

Welfare Effects of Housing Transaction Taxes

Niku Määttänen
ETLA and Aalto University*

Marko Terviö
Aalto University†

January 19, 2017‡

Preliminary Paper

Abstract

We evaluate the welfare cost of ad valorem housing transaction taxes, focusing on the distortions in the matching of different houses with different households as the channel of welfare effects. We present a one-sided assignment model with imperfectly transferable utility, where households are heterogeneous by incomes, houses are heterogeneous by quality, and housing is a normal good. We calibrate the model with data from the Helsinki metropolitan region to assess the welfare impact of a counterfactual tax reform, where the transaction tax is replaced by a revenue equivalent ad valorem property tax. The aggregate welfare gain would be modest at the current 2%, but the welfare cost of transaction taxation increases rapidly with the tax rate, with Laffer curve peaking at about 9%-11% rate.

JEL: D31, R21.

*niku.maattanen@etla.fi, The Research Institute of the Finnish Economy (ETLA) and Aalto University School of Business.

†marko.tervio@aalto.fi, Aalto University School of Business.

‡We thank Essi Eerola, Miklos Koren, and seminar participants at ANU, Bocconi, CEU, EIEF, Heidelberg, Leicester, UBC, Melbourne, UNSW, and UPV/EHU for useful comments. We both thank the Academy of Finland and Terviö thanks the European Research Council (grant 240970) for financial support.

1 Introduction

Economists tend to see transaction taxes as particularly inefficient. This is especially true for housing market transaction taxes, sometimes known as “stamp duties”. The usual argument is that transaction taxes distort the allocation of houses across different households. For instance, the highly-regarded Mirrlees review (Mirrlees et al. 2011) states that “[...] transactions taxes are particularly inefficient: by discouraging mutually beneficial transactions, stamp duty ensures that properties are not held by the people who value them most”.

Our main aim is to quantify the aggregate welfare cost of a transaction tax in the market for owner-occupied housing, with the focus on distortions in the matching of houses with households as the channel of welfare effects. In our view the existing literature does not provide a satisfying quantitative evaluation of the economists’ main qualm with transaction taxes, because existing welfare analyses do not account for how transaction taxes affect the allocation of heterogeneous houses across heterogeneous households. This point is similar in spirit to that of Glaeser and Luttmer (2003) who noted that standard welfare analysis does not take into account the mismatch caused by rent control: not only does a price ceiling prevent trade, but those who do end up renting are not necessarily those with the highest valuations.

We set up a model of a closed economy with concave utility over two goods: houses and a composite good or “money”. It builds on the one-sided assignment model in Määttänen and Terviö (2014), which we augment with transaction costs. All households are endowed with an income and an indivisible house of a given quality. Not living in any house is not an option, but staying in the current house is. The set of houses is exogenous, so the inefficiency caused by a transaction tax is that the matching between houses and households may not be optimal. Since pre-existing policies are capitalized to initial house prices, changes in the taxation have also distributional effects via their impact on house prices.

The heterogeneity of demand for housing arises from differences in income (or, equivalently, from preference parameters that are additive with income). The key simplification is that households agree on the quality of houses but differ in how they view the trade-off between house quality and other consumption. While this is a stark simplification, we think it is a reasonable way to gain traction on a very complicated problem. Housing quality (which includes location and size) is a normal good, so the most important reason why some households choose to live in more expensive houses is that they can better afford them.

In the absence of ad valorem transaction costs the general price level of houses would be just “paper wealth” in our setup. Everyone has to live somewhere, so in the absence of ad valorem taxes the level or prices would be inconsequential: if all prices go up by a million this has no real effects, because the million just washes out of all transactions. Only the price differences between different types of houses are “real,” in the sense that they have implications for consumption and welfare. The right way to think about prices in a one-sided matching model is in terms of the swapping costs. For example, how much does it cost to move from a house in the 10th percentile in the quality distribution to one in the 50th percentile? In the absence of taxes this is just the difference between the two house prices. Taxes affect welfare by affecting these swapping costs. With ad valorem taxes, even the common “paper wealth component” in prices gets taxed, so the price level matters for welfare. In our model the price of the lowest quality house is exogenous; it can be interpreted as the opportunity cost of the marginal house. In a classic monocentric city model it typically represents the cost of constructing an additional unit and the opportunity cost of marginal land at the urban margin.

The reason why a household wants to trade is that something has changed since the current house was chosen. We model this something as a shock that is additive with income. The most straightforward interpretation is that the shock captures a change in permanent income, but it can also be interpreted as a preference shock that affects the trade-off between housing quality and other goods. So what we refer to as “income shocks” for brevity can be understood as including any changes in household circumstances that alter their utility trade-off between housing and other goods.

We calibrate our model to income and house value data from the Helsinki metropolitan area. For a given elasticity of substitution, we specify the distributions of housing quality and income shocks so that the resulting equilibrium distributions, together with the transaction volume, match the data closely given the current levels of transaction tax and other transaction costs. We then consider a counterfactual where the transaction tax is eliminated and replaced with a revenue-equivalent property tax. We also experiment with different transaction tax rates to see the effects of higher tax rates. These counterfactuals generate a relation between the transaction tax rate and the trading volume that is realistic in light of available empirical evidence.

The distortionary effects of a particular tax can be measured as the marginal cost of public funds (MCPF), defined as the ratio of marginal welfare cost and marginal tax revenue. We estimate a MCPF of about 1.3 at the current 2% transaction tax. The associated welfare cost is also low in absolute terms - only about 40 Euros in terms of annual non-housing consumption per household. Hence, according to our analysis,

a relatively low transaction tax is not very distortionary. However, distortions increase rapidly at higher tax rates. At a transaction tax rate of 7% the MCPF is already about 3, and the Laffer curve peaks between 9–11%. Several European countries have transaction tax rates close to these rates, so our results suggest that lowering the transaction tax rate could increase tax revenue in those countries.¹

Our welfare results are not sensitive to the assumed elasticity of substitution. This is related to the fact that we infer the quality distribution separately at level of elasticity. The quality distribution must change with the elasticity for the empirical price distribution to be consistent with our model. If house quality matters less in the utility function (lower price-elasticity) then correspondingly the quality differences between houses must be larger to rationalize the observed price dispersion as the equilibrium outcome in our model.

We also consider the distributional effects of replacing the transaction tax with a revenue equivalent property tax. Property taxes are virtually non-distortionary in our setup so the aggregate welfare effect of this reform is unsurprisingly always positive. Nevertheless, we find that a large share of households are worse off under a property tax than under a transaction tax; this is also true at relatively high transaction taxes. This could help explain the prevalence of transaction taxes in housing markets.

