The Incidence of Rent Subsidies: Evidence on Rents, Housing Choices and Supply*

Essi Eerola[†] Teemu Lyytikäinen[‡] Tuukka Saarimaa[§]
Tuuli Vanhapelto[¶]

Abstract

This paper asks if, and in particular *how* and *when*, rent subsidies can have the adverse effect of increasing rents. Our research design is based on a housing allowance reform in Finland in 2015 that made the allowance scheme substantially more generous for some types of housing units but not for others. We find that large increases in housing allowances for affected housing units had little or no effect on their rents relative to other units. Thus, the incidence of the reform was largely on allowance recipients and not on their landlords. To understand these small rent effects, we analyze the underlying changes in rental demand and supply. The reform led to a statistically significant but economically very small change in recipients' housing choices, and at most a modest change in rental supply. These observations suggest that housing allowances can be an effective policy tool even in supply-constrained contexts, if rental demand is relatively inelastic with respect to the policy.

Keywords: Housing allowance, rent subsidy, incidence, housing demand, housing supply.

JEL codes: H22, R28

*The authors thank Oskari Harjunen, Christian Hellwig, Thierry Magnac, Daniel Schmidt, Olmo Silva and Miguel Zerecero, and audiences at the Meeting of the Urban Economics Association in London, International Association for Applied Econometrics annual conference in 2022, Helsinki GSE Public and Labour Economics seminar, Finnish Economic Association meeting in Tampere, Echoppe Conference on The Economics of Housing and Housing Policies in 2022, University of Turku, TSE Econometrics and PhD Workshop as well as Nordic Urban Economics Workshop for useful comments. Lyytikäinen thanks the Academy of Finland for funding (grant no. 315591 and 364189). Vanhapelto thanks Yrjö Jahnsson Foundation, Jenny and Antti Wihuri Foundation and Nordea Pankin säätiö for funding. The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the Bank of Finland.

[†]essi.eerola@bof.fi, Bank of Finland.

[‡]teemu.lyytikainen@vatt.fi, VATT Institute for Economic Research.

[§]tuukka.saarimaa@aalto.fi, Aalto University and Helsinki Graduate School of Economics.

[¶]tuuli.vanhapelto@glasgow.ac.uk, Adam Smith Business School, University of Glasgow.

1 Introduction

Housing affordability is a pressing issue in urbanized societies. To alleviate this challenge, many governments support low-income households through large targeted rent subsidy programs, such as housing allowances or housing vouchers.[1] These transfers can shield households from financial distress, evictions and homelessness (Fetzer et al., 2023). However, whether these transfers work as intended depends crucially on the extent to which they increase rents and end up benefiting landlords instead of subsidy recipients. The existing estimates of the rent effects of these transfers from various countries range from a pass-through of zero to a pass-through of more than 50% (see, for example, Gibbons and Manning 2006; Fack 2006; Brewer et al. 2019; Eerola and Lyytikäinen 2021). We still do not understand why these estimates vary so much depending on the context. Theoretically, the pass-through of subsidies to rents depends on the induced changes in rental demand and supply, but so far, there is very little empirical evidence on the effects of housing subsidies on quantities supplied and consumed. Without understanding these mechanisms, it is difficult for policymakers to understand what pass-through effects they can expect when considering policy reforms.

In this paper, we combine rich population-wide register data with a compelling quasi-experimental research design to shed light on whether and in which contexts rent subsidies can have the adverse effects of increasing rents. In particular, we provide evidence on the mechanisms through which rent subsidies can be expected to increase rents. Our research design is based on a major reform of the Finnish housing allowance (HA) system in 2015. The reform substantially increased allowances for some types of housing units, while changes in other unit types were small. These changes ranged from 0 to up to 200 euros per month depending on unit and household type (mean HA payment in the sample period was around 300 euros per month). We leverage this variation to identify the rent effects of allowances using a differences-in-differences design.

We start by using register data to compare the changes in the rents paid by HA recipients in these different types of units. We then analyze the demand- and supply-side responses to understand the drivers of the rent effects. To study demand-side

[1] In 2020, the annual rental subsidy spending was 0.9% of GDP in Finland, 1.3% in the UK, 0.73% in Germany and 0.69% in France. The average of the OECD-25 countries was 0.3%. For more details, see the OECD Affordable Housing Database: https://www.oecd.org/housing/data/affordable-housing-database/.

changes, we use register data on households' moving patterns and housing choices and assess whether the HA reform led to changes in recipients' housing choices. To study supply-side changes, we use register data on both new housing construction and on the conversion of units from the owner-occupied sector to the rental sector and assess how the availability of rental housing changed during our analysis period.

Throughout our analysis, we do not find evidence of the increased HA passing through to landlords in the form of higher rents. In our preferred specification, we measure treatment exposure by calculating the predicted change in HA for a given housing unit caused by the changes in policy parameters, holding constant the pre-reform recipient and rental contract characteristics. An additional euro of predicted HA change translates to 0.9 euros of actual HA change (standard error 0.034 and F-statisic >700, suggesting a very strong first-stage effect). The point estimate for the pass-through, the effect of a one euro increase in HA on rent, is roughly 0.03 euros. The standard error is approximately 0.018, suggesting that we can rule out even moderate effects of the reform on rents with a high degree of confidence. These regressions control for unobserved quality through housing unit fixed effects and are estimated using new rental contracts only, ruling out rent stickiness as a potential reason behind the small rent effects. We also verify that our results are robust to a wide range of alternative ways to compute the treatment exposure.

To understand why the rent effects are so small, we then zoom in on the behavioral changes in the demand- and supply-sides of the rental market. We write down a stylized conceptual framework to illustrate how the rent effects of allowances depend on demand- and supply-side elasticities. Housing allowances can increase rents if rental demand by recipients is relatively elastic and if rental supply is relatively inelastic, so that the increased demand is reflected in higher rents instead of higher quantities. Thus, our observation of the low pass-through of HA increases to rents could be either due to large changes in supply of treated units or due to small changes in the demand for treated units. Which is the case is of first-order importance for understanding how rent subsidies work also in other contexts.

When studying HA recipients' housing choices, we ask (*i*) are HA recipients more likely to stay in units that received a large increase in HA, and (*ii*) do they start choosing units with large HA increases when they move after the reform. We do not observe changes in the propensity of recipients to move out of units that received a large HA increase after the reform, but conditional on moving, households start choosing units with slightly larger HA increases. The effect kicks in sharply after the reform and is precisely estimated but economically small: after the reform, households

choose units with approximately 4 euros, or less than 0.2 standard deviations, larger HA changes. Moreover, the average change in recipient choices is driven by changes at the top of the HA change distribution while the bottom quantiles move very little. We interpret this as evidence that only a subset of recipients (instead of the full recipient pool) respond to the changes in incentives to choose different types of units generated by the reform.

We then analyze the changes on the supply side of the rental market, emphasizing that rental supply could adjust through new construction and also through the conversion of existing owner-occupied units into rentals. We find that the construction of units with larger HA increases started growing after the reform relative to the construction of units with smaller HA increases. However, these changes are small and become detectable only some years after the reform took place. This is important for the interpretation of our rent results since the changes in demand that we document (differential choices by recipient households) kick in immediately after the reform. Thus, any demand shift from the reform materializes very quickly, while changes in construction show up in the data some years after the reform. We interpret this as evidence that changes in construction can have at most a limited role in explaining the rent effects. The supply of rental housing could potentially increase also very quickly through the conversion of owner-occupied units to rental units. However, we do not observe changes in the types of owner-occupied units that are converted to rentals. We also verify that the changes in construction and conversion contribute very little to the overall composition of the rental housing stock in the medium-run.

We conclude from this evidence that a large supply response is *not* the main driver of our observed rent effects. Rather, our interpretation is that for the majority of recipients, their housing choices were relatively inelastic with respect to the incentives generated by the reform. If household willingness-to-pay for different types of units is not much affected by HA changes, the incidence of HA increases is largely on the recipients, whether supply is very elastic or not. There are many potential explanations for this finding. An important reason might be that recipients do not assign a large weight to HA when optimizing their housing choices. For example, less than half of the HA spells that started during our analysis period lasted more than a year. Thus, most HA recipients can expect to at some point bear the full rental burden of their unit when moving to a new housing unit. This suggests that other reasons than HA, such as the suitability of the unit, are the most important drivers of housing choices.

Related literature. Prior empirical research is inconclusive about the extent to which rent subsidies are passed on to rents. Early studies from Europe find that more

than 50% of the HA accrued to landlords (Gibbons and Manning, 2006; Fack, 2006; Kangasharju, 2010; Viren, 2013). More recent evidence from Europe, Israel and New Zealand finds more moderate, but still economically significant rent effects (Hyslop and Rea, 2019; Sayag and Zussman, 2020) or very small or even zero rent effects (Brewer et al., 2019; Eerola and Lyytikäinen, 2021). The heterogeneity in the existing pass-through estimates can reflect either differences in economic conditions across different contexts or statistical uncertainty related to the estimates. In particular, the earlier studies use relatively small datasets either from survey sources or samples of register data. While the more recent studies use register data on rents, they lever only relatively small variation in rent subsidies, leading to relatively imprecise estimates.

The US has in place a housing voucher program, which differs from entitlement programs in two important ways. First, the program has a fixed budget, and only a minority of eligible households receive a voucher, presumably mitigating rent effects relative to entitlement programs where all eligible households get the benefit. Second, the voucher recipient must live in an eligible unit, which directly incentivizes households to make changes to their housing consumption and often results in moving after receiving the voucher (Eriksen and Ross, 2013). As is the case with the policy evaluations of the entitlement programs, the findings about the rent effects of different expansions of the voucher program are heterogenous. Susin (2002) and Collinson and Ganong (2018) find that program expansions or rent ceiling increases are associated with higher rents, whereas Eriksen and Ross (2015) does not.

We contribute to the literature in the following ways. Our first contribution stems from the high-quality data and the large-scale reform that we leverage. Together they result in pass-through estimates that are very precise compared to prior literature, and allow us to provide compelling evidence on the validity of our research design. Furthermore, our data allows us to address the problem of observing housing quality. The existing pass-through estimates allow for two alternative (not mutually exclusive) interpretations. The point estimates may reflect either higher quality-adjusted rents and/or better housing quality which is then reflected in higher rents. Often these two effects cannot be reliably disentangled. We can observe the same housing units before and after the reform, which allows us to control for unobservable time-invariant housing quality.

Our second contribution is to use register data on HA recipients and the overall rental market to shed light on the key demand- and supply-side mechanisms through which changes in HA can affect rents. As Collinson and Ganong (2018) note, differ-

ent subsidy programs, depending on the incentives they give rise to, can have very different rent effects. For example, a program that does not change the marginal price of housing consumption should distort housing demand less than a program which does. Yet, causal evidence on the responsiveness of subsidy recipients to the incentives created by different subsidy schemes in terms of housing consumption is almost non-existent. Exceptions are Öst (2014) and Gibbons et al. (2020) who focus on reforms that cut benefits for families occupying relatively large units. Both studies find that the benefit cut induced households to downsize conditional on moving, and Öst (2014) also finds an effect on the decision to move. However, these studies use variation or data that concerns only a subset of recipients.[2] We contribute to the literature by analyzing the effects of a general large-scale subsidy program on recipient housing choices.

In addition to housing choices, we provide evidence on the supply-side reactions in the rental market. Rent subsidies are often suspected to have larger rent effects in locations of lower supply elasticity (e.g. Susin 2002; Eriksen and Ross 2015). We are the first to study the effects of a rent subsidy program on housing supply using register data on the total housing stock, studying separately residential construction and conversions from the owner-occupied sector to the rental sector.

The rest of the paper proceeds as follows. In the next section, we present our institutional context, data and research design. In section 3, we present our results on rent effects. In Section 4, we present and discuss our results on household mobility and rental housing supply. Section 5 concludes.

[2]Gibbons et al. (2020) study rent-controlled social housing sector, where presumably mobility is affected by availability, and Öst (2014) studies only single parents, and does not provide evidence on pre-treatment trends, making it difficult to assess the credibility of the DID design.

2 Institutional context and research design

2.1 Institutional context and data

Approximately a third of Finnish households live in rental housing. The rental housing market consists of a private unregulated segment (70%) and a social housing segment (30%). In our analysis, we focus on new rental contracts in the private rental market. While rent increases of existing rental contracts are typically tied to some publicly available index, such as the official cost-of-living index, new rental contracts in the private rental market are not subject to any constraints on rent setting. New rental contracts can therefore be expected to respond to changes in housing allowances and to provide a credible benchmark for studying the rent effects of the reform.[3]

Housing allowance is an important part of the Finnish social security system with a stated aim of reducing the housing costs of low-income households. We focus on the general HA intended for working-age households. In 2020, total outlays amounted to 1.57 billion euros (0.66% of GDP) and roughly 400,000 households (15% of all households) received the general housing allowance. When the reform was implemented in 2015, roughly 30% of tenants in the private rental market received HA (see Table A1 in Appendix A).

Our main data source is the register of housing allowances from Social Insurance Institute of Finland (Kela) for years 2010–2019. The data cover the universe of monthly recipient-level HA payments and include information on the characteristics of recipient households and their housing units. We do not directly observe new contracts in our data, but we define a contract as new if the recipient received HA in another address at most 4 months ago. The construction of the analysis samples is described in detail in Appendix A.[4]

Table 1 provides summary statistics on our sample focusing on apartments with floor area between 15m^2 and 100m^2. The first column contains all HA recipients and the second the subset of new rental contracts. Most HA recipients are single-member households, their disposable monthly income before HA is around 900 euros and the average rent around 600 euros. As we will discuss below, the HA system covers rent

[3]In the social housing sector rents are regulated and should therefore not be affected by changes in housing allowances or other demand-side factors.

[4]Students were covered by a separate student housing supplement up until 2017 and became eligible for general housing allowance in 2017. Pensioners have a separate, but similar housing allowance system. We exclude all pensioners and students from our estimation sample throughout the whole sample period.

only up to a rent ceiling. For an overwhelming majority of recipients (82% and 85%), this ceiling is binding. For these households, the marginal unit of additional housing consumption is not subsidized.

Table 1: Summary statistics, HA register data.

	All payments mean	New contracts mean
Household size	1.6	1.8
Share single-member households	0.68	0.57
Apartment surface	48.1	51.6
Household income	901.6	931.1
Rent	578.5	621.7
Housing allowance received	305.7	327.3
Share rent ceiling binding	0.82	0.85
Observations	11188052	219204

Notes: HA register 1/2010-12/2019, all monetary values in 2020 euros. Throughout, the sample is restricted to private rental market residents in units of floor area between 15 and 100 m^2. The first column summarizes all month-by-household payments to HA recipients. The second column summarizes the subset of recipients with new rental contracts. The contract is defined as new if the recipient received HA in another address at most 4 months ago. Household income refers to income other than the housing allowance.

Our second data source is the population-wide household register data provided by Statistics Finland (see Appendix A). In addition to rich demographic and socio-economic information on Finnish households, the data include information on the characteristics of their housing units. For each housing unit, we observe size of the unit (m^2), construction year and current tenure status (private market or social housing tenant or owner-occupied). This information allows us to analyze changes in rental supply in relation to the HA reform, and separate between construction of new units and conversions of owner-occupied units into rental units.

2.2 Housing allowance system and the 2015 reform

The legislation governing HA payments was renewed in late 2014, with the new legislation taking effect in 2015. The stated objective of the legislative changes was to simplify the program. The most important change concerned the calculation of the

rent ceiling which sets an upper limit to the housing allowance. Prior to the reform, the rent ceiling depended in a complicated way on a number of housing unit and household characteristics. The reform substantially simplified the determination of the rent ceiling. At the same time, the reform made the HA system more generous on average.

