Vrije Universiteit Amsterdam
School of Business and Economics
Master of Science in Economics

Is it better to build expensive or cheap housing?

An assignment model approach to the Dutch housing market.

Master's Thesis

Pieter Parlevliet
Student registration number: 2649614

Supervisor: Prof. J. Rouwendal
Second Assessor: Dr. M. Mastrogiacomo

June 26, 2026

Contents

1 Introduction **2**

1.1 Research context 3

1.2 Set-up of this thesis 5

2 Methodology **6**

2.1 The assignment model 6

2.2 Application to the Dutch owner-occupied market 9

2.2.1 Income distribution 10

2.2.2 Housing quality distribution 10

2.2.3 Remaining parameters 12

2.2.4 Resulting model of the Netherlands 14

3 Results **16**

3.1 The effects of cheap and expensive construction 16

3.2 Welfare effects 19

3.3 Welfare effect for any quality 23

4 Including a supply side to the model **24**

4.1 Modeling Dutch construction costs 24

4.2 Should the government intervene in the market? 25

5 Discussion **27**

5.1 Moving chains 28

5.2 Policy should aim for long moving chains 29

5.3 Conclusion 30

Bibliography **31**

6 AI use **32**

1 Introduction

The Netherlands is widely described as being in a housing crisis. House prices have doubled between 2013 and 2025 (Biesenbeek and Rupert, 2025), there are at least 33,000 homeless people (Centraal Bureau voor de Statistiek, 2025a) and many households are spending more on housing than they can afford. There is broad political consensus not only on the urgency of these challenges but also on the central policy response: a substantial increase in housing construction. The estimated housing shortage currently amounts to 396,000 houses (ABF Research, 2025b).

Given the emphasis on increasing housing supply, an important policy question is not only how much housing should be built, but also what type of housing should be constructed. There is considerably less agreement on this topic. In particular, there is debate over whether new housing should be expensive or cheap housing. Over the past two decades, most housing construction in the Netherlands has taken place in the higher price segments (Sniekers, 2024). Housing developers and investors consistently indicate that these dwellings are also the most profitable to build, but there are many politicians and advocacy groups arguing for more cheap and middle-quality housing.

One argument for prioritizing affordable housing is that lower- and middle-income households face the greatest pressure in the current housing market. Existing research suggests that starters and middle-income households experience particular difficulties in accessing affordable housing (Planbureau voor de Leefomgeving, 2017). Since housing shortages are concentrated in the lower and middle segments of the market, this perspective argues that construction should target these segments.

An alternative perspective emphasizes the indirect effects of new construction through moving chains. Households moving into newly constructed homes generally leave behind their previous dwelling, which subsequently becomes available to another household. Through this process, the construction of an expensive house may ultimately affect lower segments of the market as vacancies propagate through the housing stock. For example, a dwelling vacated by a high-income household may be occupied by a middle-income household, which in turn leaves behind another dwelling for a lower-income household.

This thesis provides a theoretical perspective on exactly this discussion and asks: Who benefits from cheap and expensive housing construction? Here, cheap and expensive refer to the qualities that determine a house's price in an owner-occupied market and should not be confused with the regulated rental market, which is what is usually meant with 'affordable housing' in the Dutch debate.

In order to answer the research question, I simulate the Dutch housing market with an

assignment model. Assignment models are particularly suitable because they allow for differentiation in price per house and explicitly determine which households occupy which houses. This makes it possible to trace how new construction affects both prices and housing allocation.

With this assignment model in hand, I show that in the case of perfect moving chains, expensive construction and cheap construction have the exact same effect on low-income households. Expensive construction makes more households move up in the housing quality ladder and thereby creates more welfare increase per house. If we look at welfare increase per square meter added, however, more middle priced houses are best.

1.1 Research context

The standard economics of housing markets typically treats housing as a homogeneous good over which consumers choose the quantity to consume. This framework is not well suited to our research question, as it does not allow for differentiation across housing types in terms of prices per unit of quality. In contrast, assignment models endogenise the sorting of households across heterogeneous housing types and yield a tractable equilibrium in which prices are expressed as a function of housing quality. This feature makes assignment models particularly well suited to studying how changes in the composition of housing supply affect prices and welfare across the distribution.

The study closest to this thesis is Nathanson (2019). He also analyses the effects of cheap and expensive construction in a model that is closely related to an assignment framework, in the sense that housing is indivisible and dwellings are heterogeneous in quality. His analysis takes a local perspective and examines the effects of new construction on residents within a city. In his model, households are divided into high- and low-education groups that can move between cities. The construction of low-quality housing increases welfare for low-income households but reduces welfare for high-income households.

These results crucially depend on geographic spillovers and heterogeneous preferences. As additional low-quality housing is constructed, more low-income households are attracted to the city. This reduces local amenity levels, which are valued more highly by high-income households, thereby lowering their welfare.

My model does not include these spillover effects. Even though there is empirical evidence for these effects, it is useful for policy analysis to isolate the direct market effects of housing construction from the effects operating through amenities and neighbourhood composition.

There are several reasons why it is debatable whether policy conclusions should take such spillovers into account. First, policymakers themselves can influence local amenities through

public investment in neighbourhoods. Second, while highly educated households may prefer neighbourhoods with a larger share of similarly educated residents, low-educated households may have different preferences. However, these preferences are less visible in equilibrium outcomes because they are constrained by purchasing power. Finally, using such equilibrium preferences as a basis for policy risks reinforcing residential segregation, potentially generating long-run social costs that are not captured by the model's welfare measure.

Another paper related to this work is Thiel and Zahnbrecher (2024). They do not use an assignment model but instead employ a structural model to study how housing prices in the Netherlands respond to construction in different locations, given a fixed supply of new dwellings. Their research question is closely related to the one studied in this thesis, but their notion of housing quality is restricted to location.

In contrast, this thesis interprets housing quality as dwelling size, although the framework can be seen as capturing more general forms of quality. Thiel and Zahnbrecher find that prices decrease most strongly when new construction takes place in already high-priced areas. This is because construction in these locations affects prices not only locally but also in other locations through spatial spillover effects.

As will be shown in this thesis, this mechanism is closely related to the concept of moving chains. However, the assignment framework makes these channels more transparent by explicitly determining which households move where and how they value these changes. This tractability also allows for welfare comparisons across different housing supply scenarios.

The assignment model presented here was originally used by Määttänen and Terviö (2014) to study the effects of changes in the income distribution on house prices. This thesis examines the complementary question: how do changes in the distribution of housing quality affect house prices and welfare? A variant of the assignment model has also been used by Rouwendal, Sniekers and Jia (2023) to investigate the price effect of the buy-to-let sector. This thesis leaves out investors and only studies the owner-occupied market.

To summarize, the contribution of this thesis is to apply an existing assignment framework to study how housing construction affects consumer welfare while abstracting from spillover effects. In doing so, the model provides a transparent mapping from housing quality to equilibrium prices, consumption choices, and welfare at the individual household level. This allows for a detailed analysis of distributional effects across households that is typically not tractable in standard housing market models.

1.2 Set-up of this thesis

This thesis is organized as follows. In chapter 2, I introduce the assignment model. After having explained how the model takes an income and housing quality distribution to determine a housing assignment and price function, I fine-tune its parameters to the Dutch owner-occupied housing market. I report on what outcomes the model gives with these parameters.

In chapter 3, I simulate housing construction by modifying the Dutch housing stock. I compare the outcomes to the initial situation, studying scenarios of cheap and expensive construction. Overall, cheap construction causes prices to decrease less than in the case of expensive construction. Taking the equivalent variation as measure for welfare increase, expensive construction also causes more welfare increase than cheap construction.

In chapter 4, I extend the model with a supply side. Developers face a linear cost function. Since the model gives us a price function, we can easily derive a profit function. In this set-up, developers make most profits for expensive houses. In the case that a fixed amount of houses per year is constructed, the market is efficient. If a fixed amount of square meters per year is constructed, there is a market failure. Developers would like to build expensive houses, but most social welfare is added by construction of a lower quality.

Finally, I discuss these results in chapter 5. I argue that the main lesson from the assignment model is that longer moving chains cause more welfare per added house. I put this in the context of experimental literature on how moving chains are not perfect and discuss how more extended models could take chain breakage into account.

