

Welfare Effects of Housing Transaction Taxes

Niku Määttänen

ETLA and Aalto University*

Marko Terviö

Aalto University†

October 10, 2017‡

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 10%.

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, LSE, UBC, Melbourne, UNSW, and UPV/EHU for useful comments. We both thank the Academy of Finland and Terviö thanks the European Research Council (grant ERC-240970) for financial support.

1 Introduction

Economists tend to see transaction taxes as a particularly inefficient form of taxation. This is especially true for housing market transaction taxes, sometimes known as “stamp duties”. The usual argument is that they 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.

We set up a model of an urban area where there is a distribution of house types of different qualities and a population of households with different housing demands. 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, and utility is concave over two goods: houses and a composite good or “money”. 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.

The heterogeneity of demand for housing arises from differences in income (or, equivalently, from preference parameters that are additive with income). Our key simplifying assumption is that households agree on the quality of houses but differ in how they view the trade-off between housing 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 subsumes 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. The price of the lowest quality house is exogenous; it can be interpreted as the opportunity cost of houses at the urban margin.

The reason why a household wants to trade is that something has changed since it chose its current house. 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.

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 low in absolute terms - only about 30 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 10–11%. 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 for each level of elasticity. For the empirical price distribution to be consistent with our model, the implied quality distribution must change with the assumed elasticity. 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. However, there are some limitations to what welfare questions we can answer in this setup. The reason is that we can only reasonably estimate relative house quality levels but not their the absolute level; hence we need to use welfare measures from which the absolute quality level factors out. In practice this means that we have to evaluate the welfare impact of a policy “ex post” at a point where each household knows its income and thus what its gains from trade would be under each tax regime. The unit elastic case is an exception; there we can estimate expected “ex ante” welfare for households that face income risk and do not yet know whether they will want to trade.

We also analyze the distributional effects of replacing a transaction tax with a rev-

¹European Commission (2012)

enue equivalent property tax. Property taxes are non-distortionary in our setup so the aggregate welfare effect of this reform is unsurprisingly positive and increasing in the tax rate. Nevertheless, we find that “ex post” a large share of households are worse off under a property tax. Furthermore, this share of losers is increasing in the tax rate. The reason is that with a high transaction tax rate the commensurate property tax is also high; most “additional” trades that are enabled by the tax reform produce only marginal gains from trade but households benefiting from these gains have to contribute a full share of the property tax. Many would be better off as non-traders under a transaction tax, more so when the tax burden is high. This could help explain the prevalence of transaction taxes in housing markets.

Several empirical studies have demonstrated that transaction taxes have a substantial effect on the trading volume. Examples include Best and Kleven (2016), Dachis et al. (2011), Hilber and Lyytikäinen (2017), Van Ommeren and Van Leuvensteijn (2005), and Kopczuk and Munroe (2015). Importantly, our model generates a relation between the transaction tax rate and the trading volume that is realistic in light of this evidence. Only a few studies have taken the next step and tried to quantify the welfare cost of housing market transaction taxes. Moreover, the existing analyses do not account for an empirical quality distribution of houses and their matching with different households. Dachis et al. (2011) use a two-agent partial equilibrium model together with their empirical estimate of the impact of transfer 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. This is relatively close to our estimate for such a low tax transaction tax rate. 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. O’Sullivan et al. (1995) and Stokey (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.

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 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 Määttänen 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 Määttä 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 $[\underline{y}_k, \bar{y}_k]$ 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.

2.2 Some Notes on the Model

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




Figure 1: Consumption space, with quantile of house quality on horizontal and non-housing consumption 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 shows an example endowment and the blue dots 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.

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.

Real vs nominal wealth In our closed economy the general price level of houses would be just “paper wealth” in the absence of ad valorem transaction costs. Everyone has to live somewhere, so across-the-board changes in houses prices are inconsequential: if all prices go up by a million this has no real effects, because the million just washes out of all possible 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.

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.

CES utility For the quantitative exercise we assume CES utility,

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

with the unit-elastic case defined in the usual fashion at $\varepsilon = 1$. When p and y are observed, then x can be solved for (up to a constant) under a given elasticity parameter ε . The

other CES parameter, α , is absorbed by the units of x and can thus be normalized away. We derive the formula for inferring x from the data in Appendix A.

Solving the model 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 endowment. We determine the equilibrium numerically. Given the initial allocation, we first determine the post-trade curve and no-trade regions depicted in Figure 1. We then aggregate to find the demand for each house type given the price vector, which we find using a standard root-finding algorithm. We explain the procedure in detail in Appendix B.

2.3 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 differ between the two economies. We measure aggregate welfare by taking the average of this compensating variation over all owner-occupied households.

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 aggregate household welfare as total welfare effects.

Consider a household with an endowment $\{x_h, y\}$. We keep endowments and preferences fixed throughout, and compare economies that differ by tax regime and their associated equilibrium price vectors p . In each tax regime either $\tau_T = 0$ or $\tau_P = 0$. In the baseline economy where $\tau_T > 0$ the household will in equilibrium consume some $\{x(y|h), c(y|h)\}$. In the comparison economy where $\tau_P > 0$ the same household would in equilibrium consume some other bundle $\{x_*(y|h), c_*(y|h)\}$. This household's welfare gain from a policy reform of switching from transaction tax regime to property tax regime is

$$M(y|h) = m \text{ s.t. } \left\{ u(x(y|h), c(y|h) + m) = u(x_*(y|h), c_*(y|h)) \right\}. \quad (4)$$

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

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

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

It is not possible to empirically identify the “quality units” of housing in our setup, only their quality relative to each other. In practice this means that we have to assign one house type some arbitrary positive quality level against which all other house qualities are measured. This puts some limits on what kind of welfare questions the calibrated model can be used to answer. Naturally, we can only make empirical claims about questions where the answer does not depend on our arbitrary choice of $x_1 = 1$.

