

Welfare Effects of Housing Transaction Taxes: A Quantitative Analysis with an Assignment Model

Niku Määttänen

University of Helsinki and Helsinki GSE*

Marko Terviö

Aalto University and Helsinki GSE[†]

August 18, 2021[‡]

Abstract

We evaluate the welfare cost of housing transaction taxes with a new assignment model based framework, where welfare effects are driven by distortions in the matching of houses and households. We calibrate the model with data from the Helsinki metropolitan region to assess the impact of a reform where an ad valorem transaction tax is replaced with a revenue equivalent property tax. The aggregate welfare gain from this reform increases rapidly with the initial transaction tax rate, with the Laffer curve peaking at about 10%. The proportion of households that lose out from the reform is nevertheless increasing in the tax rate. We compare our model-based counterfactual aggregate welfare results with welfare calculations based on reduced-form estimates from previous policy evaluation studies; they are broadly in line, despite the latter using data from different housing markets at various levels and changes of the tax rate.

JEL: D31, H20, R21.

Keywords: transaction tax, housing market, assignment model.

*niku.maattanen@helsinki.fi, University of Helsinki and Helsinki Graduate School of Economics.

[†]marko.tervio@aalto.fi, Aalto University School of Business and Helsinki Graduate School of Economics.

[‡]We thank Essi Eerola, Miklos Koren, Alain Trannoy, and seminar participants at ANU, Bocconi, CEU, EIEF, Heidelberg, Leicester, LSE, UBC, Melbourne, UNSW, Uppsala, UPV/EHU, and UEA-2018 in NYC 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 view 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 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.”¹ Transaction taxes on housing stand out also because housing could be taxed in a relatively efficient manner using property taxation, yet they are common and as high as 10% of the transaction price in some countries.

In this paper, our aim is to estimate the welfare cost of transaction taxes at various tax rates and to quantitatively characterise the associated Laffer curve. We also consider the distributional effects of tax reforms that would replace the transaction tax with a revenue neutral property tax. Our approach is to construct model-based counterfactuals for a wide range of tax rates. This is in the same vein as, e.g., Trabandt and Uhlig (2011) who characterise Laffer curves for labour income, capital income and consumption taxes using the neoclassical growth model, or Holter et al. (2019) who use an overlapping generations model with idiosyncratic income risk to analyse how tax progressivity affects the labour income tax revenue. Our structural model accounts for how transaction taxes affect the equilibrium allocation of heterogeneous houses across heterogeneous households.

Our model depicts an urban area with a distribution of house types of different qualities and a population of households with different housing demands. The model builds on the one-sided assignment model in Määttänen and Terviö (2014), which we augment with transaction costs. (In a one-sided assignment model the same households are both buyers and sellers.) 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”. The set of houses is exogenous. Not living in any house is not an option, but staying in the current house is. 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

¹Chapter 16 in Mirrlees et al. (2011).

good, so the most important reason why some households choose to live in more expensive houses is that they can better afford them.

In order to focus on the misallocation of house types across households as the source of welfare effects, we abstract in our baseline model away from features such as credit constraints, search frictions, and life-cycle behavior that many other housing market models include.² This allows us to keep the model tractable despite featuring an unrestricted type distribution of indivisible houses which we can map to the data. Earlier structural analysis of housing market transaction taxes has not accounted for inherent heterogeneity in house types: Lundborg and Skedinger (1999) employed a search-and-matching model where houses are observationally identical whereas O’Sullivan et al. (1995) and Stokey (2009) focused on distortions to life cycle consumption behaviour as the source of welfare effects.

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 region. Given preferences, we specify the distributions of housing quality and income shocks so that the resulting equilibrium distributions and transaction volume match the data closely at the current level of transaction tax and other transaction costs. We experiment with different transaction tax rates to study the impact of changing tax rates.

There are some limitations to what welfare questions can be answered in this setup. The reason is that we can only reasonably estimate differences in house qualities but not their levels; 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.

Our baseline estimate of the aggregate welfare gain from the tax reform is about 13% of the tax revenue at the current 2% tax rate. The marginal cost of public funds (MCPF) for the transaction tax, i.e., the ratio of marginal welfare cost to marginal

²See Piazzesi and Schneider (2016) for a review of housing market models that incorporate some of these features.

tax revenue, is about 1.3 at the current 2% rate. Hence, according to our analysis, at the current rate the 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%. Some 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.³

Several empirical studies have exploited time-variation or discontinuities in transaction tax schedules to estimate how a given change in transaction taxation affects the transaction volume. Most of these papers find lasting effects of transaction taxes on transaction volume or transaction tax revenue. These studies include Buettner (2017), Dachis et al. (2011), Fritzsche and Vandrei (2019), and Eerola et al. (2021) who use time-variation in transaction taxes, Hilber and Lyytikäinen (2017) who exploit discontinuities in the UK transaction tax schedule, and Best and Kleven (2018) who exploit both discontinuities and time-variation in the UK tax schedule. However, Bérard and Trannoy (2018) who use regional variation and changes in transaction taxes in France, Besley et al. (2014) who study the UK stamp duty holiday in 2008 and 2009, and Slemrod et al. (2017) who exploit changes in notched transaction tax rates in Washington D.C., find only weak or no long-run effect.

These empirical estimates can be used to evaluate the welfare effects of (small) tax changes using the Harberger triangle approach. In studies with statistically significant findings this implied marginal cost of public funds varies between 1.06 and 1.84. By and large the incremental welfare costs appear higher at higher levels of the tax rate. We consider the same tax rate changes in our setup and show that our welfare results are roughly in line with the estimates that are based on the empirical results and the Harberger approach. We also show that the Harberger approach, which requires assuming linear transaction demand and constant house prices, provides a reasonably accurate approximation of our model-based welfare estimate.

In order to obtain quantitative results we need to specify the elasticity of substitution between housing and non-housing consumption, which is hard to pin down. Fortunately our results are not very sensitive to the assumed elasticity. This is no coincidence: if house quality matters less in the utility function (higher elasticity of substitution) then correspondingly the quality differences between houses must be larger to rationalize the observed price distribution as the equilibrium outcome in our model.

By contrast, the level of non-fiscal transaction costs makes a clear difference to our welfare results. Higher non-fiscal costs reduce the extent to which the transaction tax alone can distort the allocation, which in turn reduces the welfare cost of transaction

³See European Commission (2012) for tax rates on property transactions in euro area countries.

taxation relative to property taxation at any given tax rate. For instance, increasing the non-fiscal transaction costs from 4% of house value to 6% decreases our estimated welfare gain from replacing the current 2% transaction tax with a property tax by about 25%. Fortunately it is relatively simple to pin down a reasonable range for non-tax transaction costs in the Finnish housing market (our baseline assumption is 4%).

Property taxation is non-distortionary in our setup so the aggregate welfare effect of using it to replace transaction taxation is, unsurprisingly, positive and increasing in the tax rate. Nevertheless, we find that not only is a large share of households worse off “ex post” under a property tax, but that this share of losers is increasing in the tax rate. In other words, the higher the initial transaction tax rate the larger the fraction of households who would lose out from a reform that replaces it with a higher property tax. 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 a marginal welfare gain but households benefiting from these gains have to contribute a full share of the property tax. At low tax rates most households are worse off even “ex ante” before knowing their gains from trading.

Our study focuses on transaction costs—in the form of a transaction tax and non-tax transaction costs—as the friction distorting the matching of different houses and households. By abstracting from other trading frictions we may exaggerate the importance of the transaction tax in distorting the allocation of housing. Therefore, in an extension to our model, we add another friction by assuming that there is randomness in which subsets of households (buyers and sellers) are able to interact in the market. This means that households may not be able to buy the house or find the buyer that they would in the fully integrated market. This way of modeling matching frictions allows the assignment model framework still to be used in analyzing the data and our policy counterfactuals. By varying the number of houses that are drawn into each submarket, we can vary the degree of the matching friction or “market thickness”. With a smaller number of houses the realised house quality distributions become increasingly coarse, which shows up as an increasing dispersion of realised prices for houses of the same type. This matching friction indeed does reduce the welfare cost of the transaction tax. However, to make a significant difference to the welfare results, the market thickness has to be set so low that the resulting price variation between exactly similar houses becomes unrealistically large.

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 $\mathbf{p} = (p_1, \dots, p_n)$ 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 analysed 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 house quality x while taking house prices 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 types k , the fraction of households choosing to live in a house of quality 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, 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 end up in the better house, or else the two could engage in a mutually profitable trade. (In the absence of transaction costs this would cover 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 A, *ibid.*)

The equilibrium allocation is illustrated in Figure 1, with house quality on horizontal and income and non-housing consumption on vertical axes. All households that trade, trade to a region in consumption space between the two innermost “curves”. For each house type k , there is a black vertical line between these curves that depicts the range of non-housing consumption levels for households that bought a house of that type, $[\underline{y}_k, \bar{y}_k]$. These bounds are increasing in the sense that equilibrium wealth and therefore the level of utility (mapped with the gray indifference curves) is higher at a higher quality house. This follows directly from Lemma 1: wealth and utility must be increasing in house quality; with continuous income distributions this holds as an equality for those at the margin. The wealthiest household choosing to buy a type k house has the same wealth as the poorest household choosing to buy a type $k + 1$ house: $p_k x_k + \bar{y}_k \leq p_{k+1} x_{k+1} + \underline{y}_{k+1}$.

In the absence of transaction costs everyone except those “born” inside the “trade-to region” would trade. The existence of a thick trade-to region would then only be due to the discreteness house types (and with a continuum of house types it would be a curve without thickness, $\underline{y}_k = \bar{y}_k$). However, due to transaction costs, the “no-trade region” is wider than the trade-to region; it is depicted in Figure 1. The vertical lines between the outermost curves depict the range of incomes for households that are endowed with a house of type k and choose not to trade, $[\underline{Y}_k, \bar{Y}_k]$. Households in the no-trade region do not trade because it is not worth paying the transaction cost for what would be a relatively short move in consumption space.

