Erasmus University Rotterdam

Master Thesis

Optimal pension fund allocation in the presence of housing

Author:

Victor Jong

University Supervisor:

Dr. J.W.N. Reuvers

Second Assessor:

Dr. H.J.W.G. Kole

April 30, 2022

Abstract

Primary homes compose a large part of household assets in the Netherlands and can be a major source of household financial risk. We show that pension funds can help manage this risk and increase household welfare by implementing an investment strategy that differentiates between home owners and renters. The advantage of differentiation grows when households live in homes beyond their optimal risk profile and shrinks when households do not have precautionary savings. We further find that the historical growth in house prices is unlikely to be sustained in the future because historical returns can in part be attributed to market cyclicality.

1 Introduction

According to the Dutch Central Bureau of Statistics (CBS), house prices in the Netherlands have risen 63.5% since their lowest point in June 2013 (CBS, 2021a). In April 2021, Dutch house prices recorded their highest year-on-year growth in 20 years. The rising prices are frequently attributed to a shortage in housing supply and cheap credit, and can cause problems on two fronts. On one hand, those that are unable to buy a house may find themselves at the lower end of increasing inequality between those who can and cannot afford a home (Boelhouwer, 2020). On the other hand, those who do buy a house are more likely to take on large amounts of debt to finance the purchase (Fischer and Stamos, 2013), which could cause problems if prices were to fall substantially. The impact of changing housing prices is amplified by how much of household wealth is stored in primary homes. In 2019, primary homes accounted for 1.4 trillion out of 2.5 trillion euros in household assets (CBS, 2021b).

The other large store of value in the Netherlands is pension wealth. The total pension fund assets in the Netherlands, which are excluded from household assets, added up to over 1.5 trillion euros in 2019 (DNB, 2021). Given that these two pillars account for such a large portion of wealth and are both subject to substantial fluctuations in asset value, it seems plausible that the optimal investment allocation in the one depends on the other. Specifically, households that own their house, often with a substantial mortgage, are already subject to a lot of market risk. These households may prefer their pension assets to be invested with less risk than they would have otherwise. Similarly, renting households that do not benefit from an expected rise in property price may desire a higher expected return on their pension money. If pension wealth is invested identically for renters and home owners, then it is likely that the investment allocation is sub-optimal for at least one of the two parties.

In this research, we investigate how optimal pension investments vary between home owners and renters and assess the benefits of a pension fund allocation that differentiates between these two groups. We construct a model that describes the dynamics of the Dutch housing market in combination with other core economic variables. We then use this model to simulate households in the Dutch economy and optimise their pension investments.

The first component in our research is a model for the economy. We calibrate a model that includes bonds, stocks, inflation, and housing on the Dutch economy. Van Hemert et al. (2006) and De Jong et al. (2008) have previously calibrated economic models in the presence of housing. However, they assume deterministic term structures, while Sangvinatsos and Wachter (2005) and Koijen et al. (2010) show that time-varying risk premia substantially affect investment strategies. Furthermore, they use data from the US economy, so their results can not necessarily be extrapolated to the Dutch economy. For research into the Dutch economy, the KNW model (Koijen et al., 2010) has become the workhorse model used by Dutch institutions such as De Nederlandsche Bank (DNB) and the Netherlands Bureau for Economic Policy Analysis (CPB). The KNW model of the economy consists of two latent factors that drive the dynamics of the time-varying bond market, the stock market and a price index. We expand the KNW model by adding a housing asset. Because this has not previously been done for Dutch data, we first calibrate the model, which we tailor to the parameter restrictions laid out by the Committee Parameters (2019). We further add a time-varying risk premium for housing returns to investigate if it can help model cycles in the housing market. This builds on previous research by Fischer and Stamos (2013) and Corradin et al. (2014), who analyse cyclicality in the US housing market by constructing models with simplified dynamics of the other asset returns. We propose an adaptation of the Fischer and Stamos (2013) model that only requires us to estimate 1 additional variable for the introduction of a dynamic risk premium.

We find that housing investments have historically been well-rewarded, but that future returns are expected to be lower. Our estimates for long-run returns from house value appreciation are below the risk free rate. This does not conform with the recent trend of double digit increases in house prices every year, but it is in line with past calibrations of the US housing market by Van Hemert et al. (2006) and De Jong et al. (2008). We show that the continued growth in house value can partly be attributed to cycles in the housing market. The addition of a time-varying risk premium reduces residual variance of the fitted housing market by 30%.

The second component in our research is a model for household behaviour. We investigate how households make decisions for consumption and housing, and how that affects their optimal pension investments. This relates our research to the large literature on life-cycle investing, an area that considers how people should invest their funds based on their changing income and preferences throughout their life. The extraordinary size of Dutch pension funds has led to its own line of research on how to invest these funds, examples of which are Chen et al. (2019) and Metselaar et al. (2020). This research generally does not consider housing in the models. On the other hand, there is a class of research that studies housing decisions, which generally assumes that households manage their own retirement assets rather than through a pension fund. Cocco (2005) find that introducing housing in the model decreases investment in other financial assets. Yao and Zhang (2005) also considers the option of renting a house and find significant welfare losses when households are permanent home owners or renters, rather than choosing based on the circumstances. They also find that when investors own a house, they decrease their equity holdings as a fraction of their net worth but actually increase their equity holdings as a fraction of their financial assets, because of the diversification effect. Fischer and Stamos (2013) find that a cyclical housing market significantly affects the choice of tenure and the size of the investment.

Our research bridges the gap between pension fund strategies and housing decisions by constructing a household model that is tailored to the Dutch environment. We first optimise household decisions on consumption and housing by using backward induction. This allows us to find the optimal decisions for every financial position of the households. We then optimise the pension fund investments and investigate how these should incorporate the tenure of a household.

We find that when households own houses that exceed their risk profile, pension differentiation based on housing tenure can yield an increase in welfare of up to 2.5%, a yearly increase in certain consumption of €600 for the median Dutch household. This result is particularly interesting in light of the current housing market, where many people maximise their mortgage to have a chance at buying a home. In a setting where households structurally buy the biggest house they can afford, the optimal pension fund allocation into risky assets is 49% point higher for renters than for owners. However, this difference decreases when households don't hold personal savings, in which case a house serves as a buffer that renters do not have.

Our research confirms that the investment decisions for housing and retirement fund are connected in the Dutch economy. Furthermore, our economic model helps to understand the dynamics of property prices and can be used in future research into the impact of policy changes.

The structure of this paper is as follows. Section 2 considers the economic model and its calibration on the Dutch economy, including an analysis of the resulting parameters. In Section 3 we construct the household model and discuss how we optimise household behaviour and pension investments through a simulation. We discuss the results of this simulation in Section 4. Finally, Section 5 concludes.

2 Economic model

To generate simulations, we construct a model of the economy that captures the most relevant aspects of the economy: the bond market, a stock index, inflation and the housing market. The model we use is closely related to the work by, among others Brennan and Xia (2002) and Koijen et al. (2010). We use uppercase letters for nominal variables and lowercase letters for their real counterparts, which are adjusted for inflation.

2.1 Model dynamics

2.1.1 Two-factor model

The dynamics of the term structure are modeled by two state variables, $\boldsymbol{x}_t = (x_{1t}, x_{2t})'$, which are assumed to be mean-reverting around zero to accommodate for first-order autocorrelation in the interest rate and expected inflation:

$$\mathrm{d}\boldsymbol{x}_t = -\boldsymbol{K}\boldsymbol{x}_t\,\mathrm{d}t + [\boldsymbol{\mathcal{I}}_{2\times 2}\ \mathbf{O}_{3\times 2}]\,\mathrm{d}\boldsymbol{z}_t,$$

with $\boldsymbol{z}_t \in \mathbb{R}^5$ a vector of independent Brownian motions driving the uncertainty in the economy. As advocated by Dai and Singleton (2000) we normalise $\boldsymbol{K}$ to be lower triangular, ensuring that the factors do not rotate. The states $\boldsymbol{x}_t$ dictate the movement of two hidden variables: the instantaneous nominal interest rate R_t and the instantaneous expected rate of inflation π_t. Both are affine in the factors:

$$R_t = \delta_{0R} + \boldsymbol{\delta}_{1R}'\boldsymbol{x}_t, \qquad \pi_t = \delta_{0\pi} + \boldsymbol{\delta}_{1\pi}'\boldsymbol{x}_t, \qquad\qquad \delta_{0R} > 0,\ \ \delta_{0\pi} > 0.$$

To derive bond prices we model the nominal stochastic discount factor ϕ_t^{N} as

$$\frac{\mathrm{d}\phi_t^{\mathrm{N}}}{\phi_t^{\mathrm{N}}} = -R_t\,\mathrm{d}t - \boldsymbol{\lambda}_t'\,\mathrm{d}\boldsymbol{z}_t,$$

where the price of risk $\boldsymbol{\lambda}_t$ is equal to

$$\boldsymbol{\lambda}_t = \boldsymbol{\lambda}_0 + \boldsymbol{\Lambda}_1\boldsymbol{x}_t \text{ with } \boldsymbol{\lambda}_t, \boldsymbol{\lambda}_0 \in \mathbb{R}^5 \text{ and } \boldsymbol{\Lambda}_1 \in \mathbb{R}^{5\times 2}.$$

For future reference we define $\tilde{\boldsymbol{\lambda}}_0$ as the first 2 elements of $\boldsymbol{\lambda}_0$ and $\tilde{\boldsymbol{\Lambda}}_1$ as the upper 2×2 matrix of $\boldsymbol{\Lambda}_1$. The price of a zero coupon bond with a payout at $t+\tau$ is exponentially affine in the state variables $\boldsymbol{x}_t$: $P_t(t+\tau) = \exp\big(a(\tau) + \boldsymbol{b}(\tau)'\boldsymbol{x}_t\big)$, such that annual yields are given by

$$y_t(\tau) = -\frac{a(\tau)}{\tau} - \frac{\boldsymbol{b}(\tau)'}{\tau}\boldsymbol{x}_t.$$

We use the analytical expressions derived by Muns (2015) for $a(\tau)$ and $\boldsymbol{b}(\tau)$. The other observable processes are the price index, the stock index and the housing market. The price index, Π_t, evolves as follows:

$$\frac{\mathrm{d}\Pi_t}{\Pi_t} = \pi_t\,\mathrm{d}t + \boldsymbol{\sigma}_\Pi'\,\mathrm{d}\boldsymbol{z}_t, \qquad \Pi_0 = 1,$$

and the stock index, S_t, according to

$$\frac{\mathrm{d}S_t}{S_t} = (R_t + \eta_S)\,\mathrm{d}t + \boldsymbol{\sigma}'_S\,\mathrm{d}\boldsymbol{z}_t, \qquad S_0 = 1.$$

As the final process we introduce the housing market through the price of one housing unit Q_t. We adopt the equation from (Van Hemert, 2010, p. 473):

$$\frac{\mathrm{d}Q_t}{Q_t} = (R_t + \eta_Q - r^{\mathrm{imp}})\,\mathrm{d}t + \boldsymbol{\sigma}'_Q\,\mathrm{d}\boldsymbol{z}_t, \qquad Q_0 = 1,$$

with r^{imp} the imputed rent. This derives from the idea that the expected benefit for a landlord or a home owner is fairly compensated by the time value of money and the price of risk:

$$\mathrm{E}[\frac{\mathrm{d}Q_t}{Q_t} + r^{\mathrm{imp}}\,\mathrm{d}t] = [R_t + \eta_Q]\,\mathrm{d}t.$$

For the purposes of this paper, net rental income and imputed rent are interchangeable.