We evaluate household welfare as a compensating variation for switching from a world with one tax regime to another. The average over households, equation (5), evaluated after the income shocks have been realized, is independent of x_1 . By contrast, the compensating variation that households would demand before knowing their shock depends on x_1 , except in the special case of unit elasticity (log utility). We now introduce some notation that helps us explain this. Consider households endowed with house type x_h and income y_h . After the income shocks are realized—but before the possible trading—their income is distributed according to the density f_h . Conditional on the realization y , it will live in house type $x(y)$ under “status quo” and in house type $x_*(y)$ under the alternative “reformed regime”, and will have (after possible swapping costs and taxes) $c(y)$ or $c_*(y)$ for other consumption.

With CES-utility and $\rho = (\varepsilon - 1)/\varepsilon$, the compensating variation from equation (4) becomes

$$M(y|h) = \left(x_*(y|h)^\rho - x(y|h)^\rho + c_*(y|h)^\rho \right)^{\frac{1}{\rho}} - c(y|h). \quad (6)$$

To see how this depends on the estimated quality levels, use the inference formula for x_h (11) on page 28 in the Appendix. The quality estimate for each house type $h = 2, \dots, n$ depends on data $y_1, \dots, y_n, p_1, \dots, p_n$, the assumed elasticity of substitution via ρ , and the quality level x_1 . It is easy to see that the constant x_1^ρ cancels out of all quality differences $x_i^\rho - x_j^\rho$. Therefore the realized welfare gain (6) is independent of x_1 ; it only depends on the data $y_i, \dots, y_j, p_i, \dots, p_j$, and the elasticity.

Now consider the expected welfare gain evaluated before the income shocks are realized. This “ex ante” definition of compensating variation has to take into account that the level of house quality and other consumption depends on the realization, and in potentially

different ways under each regime.

$$\bar{M}(h) = m \text{ s.t. } \left\{ \int (x(y)^\rho + c(y)^\rho + m)^{\frac{1}{\rho}} f_h(y) dy = \int (x_*(y)^\rho + c_*(y)^\rho)^{\frac{1}{\rho}} f_h(y) dy \right\}. \quad (7)$$

The solution of M depends in general on levels of x , because the x^ρ terms cannot be grouped in a way that would eliminate x_1 . Log-utility is the important exception: when $\rho \rightarrow 0$ then the inference formula yields the ratios \hat{x}_h/x_1 as functions of the data; see equation (12) in the Appendix. The expected welfare gain in (7) becomes $\bar{M}(h) = m$ s.t.

$$\int (\alpha \log x(y) + (1-\alpha) \log c(y) + m) f_h(y) dy = \int (\alpha \log x_*(y) + (1-\alpha) \log c_*(y)) f_h(y) dy \quad (8)$$

Subtracting $\alpha \log x_1$ from both sides shows that expected welfare depends on ratios \hat{x}_k/x_1 which in turn only depend on the data and not on x_1 . This also demonstrates that the weight parameter $\alpha \in (0, 1)$ is absorbed by the undefined x_1 and can thus be ignored.

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 regimes. 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 Appendix A). 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_h receive an income shock drawn from a smooth distribution F_h . 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 again on the post-trade curve.

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 at 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 a world with 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.

Another issue in accounting for transaction costs concerns 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 assumptions, 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. We consider owner households in the Helsinki metropolitan area. The data include register based data on household savings, debts and income and self-reported estimate of the market value of a household's main residence, which we take as our house price measure.² We proxy non-housing consumption by disposable monetary income. In the data, 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

²The 2004 survey is the last that includes self-reported house value. In later surveys, house values are estimated by the Statistics Finland. We believe that estimated house values are less accurate proxies of true market value than the self-reported values. The estimated house values are in many case smaller than the associated mortgage loans, which seems unlikely given recent housing market developments in Finland.

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

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. (The same interest rate is used when computing the implied cost of debt service that is deducted from disposable income.) 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 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.

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.

⁶See Määttä and Terviö (2014) for details.

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_h denote the stochastic non-housing endowment of a household owning a house of type h in the post-shock distribution. We assume that it is determined as $y_h = \bar{y}_h(1 + \delta_h)(\exp(\eta)/s)$, 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_h represents a systematic component of income dynamics. We further assume that δ_h can be described as a third order-polynomial in the percentile k , so that $\delta_h = a_0 + a_1h + a_2h^2 + a_3h^3$.

We normalize $x_1 = 1$ and set p_1 exogenously at its empirical value. We are left with the polynomial coefficients a , 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 distribution \bar{p} , ii) the average non-housing consumption for households with different house types is close to the empirical relation, iii) the share of households that engage in trade is approximately 33%, and iv) average income equals the average income in the data.

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 \bar{y} , and the share of households that engage in trade. We select the remaining five parameters (the standard deviation of the income shock σ_η and the polynomial coefficients a) 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%, and by requiring that average income equals the average income in the data. The latter constraint pins down one of the polynomial coefficients, given the other ones. By taking prices as given in this stage, we avoid the need to solve for equilibrium prices over and over again when varying these parameters. If we are able to closely replicate the empirical non-housing consumption curve, the associated equilibrium prices will also be close to the observed prices.

⁷The empirical estimates of this elasticity vary considerably. See for instance Li et al. (2015) and the references therein. However, as we show below, our main results are not very sensitive to the assumed elasticity.