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




Figure 1: Consumption space, with non-housing consumption on vertical and quantile of house quality on horizontal axes. The no-trade region is the shaded area between the two outermost curves, and the post-trade consumption bundles of traders are between the two innermost curves. The decreasing curve is an example of a budget curve and the slightly separated dot above it an associated endowment. The resulting post-trade bundle is the highlighted point in the intersection of the budget curve and the post-trade curve. This example was solved for a joint log-normal distribution with $\sigma_x = \sigma_y = \text{Corr}(x, y) = 0.5$, log-utility, $n = 100$ and transaction costs equal to 4% of the house value.

composite good. Figure 1 also depicts a budget curve for an example household. The endowment 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 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. Differences in household preferences 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. However, in this paper we will use the income-shock interpretation for ϵ_h .

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: all prices going up by a million 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 transaction costs this is just the difference between the two house prices. Transaction 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 in notation, 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 buyer, as is the case in Finland.

Efficiency of property taxation In the assignment model with an exogenous housing supply a property tax acts just like a lump sum tax on wealth. To see this, consider first the equilibrium conditions in the absence of transaction costs. The marginal household endowed with a house of quality x and monetary wealth y is indifferent between staying and moving to house x' and paying the price difference $p' - p$. In equilibrium the price difference must satisfy $u(x, y) = u(x', y - (p' - p))$.

Now consider a proportional property tax τ_p . If a household stays put it pays the property tax $\tau_p p$; if it moves it pays $\tau_p p'$. The indifference condition becomes

$$u(x, y - \tau_p p) = u(x', y - (p' - p) - \tau_p p'). \quad (3)$$

Denote the disposable wealth net of the property tax burden as $\tilde{y} = y - \tau_p p$ and the gross-of-property-tax house price as $\tilde{p} = (1 + \tau_p)p$. Now the marginal traders' indifference condition (3) can be written as

$$u(x, \tilde{y}) = u(x', \tilde{y} - (\tilde{p}' - \tilde{p})). \quad (4)$$

This is just the original indifference condition, but with different levels of non-housing wealth and a relabeled price variable. The only real difference is lower wealth levels; equilibrium prices and allocation are determined as before. Hence, just like for a lump sum tax, the equilibrium allocation can change due to income effects but these are not distortions in the ordinary sense.

Transaction costs distort the equilibrium in the sense that some trades that would've increased household welfare fail to take place. A property tax will in effect change the non-housing wealth distribution, but the distortions caused by transaction costs remain a separate issue. There is no reason why a change in the wealth distribution would tend to interact positively or negatively with transaction costs to make their distortions better or worse.

CES utility For the quantitative exercise we assume CES utility,

$$u(x, y) = \left(\alpha x^{\frac{\varepsilon-1}{\varepsilon}} + (1 - \alpha) y^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}, \quad \text{where } \alpha \in (0, 1), \quad (5)$$

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 normalised away. We derive the formula for inferring x from 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 numerical solution procedure in detail in Appendix B.

We don't have a proof for the existence of the equilibrium with transaction costs. However, our numerical solution method works reliably as we find the same equilibrium for given parameter values starting from a large set of initial price vectors.⁴

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 households.

A natural baseline economy for our quantitative analysis will be the actual economy with its 2% transaction tax and no property tax, and the comparison with the counterfactual of a revenue equivalent 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.

The timing of the welfare measurement turns out to be an important consideration here. In the quantitative exercise (which we describe in the next section) we consider a world where *ex ante* households are located on the equilibrium “trade-to” region but then simultaneously receive income shocks, which causes them to spread out in consumption space, which motivates a round of trading. We generally measure household welfare after households already know the realization of their shock, i.e., *ex post*, whereas measuring welfare *ex ante* before the income shock is realised is possible only in special cases.

Now 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 consume some $\{x(y|h), c(y|h)\}$ in equilibrium. 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 a transaction tax

⁴A Matlab code package to solve the model is available as a supplementary file.

regime to a property tax regime is

$$M(y|h) = m \text{ s.t. } \left\{ u(x(y|h), c(y|h) + m) = u(x_*(y|h), c_*(y|h)) \right\}. \quad (6)$$

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

$$W = \frac{1}{n} \sum_{h=1}^n \int_y M(y|h) f_h(y) dy, \quad (7)$$

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 absolute “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 (7), evaluated after the income shocks have been realised, 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 why this is so. Consider households endowed with house type x_h and income y_h . After the income shocks are realised—but before the possible trading—their income is distributed according to the density f_h . Conditional on the realization y , they 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 (6) becomes

$$M(y|h) = \left(x_*(y|h)^\rho - x(y|h)^\rho + c_*(y|h)^\rho \right)^{\frac{1}{\rho}} - c(y|h). \quad (8)$$

To see how this depends on the estimated quality levels, use the inference formula for x_h (14) in Appendix A. The quality estimate for each house type $h = 2, \dots, n$ depends on data $y_1, \dots, y_h, p_1, \dots, p_h$, the assumed elasticity of substitution via ρ , and the quality level x_1 .⁵ The constant x_1^ρ cancels out of all quality differences $x_i^\rho - x_j^\rho$ (see Appendix A). Therefore the realised welfare gain (8) is independent of x_1 ; it only depends on the data $y_i, \dots, y_j, p_i, \dots, p_j$, and the elasticity.

⁵Recall that any $\alpha \in (0, 1)$ can be absorbed into the units of x ; here we set $\alpha/(1 - \alpha) = 1$ for notational convenience.

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

$$\bar{M}(h) = m \text{ s.t. } \left\{ \int (x(y|h)^\rho + c(y|h)^\rho + m)^{\frac{1}{\rho}} f_h(y) dy = \int (x_*(y|h)^\rho + c_*(y|h)^\rho)^{\frac{1}{\rho}} f_h(y) dy \right\}. \quad (9)$$

In general, the solution of M depends on the 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 (15) in the Appendix. The expected welfare gain in (9) becomes $\bar{M}(h) = m$ s.t.

$$\int (\alpha \log x(y|h) + (1-\alpha) \log c(y|h) + m) f_h(y) dy = \int (\alpha \log x_*(y|h) + (1-\alpha) \log c_*(y|h)) f_h(y) dy \quad (10)$$

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” (6) 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.

Throughout this section, we assume that the number of house types is very large so that the trade-to region (see Figure 1) can be thought of as a curve. We first estimate the trade-to curve from the data by estimating the relation between average non-housing consumption (or a proxy measure for it) and house prices. For sure, in the presence of transaction costs, the relation between average non-housing consumption and house prices does not exactly correspond to the trade-to curve, because some households are off the trade-to curve in the no-trade region. However, since the trade-to curve is strictly contained in the no-trade region, the relation between the average non-housing consumption and house prices should give us a good approximation of the true trade-to curve.

We then use the estimated trade-to curve to infer house qualities. 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 model trade, we consider an expanded model period which has the following three stages. In stage 1, households are all located on the estimated trade-to curve. 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 non-degenerate 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 what follows, we refer to this distribution as the “post-shock” distribution. 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 again content with their bundle of house type and non-housing consumption. Those who trade, are again on the trade-to curve. However, this curve is in general different than the estimated curve in stage 1.

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 trade-to curve at the end of stage 3 is close to the estimated curve in stage 1, while requiring that the share of households that choose to trade, given realistic transaction costs, matches the observed level of trading in the data.

In other words, we assume that the estimated trade-to curve reflects a stationary equilibrium of the model that would be repeated if the income shocks were drawn

from a time-invariant distribution. However, since we assume, for simplicity, that all households are on the estimated post-trade curve in stage 1 (instead of somewhere in the no-trade region), the distribution in the end of stage 3 is necessarily different from the stage 1 distribution (or curve), 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.

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 (MA). The 2004 survey was the last one to include self-reported house value and the length of stay in the current dwelling. In later surveys house values are estimated by the Statistics Finland. We believe that the self-reported house values are generally more accurate proxies of true market value than the estimated values. The estimated house values in later surveys are in many cases smaller than the associated mortgage loans, which seems unlikely given recent housing market developments in Finland.

The data include register-based values for household savings, debts and income as well as self-reported estimate of the market value of a household's main residence, which we take as our house price measure. 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. 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.⁶

Before estimating the relation between house prices and income, we need to make the units of annual 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 wealth survey. 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,

⁶The same problem afflicts the equivalent U.S. data (AHS). However, here, unlike in the AHS, house values are not top-coded. For more detail on the Helsinki data see the description in Määttä and Terviö (2009).

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.) For simplicity, we set the property tax rate at zero and the transaction tax rate at exactly 2%.⁷

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. 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}_h standing in for the h :th percentile. We reduce the relation of income and house value to a curve by using a kernel regression to estimate \hat{y}_h as $E[y \mid F_{\bar{p}} = (h - 1/2)/100]$, where $F_{\bar{p}}$ is the empirical CDF of house values.⁸ The resulting vector \hat{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 in our one-city model. While our data does not allow us to identify households that have moved from other regions, we know from separate population statistics that these movers represent about 30% of the overall population. We therefore target a share of households that engage in trade equal to $63\% - 30\% = 33\%$ at the actual transaction tax rate of 2%.

Here we are implicitly assuming that interregional movers do not affect the equilibrium allocation in the model city. This holds exactly if we assume that for each household that remains in the city, there are (a mass of) 0.3 households who move out of the city and sell their houses to movers from other regions who then end up on the same trade-to curve. The impact in terms of Figure 1 would then be the same as replacing a fraction of households on the no-trade curve by incoming households but without changing the shape of the curve. (The simplest further assumption would be that the shock of potential welfare gain on offer outside the model economy is orthogonal to the explicitly modeled income shock; in the end only the shock reflecting the choice made is realised.) As our model abstracts away from interregional moves it mechanically scales down the level of aggregate transaction tax revenue. We will

⁷There is a municipal property tax in Finland, but effective tax rates for dwellings are very low, partly because 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%. We also abstract away from a higher transaction tax rate for detached single-family homes and an exemption for young first-time buyers.

⁸See Määttä and Terviö (2014) for details.

return to this issue when interpreting our quantitative results.