2 Methodology

In this section, I describe the model used in this research. First, I explain how the assignment model determines the price of housing. Second, I apply the model to the Dutch owner-occupied housing market by calibrating the model parameters to the Netherlands. This application will be the base scenario from which I will evaluate the effect of housing construction in later chapters.

2.1 The assignment model

The model matches houses to households. Households only differ in income y and houses only differ in quality, determined by a single parameter q. Households can choose between consumption of housing q and other consumption c. All houses have the same Cobb-Douglas utility function:

$$u(c, q) = q^{\alpha} c^{\beta} \tag{1}$$

where

$$\alpha + \beta = 1 \tag{2}$$

The distribution of households is given by the continuously differentiable function $F(y)$. It has positive support between the lowest and highest income y_{min} and y_{max} and gives the cumulative amount of households with that income, so that

$$F(y_{min}) = 0 \tag{3}$$

$$F(y_{max}) = B \tag{4}$$

Where B is the total number of households.

Similarly, the distribution of the housing stock is described by the strictly increasing, continuously differentiable function $G(q)$. It is defined only on the range $[q_{min}, q_{max}]$, with S the total amount of houses, so that we have

$$G(q_{max}) = S \tag{5}$$

Households can spend their income y on consumption or housing, leading to the budget

constraint:

$$c + p(q) = y \tag{6}$$

Where $p(q)$ gives the price of a house of quality q.

In the model, there are not enough houses for every household to have a house. In terms of the parameters, this means that we assume that $S < B$.

The sorting principle is simple: the richest households get the best households. That means that the $S - B$ households that do not get a house are the lowest income households. They have to resort to an option outside the housing market. Depending on what the model exactly represents, this can be the social housing sector, the private rental sector, or the housing market in another region. We assign a quality q_{out} to the outside option, with

$$q_{out} < q_{min} \tag{7}$$

.

The price that households pay for the outside option is p_{out}.

The rest of the households participate in the housing market. We call the income of the first household with a house the critical income y_c:

$$F(y_c) = S - B \tag{8}$$

The household with the critical income gets the house of the lowest quality. Afterwards, the household with the next income level gets assigned to the next quality house, and so forth. This mechanism is formalized in the following assignment function which couples incomes to housing quality:

$$y(q) = F^{-1}(F(y_c) + G(q)) \tag{9}$$

With the matching of households to dwellings in place, the hedonic price function of housing can be derived by starting at the household with the critical income. The price of this house is determined by the condition that this household should be indifferent between the outside option and the house with the lowest quality:

$$u(y_c - p_{out}, q_{out}) = u(y_c - p(q_{min}), q_{min}) \tag{10}$$

This pins down $p(q_{min})$. This is the starting point from which the prices of all higher housing qualities are determined. The willingness to pay for extra housing quality by households in a

higher quality house, is determined by the marginal rate of substitution. This gives a value to the derivative of the price:

$$\frac{dp}{dq} = MRS(q, y(q) - p(q)) \tag{11}$$

In the case of our utility function, we have:

$$MRS(q, c) = \frac{\alpha}{\beta}\frac{c}{q} \tag{12}$$

Evaluating at

$$(q, c) = (q, y(q) - p(q)) \tag{13}$$

gives us the following differential equation that determines the price function:

$$\frac{dp}{dq} = \frac{\alpha}{\beta}\frac{y(q) - p(q)}{q} \tag{14}$$

To derive the hedonic price function, we thus start from the price of the lowest-quality dwelling and integrate upwards. The derivative of the price with respect to housing quality depends on the income of the household matched to a dwelling of that quality.

The remainder of the paper focuses primarily on numerical results obtained for specific parameter values. Before turning to these simulations, however, it is useful to highlight several general properties of the model.

First, because the price function is obtained by integrating upwards, changes in the incomes of high-income households do not affect the prices of low-quality dwellings, provided that the ordering of households remains unchanged. For example, suppose that all households occupying dwellings of quality $q > q'$ receive the same absolute income increase. Prices will then rise for all dwellings with quality above q', while prices for dwellings with $q < q'$ remain unaffected.

Second, consider a change in the housing quality distribution that assigns households to higher-quality dwellings while leaving the income distribution unchanged. In this case, a given housing quality is matched to a household with a lower income than the previous occupant. Since the slope of the price function depends on the income of the matched household, prices fall. Intuitively, an increase in the overall quality of the housing stock reduces the price associated with a given level of housing quality.

Finally, it is important to note that the model generates perfect moving chains. When a dwelling of quality $q_{\max}$ is added to the housing stock, every household moves one step up the

quality ladder. As a result, a dwelling of quality $q_{\min}$ becomes available and is occupied by the household with the highest income below the previous critical income y_c. By contrast, when a dwelling of quality $q_{\min}$ is added, it is occupied directly by the new marginal homeowner, while the assignment of all other households remains unchanged.

Equation 10 shows that, holding $q_{\min}$ and the outside option constant, changes in the minimum house price are determined solely by changes in the critical income y_c. Consequently, the price paid by the lowest-income homeowner is identical in the two scenarios described above.

2.2 Application to the Dutch owner-occupied market

I model both the income and housing quality distributions in the Netherlands using log-normal distributions. Log-normal distributions are right-skewed and exhibit a long upper tail, making them a natural choice for modeling income and housing quality, both of which are unevenly distributed.

A variable is log-normally distributed if its logarithm is normally distributed. Consider, for example, a log-normal random variable (Y) representing household income. Then

$$\ln Y \sim \mathcal{N}(\mu, \sigma^2) \tag{15}$$

where μ and σ^2 denote the mean and variance of the underlying normal distribution. It is standard practice to parameterize a log-normal distribution in terms of these parameters. They are related to the mean of the log-normal variable as follows:

$$E[Y] = e^{\mu + \frac{\sigma^2}{2}} \tag{16}$$

The mode of a log-normal distribution is given by

$$\text{mode} = e^{\mu - \sigma^2} \tag{17}$$

The model uses cumulative distribution functions. To calibrate these to the Dutch situation, I specify a lower and upper cutoff, the total number of houses or households (N), and the desired mean and mode. The cumulative distribution is then constructed in three steps.

First, μ and σ are determined from the target mean and mode using equations 16 and 17. These parameters define the standard log-normal cumulative distribution function $LN_{\mu,\sigma}(x)$.

Second, the distribution is shifted such that the cumulative amount at the lower cutoff equals zero:

Defining parameters for $F(y)$	
Parameter	Value
B	7,957,600
y_{min}	11,000
y_{max}	140,000
mean	60,000
mode	30,000

Table 1: Input parameters for the household distribution

$$z(x) = LN_{\mu,\sigma}(x) - LN_{\mu,\sigma}(x_{min}) \tag{18}$$

Finally, the distribution is rescaled such that the cumulative amount at the upper cutoff equals the desired total (N):

$$\tilde{z}(x) = \frac{N}{z(x_{max})} z(x) \tag{19}$$

$\tilde{z}(x)$ is then the calibrated cumulative distribution used in the model.

2.2.1 Income distribution

I now apply this procedure to the Dutch income distribution. There are 7,957,600 households in the Netherlands (WoON, 2025b). The mean disposable household income is approximately €60,000 per year, while the mode is approximately €30,000 (Centraal Bureau voor de Statistiek, 2026). I set the minimum income at €11,000 and the maximum income at €140,000 per year.

These cutoff values are somewhat arbitrary, but they have only a limited effect on the results. Households below the critical income do not occupy a dwelling in the model, so the lower cutoff affects prices only through its influence on the critical income. The upper cutoff matters through the assignment process. It is therefore important that the maximum income is consistent with the maximum housing quality in the model: a household with the highest income should plausibly be matched to a dwelling with the highest quality. I return to this consideration when discussing the calibration of the housing quality distribution.

All input parameters for the household distribution are summarized in table 1. The resulting distribution is plotted in figure 1.

2.2.2 Housing quality distribution

We now turn to the housing quality distribution. The dwellings in this distribution represent the owner-occupied housing sector. Since 59% of all dwellings are owner-occupied (WoON,

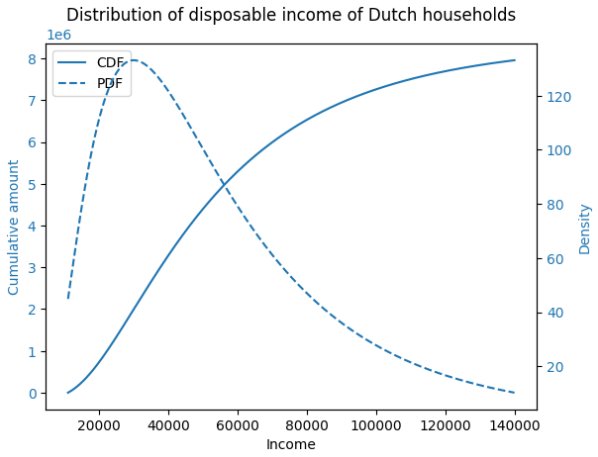

Figure 1: The cumulative and density distribution of the disposable income of households used as input for the assignment model.