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, instead of reflecting the 10 year model period. The top-left panel shows the empirical price distribution p . The top-right panel shows the calibrated δ . In the calibrated model, δ_h 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 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 calibration requires a different standard deviation 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 of transaction taxes

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-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 43% to 38%, or by about 12% in relative terms. Increasing the tax rate from, say, 2% to 4% decreases the trade volume by 21%. The relation between the transaction tax rate and the trade volume is almost the same at different elasticities. Reassuringly, it also seems realistic in light of previous empirical evidence. For example, Van Ommeren and Van Leuvensteijn (2005) found that a one percentage-point increase in the value of transaction




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.

costs in the Netherlands decreased residential mobility by at least 8 percent. Dachis et al. (2011) found that an average transaction tax of 1.1% led to a 15% decrease in the transaction volume of single-family houses in Toronto. Bérard and Trannoy (2017) analyzed an increase in the real estate transfer tax rate from 3.8% to 4.5% in certain French departments. Their estimated effect is a 4.6% decrease in the transaction volume. Hilber and Lyytikäinen (2017) estimated from UK data that increasing the transaction tax rate from 1% to 3% reduced the transaction volume by 27%, while Best and Kleven (2016) found that a lowering of the transaction tax from 1% to 0% over a certain price range increased transactions (in the relevant house price range and after the reversal of anticipation effects) by about 12%.

The top-left panel shows the annual transaction tax revenue (per owner household) 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 10% in all cases, which is not far from actual rates in some European countries. Our results suggests that in those countries lowering the tax rate might not decrease tax revenue at all.

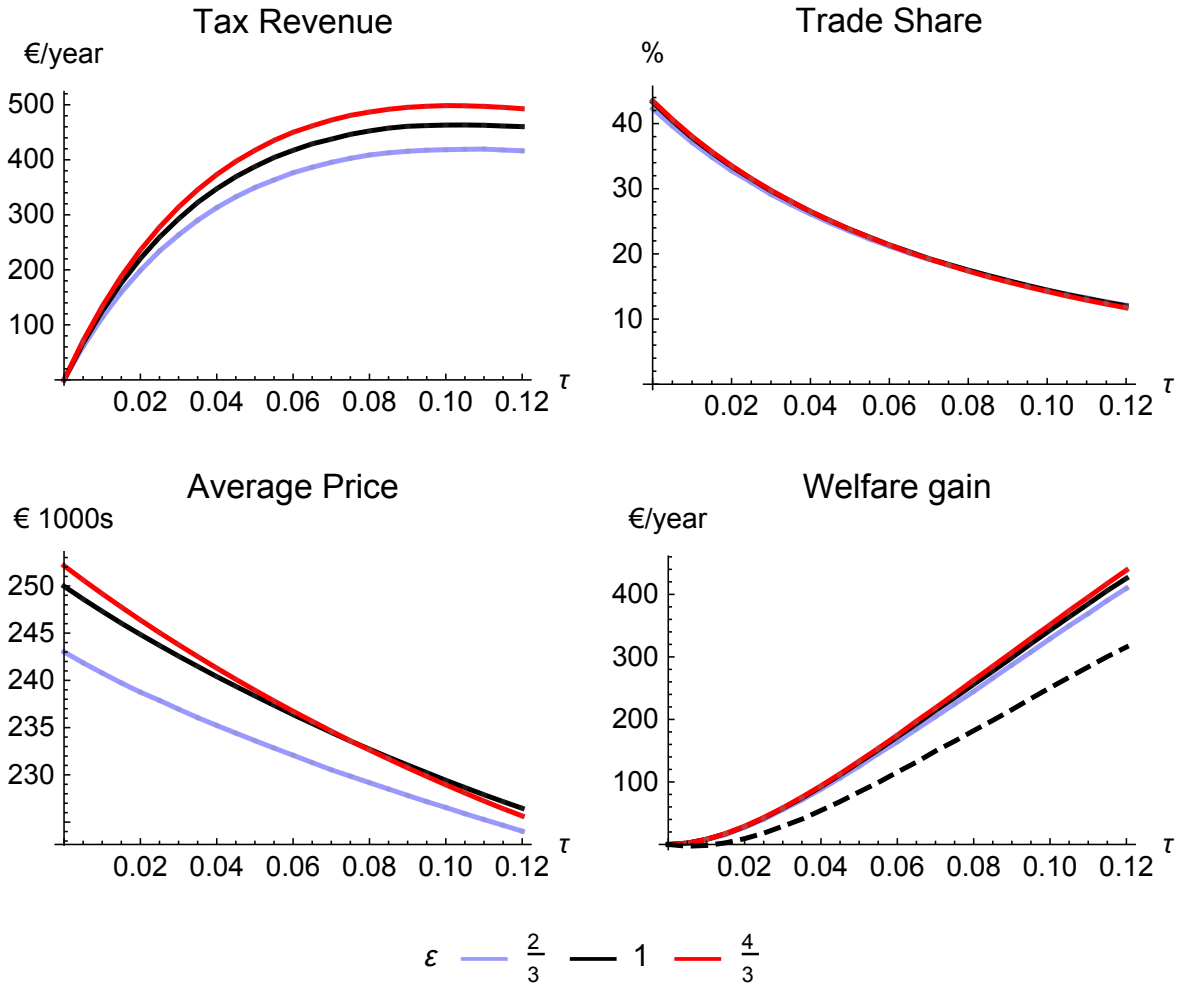


Figure 3: Aggregate effects of a transaction tax τ , at selected levels of the elasticity of substitution ε between housing and other goods. Welfare gain is the average compensating variation for avoiding the reform; dashed curve shows the expected welfare gain in the unit elastic case (see section 2.3).