There are surprisingly few previous studies attempting to quantify the welfare cost of housing market transaction taxes. Moreover, they do not account for the quality distribution of houses and their matching with households. For example, Lundborg and Skedinger (1999) consider transaction taxes in a search-and-matching model where houses are observationally identical. Their estimate of the excess burden of transaction taxes is only a fraction of our estimate. Dachis, Durantou, and Turner (2012) use a two-agent partial equilibrium model together with an observed change in tax rate to estimate the impact of a transaction tax on the transaction volume in Toronto. They find that the welfare loss associated with a transaction tax of 1.1% is about 1\$ for every 8\$ in tax revenue, relative to a revenue equivalent property tax, which is close to our estimate for such a low tax transaction tax rate. O’Sullivan, Sexton, and Sheffrin (1995) and Lucas (2009) in turn focus on distortions in life cycle consumption behaviour and portfolio choice, rather than misallocation of different houses across different households, as the source of welfare effects.

¹European Commission 2012, chapter 2.2.

2 Model

The model features 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 concave utility function u . Houses are indivisible, and utility depends on the exogenous quality of the house, denoted by x . Every household is endowed with and wishes to consume exactly one house. A household's endowment of the composite good y can be interpreted as its income or "money". There are no informational imperfections, or other frictions besides transaction costs and the indivisibility of houses.

The aggregate endowment is described by the joint distribution of households over the consumption space, $\mathcal{S} = X \times \mathbb{R}_+$, where $X = \{x_1, x_2, \dots, x_n\}$ is the set of house quality levels, each owned by a mass $1/n$ of households. The distribution of income for households endowed with house type x_k has cumulative distribution $F_k(y)$, which has full support over some interval $[y_{min}, y_{max}]$, where $y_{min} > 0$, for all k . Households take prices $p = (p_1, \dots, p_n)$ as given. While u is the same for all households, its concavity combined with wealth heterogeneity implies that wealthier households have higher demand for house quality.

Consider first the problem of an individual household with endowment $\{x_h, y\}$. Denote the rate of ad valorem transaction tax by τ_T and property tax by τ_P . (In our quantitative analysis, only one of the taxes will be held nonzero at any time). There is also a fixed non-tax transaction cost ξ_k , which can depend on house type in a non-decreasing way. (The special case without taxes and transaction costs is essentially the model analyzed in Määttä and Terviö (2014).) Household h selects house type k to maximize

$$u(x_k, y + p_h - (1 + \tau_P)p_k - (\xi_k + \tau_T p_k) 1_{\{k \neq h\}}) \quad (1)$$

where the indicator function $1_{\{k \neq h\}}$ gets value zero if the household selects to live in its endowed house. Notice that household wealth $y + p_h$ is endogenous, as it depends on the price of the endowed house.

2.1 Equilibrium

In equilibrium *i*) all households choose their utility-maximizing x_k while taking p as given and *ii*) the resulting allocation is feasible. The indivisibility of houses means that the distribution of house types cannot be altered by trading, so feasibility requires that, for all k , the fraction of households choosing to live in house type x_k is equal to the fraction of households endowed with x_k .

The price of the lowest quality house p_1 is pinned down by the opportunity cost of the marginal house, which is exogenous in the model. While land use inside the urban area is heavily restricted by zoning, building at the urban-rural fringe of the metropolitan region is possible; the value of the marginal house can be interpreted as the value at best available unbuilt location.

The following lemma is useful for understanding the model.

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

That is, in equilibrium, for households that choose to trade, the ranking by wealth and by house quality must be the same. For proof, see the Appendix of Määttänen and Terviö (2014). In short, diminishing marginal rate of substitution guarantees PAM: of any two households that trade, the wealthier must live in the better house, or else the two could engage in a mutually profitable trade. (In the absence of transaction costs this means all households.) The twist here is that the ordering by wealth is endogenous, because it depends on house values. So, despite PAM, the equilibrium matching is not obvious and depends on the shape of the joint distribution of endowments. (For a proof of existence see Appendix *ibid.*)

Lemma 2 *In equilibrium, all households that trade, trade to a curve in consumption space.*

This follows directly from Lemma 1: wealth and therefore utility must be increasing in house quality or else there would still be mutually profitable trading opportunities left. Thus, for each house type x_k , there is a single interval of income levels $[y_k, \bar{y}_k]$ where i) those whose endowment is in the interval do not trade and ii) everyone who bought a house of type k is located in this interval after the trade, meaning their money (non-housing consumption) is in the same interval.

The equilibrium allocation is illustrated in Figure 1. The black curve shows where all those who trade end up; in the absence of transaction costs everyone except those “born” on the equilibrium curve would trade. The shaded area around the curve is the no-trade region: households born in the shaded “no-trade region” do not trade because it is not worth paying the transaction costs. Households above the no-trade region are relatively well endowed in money and will give up some of it in order to trade up to a better house; conversely, households below the curve are the net suppliers of quality: they are endowed with a relatively high quality house and will trade down in order to increase their consumption of the composite good. Figure 1 also depicts a budget curve for an

example household. The endowment (green dot) is above the rest of the budget curve, because even trading to a very similar quality house would entail a significant transaction tax burden.




Figure 1: Consumption space, with quantile i of house quality on horizontal and non-housing consumption y on vertical axes. The no-trade region is depicted by red lines, and the post-trade consumption bundles of traders by the black “curve.”. The green dot above the no-trade region depicts an example endowment and the blue dots depict the entire budget set for this household, with the post-trade bundle highlighted in green. This example is solved for a joint log-normal distribution with $\sigma_x = \sigma_y = \text{Corr}(x, y) = 0.5$, log-utility, transaction tax $\tau_T = 0.04$, and with no other transaction costs.

Preference heterogeneity The model admits a simple type of preference heterogeneity with almost just a relabeling. The second argument of the utility function can be interpreted as including an additive household-specific preference parameter. The model and equilibrium conditions remain the same. In terms of the common utility function u

the utility of household h is

$$u_h(x, y) = u(x, y + \epsilon_h) \quad (2)$$

This formulation allows households of the same income level to have different demand for housing versus non-housing, while still agreeing on the relative quality of different houses. Household “preference” can be due to demographic factors, such as family size, as well as tastes. A positive preference shock will have the same effect on housing demand as a positive income shock: it moves the household higher up in endowment space and so makes it demand higher quality housing.

CES utility For the quantitative exercise we assume CES utility,

$$u(x, y) = \left(\alpha x^{\frac{\epsilon-1}{\epsilon}} + (1-\alpha) y^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}, \quad \text{where } \alpha \in (0, 1), \quad (3)$$

with Cobb-Douglas utility defined in the usual fashion at $\epsilon = 1$. When p and y are observed, then x can be solved for under a given elasticity parameter ϵ . The other CES parameter, α , is absorbed by the units of x and can thus be normalized away.

Nominal incidence Whether a transaction tax is levied on the buyers or sellers does not matter for real outcomes. However, the equivalent tax rate depends on the incidence. Recall that the marginal house is priced at an exogenous opportunity cost: there are outside owners ready to sell potential houses of type x_1 at p_1 . If buyers pay the tax then outside owners get the pre-tax price p_1 , whereas if sellers have to pay the tax at rate τ_s then the pre-tax price has to be $p_1/(1 - \tau_s)$ for the outside owners to be indifferent. Therefore a tax levied on the sellers at rate τ_s is nominally equivalent to a tax levied on the buyers at rate $\tau_b = \tau_s/(1 + \tau_s)$. It is straightforward to check that, after this adjustment, all after-tax prices and tax revenues are unaffected by nominal incidence (for all possible trades, not just those involving x_1). In our notation the tax is paid by the seller, as is the case in Finland.