We also need to specify other (non-fiscal) transaction costs. In Finland, the legal and administrative costs of buying and selling a house are relatively low. This may explain why typical broker fees are quite low as well, around 2% – 3% of house value in Helsinki and why it is common for households to sell their house without an agent (buyer agents are unheard of). We set the vector of house-type specific transaction costs ξ so that it corresponds to 4% of empirical house values. We conduct a sensitivity analysis with respect to non-fiscal transaction costs and the interest rate. (We leave non-tax transaction costs fixed when varying taxes; 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 transform taxes and transaction costs to reflect the model period in the same as 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 (5)), and consider elasticity values ε at 2/3, 1, and 4/3.⁹ 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 $y_h = \hat{y}_h(1 + \delta_h)(e^\eta/s)$, 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 e^η/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 h th percentile, 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 relation of average non-housing consumption and house price is close to the empirical relation \hat{y} , iii) the share of households that engage in trade is 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

⁹The 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.

distribution, determined by δ_h and σ_η , we can aggregate to find the post-trade relation between average non-housing consumption and housing, which we denote by $\tilde{\mathbf{y}}$, 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 $\hat{\mathbf{y}}$ and $\tilde{\mathbf{y}}$, 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.

3.4 Evaluation of fit

Figure 2 illustrates the data and the calibrations with different elasticities of substitution between housing and non-housing consumption. 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. 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 of Figure 2 shows the empirical price distribution \mathbf{p} . The top-right panel shows the calibrated mean reversion of income. 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 could not be stable 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 quality houses in the post-shock distribution tend to move upwards in the quality ladder, and vice versa.

The bottom-left panel of Figure 2 compares the equilibrium price distribution in the model (prices at the end of stage 3) 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 matches closely the empirical price distribution, except for the most valuable houses where the highest absolute deviations are about 5%. The bottom-right panel in turn shows the estimated relation of disposable income and house quality, $\hat{\mathbf{y}}$, and the relation of average post-trade consumption and house quality in the model, $\tilde{\mathbf{y}}$. Again, the calibrated model replicates the empirical

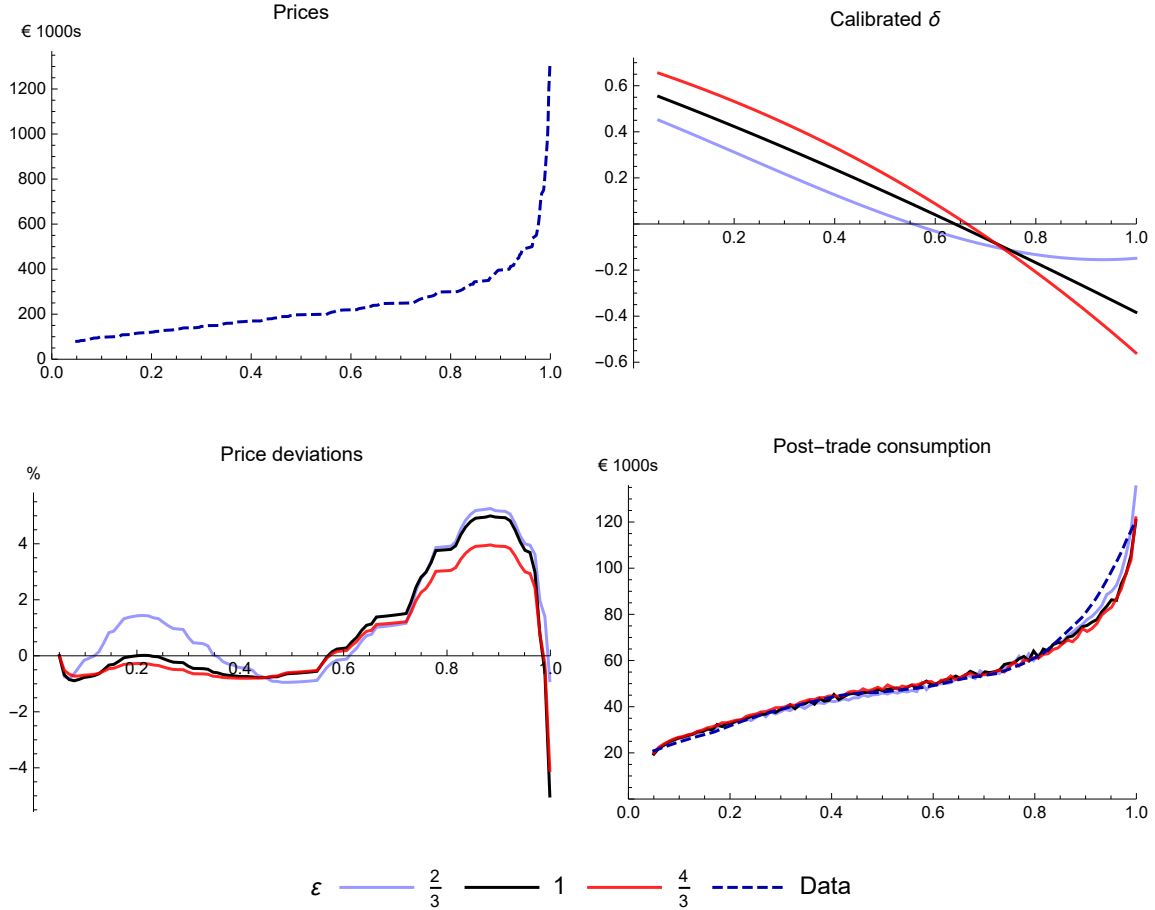


Figure 2: Empirical price distribution (top-left), calibrated δ_h (top-right), prices in the data vs. model (bottom-left), and average post-trade consumption implied vs. empirical relation of disposable income and house quality.

relation quite closely, especially below the 90th percentile or so. The mismatch near the top end could be a sign of non-homothetic preferences. It could also be related to the fact that our data is not top coded and may thus include erroneous outliers. We explore the sensitivity of our results to dropping the top end of the house price distribution in section 7.

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 13%. 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.

The top-left panel shows the annual transaction tax revenue 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 between a tax rate of 10% and 11% in all cases, which is not far from actual rates in some European countries. Our results suggest that in those countries lowering the tax rate might not decrease tax revenue at all.

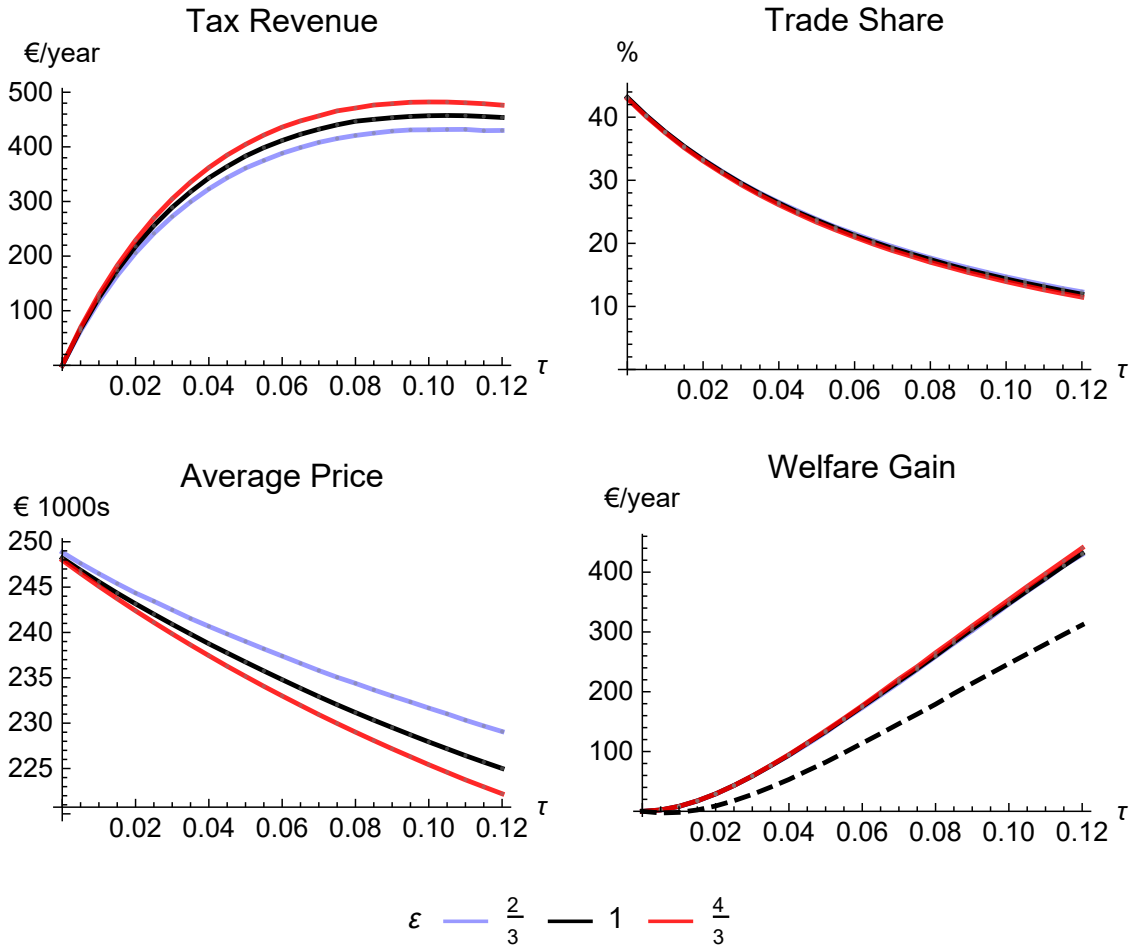


Figure 3: Tax revenue, trade share and average price at different transaction tax rates τ , for selected levels of the elasticity of substitution ϵ between housing and non-housing consumption, and welfare gain from a revenue neutral reform that replaces the transaction tax with a property tax. The black dashed curve in the bottom right panel 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 post-shock allocation 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 realised. For reasons explained in section 2.3, we consider ex ante welfare only in the unit elasticity case.

