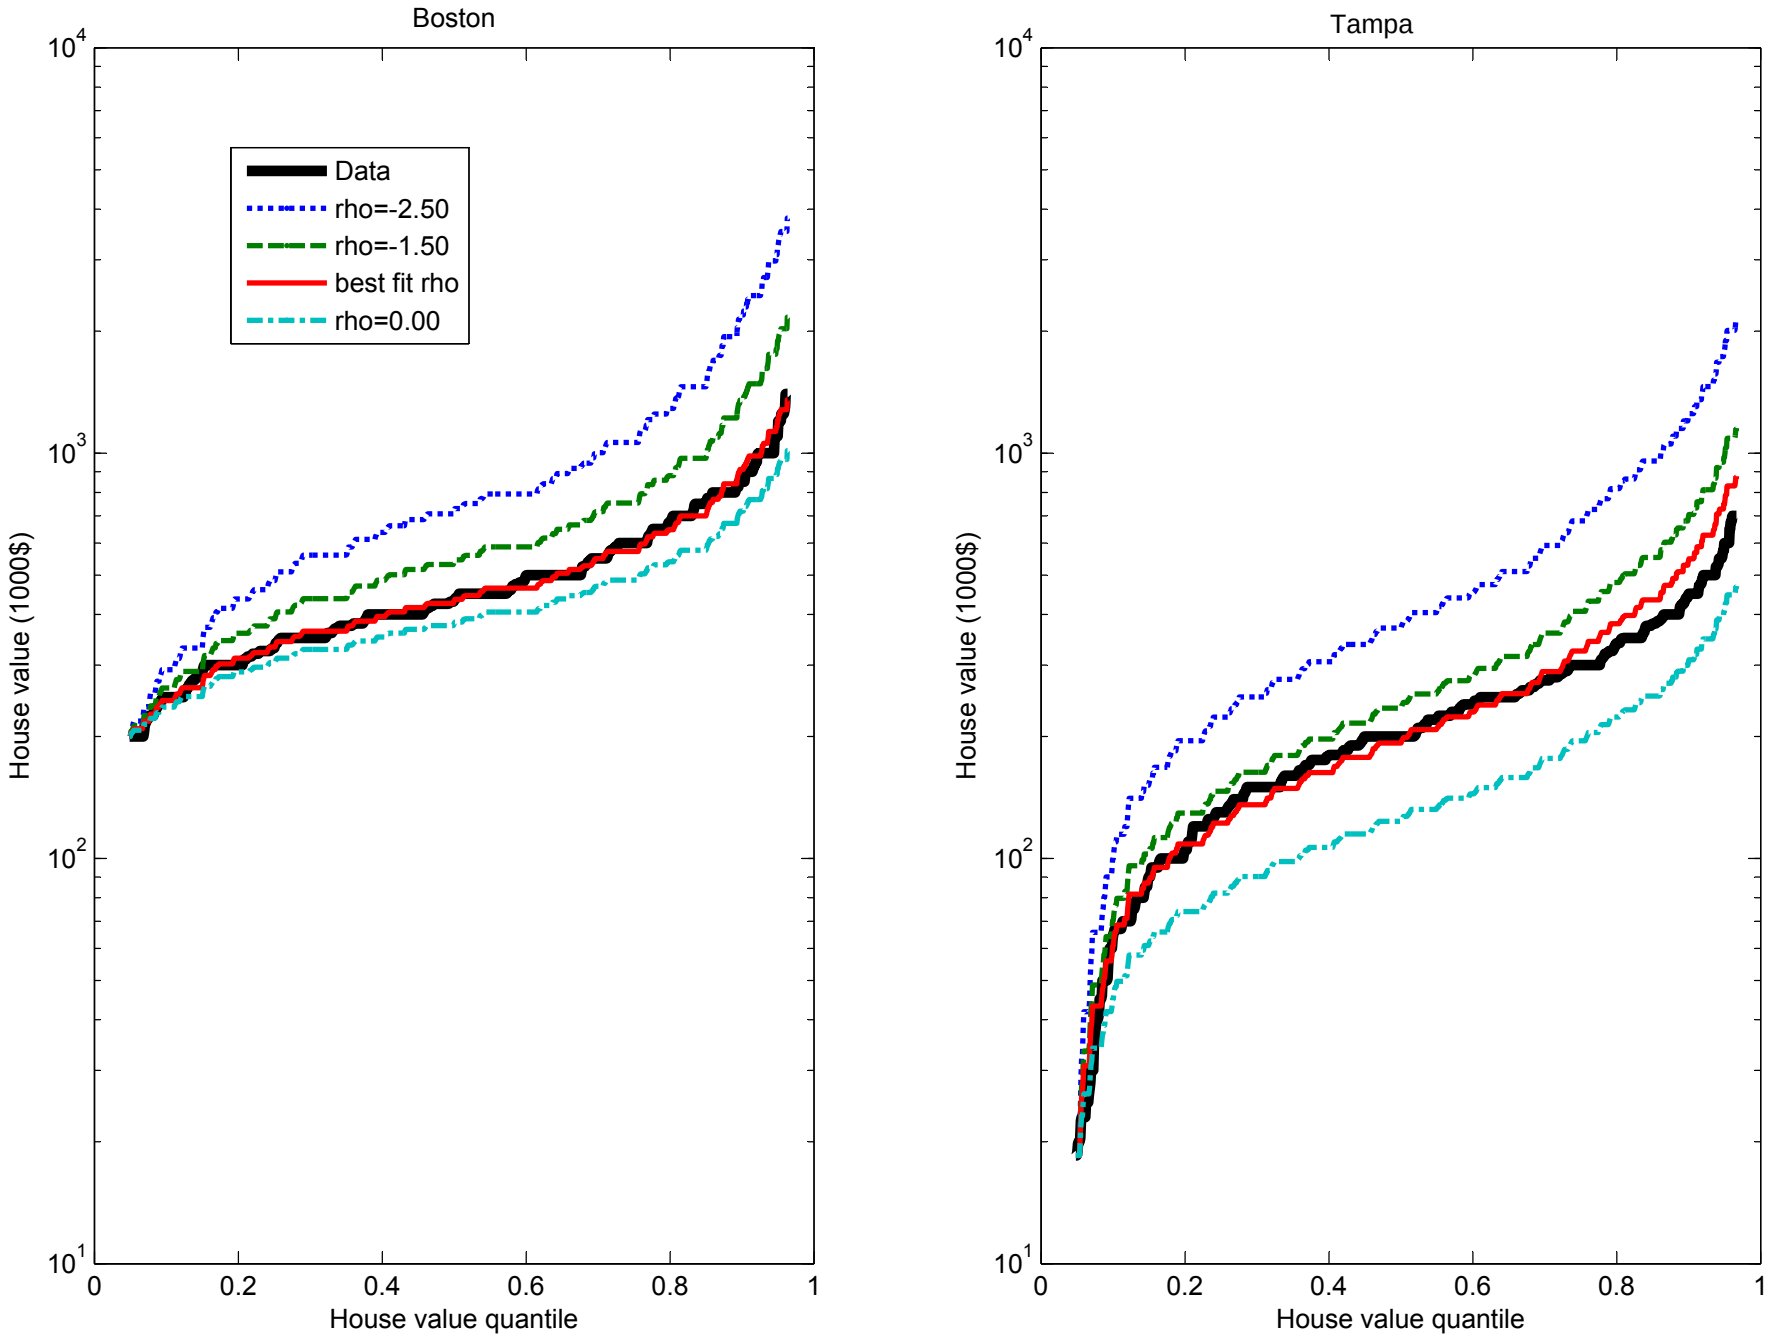
Boston
Tampa
Data
rho=-2.50
rho=-1.50
best fit rho
rho=0.00
House value (1000$)
House value quantile
10^4
10^3
10^2
10^1
0
0.2
0.4
0.6
0.8
1

Figure 7.