2.2 Welfare

Our measure of welfare effects is based on compensating variation for changes between tax regimes. That is, we take a baseline tax regime and its associated equilibrium prices and ask how much additional money a household with a given endowment would have to be given in the baseline economy to be equally well off as in the comparison economy, taking into account that not just taxes but also equilibrium prices will be different in the two economies.

A natural baseline economy for our quantitative analysis will be the actual with 2% transaction tax and no property tax, and the comparison with the counterfactual of a revenue neutral property tax is of specific interest to us. However, most comparisons we make are between two counterfactuals: we consider the baseline economy at a range of counterfactual transaction tax rates and compare them each with their revenue equivalent property tax regime. These revenue neutral comparisons allow us to interpret changes in aggregated household welfare as total welfare effects.

Consider a household with an endowment $\{x_h, y\}$. We keep endowments and preferences fixed throughout, and use \mathcal{T}^i to denote economies with particular tax regimes $\{\tau_T, \tau_P\}$ and their associated equilibrium price vectors p . In the baseline economy where $\mathcal{T}^0 = \{\tau_T^0, 0, p^0\}$ the household will in equilibrium consume some $\{x^0, c^0\}$. In the comparison economy where $\mathcal{T}^* = \{0, \tau_P^*, p^*\}$ the same household would in equilibrium consume some other bundle $\{x^*, c^*\}$. This household's welfare gain from a policy reform of replacing \mathcal{T}^0 with \mathcal{T}^* is

$$M_h(c^0) = M \text{ s.t. } \left\{ u(x^0, c^0 + M) = u(x^*, c^*) \right\}. \quad (4)$$

Our measure of the aggregate welfare effect of a change in regime is simply the average of compensating variation M_h over all households,

$$W = \frac{1}{n} \sum_{h=1}^n \int_y M_h(y) f_h(y) dy, \quad (5)$$

where f_h is the income distribution for those endowed with house type h .

2.3 Finding the equilibrium

We solve for the equilibrium numerically. Finding the equilibrium is complicated by the fact that transaction costs create a discontinuity in the budget set. Households can avoid the transaction costs by choosing to consume their endowments. Appendix A explains how to determine the demand for each house type given the price vector. We find the equilibrium price vector numerically by using a standard root-finding algorithm.

3 Calibration

The main purpose of our calibration is to quantify the aggregate welfare cost of housing transaction taxes at various levels of the tax rate. We calibrate the model using data from the Helsinki metropolitan area and then use it to conduct policy experiments with counterfactual tax rates. We solve for the equilibrium allocation at each transaction tax

rate and define the welfare cost of taxation for each household as its “willingness-to-pay” (4) to switch from its equilibrium allocation in a world with the transaction tax to its equilibrium allocation in the world without the tax.

Calibrating the model means specifying an initial joint distribution of incomes and house qualities that is realistic and conforms to the assumptions about the initial endowments of the model in Section 2. Mapping our static model to the dynamic world requires some interpretation. In what follows, we first describe the general idea of the calibration and then provide details on the data and the implementation of our calibration procedure.

3.1 General idea

The starting point is that we observe the joint distribution of house prices and household incomes. We interpret the cross-sectional data as reflecting the equilibrium of our model. In order to understand the general idea behind our calibration, it is helpful to first understand how the calibration would work if there were no transactions costs. Throughout this section, we also assume that the number of house types is very large so that the post-trade region (the shorter vertical intervals in Figure 1) can be thought of as a curve.

Our first task is to infer the house qualities. In the absence of transaction costs, all households should be located on the post-trade curve. This curve can be estimated as the empirical relation between house prices and a proxy measure for non-housing consumption. For a given elasticity parameter ε of a CES utility function, there exists a unique distribution of relative house qualities that rationalizes the observed relationship between house prices and non-housing consumption as the competitive equilibrium of our model (see Määttänen and Terviö (2014) for more detail). For a given value of ε , we can thus infer the implied house qualities (up to a multiplicative constant, which does not affect the welfare analysis).

In order to determine the initial “pre-trade” distribution, now consider an expanded model period which has the following three stages. In stage 1, all households are content with their current bundle of house type and non-housing consumption. In the absence of transaction costs, this means that they are all located on the post-trade curve, which we have already estimated. In stage 2, every household with a house of type x_k receive an income shock, which is drawn from a smooth distribution F_k . After the shock, households with a given house type will have a nondegenerate distribution of incomes y . At this stage the situation conforms to the assumptions about the initial endowments of the one-period model in Section 2 and we can use it to conduct our policy experiments. In stage 3, households have the opportunity to trade and the market for every house type clears at

equilibrium prices. (The model is static, so households do not take into account that they face more shocks in the future.) In the end, all households are content with their bundle of house type and non-housing consumption, i.e. they are all on the post-trade curve again.

The question is then how to determine the income shocks. The key idea behind our calibration is to choose the income shocks so that the resulting post-trade distribution in the end of stage 3 is close to the distribution in stage 1. In other words, we assume that the equilibrium distribution reflects a stationary distribution that would be repeated if the income shocks were drawn from a time-invariant distribution.

The data has of course been generated in the presence of transaction costs. Transaction costs make the calibration procedure conceptually more complicated and we need to make certain simplifying assumptions to deal with them. The first issue relates to the estimation of the post-trade curve. In the presence of transaction costs, the relation between average non-housing consumption and house quality in the model does not exactly correspond to the post-trade curve, because not all households trade. While non-traders should all be in the no-trade region, many of them are off the post-trade curve (see Figure 1). However, since the post-trade curve is strictly contained in the no-trade region, the average non-housing consumption should still be a good approximation of the post-trade curve. We assume that the estimated relation between house quality and non-housing consumption gives us the post-trade curve.

The second issue concerns the inference of the house qualities. Accounting for transaction costs causes a small change in the inferred house qualities, given the estimated post-trade curve (see Appendix B). A third issue relates to the distribution of endowments in stage 2. For simplicity, we assume that all households are on the estimated post-trade curve in stage 1. Due to transaction costs and income shocks, the stage 3 distribution is then necessarily different from the stage 1 distribution, as some households (those who receive a relatively small income shock) choose not to trade. That is, we can only approximate the stationary equilibrium of the model. We think this is a reasonable simplification, partly because the actual transaction costs in Finland are relatively low.

While transaction costs force us to make some stark simplifying assumption, they also allow us to pin down the size of the income shocks. Specifically, given empirically plausible transaction costs, we specify the variance of the income shocks in stage 2 so that the share of households that choose to trade in the model matches the observed level of trading in the data.