According to the ex post measure of welfare, 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 and shows that the reform distributes resources in a way that to some extent tends to benefit households that receive a positive income shock more than those that receive a negative shock. In particular, households that move to a very expensive house due to a large positive income shock save the most in absolute terms when the transaction tax is replaced with a property tax.

So how distortionary is the transaction tax? One way to measure it is the ratio of the welfare loss to the gain in tax revenue. At the current 2% tax rate this ratio is between 12% and 15%, depending on the elasticity and using the ex post welfare measure. The relative welfare loss is steeply increasing in the tax rate; for example, going from a 2% tax to a 4% tax, the ratio of the increase in welfare loss to the increase in tax revenue is about 50%. Another way to measure the distortion is 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. 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 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 well above 3 in all cases. Naturally, as the tax rate approaches its revenue maximizing level the MCPF approaches infinity.

To interpret our aggregate results it must be kept in mind that the closed-economy




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

model does not include interregional movements. The simplest assumption we can make is that interregional moves are equally responsive to transaction costs as intra-region moves. Under this assumption our results for tax revenues and trade volumes are otherwise unaffected by interregional movers, but their magnitudes have to be scaled up by the ratio of all movers to within-city movers (i.e., by $63/33 = 1.9$, see section 3.2). Only with further assumptions does the same rescaling apply to euro-valued measures of welfare gains, whereas MCPF (which is a dimensionless measure of welfare costs) would then not need rescaling. This is because, even if interregional movers were equally responsive to transaction costs, their welfare gains from more trading could be different as they depend on the shape of the shock distribution that determines the inframarginal gains from trade.

Interregional moves are more likely to be driven by job opportunities than are moves within commuting zones. Hilber and Lyytikäinen (2017) show that this is indeed the case, at least in the UK. They also find that the transaction tax reduces long-distance and job-related moves less than short-distance and housing-related moves. To the extent that job-related moves are less responsive to transaction costs than housing-related moves, the fact that we abstract from job-related moves makes us more likely to overestimate the welfare costs of transaction taxation (relative to the tax revenue it generates). On the other hand, since income levels are exogenous in our model, we effectively rule out fiscal externalities related to job-related moves. This in turn makes us more likely to underestimate the true welfare cost of transaction taxation.

All the above results are relatively insensitive to the assumed elasticity of substitution. To understand why, recall that we infer the quality distribution separately for each assumed elasticity. The inference is based on the observed relation between house prices and income. Intuitively, if house quality matters less in the utility function (higher 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 investigate the sensitivity to selected parameters in section 7.

5 Comparison with earlier literature

We compare our results with earlier empirical literature based on the welfare effect of a transaction tax change relative to the change in tax revenue, also known as the cost of public funds. Most studies do not explicitly calculate this cost, but as long as changes in tax rate and trading volume are reported we can estimate the cost of public funds using the Harberger triangle approach. The Harberger triangle approach starts from the assumption of a linear demand curve for transactions. (For aggregate welfare it is, of course, immaterial how this reservation value or “gains from trade” is divided between buyers and sellers.) There is only one type of housing, or taxes don’t affect the composition of houses traded and p is their average price. For simplicity and consistency with earlier literature, let’s suppose that the pre-tax price level is constant at p throughout. There are two observed ad valorem transaction tax rates, $\tau_1 < \tau_2$, and two observed transaction volumes $q_1 > q_2$. The change in tax revenue is then $\Delta T = p(\tau_2 q_2 - \tau_1 q_1)$. The change in consumer surplus is $\Delta CS = p\bar{r}(q_2 - q_1)$, where $p\bar{r}$ is the average reservation value for those potential transactions that are realised only at the lower tax rate. If those reservation values are uniformly distributed, i.e, if the demand for transactions is linear, then the average is $p\bar{r} = p(\tau_1 + \tau_2)/2$ and the loss in consumer surplus is literally a triangle in quantity-value space.

The incremental welfare cost of increasing the transaction tax from τ_1 to τ_2 can be measured as a percentage of the incremental tax revenue,

$$-\frac{\Delta CS}{\Delta T} = -\frac{p(\tau_1 + \tau_2)}{2} \frac{(q_2 - q_1)}{p(\tau_2 q_2 - \tau_1 q_1)} = -\frac{1}{2} \left(1 + \frac{\tau_1 q_2 - \tau_2 q_1}{\tau_2 q_2 - \tau_1 q_1} \right) \quad (11)$$

Similarly, $1 - \Delta CS / \Delta T$ can be interpreted as the discrete equivalent of the marginal cost of public funds.

Table 1 lists papers that provide results that can be compared to ours relatively easily and that find statistically significant long-run effect of transaction taxes on the volume of transactions (or moves). For each paper in the table, we show the transaction tax rate change in their study, the associated estimated change in the volume of transactions or moves, and the resulting cost of public funds we computed using (11) in the column “CPF (Harberger)”. We have tried to take the authors’ preferred estimate of the medium or long-run effect that should exclude short-run timing or anticipation effects.

Best and Kleven (2018) exploit discontinuities and time-variation in the UK transaction tax schedule focusing on a “stamp duty holiday” that eliminated a 1% transaction tax in a certain house price bracket. Dachis et al. (2011) compare the mobility of

Toronto residents to the mobility of their suburban neighbors following the introduction of a “Land Transfer Tax” on real estate purchases in Toronto. Eerola et al. (2021) exploit a tax reform in Finland where the transaction tax rate was increased for housing co-operatives, but tax treatment of directly owned single-family houses remained unchanged. The authors complement their empirical analysis using a theoretical model to account for spillovers between the market for directly-owned houses and co-operatives.¹⁰ Fritzsche and Vandrei (2019) use time variation in transaction tax rates by regressing transaction volumes for single-family homes on transaction tax rates in different German states while controlling for short-run timing effects; Table 1 shows the estimated effect of a one percentage point tax rate change around the average tax rate in their data. Finally, Hilber and Lyytikäinen (2017) estimate the effect of the transaction tax on household mobility by comparing households reporting house values below and above a cut-off value where the UK stamp duty tax rate jumps from 1 to 3 percent.

The last two columns show our model-based results. “CPF-Model Direct” refers to the aggregate ex post compensating variation and “CPF-Model Harberger” to the approximation based on (11). The fact that two measures are very close to each other, shows that the Harberger approach gives a very accurate approximation of the direct model-based welfare effect at these relatively small changes in the tax rate found in the empirical literature. For larger changes in taxation the Harberger approximation gets further from the model-based welfare calculation; for example, going from 0% to 8% transaction tax results in an average cost of public funds at 1.58 when calculated directly within our model, whereas applying the Harberger formula to the prices and volumes predicted by our model yields a CPF of 1.74.

Welfare estimates from the empirical literature are broadly in line with our model-based estimates. In particular, both estimates suggest that the cost of public funds increases quite rapidly with the tax rate. The model-based estimates are actually remarkably close to those based on Best and Kleven (2018), Dachis et al. (2011), Eerola et al. (2021).¹¹ However, the baseline results in Hilber and Lyytikäinen (2017) imply a substantially higher welfare cost and the results in Fritzsche and Vandrei (2019) a much lower welfare cost than the model.

We can also compare our model-based results to Buettner (2017). He regresses transaction tax revenues on transaction tax rates in various German states and uses

¹⁰The tax rate for housing co-operatives was increased from 1.6% to 2%. However, the tax base was also broadened to include housing co-operative loans. The authors estimate that the initial effective tax rate was therefore only about 1.5%.

¹¹Dachis et al. (2011) use a somewhat different approach to estimate the welfare effect. The figure they provide in the paper is 1 dollar for every 8 dollar in revenue, which translates into a cost of public funds at 1.125.

	Tax rates		Change in volume	CPF	CPF-Model	
	lower	higher		Harberger	Direct	Harberger
Best and Kleven (2018)	0	0.01	0.10	1.055	1.073	1.067
Dachis et al. (2011)	0	0.011	0.15	1.088	1.080	1.076
Eerola et al. (2021)	0.015	0.02	0.07	1.340	1.273	1.267
Fritzsche and Vandrei (2019)	0.037	0.047	0.07	1.438	1.863	1.875
Hilber and Lyytikäinen (2017)	0.01	0.03	0.37	1.839	1.320	1.305

Table 1: Implied estimates for the welfare cost of public funds per unit of tax revenue (CPF) from previous studies, for a change between two tax rates; our calculation based on the change in volume reported by the authors and the Harberger formula. CPF is the cost in euros per euros of revenue. Last two columns show the CPF from the same change in tax rates in our calibrated setup, using the welfare results from the model (“Direct”) and by applying the Harberger formula to the prices and volumes predicted by the model.

the resulting estimate of the tax revenue elasticity to estimate the associated welfare cost around a 4% average tax rate. His preferred estimate, which includes possible tax sheltering effects, translates into a marginal cost of public funds equal to 1.67. The corresponding figure is 1.8 in our model with log-utility. However, we cannot compare our results directly with Borbely (2018) which studies the impact of progressive transaction taxes in Scotland using both price and time notches; he finds that a unit increase in effective tax rates leads on average to a 5.6% reduction in transaction volume. Equation (11) cannot be applied because this study does not report the average tax rate.