Defining parameters for $G(q)$	
Parameter	Value
S	4,694,984
q_{min}	$70m^2$
q_{max}	$200m^2$
mean	$143m^2$
mode	$90m^2$

Table 2: Input parameters for the housing distribution

2025a), this corresponds to 4,694,984 houses.

Specifying the distribution over q is more challenging than for income, since housing quality is not a directly observable quantity. House prices are determined by a combination of factors, most notably size and location. In this paper, I interpret q as the physical size of a dwelling measured in square meters. This choice has the advantage of being clearly measurable, but it abstracts from other important determinants of house prices. The underlying assumption is these omitted characteristics are implicitly averaged over.

The mean house size in the owner-occupied sector in the Netherlands is $143m^2$. No data are available for the mode, so this requires an educated guess. A value of 90 m^2 appears reasonable for the owner-occupied segment. I set the lower bound of the distribution at 70 m^2 and the upper bound at 200 m^2, which is consistent with the maximum income of €140,000; That is a reasonable dwelling size for a household with that disposable income. The resulting distribution is shown in Figure 2 and Table 2.

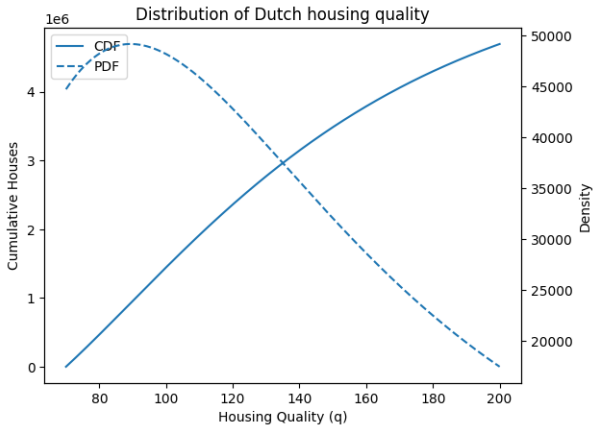

Figure 2: The cumulative and density distribution of the quality of Dutch houses used as input for the assignment model.

2.2.3 Remaining parameters

The final step in solving the model is to specify Cobb–Douglas preferences and the outside option. The outside option plays a key role in pinning down the price level by providing the marginal household with an alternative to housing consumption within the model. It can be interpreted in several ways, such as living with parents, living in social housing or private rent. These interpretations would imply different levels for both the price and the size of the outside option. However, the purpose of the outside option in the model is not to represent a specific real-world alternative, but to capture the idea that households derive less utility outside the housing market.

This interpretational ambiguity implies that the level of the outside option is not directly observable and should not be assigned a precise empirical counterpart. For instance, interpreting the outside option as the average size of a social rental dwelling would likely overstate its utility value, since it ignores the additional utility households derive from ownership.

Instead of forcing a literal interpretation on this parameter, I exploit this flexibility by calibrating both the Cobb–Douglas parameter and the outside option to best match observed features of the Dutch housing market. In figure 3, I illustrate how these parameters affect the price function. The parameter α primarily determines the slope of the hedonic price function, as can also be seen directly from equation 14. The outside option, in contrast, mainly affects the level of the price function.

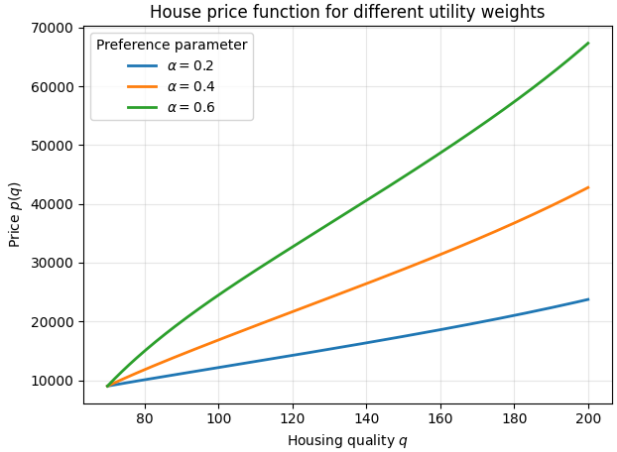

(a) Utility sensitivity

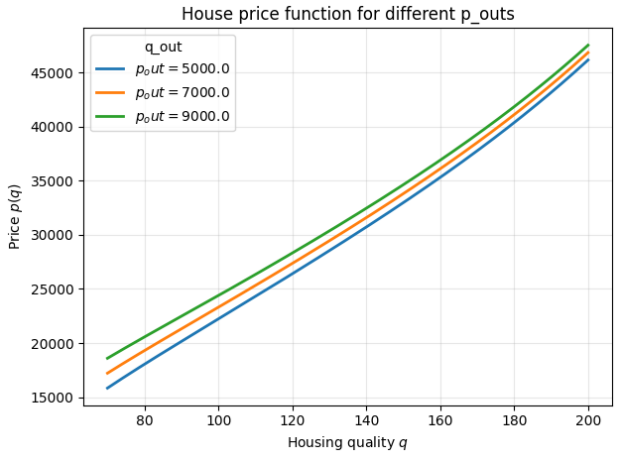

(b) Price sensitivity

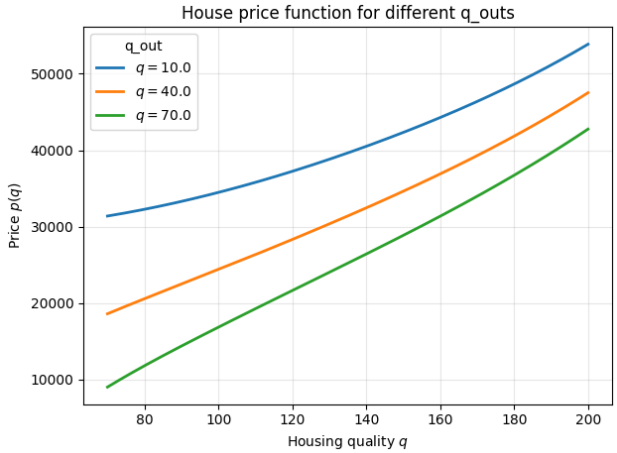

(c) Quantity sensitivity

Figure 3: Sensitivity analysis across different dimensions. In the upper diagram, $q_{out} = 70$ and $p_{out} = 9000$. In the middle diagram, $\alpha = 0.4$ and $q_{out} = 40$. In the bottom diagram, $\alpha = 0.4$ and $p_{out} = 9000$.

In the Netherlands, the average house price per square meter is approximately €4,500

(Huisverkopen.nl, 2025). Under the strong assumption that this marginal price is constant, I calibrate the model to match this value. To convert house prices into annual housing costs, I assume an interest rate of 5% and a 30-year annuity mortgage. This yields an annuity factor:

$$\text{yearly price} = \text{upfront price} * \frac{0.05}{1 - (1 + 0.05)^{-30}} \tag{20}$$

The parameter values that result from this matching procedure are reported in table 3.

Parameter	Value
α	0.5
p_{out}	7000
q_{out}	40

Table 3: Calibrated parameter values

2.2.4 Resulting model of the Netherlands

Figure 4 presents the equilibrium outcome of the model given the parameter values specified above. The upper-left panel shows the hedonic price function. The upper-right panel shows the housing expenditure of households as a function of income. The lower-left panel orders households by income on the horizontal axis and shows the housing quality they consume. The lower-right panel shows housing expenditure and income.

We observe that households in the lowest-quality segment spend more than half of their income on housing. While this may appear high, it is consistent with the observation that first-time buyers often rely on savings or intergenerational transfers in addition to income. As income increases, the housing expenditure share declines, which is consistent with empirical patterns.

I point out that no meaning should be attached to the values of the graphs at the outside option. Since the outside option does not correspond to a well-defined empirical counterpart, levels near the boundary of the model should not be given a direct structural interpretation.