Before the 2015 reform, HA was determined according to the following formula:

$$HA = 0.8[\min(Rent/m^2, MaxRent\_m^2) \cdot \min(FloorArea, Max\_m^2) - d], \quad (1)$$

where $Rent/m^2$ denotes the actual monthly rent per square meter of the unit and $MaxRent\_m^2$ the ceiling on the monthly rent per square meter. The ceiling varied depending on construction year, floor area, and heating system of the building and affordability group of the municipality.[5] This ceiling was binding for a large majority of recipient households before the reform (see Eerola and Lyytikäinen, 2021). $FloorArea$ denotes the actual size of the unit, and Max_m^2 denotes a ceiling on the size of the unit. This ceiling varied by household size.[6] Finally, d denotes a deductible which was increasing in household income and governed the phase-out of the allowance. The HA covered up to 80% of the rent of the unit.

The HA reform replaced the ceiling on rent per square meter and the ceiling on unit size with a single ceiling on total rent. Since January 2015, the HA is determined as

$$HA = 0.8[\min(Rent, MaxRent) - d], \quad (2)$$

where $MaxRent$ denotes a ceiling on the total monthly rent, $Rent$. $MaxRent$ depends on household size and and the affordability group of the municipality, but not on other housing unit characteristics.

The reform treated apartments of different sizes and rents differently depending on whether the ceiling on rent per square meter, the ceiling on floor area, or both were binding before the reform. As an example, the left panel of Figure 1 illustrates the maximum possible HA before and after the reform for a low-income (zero deductible) single-member household in Helsinki. Before the reform, maximum possible

[5]Municipalities are divided into four affordability groups depending on the local rent level. The city of Helsinki constitutes one affordability group with the highest rent ceiling. Other groups in descending order of the rent ceiling are the rest of the Helsinki Metropolitan Area (3 municipalities), roughly 30 large and mid-sized cities, and finally, all other municipalities consisting of small towns and rural municipalities.

[6]For example, the ceiling on floor area was 37 m^2 for singles, 57 m^2 for two-person households and 77 m^2 for three-person households.

HA increases with unit size up to a limit after which maximum HA is constant (for units smaller than 37 m^2, only the ceiling on rent per square meter is binding, but for units larger than 37 m^2, also the ceiling on floor area is binding). After the reform, the maximum HA is independent of unit size. Thus, the maximum possible HA increased significantly in small units, but only a little in larger units. On the other hand, allowances could also increase in larger units (even if the maximum possible allowance did not increase), if the rent-per-square-meter limit was not binding prior to the reform.

To illustrate how the reform affected actual HA amounts, we calculate predicted changes in allowances, implied by the differences between equations (1) and (2). To do so, we take our main estimation sample in the years prior to the reform (2010-2014), and for each unique observation of a new rental contract (a unit-household combination), we compute the HA that the household would have received using 2014 policy parameters and using 2015 policy parameters. We call the difference between these two (hypothetical) allowance levels the *predicted HA change*. These predicted changes would accurately describe the changes in allowances paid to households if there were no other changes taking place other than the change in the policy. We use these predicted HA changes as a continuous treatment variable in a DID framework, as we explain in more detail later.

The right panel of Figure 1 shows the resulting change in HA, averaged over the floor area groups. As the figure shows, there is a systematic pattern in the variation generated by the reform that follows the change in the maximum HA of the left panel of the figure. The predicted HA change was large in small units, close to zero in mid-sized units and positive again for larger units.

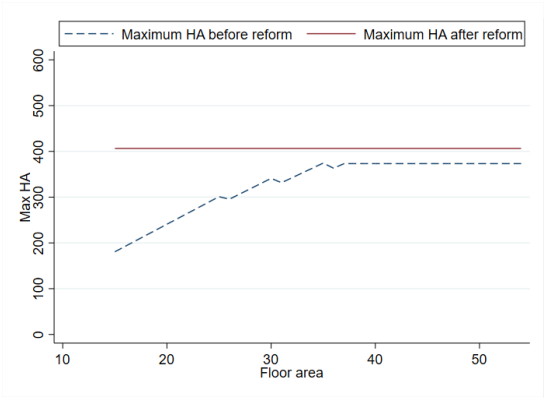

(a) Maximum HA for single-member households in Helsinki

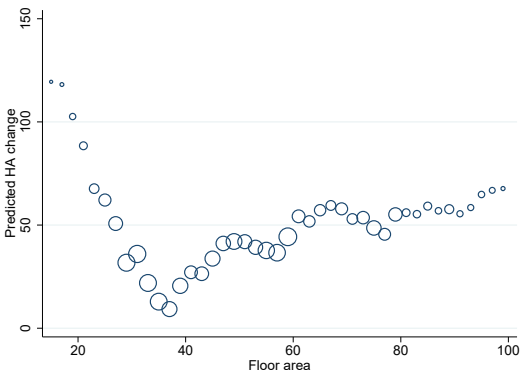

(b) Predicted HA change

Figure 1: Maximum HA before and after the reform in Helsinki (left) and average predicted HA change in HA register data (right), both by floor area.

Notes: The left-hand graph illustrates the maximum HA before and after the reform in Helsinki for a single-member household with a zero deductible renting a housing unit built before 1986. The right-hand graph illustrates average predicted HA changes by unit floor area based on HA register data. The size of the circle is proportional to the number of units in each bin.

The reform also simplified the formula for calculating the deductible, but this had little practical relevance for actual HA amounts. Furthermore, in September 2015, an earnings deduction of 300 euros/month was introduced. This made the HA more generous to existing recipients with labor earnings and also enlarged the pool of eligible households. We do not exploit this variation in our analysis.

2.3 Research design

Graphical analysis. The reform changed the HA in different ways for different types of rental units. This variation is the starting point for our empirical analysis. We start our analysis by dividing housing units into discrete floor area groups based on Figure 1b, and describing the evolution of HAs and rents in these groups over time. This graphical analysis allows us to transparently examine and assess the magnitude of the changes in the HA caused by the reform in these groups and possible coinciding changes in rents. We show separately the pool of all contracts and the subset of new rental contracts.

Continuous treatment DID. In our econometric analysis, we use a continuous treatment DID strategy, which exploits the full variation in HA changes induced by

the reform. For this analysis, we use a repeat observation sub-sample where housing units are observed with new tenants at least once before and once after the reform. For each unit, we compute a continuous *treatment exposure*: We use the pre-reform unit and tenant characteristics and compute predicted HA changes implied by the changes in the allowance formula, as described in equations (1) and (2). In other words, the treatment exposure of apartment j measures the change in HA that would have occurred due to the reform, if there had not been any changes in the characteristics of the unit, the tenant or the rental contract.

We use the treatment exposure in event study regressions as well as DID and DID-IV regressions. We start with the following event-study style regressions:

$$y_{it} = \sum_{\substack{s=2010q1 \\ s\neq 2014q4}}^{2019q4} \theta_s \, \Delta pred\_HA_j \; + \gamma_t + \omega_j + u_{it}, \tag{3}$$

where the outcome y_{it} is either HA or rent, i indexes the rental contract, t time and j the housing unit, and $\Delta pred\_HA_j$ is the predicted HA change of the unit. The regression includes fixed effects for the time period (γ_t) and for the housing unit (ω_j).

The corresponding continuous treatment DID specification is

$$y_{it} = \beta \times \Delta pred\_HA_j \times post_t + \gamma_t + \omega_j + u_{it}, \tag{4}$$

where the outcome y is either HA or rent, i indexes the rental contract, t time and j the housing unit. The regression includes fixed effects for the time period (γ_t) and for the housing unit (ω_j). $\Delta pred\_HA_j$ is the predicted HA change of the unit. This regression compares the change in the HA or rent of a given housing unit from pre-treatment to post-treatment time period as a function of the predicted change in HA induced by the reform. Since all units are treated at the same time, there are no issues with a staggered treatment and we estimate the regression using two-way fixed effects.

Following Callaway et al. (2024), we can interpret the coefficient of interest in the continuous treatment DID regression as an average causal response to treatment on the treated under a strong parallel trends assumption, which demands that low-exposure units provide a good counterfactual for what would have happened to outcomes in high-exposure units had they received a small exposure.[7] The assumption

[7] If the strong parallel trends assumption holds, then β estimated using a two-way fixed effects regression is a weighted average of the average causal response on the treated of a given exposure for units who get that exposure, with all weights positive (although the weights do not correspond to the population distribution of the exposure).

would fail if, for example, small units received on average larger treatment exposure and the effect of HA on rents was different for small units and for other units. We thus have to assume some degree of homogeneity in the effects of HA on rents across units with different exposures.

We also interpret our DID estimates through a DID-IV regression, which relates the size of the rent increase to the size of the HA increase by estimating the effect of a one euro change in HA on rents. The DID-IV estimates are informative about the incidence of changes in HA between tenants and their landlords and also facilitates comparison with previous studies. Here, our regression of interest writes

$$Rent_{it} = \beta HA_{it} + \gamma_t + \omega_j + \epsilon_{it}, \tag{5}$$

where the outcome variable is monthly *Rent* in rental contract i in time t and the parameter of interest is β. Similar to the event study specification, we include fixed effects for the time period (γ_t) and housing unit (ω_j). The endogeneity concerns in equation (5) are addressed by instrumenting HA with the treatment exposure interacted with a post-reform indicator ($\Delta pred\_HA_j \times post_t$). The coefficient β in this regression will simply be the DID-estimate for rents divided by the DID-estimate for allowances (similar to how the typical IV estimator amounts to scaling the reduced-form parameter by the first-stage parameter).

Regarding heterogeneous treatment effects in the context of a DID-IV specification, as summarized by De Chaisemartin (2010), we can interpret our estimates for β as local average treatment effects even if the conventional IV assumption of instrument exogeneity is not satisfied as long as the instrument is uncorrelated with potential outcomes, accompanied with two parallel trends assumptions: one related to the first-stage and another related to the second-stage outcome. These are the same parallel trends assumptions we make in our DID estimation. We inspect the credibility of these parallel trends assumptions by comparing the pre-reform trends in units receiving varying treatment intensity as a part of our event study analysis.

Finally, both the DID and DID-IV specifications rely on the stable unit treatment value assumption (SUTVA), which requires that the increase of HA in some housing units does not affect rents in other housing units. This assumption would be violated if, for example, housing allowances had increased for small units enough to have decreased the demand for medium-sized units, which is possible at least in principle. However, the failure of the SUTVA assumption in our case would most likely lead to an upward bias in our estimate on the effects of HA on rents (if there was an important shift in demand away from the units with small predicted HA changes, we

would expect a rent decrease in this group). Given that we find point estimates which are close to zero, the concern of an upward bias is not too worrying in our context.

3 Rent effects

3.1 Main results

We first report a set of descriptive graphs in which we divide housing units into groups by their floor area. Panel (a) of Figure 2 shows the evolution of mean monthly HA payments to recipient households (left) and their rents (right) in the different floor area groups. As expected based on Figure 1b, housing allowances in units of 15–25m^2 (highlighted in red) increased dramatically in 2015, while increases in other floor area groups were moderate. Especially in medium-sized units (35–45m^2, highlighted in green), HA deviates in 2015 only slightly from its pre-reform trend. The development of mean rents in the different floor area groups in turn is stable around the reform period, and there are no visible differences between the groups.

Panel (b) in Figure 2 shows the evolution of monthly HA payments and mean rents in new rental contracts only. Rents in existing contracts can be rigid, but in new rental contracts there are no legal or other reasons not to expect immediate rent effects if HA increases are indeed passed through to rents. The bottom panels are very similar to the top panels of Figure 2, and indicate no changes in the relative rents between groups. For the remainder of our analysis, we focus exclusively on new rental contracts.

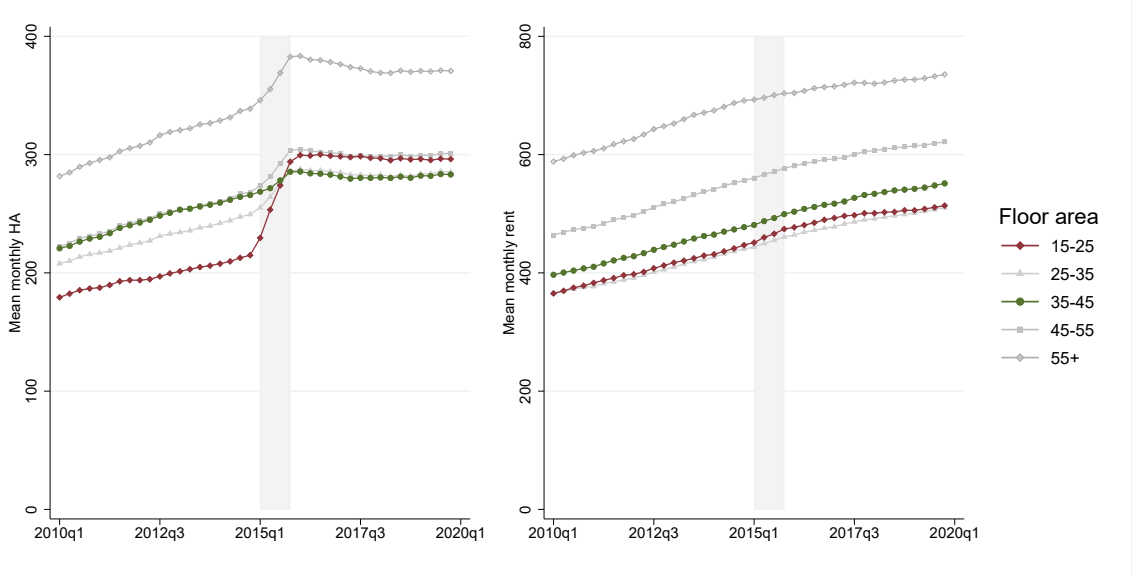

(a) All payments.

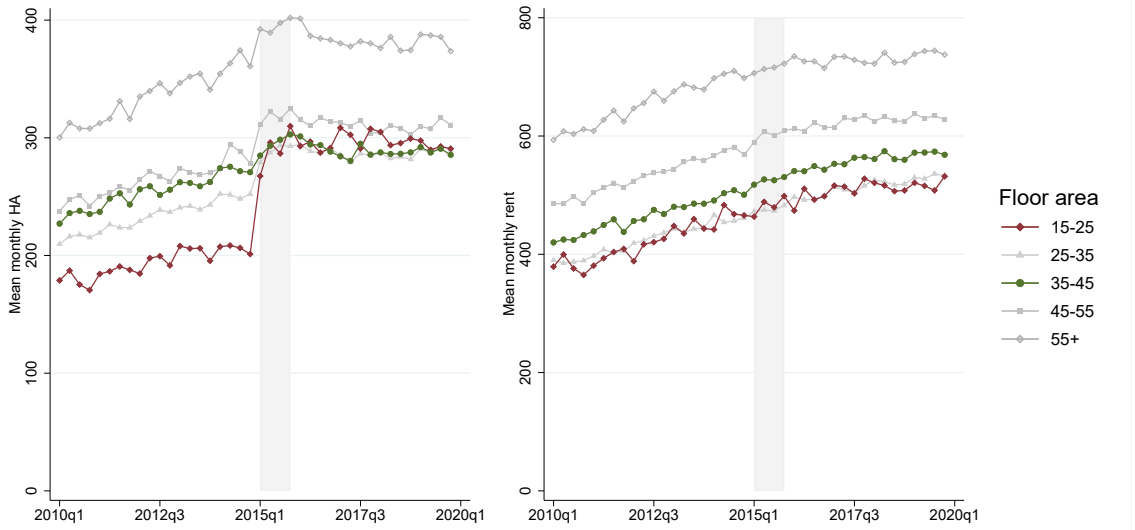

(b) New contracts

Figure 2: Mean HAs and rents by floor area group.