6 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 different 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, but 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

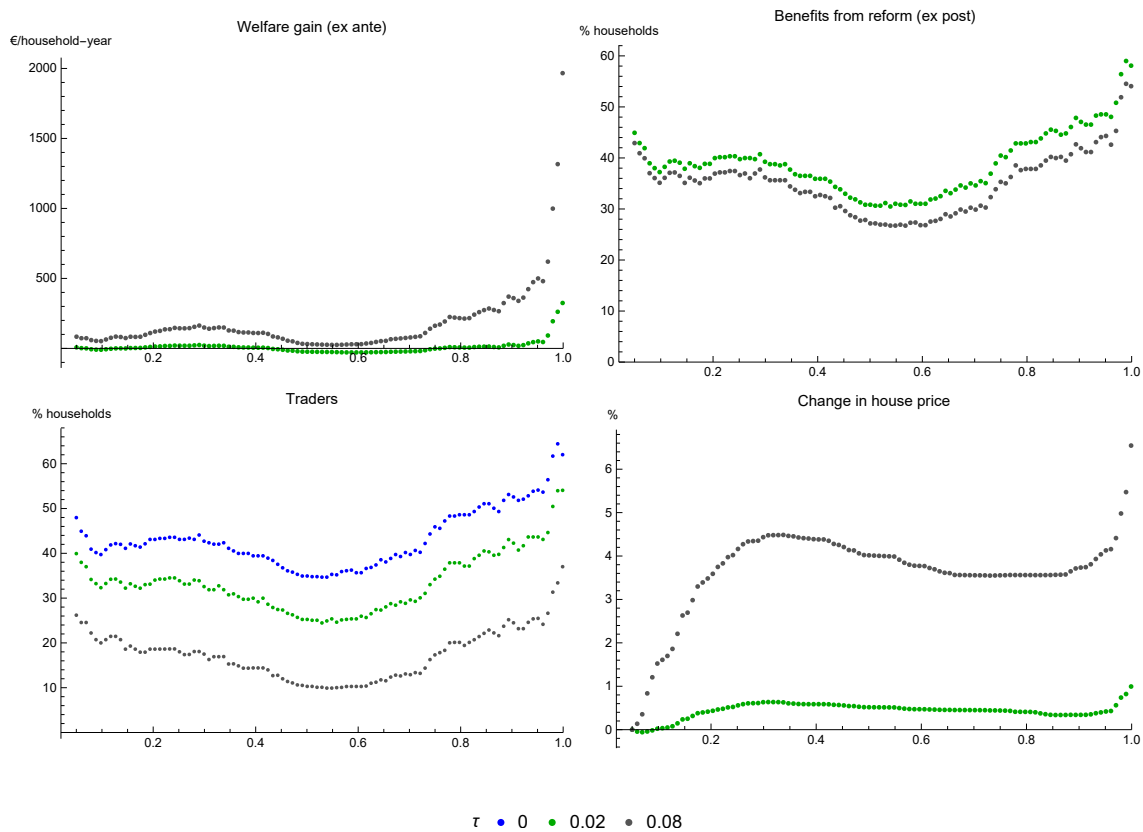


Figure 5: Distributional effects from replacing a transaction tax with a revenue equivalent property tax, with quantiles of house quality on horizontal axis. The case with zero transaction tax in the bottom-left panel shows also the share of traders under any property tax rate.

households have on average the highest levels of welfare. 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. 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 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

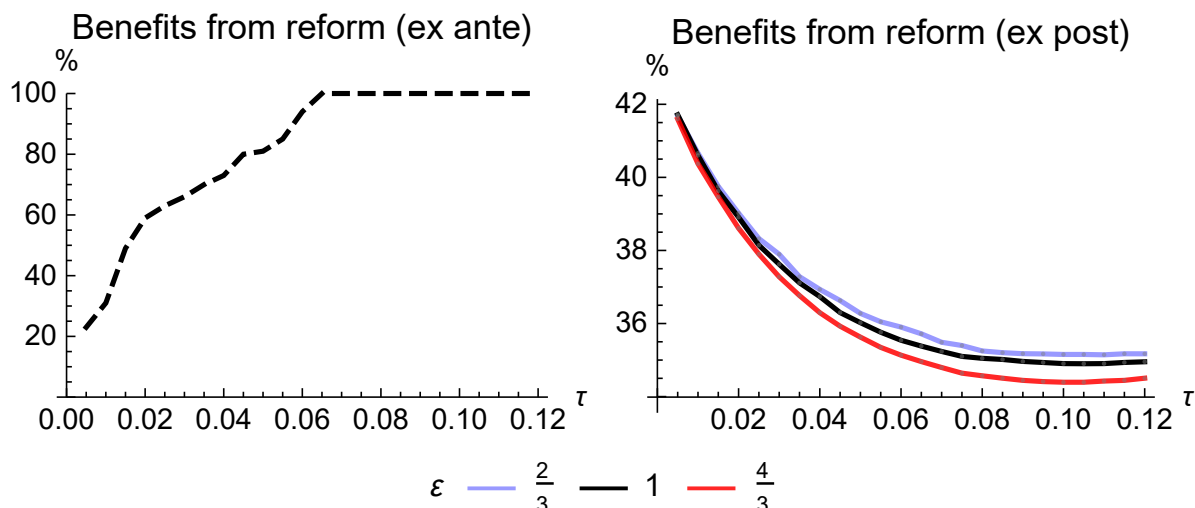


Figure 6: Proportion of households benefiting from a revenue neutral reform that replaces a transaction tax τ with a property tax. Left panel shows the proportion ex ante before households know the realization of their income shock; this is calculated only for the unit elastic case. The right panel shows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

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 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 in our data). The top panel shows a case where the transaction tax is 2% and the bottom panel a case where it is 8%.

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 shade shows the no-trade region under the transaction tax, it is naturally wider in the high-tax case. The darker 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. This no-trade region stems only from the non-fiscal transaction costs and (to a small extent) from the discrete housing quality distribution. 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 going to trade in any case clearly benefit

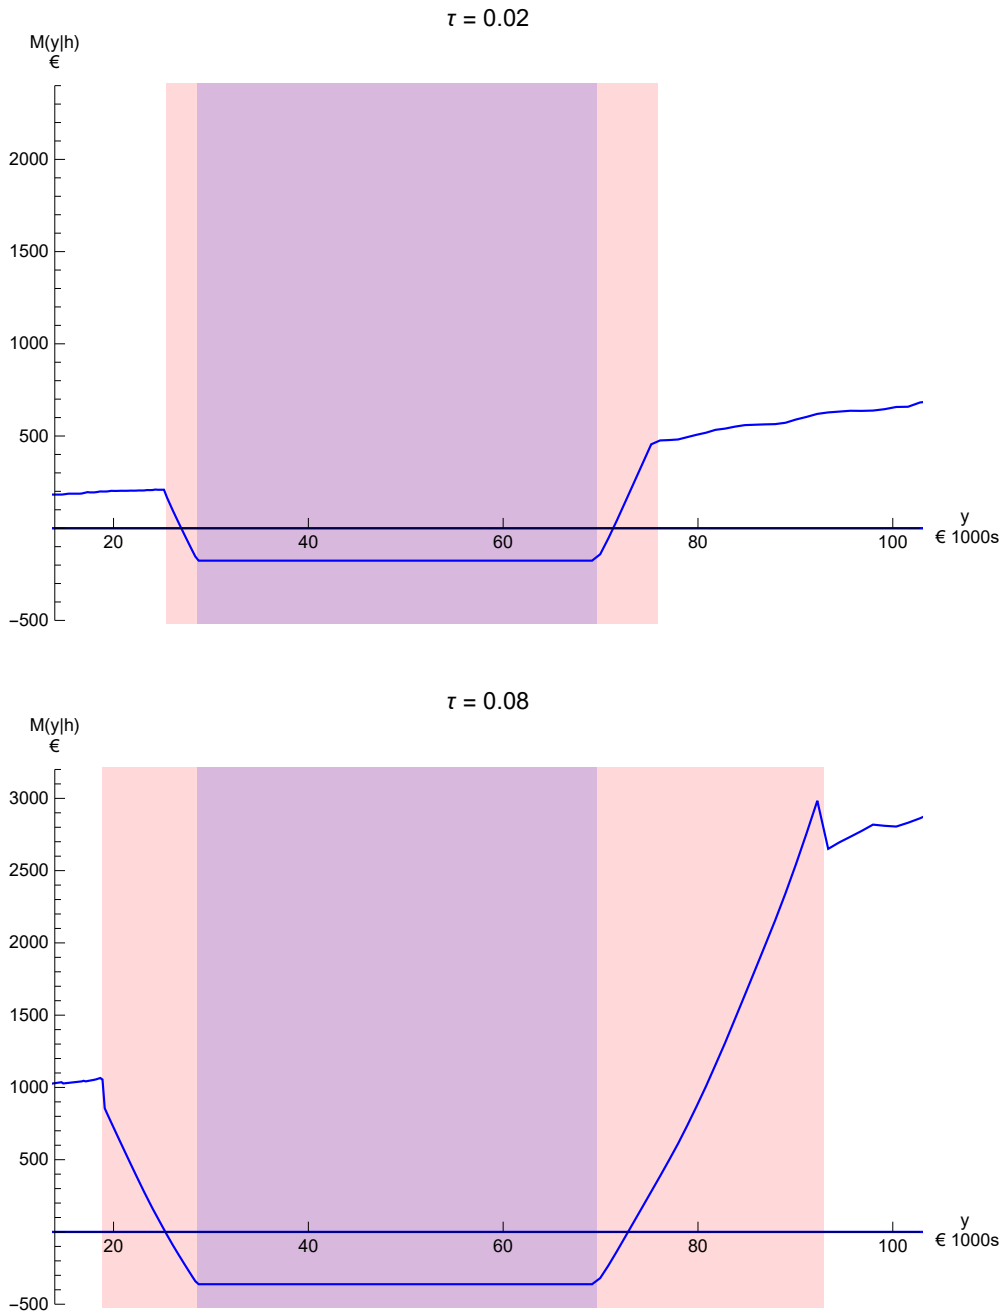


Figure 7: Difference in welfare relative to the equilibrium under a transaction tax of 2% (top) and 8 % (bottom) for households endowed with median house type. Household post-shock income on horizontal axes. The curves show the change in utility caused by a shift to a property tax.

from the change that divides the same aggregate tax burden over a larger number of taxpayers.

Now consider those who are induced to trade by the tax reform. The new marginal traders, namely 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.

7 Sensitivity

In calibrating the model we set the non-fiscal transaction costs exogenously at 4% of the empirical house prices and the interest rate at 5%; next we explore how our results are affected by changing these parameters. We consider “low” and “high” cases where two percentage points are either subtracted from or added to the baseline values of these parameters.

When varying these parameters (one at a time), we set the elasticity of substitution between housing and non-housing consumption at one and recalibrate the remaining parameters with the same targets as above. For instance, higher non-fiscal transaction costs require a higher variance for the income shocks in order to generate the same trade share as before.