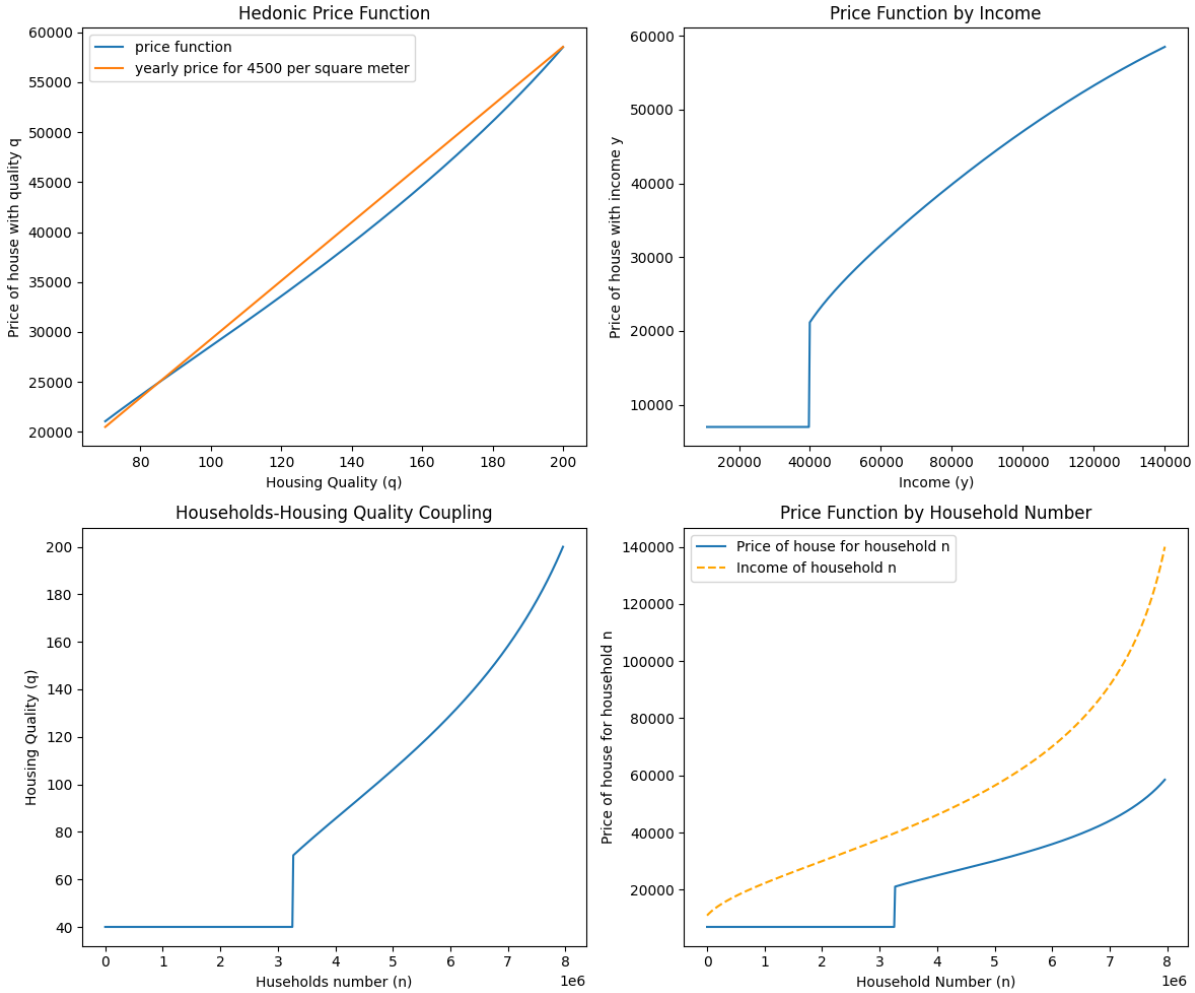

Figure 4: The outcome of the model in the baseline scenario. The left upper figure shows the predicted hedonic price function. For reference, the yearly price of a house of $4500/m^2$ is shown in orange. The right upper panel shows the price paid by houses with a certain yearly income. In the bottom two panels, households are ordered from low to high income on the horizontal axis. The bottom left panel than shows what quality these houses are coupled to. The bottom right panel shows the price that these households pay for their house. For reference, the household income is plotted in orange.

3 Results

3.1 The effects of cheap and expensive construction

With the base case calibrated to the Netherlands, we now turn to the effects of housing construction on prices and welfare. I compare two scenarios in which a fixed number of new dwellings is added to the Dutch housing stock: one in which only small houses are constructed and one in which only large houses are constructed.

Housing construction of quality q^{new} is modeled as a discrete increase in the housing distribution function at that quality:

$$G_{\text{new}}(q) = \begin{cases} G_{\text{old}}(q), & q < q_{new} \\ G_{\text{old}}(q) + N_{new}, & q \geq q_{new} \end{cases}$$

where N_{new} is the number of new houses added. Note that this definition violates the assumption that $G(q)$ is a smooth function. This could be addressed by smoothing the jump, for instance by adding a narrow logistic function. For the numerical simulations, however, this is not necessary, as the integration algorithm effectively treats the function as being smooth.

The new house prices are obtained by solving the assignment model again using the updated housing distribution function. This implies that a new equilibrium is reached after the N_{new} houses have been constructed. It is, of course, unrealistic that all 8 million households move within a given period. An argument for why we can neverthelss trust that the model provides a reasonable approximation is that households with similar incomes experience similar price effects. Consider, for example, a household with an income of €30,000 that moves. In the model, the household earning €29,999 then occupies the vacated dwelling. In reality, that household may choose not to move, and the dwelling may instead be occupied by a household earning €29,000. Since these households have similar purchasing power, the resulting price effects will even out.

Figure 5 presents the price effects of constructing 500,000 new dwellings. Cheap housing is defined as housing with a floor area of 80 m^2, while expensive housing is defined as housing with a floor area of 180 m^2.

The first thing I would like to point out is that the prices decrease by surprisingly small amounts, given that 500,000 new houses were added. This is more than the Dutch housing shortage, so one would expect prices to decrease greatly. What is happening here, however, is that all new houses are owner-occupied houses while the outside price and quality remain the same. In reality, we would also expect an increase in the outside utility, perhaps because social

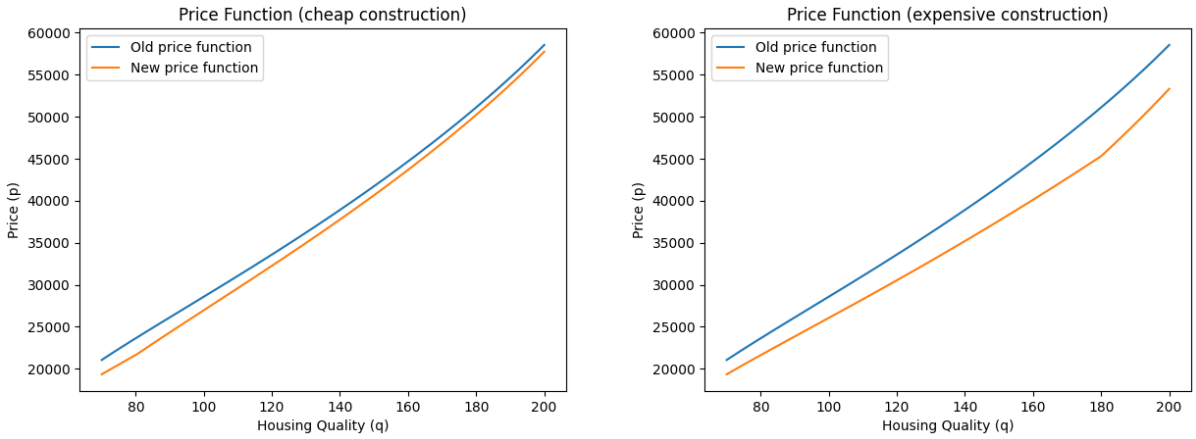

Figure 5: The price functions before and after cheap and expensive construction.

housing was improved simultaneously with the construction of owner-occupied housing. This would make prices decrease more.