3.2 Data and transformations

We use Statistics Finland’s 2004 Wealth Survey to estimate the empirical relation between house prices and income or non-housing consumption.² We consider owner households in the Helsinki metropolitan area. The data include information on household savings, debts and income, and is mostly based on register data. Our house price measure is the current market value of a household’s main residence. We proxy non-housing consumption by disposable monetary income net of interest expenses on household debt. Household disposable income accounts for wage income, transfers, taxes and capital income, but excludes interest expenses. We take debts into account by deducting implied cost of debt service from disposable income.

Before estimating the relation between house prices and income, we need to make the units of yearly income comparable with house prices. This amounts to fixing the time horizon and the interest rate. We set the time horizon equal to the average length of stay in the current house for home owners, which is about 10 years in the data.³ Thus we measure income as the present value of 10 year’s annual income. by multiplying the annual disposable income in the data by $R = \sum_{t=0}^{T-1} (1+r)^{-t}$, where r is the annual interest rate, which we set at $r = 5\%$, and $T = 10$. This results in the empirical counterpart of the non-housing consumption y in the model. Similarly, we multiply the nominal house value by rR , to obtain the capital cost of housing for the 10-year period. We set the property tax rate at zero.⁴

In order to infer house qualities, we need a single-valued relation between house prices and non-housing consumption, which we proxy by disposable income. We first sort households according to the value of their house. There appear to be problems with data quality at the bottom of the price distribution, with some house prices observed in the range of a few thousands of Euros. For this reason, we exclude the bottom 5% of houses from the data.⁵ We lump houses to discrete quality types that represent percentiles in our data. We use \bar{p} to denote the vector of house values, with typical element \bar{p}_k standing in for the k :th percentile. We reduce the relation of income and house value to a curve by using a kernel regression to estimate \bar{y}_k as $E[y \mid F_{\bar{p}} = (k - 1/2)/100]$, where $F_{\bar{p}}$ is the empirical

²The 2004 survey is the last that includes self-reported house value. In later surveys, house values are estimated by the Statistics Finland, and we believe they are less accurate.

³Households were asked to report how long they have stayed in their current dwelling.

⁴There is a municipal property tax in Finland, but effective tax rates for dwellings are very low, partly because the taxable values are only a fraction of the market values. According to Peltola (2014), the average annual effective property tax rate in Helsinki is about 0.12%.

⁵The same problem afflicts the equivalent U.S. data (AHS). However, here, unlike in the AHS, house values are not top-coded.

CDF of house values.⁶ The resulting vector \bar{y} is the calibration target for the post-trade relation of housing and average non-housing consumption.

Assuming that the timing of trades is a Poisson process at household level, the 10-year average duration between moves implies that the share of households that engage in trade within a model period is 63%. However, the data include households that have moved to the Helsinki MA from other regions and these households are not accounted for by our one-city model. Currently these movers represent about 30% of the overall population. We therefore target a share of households that engage in trade equal to $63\% - 30\% = 33\%$. In the calibration we set the transaction tax at the actual 2% level. We assume that other transaction costs amount to 4% of the house value. That is, we set the vector of house-type specific transaction costs ξ to 4% of actual house price. (We leave non-tax transaction costs fixed at these levels in all counterfactuals; while broker fees may in reality change in response to changes in house prices we keep them fixed in order to have a clean interpretation of our estimated welfare effects.) We also transform taxes and transaction costs to reflect the model period and the way we have transformed house values and annual incomes. For instance, an ad valorem transaction tax τ translates into a transaction tax equal to τ/rR in the model.

3.3 Implementation

We assume CES-utility (expression (3)), and consider elasticity values ε at $2/3$, 1 , and $4/3$.⁷ We parameterize the distribution of income shocks as follows. Let y_k denote the stochastic non-housing endowment of a household owning a house of type k in the post-shock distribution. We assume that it is determined as $y_k = \bar{y}_k(1 + \delta_k)(\exp(\eta)/s)$, where η is normally distributed with mean zero and standard deviation σ_η , and s is a scaling term that is chosen so that the expected value of $\exp(\eta)/s$ equals one. Parameter δ_k represents a systematic component of income dynamics. We further assume that δ_k can be described as a third order-polynomial in the percentile k , so that $\delta_k = a_0 + a_1k + a_2k^2 + a_3k^3$.

We normalize $x_1 = 1$ and set p_1 exogenously at its empirical value. Since average income must equal the average income in the data, we are left with three polynomial coefficients, the shock variance σ_η , and house qualities x_2, \dots, x_{100} . We choose these parameters so that i) the resulting equilibrium house prices p are close to the empirical

⁶See Määttä and Terviö (2014) for details.

⁷The empirical estimates of this elasticity vary considerably and some of them are smaller than $2/3$. See for instance Li et al. (2015) and the references therein. We faced computational problems when trying to solve for the equilibrium with a very low elasticity of substitution. However, as we show below, our main results are not very sensitive to the assumed elasticity.

distribution \bar{p} , ii) the average non-housing consumption for households with different house types is close to the empirical relation by \bar{y} , and iii) the share of households that engage in trade is approximately 33%.

Solving for the equilibrium prices is time consuming. Fortunately, it is not necessary to solve for the prices over and over again when calibrating the model. We first infer housing qualities based on the estimated relation between household income and house prices. In the next step, we take the observed house prices as given, and find the optimal trading pattern for households with different initial housing and non-housing endowments. Given these household policies, and for any given post-shock distribution, determined by δ and σ_η , we can aggregate to find the post-trade relation between average consumption and housing, which we denote by \tilde{y} , and the share of households that engage in trade. We select the remaining four parameters so as to minimize the sum of squared differences between the elements of \bar{y} and \tilde{y} , subject to the constraint that the share of households engaging in trade equals 33%. If we are able to closely replicate the empirical non-housing consumption curve, the equilibrium prices must be close to the observed prices, which we took as given in this calibration process.

3.4 Evaluation of fit

Figure 2 illustrates the data and the calibrations with different elasticity of substitution between housing and non-housing consumption. For this and other figures that follow, we have rescaled house values, non-housing consumption, tax revenues, and welfare gains so that they are comparable with actual nominal house prices and annual consumption. The top-left panel shows the empirical price distribution p . The top-right panel shows the calibrated δ . In the calibrated model, δ_k is positive in the left-hand side of the distribution and negative in the right-hand side. Intuitively, there must be some regression toward the mean, or else the income distribution would widen with the shocks and we would not be able to replicate the estimated relation between household income and house values. The calibrated income shocks also imply that households with relatively low value houses in the post-shock distribution tend to move upwards in the quality ladder, and vice versa for those with high house values.