Notes: The figure shows mean monthly HA paid to recipient households and mean monthly rents paid by recipients in our estimation sample, at quarterly level. The light gray shaded area refers to year 2015. For existing HA spells, the reform was rolled out during 2015, but for new rental contracts, it became effective immediately in the beginning of 2015. "All payments" refers to all monthly payments to recipients, and "new contracts" refers to the first payment made to a recipient who has changed addresses. For details on sample selection and identifying new rental contracts, see Appendix A.

Next, we turn to the continuous-treatment DID analysis using our repeat observation sub-sample with only housing units that we observe at least once before and after the treatment. This analysis compares changes in rents across housing units which received different-sized treatment exposures, where the treatment exposure is defined as a predicted change in HA assuming that only the parameters of the HA system changed, but other characteristics of the unit and the recipient remained the same (see Section 2.3 for details). First, we report event-study-type evidence in Figure 3. The left panel shows that a one-euro increase in the predicted HA change (treatment exposure) led to approximately one-euro increase in the actual HA payment. This means that we get significant exogenous variation in HAs, although we cannot perfectly predict the new HA received by the new tenant (since the tenant characteristics also changed, which we do not control for, as did possibly the rent). The right panel shows that prior to the reform rents developed similarly in units that received different-sized treatment exposures lending support to our parallel trends assumption. The right panel also shows that increases in treatment exposure do not lead to increases in rents paid by HA recipients.

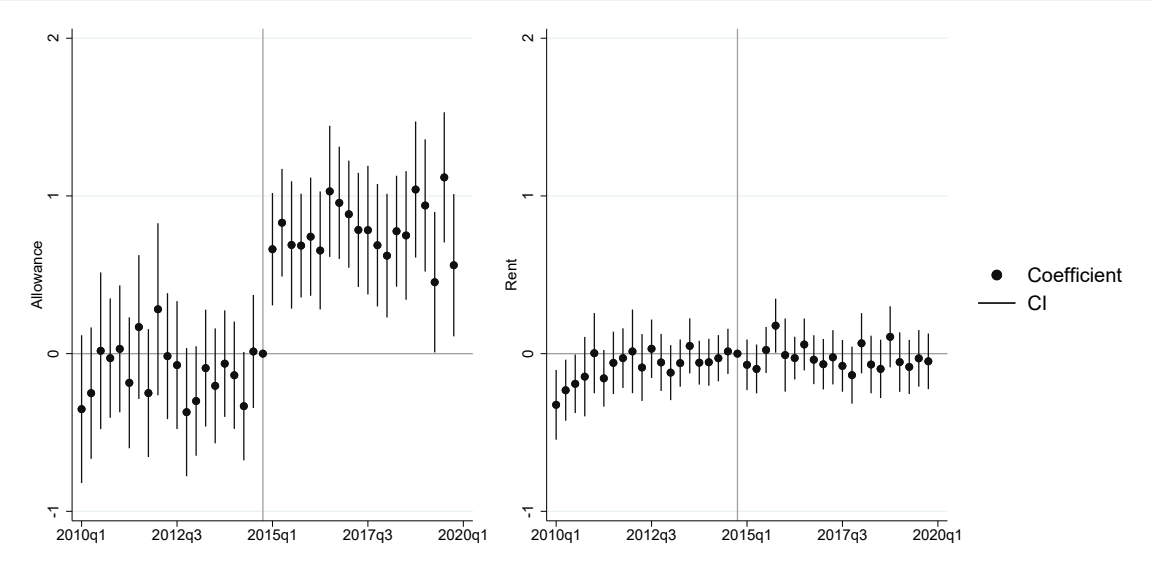

Figure 3: Comparison of monthly HAs and rents of units that received different-sized treatment exposures, with housing unit fixed effects.

Notes: The figure plots coefficients from an event study regression in our fixed effects sample, where the outcome (HA or rent) is regressed on quarter fixed effects, housing unit fixed effects and treatment exposure × quarter fixed effects, omitting the last quarter before the reform. Dots and whiskers illustrate the point estimate and the 95% confidence intervals of the treatment × quarter coefficients. Standard errors are clustered at the housing unit level. N=22,346.

The corresponding continuous treatment DID regression results are reported in Table 2. First, in column 1, we regress actual HA on our measure of predicted HA change. The point estimate suggests that a one euro increase in the predicted HA change is associated with an increase in the actual HA of 0.90 euros with a standard error below 0.04. Thus, our predicted HA change is highly correlated with changes in actual HA. Column 2 reports estimates from a similar regression where the outcome is the rent. The point estimate suggests that the average rent increase following a one euro increase in the predicted HA was very modest, roughly 2.5 cents. For completeness, in column 3 of Table 2, we present the DID-IV results in which HA is instrumented for by our measure of predicted HA change. The estimated effect of HAs on rents, roughly 3 cents per an additional euro of HA, is small and not statistically significantly different from zero. The standard error is 0.02 and implies that we can rule out even moderate rents effects with a high degree of confidence.

Table 2: Continuous-treatment DID-IV estimates, with housing unit fixed effects.

	DID		IV
	(1)	(2)	(3)
	Allowance	Rent	Rent
Predicted HA change	0.899	0.0243	
	(0.0339)	(0.0161)	
Allowance			0.0270
			(0.0177)
Month × year FEs	✓	✓	✓
Unit FEs	✓	✓	✓
Outcome mean	314.1	577.3	577.3
N	22346	22346	22346
SE clustered by	Unit	Unit	Unit
First-stage F			705.1

Notes: The table reports results from DID and IV regressions where the treatment or instrument is defined to be a predicted change in HAs as described in Section 2.3. Columns 1 and 2 report coefficients from a regression of the outcome on our measure of predicted HA change × post indicator. Column 3 reports the second-stage of an IV regression, where HAs are instrumented for by a predicted HA change × post indicator. The first stage of this regression corresponds to column 1. All specifications contain month-by-year and housing unit fixed effects. Standard errors are clustered at the housing unit level.

3.2 Robustness

Taken together, our results so far imply that increases in HA due to the 2015 reform did not increase rents. Next, we provide additional analyses with alternative treatment specifications and robustness checks.

Household characteristics. Since our main analysis holds constant the pre-reform characteristics of units and tenants, a factor that could complicate the interpretation of our results is the potential effect of the reform on household sorting into different types of housing units. We study this potential issue in Appendix B.1. We show that there were no large changes in household characteristics across different types of units

after the reform. This is consistent with the high predictive power of the first-stage regression in Table 2.

Rent results with alternative treatment definitions. We analyze the robustness of our main results to alternative treatment definitions in Appendix B.2.

We begin by testing whether our results hold when we *do not use a parametric measure treatment exposure* in the regression. We do this simply by comparing the development of rents and allowances in different floor area groups. We use a two-group specification where the 15–25m^2 units act as the treatment group and the 35–45m^2 units as the control group. This specification is based on Figure 2 showing that the average increase in HA is much larger in housing units with floor area 15–25m^2 than in housing units with floor area 35–45m^2. This specification yields very similar results to those obtained by using the predicted HA change. Event study results in Figure B3 show that rents in the treatment and control groups developed similarly prior to the reform and that the rents did not increase in the treatment group relative to the control group after the reform despite large differences in HA increases.

Second, we consider alternative ways to compute the continuous treatment exposure. The treatment exposure definition in our main results reported in Figure 3 and Table 2 holds constant pre-reform observed characteristics of the unit and the tenant. As described above, a potential concern related to this is instrument weakness in the case where tenant characteristics in given types of units change after the reform. We are not particularly concerned by this given the strong predictive power of the first-stage regression. Also, as discussed above, household characteristics do not change in different floor area groups after the reform. Nonetheless, we address this concern by considering two alternative methods for computing the predicted HA changes.

We first approximate the predicted changes in HA by the average predicted changes for similar households in similar housing units. For all new contracts before the reform (2010–2014), we compute the predicted HA change given the pre-reform characteristics of each unit and household living in the unit. We then average this measure by household and unit characteristics by splitting the data into groups by floor area (2m^2 brackets), city size (3 groups[8]), and household size (1, 2 or 3 members[9]). We compute

[8]These are Helsinki Metropolitan Area (Helsinki, Espoo and Vantaa), six other largest cities (cities which have a population of at least 100 000, excluding Helsinki MA), and the rest of Finland.

[9]We exclude households with 4 or more members, because there are relatively few such households in our main estimation sample and because these households likely consider also units with floor areas larger than 100m^2, thus excluded from our mobility analysis.

the average predicted HA change in these cells and label this variable as $\overline{\Delta pred\_HA}_1$. Using this alternative exposure definition, we can let the exposure depend on *actual* household characteristics at the time of observation instead of holding constant the pre-reform characteristics. We will use this treatment definition also in subsequent analysis when we analyze the rent effects by city size and when we explore households' moving behavior in Section 4.2. Results in Figure B5 show that the effects of the reform on HAs and rents are very similar to our main results. Again, we observe a sharp increase in actual HA payments as a response to our alternative definition of predicted HA change, and we observe no differential changes in rents of units with more or less predicted HA change.

Moreover, to make the instrument completely blind to household characteristics, we also compute the average predicted HA changes without household characteristics. To measure changes in HA as a function of the characteristics of housing units only, we proceed by averaging predicted HA changes similarly as before, but instead of averaging at the household level, we calculate averages at the same floor area ($2m^2$ brackets) and city size (3 groups) cells as before. In other words, each unit is mapped to belong to a cell defined by the floor area and municipality group of the unit, and the treatment exposure is the average predicted HA change in the corresponding cell.[10] We label this variable $\overline{\Delta pred\_HA}_2$ and we will use this variable later on when analysing rental housing supply in Section 4.3. Rent effects using this alternative treatment definition are reported in Figure B6 and are again very similar to our main results.

With these alternative treatment definitions the estimation sample of new rental contracts consists of more than 200,000 observations. Therefore, they allow us to use a substantially larger sample than in the specification with housing unit fixed effects. Although we are not able to control for unit fixed effects in these specifications, our main results carry through, and none of the specifications indicate growth of rents in units with higher predicted HA changes.

Rent effects by city size. We also analyze changes in HAs and rents separately in different-sized cities, ranking city sizes from very large to small. For this analysis, we use the alternative treatment definition ($\overline{\Delta pred\_HA}_1$) allowing for a sufficiently large estimation sample. Figures B7 and B8 report our descriptive and event-study graphs separately for the above defined three groups: Helsinki MA, other largest cities, and the rest of Finland.

[10] Again, we exclude households with 4 or more members.

Across the board, in these location-specific regressions where time effects are estimated separately for each municipality group, the estimates for the rent effects of HA are negligible. This is true even for the Helsinki metropolitan area where HA increases were clearly the largest and housing supply is presumably most inelastic (Oikarinen et al. 2015).

Social assistance. Housing costs of low-income households are also covered through social assistance which is the last-resort form of financial assistance in the Finnish social security system. For social assistance recipients, a change in HA may be partly offset by changes in social assistance. It is therefore possible that the rent effects of the HA reform would be more pronounced for households who are not eligible for social assistance. We address this issue in Appendix B.4.

First of all, we ask to what extent were changes in HA offset by changes in social assistance. Unfortunately, our HA register data do not include information on social assistance. Therefore, we use total population register data from Statistics Finland to compare annual HA and social assistance payments across different floor area groups. Although there is evidence of social assistance responding to the HA reform, there remains substantial variation in the total subsidies by floor area group also after accounting for these changes. The *sum* of the two subsidies increased by 271 euros more per year in 15–25 m^2 units than 35–45 m^2 units from before to after the reform. This number includes non-recipient tenants as zeros, whereas recipients correspond to only approximately 28% of the private rental market. This suggests that despite changes in social assistance, there were large and clear changes in subsidy payments across different floor area groups.

Next, we check whether our estimated rent effects depend on social assistance eligibility. Since we do not observe social assistance recipient status in the HA register, we do this by dividing the sample into two groups based on income, as households with sufficiently high incomes are very unlikely to be eligible for social assistance. Figure B9 shows the evolution of mean monthly HA payments and mean rents in different floor area groups and Figure B10 event-study graphs separately for the two income groups. In order to increase sample size, in the event-study design we again use the alternative treatment specification $\overline{\Delta pred\_HA}_1$. The rent effects for the two groups are very similar to one another (and overall similar to our main results). Therefore, it does not seem likely that our main results are mitigated by the countervailing incentives generated by the social assistance program.

4 Why do we observe small rent responses?

The observed effects of HA on rents are an equilibrium outcome that stems from rental housing demand and supply. Our estimates of the rent effects of the HA changes are small in economic magnitudes and also small relative to some earlier findings in the literature. To understand why we observe such small rent effects, we next zoom in on the underlying changes in rental demand and supply that can help explain the observed results.

We start by setting up a conceptual framework to illustrate the different mechanisms that we are interested in capturing. Having discussed the conceptual framework, we analyze empirically the changes along these different margins.

4.1 Conceptual framework

This section presents a stylized conceptual framework to illustrate the forces that govern the extent to which housing allowances affect rents. The purpose is to clarify the different margins of adjustment and to discuss their quantitative importance for the observed rent effects. For derivations, see Appendix C.

Consider a competitive rental market where $D(r)$ is aggregate rental demand and $S(r)$ aggregate rental supply as a function of the quality-adjusted rent r.[11] In an initial equilibrium, the market is cleared by r, so that $S(r) = D(r)$. The HA creates a wedge between rent paid by recipients and rent received by landlords, shifting the equilibrium quantity. The pass-through of HA changes to rents can be low if either the supply elasticity is very high or the demand elasticity is very low.

The demand for rental housing aggregates over demand by recipients (R) and non-recipients (N):

$$D(r) = D^R(r) + D^N(r).$$

The supply of rental housing aggregates over construction of new units (S^C), conversion of existing owner-occupied units to rental units (S^I), and stock of preexisting rental units net of depreciation $S^O(1-\delta)$:

$$S(p) = S^C(p) + S^I(p) + S^O(1-\delta).$$

Considering an increase in HA, ds, equating the change in quantity demanded

[11] Competitive rental market implies in particular that landlords cannot charge different prices from recipients and non-recipients. Therefore, any rent effects faced by HA recipients also apply to non-recipients.

with the change in the quantity supplied yields

$$\underbrace{D^{R'}(r)\cdot[dr-ds]}_{\substack{\text{Change in quantity}\\ \text{demanded by recipients}}}+\underbrace{D^{N'}(r)\cdot dr}_{\substack{\text{Change in quantity}\\ \text{demanded by non-recipients}}}=\underbrace{S^{C'}(r)\cdot dr+S^{I'}(r)\cdot dr.}_{\text{Change in quantity supplied}}$$

The rent effect in terms of the various elasticities can be expressed as

$$\frac{dr}{ds}=-\frac{\frac{D^R}{q}\epsilon_D^R}{[\frac{S^C}{q}\epsilon_S^C+\frac{S^I}{q}\epsilon_S^I]-[\frac{D^R}{q}\epsilon_D^R+\frac{D^N}{q}\epsilon_D^N]} \tag{6}$$

where q denotes the initial equilibrium quantity. Equation (6) highlights the relative importance of the different margins. First, in a given period of time, only a small fraction of the units supplied for rent are newly built or newly converted to rentals. This mitigates the importance of changes in new rental supply for the rent effects of the subsidy. Even if new rental supply was very elastic, either due to construction (high ϵ_S^C) or because of conversions (high ϵ_S^I), their importance for the overall stock would be dampened by the fact that $\frac{S^I}{q}$ and $\frac{S^C}{q}$ are small and a large share of rental supply is inherited from the past.