		Tax revenue		Trade share		Avg. price	
		€/hh		%		€1000s	
τ		0.02	0.08	0.02	0.08	0.02	0.08
Elasticity	Low	205.8	420.8	33.3	17.7	244.3	234.4
	High	229.0	471.0	33.1	17.0	242.4	229.0
Fixed cost	Low	208.3	322.1	33.1	13.1	242.2	231.3
	High	226.5	535.7	33.3	20.2	243.8	230.9
Interest rate	Low	223.2	469.5	32.9	17.5	242.8	220.8
	High	222.5	453.8	33.3	17.4	242.2	233.7
Calibrated		217.8	446.9	33.3	17.4	243.2	231.2

Table 2: Sensitivity Analysis: Tax revenue, trading volume, and prices.

Tables 2 and 3 display selected results at transaction tax rates of 2% and 8%. “Fixed cost” refers to the non-fiscal transaction cost. For completeness, we also present the results for the higher and lower elasticities of substitution that were already considered throughout our main analysis. The bottom rows refer to the baseline calibration at unit elasticity (log-utility).

Table 2 displays the average (annual) tax revenue per household, the share of households that trade, and the average house price. The average house price at a 2% tax

rate is very similar across all cases, because we set the tax rate at 2% when calibrating the model and always target the same empirical price distribution. However, the variation in the average price is fairly limited also at an 8% tax rate.

As already seen in Section 4, the relationship between the transaction tax rate and trade share or tax revenue is not much affected by the elasticity. The same is true of the interest rate. By contrast, non-fiscal transaction costs make a clear difference to our results. The higher they are the smaller is the share of a given transaction tax of the overall transaction costs. As a result, increasing the transaction tax rate from 2% to 8% lowers the transaction volume by less and increases the tax revenue by more when non-fiscal transaction costs are higher.

		Welfare gain		Winners		E Welfare gain		E Winners	
		€/hh		%		€/hh		%	
τ		0.02	0.08	0.02	0.08	0.02	0.08	0.02	0.08
Elasticity	Low	29.0	258.9	39.0	35.2				
	High	29.1	265.2	38.6	34.6				
Fixed cost	Low	44.6	318.3	42.4	39.2	32.0	264.8	96	100
	High	22.4	224.0	37.3	33.6	-4.7	117.1	34	70
Interest rate	Low	26.1	232.1	38.1	33.9	8.0	157.6	18	58
	High	30.9	277.0	38.9	35.3	10.3	193.6	60	100
Calibrated		29.2	260.2	38.9	35.0	8.9	179.1	59	100

Table 3: Sensitivity Analysis: Welfare.

Table 3 displays the aggregate welfare gain associated with switching to a revenue equivalent property tax and the share of households who would be better off 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”) to the ex ante comparison. For reasons explained in Section 2.3, we conduct the ex ante welfare comparison only for the unit-elastic case.

One might expect other transaction costs to make transaction taxes more harmful. However, here higher non-fiscal transaction costs are associated with lower aggregate welfare gain from the reform and also smaller share of households benefiting from it. Higher non-fiscal costs reduce the extent to which the transaction tax alone can distort the allocation, which reduces both ex ante and ex post welfare gains of the reform (also relative to tax revenue). With a high non-fiscal transaction cost and a relatively low initial transaction tax rate, the ex-ante welfare gain is even negative. This reflects the observation, discussed in section 6, that here replacing a transaction tax with a property tax turns out to shift resources between households in a way that tends to decrease expected utility (as evaluated before the realization of the income shocks),

and this effect can dominate at a low enough tax rate.

By contrast, increasing non-fiscal transaction costs decreases the share of (ex-post) winners. This again illustrates that a higher aggregate welfare cost of the transaction tax may be associated with a lower share of households that would benefit from a reform that replaces it with a property tax.

Changing the interest rate has smaller effects on the aggregate welfare gains from switching to a revenue neutral property tax. The share of ex-ante winners from the reform appears to be very sensitive to both the non-fiscal transaction cost and interest rate.

In some countries, the transaction tax is progressive rather than linear as in Finland and in our baseline model. To see whether a progressive transaction tax has very different welfare effects than a linear tax, we compare the linear 2% transaction tax with a simple progressive transaction tax scheme where the tax rate is 0% up to a threshold and 4% for the portion of the house value exceeding the threshold. We choose the threshold so that the progressive tax generates exactly the same tax revenue as the linear one. The threshold is close to the value of the house in the 23rd percentile of the price distribution. We also compare a 4% linear tax with a similarly constructed progressive tax with a 8% tax rate. In this case, the threshold is close to the 20th percentile in the quality distribution. We assume log-utility here and stick to a linear property tax as the reform option.

	Linear 2%	Prog. 4%	Linear 4%	Prog. 8%
Welfare gain (€/hh)	29.2	31.3	94.0	101.3
Winners (%)	38.9	34.8	36.7	34.6
E Welfare gain (€/hh)	8.9	9.4	53.3	55.6
E Winners (%)	59	24	73	27

Table 4: Welfare under linear and progressive transaction taxes.

Table 4 displays the results. The progressive schemes with 4% and 8% tax rates generate the same tax revenue as linear schemes with 2% and 4% tax rates, respectively. The first row reveals that, in terms of aggregate welfare, a progressive transaction tax is more distortionary than a revenue equivalent linear tax. However, the difference is not very large. The average annual ex-post welfare gain associated with replacing the transaction tax with a revenue equivalent property tax increases either from 29 to 31 euros, or from 94 to 101 euros, when moving from a linear to a progressive transaction tax. The result that changes the most is the share of ex-ante winners, i.e., the share of households that would prefer to replace the transaction tax with a revenue equivalent

property tax before knowing the realization of their income shock. This share is much smaller under a progressive transaction tax scheme than under a linear scheme.

Our data is not top coded. It may therefore include erroneous outliers at the top of the distribution, which may affect our results. In order to see the extent to which this matters for our results, we alternately removed the top 3 and the top 5 percent of the data (based on house values), then re-estimated the relation of income and house value \hat{y} , and then recalibrated the model following the same procedure as in the baseline case. To briefly summarise our findings, removing the top end of the distribution improves the fit of the model somewhat. With the top 3 percent of observations dropped, the maximum (absolute) deviation of equilibrium house prices from their empirical counterparts is about 4% percent (positive), compared to around 5% (negative) in the baseline case. To be sure, the top percentiles of the distribution make a disproportionate difference to the average tax revenue per household. The average tax revenue generated by a 2% transaction tax falls from 218 euros to 188 euros (when top 3% dropped) or 171 euros (5% dropped). The average welfare costs of the transaction tax fall by roughly the same proportion as aggregate tax revenue. For instance, the average ex post welfare gain from replacing a 2% transaction tax with a property tax falls from 29 euros to 27 euros when we drop the top 5%. The share of ex ante winners decreases by two percentage points while the share of ex post winners decreases by about one percentage point.

8 Matching frictions

So far we have focused on transaction costs as the only friction distorting the matching of different houses and households. However, there are also other type of frictions in housing markets. In particular, potential buyers usually want to visit and inspect a house before possibly buying it. Since this takes time, they cannot visit all houses on the market. Moreover, at any point in time, only a subset of all houses are on the market in the sense that potential buyers could even consider buying them.

The implications of this type of frictions are analysed in search and matching models of the housing market.¹² These models tend to focus on price determination, transaction volume and time-on-the market rather than on the matching of house and household types. However, search frictions may also distort the allocation of distinct house types across households—the focus of our study. By abstracting from other frictions we may therefore be overstating the distortionary impact of the transaction tax.

¹²For recent major contributions to this literature see for instance Ngai and Tenreyro (2014) and Piazzesi et al. (2020)

In this section we add another type of friction to the model on top of transaction costs. Specifically, we assume that only a subset of house types are on the market at a given point time. This is a straightforward way of capturing other frictions that distort the cross-sectional matching without departing from the the assignment model framework that is useful in capturing the matching of houses and households. The model is still perfectly competitive: all submarkets clear separately just like the only market in our main model; all households take prices as given, because there is still a continuous distribution of households incomes. This friction captures random variation in house types that happen to be in the same market at the same time, but not frictions coming from variation in household bargaining strength over match-specific rents. (The latter would require a fundamentally different type of a model.)

We implement the submarkets by drawing random samples of s house types from the calibrated house type (quality) distribution of the baseline model, where the size of submarkets captures “market thickness”. We assume that households with a given house type have also the same distribution of incomes y as in the baseline calibration. The subsamples are drawn with replacement, so some house types may be drawn several times for a given sample, in this case the mass of households endowed with the “repeated” house type is correspondingly higher. Otherwise the calibration is the same as in the baseline model with log-utility described in section 3, and market equilibrium in each submarket is again solved as described in Appendix B.

We treat each random sample as giving rise to a separate and independent submarket. For each submarket, we again solve the equilibrium allocation and prices for a range of transaction tax rates and their associated revenue equivalent property taxes.¹³ The price of a given house type varies across realised submarkets where it is present because equilibrium prices depend on the entire initial distribution in the market.

As the number of house types s gets smaller the house quality distribution becomes increasingly coarse, implying that many households are not able to buy a house that would be even close to the one they would buy if all house types were on the market. We repeat the exercise for market sizes $s = 5, 15, 30, 60, 100$.¹⁴ For each “submarket size” s we repeat the random draw so as to keep the number of draws at 2000 for each s (up to integer constraint), thus we draw more submarkets for smaller s . We solve the equilibrium prices and allocations in each submarket at a range of transaction tax rates. At $s = 100$ the only difference with our baseline setup is due to the sampling with replacement.

One might expect this new matching friction to cause a large drop in the trade

¹³We set the property tax rates separately for each submarket to ensure that they generate the same tax revenue as the corresponding transaction tax.

¹⁴The corresponding average numbers of distinct house types in a submarket turn out to be 5, 14, 26, 45, and 65.

share. However, the average trade share across submarkets decreases only slowly as we decrease market thickness s . For instance, compared to the baseline model at $\tau = 0.02$, frictions at level $s = 30$ reduce the trade share by about one percentage point. Intuitively, as long as there are any transaction costs (whether fiscal or otherwise), it is not attractive to move to a house that is too similar to one's current house. Therefore making the house distribution coarser has little effect on the share of households who trade, because even when only a small number of house types are drawn into a submarket it is unlikely that they would be very close to each other in the quality distribution.