We discretise the model and its parameters to be able to calibrate them. Here we follow Pelsser (2019) in augmenting the state vector with the observed variables:

$$\tilde{\boldsymbol{x}}_t = \begin{pmatrix} \boldsymbol{x}_t \\ \ln \Pi_t \\ \ln S_t \\ \ln Q_t \end{pmatrix}.$$

This state vector then follows an Ornstein-Uhlenbeck process $\mathrm{d}\tilde{\boldsymbol{x}}_t = (\boldsymbol{\theta}_0 + \boldsymbol{\Theta}_1\tilde{\boldsymbol{x}}_t)\,\mathrm{d}t + \boldsymbol{\Sigma}_X\,\mathrm{d}\tilde{\boldsymbol{z}}_t$, which can be discretised to

$$\tilde{\boldsymbol{x}}_t = \boldsymbol{\phi}^{(h)} + \boldsymbol{\Phi}^{(h)}\tilde{\boldsymbol{x}}_{t-h} + \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \sim \mathrm{N}(\mathbf{0}, \boldsymbol{Q}^{(h)}). \tag{1}$$

Details on the discretisation process can be found in Appendix A.

2.1.2 Dynamic risk premium

As an adaptation to the model we introduce a time-varying risk premium for the housing market, which can help explain potential cycles in the data. Next to the unconditional risk premium η_Q, we add a sensitivity β_Q to some predictive signal ζ_t, such that the housing price evolves according to

$$\frac{\mathrm{d}Q_t}{Q_t} = (R_t + \eta_Q + \beta_Q\zeta_t - r^{\mathrm{imp}})\,\mathrm{d}t + \boldsymbol{\sigma}'_Q\,\mathrm{d}\boldsymbol{z}_t. \tag{2}$$

This is similar to the model from Fischer and Stamos (2013), who then let ζ_t follow an Ornstein-Uhlenbeck process with a unique drift and stochastic process. However, adopting their method would introduce various unknown parameters, which is undesirable in our already complex model. Instead, we propose dynamics for the drift and shocks in ζ_t that are mostly governed by the previously defined parameters.

First, we note that Fischer and Stamos (2013) find a correlation of more than 98% between shocks in the predictive signal and shocks in the housing market. As such, we propose to model the predictive signal as a univariate Ornstein-Uhlenbeck process with identical shocks to the housing market:

$$\frac{\mathrm{d}\zeta_t}{\zeta_t} = -k\,\mathrm{d}t + \sigma_Q'\,\mathrm{d}\boldsymbol{z}_t,$$

which corresponds to the AR(1) process $\zeta_t = e^{-kh}\zeta_{t-h} + \varepsilon_{t,5}$. In this setting, we do not need to estimate additional parameters for the shocks in ζ_t.

As a second measure to reduce the number of parameters to estimate, we let the mean reversion be equal to the senstivity from (2): $\beta_Q = e^{-kh}$, such that ζ_t evolves as

$$\zeta_t = \beta_Q\zeta_{t-h} + \varepsilon_{t,5}. \tag{3}$$

We can show that this leads to an intuitive interpretation of the predictive signal. For this we note that the discrete housing price dynamics, as indicated by the 5th row of the discretisation from (1), is revised when adding the dynamic risk premium:

$$\ln Q_t = \boldsymbol{\phi}_5^{(h)} + \boldsymbol{\Phi}_{5\boldsymbol{\cdot}}^{(h)}\,\tilde{\boldsymbol{x}}_{t-h} + \beta_Q\zeta_{t-h} + \varepsilon_{t,5}. \tag{4}$$

Here $\boldsymbol{\phi}_5^{(h)}$ and $\boldsymbol{\Phi}_{5\boldsymbol{\cdot}}^{(h)}$ denote the 5th row of $\boldsymbol{\phi}^{(h)}$ and $\boldsymbol{\Phi}^{(h)}$ respectively. Combining (3) and (4) then provides the definition for the predictive signal:

$$\zeta_t = \ln Q_t - \boldsymbol{\phi}_5^{(h)} - \boldsymbol{\Phi}_{5\boldsymbol{\cdot}}^{(h)}\,\tilde{\boldsymbol{x}}_{t-h},$$

which can be interpreted as the change in house price that could not be explained by the unconditional drift or the factor dynamics. Because the signal follows a mean-reverting process, it can help explain cyclicality in the data. Furthermore, with this definition of ζ_t we only need to estimate one additional parameter, k, when adding a dynamic risk premium.

2.1.3 Restrictions

We assume the unconditional risk premia and the imputed rent to be constant. Because the drift must be 0 under the $\mathcal{Q}$-measure, the risk premia are restricted to

$$\boldsymbol{\sigma}_S' \boldsymbol{\lambda}_t = \eta_S, \qquad \boldsymbol{\sigma}_Q' \boldsymbol{\lambda}_t = \eta_Q - r^{\text{imp}}.$$

We follow Muns (2015) and Committee Parameters (2019) in adding three additional restrictions. The fixed values 5.6% and 1.9% are imposed on the unconditional expected returns for the stock index and inflation. The Ultimate Forward Rate (UFR) is set to 2.1%. The restrictions that follow are:

$$\begin{aligned}
\eta_S &= \ln(1.056) - \delta_{0R} + \frac{1}{2}\boldsymbol{\sigma}_S' \boldsymbol{\sigma}_S, \\
\delta_{0\Pi} &= \ln(1.019) + \frac{1}{2}\boldsymbol{\sigma}_\Pi' \boldsymbol{\sigma}_\Pi, \\
\delta_{0R} &= \ln(1.021) + \boldsymbol{\lambda}_0' \boldsymbol{b}_\infty + \boldsymbol{b}_\infty' \boldsymbol{b}_\infty.
\end{aligned}$$

with $\boldsymbol{b}_\infty = \lim_{\tau \to \infty} \boldsymbol{b}(\tau) = (\boldsymbol{K} + \tilde{\boldsymbol{\Lambda}}_1)'^{-1} \boldsymbol{\delta}_{1R}$.

We also consider a restriction on house value increases. Although Committee Parameters (2019) does not incorporate housing in their model, they do advise a value of 4.1% for the unconditional return on non-public real estate. This takes into account both price increases and net rental income, so we first must separate the two figures. We consider a simplified model where the value of a house, H, is the discounted sum of all future net rental income. In an environment with stable growth q and discount factor ρ, the value of a house is $H = \frac{H r^{\text{imp}}}{\rho - q}$. Rearranging gives $r^{\text{imp}} = \rho - q$. CPB (2020) estimate that $\rho = 3.95\%$ and $q = 0.7\%$. Consequently, we set the imputed rent (net rental income) at 3.25% of the house value. This means that, under the aforementioned restriction, the remaining 0.85% forms the unconditional yearly return on house value r_Q^{h}. From this we can construct our final restriction. The unconditional continuously compounded return on house value is given by

$$\ln(1 + r_Q^{\text{h}}) = \lim_{t \to \infty} \mathrm{E}\left[\ln \frac{Q_{t+1}}{Q_t}\right] = \delta_{0R} + \eta_Q - r^{\text{imp}} - \frac{1}{2}\boldsymbol{\sigma}_Q' \boldsymbol{\sigma}_Q.$$

The ensuing restriction is

$$\eta_Q - r^{\text{imp}} = \ln(1.0085) - \delta_{0R} + \frac{1}{2}\boldsymbol{\sigma}_Q' \boldsymbol{\sigma}_Q.$$

This restriction does not affect the time-varying risk premium, because its unconditional expected value is 0.

2.2 Calibration

We use the Kalman filter to estimate the model, adopting the augmented formulation from Pelsser (2019) that allows us to use the standard version of the Kalman filter. The Kalman filter is based on a discrete model, but our economic model is set up in continuous time. As such, we first discretise the model and its parameters. Then, we run the Kalman filter in combination with Quasi-Maximum Likelihood to find the model parameters that best fit the Dutch economy. Appendices A and B describe the state space formulation, the discretisation and the Kalman equations in more detail. De Jong (2000) provides further discussion on how to estimate affine term structure models with a Kalman filter.

2.3 Data

Our data consists of quarterly observations for bond yields, stock returns, inflation rates and house prices, spanning from March 1973 until March 2021. It is an updated version of the data used by Draper (2014). In contrast, the Committee Parameters (2019) uses monthly observations from the last 20 years, which have lower measurement error and are likely a more accurate reflection of the current dynamics. However, housing data is only available in quarterly frequency, so the Committee Parameters (2019) set would leave us with less than 100 observations. This could lead to substantial estimation errors due to the large number of parameters to be estimated, which is why we opt for the larger data set.

We consider bonds with 6 different maturities: 3 months, 1 year, 2 years, 3 years, 5 years and 10 years. From Jan-1973 until Dec-1998, the shortest maturities are 3-month money market rates from Frankfurt banks. From Dec-1999, they are the 3-month Euribor rates. Both are available through the Bundesbank. For the longer maturities, we use German government bonds up to Dec-2003, also available through the Bundesbank. From Jan-2004 onward, we use the zero-coupon bond rates constructed by De Nederlandsche Bank.

The price index is a combination of German CPI figures as published by the IMF and, from Jan-1999 onward, the Harmonized Index of Consumer Prices that the ECB publishes. The MSCI world index in euros serves as our stock market, available through Eikon. We complete the data with house prices. Dutch residential property prices come from the Bank for International Settlements.