The bottom-left figure compares the equilibrium price distribution in the model with the empirical one by showing the percentage difference between the data and the model (a negative deviation means that the price is lower in the model). The calibrated model closely matches the empirical price distribution, except for the most valuable houses. The bottom-right figure in turn shows the estimated relation of disposable money income and

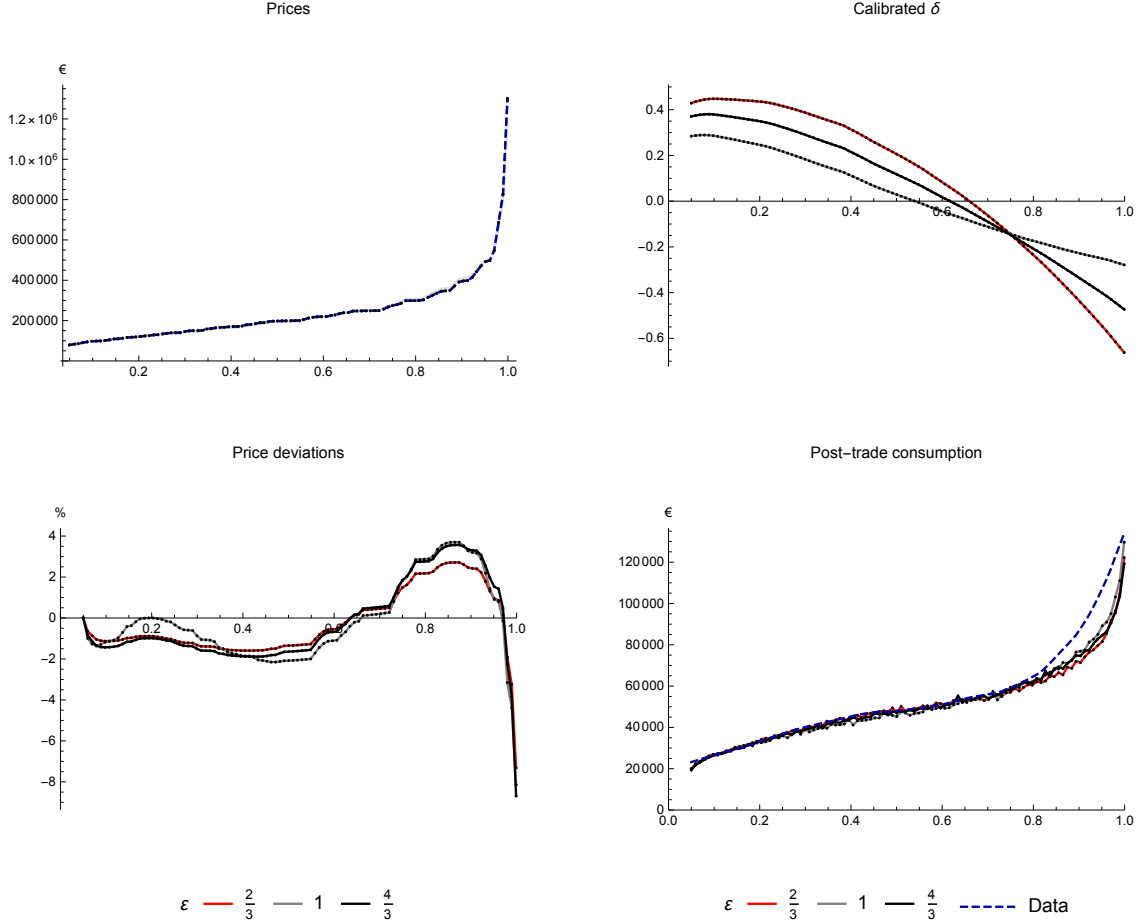


Figure 2: Empirical price distribution (top-left), calibrated δ (top-right), prices in the data vs. model (bottom-left), and average post-trade consumption implied vs. empirical relation of disposable income (annualized) and house value.

houses, \bar{y} , and the average post-trade consumption in the model, \tilde{y} . Again, the calibrated model replicates the empirical relation quite closely, especially below percentile 90 or so. Each calibrations requires a different standard deviations σ_η . The standard deviations associated with $\varepsilon = 2/3$, 1, and $4/3$ are approximately 0.40, 0.47 and 0.51, respectively. The share of households that trade matches the target 33% in all cases.

4 Aggregate effects

Figure 3 displays the main aggregate effects of transaction taxes for the three calibrations with different elasticities of substitution between housing and non-housing consumption. The top-left panel shows the annual transaction tax revenue (per owner households) as a function of the tax rate. The assumed elasticity of substitution makes a difference to tax

revenue only at higher tax rates. The higher is the elasticity, the lower is the tax revenue. However, the Laffer curve peaks around a tax rate of 9 – 10% in all cases. As mentioned in the Introduction, some European countries have housing transaction tax rates close to this range. Our results suggests that in those countries lowering the tax rate might not decrease tax revenue at all.




Figure 3: Aggregate effects of a transaction tax τ .

The top-right panel shows how the transaction tax rate affects the trade volume. For instance, increasing the tax rate from 0 to 1% lowers the trade share from about 40% to about 36%, or by about 10% in relative terms. Increasing the tax rate from, say, 2% to 4% decreases the transaction volume by about 25%. The relation between the transaction tax rate and the trade volume is almost the same at different elasticities, and seems realistic in the light of the available empirical evidence. For example, Dachis et al. (2012) find that a

1.1% transaction tax led to a 15% decrease in transactions in the first eight months after its introduction in Toronto. Hilbert and Lyytikäinen (2012) estimate from UK data that increasing the transaction tax rate from 1% to 3% would decrease the transaction volume by about 40%, while Best and Kleven (2016) find that a lowering of the transaction tax from 1% to 0% over a certain price range increases transactions (in the relevant range) by about 12%.⁸

The top left panel of Figure 4 shows the aggregate annual welfare gain from replacing the transaction tax with a revenue equivalent property tax. Hence, this curve displays the welfare cost of transaction taxes relative to property taxes, which in turn are virtually non-distortionary in the model. As explained in section 2.2, the welfare gain is measured as the average increase in non-housing consumption that would make households in the initial distribution indifferent between the equilibria associated with a given transaction tax or a revenue equivalent property tax.

Replacing the current 2% transaction tax by a revenue equivalent property tax would increase household welfare on by about 35 Euros in terms of non-housing consumption (in 2004 Euros). This is about 20% of the associated tax revenue. Arguably, then, the current transaction tax does not create large distortions. On the other hand, the welfare cost increases rapidly as we increase the tax rate. For instance, a 6% transaction tax rate generates an average annual welfare cost of around 180 euros relative to a revenue equivalent property tax.

Figure 4 also shows that estimated welfare costs of the transaction tax are very insensitive to the assumed elasticity of substitution. In order to understand this result, recall that the quality distribution is inferred separately for different elasticities. The inference is based on the relation of house prices and income. Intuitively, if house quality matters less in the utility function (lower price-elasticity) then correspondingly the quality differences between houses must be larger to rationalize the observed price dispersion as the equilibrium outcome in our model.

The top right panel shows how replacing the transaction tax with a revenue equivalent property tax affects average pre-tax house prices. Such a reform would increase the average house price. The effect is larger for higher assumed level of the elasticity of substitution. The bottom-left panel shows the share of households that would be better off if the transaction tax was replaced by a revenue neutral property tax: even for very high transaction tax rates, that share is still about one third. This is not surprising, because at least households that would not move even without a transaction tax, do not

⁸Best and Kleven (2016) analyse a temporary tax break. Their estimated overall short-run effect is larger because some of the effect is just time-shifting.