The bottom-left panel shows the average pre-tax house prices as a function of the transaction tax rate. Naturally, a higher transaction tax implies a lower average house price. Finally, the bottom-right panel 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

non-distortionary. The ex-post welfare gain is measured as the average increase in non-housing consumption that would make households in the initial post-shock distribution indifferent between the equilibria associated with a given transaction tax or a revenue equivalent property tax. The welfare gain represented by the dashed line is based on expected or ex ante welfare before the income shocks are realized. For reasons explained in section 2.3, we consider ex ante welfare only in the unit elasticity case.

According to the ex post measure, replacing the current 2% transaction tax by a revenue equivalent property tax would increase household welfare by about 30 Euros in terms of non-housing consumption (in 2004 Euros). The welfare cost increases rapidly as we increase the tax rate. For instance, according to the ex post measure, replacing a 6% transaction tax rate with a property tax would generate an average annual welfare gain of around 180 euro, for elasticities considered. Since the property tax is essentially a lump-sum tax, these welfare gains can be interpreted as the overall welfare cost of the transaction tax. The ex ante aggregate welfare gain is always smaller than the ex post gain. This reflects the concavity of the utility function together with the fact (illustrated below) that in absolute terms the reform tends to benefit those with a positive income shock the most.

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 4 displays the approximated MCPF 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. The welfare cost is again measured using the ex post comparison. Since the welfare cost (or the private cost of public funds) includes the tax revenue, MCPF of a non-distortionary tax would equal to one by definition.

At a 2% tax rate, for instance, the MCPF is about 1.3. Hence, according to the model, the current transaction tax is not very distortionary. However, the MCPF increases rapidly with the tax rate. At a 7% tax rate the MCPF is already above 3 in all cases. 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.)

All the above results are relatively insensitive to the assumed elasticity of substitution. In order to understand this feature, 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




Figure 4: Marginal cost of public funds as a function of the transaction tax rate.

rationalize the observed price dispersion as the equilibrium outcome in our model. We vary some other parameters in section 6 below.

5 Distributional effects

Despite aggregate welfare gains from replacing the transaction tax with a property tax, some households may of course be worse off with such a reform. Figure 5 displays the ex ante welfare gain (top-left panel) and the share of households that are ex post better off (top-right) across the distribution of initial house endowments for initial transaction tax rates equal to 2% and 8%. It also displays the trade share (bottom-left) and the distributional house price impact of the reform. The top panels reveal that when the initial the transaction tax rate is 2%, many households are worse off both in the ex post and in the ex ante welfare comparison. However, the ex ante losses are all very small. When the initial transaction tax is 8%, all households are better off ex ante, whereas many households are still worse off ex post.

The largest ex ante gains accrue to households endowed with the very best houses. This is natural since the welfare gains are measured in absolute terms and those households have the highest average utility levels. The ex ante welfare gain and the share of ex post winners are also positively correlated with the trade share, which varies somewhat across the initial house endowments. There are also price effects. In particular, the increase in the value of the very best houses benefits those who initially own them, as those households are more likely to trade down than up.

Figure 6 shows the overall share of households that would be better off with the reform