Figure 1 depicts the gradual decline of bond yields throughout the years. Since we model the yields as an affine function of the state variables, this decline makes at least one of the

Figure 1: Bond yields (%)

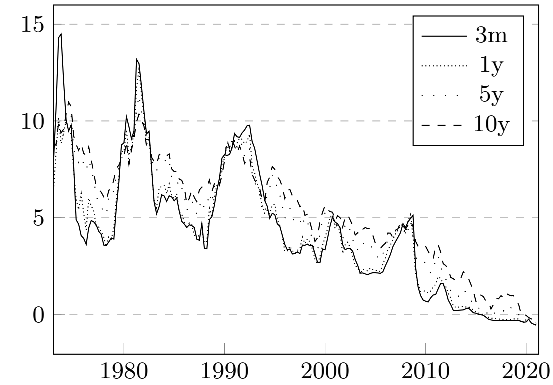

Figure 2: Inflation rate (%)

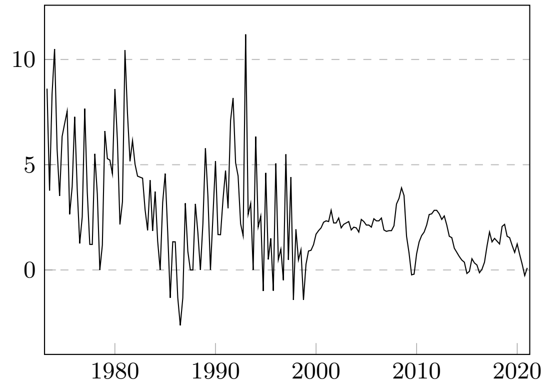

states likely to be close to nonstationary. This could hamper parameter estimation if we did not put restrictions on the unconditional interest rate, because the high unconditional variance of the state variable obscures the steady-state value of the interest rate. Table 1 depicts how the yield curve is on average upward-sloping, with short-term yields more volatile than long-term yields. Table 2 shows that correlations across yields are generally above 90%, suggesting that a two-factor model could be sufficient to explain a large part of the yield curve. We also see a substantial correlation between the yields and the inflation rate, supporting the applicability of a simple factor structure for these time series.

Figure 3: Equity returns (%)

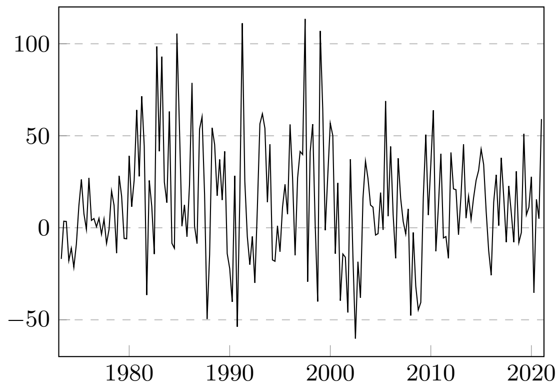

Figure 4: Housing returns (%)

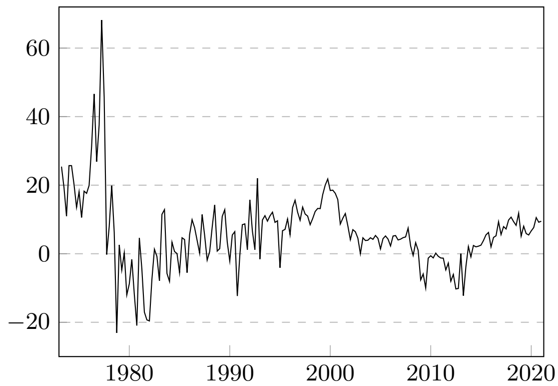

Figures 3 and 4 show returns for the stock index and the housing market respectively. The stock returns exhibit no obvious pattern, substantiated by low autocorrelation and correlation with the other variables, as depicted in Tables 1 and 2. As such, the two-factor model will likely be of little help in predicting stock behaviour. However, stock returns are notoriously hard to predict, so our model is likely a sufficient approximation. Unlike stocks however, housing prices appear highly predictable. Indeed, the quarterly returns have a first-order autocorrelation (FOAC) of 68%. While our model does allow for autocorrelation through the two factors, the housing returns are mostly uncorrelated with the other variables. Thus, it is unlikely that the two-factor model can explain the high autocorrelation in housing returns, and the model with a constant risk premium for house price dynamics is likely not a perfect fit. The housing and inflation data also depict the trade-off between sample size and accuracy, with the older data seemingly noisier than the more recent data.

Table 1: Summary statistics

	Mean (%)	Variance (%)	Skewness	Kurtosis	FOAC
3-month Yield	4.199	3.358	0.670	0.130	0.945
1-year Yield	4.204	3.061	0.323	-0.541	0.959
2-year Yield	4.402	3.033	0.133	-0.761	0.961
3-year Yield	4.595	3.022	0.005	-0.843	0.963
5-year Yield	4.905	2.969	-0.138	-0.870	0.966
10-year Yield	5.345	2.796	-0.290	-0.787	0.967
Inflation rate	2.438	2.263	1.060	1.624	0.475
Equity returns	10.429	29.658	-0.249	0.501	0.163
Housing returns	5.380	10.289	0.621	4.031	0.678

Table 2: Correlation matrix

	3-m Yield	1-year Yield	2-year Yield	3-year Yield	5-year Yield	10-year Yield	Inflation rate	Equity returns	Housing returns
3-m Yield	X	0.979	0.966	0.951	0.928	0.897	0.596	-0.060	-0.017
1-year Yield	0.979	X	0.995	0.986	0.969	0.941	0.581	-0.050	-0.015
2-year Yield	0.966	0.995	X	0.997	0.987	0.966	0.572	-0.045	0.009
3-year Yield	0.951	0.986	0.997	X	0.996	0.981	0.562	-0.039	0.028
5-year Yield	0.928	0.969	0.987	0.996	X	0.994	0.547	-0.032	0.051
10-year Yield	0.897	0.941	0.966	0.981	0.994	X	0.529	-0.029	0.076
Inflation rate	0.596	0.581	0.572	0.562	0.547	0.529	X	0.023	-0.015
Equity returns	-0.060	-0.050	-0.045	-0.039	-0.032	-0.029	0.023	X	-0.015
Housing returns	-0.017	-0.015	0.009	0.028	0.051	0.076	-0.015	-0.015	X

2.4 Estimation results

We calibrate four different models that are each an adjustment to the previous model. Model (i) omits the housing data and is most comparable to previous calibrations of the KNW model. Model (ii) adds housing data without additional restrictions. Model (iii) expands on that with an additional restriction on the unconditional housing return. Finally, model (iv) incorporates a time-varying risk premium.

Table 3: Model calibration on quarterly data from Q1-1973 to Q1-2021

	(*i*)		(*ii*)		(*iii*)		(*iv*)	
Parameter	Estimate	SE	Estimate	SE	Estimate	SE	Estimate	SE
Instantaneous expected inflation $\pi_t = \delta_{0\pi} + \boldsymbol{\delta}'_{1\pi}\boldsymbol{x}_t$								
$\delta_{0\pi}$	1.89%	-	1.89%	-	1.89%	-	1.89%	-
$\delta_{1\pi(1)}$	0.57%	(0.08%)	0.07%	(0.10%)	-0.06%	(0.16%)	0.03%	(0.40%)
$\delta_{1\pi(2)}$	0.39%	(0.12%)	-0.72%	(0.13%)	-0.71%	(0.11%)	-0.73%	(0.11%)
Instantaneous nominal interest rate $R_t = \delta_{0R} + \boldsymbol{\delta}'_{1R}\boldsymbol{x}_t$								
δ_{0R}	2.30%	-	2.32%	-	1.73%	-	2.39%	-
$\delta_{1R(1)}$	1.32%	(0.12%)	0.23%	(0.13%)	0.10%	(0.30%)	0.15%	(0.84%)
$\delta_{1R(2)}$	0.82%	(0.12%)	-1.56%	(0.15%)	-1.53%	(0.09%)	-1.56%	(0.15%)
State variable dynamics $\mathrm{d}\boldsymbol{x}_t = -\boldsymbol{K}\boldsymbol{x}_t\,\mathrm{d}t + [\boldsymbol{\mathcal{I}}_{2\times2}\ \mathbf{O}_{3\times2}]\,\mathrm{d}\boldsymbol{z}_t$								
$K_{(11)}$	4.30%	(3.65%)	3.59%	(3.51%)	4.06%	(3.43%)	4.28%	(4.11%)
$K_{(22)}$	16.35%	(9.26%)	21.35%	(12.19%)	18.20%	(11.08%)	19.10%	(12.23%)
$K_{(21)}$	5.58%	(5.16%)	6.93%	(5.93%)	10.36%	(7.87%)	8.50%	(9.22%)
Realised inflation process $\frac{\mathrm{d}\Pi_t}{\Pi_t} = \pi_t\,\mathrm{d}t + \boldsymbol{\sigma}'_\Pi\,\mathrm{d}\boldsymbol{z}_t$								
$\sigma_{\Pi(1)}$	-0.02%	(0.08%)	-0.01%	(0.08%)	0.00%	(0.08%)	-0.02%	(0.08%)
$\sigma_{\Pi(2)}$	-0.02%	(0.07%)	0.03%	(0.07%)	0.03%	(0.07%)	0.03%	(0.07%)
$\sigma_{\Pi(3)}$	-0.92%	(0.05%)	-0.92%	(0.05%)	-0.92%	(0.05%)	0.92%	(0.05%)
Stock return process $\frac{\mathrm{d}S_t}{S_t} = (R_t + \eta_S)\,\mathrm{d}t + \boldsymbol{\sigma}'_S\,\mathrm{d}\boldsymbol{z}_t$								
η_S	4.24%	-	4.22%	-	4.83%	-	4.17%	-
$\sigma_{S(1)}$	-0.62%	(1.24%)	-1.84%	(1.29%)	-1.98%	(1.31%)	-2.07%	(1.35%)
$\sigma_{S(2)}$	1.82%	(1.15%)	-0.73%	(1.13%)	-0.82%	(1.18%)	-0.55%	(1.62%)
$\sigma_{S(3)}$	-1.37%	(1.08%)	-1.12%	(1.09%)	-1.23%	(1.09%)	-1.32%	(1.09%)
$\sigma_{S(4)}$	14.53%	(0.73%)	14.61%	(0.75%)	14.67%	(0.75%)	14.73%	(0.76%)
Prices of risk $\boldsymbol{\lambda}_t = \boldsymbol{\lambda}_0 + \boldsymbol{\Lambda}_1\boldsymbol{x}_t$								
$\lambda_{0(1)}$	-0.126	(0.029)	-0.213	(0.014)	-0.226	(0.012)	-0.212	(0.012)
$\lambda_{0(2)}$	0.175	(0.029)	-0.010	(0.062)	-0.005	(0.053)	-0.007	(0.140)
$\Lambda_{1(11)}$	0.023	(0.042)	0.092	(0.044)	0.106	(0.088)	0.104	(0.237)
$\Lambda_{1(12)}$	0.278	(0.044)	0.102	(0.038)	0.102	(0.046)	0.109	(0.127)
$\Lambda_{1(21)}$	-0.006	(0.055)	0.250	(0.085)	0.244	(0.076)	0.250	(0.074)
$\Lambda_{1(22)}$	0.277	(0.095)	0.163	(0.125)	0.164	(0.134)	0.163	(0.202)
Housing return process $\frac{\mathrm{d}Q_t}{Q_t} = (R_t + \eta_Q - r^{\mathrm{imp}})\,\mathrm{d}t + \boldsymbol{\sigma}'_Q\,\mathrm{d}\boldsymbol{z}_t$								
η_Q	-	-	4.85%	(2.47%)	2.45%	-	1.78%	-
$\sigma_{Q(1)}$	-	-	-0.96%	(0.47%)	-0.89%	(0.49%)	-0.35%	(0.39%)
$\sigma_{Q(2)}$	-	-	0.23%	(0.46%)	0.25%	(0.48%)	0.48%	(0.35%)
$\sigma_{Q(3)}$	-	-	-0.19%	(0.43%)	-0.10%	(0.45%)	-0.39%	(0.29%)
$\sigma_{Q(4)}$	-	-	0.19%	(0.39%)	0.25%	(0.39%)	0.44%	(0.28%)
$\sigma_{Q(5)}$	-	-	5.36%	(0.27%)	5.46%	(0.28%)	3.78%	(0.19%)
Dynamic risk premium $\frac{\mathrm{d}\zeta_t}{\zeta_t} = -k\,\mathrm{d}t + \sigma'_Q\,\mathrm{d}\boldsymbol{z}_t$								
k	-	-	-	-	-	-	34.93%	(2.46%)
β_Q	-	-	-	-	-	-	70.52%	(4.97%)
$\log L$	7861.3		8459.7		8453.7		8525.6	