The amount of 500,000 was chosen to give sizable results, but the shape of the difference curves is identical for lower amounts. Let us now try to understand the shape of these curves. In both scenarios, the largest price effect occurs at the quality level at which new housing is constructed. This is intuitively plausible: the impact on prices is strongest where the additional supply enters the market. The mechanism can be explained as follows. Every household that was previously housed below q_{new} moves up the quality ladder. Consequently, for all $q < q_{new}$, $y(q)$ is lower than before. The households now occupying these dwellings have lower incomes and therefore a lower willingness to pay. From equation 14 it follows that prices for housing below the constructed quality increase less rapidly than in the original situation.

This effect accumulates up to the point where the housing assignment becomes identical to the original assignment. Beyond that point, the new price function gradually converges to the old one. The reason is that the same households occupy the dwellings, but housing prices are now lower. As a result, these households wish to consume more housing, increasing the marginal cost of housing consumption.

Combining everything, we have the following form for the price-difference curve in the case of a fixed amount of construction: It begins at a value fixed by the outside option, declines until the point of construction, and subsequently rises again.

To understand why some price differences are larger in the case of expensive construction, it is informative to plot both difference curves in one plot. This is done in figure 6. We now see that prices decline by exactly the same amount or more across the entire quality range when expensive housing is added.

This is because the households that have moved in the scenario of cheap construction,

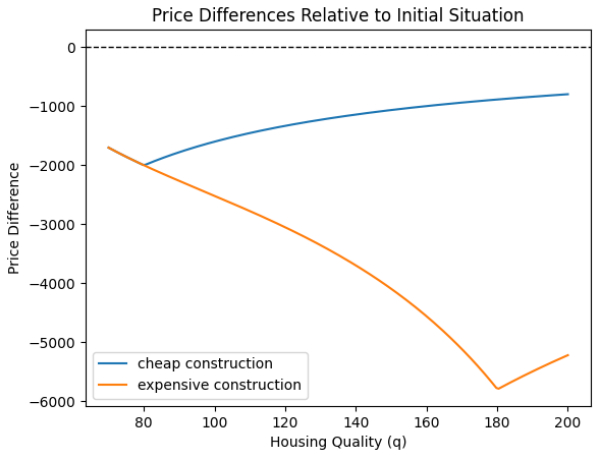

Figure 6: The price difference curves of expensive and cheap construction.

would have all moved to exactly the same house in the scenario of expensive construction. The assignment is thus exactly the same for $q < 80$. In the case of expensive construction, more households move up the quality ladder so that a larger share of dwellings becomes occupied by households with lower willingness to pay, resulting in a stronger cumulative effect.

To make this mechanism even more clear, I have constructed a toy model in figure 7. It shows a base case and two scenarios, one with cheap construction and one with expensive construction. In the old situation, the first three households are in the outside option, represented by X, and households 4 to 8 live in qualities 1 to 5. After construction, two houses are added in red. In the case of cheap construction they are of quality 3, in the case of expensive construction of quality 5. Now we can see very clearly that the situation for the first three households is the same in both scenarios. This is because everyone below the lowest-income house that occupied one of the new houses moved up two houses. It is the same houses that they can move up to.

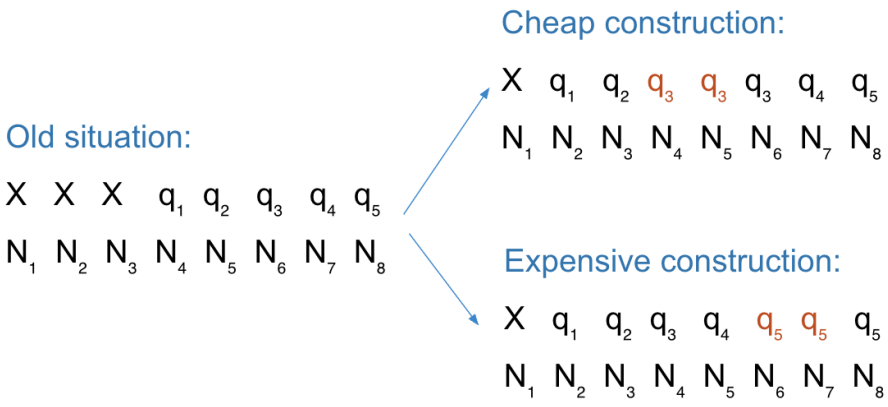

Figure 7: A toy model illustrating how low-households are indifferent between cheap and expensive construction.

Price changes across housing qualities are not the whole story, since households also move to higher-quality dwellings. It is therefore informative to examine the prices paid for housing by households after construction. These are presented in figure 8. Interestingly, more households experience a reduction in housing expenditure in the cheap-construction scenario than in the expensive-construction scenario, even though prices decline more strongly in the latter. The reason is that most households move to higher-quality housing in the expensive-construction scenario, whereas relatively few households move in the cheap-construction scenario. As a result, many households in the latter case benefit only from lower prices, not from more consumption.



Figure 8: Comparison of price paid for housing by households in the scenarios of cheap and expensive construction

3.2 Welfare effects

Assessing household welfare directly from the above results is complicated by the fact that both prices and housing consumption change. It is therefore useful to examine welfare effects directly. Two common measures for comparing scenarios are the equivalent variation and the compensating variation.

The equivalent variation asks: How much income would need to be given to a consumer in the initial situation to make them as well off as they are in the new situation? It therefore measures the consumer's willingness to pay for the change.

The compensating variation asks the reverse question: How much income would need to be given to a consumer in the new situation to make them as well off as they were initially? The two measures thus take different situations as their baseline. In general, they are not equal because the marginal utility of income differs between the two situations.

I use the equivalent variation because the objective is to compare different policy interventions. It is thus best to take the initial situation as the base scenario. Using compensating vari-

ation would require comparing the value of an additional euro under the expensive-construction scenario with its value under the cheap-construction scenario.

Suppose that we have solved for the price function $p_1(q)$ in the initial situation and $p_2(q)$ in the new situation. We compute the equivalent variation for a household with income y as follows. First, we determine the bundle consumed in the new situation, (q_2, c_2) . Next, we calculate the utility levels associated with different incomes in the initial situation, yielding a utility function $u = u_1(y)$. We then seek the income y^* that would provide the same utility in the initial situation as the household attains in the new situation, where $u_2 = u(q_2, c_2)$. To find this income, we solve

$$u_1(y^*) = u_2 \tag{21}$$

The equivalent variation then is the difference with the initial income:

$$EV = y^* - y \tag{22}$$

This procedure contains a subtle assumption. The hedonic price functions are computed using the income distribution. However, when evaluating the utility associated with a hypothetical income y^*, the price function is held fixed. In other words, although prices are determined by the aggregate income distribution, individual households are assumed to be price takers. Given that the model contains approximately 8 million households, this assumption is reasonable, as a change in the income of a single household has a negligible effect on equilibrium prices.

A second subtlety arises at the top of the income distribution. Consider the highest-income household, which occupies the highest-quality dwelling. When calculating the hypothetical income required to attain equivalent utility, the price function can no longer be used once the required housing quality exceeds the maximum quality available in the model. This is not merely a numerical issue but a conceptual one: housing of higher quality does not exist, so the richest household cannot further increase its housing consumption. To remain consistent with the original setting while ensuring that the hypothetical income change does not affect prices, I assume that households cannot consume more housing than is available. Any additional income is therefore spent entirely on the composite consumption good c.

Before comparing cheap and expensive construction, let us investigate how welfare effects develop for different households. Figure 9 shows what welfare effects households experience when housing of 180 m^2 is added to the housing stock.

Let's investigate the shape of the curve from left to right. The first thing to notice is

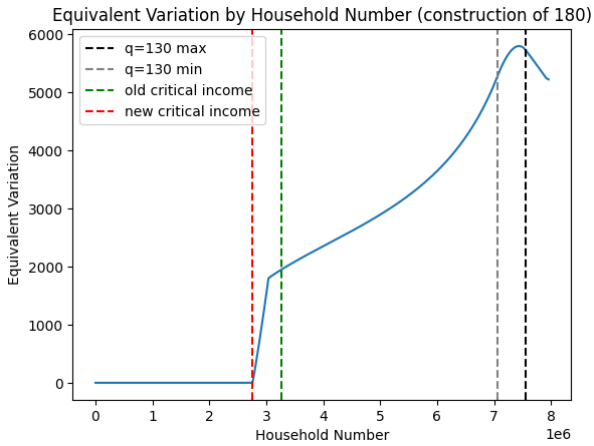

Figure 9: Equivalent variation for household number. Vertical lines are added to indicate the properties of certain households

that change in welfare at the critical income is not discrete, even though the lowest income households have made a discrete jump in housing quality from the outside option to a house. The fact that the graph increases sharply, but not discretely, is because the utility of the new critical household has not increased. This is enforced by definition, because the price of the lowest quality house is set such that the critical household is indifferent between the outside option and the house. The critical household thus experiences no increase in welfare.