This observation also highlights the difference between the effective short-run and long-run supply elasticities. In the very short run, overall housing supply is almost perfectly inelastic. In the long run, the aggregate supply elasticity depends mainly on construction. Moreover, while new construction has a limited role for supply in the short run, the conversion of units from owner-occupied units to rentals by investor-landlords ($\epsilon_S^I > 0$) could potentially be more important also in the relatively short run.

Second, as we consider separately HA recipients and non-recipients, we see that the rent effect are strictly bounded from above, even in the case where the supply is completely inelastic, as long as non-recipients do not have a perfectly inelastic demand. With completely inelastic supply, we have

$$\frac{dr}{ds}=\frac{\epsilon_D^R}{\epsilon_D^R+\frac{D^N}{D^R}\epsilon_D^N}$$

Thus, for example, in the case where the price elasticities of demand of recipients and non-recipients are approximately similar, the rent effects are bounded from above at $\frac{dr}{ds}\approx-\frac{D^R}{q}$, even in the extreme case where supply is fully inelastic. In other words, if the recipients' share of the rental market is, for example, 30%, then a one euro increase in the subsidy paid to recipients *cannot* increase rents by more than 30 cents if the rental market is competitive.

4.2 Housing choices of HA recipients

To understand whether the reform affected recipients' demand for different types of units, we next analyze their housing choices. The reform increased recipient households' financial incentives to occupy units that received larger treatment doses, since allowances in those units increased, but rents did not. We analyze recipient households' responses along two dimensions: household mobility (the propensity to move out of units with smaller or larger exposure to treatment), and housing choices conditional on moving (whether households who move choose units with larger treatment exposure). Estimating demand and supply elasticities is beyond the scope of this paper.

Defining treatment exposure. To be able to assess whether households' choices were affected by the reform, we need to be able measure changes in HA induced by the reform also for housing units that we do not necessarily observe both before and after the reform. To do so, we approximate predicted changes in allowances by the average predicted changes calculated within cells (c) of unit, household and city size. This is the variable $\overline{\Delta pred\_HA}_1$ as explained Section in 3.2.

Moving out: mobility rates. To measure whether households are more likely to move out of units where the predicted allowance increase was smaller, or less likely to move out of units where the predicted allowance increase was larger, we run the following event-study regression in the sample of all payments:

$$y_{it} = \theta \, \overline{\Delta pred\_HA}_{1c} + \sum_{\substack{s=2010q1 \\ s \neq 2014q4}}^{2019q4} \theta_s \, \overline{\Delta pred\_HA}_{1c} + \delta_q + \gamma_m + \epsilon_{it}, \tag{7}$$

where the outcome is $y_{it} = 1$ if during the quarter, recipient household i moves to a new unit, and 0 otherwise.[12] $\overline{\Delta pred\_HA}_{1c}$ refers to the average treatment exposure of the *current* unit (before the move). The regression includes quarter fixed effects (δ_q) and municipality fixed effects (γ_m) for the current location. If, after the reform, households are less likely to move out of units with larger treatment exposure (or

[12]We aggregate the data at quarterly level so the outcome indicator variable takes value 1 if household moves in any month of the quarter. We define moving households as households who will receive HA in a different address in their next payment, at most 4 months later. We look only at moves where the unit before the move is on the private rental market, and the household size is at most 3.

more likely to move out of units with smaller treatment exposure), we would expect the coefficients θ_s to be negative after the reform.

Moving in: characteristics of new units, conditional on moving. Next, we consider if, conditional on moving, households choose units that receive larger treatment exposure after the reform. To measure whether household choices shift towards apartments where HA increased on average, we estimate the following regression in the sample of households who *do* move:

$$\overline{\Delta pred\_HA}_{1c} = \sum_{\substack{s=2010q1 \\ s\neq 2014q4}}^{2019q4} \theta_s + \omega_m + \epsilon_i. \tag{8}$$

Now, the *outcome of the regression* is the average predicted change ($\overline{\Delta pred\_HA}_{1c}$) in the new unit of the household.[13] We explain this outcome variable simply by the time fixed effects describing the timing of the move, omitting the last period before the reform. The analysis includes fixed effects at the municipality level by the arrival municipality (ω_m). Thus, the regression aims at showing if, conditional on moving and given the choice of municipality, households choose different types of units after the reform than before. If, after the reform, households start choosing units with larger treatment exposure, we would expect to see positive post-reform time effects in this regression. Note that this regression should be interpreted as a simple "differences" estimation (changes in the characteristics of units chosen by households over time) as opposed to a differences-in-differences estimation (differential changes over time across more and less treated units).

Results. Figure 4 describes the results from the two regressions analyzing housing choices. The left panel illustrates the propensity to move out of units with different-sized treatment exposure, as summarized in equation (7). There are no differences in moving out-patterns before and after the reform between units with different treatment exposure. If anything, in the first year after the reform households seem slightly more likely than before to move out of units with larger treatment exposure. In other words, the evidence does not suggest that households would leave units with smaller exposure more frequently or leave units with larger exposure less frequently after the reform than before the reform.

[13] We only look at moves where household's new unit after the move is in private rental market and the household size is at most 3.

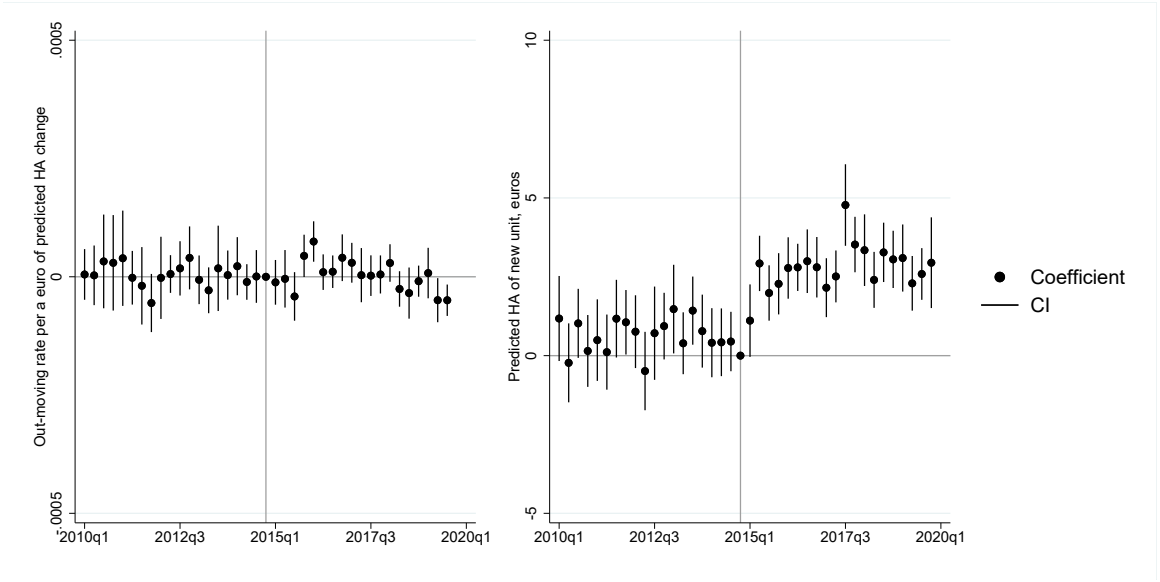

Figure 4: Household choices.

Notes: The left-hand graph depicts the propensity to move out of units with different treatment doses by plotting the event study estimates from equation (7). The baseline mean quarterly mobility rate (the share of observations where a recipient household changes address from one quartile to the next) is 3.46% in the pre-reform period. The right-hand graph depicts, conditional on moving, the size of the treatment dose of the new unit of the household by plotting the event study estimates from equation (8). In both panels, we cluster the standard errors at the municipality level. N = 3,713,762 in the left figure and N = 201,293 in the right figure.

Even if the changes in HA do not affect the decision to move, they might influence the choice of the new unit for those who do move. The right panel of Figure 4 describes the types of units that households choose *conditional on moving*, as measured by the time effects of equation (8). The graph reveals that there is indeed a small but statistically significant response by recipient households to the reform. The point estimates from 2010 until 2014 are all similar to the end-of-2014 level (reference category). Thus, in the pre-reform period, there are no systematic changes in the types of units chosen by households. However, immediately after the reform, there is an increase in the point estimates. This indicates that after the reform, households who do move start choosing units with larger average treatment exposure (that is, larger predicted HA change). The effect size, however, remains small, stabilizing at approximately 3-4 euros per unit. The effect stabilizes almost immediately after the reform, suggesting that there was no transition during which households would, for example, learn about the reform.

Overall, we interpret this as evidence that even if the reform did not change moving

frequencies at the extensive margin, the reform did induce a small change in the types of units chosen by households *conditional on moving*. However, the effect size, although statistically significant, is relatively small (the pre-reform standard deviation of $\overline{\Delta pred\_HA}_{1c}$ in the sample is 29, suggesting that households start choosing units with less than 0.2 standard deviations higher predicted HA changes). Moreover, we find that the change is not uniform throughout the distribution: the observed average change ($\approx$ 3-4 euros) is driven mainly by changes at the top of the distribution. When comparing moves before and after the reform, the percentiles at the bottom of the distribution move very little compared to the top of the distribution (the 5^{th} percentile does not change, the median increases by approximately 2 euros, and the 95^{th} percentile increases by more than 12 euros).

These findings indicate that only few recipient households take into account HA changes when choosing their units. One potential reason that could explain why many recipients do not take HA changes into account is the expected duration of allowance spells relative to expected tenure spells in the new housing unit. For example, less than half of the HA spells that started during our analysis period lasted more than a year. Thus, when moving to a new housing unit most HA recipients can expect to at some point bear the full rental burden of their unit. Another explanation for the results relates to the availability of different types of units in the private rental market. More generally, our results suggest that the HA paid to a given housing unit is not a major factor affecting the choice of the bundle of housing and non-housing consumption of the recipients households.

4.3 Developer and landlord choices

Next, we describe changes in the supply side of the rental market.

Defining treatment exposure. To assess whether more units with larger treatment exposure were supplied on the rental market after the reform, we need to measure changes in allowances induced by the reform for units that may not have existed or may not have been available to rent before the reform. Moreover, to be conservative, we assume that landlords and developers cannot anticipate what type of a household will move into the unit. Therefore, to measure changes in allowances as a function of the characteristics of apartments, we proceed by averaging treatment exposure in cells defined by floor area ($2m^2$ brackets) and municipality group (3 municipality groups), making the exposure blind to household size. This is the variable

$\overline{\Delta pred\_HA_2}$ as explained in Section 3.2.

Construction. To measure changes in the composition of construction of new private-market rental units, we use the population-wide data on occupied housing units from Statistics Finland, as described in Section 2.1. The data is annual and each annual observation summarizes the end-of-year situation (as opposed to our monthly HA register data). We define new rental units as units which are built in year t and where someone is living in the unit at the end of year t as a private-market tenant. Our estimation equation is similar to equation (8), since we are interested in an intensive margin response to the reform. Using the sample of newly constructed rental units we run the following regression:

$$\overline{\Delta pred\_HA}_2 = \sum_{\substack{s=2010 \\ s \neq 2014}}^{2019} \theta_s + \omega_m + \epsilon_{it}, \tag{9}$$

where the outcome ($\overline{\Delta pred\_HA}_2$) is the average treatment exposure in the floor area - city size - cell (c) of the rented unit. The explanatory variable of interest is the time effects, and we control for municipality fixed effects. If the time effects in this regression were positive after the reform, we would interpret that as evidence of new rental construction shifting towards units with higher average treatment doses.

Conversion of units to rentals. To see whether there have been changes in the types of units that are taken away from the owner-occupied market and offered to rent in the private rental market, we again run a specification which is identical to equation (9). This time we use the pool of units where the unit was occupied by the owner in $t-1$ and is occupied by a private-market tenant in year t. If investors start converting units with larger treatment exposure into rentals, we would expect to see positive time effects in this regression after the reform.

Composition of the rental stock. Finally, we examine how the overall stock of private rental units has evolved. To measure changes over time in the types of units occupied by private-market tenants, we take as given the aggregate number of units rented and assess whether the composition of units has shifted towards units with larger treatment exposure. We run a regression which is identical to equation (9), but this time using the overall private rental housing stock. Again, if the time effects in this regression were positive after the reform, we would interpret that as evidence of the overall rental stock shifting towards units with larger average treatment exposure.

Note that if the rental stock grows overall, but the distribution of units with different treatment exposure does not change, then we would not expect to see much of a change in the time effects of this estimation equation after the reform.

Results. Starting with construction, Panel (a) of Figure 5 shows that in the pre-reform period, there were no systematic changes in the types of units constructed with respect to average treatment exposure. After the reform, there is a change in the composition toward units with larger treatment exposure, likely reflecting a relative increase in the construction of small units. However, this change is statistically significant only five years after the reform and the size of the estimate is small at roughly 2.5 euros per unit.

Panel (b) of Figure 5 shows the results on the composition of owner-occupied units converted into private rental units. In this case, there is a small compositional shift toward larger treatment exposure units two years prior to the reform, but the composition remains largely constant during and after the reform. Given the very small point estimates, we conclude that throughout the time period there have been no economically significant changes in the types of units that are converted into rentals.

Even if we do observe a small shift in the types of units that are built over time, as illustrated in panel (a) of Figure 5, this does not seem to translate into significant changes in the overall rental stock, since only a small share of rental units are recently built. Panel (c) of Figure 5 shows the evolution of the total rental stock. This aggregates over all possible margins through which the rental stock can change, including also conversions of units from rentals to owner-occupied units and units becoming unoccupied (depreciation). Prior to the reform, the total stock has not been changing toward units with larger treatment exposure. After the reform, there is a small statistically significant change toward units with smaller treatment exposure, but this change is tiny in magnitude. Given that the point estimates are very close to zero throughout, we interpret that the rental stock has not been shifting towards units with larger treatment exposure.

For completeness, Panel (d) of Figure 5 shows how the full housing stock evolved before and after the reform by floor area group. Especially the stock of smaller and medium-sized housing units grew during the period we analyze, but there are no sharp changes coinciding with the reform. Overall, based on the evidence summarized in Figure 5, supply responses are unlikely to be important drivers of our rent results.