Parameter estimates and standard errors of (*i*) estimation without housing data, (*ii*) estimation with housing data without a constraint on $\eta_Q - r^{\mathrm{imp}}$, (*iii*) estimation with a constraint on $\eta_Q - r^{\mathrm{imp}}$, and (*iv*) estimation with a time-varying risk premium for housing. The log likelihood of model (*i*) is incomparable with the other three, as it uses a different data set.

Table 3 displays the estimated parameters of the four models. To understand the impact of adding housing data we examine the parameters of models (*i*) and (*ii*). The addition of housing data leads to an estimate of 4.85% for the risk premium η_Q, which is 0.63% higher than the equity risk premium. This is an interesting observation, considering that we estimate the volatility of equity from σ_S at almost 3 times that of housing. Furthermore, housing returns are, at least on the surface, more predictable. Because of these factors, one might expect that housing investments would have a lower price of risk. We offer two potential explanations why that is not the case in this unconstrained calibration. The first explanation is that certain obstacles, such as low liquidity, entry costs, and maintenance costs, lead investors to require a higher return on investment. The second explanation is that the data period from 1973 to 2021 is not representative of the long term behaviour of house prices. Other forces such as the shortage in housing supply, the increase in available credit and a decreasing discount rate may have caused positive price shocks that cannot (fully) be attributed to the factors in the current model. In either case, the parameters confirm that housing has historically been a well-rewarded investment.

However, the Committee Parameters (2019) does not expect similar returns in the future. In model (*iii*), we apply their restriction on the unconditional return on real estate investments. Because the unconditional return is a function of the steady state interest rate δ_{0R} and the risk premium η_Q, this restriction leads to a lower estimate of both. This decrease in δ_{0R} in turn leads to an increase in η_S to explain the high historical return on equity. The difference between the restricted and unrestricted estimates shows the discrepancy between the housing market's historical performance and the future projection by the Committee Parameters (2019). The restriction leads to a drop in log-likelihood of 6.0.

In model (*iv*) we allow for a time-varying risk premium. The mean-reversion parameter k is estimated at 0.349, corresponding to an autoregressive coefficient β_Q of 0.705, which is roughly equal to the autocorrelation in the data. The log likelihood increases by 71.9, which corresponds to a Likelihood-Ratio Test statistic of 143.8. Because we only require 1 additional parameter, the test statistic follows a $\chi^2(1)$ distribution. As such, the increase is significant for any practical confidence level.

We take a closer look at how model (*iii*) fits the data in Table 4. For conciseness, we only show the statistics of the the 5-year yield, as they are comparable with those of the other maturities. We find that the model is able to explain most of the variance and first-order autocorrelation (FOAC) in the 5-year yields. It is also able to partly follow the inflation rate. However, it has no explanatory power for any of the variance in the equity and housing returns. In fact, the

Table 4: Model (*iii*) fit statistics

	Mean (%)	Variance (%)	Skewness	Kurtosis	FOAC
5-year Yield					
Data	4.905	2.969	-0.138	-0.870	0.966
Residuals	-0.040	0.482	0.227	0.579	0.156
Inflation rate					
Data	2.438	2.263	1.060	1.624	0.475
Residuals	-0.137	1.851	0.530	1.500	0.211
Equity returns					
Data	10.429	29.658	-0.249	0.501	0.163
Residuals	2.829	30.033	-0.321	0.421	0.178
Housing returns					
Data	5.380	10.289	0.621	4.031	0.678
Residuals	2.443	10.967	0.043	3.631	0.701

residual variance is larger than the variance in the original data, which is likely due to the extra restrictions on their respective risk premia. These restrictions also explain why residuals of both returns have a substantial positive mean. Furthermore, the residuals in the housing data are still heaviliy serial correlated, so the two-factor structure does not pick up on the housing trends. We can see this confirmed by the model fit of equity and housing in Figure 5, which does not appear to track the data. This is not surprising, considering that we only have two factors to explain nine variables. Equity returns in particular are notoriously hard to predict. However, the serial correlation in housing returns suggests that we can do better on that front.

Figure 5: Model fit with the calibration of model (*iii*)