Figure 4: Welfare and price effects of a tax revenue neutral reform that replaces a transaction tax τ with a property tax. The bottom left panel shows the proportion of households with a positive welfare gain from this reform.

benefit from a reform that replaces the transaction tax with a property tax.

So how distortionary is the transaction tax? The standard way to measure this is to consider the marginal cost of public funds (MCPF), defined as the marginal welfare cost per euro of tax revenue. Figure 5 displays the approximated marginal cost associated with the transaction tax in the model. It is the rate at which the aggregate welfare cost and the tax revenue increase, as we increase the transaction tax rate by one percentage point. Here the welfare cost (or the private cost of public funds) includes the tax revenue so in the MCPF of a non-distortionary would equal to one by definition.

At a 2% tax rate, for instance, the MCPF is about 1.3. Hence, according to the model,




Figure 5: Marginal cost (left) and average cost (right) of public funds.

low transaction taxes are not very distortionary. However, the MCPF increases rapidly with the tax rate. At a tax rate equal to 6%, which is by no means an exceptionally high transaction tax rate in international context, the MCPF is already between 2 and 3, depending on the elasticity. Naturally, the MCPF approaches infinity, as the tax rate approaches its revenue maximising level. (We don't display the MCPF beyond the revenue maximising rate, where it would be negative.)

5 Distributional effects

Figure 6 displays some distributional effects of a transaction tax from the calibration where the elasticity of substitution between housing and non-housing consumption equals 1 and the current 2% transaction tax. The top-left panel shows the welfare gains to households from replacing the transaction tax by a revenue equivalent property tax or from simply eliminating it without increasing the property tax. The gains are averages by house quintiles in the initial “post-shock” distribution, which are also the same as the net wealth quintiles in the “pre-shock” distribution. Unsurprisingly, welfare gains are all positive when the transaction tax is simply eliminated without increasing other taxes. In absolute terms, households with the very best houses benefit much more than those with median quality houses. This is simply because the best houses are much more expensive than median quality houses. The average welfare gain associated with replacing the transaction tax with a revenue equivalent property tax is positive for households in house quantiles 0-60 or so and negative for households in most of the higher quintiles. These distributional effects reflect changes in house prices and the variation in the share of households that engage in trade.

The top-right panel shows the share of households that engage in trade by households' initial house quantile. The trade share varies quite a lot over the housing distribution. This is because in the presence of transaction costs, these trade shares depend on the entire joint distribution of house quality and wealth. However, the effect of removing the transaction tax on the transaction volume is similar across the entire distribution. The figure also confirms that the property tax has virtually no effect on the share of transacting households.

The bottom-left panel shows the average transaction and property tax revenue generated by each house quality percentile. Obviously, the average property tax paid increases with house quality. The relation between the house quality and the average transaction tax rate has some decreasing regions in the middle where the trade share is decreasing in the quantile. At lowest quantiles the average transaction tax paid is larger than the average property tax. Finally, the bottom-right panel shows the impact of tax regimes on house prices. Replacing the transaction tax with a property tax increases most pre-tax house prices slightly. Naturally, both the property tax and the transaction tax reduce pre-tax house prices relative to the case without taxes.

6 Conclusion

Economists often view housing transaction tax as an example of a particularly inefficient tax. The main argument is that by discouraging mutually beneficial transactions it implies that houses are not held by households that value them the most. Yet, few studies have attempted to quantify the efficiency cost of housing transaction taxes. Moreover, in the small related literature, the focus has been on how transaction taxes distort life cycle consumption behavior or portfolio choices, rather than the allocation of different houses across households.

In this paper, we focus on cross-sectional misallocation of houses as the channel of welfare effects. Our main question is what is the efficiency or welfare cost of an valorem transaction tax in the market for owner-occupied housing. In addition to efficiency, we are also interested in the distributional impact of tax regimes.

We develop a one-sided assignment model, with a fixed distribution of different houses. Households own the houses even before any changes in policy, and pre-existing policies are capitalized to house prices. Because of income shocks, some households want to move. A transaction tax affects the equilibrium allocation and prices by affecting the swapping costs.

We calibrate the model to data from the Helsinki metropolitan area housing market.

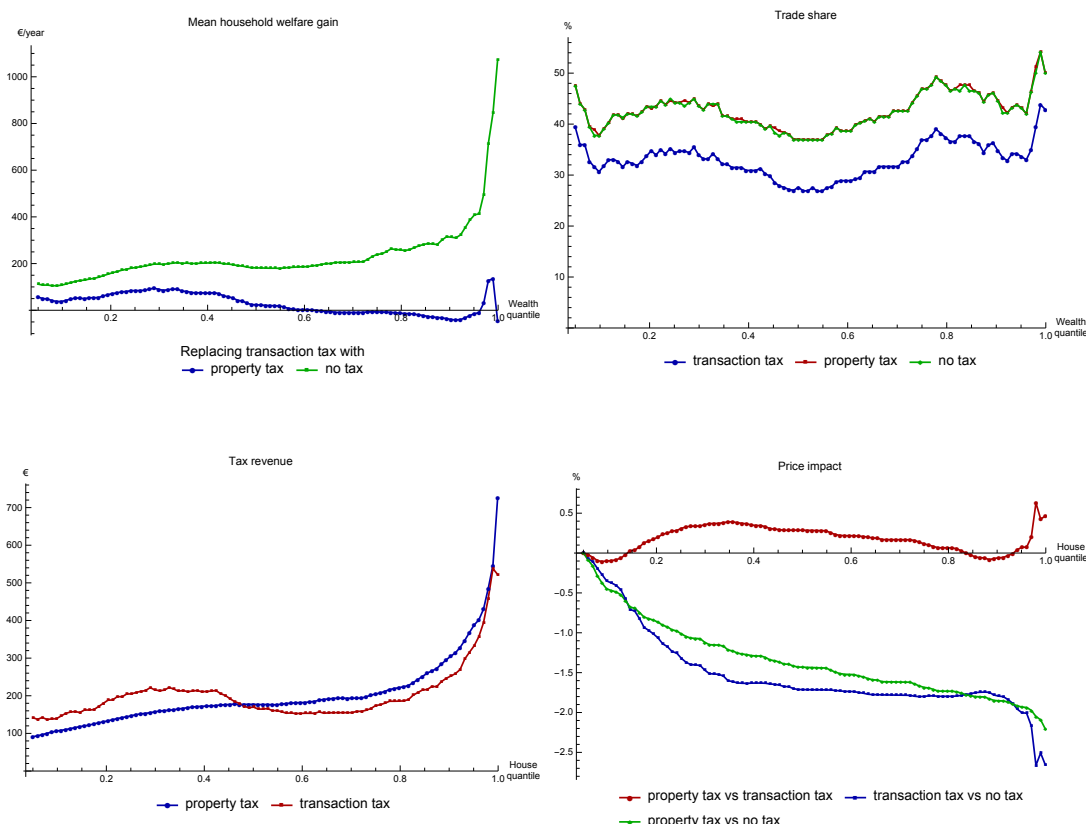


Figure 6: Distributional effects of transaction and property taxes.