Figure 8 displays the aggregate welfare results (averaged over all submarkets) for each level of market thickness s and over a range of transaction taxes. The figures in the top row show the average euro-valued ex-ante and ex-post welfare gains from replacing the transaction tax rate with a revenue equivalent property tax. The figures in the bottom row show the corresponding share of households that would be better off due to the reform. The corresponding results based on the log-utility case in the baseline model reported in sections 4 and 6 are displayed with dashed curve.

Figure 8 shows that the aggregate welfare results are qualitatively similar to those obtained with the baseline model, and have a similar dependence on the tax rate. Both the ex-post and ex-ante welfare gains from replacing the transaction tax rate with a property tax are diminished at lower levels of market thickness s . The results are thus indeed quantitatively sensitive to this new friction. Intuitively, the relative importance of the transaction tax in distorting the allocation diminishes with other frictions. However, with s equal to 100 or 60, the results are still very close to baseline results.

The question then is what range for the market thickness parameter s could be considered realistic. As we decrease s the equilibrium prices vary more and more across the submarkets, reflecting the fact that there is effectively more "randomness" in the initial distribution. It seems unlikely that the price of a house of the exact same type would vary by more than 10% just due to random matching frictions. We thus look for price variation across submarkets for different s to determine a reasonable range for it.

Figure 9 plots the variability of prices across the realizations at selected levels of market thickness. We measure price variability as the average absolute deviation relative to the price of the same house type in the baseline calibration. For graphical clarity, these averages are further averaged over deciles of house quality. The price dispersion is already noticeable at $s = 100$ and gets higher the less thick is the market. The average absolute deviation of a given house over the whole quality distribution type is 6% at $s = 30$ and 10% at $s = 15$, and 22% at $s = 5$. And these are just

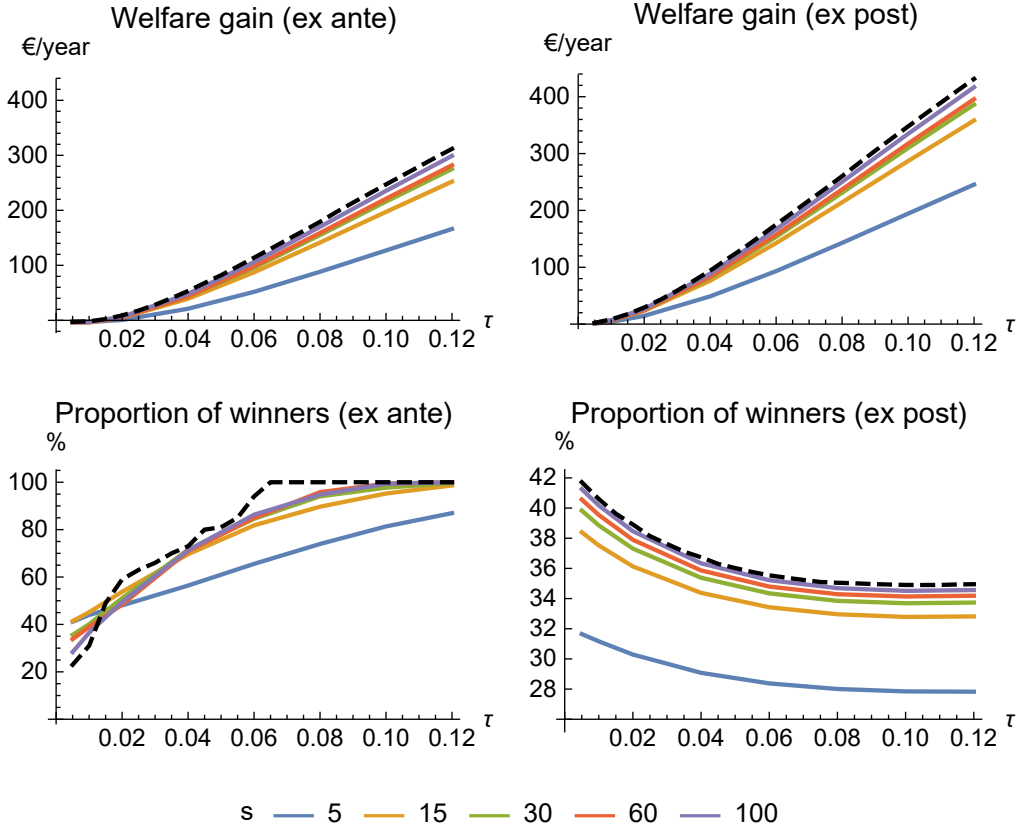


Figure 8: Welfare results at various levels of transaction tax rates τ and market thickness s . The top panel shows average welfare gain per household from replacing the transaction tax with a revenue equivalent property tax. The bottom panel shows the corresponding percentage of households that would be better off under the reform. The black dashed curves show the same results in the baseline calibration (see section 2.3).

the average deviations; at $s \leq 30$ variations over 10% for the exact same house type are common. Given that here there is no measurement error for house quality, the price dispersions generated at $s = 5$ and $s = 15$ seem unrealistically high to us. On the other hand, as noted above, at $s = 15$ the welfare effects are still relatively close to our baseline results. We therefore conclude that other frictions that distort the matching between households and houses reduce the importance of the transaction tax for welfare, but seem unlikely to substantially diminish the welfare cost of the transaction tax compared to our baseline results.

9 Conclusion

We evaluated the welfare effects of housing transaction taxes within a new one-sided assignment model framework. We used data from the Helsinki metropolitan region




Figure 9: Average absolute deviation of price realizations from baseline at transaction tax rate $\tau = 0.02$, by house quality decile and level of market thickness s . Deviations are calculated relative to the price of the same house type in the baseline calibration.

and considered a counterfactual tax reform, where the transaction tax was replaced with a revenue equivalent ad valorem property tax. The welfare gain from the reform is moderate at the 2% rate currently used in Finland, but steeply increasing in the tax rate.

Despite clear aggregate welfare gains from replacing the transaction tax with a property tax, many households are worse off with 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. The ex post perspective naturally leads to an uneven distribution of the gains and losses from the tax reform as households that know they would not be trading even with the reform can only lose. In addition, the higher the tax burden the harder it is for the marginal traders who are induced to trade as a result of the tax reform to end up better off. This is because under a property tax they all need to contribute their share to the tax burden. It is not essential that the replacing tax be a property tax; the burden of practically any alternative tax is going to be more evenly shared between traders and non-traders.

The ex post view may seem contrived, as everyone knows for sure whether they

are traders or not. However, the ex ante perspective to welfare gains is also quite “extreme”, as everyone is equally ignorant about their trading propensity. In reality, some households are more likely to trade than others, yet few are completely certain; this is consistent with political support for transaction taxes even at very high and distortionary rates.

Introducing life cycle features to our framework could probably provide more insights into this issue. For instance, young households, who realise that they are likely to want to move several times over the rest of their lives, should presumably be in favor of eliminating the transaction tax based on their expected life time utility. More generally, buying and selling in the housing market also reflects life cycle considerations that are often anticipated, and not just unexpected shocks to (permanent) income or preferences. A life cycle model would also provide a framework for accounting for the role of credit constraints in the housing market. On the other hand, introducing life cycle aspects would have forced us to simplify the cross-sectional aspects of the housing market which we are now focusing on. This is because it would require a setup with intertemporal optimization on the part of households.

In order to focus on the allocational effects of the transaction tax in the housing market, we abstracted from other mechanisms that are also likely to be important channels of welfare effects. To be sure, our results on the welfare costs of transaction taxation need to be interpreted with these limitations in mind. For one thing, our model does not address the impact of moving to better job opportunities; we treat incomes as exogenous so the model covers only one housing market with a common labor market. In reality, transaction taxes may also increase mismatch in the labor market.

The rent-or-buy decision is not part of our model. This margin can also be expected to be distorted by a transaction tax, because there is no tax on changing tenants. It is not clear whether incorporating the tenure choice would increase or decrease the aggregate welfare cost of the transaction tax. On the one hand, it would allow households to avoid high transaction taxes by choosing renting over owning, thus improving the allocation. On the other hand, it would also make the transaction tax base more elastic.

We believe that at moderate transaction tax rates the tenure choice is not a very significant determinant of the welfare effects of a transaction tax. The rental housing stock in the Helsinki MA consists largely of subsidised social housing, which is allocated to households with relatively low income. Single-family homes and large family apartments are virtually non-existent in the rental market, possibly because the tax system incentivizes wealthier households, whose tenure choice is not limited by borrowing constraints, to own rather than rent. Most households are thus unlikely to change their tenure choice following moderate changes in the transaction tax rate.

One aspect of taxation that we did not consider is administrative costs. The administrative costs of a transaction tax are arguably lower than those of a property tax, because the taxable value of a house is easy to observe when it is based on a recent market transaction but may require a separate periodical assessment under a property tax. At the same time, reported transaction prices may be easier to manipulate than official assessed values used in property taxation. Nevertheless, at a sufficiently low tax rate the aggregate benefit from using a property tax could conceivably be overturned by a difference in administrative costs.

References

- Bérard, G., Trannoy, A., 2018. The impact of the 2014 increase in the real estate transfer taxes on the french housing market. *Economie et Statistique* 500, 179–200.
- Besley, T., Meads, N., Surico, P., 2014. The incidence of transaction taxes: Evidence from a stamp duty holiday. *Journal of Public Economics* 119, 61–70.
- Best, M., Kleven, H., 2018. Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK. *Review of Economic Studies* 85, 157–193.
- Borbely, D., 2018. Limiting the distortionary impacts of transaction taxes: Scottish stamp duty after the Mirrlees review. *Strathclyde Discussion Papers in Economics* .
- Buettner, T., 2017. Welfare cost of the real estate transfer tax. *CESifo Working Paper* 6321.
- 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, 327–354.
- Eerola, E., Harjunen, O., Lyytikäinen, T., Saarimaa, T., 2021. Revisiting the effects of housing transfer taxes. *Journal of Urban Economics* 124.
- European Commission, 2012. Quarterly report on the Euro area, Vol. 11.
- Fritzsche, C., Vandrei, L., 2019. The German real estate transfer tax: Evidence for single-family home transactions. *Regional Science and Urban Economics* 74, 131–143.
- 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.
- Holter, H.A., Krueger, D., Stepanchuk, S., 2019. How does tax progressivity and household heterogeneity affect Laffer curves? *Quantitative Economics*, Forthcoming.