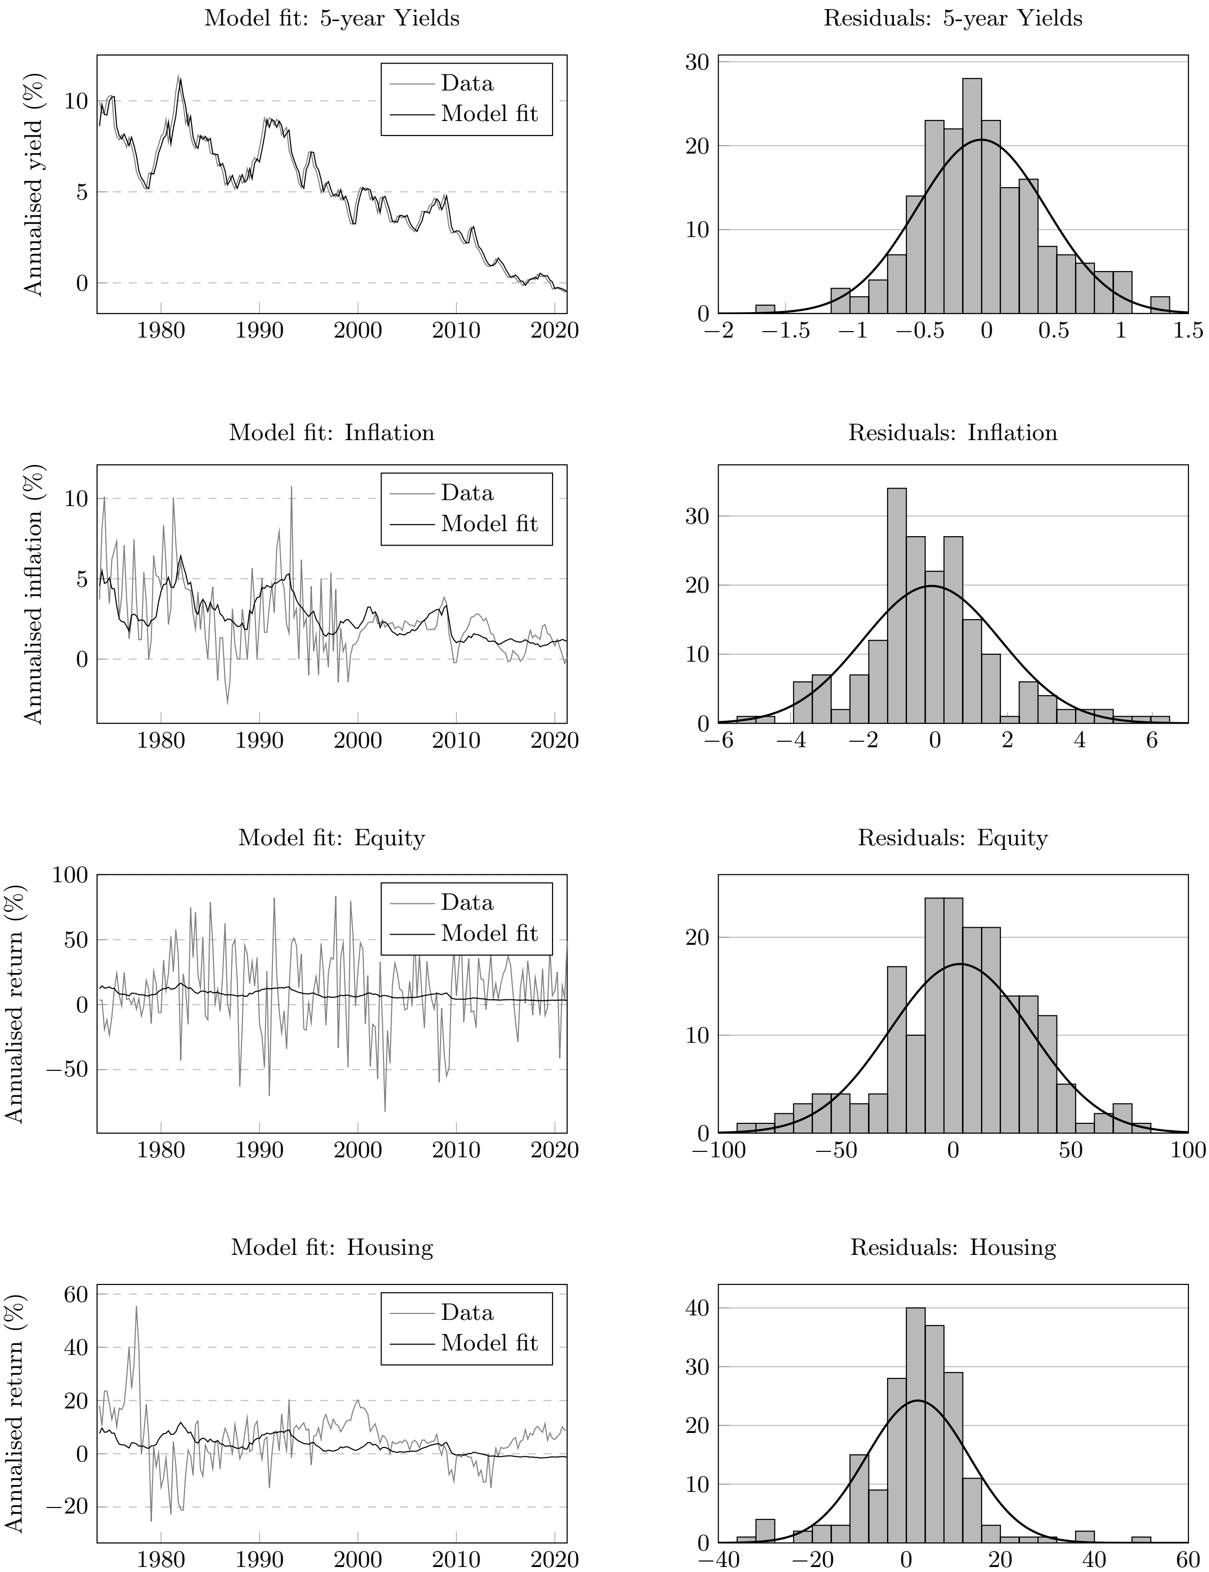

Figure 6 and Table 5 describe the fit on housing returns of model (*iv*), which includes the time-varying risk premia. There are no noticeable difference in the fit of the other variables as compared to model (*iii*), so we omit them from the analysis.

Figure 6: Model (*iv*) fit

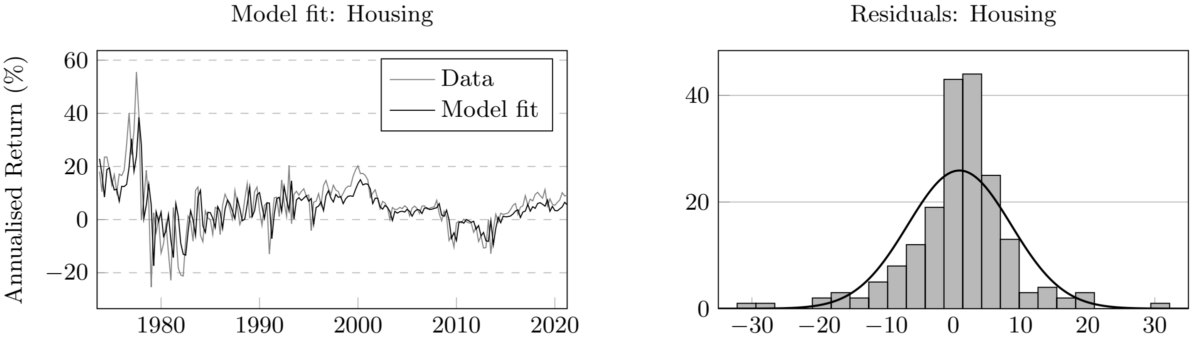

Table 5: Model (*iv*) fit statistics

	Mean (%)	Variance (%)	Skewness	Kurtosis	FOAC
Housing returns					
Data	5.380	10.289	0.621	4.031	0.678
Residuals	0.909	7.700	-0.443	3.613	-0.066

The addition of a time-varying risk premium leads to a better fit on housing data. Here the model matches the trend in the data, especially as it becomes less noisy. Indeed, Table 5 shows that the residual autocorrelation in model (*iv*) is close to 0, and the residual variance drops by 30% as compared to model (*iii*). Furthermore, the residual mean decreases from 2.4% to 0.9%. The dynamic risk premium thus helps explain the high housing returns that do not match the Committee Parameters (2019) restrictions. Given the substantial improvement in fit for the housing market, as also indicated by the jump in log likelihood, we use model (*iv*) and its parameters in the second part of this paper.

3 Household model

Within the economy, we construct a household consumption model. We can infer which pension fund investment strategy maximises household welfare by analysing the utility a household obtains from consumption. We simulate $N_{sim} = 20000$ households that enter our model at age 25, at which point $t = 1$. They first work for $T^{\mathrm{W}} = 40$ years and then spend $T^{\mathrm{P}} = 20$ years in retirement, for a total of $T = 60$ years in our model. Households obtain utility from real consumption of non-durable goods c and the number of housing units h, which represents the various valuable aspects of houses, such as location, size and appearance. In line with among others Cocco (2005), households have a Cobb-Douglas utility function with Constant Relative

Risk Aversion (CRRA):

$$u(c,h) = \frac{(c^{1-\lambda}h^{\lambda})^{1-\gamma}}{1-\gamma}.$$

Here λ describes the relative preference between the two goods and γ is the coefficient of relative risk aversion. The goal for households is to maximise their lifetime utility

$$\max \sum_{t=1}^{T} \rho^t \frac{(c_t^{1-\lambda}h_t^{\lambda})^{1-\gamma}}{1-\gamma}. \tag{5}$$

Although we are specifically focused on finding optimal pension investments, we still require a realistic model for the housing and consumption decisions. Optimising all decision variables at once is theoretically possible, but this quickly becomes computationally infeasible, especially when adding more complexities to the model. We deal with this by first constructing a simplified model which can be solved numerically. We call this model the 'Baseline Model' and discuss its optimisation in Section 3.1. Then we use the optimised household choices from the Baseline Model as inputs in more complex models, in which we focus solely on optimising the pension fund strategy. We discuss this process in Section 3.2.

3.1 Baseline Model

3.1.1 Income and Pension

The income for a household between 25 and 64 years old comes from labour L_t. Real labour income l_t is constant and equal for each household. The pension system is based on the new Dutch pension contract, the details of which are outlined in Metselaar et al. (2020). To simplify the model, we assume no solidarity buffer, which means there is no inter-generational risk-sharing. In this system, each household has its own pension depository, so we can model one individual household at a time, rather than a whole population.

Each year, households allocate a fraction p of their income to their pension depository W^{P}, which starts at 0. When households reach age 65 at $t = T^{\mathrm{W}}$, they spend the next 20 years in retirement and start receiving pension benefits according to

$$B_t = \frac{W_t^{\mathrm{P}}}{\sum_{\tau=0}^{T-t} DF_{\tau,t}},$$

where $DF_{\tau,t}$ denotes the discount factor between t and $t+\tau$, which we set equivalent to a bond price at time t with maturity τ. The pension assets of a household are invested into stocks, long-term bonds that match the expected pension benefits, and single-period bonds with fractions x^{S}, x^{B} and x^{C} respectively, where $x^{\mathrm{S}} + x^{\mathrm{B}} + x^{\mathrm{C}} = 1$. A household has a return

$R_t^{\mathrm{P}} = x^{\mathrm{B}} R_t^{\mathrm{B}} + x^{\mathrm{S}} R_t^{\mathrm{S}} + x^{\mathrm{C}} R_t^f$ on its pension, such that

$$W_{t+1}^{\mathrm{P}} = R_t^{\mathrm{P}}(W_t^{\mathrm{P}} - B_t) + pL_{t+1}. \tag{6}$$

3.1.2 Consumption and Housing

Households spend their wealth on the consumption of non-durable goods, C_t^{NDG}, and housing. The housing decision consists of their tenure 1_t^{own} and number of housing units h_t, which is the current house price H_t normalised by the housing index:

$$h_t = \frac{H_t}{Q_t}.$$

Renting households spend a fixed fraction ξ_{R} of the house value on housing services. Any financial wealth that remains after consumption is invested in the single-period bond with return R_t^f. The dynamics of a renting household's financial wealth W_t are then

$$W_{t+1} = (W_t - C_t^{\mathrm{NDG}} - H_t \xi_{\mathrm{R}}) R_{t+1}^f + B_{t+1} + (1-p)L_{t+1}, \quad \text{if } 1_t^{\mathrm{own}} = 0. \tag{7}$$

Because the elasticity of substitution for Cobb-Douglas utility is 1, a renting household will always spend the same fraction λ on housing services and (1-λ) on non-durable goods, irrespective of current price levels. This can intuitively be thought of as a balance between two effects when a good increases in price. On the one hand, a household wants to maintain its consumption so is inclined to spend more. On the other hand, the good has become less attractive so the household is inclined to spend less. These effects even out exactly in Cobb-Douglas utility. This simplification also holds in multiple periods because the wealth in future periods is not affected by the current consumption (Yao and Zhang, 2005). However, it does not hold for home owners, because their expenses not only count as consumption but also as a risky investment in a housing asset.

Home owners also spend a fixed fraction ξ_{O} of the house value on housing services. In addition, they pay a minimum fraction δ as a down payment. They receive return $R_t^{\mathrm{H}} = Q_t/Q_{t-1}$ on their house value and pay R_t^f on the mortgage. Equating mortgage interest to savings interest greatly reduces the complexity of the model because it eliminates separate decisions on the size of the mortgage and repayments. We further reduce the number of state and decision variables by assuming that an investor can adjust mortgage size and tenure flexibly every year without costs. This is equivalent to the household selling its house at the end of each year and buying a new house without transaction costs. The financial wealth of a home owner then evolves according to

$$W_{t+1} = (W_t - C_t^{\mathrm{NDG}} - H_t(\xi_{\mathrm{O}} + \delta))R_{t+1}^f + H_t(R_{t+1}^{\mathrm{H}} - (1-\delta)R_{t+1}^f) + B_{t+1} + (1-p)L_{t+1}, \quad \text{if } 1_t^{\mathrm{own}} = 1. \quad (8)$$

Consumption and house size must be positive in each period. Furthermore, the savings, as indicated by the first term on the right-hand side in (7) and (8), must be non-negative. We impose each of these restrictions when optimising the decisions.