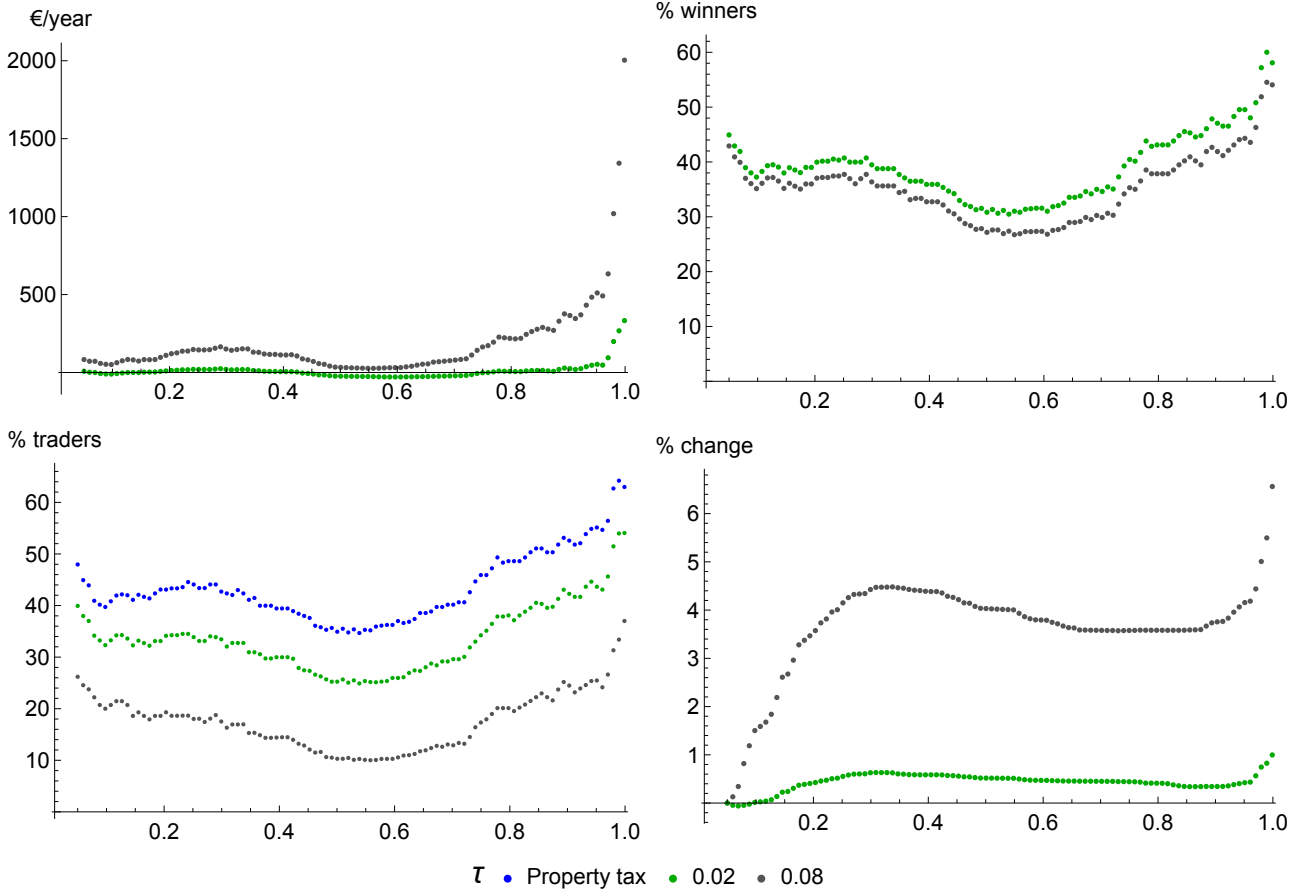


Figure 5: Distributions associated with replacing a transaction tax with a revenue equivalent property tax, for two transaction tax rates: ex ante welfare gain (top-left), share of ex post winners (top-right), and change in house price (bottom-right). The share of traders (bottom-left) is depicted at both transaction rates and also for the case of no transaction tax (which is the same at any property tax).

ex ante (left panel) and ex post for different initial transaction tax rates. Clearly, the two welfare comparisons provide very different results. In the ex ante comparison, the share of households that are better off with the property tax is increasing in the transaction tax rate (or, equivalently, the tax revenue requirement) and reaches 100% for transaction tax rates above 7%. In the ex post comparison, in contrast, the share of winners is decreasing with the tax rate.

The fact that the share of ex post winners is decreasing in the transaction tax rate is perhaps surprising, especially given that we already showed that high transaction taxes are very distortionary in the model. However, while the exact share of households that benefit or lose from the reform certainly depends on the details of the calibration, this feature is robust; Figure 7 helps understand why. There we depict the impact of reform




Figure 6: Proportion of households that benefit from a revenue neutral reform that replaces a transaction tax τ with a property tax? Left panel shows the proportion ex ante before households know the realization of their income shock; this is calculated only for the unit elastic case. The right panel shows the proportion of winners after the realization of the income shock for different elasticities.

on the welfare for owners of one house type, with the pre-trade endowment of income represented on horizontal axes. So this picture captures households on one vertical slice of Figure 1 (in fact those at the median house type). The top panel shows a case where the transaction tax is 2% and the bottom panel a case where it is 2%.

The shaded areas show no-trade regions, i.e., households that don't find it worthwhile to trade as their post-shock income is sufficiently "in balance" with their house type. The lighter (red) shade shows the no-trade region under the transaction tax, it is naturally wider in the high-tax case. The darker (blue) shade shows the no-trade region under a property tax; it is the same in both cases because property taxes don't distort trading decisions. The impact of tax reform on trading is thus naturally higher in the high-tax case, where the no-trade area contracts from a wider starting point.

The first thing to understand is that everyone who does not trade is worse off under a property tax, and this loss of utility is larger the higher the tax rate because the revenue-equivalent property tax must be higher in the high tax case. For those who are still inside the no-trade region under the property tax this reduction welfare is the only change caused by the policy. Those who are trading in any case clearly benefit from the change: the same aggregate tax burden is divided over a wider base of houses.