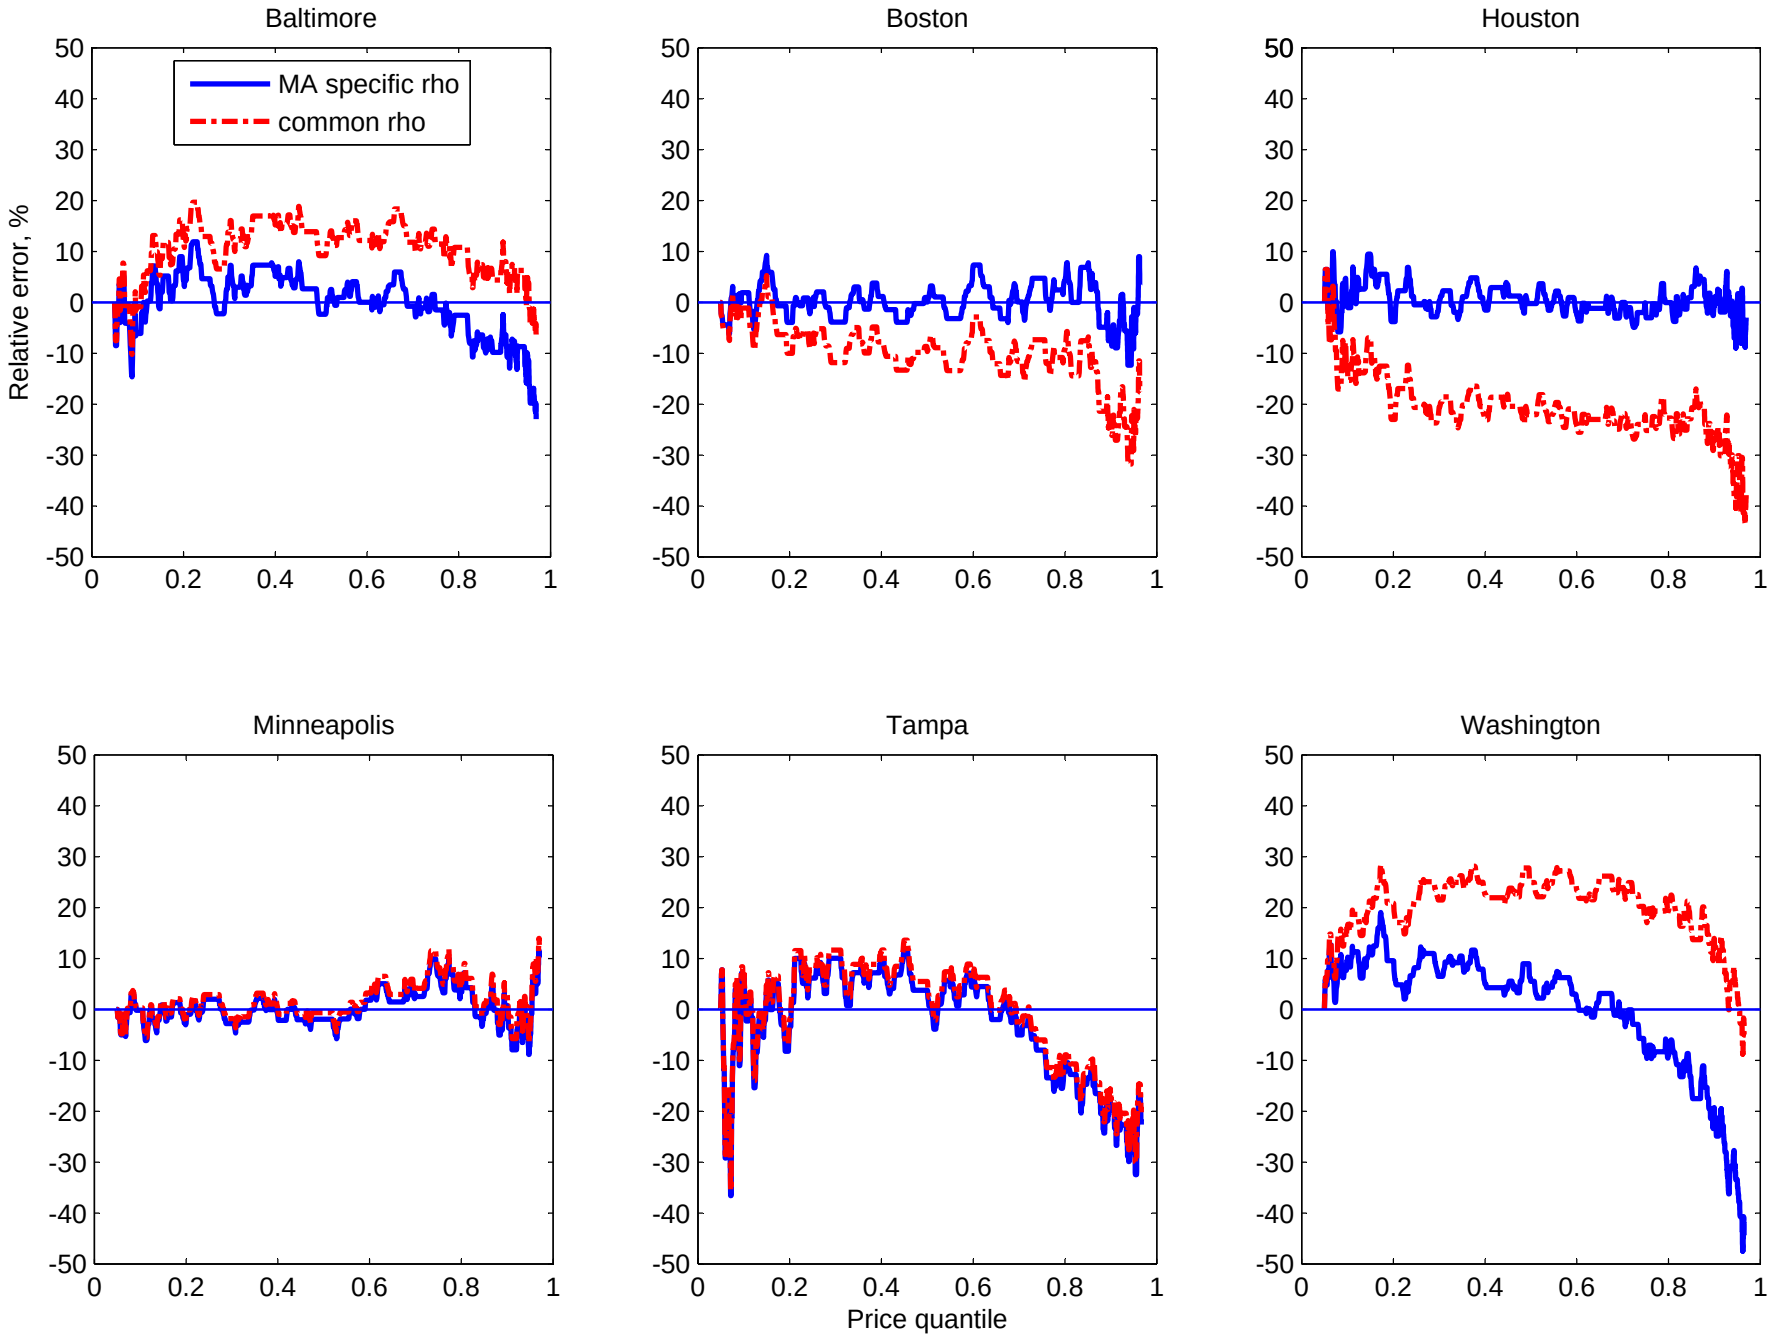
Baltimore
Boston
Houston
Minneapolis
Tampa
Washington
MA specific rho
common rho
Relative error, %
Price quantile

Figure 8.