We let half of the simulated households be permanent owners and the other half permanent renters. It is feasible to optimise the tenure decision, but initial results show that this is a somewhat futile exercise. Specifically, we find that households beyond the age of 75 always prefer renting due to the reduced risk. Before that, households almost always prefer to own their house, except when their current wealth is very low compared to expected wealth in future periods. These choices are not an accurate reflection of the real world, where many younger households rent, and some older households still own their house in their final years. Important reasons for this discrepancy between our model and the real world are lack of a bequest motive and of a lower bound for house prices. Because we are also interested in the impact on young renters and on old owners, we let tenure be decided exogenously in that households are either permanent owners or permanent renters. The set d of decision variables to be optimised is then $d = \{c_t, h_t, x_t\}_{t=1}^T$. We discuss the optimisation technique in the next section.

3.1.3 Solving with backward induction

We refer to the optimal decision variables as the solution to the model, d^*. Finding the solution is complicated because it is different for each combination of the state variables. In the Baseline Model the state variables are time, wealth, pension wealth and income $\{t, w_t, w_t^{\mathrm{P}}, l_t\}$. We deal with this dependency by constructing a grid over the state variables and finding the solution $d_t^*(\{t, w_t, w_t^{\mathrm{P}}, l_t\})$ for each gridpoint numerically, where d_t denotes the set of decision variables at time t. This presents two new issues, however: (1) long computation time can make solving infeasible and (2) the gridpoints can not be solved separately because the solutions across gridpoints depend on each other.

First, computation time rapidly increases with the number of state variables since each state variable adds an additional dimension to this grid. For instance, allowing households to carry over their house to the next period would add two additional state variables: the house value and the mortgage. If we include 20 gridpoints for each, this simple modification would increase the computation time by 400.[1] Evidently, minimising the number of state variables is the key to a tractable model. Next to simplifying the model, we normalise the variables by w_P, which further reduces the number of state variables by 1. We then denote the set of remaining state variables by $z_t = \{t, w_t/w_t^{\mathrm{P}}, l_t/w_t^{\mathrm{P}}\}$ Even though labour income is assumed constant, we can not ignore it as a state variable because the ratio l_t/w_t^{P} can still vary across simulations.

The second issue is that the solutions across time are interdependent. Specifically, for a given state z_t, each combination of decisions leads to a realised utility at time t and a new state z_{t+1}. Thus, to find the solution at t we first need to know the expected utility when reaching any state z_{t+1}. We solve this by splitting the problem into many two-period problems. We start in the final period, find the utility of each state and then iterate backwards through a process called backward induction. Judd (1998) writes extensively on this topic.

Recall from equation (5) that households maximise the expected utility over the remaining years. Thus, we can define the value of being in state z_t as the expected utility over the remaining years, given that the household makes the optimal decisions. We refer to this as the value function:

$$v_t(z_t) = \max_{\{c_s,h_s,x_s\}_t^T} \sum_{s=t}^{T} \rho^{s-t}\mathbb{E}_t(u(c_s,h_s)).$$

The value function can be written recursively, commonly known as the Bellman equation:

$$\begin{aligned} v_t(z_t) &= \max_{\{c_s,h_s,x_s\}_t^T} \sum_{s=t}^{T} \rho^{s-t}\mathbb{E}_t(u(c_s,h_s)) \\ &= \max_{\{c_s,h_s,x_s\}_t^T} u(c_t,h_t) + \sum_{s=t+1}^{T} \rho^{s-t}\mathbb{E}_t(u(c_s,h_s)) \\ &= \max_{\{c_t,h_t,x_t\}} u(c_t,h_t) + \rho\mathbb{E}_t(v_{t+1}(z_{t+1})). \end{aligned}$$

This representation shows how we can split the large problem into many two-period problems, using the value of any given state in the next period to calculate the value in the current period.

The final period T is easily solved because households obtain no further utility in $T+1$, so the expected value term disappears. We construct a grid G_T over the state variables w/w^{P} and l/w^{P}. For each gridpoint $Z_T \in G_T$, we find the decision variables that maximise $u(c_T,h_T)$. We then save this solution $d_T^*(Z_T)$ along with the value $v_T(Z_T)$.

Next, we work backwards from $T-1$ to 1. At each point in time we construct a grid G_t and optimise the decision variables for all gridpoints. Here we use the wealth dynamics from equations (6), (7) and (8) to determine what state z_{t+1} will be reached with the decisions d_t and

[1] For reference, Van Hemert (2010) uses 60 parallel-connected computers to solve a comparable model with six decision variables and five state variables and has a computation time of 10 hours for a single run.

current state Z_t. We can then assess the value of this state because we have already determined $v_{t+1}(Z_{t+1})$ for all $Z_{t+1} \in G_{t+1}$. Since the state variables for wealth are on a continuous scale, the final step is to interpolate the value function between the gridpoints. We use cubic Hermite interpolation, as is frequently done in the literature. For a comprehensive review of different interpolation techniques, we refer to Judd (1998). We set up a grid of 20 values for w_t/w_t^{P}, 20 for l_t/w_t^{P} and 60 for t. Since labour income is 0 in the final 20 years, this gives a total of 16,400 individual optimisations.

3.2 Extended models

While the Baseline Model allows us to find the optimal solution for every possible state, it relies on unrealistic assumptions. In particular, real households are not able to switch houses every year without costs and there may not be any houses available in the lower price ranges, even if the household would theoretically prefer these. These frictions mean that in reality home owners are at higher risk than the Baseline Model accounts for. We construct four models to understand how this impacts the optimal pension strategy. The first model, which we refer to as the 'Flexible Housing Model', uses the same assumptions as the Baseline Model, but with a new estimated pension strategy to make it comparable with the other models. The second model, which we refer to as the 'Inflexible Housing' model, reduces the flexibility in housing transfers. Rather than changing every year, households have a holding period of 5 years, which they can only break if their net worth before income would become negative. In this case, they sell their house and use their income to finance an alternative option. The third model, which we refer to as the 'Maximised Housing' model, also has this inflexibility. In addition, home owners choose the biggest house they can finance after their non-durable goods consumption. This maximum is restricted by the minimum down payment δ. These models all assume households may have some cash saved up for a rainy day. In reality, some households live mostly of their pension income. In the fourth model, 'Maximised Consumption', households spend all wealth that is left after their housing expenses on non-durable goods consumption, so they have no precautionary savings.

Due to the additional state variables, it is no longer feasible to solve these extended models through backward induction. Instead, we use the solutions from the Baseline Model as input for household choices c_t and h_t. We then find the pension strategy by optimising a parametric function that accounts for the age of the household. The age-dependent functions are captured in the vector $\boldsymbol{\theta}(t)$:

$$\boldsymbol{\theta}(t) = \begin{pmatrix} 1 \\ T-t \\ \frac{T}{t}-1 \end{pmatrix}.$$

We further allow the allocation to depend on some indicator of home ownership status $\omega_{m,t}$. We optimise the the model-dependent parameters $\boldsymbol{B}_m \in \mathbb{R}^{2\times 3}$, which govern the interaction such that the allocation to risky assets is

$$x^{\mathrm{S}}_{m,t} = \begin{pmatrix} 1 & \omega_{m,t} \end{pmatrix} \boldsymbol{B}_m \boldsymbol{\theta}(t).$$

We consider three investment strategies that depend on home ownership through $\omega_{m,t}$, where $\omega_{m,t}$ differs for each model m. The first strategy does not differentiate through home ownership: $\omega_{1,t} = 0$. In the second strategy we add a linear relation with the dummy variable $\omega_{2,t} = 1^{\mathrm{own}}_t$, which indicates home ownership status. The third strategy uses the value of the house owned as a fraction of total wealth: $\omega_{3,t} = H_t/(w_t + w^{\mathrm{H}}_t + w^{\mathrm{P}}_t + hc_t)$. Here w^{H}_t denotes net housing wealth and hc_t the human capital at t: the sum of discounted future labour income. This strategy more accurately reflects the sensitivity of household wealth to shocks in the housing market, but would be more complex to implement in practice.

We find the optimal parameters by maximising the certainty equivalent across the pension years. For an average utility $\bar{U}$, the certainty equivalent measures the corresponding value for the Cobb-Douglas term $c^{1-\theta}h^{\theta}$:

$$CE = \left(\frac{(1-\gamma)\bar{U}}{\sum_{t=T^{\mathrm{W}}}^{T^{\mathrm{W}}+T^{\mathrm{P}}} \rho^t} \right)^{1/(1-\gamma)},$$

which follows from inverting the function for lifetime average utility (5). We separately measure the utility of the home owners and renters, $\bar{U}_{\mathrm{O}}$ and $\bar{U}_{\mathrm{R}}$, to find their corresponding certainty equivalents CE_{O} and CE_{R}. We then obtain the parameters by maximising the equally weighted sum

$$\max_{\boldsymbol{B}_m} \frac{CE_{\mathrm{O}} + CE_{\mathrm{R}}}{2}.$$

This approach ensures that the interests of home owners and renters are weighted equally.

3.3 Parameters

Finally, we discuss the parameters of the model. The economy of each simulation is generated with the variables from Section 2. Previous research on the US market frequently scales the house price volatility upward because inidividual house prices are more volatile than the aggregate. For instance, Van Hemert (2010) scales annual volatility to 15%. However, we have no data on idiosyncratic volatility in the Netherlands, in particular in the presence of market cycles, so we refrain from this scaling. It is possible that this underestimates the housing risk that households are subject to.

Table 6 contains the parameters that describe the household preferences. The relative risk aversion is 5, as is common in the literature. We follow Corradin et al. (2014) and set the Cobb-Douglas parameter to 0.3. We set the subjective discount rate to discount according to the annualised steady-state nominal interest rate δ_{0R}. The years spent in each bracket are in line with previous research into the Dutch pension system (Chen et al., 2019). Real labour income is set equal to the median income for a 25-year-old household (CBS, 2021c). We calibrate the annual pension contribution such that, on average, the pension income in the first year of pension is equal to the labour income in the last year of work.