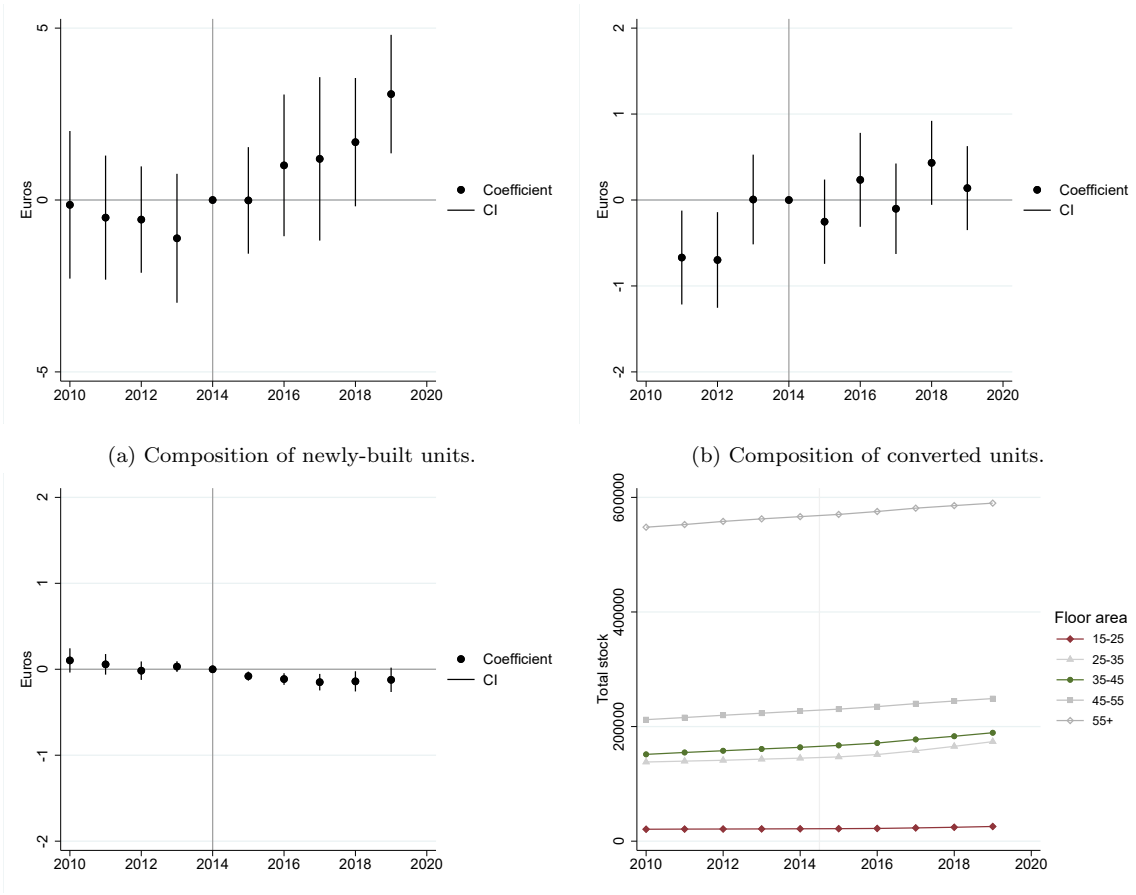

(a) Composition of newly-built units.

(b) Composition of converted units.

(c) Composition of total rental stock.

(d) Total housing stock.

Figure 5: Rental housing supply and housing stock.

Notes: The top-left graph describes the composition of newly built and privately rented units annually (units such that someone lives in the unit at the end of year t and unit is built during t). The top-right graph describes the composition of the units converted from owner-occupied units to privately rented units during the year. The bottom-left graph illustrates the composition of the total private-market rental stock. The bottom-right graph illustrates the full apartment stock (not only the private rental market) in levels by floor area group. Panels (a)-(c) exclude units that are held by households where at least one member was receiving student allowance. In panels (a)-(c), regressions include fixed effects at the municipality level. Standard errors are clustered at the zipcode level. The whiskers describe the 95% confidence intervals. For details on the samples used for each graph, see Appendix A. N = 41,793 in panel (a), N =70,907 in panel (b) and N= 2,823,964 in panel (c).

5 Conclusions

This paper addresses the question of how rent subsidies affect rents by exploiting exogenous variation generated by a large housing allowance reform in Finland and by using rich register data on allowance recipients and the full population. The reform led to differential increases in allowance payments depending on unit characteristics. We find that despite large differences in housing allowance increases by unit types, the reform did not have economically significant effects on the relative rents between different types of units. This implies that the changes in housing allowances mainly benefited the recipients rather than their landlords.

The mechanisms via which rent subsidies affect rents, and how the rent effects depend on the characteristics of the housing market and of the subsidy program, have received little attention in the literature on the incidence of rent subsidies. We study not only the rent effects of the allowance reform but also the housing consumption choices of recipient households and the quantities of rental housing supplied. Despite large changes in financial incentives to choose different types of units, we observe only modest changes in recipient households' housing choices. Recipients do not stay longer in units with large increases in allowances after the reform. However, conditional on moving, they start choosing units with slightly larger allowance increases than before the reform. We argue that one explanation for small demand responses is that allowance spells are often short.

While the construction of small units increased after the reform, this led to very slow changes in the stock of private-market rental units, given that only a small share of the housing stock is newly built. Taken together, these observations about recipient's housing choices and quantities supplied in the rental market suggest that the small rent effects are more likely due to relatively small demand responses to the reform rather than a strong supply response.

Is our pass-through estimate externally valid? We argue that there is no single structural parameter that describes the pass-through of rent subsidies to rents. This pass-through is always context-specific and depends on the details of the program and housing market conditions. We address this point by carefully documenting demand- and supply-side changes in the housing market that are likely contributing to the small rent effects we uncover, thereby improving the external validity of our study relative to the existing literature. Future research on rent subsidies should aim not only at estimating pass-through but also at characterizing the demand- and supply-side mechanisms driving the results.

References

Brewer, M., Browne, J., Emmerson, C., Hood, A., Joyce, R., 2019. The curious incidence of rent subsidies: Evidence of heterogeneity from administrative data. Journal of Urban Economics 114, 103198.

Callaway, B., Goodman-Bacon, A., Sant'Anna, P.H., 2024. Difference-in-differences with a continuous treatment. Technical Report. National Bureau of Economic Research.

Collinson, R., Ganong, P., 2018. How do changes in housing voucher design affect rent and neighborhood quality? American Economic Journal: Economic Policy 10, 62–89.

De Chaisemartin, C., 2010. A note on instrumented difference in differences. Unpublished manuscript .

Eerola, E., Lyytikäinen, T., 2021. Housing allowance and rents: Evidence from a stepwise subsidy scheme. The Scandinavian Journal of Economics 123, 84–109.

Eriksen, M.D., Ross, A., 2013. The impact of housing vouchers on mobility and neighborhood attributes. Real Estate Economics 41, 255–277.

Eriksen, M.D., Ross, A., 2015. Housing vouchers and the price of rental housing. American Economic Journal: Economic Policy 7, 154–76.

Fack, G., 2006. Are housing benefit an effective way to redistribute income? Evidence from a natural experiment in France. Labour Economics 13, 747–771.

Fetzer, T., Sen, S., Souza, P.C., 2023. Housing insecurity and homelessness: Evidence from the united kingdom. Journal of the European Economic Association 21, 526–559.

Gibbons, S., Manning, A., 2006. The incidence of UK housing benefit: Evidence from the 1990s reforms. Journal of Public Economics 90, 799–822.

Gibbons, S., Sanchez-Vidal, M., Silva, O., 2020. The bedroom tax. Regional Science and Urban Economics 82, 103418.

Hyslop, D.R., Rea, D., 2019. Do housing allowances increase rents? Evidence from a discrete policy change. Journal of Housing Economics 46, 101657.

Kangasharju, A., 2010. Housing allowance and the rent of low-income households. Scandinavian Journal of Economics 112, 595–617.

Oikarinen, E., Peltola, R., Valtonen, E., 2015. Regional variation in the elasticity of supply of housing, and its determinants: The case of a small sparsely populated country. Regional Science and Urban Economics 50, 18–30.

Öst, C.E., 2014. Economic incentives, housing allowance, and housing consumption: An unintended consequence of a shift in housing policy. Journal of Housing Economics 23, 17–27.

Sayag, D., Zussman, N., 2020. Who benefits from rental assistance? Evidence from a natural experiment. Regional Science and Urban Economics 80.

Susin, S., 2002. Rent vouchers and the price of low-income housing. Journal of Public Economics 83, 109–152.

Viren, M., 2013. Is the housing allowance shifted to rental prices? Empirical Economics 44, 1497–1518.

Appendix

A Data and sample selection

Kela HA register data and sample selection. The HA register of Social Insurance Institution of Finland (Kela) covers years 2008-2019 on a monthly basis, and each regular monthly HA payment is a separate observation. The data contain an ID for the individual to whom the payment was made and the ID of their spouse (if there is one), since the housing allowance is determined at household-level.

We make the following restrictions in selecting our main estimation sample. First, we only include regular monthly payments (excluding for example overpayment recoveries). Second, we exclude the following observations: 1) observations from Åland Islands, as it is a very specific region both in terms of geography and demographics, 2) observations for which either address or zipcode is missing, and 3) observations that are clearly outliers in terms of their rent per m^2 (below 3 euros/m^2 or above 80 euros/m^2). We also exclude all housing units with floor area either below 15m^2 or above 100m^2. Third, we leave out years 2008 and 2009 to avoid any confounders stemming from the financial crisis.

Furthermore, as we want to focus on regular rental contracts, we exclude certain types of observations. First, we exclude recipient households who are owner-occupiers. Second, we exclude apartments from publicly subsidised right-to-occupy apartments ('*asumisoikeusasunnot*' & '*osaomistusoikeusasunnot*'). Third, we exclude social rental housing where rents are regulated and determined based on maintenance and capital costs (identifying these units in the data is based on the information that they benefit from the government-subsidised debt program). Finally, we exclude shared housing units (either if the unit is defined as shared by Kela or if there are more than 12 monthly payments per year for the same unit in our data). This is because the housing units that were occupied by more than one household were subject to special rules in the pre-reform HA system.

We also exclude students from our sample throughout the time period because the housing benefits of students change over time. Before 2017, most students were covered by separate student housing subsidy program (*asumislisä*). In 2017, students became eligible for general HA. To exclude students we use a separate Kela register on student allowance payments (*opintoraha*). We assume that the student status mostly changes end of term in December and June and we classify all individuals who receive student allowance at least once during the half-calendar year (January-June or

July-December) as students for the six month period in question. Thus, for example, someone who received student allowance in February 2018, will be excluded from our estimation sample throughout the first half of 2018.[14]

Pensioners have a separate HA program. General HA is not granted to a couple (married or co-habiting) if one of them is entitled to the pensioner's HA or to an individual who is entitled to the pensioner's HA. Before the 2015 reform, families with children entitled to both pensioner's HA and general HA were allowed to choose their program. After the 2015 reform these families have been allocated to the general HA system. The government proposal estimated that the change concerns roughly 2,500 families with children.

We do not directly observe new rental contracts in the data. To determine new rental contracts we proceed as follows: If the individual received HA in another address at most 4 months before, we classify the first observation in the new address as a "new rental contract" (this is done before other sample restrictions, so, for example, if someone moves from social rental housing to private rental housing, the observation in private rental housing is registered as a new rental contract). Misclassification can occur in two ways: First, some contracts can be labeled as new even if in reality they are not. For example, an individual who receives HA may move to a unit in which someone was already living in and therefore the rental contract may benefit from terms that had been set prior to the move. Secondly, some rental contracts that are new will not be identified as such. For example, when an individual appears in the HA register for the first time, we will not classify the observation as a new contract although the individual could well have moved at the same time.

For the housing unit fixed effects analysis, we identify repeat observations of units using the exact street address including the unit number. Apartment floor area is self-reported and there are some repeat observations where the reported floor area group varies, for example if someone living in a shared unit misreports the floor area. We exclude these observations.

Our main outcome variables are the allowances paid to the household and the rent paid by the household (rent excludes other costs such as the water charge). These and other sample characteristics are summarized in Table 1 in the main text. In reporting summary statistics, we report all monetary values in 2020 euros, where the deflator is obtained from the Statistics Finland CPI.[15] To describe household incomes, we use

[14] A small fraction of the observations in our main estimation sample are individuals whose spouse is a student (approximately 0.7% before 2017 and 2% after 2017.)

[15] Official Statistics of Finland (OSF): Consumer price index [e-publication]. ISSN=1799-0254.

variable "Household income". This variable is used in the summary statistics table as well as in Appendices B.1 and B.4. To compute the predicted HA change used in Table 2, we use variable "Household income net of deductions" (income net of earnings deductible).

Treatment exposure. We compute the treatment exposure for all new rental contracts observed in 2010–2014. We exclude households with more than eight members as they are unlikely to live in units with at most 100m^2. For all units, we use the information on pre-reform unit and tenant characteristic. We first deflate all pre-reform housing costs from different years to 2014 euros (we use the HA register variable *housing cost used to determine HA* which includes the rent and other necessary costs such as water). We then use the housing cost, floor area of the unit, municipality group and construction year of the building (assuming that all buildings have central heating, which is very likely to be the case), together with the observed pre-reform household size and income (deflated to 2014 euros), to predict the level of HA that the household should have using the 2014 HA policy parameters. Next, for the same observations, we predict the HA that the household should have had in the same unit with the same housing cost and income deflated to 2015 euros with the 2015 HA policy parameters. This *predicted change in HA* for each unit is our measure of treatment exposure.

Statistics Finland register data. We use population-wide register data from Statistics Finland ready-made research data modules (*Folk Basic* and *Income*). In addition, we have obtained data on HAs from the register-based total statistics on income distribution (*Tulonjaon kokonaistilasto*). These data include the annual amount of general HA, pensioners' HA and students' housing supplement. The ready-made data and the tailored data are combined with secured individual identifiers. We also use Statistics Finland's housing unit and building register data. This data covers the universe of buildings and housing units in Finland, with unique identifiers. The data includes information on the building, such as construction year, as well as data on each unit, such as the floor area. The unique unit identifiers can be linked to individuals.

Using the information on units and individuals, we construct a dataset at the level of households (individuals who share the same unit). Throughout, we focus

Helsinki: Statistics Finland [referred: 7.6.2022]. Access method: `http://www.stat.fi/til/khi/index_en.html`

only observations where the building type is a multi-unit building (referring to blocks of flats in residential use). We focus on units with floor area between 15 and 100 m^2 and exclude units where the floor area is missing. For most of the analysis, we exclude student households based on whether at least some member of the household has been receiving student allowance during the year. For most of the analysis, we only focus on units on the private unregulated rental sector.

Table A1 summarizes the Statistics Finland data, both household and unit characteristics, at the household level, for the sample of households living in blocks of flats in the private rental sector, excluding student households. Household income refers to the total disposable income at the household level (after taxes and transfers) and HA refers to general HA (not to students' or pensioners' HA). Households receiving at least 100 euros of HA during the year are classified as HA recipients.

Table A1: Household characteristics, private rental market.

	Non-recipients	Recipients	Both
	mean	mean	mean
Household size	1.4	1.5	1.4
Floor area	49.0	46.2	48.2
Brand-new unit	0.015	0.013	0.015
Income (excl. HA)	29,209	15,998	25,501
HA		3036	845.7
Share HA recipients			0.28
N	2,035,606	794,329	2,829,935

Notes: Statistics Finland register data 2010-2019. Table summarizes the tenants in the private rental market in apartments of floor area between 15 and 100 m^2, excluding students. Households receiving at least 100 euros of HA during the year are classified as HA recipients.

Our main use of the Statistics Finland register data is to describe changes in the supply of rental units. Graph (d) in Figure 5 gives an overview of the growth of the stock of all units. This includes all units in permanent use in multi-unit buildings at the end of each year, also containing units that are not on the private rental market or that are held by students. For the purposes of graphs $(a) - (c)$ in Figure 5, we look at households living in multi-unit buildings in the private rental sector, excluding student households. Graph (a) only looks at new construction of private-market rental units. We define new construction of private-market rentals as units which

are built in year t and where a private-market tenant is residing at the end of year t. This means that we underestimate the level of new construction slightly, as some units which are completed during year t might not be held by tenants by the end of year t. However, we only focus on the composition changes in new construction instead of levels. To the extent that the units which are occupied at the end of the year are not different from the units which are unoccupied, this is not an issue for the analysis of the composition of new construction. Furthermore, graph (b) only looks at units which are converted from units held by owner-occupiers in t-1 to units held by private rental market tenants in year t.