After the critical income, the added welfare increases sharply until it reaches a discontinuity in its derivative. The precise structural cause of this directional change remains an area for further investigation.

After the bump, the equivalent variation continues to increase. At one point, we reach the lowest-income household that got one of the new constructed houses. After this household, there are 500,000 households which consume the exact same quality. The grey dotted line indicates the lowest income household in this quality. This household used to be in a lower quality house. The black dotted line represents the highest income households in this quality. This household did not move and already had this house.

In between these dotted lines, the welfare first increases and then decreases again. This can be explained because there are two effects with opposite effect. The first effect is that as we move to higher income households, the difference between the willingness to pay and the price paid increases. The opposite effect is that as the incomes in this regime get higher, the inefficiency of the consumption bundle becomes greater. The last household would much prefer to use the price decrease to consume more housing than spend less on housing. The interplay of these effects explains the shape of the equivalent variation curve between the dotter lines.

After the black dotted line, all households are assigned to the same house as before the

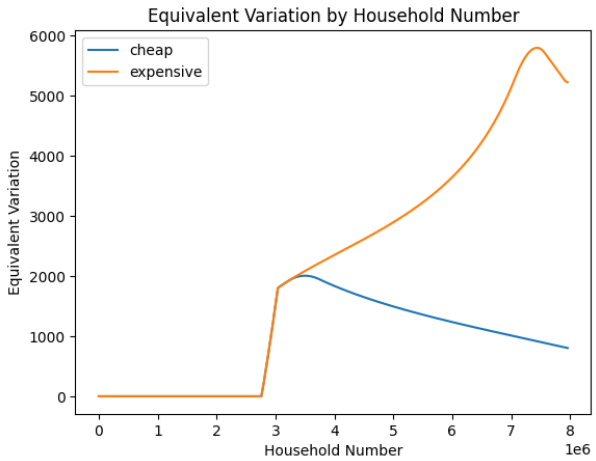

Figure 10: The equivalent variation for expensive and cheap construction.

construction. All welfare increase for these households thus comes from increased other consumption due to the decrease in house price. As the difference in price decreases as quality grows in this regime, the welfare increase also decreases.

Now that we understand the form of the equivalent variation curve, we move to comparing the scenarios of cheap and expensive construction. Figure 10 presents the equivalent variation for households under both construction scenarios. We once again see that both scenarios have exactly the same effect for lower income households. Expensive construction, however, has a greater welfare effect for richer households.

We can also quantify the aggregate welfare effects by integrating the equivalent variation over the population. This yields the following results:

$$\text{Total welfare increase for cheap construction} = 7.1 \times 10^9 \tag{23}$$

$$\text{Total welfare increase for expensive construction} = 17.7 \times 10^9 \tag{24}$$

The construction of 500,000 expensive houses thus causes twice as much welfare increase as the construction of the same amount of cheap houses. It is, however, important to note that the whole difference comes from the extra welfare increase of rich households. Poor households are indifferent between both scenarios.

Policy makers that are interested in helping poor households, should therefore not only look at the total welfare increase. Instead, they should also consider the opportunity costs of constructing certain types of housing. This will be addressed in the next sections. First, I investigate the increased welfare per square meter and afterwards I consider construction costs.

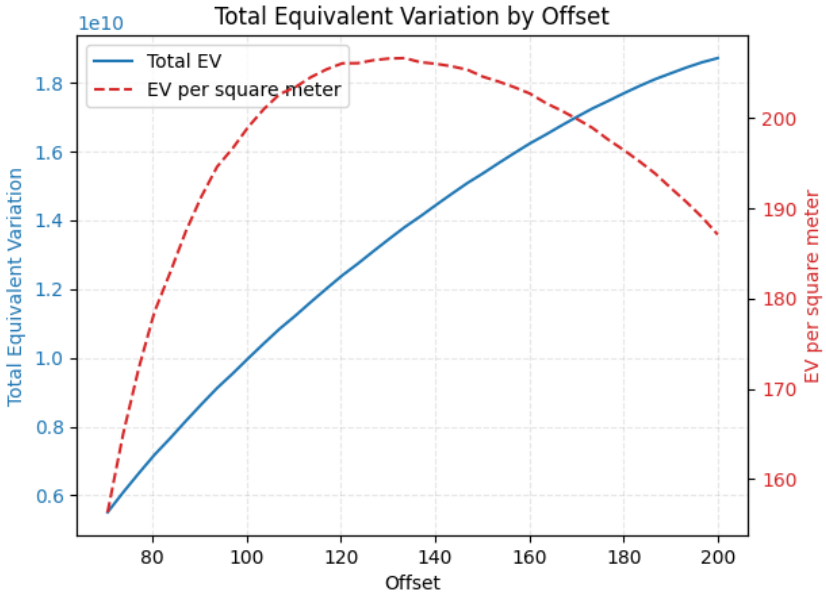

Figure 11: Total equivalent variation per house and per square meter for many different qualities constructed.

3.3 Welfare effect for any quality

So far, we have researched the welfare increase per house. However, one might argue that comparing the total increase in welfare between scenarios is no right metric. After all, when we build 500,000 large houses we add much more quantity of housing to the market than when we add 500,000 small houses. If in reality 2 houses of $100m^2$ can be made instead of 1 house of $200m^2$, then a better metric would be the welfare increase per square meter of housing added.

This is addressed in figure 11. The horizontal axis plots the quality which is constructed. The vertical axes represent the total welfare increase and the welfare increase per added square meter. We see that the total amount of welfare keeps increasing for larger houses, which is in line with the previous results. The second derivative of this function, however, is negative. This makes that the welfare increase per square meter constructed actually peaks at $q = 133$.

We thus come to the conclusion that if we are trying to maximize welfare per house, it is best to build expensive houses and that if we are trying to maximize welfare per added square meter, it is best to build houses of $q = 133$. Which metric a policy maker should maximize, depends on what determines how much housing is constructed. This will be discussed in the next section by considering the producer side of the housing market.

4 Including a supply side to the model

So far, we have altered housing supply without considering who actually builds these houses. In the Netherlands, however, the government does not directly construct owner-occupied housing. Instead, housing production is delegated to a regulated market in which private developers play a central role. As a result, policy makers cannot simply choose which types of houses should be built; rather, policy influences the market outcomes.

In this section, I extend the assignment model by introducing a supply side. A cost function determines the cost of constructing a house of quality q. Given an existing housing stock and the corresponding price function, developers decide where they prefer to build by comparing expected profits across housing types. Naturally, developers will choose to construct houses for which profits are highest.

4.1 Modeling Dutch construction costs

I model construction costs using a linear cost function. For most housing types, a construction cost of approximately €2,000 per square meter is realistic (iTX BouwConsult, 2025). Next to the cost per square meter, every house also has a fixed cost. This fixed cost represents development expenses that do not depend on housing size, such as planning costs, negotiations with local authorities, permit procedures, and other administrative expenses. I assume a fixed cost of €70,000 per dwelling. Since all prices in the model are expressed in annual terms, these values are converted using the same annuity factor as in equation 20.

The resulting cost function for constructing a house of size q is then:

$$c(q) = 130 * q + 4550 \tag{25}$$

Figure 12 shows the housing price function introduced earlier, together with the cost and implied profit function. The figure demonstrates that profits increase with housing size. This is broadly consistent with observations from the Dutch housing market, where developers often report that larger dwellings generate higher profits. Indeed, the average size of newly built owner-occupied housing in the Netherlands was approximately 184 m^2 in 2025 (Centraal Bureau voor de Statistiek, 2025b), although it is difficult to determine to what extent this reflects market incentives rather than policy restrictions.