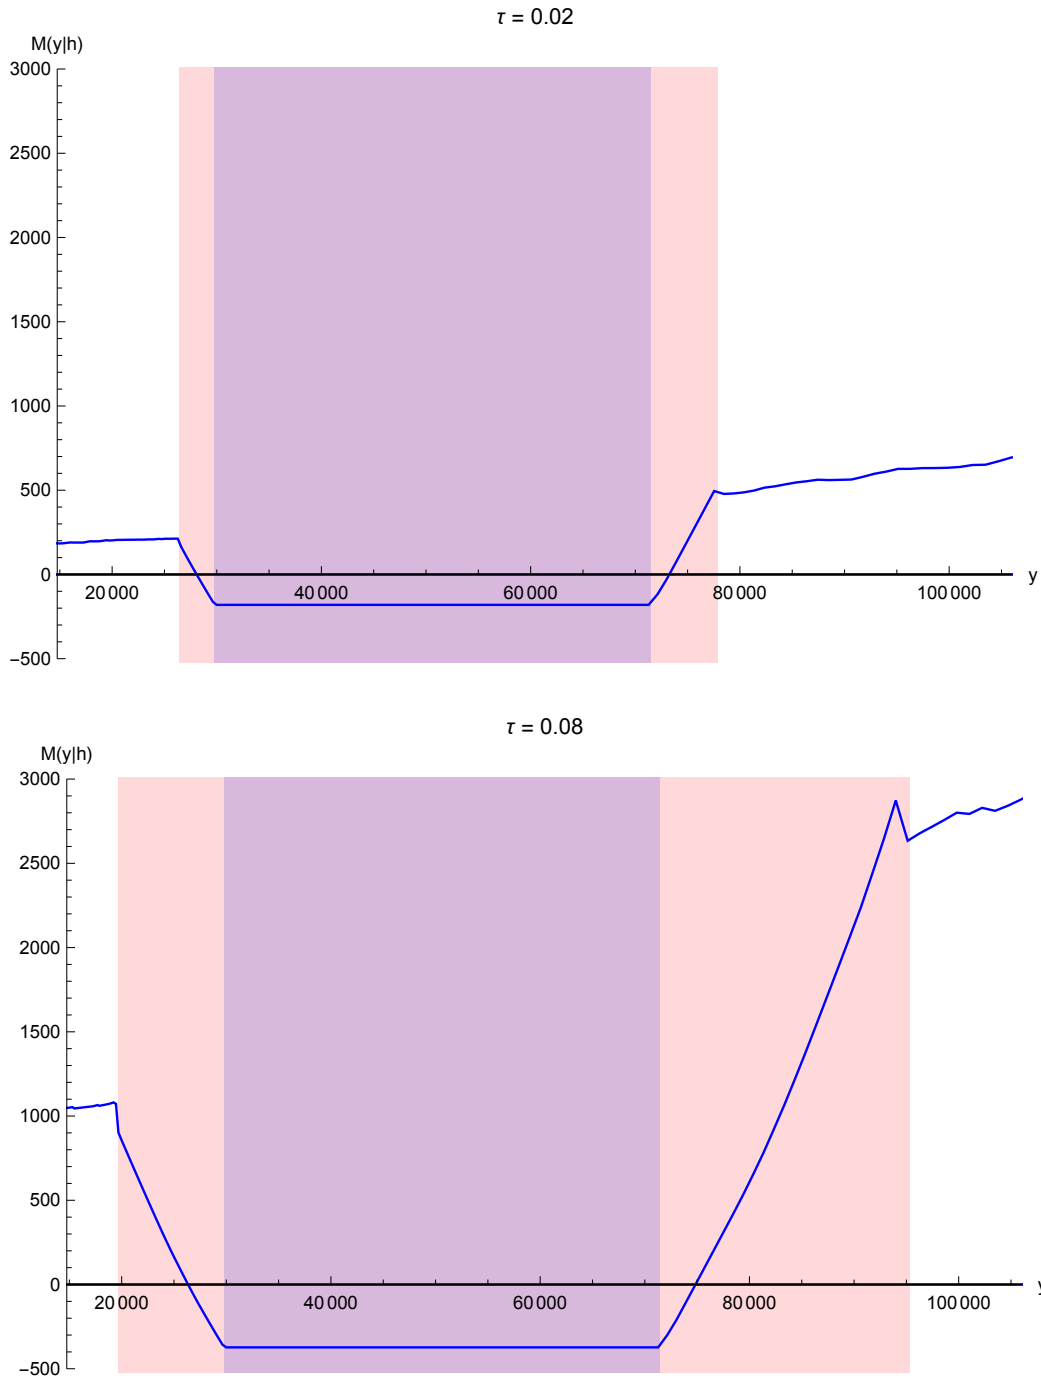


Figure 7: Difference in welfare relative to the equilibrium under a transaction tax of 2% (top) and 8 % (bottom) for households endowed with median house type. Household income on horizontal axes. Blue curves show the change in utility caused by a shift to a property tax.

Now consider those who are induced to trade by the tax reform. The new marginal traders, those at the boundary of the "new" no-trade region suffer the same reduction in

welfare as the marginal non-trader, so they are clearly made worse off by the reform, and more so in the high tax case. As we move further out from the no-trade boundary the gain from trade increases and, eventually, it is large enough to catch up with the burden of the property tax. But the deeper the starting point—the higher the tax burden—the further out into the trading region is the point of catch-up. This means that there are fewer households beyond the break-even points where the curve crosses the horizontal axes and thus winners from the tax reform.

The ex post welfare comparison could help explain why there appears to be political support for retaining even relatively high transaction tax rates that are likely to be close to the top of the Laffer curve, despite an obvious and efficient alternative for taxing the same tax base. In our view, the ex post comparison may well be more relevant when thinking about the political economy of transaction taxes than the ex ante comparison. At any given point in time, a majority of households would not benefit a lot from moving. Therefore, they have little incentive to support the elimination of the transaction tax at least in the short-run. While households should anticipate that they may want to move sometime in the future, they may want to at least postpone the elimination of the transaction tax until they themselves contemplate moving.

6 Sensitivity

We now consider briefly how changes in parameters other than the elasticity of substitution affect the results. We vary the non-fiscal transaction costs and the variance of the income shock. We also consider changes in the quality distribution. In the benchmark calibration, the non-fiscal transaction costs are 4% of the empirical house prices while the variance of the income shock is 0.47. Here we consider transaction costs equal to 2% and 8% of the empirical prices and income shock variances equal to 0.37 and 0.57. As for the quality distribution, we first decrease and then increase each quality increment $x_j - x_{j-1}$, $j > 1$, by 25%. Decreasing (increasing) these quality increments implies a less (more) dispersed quality distribution and a lower (higher) average quality. All these changes are made one at a time around the calibration with a unit elasticity.

Tables 1 and 2 display selected statistics for transaction tax rates of 2% and 8%. Table 1 displays the average (annual) tax revenue per household, the share of households that trade, and the average house price. Non-fiscal transaction costs (“Fixed cost”) have a large effect on the trade share, as could be expected. Therefore higher transactions costs also imply a lower transaction tax revenue, and vice versa. A higher income shock variance and increased quality dispersion also increase the trade share and the tax revenue

Table 1. Sensitivity Analysis: Tax revenue, trading volume, and prices.