B Additional Results

In this appendix, we provide additional results supporting our main analyses. First, we explore whether the evolution of household characteristics is balanced across different types of units. Second, we show that our main results are robust to alternative treatment exposure definitions. These alternative treatment definitions both provide a robustness check for our main results and also will be used for the analysis of different mechanisms. Third, we analyze changes in HA and rent separately in different-sized cities. Finally, we analyze the role of the social assistance system.

B.1 Household characteristics

In this appendix section, we inspect the evolution of household characteristics that are used to determine HA levels. The aim is to assess whether there were changes in household sorting to treated and non-treated apartments after the 2015 reform.

First, Figure B1 describes the average household size in different floor area groups. The figure indicates that in our two-group DID estimation with 15–25m^2 and 35–45m^2 apartments majority of the recipients were single-member households, and the reform did not have any effect on this.

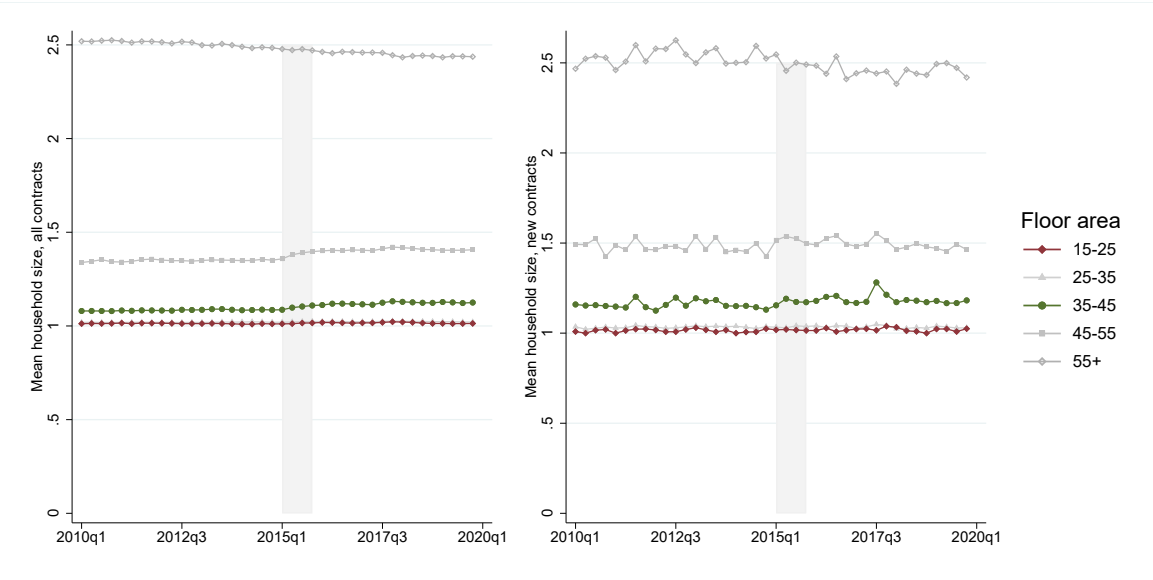

Figure B1: Mean household size in different floor area groups. All payments (left) and new rental contracts (right).

Notes: Mean household size for all payments in our estimation sample (left) and new rental contracts (right), aggregated to quarterly level. The light gray shaded area refers to year 2015. For details on identifying new rental contracts, see Appendix A.

Figure B2 describes the evolution of household incomes in different floor area groups. Our measure of household income excludes any HA payments or social assistance payments (for details on social assistance, see appendix B.4). While average incomes increase over time, there are no significant changes in the time trends after the reform.

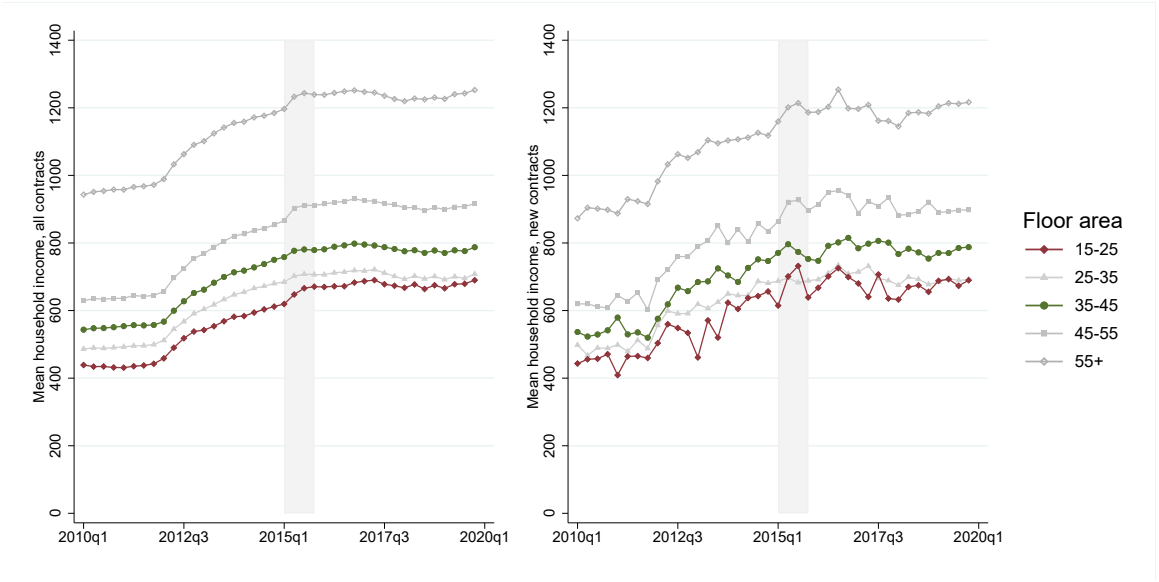

Figure B2: Mean household income in different floor area groups. All payments (left) and new rental contracts (right).

Notes: Mean monthly household income in all payments in our estimation sample (left) and new rental contracts (right), aggregated to quarterly level. The light gray shaded area refers to year 2015. For details on identifying new rental contracts, see Appendix A.

Together, Figure B1 and Figure B2 suggest that household composition in units of different sizes (household sorting into units) did not change significantly after the reform.

B.2 Rent effects using alternative treatment definitions

This appendix section verifies that our main HA and rent results are robust to different definitions of the treatment. In particular, the results hold also when we use the alternative treatment definitions that are used to analyze the demand and supply responses in Section 4.

Discrete treatment. We begin by verifying that our results hold when we *do not compute a variable describing treatment exposure at all.* We can do this by comparing HAs and rents in different floor area groups (as motivated by Figures 2a and 2b). In this approach, we compare HAs and rents in very small housing units (15–25m^2) to those in medium-sized units (35–45m^2). We run the following event-study style

regressions using the sample of new rental contracts:

$$y_{it} = \sum_{\substack{s=2010q1 \\ s\neq 2014q4}}^{2019q4} T_i\,\theta_s + \delta_q + \omega_z + T_i + u_{it} \tag{10}$$

where the outcome variable y_{itz} is either HA or rent. Subscript i indexes rental contracts (a combination of a household and a housing unit) and t time. T_i denotes an indicator variable that takes value 1 for treated units (15–25m^2) and value 0 for control units (35–45m^2). The coefficients of interest in equation (10) are treatment group times quarter fixed effects $T_i\,\theta_s$. The last quarter before the reform is the omitted category implying that the other coefficients measure the differences in the group difference relative to the pre-reform value. We estimate this model including fixed effects at the quarter (δ_q) level and either zipcode (ω_z) or housing unit level.

In Figure B3, we show the event study estimates from the discrete treatment design, comparing floor area groups 15–25m^2 (treated units) to 35–45m^2 (control units). Panel (a) of the figure indicates that HAs and rents developed in parallel before the reform. After the reform, there was a sizable increase in HA in the treatment group relative to the control group, amounting to approximately 70 euros of additional HA per month. Despite this increase, there were no changes in the relative rents between the two groups. Panel (b) in Figure B3 shows the event study estimates with unit fixed effects. Confidence bands are now wider, but the patterns of average HAs and rents are very similar. After the reform, HAs increased substantially in 15–25m^2 (treated units) relative to 35–45m^2 (control units), but rents were unaffected.

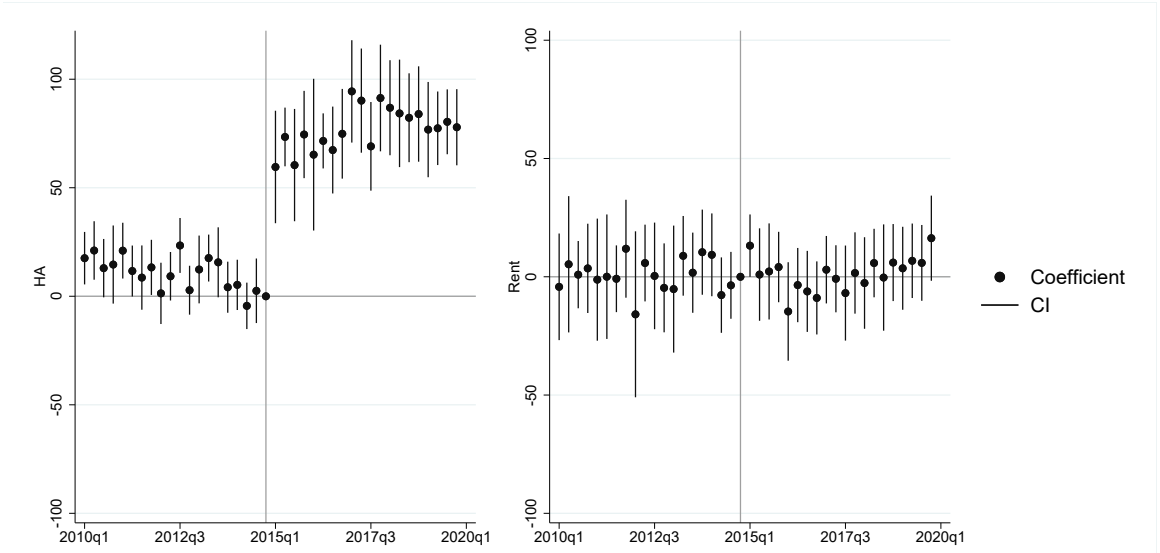

(a) Zipcode fixed effects.

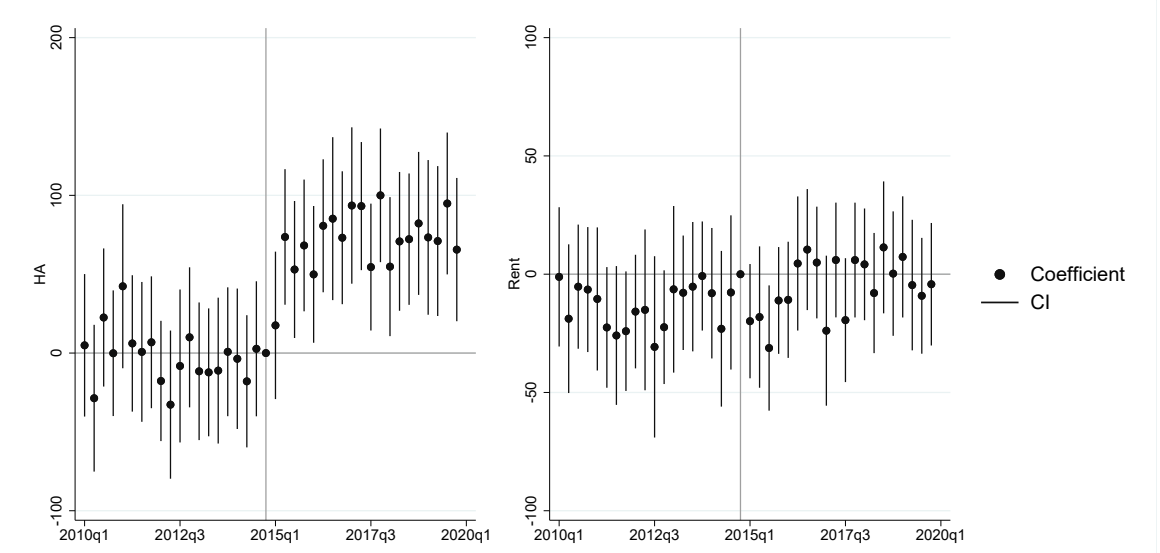

(b) Unit-level fixed effects.

Figure B3: Comparison of monthly HAs and rents between discrete treatment (15–25m^2) and control groups (35–45m^2), new rental contracts.

Notes: The figure plots coefficients from an event study-regression, where the outcome (HA or rent) is regressed on a treatment group indicator, quarter fixed effects, zipcode *or* unit fixed effects and treatment × quarter fixed effects, omitting the last quarter before the reform. Dots and whiskers illustrate the point estimate and the 95% confidence intervals of the treatment × quarter coefficients. Standard errors are clustered at the municipality level in Panel (a) and at the unit level in Panel (b). N = 45,416 in Panel (a) and N = 11,045 in Panel (b).

Average treatment exposure by household size and unit characteristics. Next, we verify that our results hold for two alternative treatment definitions. First, we report results from a regression which uses as the treatment variable the average predicted HA changes by household and unit characteristics ($\overline{\Delta pred\_HA}_1$). We use this treatment definition also in our mobility analysis in Section 4.2 and in Appendix B.3 where study the rent effects by city size. Second, we exclude household characteristics and only use unit characteristics in the assignment of treatment exposure to housing units ($\overline{\Delta pred\_HA}_2$). This treatment definition is also used to analyze the evolution of the housing stock in Section 4.3.

In the sample of new rental contracts in the private rental sector, for households with at most 3 members, observed in 2010–2014, we compute for each observation the predicted changes in HA as described in Appendix A. For treatment definition $\overline{\Delta pred\_HA}_1$ we average the treatment exposure by household size (1, 2, or 3 members), floor area ($2m^2$ brackets) and city size (3 groups).[16] This gives us the predicted average change in HA for example for a 2-member household living in a 40-m^2 unit in the Helsinki MA. We will use this average predicted change in HA as the treatment variable for any 2-member household with a new rental contract in a 40-m^2 unit in Helsinki MA. For treatment definition $\overline{\Delta pred\_HA}_2$ we average the predicted HA changes only across floor area ($2m^2$ brackets) and city size (3 groups). This gives us the predicted average change in HA for example for a 40-m^2 housing unit in the Helsinki MA to be used as the treatment variable for a new rental contract in a 40-m^2 housing unit in Helsinki MA. In calculating the averages, we only use households with at most 3 members.

The resulting average predicted HA changes are illustrated in Figure B4 where Panels (a)-(c) relate to treatment definition $\overline{\Delta pred\_HA}_1$ for different-sized households and Panel (d) to definition $\overline{\Delta pred\_HA}_2$.

[16]The groups are: 1) Helsinki Metropolitan Area (Helsinki, Espoo and Vantaa), 2) six other large cities (Tampere, Turku, Oulu, Jyväskylä, Kuopio and Lahti), and 3) rest of Finland which pools together the remaining municipalities.

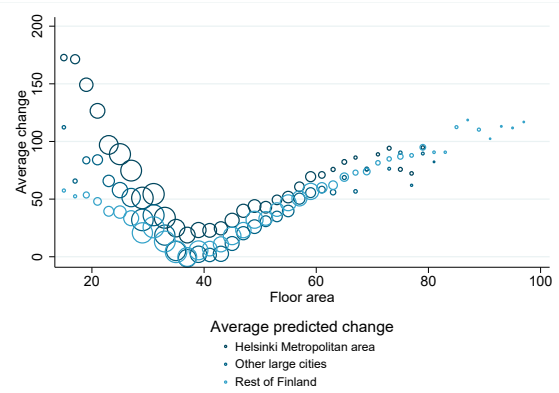

(a) Average predicted HA change ($\overline{\Delta pred\_HA}_1$) for single-member households.