Table 6: Household parameters

Parameter	Symbol	Value
Relative risk aversion	γ	5
Cobb-Douglas housing weight	λ	0.3
Subjective discount rate	ρ	98.2%
Years spent working	T^{W}	40
Years spent in retirement	T^{P}	20
Real labour income	l	24,100
Pension contribution	p	13%
Recurring costs of ownership	ξ_{own}	0.85%
Recurring costs of rental	ξ_{rent}	4.10%
Minimum down payment	δ	30%

Table 6 also shows the parameters that affect the housing choice. CPB (2020) calibrate the steady-state rental price at 4.10% of the house price. In Section 2, we derived the net rental income at 3.25%. It follows that the costs of ownership are then 0.85%. Finally, we set the minimum down payment at 30%.

4 Results

We present the optimal pension strategies and household choices. All figures in this section are inflation-corrected.

4.1 Pension strategy

In this section, we discuss the optimised pension strategies from the extended models. Table 7 shows how a differentiating strategy impacts the certainty equivalent across the four models. As expected, the benefit of accounting for tenure is relatively small in the Flexible Housing model, where households switch to a new house every year. This indicates households can better manage their risk from housing shocks by moving to smaller houses when this option is always immediately available. This is likely further aided by the minimum down payment of 30%, which keeps the housing net worth positive in most scenarios.

The introduction of a holding period of 5 years in the Inflexible Housing model substantially decreases the certainty equivalent. Throughout this holding period, house sizes can deviate from their optimum and leave both home owners and renters with a house that is either smaller than they would choose or one that is worth too much, such that the house no longer matches the risk profile of the household. In this case, accounting for tenure increases welfare by a full percentage point, the equivalent of a certain €240 in consumption each year for a median income household. Households in the Maximised Housing model have an even bigger gap between their chosen and optimal house. Here accounting for tenure increases welfare by 2.5%. The advantage of differentiation disappears when households maximise their consumption, in which case it is primarily the renters that take on additional risk because they lack a buffer.

Table 7: Certainty Equivalent for different housing models and pension strategies

	Equal allocation	Tenure	Housing wealth
Flexible Housing			
CE	27,626	27,735	27,701
Welfare increase		0.39%	0.27%
Inflexible Housing			
CE	24,278	24,526	24,511
Welfare increase		1.02%	0.96%
Maximised Housing			
CE	20,781	21,300	21,104
Welfare increase		2.50%	1.55%
Maximised Consumption			
CE	12,418	12,430	12,423
Welfare increase		0.10%	0.04%

The equal allocation strategy does not consider a differentiating strategy. The tenure strategy allocates differently based on tenure. The housing wealth strategy allocates based on the fraction of total household wealth that is invested in housing.

In all housing models, accounting for tenure is more effective than accounting for the fraction of total wealth invested in housing, even though the latter is arguably a better indicator of the risk a household is exposed to. This may be explainable by a disparity between our model and the real world. In our model, all households base their housing choices on the same solution and have the same down payment, which may decrease the value of differentiating between households in this way. Furthermore, households in our model are either permanent owners or permanent renters, making the tenure strategy especially resilient to changes in the future. We can thus not be sure that the more complex housing wealth strategy would also be inferior in the real world. Still, the results suggest that the relatively simple tenure strategy can bring substantial welfare increases.

Figure 7 shows the optimal risky asset allocation for the 'Equal allocation' and the 'Tenure' strategy. The general picture is that allocations trend downwards as households age, in line with the consensus in life-cycle investing. Households would optimally take a highly leveraged position in the first few years, which rapidly decreases as they build up pension wealth. This leverage means that pension wealth occasionally dips below 0 in our simulation, although the new contributions allow it to recover relatively quickly. Chen et al. (2019) find similar results at the start of the curve. However, through the way they parameterise their curve, they impose that that pensioned households are largely treated the same across their pension years. Our curves, however, show that the optimal allocation still decreases in the last years of the life cycle, in particular for renters.

A comparison between the four housing models shows substantial differences in the allocation. When households can flexibly change houses, we see a relatively small difference between the investment strategy for renters and owners. The strategies deviate considerably when households are under increased risk in the Inflexible and Maximised Housing models. Interestingly, the strategies converge again at the end of the life cycle. This is because home owners decrease their house size towards the end of their life to reduce risk, which brings their risk profile closer to that of renters. Finally, the Maximised Consumption model shows that the difference in strategies disappears when households immediately spend all their savings. This has a larger effect on renters, who do not have the house value as a buffer when pension returns are negative. As such, they prefer a substantially safer strategy than in the other models. This is in line with the results from Chen et al. (2019), whose allocations are generally less aggressive than ours. Households in their model also have no personal savings, so their stable consumption relies entirely on the pension depository.

4.2 Household finances

Our simulation also offers insights into the financial differences between renters and owners. Figure 8 shows how their average asset composition evolves. In the first years, home owners quickly grow their assets compared to renters. However, as the years pass, renters partly make

Figure 7: Risky asset allocation in the equal allocation and tenure strategy

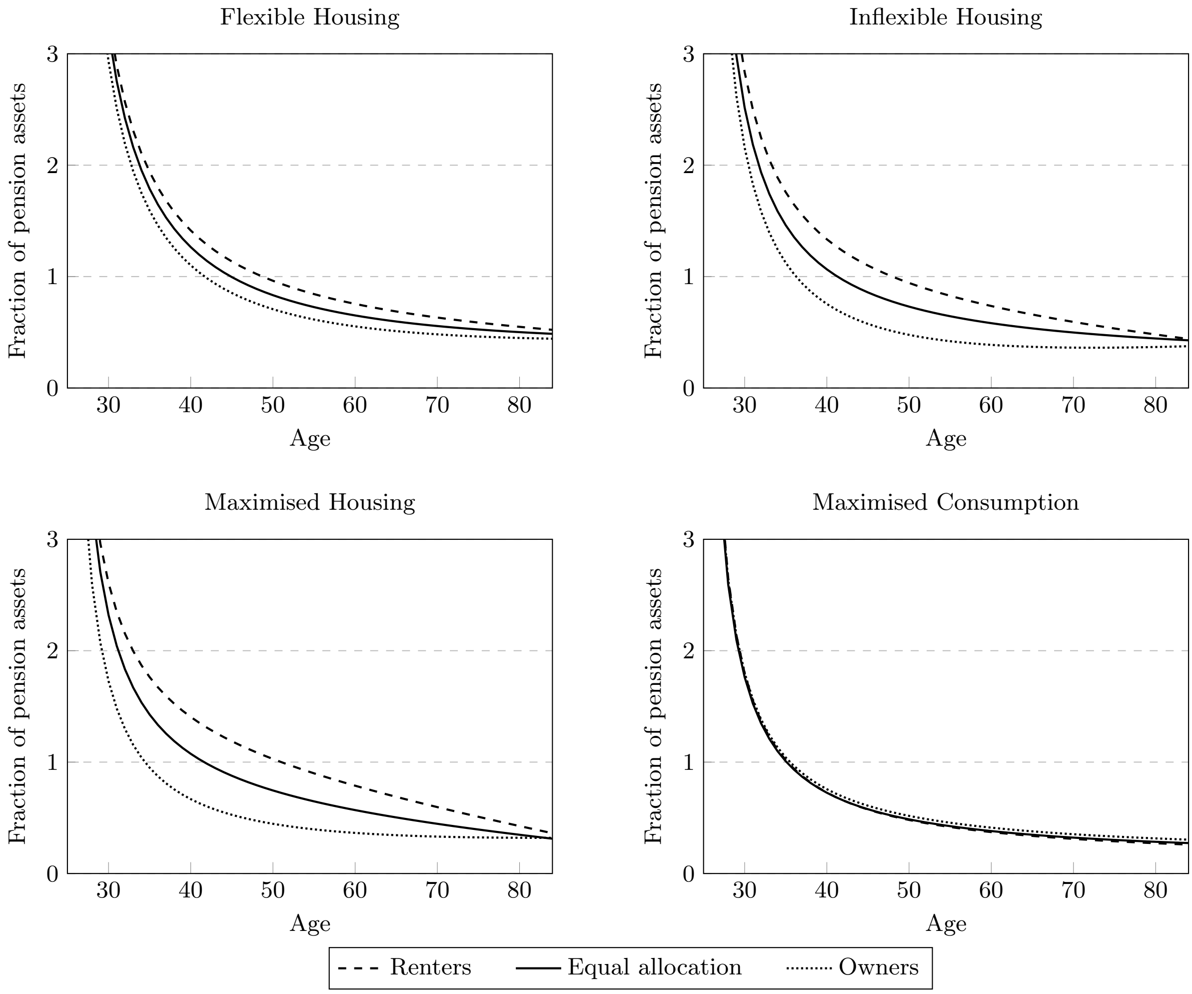

up for this through their riskier pension allocation. Both types of households retain a buffer of cash that they start to consume in their pension years. Our model assumes that this cash is invested safely in the short rate. If households were to invest part of this cash in risky assets, we would likely see an overall decrease in risky pension allocation.

Figure 9 further explores how the household choices vary among renters and owners. As households reduce their built-up buffer, the average consumption increases gradually over time. Rental home values increase for the same reason. Neither of these is related to inflation, as the depicted figures are inflation-corrected. On the other hand, home owners start decreasing their home value after reaching 50, with a rapid decrease at the end of their lifetime. These trends may not reflect how some households behave in real life. An explanation is that people tend to leave behind an inheritance in the real world, which we have not accounted for. In this case, renters may prefer to save some wealth later in life rather than spending it on their rental home and consumption. In a similar vein, home owners would frequently have some cash on hand to compensate for housing shocks, so they would be less averse to the housing risk.

Table 8: Average risky asset allocation

	Renters	Equal allocation	Owners
Flexible Housing	154%	143%	130%
Inflexible Housing	139%	121%	100%
Maximised Housing	136%	114%	87%
Maximised Consumption	88%	88%	91%

The 'Equal allocation' investment displays the strategy when the fund does not differentiate between renters and owners. The 'Renters' and 'Owners' investments display the strategy where the fund differentiates based on tenure.