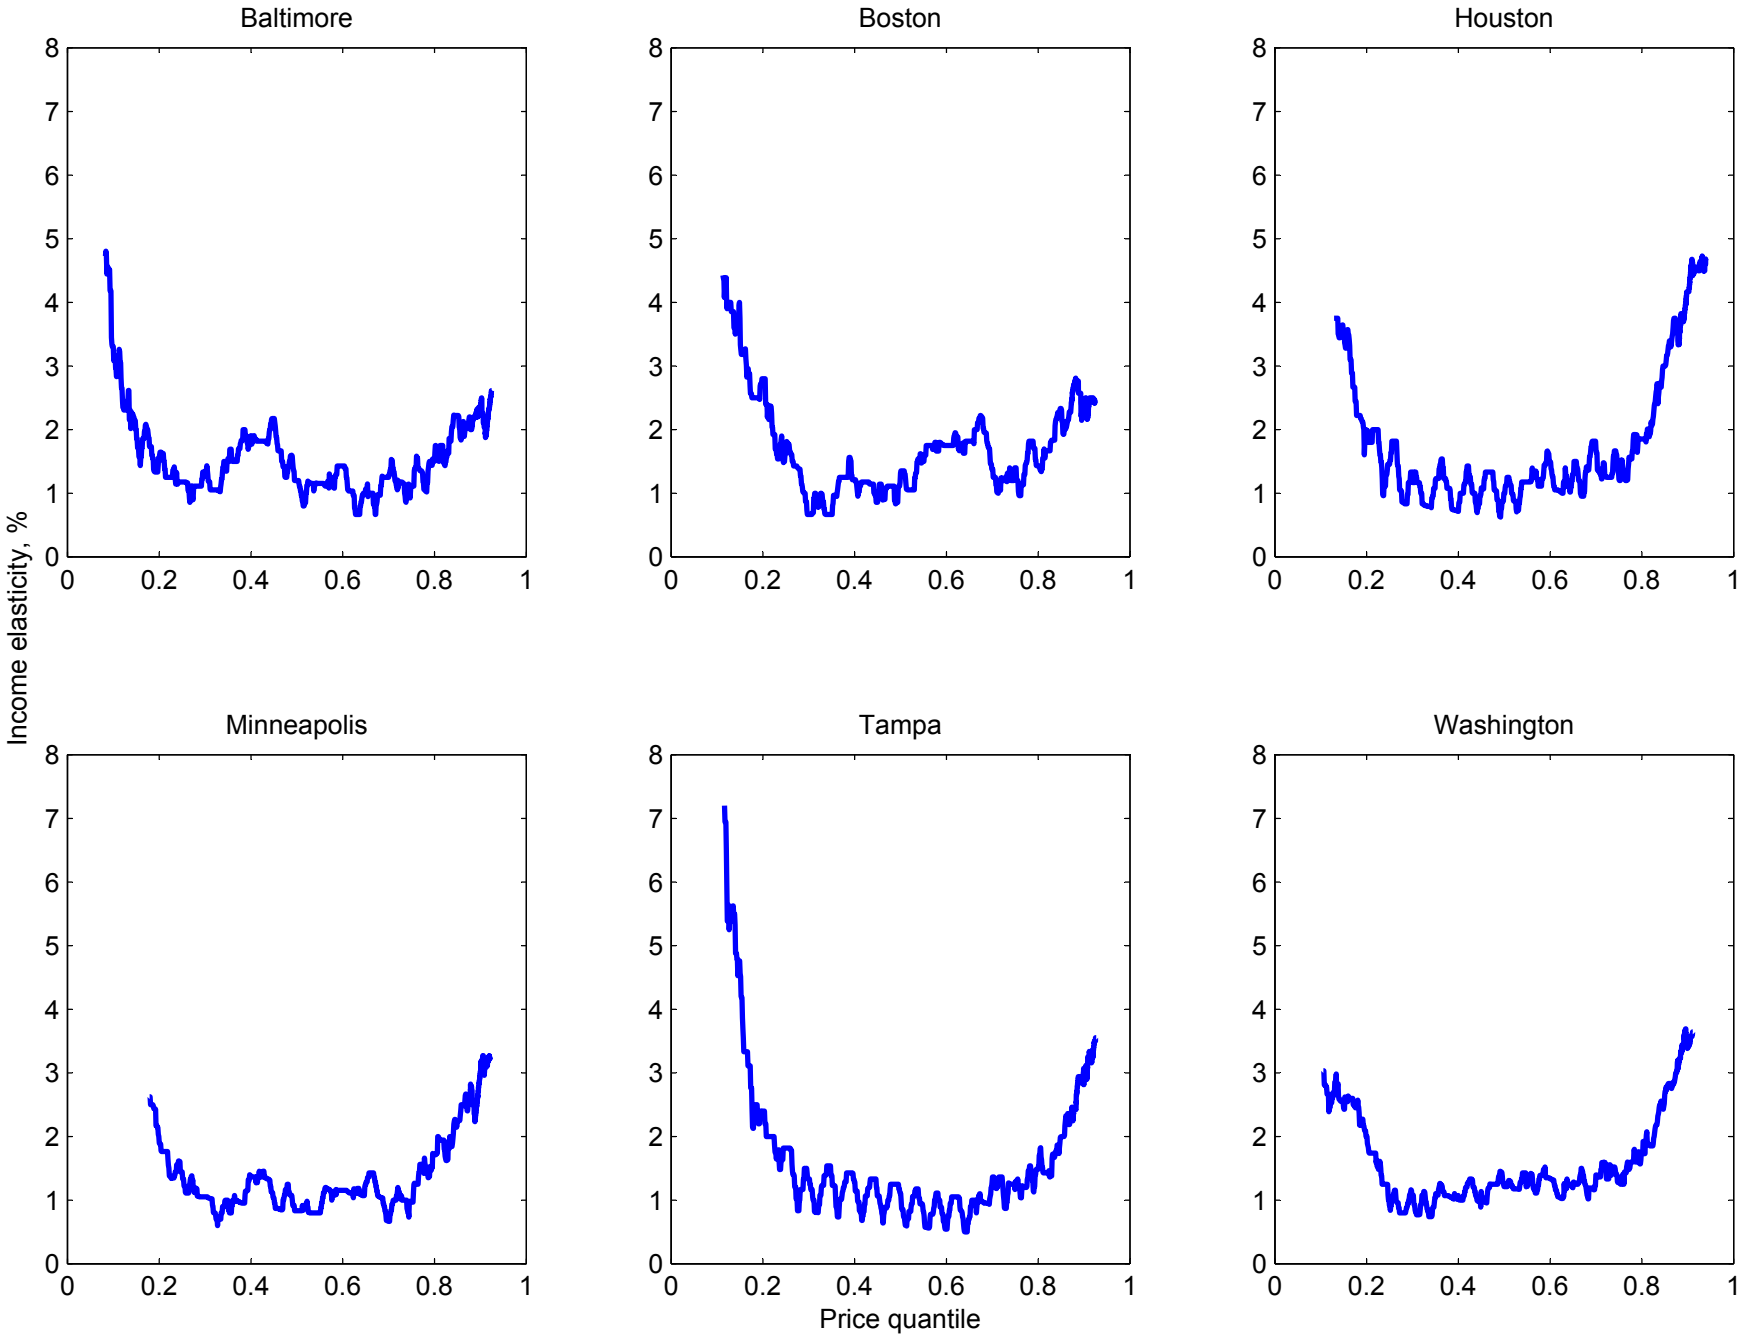
Baltimore
Boston
Houston
Minneapolis
Tampa
Washington
Income elasticity, %
Price quantile
0
1
2
3
4
5
6
7
8
0.2
0.4
0.6
0.8

Figure 9.