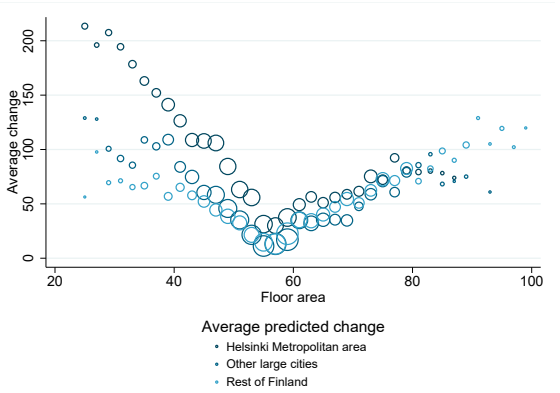

(b) Average predicted HA change ($\overline{\Delta pred\_HA}_1$) for 2-member households.

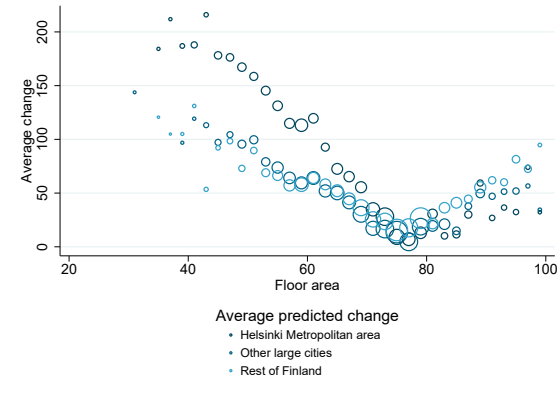

(c) Average predicted HA change ($\overline{\Delta pred\_HA}_1$) for 3-member households.

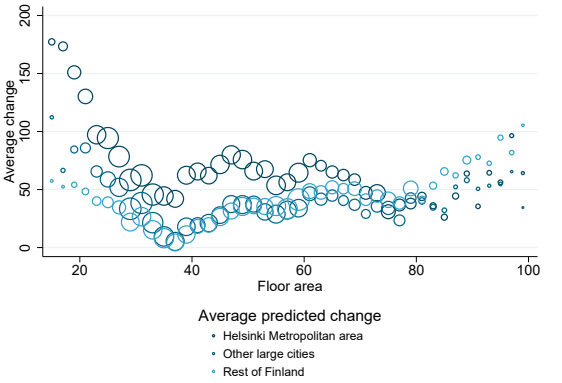

(d) Average predicted HA change ($\overline{\Delta pred\_HA}_2$).

Figure B4: Different measures of treatment exposure.

Notes: Panels (a)-(c) summarize $\overline{\Delta pred\_HA}_1$ and panel (d) summarizes $\overline{\Delta pred\_HA}_2$. Different treatment exposure correspond to predicted HA change in the sample of new contracts observed in 2010–2014, averaged at different cell levels. $\overline{\Delta pred\_HA}_1$ averages over household size, unit floor area and unit municipality group. $\overline{\Delta pred\_HA}_2$ averages over unit floor area and unit municipality group. Cells with fewer than 6 observations are excluded.

We then run the following event study regression in the sample of all new rental contracts 2010–2019 using predicted changes averaged either by k=1,2:

$$y_{it} = \theta\ \overline{\Delta pred\_HA}_{kc} + \sum_{\substack{s=2010q1 \\ s\neq 2014q4}}^{2019q4} \theta_s \overline{\Delta pred\_HA}_{kc} + \delta_q + \gamma_z + \epsilon_{it}, \tag{11}$$

where the outcome is either the HA or the rent of the unit-household pair i. Subscript k in the treatment variable refers to the treatment definition and c to the cell. The regression includes fixed effects for the zipcode (γ_z) and the quarter (δ_q).

Figure B5 and Figure B6 report the event study-graphs that correspond to estimation equation (11) with $\overline{\Delta pred\_HA}_1$ and $\overline{\Delta pred\_HA}_2$ respectively. In Figure B5 a one euro increase in predicted HA maps almost one-for-one to the actual HA increase. This is very cleanly estimated with very narrow confidence intervals suggesting that also this treatment definition is strongly associated with actual variation in HA. For rents, we do not detect any changes between units that receive smaller and larger predicted HA changes. Moreover, this sample is substantially larger than our unit fixed-effects sample (we now use information on N=200,890 new rental contracts instead of N=22,346 in the fixed-effects specification), improving the overall validity of our findings.

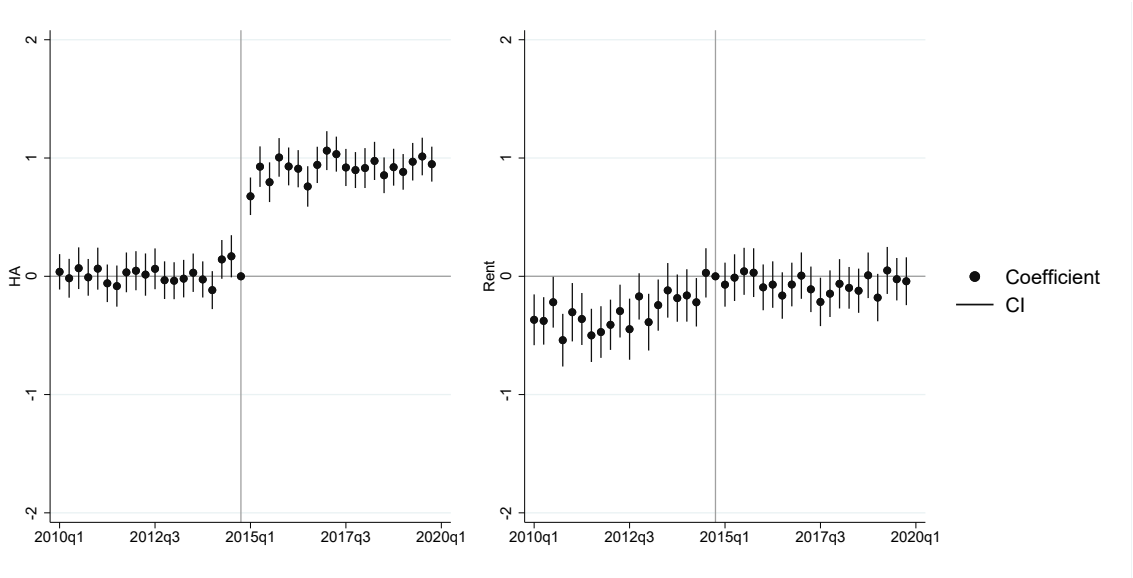

Figure B5: Average HA and rent regressions using treatments averaged by household size and unit characteristics.

Notes: The figure plots coefficients from an event study-regression, where the outcome (HA or rent) is regressed on the level of predicted HA change, quarter fixed effects, zipcode fixed effects and treatment × quarter fixed effects, omitting the last quarter before the reform using treatment definition $\overline{\Delta pred\_HA}_1$. Dots and whiskers illustrate the point estimate and the 95% confidence intervals of the treatment × quarter coefficients. Standard errors are clustered at the zipcode level. N=200,890.

Results shown in Figure B6 are again consistent with our main results, although the first-stage results are slightly less strong. A one euro increase in predicted HA maps to approximately 0.75 euro increase of actual HA. For rents, again, we do not detect any changes between units that receive smaller and larger predicted changes.

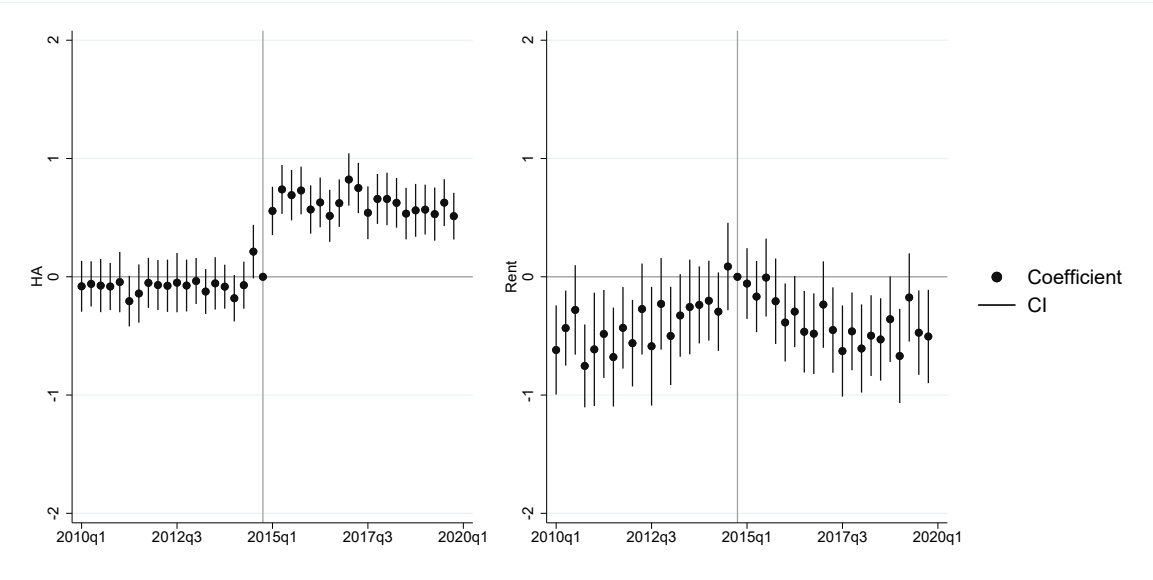

Figure B6: Average HA and rent regressions using treatments averaged by unit characteristics.

Notes: The figure plots coefficients from an event study-regression, where the outcome (HA or rent) is regressed on the level of predicted HA change, quarter fixed effects, zipcode fixed effects and treatment × quarter fixed effects, omitting the last quarter before the reform using treatment definition $\overline{\Delta pred\_HA}_2$. Dots and whiskers illustrate the point estimate and the 95% confidence intervals of the treatment × quarter coefficients. Standard errors are clustered at the zipcode level. N=201,531.

B.3 Rent effects by city size

This appendix section verifies that our rent results hold if we split the sample and run the analysis separately for different-sized cities. Repeating the analysis by city size also alleviates the concern that pre-reform trends in rents would have been potentially different depending on the region. We split the country to three groups: Helsinki Metropolitan area (the largest metropolitan area in Finland, with population above 1,000,000), 6 other large cities (each has population greater than 100,000), and other Finland.

Figure B7 first replicates the descriptive analysis by floor area group of Figure 2, focusing only on new rental contracts. The left-hand graph shows that the 2015 reform increased mean HA in small apartments especially in Helsinki metropolitan area (Panel (a)) and in other large cities (Panel (b)). The right-hand graph in turn shows again no visible effect on rents in any group. Figure B8 repeats the estimation shown in Figure B5, but again splitting the sample by city size. In these regressions, we regress allowances or rents on the average predicted change given the observed household size and unit floor area ($\overline{\Delta pred\_HA}_1$). Again, we confirm a strong first stage and no rent effects. Moreover, more and less treated units have very similar parallel trends prior to the reform across all regions.

(a) Helsinki MA

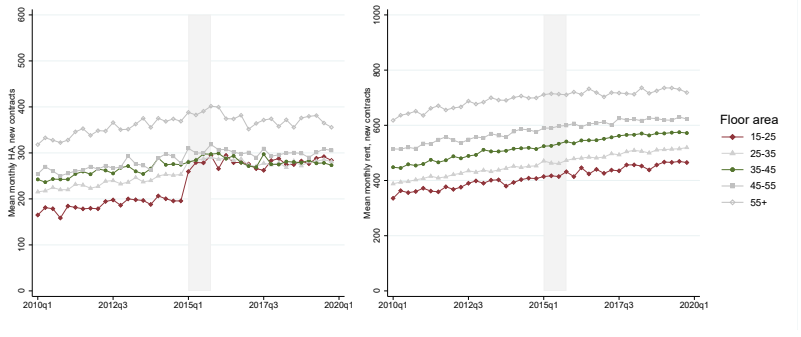

(b) Other cities

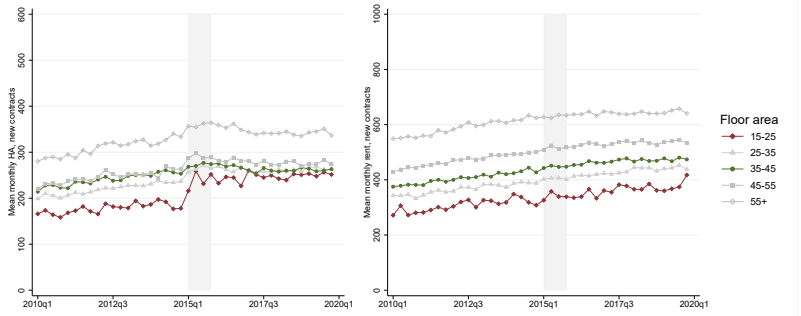

(c) Rest of Finland

Figure B7: Mean HAs and rents by floor area and municipality group, new contracts.

Notes: The figure shows mean monthly HA paid to recipient households and mean monthly rents paid by recipients in new contracts by floor area and municipality group, aggregated to quarterly level. Panel (a) shows Helsinki Metropolitan area, Panel (b) shows six other large cities and Panel (c) the rest of Finland. The light gray shaded area refers to year 2015. For details on sample selection and identifying new rental contracts, see Appendix A.

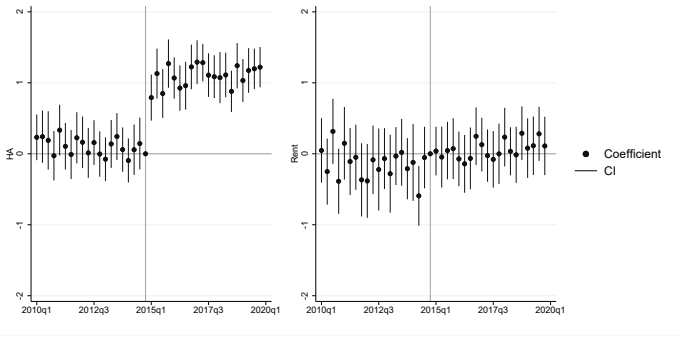

(a) Helsinki MA

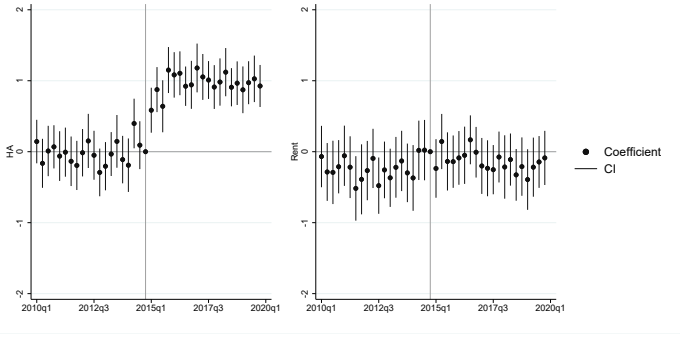

(b) Other cities

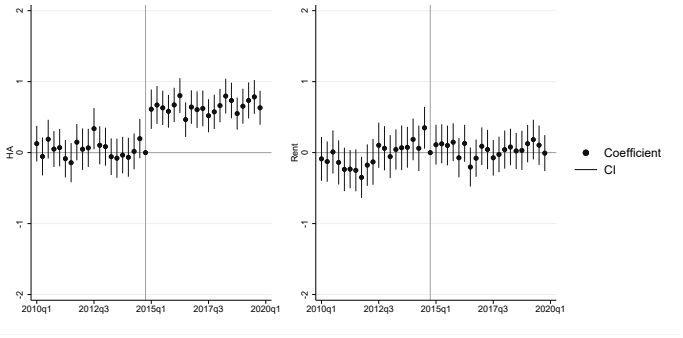

(c) Rest of Finland

Figure B8: Comparison of monthly HAs and rents by municipality group, new contracts.