If profits exist for all housing types, one may ask why developers do not simply continue building until prices decline and profits disappear. The answer is that housing supply cannot be expanded without constraint. The construction sector faces substantial labour shortages,

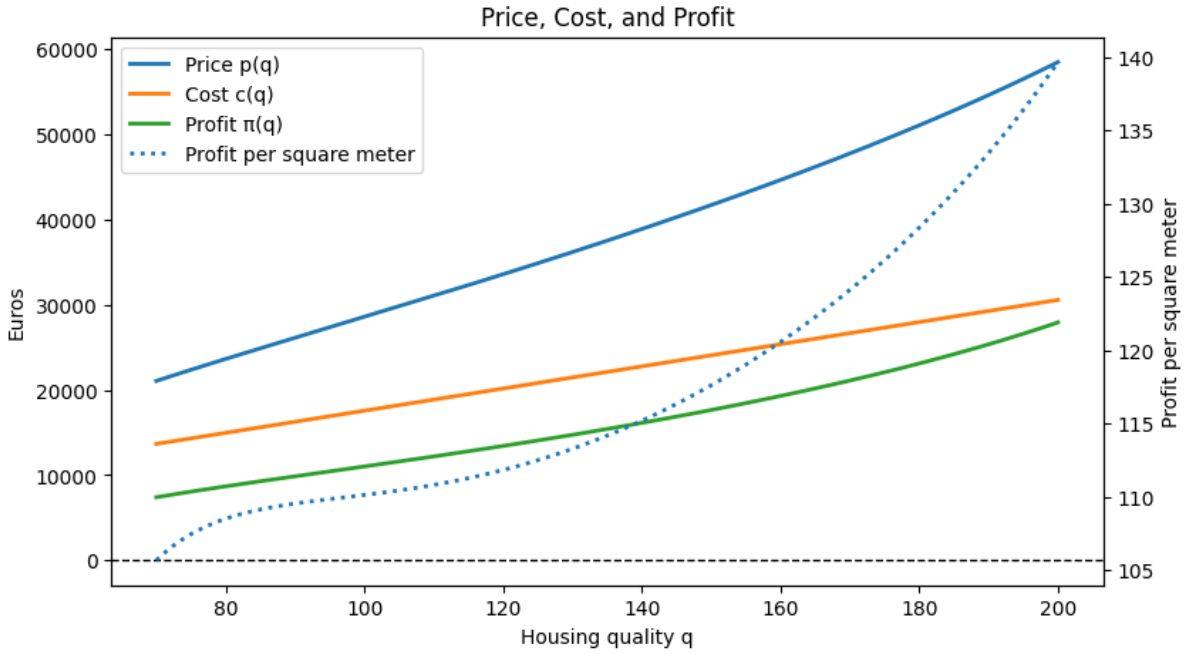

Figure 12: The price function in the base scenario, together with the cost functions and profits.

limiting the speed at which new housing can be produced. Perhaps more importantly, developers depend heavily on government actions such as spatial planning, land preparation, and permit approval. All land already has a current function that needs to change before one can build a house.

4.2 Should the government intervene in the market?

The amount of housing added to the Dutch housing market thus appears to be constrained. The key question is whether this constraint operates at the level of the number of dwellings or at the level of total floor space (square meters). This distinction is important, as it affects whether policy intervention in the market is jusitified.

If the number of houses constructed each year is fixed, no intervention is needed. Developers maximize profits by building high-quality housing, which also coincides with the construction that maximizes consumer welfare. In this case, the market outcome is efficient.

If, however, the constraint applies to the total amount of housing floor space constructed, the market outcome is inefficient. Developers continue to have an incentive to supply high-end housing, while consumer welfare is maximized at lower-quality dwellings.

Arguably, a fixed of amount of housing area is a closer representation of reality. Both construction capacity and available land are limited resources and developers frequently operate under land constraints, for example through tenders for designated construction sites with fixed surface areas.

In this case, there is a rationale for government intervention as the actions of the developers do not match what is best for the consumers. As the developers still make profit for smaller houses, the government could enforce the construction of the ideal construction size.

But what would be the social optimum, also taking profits into account? Some simple arithmetic can tell us the answer. First, let us consider two policy scenarios. One where there is a free market and all houses constructed are of quality $q = 200$. The other scenario restricts construction to $q = 133$.

The profit per square meter for houses of is 114 euros. The profit per square meter for houses of is 140 euros. The EV per meter added for $133m^2$ houses are 206. The added welfare per meter for houses of $200m^2$ is 187. Table 4 shows the welfares and profits that this translates to for 500,000 houses. Adding everything up, we see what welfare is maximized for the restricted case, with around 7% more total welfare than in the case of a free market.

	Social welfare gain	Profit	Total
Construction of only $q = 133$	985,504,000	688,896,000	1,674,400,000
Construction of only $q = 200$	894,608,000	669,760,000	1,564,368,000

Table 4: Welfare and profits for two policy scenarios

Even though houses of $q = 133$ add most welfare, there might be political reasons for the government to alllow the construction of more profitable houses too. Let's investigate how total welfare changes in that case.

Say, for instance, that x percent of all newly added housing area should be $q = 133$, and the rest is left to the free market. The added welfare is then

$$\text{welfare} = A\left[x(206 + 144) + (1 - x)(187 + 140)\right] \tag{26}$$

$$= A(350 + 23x) \tag{27}$$

Where A is the total area of housing added.

This concludes the exercise with including a supply side to the assignment model. We have seen that the market is not necessarily efficient, depending on what constraints the construction of housing. The mechanism of such a constrainnt is an empirical question, the answer to which could also differ between construction locations. My results, however, can help policy makers reason about what housing construction to demand in certain places.

5 Discussion

Let us revisit the policy question about whether to build expensive or cheap houses. In the model that I have presented, all households move when new dwellings are added to the housing stock. This makes that for the lowest income households, it does not matter whether cheap or expensive houses are added. Everyone that can move up, moves up, so if you more expensive houses are added, more people move up in the housing quality ladder. This makes that the overall decrease in price is largest for expensive construction and that the biggest welfare increase per house is achieved for expensive housing.

The implication for policy depends on who you want to help and what the reality of housing production is. If the aim is to house poorer households and there is a limited amount of housing area added, then smaller houses are more effective. If the aim is to maximize welfare and there is a fixed amount of houses added, more expensive housing is better.

However, the exact results of the model presented should not be directly translated into policy, as the model presents a highly stylized version of reality. Some of the main limitations are:

1. The housing stock and household population are only heterogeneous in one parameter. In reality, housing quality and household preferences are determined by many factors. The most notable dimension left out is location, but there are many more like accessibility, age of the house, and building style.

2. There is only one outside option. In reality, there are many outside options that also differ in quality. A high-income household could, for instance, also choose to rent a high quality house in stead of buying an owner-occupied house.

3. We consider one closed market. Housing construction does not cause households to move. This is connected to both former limitations. Our model does not allow people to enter the market from somewhere else (location) or from another type of housing (rental). All we can do is exogenously increase the household population

4. The model predicts perfect moving chains. The construction of an expensive house causes all households to move up, thereby creating a cheap house to free up.

The last limitation is the most important, as perfect moving chains are central to the results of this paper. This section discusses how this simplification departs from reality and what implications this has for the findings. In doing so, all other mentioned limitations are also addressed, as they are closely related to the functioning of moving chains.

5.1 Moving chains

In reality, new expensive housing does not always cause moving chains that reach low income households. In the Netherlands, for instance, construction of expensive housing hardly frees up space in the social housing sector (ABF Research, 2025a). Similar findings have been reported for the United States (Mast, 2023) and Helsinki (Bratu et al., 2023). The latter found that for every 100 units of new market-rate housing, only 35 houses free up in lower-quintile neighbourhoods.

Why do moving chains break? The assignment model helps us reason about this. As the two inputs to the model are the housing stock and the household population, one of them must change. On the supply side, a chain breaks when a house in the chain leaves the owner-occupied housing stock. This could be because it is destroyed, it is bought to let out or becomes a second house. On the demand side, the breakage of a moving chain must mean that there was a growth in the household population.

Let us focus on this demand side mechanism of chain breakage. The easiest way to look at this, is as an exogenous population increase, for instance because of migration or because of households splits. In the assignment model, an exogenous increase of non-poor households will always work against the goal of making houses available for poorer households. If you construct expensive housing, the new households will break the moving chains. If you construct cheap housing, the new households will crowd out other households, which then move into the new cheap housing. Exogenous population increase thus does not influence my results.

We could, however, also consider the breakage of moving chains by endogenous population increase. What I mean with that, is that it is the construction of new houses which causes the population increase. Households would not enter the market into a lower-quality house, but will move into a high-quality house when it becomes available.