- Lundborg, P., Skedinger, P., 1999. Transaction taxes in a search model of the housing market. *Journal of Urban Economics* 45, 385–399.
- Mirrlees, J., Adam, S., Besley, T., Bell, R., Bond, S., Chote, R., Gammie, M., Johnson, P., Myles, G., Poterba, J.M., 2011. *Tax by design*. Oxford University Press.
- Määttänen, N., Terviö, M., 2009. Income distribution and housing prices: An assignment model approach. HSE Working paper .
- Määttänen, N., Terviö, M., 2014. Income distribution and housing prices: An assignment model approach. *Journal of Economic Theory* 151, 381–410.
- Ngai, L.R., Tenreyro, S., 2014. Hot and cold seasons in the housing market. *American Economic Review* 104, 3991–4026.
- O’Sullivan, A., Sexton, T.A., Sheffrin, S.M., 1995. Property taxes, mobility, and home ownership. *Journal of Urban Economics* 37, 107–129.
- Piazzesi, M., Schneider, M., 2016. Housing and macroeconomics, Elsevier. volume 2 of *Handbook of Macroeconomics*, pp. 1547–1640.
- Piazzesi, M., Schneider, M., Stroebel, J., 2020. Segmented housing search. *American Economic Review* 110, 720–59.
- Slemrod, J., Weber, C., Shan, H., 2017. The behavioral response to housing transfer taxes: Evidence from a notched change in DC policy. *Journal of Urban Economics* 100, 137–153.
- Stokey, N.L., 2009. Moving costs, nondurable consumption and portfolio choice. *Journal of Economic Theory* 144, 2419–2439.
- Trabandt, M., Uhlig, H., 2011. The Laffer curve revisited. *Journal of Monetary Economics* 58, 305–327.

Appendix

A Inferring the quality distribution

Inference of house qualities \mathbf{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 the continuous income distribution, for any house type j there must be trading households that are indifferent between moving to j or its immediate neighbor in the quality order.

Consider a household endowed with a house of 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 Section 2, 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 $(p_{j+1} - p_j)(1 + \tau_T + \phi)$. For this households we have the indifference condition

$$u(x_j, \bar{y}_j) = u(x_{j+1}, \bar{y}_j - (p_{j+1} - p_j)(1 + \tau_T + \phi)). \quad (12)$$

Under CES-utility this can be solved for Under CES-utility this can be solved for

$$x_{j+1}^\rho - x_j^\rho = \bar{y}_j^\rho - (\bar{y}_j - (p_{j+1} - p_j)(1 + \tau_T + \phi))^\rho. \quad (13)$$

Everything on the right hand side is either data or parameters for which we can assume reasonable values ($\rho = \varepsilon/(\varepsilon - 1)$, ϕ). With a sufficiently fine grid $\bar{y}_j \approx \underline{y}_j$ we can treat both as approximations of the same curve y_j , which captures the average non-housing consumption of households in house type j .¹⁵ The CES inference formula under transaction costs is

$$\hat{x}_h = \left(x_1^\rho + \sum_{j=2}^h ((y_{j-1} + (p_j - p_{j-1})(1 + \tau_T + \phi))^\rho - y_{j-1}^\rho) \right)^{\frac{1}{\rho}} \quad (14)$$

where x_1 is an arbitrary positive constant. It would be hard to come up with a reasonable range of values for this abstract quality measure. Crucially, x_1 washes out of all differences of the type $\hat{x}_h^\rho - \hat{x}_j^\rho$.¹⁶ 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

$$\hat{x}_h = x_1 \prod_{j=1}^h \frac{\bar{y}_j}{\bar{y}_j - (p_{j+1} - p_j)(1 + \tau_T + \phi)} \quad (15)$$

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.

¹⁵Caveat: The non-traders of type j need not have an average y in the traders' post-trade range $[\underline{y}_j, \bar{y}_j]$, but if transaction costs are low then the no-trade region is not very wide and they are close.

¹⁶Apply the inference equation (14) separately to \hat{x}_h and \hat{x}_j , raise both to the power ρ , and take their difference; then x_1^ρ cancels out of the difference.

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 \mathbf{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. However, this discontinuity is not relevant when considering only those who trade. The bounds of the trade-to region (see Figure 1) can be found by solving for the no-trade intervals in a “continuous” world where transaction costs are incurred even if the household does not trade. We thus solve the bounds of the trade-to intervals $\bar{y}_k(\mathbf{p})$ and $\underline{y}_k(\mathbf{p})$ from

$$\bar{y}_k(\mathbf{p}) = \{y \text{ s.t. } u(x_k, y - \tau p_k - \xi_k) = u(x_{k+1}, y + p_k - (1 + \tau)p_{k+1} - \xi_{k+1})\}, \quad (16)$$

$$\underline{y}_k(\mathbf{p}) = \{y \text{ s.t. } u(x_k, y - \tau p_k - \xi_k) = u(x_{k-1}, y + p_k - (1 + \tau)p_{k-1} - \xi_{k-1})\}. \quad (17)$$

Positive assortative matching implies that, in equilibrium, $\bar{y}_N = y_{max}$ and $\underline{y}_1 = y_{min}$.

The bounds of the actual no-trade intervals extend wider because they include those who are deterred from moving by the transaction costs. Consider the k -type households, i.e., the households endowed with house x_k . As in the case of trade-to intervals, we need to find the bounds of the income interval at which a k -type will choose to not trade, denoted by $\bar{Y}_k(\mathbf{p})$ and $\underline{Y}_k(\mathbf{p})$. 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 $\bar{Y}_k(\mathbf{p})$ 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.

We use the following procedure to find out the value of $\bar{Y}_k(\mathbf{p})$. First, notice that as long as transaction costs are strictly positive, $\bar{Y}_k(\mathbf{p}) > \bar{y}_k(\mathbf{p})$. Second, notice that those who trade will end up in one of the trade-to intervals

$$x = x_j, y \in [\underline{y}_j(\mathbf{p}), \bar{y}_j(\mathbf{p})]. \quad (18)$$

Households with incomes above 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 trade-to allocations $\bar{y}_{k+s}(\mathbf{p})$. The first question is, at which income level y is a household endowed with a house k -type house exactly able to pay the price difference and the transaction tax in order to swap into a house of type $k + s$ and have just the amount of money left over to consume at the upper bound of

the trade-to interval, $\bar{y}_{k+s}(\mathbf{p})$. The answer is

$$\tilde{y}_{k,s} = \bar{y}_{k+s}(\mathbf{p}) + (1 + \tau)p_{k+s} + \xi_{k+s} - p_k. \quad (19)$$

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

$$u(x_{k+s}, \bar{y}_{k+s}(\mathbf{p})) \geq u(x_k, \bar{y}_{k+s}(\mathbf{p}) + (1 + \tau)p_{k+s} + \xi_{k+s} - p_k) \quad (20)$$

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 $\tilde{y}_{k,s} \geq y_{max}$ which would show that k -types don't trade up so that $\bar{Y}_k(\mathbf{p}) = y_{max}$.

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

$$u(x_k, y) = u(x_{k+s}, y - (1 + \tau)p_{k+s} - \xi_{k+s} + p_k). \quad (21)$$

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

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

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 y . They will buy a type- k house, where $k > j$, if the following two conditions are satisfied: 1) Their resulting non-housing consumption would be in the same range as the non-housing consumption of those type- k households who would consume their endowment under unavoidable transaction costs, that is $y + p_j - p_k - \tau p_k - \xi_k \in [\underline{y}_k(\mathbf{p}) - \tau p_k - \xi_k, \bar{y}_k(\mathbf{p}) - \tau p_k - \xi_k]$; 2) Their income level is outside the no-trade interval of type- j households.

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

$$\begin{aligned} y &\leq \bar{y}_k(\mathbf{p}) + p_k - p_j, \\ y &\geq \max\{\underline{y}_k(\mathbf{p}) + p_k - p_j, \bar{Y}_j(\mathbf{p})\}. \end{aligned} \quad (22)$$

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

$$\begin{aligned} y &\leq \min \{ \bar{y}_k(\mathbf{p}) + p_k - p_j, \underline{Y}_j(\mathbf{p}) \}, \\ y &\geq \underline{y}_k(\mathbf{p}) + p_k - p_j. \end{aligned} \quad (23)$$

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

$$\underline{Y}_k(\mathbf{p}) < y \leq \bar{Y}_k(\mathbf{p}). \quad (24)$$

Total demand for type- k houses is

$$\begin{aligned} Q_k(p) &= \sum_{j=0}^{k-1} \max \left\{ 0, F_j(\bar{y}_k(\mathbf{p}) + p_k - p_j) - F_j \left(\max \{ \underline{y}_k(\mathbf{p}) + p_k - p_j, \bar{Y}_j(\mathbf{p}) \} \right) \right\} \\ &\quad + F_k(\bar{Y}_k(\mathbf{p})) - F_k(\underline{Y}_k(\mathbf{p})) \\ &\quad + \sum_{j=k+1}^N \max \left\{ 0, F_j(\min \{ \bar{y}_k(\mathbf{p}) + p_k - p_j, \underline{Y}_j(\mathbf{p}) \}) - F_j(\underline{y}_k(\mathbf{p}) + p_k - p_j) \right\}. \end{aligned} \quad (25)$$

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

In order to find the equilibrium prices, we have written a Matlab function that returns the excess demand for each house type for a given price vector and a given initial (or post-shock) allocation of house qualities and incomes. This function first determines the trade-to 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.