		Tax revenue		Trade share		Avg. price	
		€/hh		%		€1000s	
$\varepsilon \setminus \tau$		0.02	0.08	0.02	0.08	0.02	0.08
Fixed cost	Low	284.3	558.1	42.9	21.4	249.6	236.5
	High	138.0	299.8	21.2	11.8	236.1	225.7
Shock variance	Low	156.3	249.8	24.4	10.1	246.2	234.0
	High	274.0	651.1	40.3	24.3	245.5	233.1
Quality increment	Low	160.7	330.7	29.3	15.1	200.1	191.4
	High	277.1	560.0	36.1	18.9	290.1	274.7
Calibrated		220.5	452.4	33.4	17.5	244.9	232.7

Table 2. Sensitivity Analysis: Welfare.

		Welfare gain		Winners		E Welfare gain		E Winners	
		€/hh		%		€/hh		%	
$\varepsilon \setminus \tau$		0.02	0.08	0.02	0.08	0.02	0.08	0.02	0.08
Fixed cost	Low	40.9	332.3	49.9	43.6	25.9	275.0	68	100
	High	15.0	143.7	25.0	23.3	-4.5	82.1	39	72
Shock variance	Low	28.7	219.3	31.4	29.6	16.1	181.4	71	100
	High	25.9	247.2	44.8	39.2	2.3	161.1	22	79
Quality increment	Low	19.9	181.2	34.6	31.4	5.3	131.2	41	85
	High	36.5	309.3	42.0	37.7	14.1	235.2	66	100
Calibrated		27.7	244.6	39.0	35.2	9.3	182.1	57	100

by increasing gains from trade. The average house price is relatively insensitive to the non-fiscal transaction costs (which are fixed in absolute terms) and the shock variance. In contrast, changes in the quality distribution obviously also affect the average house price via changes in the average house quality.

Table 2 displays the aggregate welfare gain associated with switching to a revenue equivalent property tax and the share of households who would be better with such a reform (“Winners”). The first two supercolumns (“Welfare gain” and “Winners”) refer to the ex post welfare comparison and the last two (“E Welfare gain” and “E Winners”) refer to the ex ante comparison.

One might reason that other transaction costs make transaction taxes more harmful. However, in the model higher non-fiscal transaction costs are associated with a lower aggregate welfare gain from the reform and also a smaller share of households that are better off with it. The reason why the share of ex post winners decreases with the transaction costs should be clear from the previous section: Higher non-fiscal transaction

costs increase the share of households that would not move even without the transaction tax. Those households cannot benefit from switching to a property tax. More generally, higher non-fiscal transaction costs reduce the extent to which the transaction tax alone can distort the allocation. This reduces both the ex ante and ex post welfare gains of the reform.

With a high non-fiscal transaction cost and a relatively low initial transaction tax rate, the ex-ante welfare gain can even be negative. Related to this, changes in the income shock variance have a different effect on ex post and ex ante welfare comparisons. A higher variance increases aggregate welfare gain and the share of winners from switching to a property tax in the ex post comparison, but decrease them in the ex ante comparison.

7 Conclusion

Our main question was what is the efficiency cost of an valorem transaction tax in the market for owner-occupied housing arising from the cross-sectional misallocation of houses. In order to answer the question we developed a one-sided assignment model, with a fixed distribution of different houses. In the model, households own the houses already before any changes in policy, and pre-existing policies are capitalized to house prices. Because of income/preference shocks, some households want to move; transaction costs affect the equilibrium allocation and prices by affecting these swapping costs.

We calibrated the model to data from the Helsinki metropolitan area housing market. We then use it to compare transaction taxes relative to revenue equivalent property taxes. We consider both ex ante and ex post welfare comparisons. Ex post comparison corresponds to a world where households know for sure whether they will trade and how much they would benefit from the opportunity to trade. Ex ante measurement corresponds to a world where they don't know whether or not they will be among the traders.

We found that the distortionary effects of the transaction tax increase rapidly with the tax rate. For instance, based on our ex post welfare comparison, for the 2% tax rate currently in place in Finland, the associated marginal cost of public funds (MCPF) is about 1.3 (whereas 1 would indicate a lump-sum tax). Hence, relatively low transaction taxes do not seem to create very large distortions. However, the MCPF increases to over 3 as we increase the tax rate to 7%, 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%.

Our efficiency results are not very sensitive to the elasticity of substitution between housing and non-housing consumption, a parameter which is very hard to estimate accu-

rately. This reflects the fact we use distributional data on house prices and income when calibration the housing quality distribution. 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.

Despite aggregate welfare gains from replacing the transaction tax with a property tax, we found that many households may be worse off with a such a reform. Moreover, in the ex post comparison, the share of households that are worse off is increasing in the initial transaction tax rate up to tax rates close to the peak of the Laffer curve. This result may explain why there appears to be political support even for high transaction tax rates that are very distortionary.

In order to 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

- Bérard, G., Trannoy, A., 2017. The impact of a rise in the real estate transfer taxes on the French housing market.
- Best, M., Kleven, H., 2016. Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK. *The Review of Economic Studies*, forthcoming.
- Dachis, B., Duranton, G., Turner, M. A., 2011. The effects of land transfer taxes on real estate markets: evidence from a natural experiment in Toronto. *Journal of Economic Geography* 12 (2), 327–354.