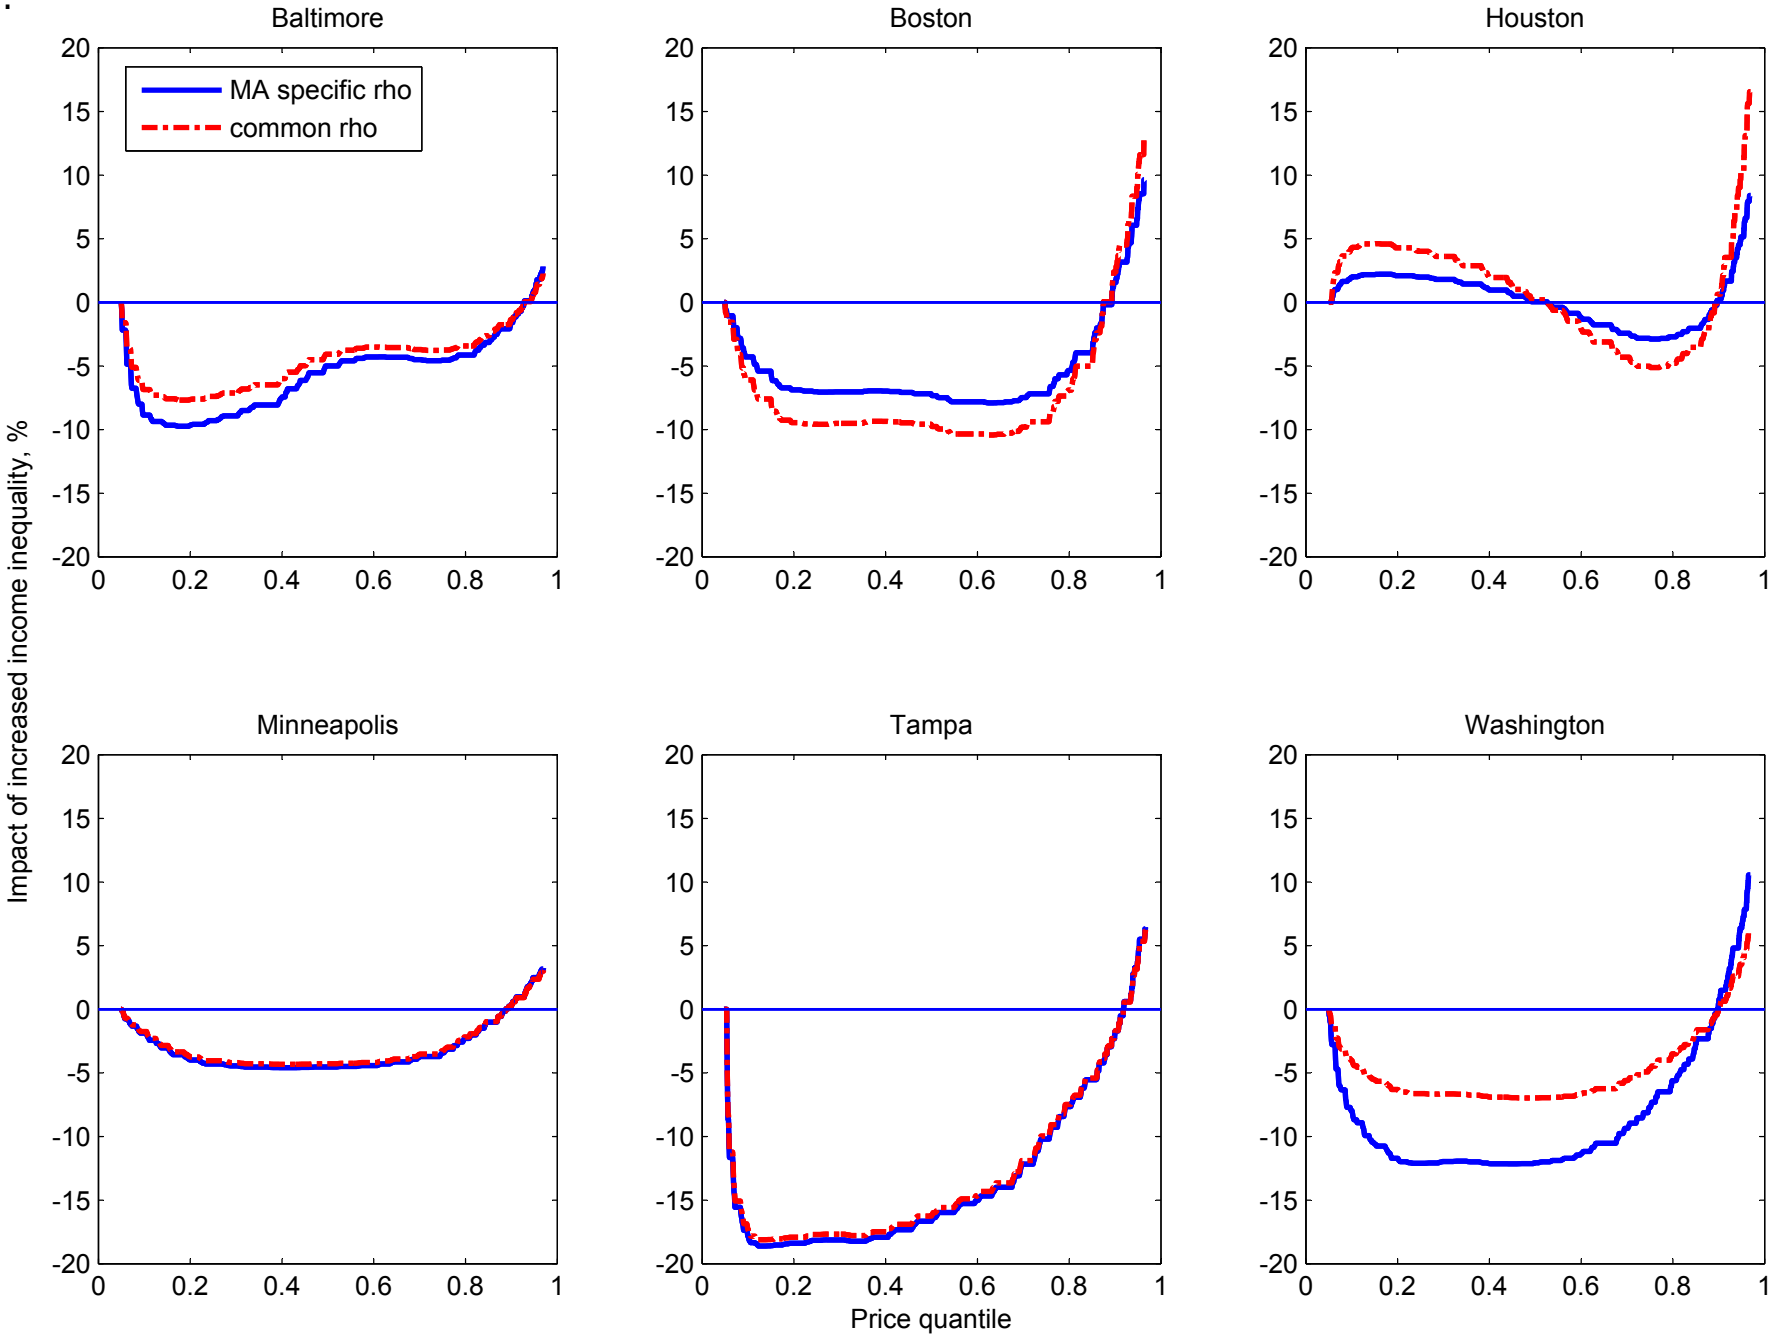
Baltimore
Boston
Houston
Minneapolis
Tampa
Washington
MA specific rho
common rho
Impact of increased income inequality, %
Price quantile
20
15
10
5
0
-5
-10
-15
-20
0
0.2
0.4
0.6
0.8
1