We then use it to compare transaction taxes relative to revenue equivalent property taxes. We find that the distortionary effects of the transaction tax increase rapidly with the tax rate. For instance, for the 2% tax rate currently in place in Finland, the associated marginal cost of public funds (MCPF) is only about 1.2 (a MCPF equal to 1 would indicate a lump-sum tax). Hence, relatively low transaction taxes do not seem to create large distortions, even relative to the rather small tax revenue they generate. However, the MCPF increases to over 4 as we increase the tax rate to 8%, which is close to the actual tax rate in several other European countries. Related to this, we also find that the Laffer curve peaks around a tax rate of 10%.

These efficiency results are not very sensitive to the elasticity of substitution between housing and non-housing consumption, which would be hard to estimate accurately in this context. This reflects the fact we infer the housing quality separately for different elasticities based on the relation of house prices and income. Intuitively, if house quality matters less in the utility function (lower price-elasticity) the quality differences between

houses must be larger to rationalize the observed price dispersion as the equilibrium outcome in our model.

The transaction tax lowers average pre-tax house prices relative to a revenue equivalent property tax. However, this effect varies a lot across the housing quality distribution, so the transaction tax also changes relative house prices and swapping costs. These price effects in turn influence the distribution of welfare gains or losses following from a tax reform replacing transaction taxes with a property tax. In the model, about one third of households are better off with very high transaction taxes (say, above 10%), rather than a revenue equivalent property tax.

In order to focus on the allocational effects of the transaction tax, we have abstracted from other mechanisms that are also likely to be important channels of welfare effects. Therefore, we think our model quantifies a lower bound for the welfare costs of transaction taxes.

In particular, our model does not address moving to better job opportunities; we treat incomes as exogenous so the model covers only one housing market with a common labor market. The rent-or-buy decision is also not part of our model. Rent-or-buy margin can be expected to be distorted by a transaction tax, because there is no tax on changing tenants. At the same time, however, there are strong tax incentives for owning over renting in many Western countries including Finland. Finally, in our model the administration of taxes is costless and perfect. Anecdotal evidence from some countries with high transaction taxes suggest that the transaction price submitted to the authorities can be an understatement of the actual price, as transaction parties try to minimize the tax base.

References

Best, Michael and Henrik Kleven (2016) Housing Market Responses to Transaction Taxes: Evidence From Notches and Stimulus in the UK, LSE working paper.

Dachis, Ben, Gilles Duranton and Matthew A. Turner (2012), *Journal of Economic Geography*, 12(2), pp. 327-354.

European Commission (2012), Quarterly report on the euro area, 11(4).

Glaeser, Edward L. and Erzo F. P. Luttmer (2003). The Misallocation of Housing Under Rent Control, *American Economic Review*, 93(4), pp. 1027-1046.

Hilber, Christian and Teemu Lyytikäinen (2012): The Effect of the UK Stamp Duty Land Tax on Household Mobility. SERC Discussion Papers 115.

Li, Wenli. Haiyoung Liu and Rui Yao(2015). Housing over Time and over the Life Cycle: A Structural Estimation, working paper, Federal Reserve Bank of Philadelphia,

forthcoming, *International Economic Review*.

Lucas, Nancy L. (2009). Moving costs, nondurable consumption and portfolio choice, *Journal of Economic Theory*, 144, pp. 2419-2439.

O’Sullivan, Arthur , Terry A. Sexton, and Steven M. Sheffrin (1995). Property Taxes, Mobility, and Home Ownership, *Journal of Urban Economics* 37, pp. 107-129.

Mirrlees, J., S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, J. Poterba (2011). Tax by Design, Oxford University Press.

Määttänen, Niku and Marko Terviö(2009). Income Distribution and Housing Prices: An Assignment Model Approach, *Journal of Economic Theory*, 151, pp. 381-410.

Peltola, Risto (2014). Maan arvo Helsingissä: kunnallistalouden näkökulma. *Kansantaloudellinen aikakauskirja.*, 110(3), pp. 358-380.

Appendix

A Finding the equilibrium

Let us first consider how to determine the aggregate demand for each house type given an initial allocation of houses and incomes and some price vector p , where $0 \leq p_k < p_{k+1}$. For simplicity, we abstract here from the property tax and use τ to denote the ad valorem transaction tax.

The tricky part of transaction costs are that they can be avoided if a household decides to consume its endowment, which creates a discontinuity in the budget set. It is helpful to first define the bounds of no-trade intervals for a “continuous” world where the transaction tax and the non-tax transaction cost need to be paid, even if the household does not trade. Denote the bounds of the resulting **post-trade intervals** $\bar{y}_k(p)$ and $\underline{y}_k(p)$. They are solved from

$$\bar{y}_k(p) = \{y \text{ s.t. } u(x_k, y - \tau p_k - \xi_k) = u(x_{k+1}, y + p_k - (1 + \tau)p_{k+1} - \xi_{k+1})\}, \quad (6)$$

$$\underline{y}_k(p) = \{y \text{ s.t. } u(x_k, y - \tau p_k - \xi_k) = u(x_{k-1}, y + p_k - (1 + \tau)p_{k-1} - \xi_{k-1})\}. \quad (7)$$

(Recall that, in equilibrium, $\bar{y}_N = \bar{\theta}$ and $\underline{y}_0 = \underline{\theta}$.) In a world with a transaction tax these will be the post-trade allocations of those who trade, but the actual no-trade intervals will extend wider. Let’s denote the bounds of the **no-trade intervals** for those endowed with house type k as $\bar{Y}_k(p)$ and $\underline{Y}_k(p)$.

Consider the k -types, i.e., the households endowed with house x_k . As in the no-tax case, we need to find the bounds of the income interval at which a k -type will choose to

not trade. The crucial difference (which makes computation slower) is that it is no longer obvious which house type is the binding outside opportunity. For example, at the upper bound $\bar{Y}_k(p)$ the binding option is to trade up, but the house might be of type higher than $k + 1$. Intuitively, it is not worth paying a transaction cost to swap to a house that is very similar to the current house.

The following procedure can be used to find out the value of $\bar{Y}_k(p)$.