- European Commission, 2012. Quarterly report on the euro area. Vol. 11.
- Hilber, C. A., Lyytikäinen, T., 2017. Transfer taxes and household mobility: distortion on the housing or labor market? *Journal of Urban Economics* 101, 57–73.
- Kopczuk, W., Munroe, D., May 2015. Mansion tax: The effect of transfer taxes on the residential real estate market. *American Economic Journal: Economic Policy* 7 (2), 214–57.
- Lundborg, P., Skedinger, P., 1999. Transaction taxes in a search model of the housing market. *Journal of Urban Economics* 45 (2), 385–399.
URL <http://EconPapers.repec.org/RePEc:eee:juecon:v:45:y:1999:i:2:p:385-399>
- Määttänen, N., Terviö, M., 2014. Income distribution and housing prices: An assignment model approach. *Journal of Economic Theory* 151, 381–410.
- Mirrlees, J., Adam, S., Besley, T., Bell, R., Bond, S., Chote, R., Gammie, M., Johnson, P., Myles, G., Poterba, J. M., Sep 2011. *Tax by design*. Oxford University Press.
URL <http://www.ifs.org.uk/docs/taxbydesign.pdf>
- O’Sullivan, A., Sexton, T. A., Sheffrin, S. M., 1995. Property taxes, mobility, and home ownership. *Journal of Urban Economics* 37 (1), 107–129.
- Stokey, N. L., 2009. Moving costs, nondurable consumption and portfolio choice. *Journal of Economic Theory* 144 (6), 2419–2439.
- Van Ommeren, J., Van Leuvensteijn, M., 2005. New evidence of the effect of transaction costs on residential mobility. *Journal of Regional Science* 45 (4), 681–702.

Appendix

A 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 B, 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

$$u(x_j, \bar{y}_j) = u(x_{j+1}, \bar{y}_j - (p_{j+1} - p_j)(1 + \tau_T + \phi)). \quad (9)$$

Under CES-utility this can be solved for

$$x_{j+1}^\rho - x_j^\rho = \bar{y}_j^\rho - (\bar{y}_j - (p_{j+1} - p_j)(1 + \tau_T + \phi))^\rho. \quad (10)$$

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 ...) With a sufficiently fine grid $\bar{y}_j \approx \underline{y}_j$, and we can treat both as approximations for y_j .⁸ The CES inference formula under transaction costs is

$$\hat{x}_h = \left(x_1^\rho + \sum_{j=2}^h ((y_{j-1} + (p_j - p_{j-1})(1 + \tau_T + \phi))^\rho - y_{j-1}^\rho) \right)^{\frac{1}{\rho}} \quad (11)$$

⁸Caveat: 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 are close.

where x_1 is an arbitrary positive constant. Crucially, x_1 washes out of all differences of the type $\hat{x}_h^\rho - \hat{x}_j^\rho$. Therefore the inferred quality differences between house types only depend on the data (prices p_i and disposable incomes y_i for i in j, \dots, h) and the elasticity of substitution between housing and other consumption.

With $\rho \rightarrow 0$, i.e., with log-utility, the inference formula becomes

$$\hat{x}_h = x_1 \prod_{j=1}^h \frac{\bar{y}_j}{\bar{y}_j - (p_{j+1} - p_j)(1 + \tau_T + \phi)} \quad (12)$$

In this case x_1 cancels out of inferred quality ratios \hat{x}_h/\hat{x}_j , which therefore only depend on observed prices and incomes.

B 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 is 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

$$\bar{y}_k(p) = \{y \text{ s.t. } u(x_k, y - \tau p_k - \xi_k) = u(x_{k+1}, y + p_k - (1 + \tau)p_{k+1} - \xi_{k+1})\}, \quad (13)$$

$$\underline{y}_k(p) = \{y \text{ s.t. } u(x_k, y - \tau p_k - \xi_k) = u(x_{k-1}, y + p_k - (1 + \tau)p_{k-1} - \xi_{k-1})\}. \quad (14)$$

(Recall that, in equilibrium, $\bar{y}_N = \bar{\theta}$ and $\underline{y}_0 = \underline{\theta}$.) In a world with transaction costs, these will be the post-trade allocations of all 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

$$x = x_j, y \in [\underline{y}_j(p), \bar{y}_j(p)]. \quad (15)$$

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

$$\tilde{y}_{k,s} = \bar{y}_{k+s}(p) + (1 + \tau)p_{k+s} + \xi_{k+s} - p_k. \quad (16)$$

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

$$u(x_{k+s}, \bar{y}_{k+s}(p)) \geq u(x_k, \bar{y}_{k+s}(p) + (1 + \tau)p_{k+s} + \xi_{k+s} - p_k) \quad (17)$$

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

$$u(x_k, y) = u(x_{k+s}, y - (1 + \tau)p_{k+s} - \xi_{k+s} + p_k). \quad (18)$$

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 $\bar{Y}_k(p) = \bar{\theta}$.

⁹If $\tau=0$ and $s = 1$, then (17) 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.

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

$$\begin{aligned}\theta &\leq \bar{y}_k(p) + p_k - p_j, \\ \theta &\geq \max\{\underline{y}_k(p) + p_k - p_j, \bar{Y}_j(p)\}.\end{aligned}\tag{19}$$

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

$$\begin{aligned}\theta &\leq \min\{\bar{y}_k(p) + p_k - p_j, \underline{Y}_j(p)\}, \\ \theta &\geq \underline{y}_k(p) + p_k - p_j.\end{aligned}\tag{20}$$

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

$$\underline{Y}_k(p) < \theta \leq \bar{Y}_k(p).\tag{21}$$

Total demand for type- k houses is

$$\begin{aligned}Q_k(p) &= \sum_{j=0}^{k-1} \max\left\{0, F_j(\bar{y}_k(p) + p_k - p_j) - F_j\left(\max\{\underline{y}_k(p) + p_k - p_j, \bar{Y}_j(p)\}\right)\right\} \\ &\quad + F_k(\bar{Y}_k(p)) - F_k(\underline{Y}_k(p)) \\ &\quad + \sum_{j=k+1}^N \max\left\{0, F_j(\min\{\bar{y}_k(p) + p_k - p_j, \underline{Y}_j(p)\}) - F_j(\underline{y}_k(p) + p_k - p_j)\right\}.\end{aligned}\tag{22}$$

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.