Notes: The figure plots coefficients from an event study-regression that corresponds to Figure B6 but is split by municipality group. The outcome (HA or rent) is regressed on a the level of treatment, quarter fixed effects, zipcode fixed effects and treatment × quarter fixed effects, omitting the last quarter before the reform. The treatment definition used is $\overline{\Delta pred\_HA}_1$. Dots and whiskers illustrate the point estimate and the 95% confidence intervals of the treatment × quarter coefficients. Standard errors are clustered at the zipcode level. Panel (a) shows Helsinki Metropolitan area (N=36,646), Panel (b) shows six other large cities (N=58,552) and Panel (c) the rest of Finland (N=105,395).

B.4 Social assistance and rents

In this appendix section, we assess the robustness of our baseline results to other elements of the social security system that could potentially mitigate the effects of the HA reform. Social assistance, or income support, is the last-resort form of financial assistance in the Finnish social security system. Households with very low incomes can receive social assistance if their income, including HA, is considered insufficient for covering their basic needs. As a result, for households eligible for social assistance, an increase in HA can be partly or entirely offset by a reduction in social assistance. Therefore, the rent effects of HA changes might be muted for this group.

The social assistance program consists of three parts: basic, supplementary, and preventive social assistance. The two latter programs are small relative to basic social assistance which we focus on. Basic social assistance can reimburse housing costs in full. It is calculated by adding together all household income, including HA, and subtracting from it housing costs as well as a "base amount" meant to cover necessary non-housing living costs.[17] Basic assistance is only granted for 1–2 months at a time as opposed to HA which is usually granted for 12 months.

HA and social assitance in the population register data. The HA register data does not contain information on social assistance. However, using data from the full population register from Statistics Finland, we observe the yearly HA payments and yearly social assistance payments separately. We can therefore verify that the total subsidies paid to households did actually change, and that the increases in HA payments were not offset by changes in social assistance.

Motivated by the graphical analysis in Figures 2a and 2b, we compare total subsidy payments across different floor area groups before and after the HA reform. In the Statistics Finland register data, the annual general HA payment to households in 15–25 m^2 units increased by 416 euros more than for 35–45 m^2 units from 2014 to 2016.[18] This is the average across all renter households in given unit types, not only

[17] There is a municipality-specific limit on the housing costs that can be covered. The limits are higher than the limits in HA system and less strictly enforced. If housing costs exceed the limit, the recipient may be instructed to seek more affordable housing. However, Kela may also reimburse the housing costs in full, if, for instance, affordable housing is not available. The declared housing costs are accepted in full in roughly 70% of the cases where housing costs exceed the municipality-specific limit.

[18] All the reported numbers are for households in blocks of flats in private-market rentals, excluding students, and includes non-recipients as zeros (thus not conditioning on being HA recipient).

for subsidy recipients.[19] Thus, the variation in HA is also clear at the level of the overall rental market, not only at the level of recipients. Taken together with the fact that around 28% of the renters in the private rental market are HA recipients, these numbers are broadly speaking consistent with the monthly HA payments reported in figures 2a and 2b. This suggests that the reform caused large variation in HA payments at the level of the overall rental market.

When we compare annual average social assistance payments in different floor area groups, we see that the basic assistance in decreased by 145 euros per year in 15–25 m^2 units relative to 35–45 m^2 units from 2014 to 2016. Together with the increase in HA payments, this implies that the *net* change in total subsidies (HA + social assistance) paid to households increased by 271 euros annually (including non-recipients) in 15-25 m^2 units relative to 35-45 m^2 units. This suggests that the total subsidy payments clearly increased in 15–25 m^2 units and the differential changes in HA were only moderately offset by social assistance changes.

Since these numbers include non-recipients as zeros as well, they suggest that the 2015 reform caused variation in (total) subsidy payments which is large enough to affect the overall rental market even after accounting for changes in social assistance payments.

Rent effects by recipient income. After having verified that the HA changes are only moderately offset by basic assistance changes, we also check if our rent effects are different for those who are potentially eligible for the social assistance. Our HA register data does not include information on social assistance. Therefore, we test the robustness of our results with respect to social assistance status by splitting the sample to low- and high-income HA recipients. We divide households in our main estimation sample to two groups depending on whether household income per consumption units (excluding HA) is below or above the median in our sample. We exclude households with exactly the median income: In most years, the median household income appears to correspond to monthly basic unemployment allowance. Across the years, 41% of HA recipients have incomes strictly below median and 37% strictly above. The households with above-median income are unlikely to be eligible for social assistance.[20]

[19]These numbers, since they are averages in the overall rental population, are of course affected not only by subsidies paid to a single household but also by changes in the number of recipients.

[20]We only look at households with at most 3 members. We construct consumption equivalence units as follows: adults in the household count for one and children under the age of 18 for 0.5 consumption units. We calculate median incomes in the sample annually.

Figure B9 shows the evolution of mean HAs and rents in different floor area groups, with the top panel describing HA recipients below and bottom panel above median income. Although the levels of mean HA received in the different groups are different, the reform resulted in very similar variation in HA by floor area group. HA increased substantially in 15–25 m^2 units and did not increase by much in 35–45 m^2 units (with the exception of below median income households whose mean HA increased slightly also in 35–45 m^2 units). The right-side graph in turn shows that there are no visible differences in mean rents of those above and below median income.

We then run an analysis which corresponds to Figure B6 but where we split the sample by household income. The resulting coefficients are reported in Figure B10. We again observe that allowances increased in both groups and that there are no meaningful rent effects in either group. Overall, even in the subset of households who are very unlikely to receive social assistance, rent effects are very similar to our main estimates.

(a) New constracts of households with below-median incomes.

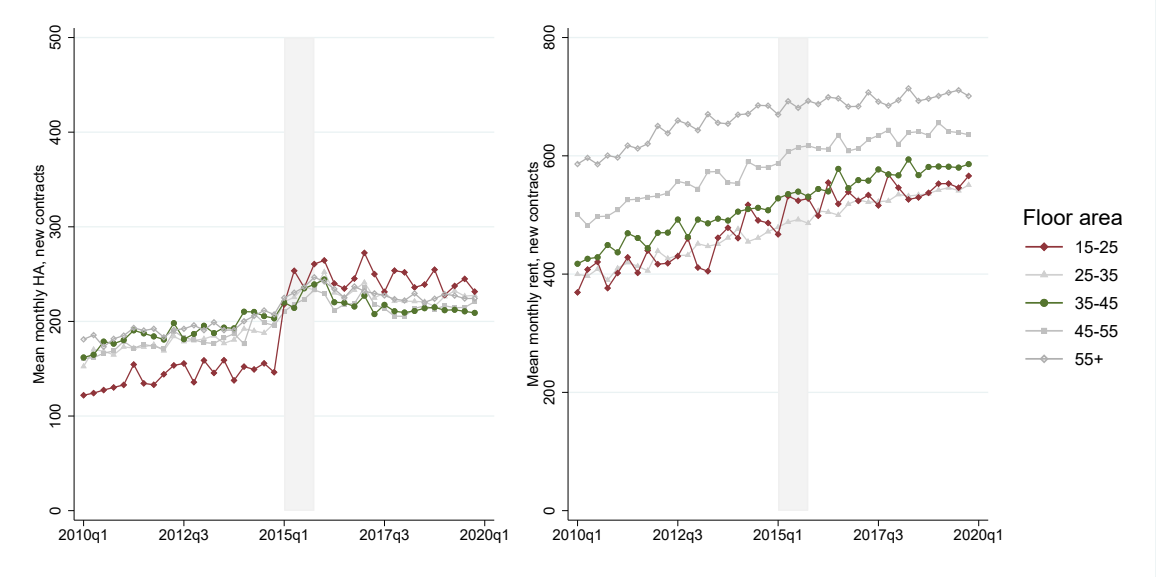

(b) New constracts of households with above-median incomes.

Figure B9: Mean HAs and rents by household income.

Notes: The figure shows mean monthly HA paid to recipient households and mean monthly rents paid by recipients in our estimation sample, aggregated to quarterly level. The light gray shaded area refers to year 2015. Panel A illustrates 1-3-member households with strictly below-median incomes. Panel B illustrates 1-3-member households with strictly above-median incomes.

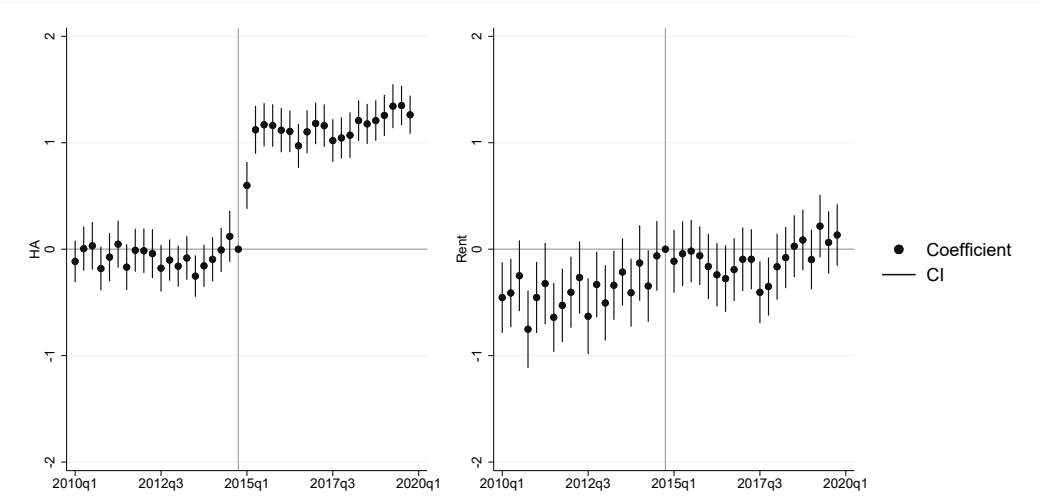

(a) Households with below-median incomes

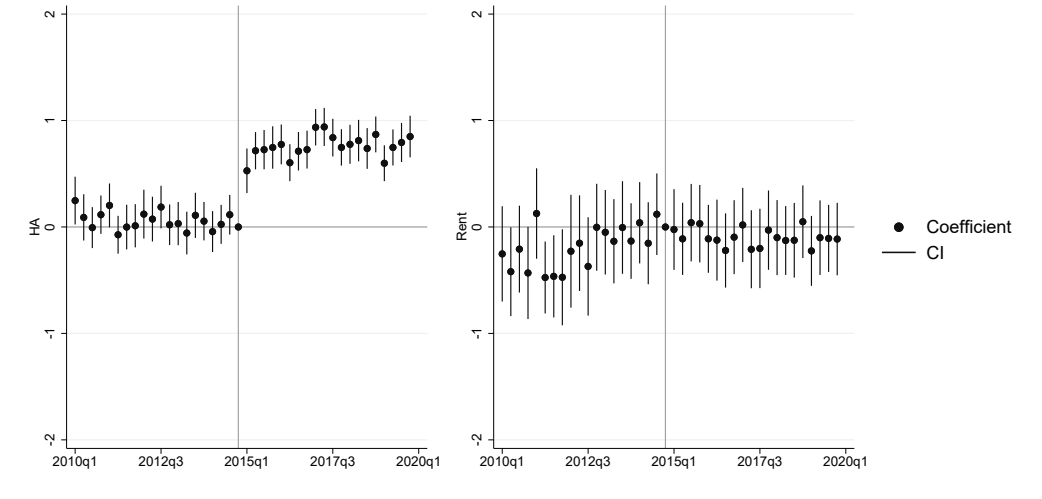

(b) Households with above-median incomes

Figure B10: Comparison of monthly HAs and rents by household income.

Notes: The figure plots coefficients from an event study-regression that corresponds to Figure B5 but is split by recipient income. Panel A illustrates 1-3-member households with strictly below-median incomes. Panel B illustrates 1-3-member households with strictly above-median incomes. The outcome (HA or rent) is regressed on the level of treatment, quarter fixed effects, zipcode fixed effects and treatment × quarter fixed effects, omitting the last quarter before the reform. The treatment definition used is $\overline{\Delta pred\_HA}_1$. Dots and whiskers illustrate the point estimate and the 95% confidence intervals of the treatment × quarter coefficients. Standard errors are clustered at the zipcode level. N=83,445 in Panel (a) and N=75,065 in Panel (b).

C Derivations for the Conceptual Framework

Consider a competitive rental market where $D(r)$ is aggregate rental demand and $S(r)$ aggregate rental supply as a function of the quality-adjusted rent r. Rental demand aggregates over demand by HA recipients (R) and non-recipients (N)

$$D(r) = D^R(r) + D^N(r).$$

Supply of rental housing aggregates over three margins: the construction of new units (C), the (net) conversion of existing owner-occupied units to rental units (I), and the stock of pre-existing rental units net of depreciation (δ)

$$S(r) = S^C(r) + S^I(r) + S^O(1-\delta).$$

Initially $D(r) = S(r) = q$ where q denotes the quantity consumed in the initial equilibrium.

Consider then a small change in the housing allowance, ds. This change affects the rent paid by HA recipients, $r - ds + dr$, and through the demand effects also rents paid by non-recipients, $r + dr$, as well as rent payments received by landlords, $r + dr$.

Taking into account changes in quantity demanded and supplied gives

$$D^R(r - ds + dr) + D^N(r + dr) = S^C(r + dr) + S^I(r + dr) + S^O(1 - \delta).$$

Rewriting and taking into account $D(r) = S(r)$ gives

$$D^{R'}(r) \cdot (dr - ds) + D^{N'}(r) \cdot dr = S^{C'}(r) \cdot dr + S^{I'}(r) \cdot dr$$

$$\frac{dr}{ds} = -\frac{D^{R'}(r)}{[S^{C'}(r) + S^{I'}(r)] - [D^{R'}(r) + D^{N'}(r)]}$$

Now, expressing terms on the RHS as $D^{R'}(r) = \dfrac{dD^R}{dr}$, and so on gives:

$$\frac{dr}{ds} = -\frac{\frac{dD^R}{dr}\frac{r}{q}}{[\frac{dS^C}{dr}\frac{r}{q} + \frac{dS^I}{dr}\frac{r}{q}] - [\frac{dD^R}{dr}\frac{r}{q} + \frac{dD^N}{dr}\frac{r}{q}]}$$

Finally, labeling the elasticities along the different margins, for example, $\epsilon_D^R = \frac{dD^R}{dr}\frac{r}{D^R}$, gives:

$$\frac{dr}{ds} = -\frac{\frac{D^R}{q}\epsilon_D^R}{[\frac{S^C}{q}\epsilon_S^C + \frac{S^I}{q}\epsilon_S^I] - [\frac{D^R}{q}\epsilon_D^R + \frac{D^N}{q}\epsilon_D^N]}$$

This shows that the rent effect of the subsidy can be expressed as a weighted sum of the elasticities along the different margins, with weights that correspond to the "market shares" of the different components.