First, notice that as long as we have strictly positive transaction costs, $\bar{Y}_k(p) > \bar{y}_k(p)$. Second, notice that those who trade will end up in one of the post-trade intervals

$$x = x_j, y \in [\underline{y}_j(p), \bar{y}_j(p)]. \quad (8)$$

Households located above the income levels of the upper bound of the no-trade interval will be trading up. We go through house types x_{k+s} , starting from $s = 1$, comparing autarky with bundles at the upper bounds of post-trade allocations $\bar{y}_{k+s}(p)$. The first question is, at which income level y is the household endowed with a house x_k exactly able to pay the the price difference and the transaction tax in order to swap into a house of type x_{k+s} and have just the amount of money left over to consume at the upper bound of the post-trade interval, $\bar{y}_{k+s}(p)$. The answer is

$$\tilde{y}_{k,s} = \bar{y}_{k+s}(p) + (1 + \tau)p_{k+s} + \xi_{k+s} - p_k. \quad (9)$$

Next we need to check whether this feasible trade is at least weakly preferred to autarky. If

$$u(x_{k+s}, \bar{y}_{k+s}(p)) \geq u(x_k, \bar{y}_{k+s}(p) + (1 + \tau)p_{k+s} + \xi_{k+s} - p_k) \quad (10)$$

holds then we have found the lowest house type to which k -types trade up to; if it does not hold then we increment s by one and redo this same procedure. We keep incrementing s until we either find the **upmarket neighbor** of type k , or until we hit $\tilde{y}_{k,s} \geq \bar{\theta}$ which would show that k -types don't trade up so that $\bar{Y}_k(p) = \bar{\theta}$.⁹

Suppose we have found the lowest $k + s$ with which any k -type will prefer trading to autarky. The preference of the candidate type $\{x_k, \tilde{y}_{k,s}\}$ will almost surely be strict. Hence, now that we know s , we still need to find the exact upper bound by solving $\bar{Y}_k(p)$ as the y from equation

$$u(x_k, y) = u(x_{k+s}, y - (1 + \tau)p_{k+s} - \xi_{k+s} + p_k). \quad (11)$$

⁹If $\tau=0$ and $s = 1$, then (10) must hold strictly by revealed preference. The equivalent condition, but with lower bounds \underline{y}_{k+1} , would hold as an equality, as it is just the same as equation.

This implies that the k -type at the upper bound of the no-trade interval will trade into the interior of the post-trade interval of house $k + s$.

It is now possible that some types k do not trade at all. Then $\underline{Y}_k(p) = \underline{\theta}$ and $\bar{Y}_k(p) = \bar{\theta}$.

Finding the lower bounds of the no-trade intervals and the downmarket neighbors is analogous, but done starting from the owners of the best house type and incrementing downwards.

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, where $k > j$, if their post-tax wealth is in the same range as of those type- k households who would consume their endowment under unavoidable transaction costs τ : $p_j + \theta \in [p_k + \underline{y}_k(p), p_k + \bar{y}_k(p)]$. (Notice that both j and k would have paid the same amount of transaction costs, τp_k , but this is already deducted from the definition of post-tax wealth). At the same time, their income level θ must be outside the no-trade interval of type j .

Combining these requirements, the bounding inequalities for the interval from where j -types trade up to house $k > j$ can be written as

$$\begin{aligned}\theta &\leq \bar{y}_k(p) + p_k - p_j, \\ \theta &\geq \max\{\underline{y}_k(p) + p_k - p_j, \bar{Y}_j(p)\}.\end{aligned}\tag{12}$$

Similarly, the bounding inequalities for $j > k$ who trade down to house k are

$$\begin{aligned}\theta &\leq \min\{\bar{y}_k(p) + p_k - p_j, \underline{Y}_j(p)\}, \\ \theta &\geq \underline{y}_k(p) + p_k - p_j.\end{aligned}\tag{13}$$

Finally, the own demand by $j = k$ (the no-traders) is from the interval

$$\underline{Y}_k(p) < \theta \leq \bar{Y}_k(p).\tag{14}$$

Total demand for type- k houses is

$$\begin{aligned}Q_k(p) &= \sum_{j=0}^{k-1} \max\left\{0, F_j(\bar{y}_k(p) + p_k - p_j) - F_j\left(\max\{\underline{y}_k(p) + p_k - p_j, \bar{Y}_j(p)\}\right)\right\} \\ &\quad + F_k(\bar{Y}_k(p)) - F_k(\underline{Y}_k(p)) \\ &\quad + \sum_{j=k+1}^N \max\left\{0, F_j(\min\{\bar{y}_k(p) + p_k - p_j, \underline{Y}_j(p)\}) - F_j(\underline{y}_k(p) + p_k - p_j)\right\}.\end{aligned}\tag{15}$$

Excess demand is $Z_k(p) = Q_k(p) - m_k$, where $m_k = F_k(\bar{\theta})$ is the mass of type- k houses. Equilibrium prices are solved by finding p such that $Z(p) = 0$.

In order to find the equilibrium prices in practice, we have written a Matlab function that returns the excess demand for each house type for a given price vector and a given initial allocation of house qualities and incomes. This function first determines the post-trade and no-trade intervals described above. Using those intervals, it then determines the excess demand for each house. Since we are assuming that the income shocks are log-normally distributed, it is easy to determine the cumulative distribution F_j . We use this function together with Matlab's `fsolve` algorithm to find the equilibrium price vector.

B Inferring the quality distribution

Inference of x is based on the idea that, with given incomes and preferences the observed price difference between two neighboring house types in the quality order can be rationalized as an equilibrium price difference only with a particular quality increment. As long as there is trading there are some households that are indifferent between two neighboring house types. If there is trading then, due to continuous income distribution, for any house type j there must be households that are indifferent between moving to j or its immediate neighbor in the quality order.

Consider a household endowed with house type $k \ll j$ (or $k \gg j + 1$) that is in equilibrium indifferent between moving to j or $j + 1$. Using the notation introduced in Appendix A, this household will have after trading either the highest level of non-housing consumption among those who traded to house type j , \bar{y}_j , or the lowest level among those who traded to house type $j + 1$, \underline{y}_{j+1} . The incremental cost of trading to $j + 1$ as opposed to j is $(p_{j+1} - p_j)(1 + \tau_T) + \xi_{j+1} - \xi_j$. When inferring the quality distribution, we assume that non-tax transaction costs are a constant fraction ϕ of the equilibrium purchase price, so this cost difference can be written as $(p_{j+1} - p_j)(1 + \tau_T + \phi)$. (However, we keep the non-tax transaction costs fixed in absolute terms in our counterfactuals.) For these households we have the indifference condition

$$u(x_j, \bar{y}_j) = u(x_{j+1}, \bar{y}_j - (p_{j+1} - p_j)(1 + \tau_T + \phi)). \quad (16)$$

Under CES-utility this can be solved for

$$x_{j+1}^\rho - x_j^\rho = \bar{y}_j^\rho - (\bar{y}_j - (p_{j+1} - p_j)(1 + \tau_T + \phi))^\rho. \quad (17)$$

Everything on the right hand side is either data or parameters for which we can assume reasonable values ($\rho = \varepsilon/(\varepsilon - 1)$, ϕ). (By contrast, it would be hard to come up with a

reasonable range of values for the abstract quality measure $x_1 \dots$) With a sufficiently fine grid $\bar{y}_j \approx \underline{y}_j$ we can treat both as approximations for y_j .¹⁰ The CES inference formula under transaction costs is

$$\hat{x}_h^\rho = \hat{x}_1^\rho + \sum_{j=2}^h ((y_{j-1} + (p_j - p_{j-1})(1 + \tau_T + \phi))^\rho - y_{j-1}^\rho). \quad (18)$$

¹⁰Caveat: The non-traders of type j need not have an average y in the traders' post-trade range $[\underline{y}_k, \bar{y}_k]$, but if transactions costs are low then the no-trade region is not very wide and they should be close.