The most obvious way this could happen, is when households move in from another region. Regional moving is part of the reason why Bratu, Harjunen and Saarimaa find that moving chains break in Helsinki: they only consider houses that are freed up within Helsinki and not through the whole of Finland.

Next to immigration, endogenous relocation can also be caused by households that in a different way are not part of the market under consideration. It could be rich households which do not split unless there is a high-quality house available or households currently renting luxurious apartments.

If such endogenous population increases occur often, it could alter our results substantially. In that case, expensive construction causes a population increase which breaks moving chains

to break. In this case, less houses free up for low-income households. In contrast, cheap construction does not create the population increase and would thus reach low-income households.

It would be very interesting to understand the size of such effects. The policy relevance of it, however, requires a careful analysis of who your policy is aimed at. A Helsinki policy maker might take the result that market-rate housing does not reach cheap neighbourhoods as a reason to build cheap housing, but a national policy maker might come to very different conclusions. A good illustration of this is the paper by Thiel and Zahnbrecher (2024). They conclude that the Dutch government can best stimulate construction in Amsterdam and not in the countryside, exactly because it will cause people from the countryside to move to Amsterdam, leaving a house there vacant. Construction in Amsterdam thus also helps households in the countryside. What is the breakage of a moving chain to one policy maker is a long moving chain to the other.

5.2 Policy should aim for long moving chains

Given that moving chains can break, it is not the numerical outputs of my simulations that should be taken as the main lesson. Instead, it is the mechanism which causes those outputs that should be the take-home message. The main insight is that long moving chains cause most welfare per house.

Often, the parameters which predict long moving chains will also be strong predictors of whether the house is expensive. Research by the Amsterdam Bureau of Statistics showed that larger houses cause larger moving chains (Booi, 2022). Similarly, the conclusion by Thiel and Zahnbrecher (2024) that building in expensive locations is efficient can be connected to longer moving chains. Size and location are strong predictors of price.

But if something is very expensive, there will often also be high costs to constructing the house. For large houses it is the construction costs and for expensive locations it is the opportunity cost of the use of the location. In the case of size, we have seen in my model that the welfare increase per square meter is actually maximized for a shorter moving chain.

What would therefore be perfect, is to find characteristics of a house which are less strong predictors of its price, but are highly correlated with long moving chains. The canonical example is housing that is accessible for elderly people. Elderly occupy the largest homes per household member but often indicate they want to move (Biesenbeek and Rupert, 2025). Building homes that they can move into would thus also create large moving chains. Elderly people perhaps do not care about size as much, but value accessibility and the availability of certain amenities.

What is happening here, is that elderly people attach more value to something that many others value less (accessibility), while owning something that others value highly (a large house).

Here we see why quality should be a vector and not a single parameter: The utility function of elderly people values the quality vector differently than other utility function.

It would be interesting to understand whether there are more such inefficiencies in the current housing stock: households which undervalue the house they currently possess. What type of construction could make them move?

I think that these are very relevant questions for policy and require more research into moving chains. What causes long moving chains? When does endogenous chain breakage occur? And what types of households would move if there was a house available?

Such research requires moving away from simplifying assumptions like homogenous preferences and a single quality parameter. One could imagine a highly dimensional assignment model, but given the many parameters and the fact that the Dutch housing market actually consists of many different regimes (social housing, controlled rent, free rent and owner-occupied), I do not think that an analytical equilibrium model is very promising. Instead, I would propose an agent-based model. Modeling what types of households live where and giving them some heuristics for when they will move, might answer the question how many people will actually move in reality.

Such a model could greatly help policy makers in making realistic choices about what to build where. Ideally, it could help policy makers decide up to a highly local level what kind of housing is needed. The hope would be that a high-resolution representative agent model might predict that one neighbourhood needs more elderly homes, while in the other longer moving chains would be caused by family homes.

5.3 Conclusion

Wrapping up, I conclude that it is exactly the simplicity of the assignment model which makes it very insightful for thinking about who profits from construction. The model makes it tractable how outcomes are influenced by the fact that housing is an indivisible stock good. By not having to consider many different parameters, it helps to make clear that expensive construction has the same effect on poor households as cheap construction. The difference between cheap and expensive construction is that higher income households also profit in the latter. That is why expensive construction causes more welfare increase per house.

The main conclusion that we should draw from this result is that longer moving chains cause more welfare per house. Translating this insight into exact policy proposals requires moving beyond the one-dimensional cheap-expensive axis. In particular, a careful analysis of why moving chains break and which kinds of new housing create long moving chains is needed

for a better determination of what housing is needed.

References

Määttänen, N., & Terviö, M. (2014). Income distribution and housing prices: An assignment model approach. *Journal of Economic Theory*, *151*, 381–410.

Planbureau voor de Leefomgeving. (2017). *Middeninkomens op de woningmarkt: Ruimte op een krap speelveld*.

Nathanson, C. G. (2019). *Trickle-down housing economics* (Working Paper). Northwestern University. Retrieved June 26, 2026, from https://charlesnathanson.com/s/TrickleDownHousingEconomics.pdf

Booi, H. (2022). *Verhuisketens van de nieuwbouw in 2020*. Onderzoek en Statistiek Amsterdam. https://onderzoek.amsterdam.nl/publicatie/verhuisketens-van-de-nieuwbouw-in-2020

Bratu, C., Harjunen, O., & Saarimaa, T. (2023). Jue insight: City-wide effects of new housing supply: Evidence from moving chains. *Journal of Urban Economics*, *133*.

Mast, E. (2023). The effect of new market-rate housing construction on the low-income housing market. *Journal of Urban Economics*, *133*. https://www.sciencedirect.com/science/article/abs/pii/S0094119021000656

Rouwendal, J., Sniekers, F., & Jia, N. (2023). *Borrowing constraints, housing tenure choice and buy-to-let investors: An assignment model* (Tinbergen Institute Discussion Paper No. TI 2023-005/V). Tinbergen Institute. Amsterdam, Rotterdam.

Sniekers, F. (2024). Nederland bouwt vooral dure woningen. *Economisch-Statistische Berichten*, *109*(4840), 536–537. https://esb.nu/536-537_sniekers2/

Thiel, J., & Zauhnbercher, H. (2024). Betaalbaar wonen vraagt om bouwen op de duurste plekken. *Economisch-Statistische Berichten*. https://esb.nu/betaalbaarder-wonen-vraagt-om-bouwen-op-de-duurste-plekken/

ABF Research. (2025a). *Bouwmix, doorstroming en betaalbaarheid*.

ABF Research. (2025b). Woningtekort dashboard [Accessed: 2026-06-18]. https://primos.abfresearch.nl/dashboard/dashboard/woningtekort

Biesenbeek, C., & Rupert, N. (2025). *Doorstroming op de nederlandse woningmarkt*. De Nederlandsche Bank. https://www.dnb.nl/publicaties/publicaties-onderzoek/analyse/doorstroming-op-de-nederlandse-woningmarkt/

Centraal Bureau voor de Statistiek. (2025a, January). *33 duizend mensen dakloos begin 2024*. Retrieved June 26, 2026, from https://www.cbs.nl/nl-nl/nieuws/2025/05/33-duizend-mensen-dakloos-begin-2024

Centraal Bureau voor de Statistiek. (2025b). Cbs statline 83673ned. https://www.cbs.nl/nl-nl/cijfers/detail/83673NED

Huisverkopen.nl. (2025). Gemiddelde vraagprijs per m2 [Accessed: 2026-06-18]. https://www.huisverkopen.nl/blog/gemiddelde-vraagprijs-per-m2

iTX BouwConsult. (2025). Gemiddelde bouwkosten per m2 [Accessed: 2026-06-26].

WoON. (2025a). Kerncijfers [Accessed: 2026-06-18]. https://woononderzoek.nl/mosaic/woon-2/1--kerncijfers

WoON. (2025b). Woningkwaliteit en woonbeleving [Accessed: 2026-06-18]. https://woononderzoek.nl/mosaic/woon-2/6--woningkwaliteit-en-woonbeleving

Centraal Bureau voor de Statistiek. (2026). Beschikbaar inkomen huishoudens 2,7 procent hoger in 2025. https://www.cbs.nl/nl-nl/nieuws/2026/13/beschikbaar-inkomen-huishoudens-2-7-procent-hoger-in-2025

6 AI use

I have used ChatGPT and GithubCopilot for writting snippets of code, Latex formatting and have used it to improve my texts.