Figure 8: Average asset composition under the flexible model

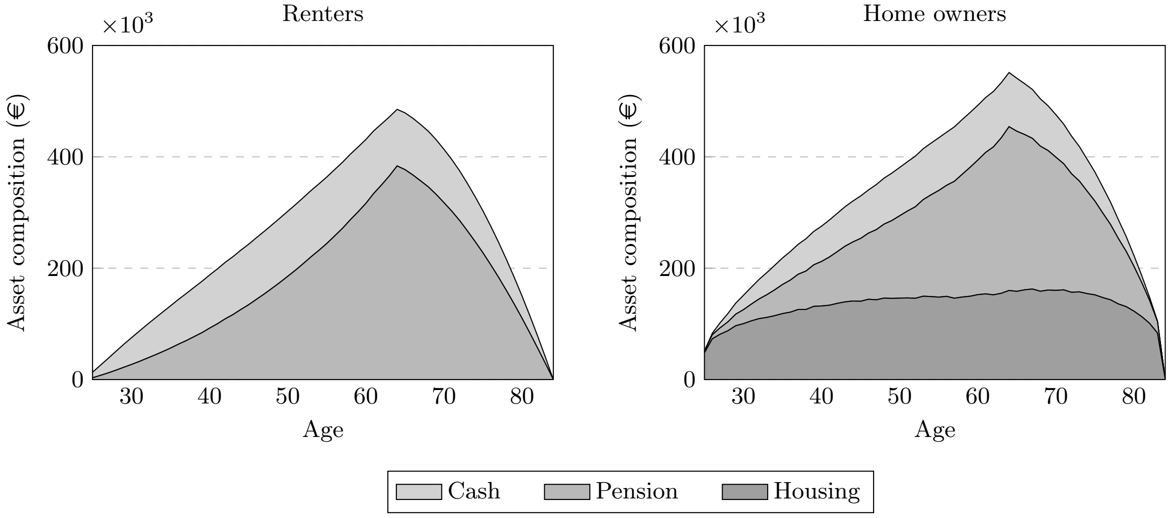

Figure 9: Average consumption and house value under the flexible model

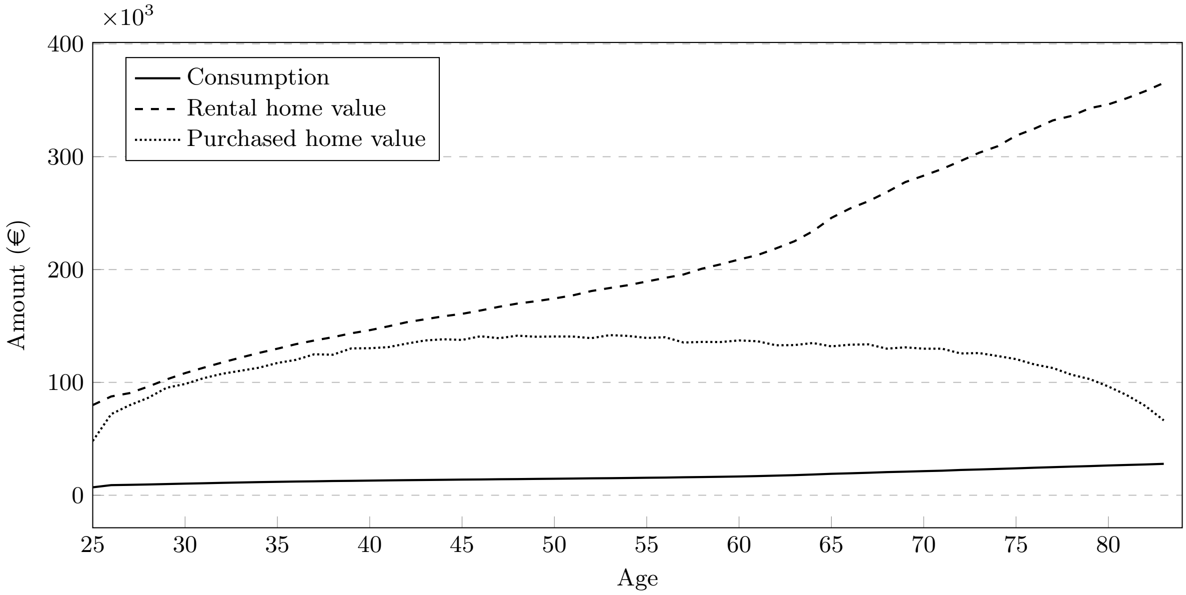

5 Conclusion

Our research sheds new light on how pension funds can personalise the risk profile of their beneficiaries. We find that differentiation based on housing tenure is useful when households are in a sub-optimal housing situation. The benefits are particularly prominent when home owners have to spend more on their house than they would ideally like to, as is the case in the current Dutch housing market. This benefit can already be obtained from relatively simple differentiation between home owners and renters, without requiring additional information.

One caveat to these benefits is that they only come to fruition in the later stages of one's life. They therefore do not provide a direct solution for younger people who can not finance a home. In fact, it would likely be ill-advised to use pension assets for this purpose, as extra funds would only further inflate the housing market.

However, our calibration of the Dutch economy indicates that the housing market itself may reach a solution for those now priced out. In our economic model, where we add a housing asset to the existing models, we expect future returns to be lower than their historical counterparts. In fact, once house prices have adjusted to a low-interest rate environment, our model expects house prices to decrease slightly on a yearly basis, with rental income as the sole profit driver for investors. We find that market cycles can help explain the surge in house prices over the last decade. A different state in the cycle, with decreasing house prices, could make the housing market more accessible for those currently priced out, although the timing of such a correction is unknown.

Future research into the Dutch housing market could use our financial market model and its calibration to analyse the impact of policy changes. Another interesting direction would be to investigate how the housing market cycles affect investment decisions.

Bibliography

Peter Boelhouwer. The housing market in the netherlands as a driver for social inequalities: proposals for reform. *International Journal of Housing Policy*, 20(3):447–456, 2020.

Michael J Brennan and Yihong Xia. Dynamic asset allocation under inflation. *The Journal of Finance*, 57(3):1201–1238, 2002.

CBS. Hoogste prijsstijging koopwoningen in bijna 20 jaar, 2021a. URL `https://www.cbs.nl/nl-nl/nieuws/2021/20/hoogste-prijsstijging-koopwoningen-in-bijna-20-jaar`.

CBS. Vermogen van huishoudens; huishoudenskenmerken, vermogensbestanddelen, 2021b. URL `https://www.cbs.nl/nl-nl/cijfers/detail/83834NED`.

CBS. Inkomensverdeling (gestandaardiseerd inkomen), 2021c. URL `https://www.cbs.nl/nl-nl/visualisaties/inkomensverdeling`.

Damiaan Chen, Maurice Doll, and Annick van Ool. De Toegevoegde Waarde van Maatwerk in Risicotoedeling. Technical report, De Nederlandsche Bank, 2019.

Joao F Cocco. Portfolio choice in the presence of housing. *The Review of Financial Studies*, 18 (2):535–567, 2005.

Committee Parameters. Advies commissie parameters. 2019.

Stefano Corradin, José L Fillat, and Carles Vergara-Alert. Optimal portfolio choice with predictability in house prices and transaction costs. *The Review of Financial Studies*, 27(3): 823–880, 2014.

CPB. Achtergronddocument Woningmarktmodel 2020. Technical report, 2020.

Qiang Dai and Kenneth J. Singleton. Specification Analysis of Affine Term Structure Models. *The Review of Financial Studies*, 55(5):1943–1978, 2000. doi: https://doi.org/10.1111/0022-1082.00278.

Frank De Jong. Time series and cross-section information in affine term-structure models. *Journal of Business & Economic Statistics*, 18(3):300–314, 2000.

Frank De Jong, Joost Driessen, and Otto Van Hemert. Hedging house price risk: Portfolio choice with housing futures. *Available at SSRN 740364*, 2008.

DNB. De nederlandsche bank, dashboards, pensioenen, 2021. URL https://www.dnb.nl/statistieken/dashboards/pensioenen/.

Nick Draper. A Financial Market Model for the Netherlands. Technical report, Centraal Planbureau, 2014.

Marcel Fischer and Michael Z Stamos. Optimal life cycle portfolio choice with housing market cycles. *The Review of Financial Studies*, 26(9):2311–2352, 2013.

Kenneth L Judd. *Numerical methods in economics*. MIT press, 1998.

Ralph S. J. Koijen, Theo E. Nijman, and Bas J. M. Werker. When Can Life Cycle Investors Benefit from Time-Varying Bond Risk Premia? *The Review of Financial Studies*, 23(2): 741–780, 2010. doi: https://doi.org/10.1093/rfs/hhp058.

Luuk Metselaar, André Nibbelink, and Peter Zwaneveld. Het doorontwikkelde contract: beschrijving van varianten, opties en resultaten. 2020.

Sander Muns. A Financial Market Model for the Netherlands: A methodological refinement. Technical report, 2015.

Antoon Pelsser. Kalman Filter Estimation of the KNW Model. *Netspar Academic Series DP 05/2019-047*, 2019.

Antonios Sangvinatsos and Jessica A Wachter. Does the failure of the expectations hypothesis matter for long-term investors? *The Journal of Finance*, 60(1):179–230, 2005.

Otto Van Hemert. Household interest rate risk management. *Real Estate Economics*, 38(3): 467–505, 2010.

Otto Van Hemert et al. Life-cycle housing and portfolio choice with bond markets. Technical report, Citeseer, 2006.

Rui Yao and Harold H Zhang. Optimal consumption and portfolio choices with risky housing and borrowing constraints. *The Review of Financial Studies*, 18(1):197–239, 2005.

A State space model

We follow Pelsser (2019) in augmenting the state vector with the observed variables, which allows us to use a standard version of the Kalman filter:

$$\tilde{\boldsymbol{x}}_t = \begin{pmatrix} \boldsymbol{x}_t \\ \ln \Pi_t \\ \ln S_t \\ \ln Q_t \end{pmatrix}.$$

The state space formulation consists of two parts: the measurement equation and the transition equation. The measurement equation dictates how the observed variables relate to the state vector:

$$\tilde{\boldsymbol{y}}_t = \begin{pmatrix} \boldsymbol{y}_t \\ \ln \Pi_t \\ \ln S_t \\ \ln Q_t \end{pmatrix} = \boldsymbol{a} + \boldsymbol{B}\tilde{\boldsymbol{x}}_t + \boldsymbol{\eta}_t, \qquad \boldsymbol{\eta} \sim \mathrm{N}(\boldsymbol{0}, \boldsymbol{H})$$

Here $\boldsymbol{\eta}$ is assumed i.i.d. and $\boldsymbol{a}$ and $\boldsymbol{B}$ contain the affine structure of the yields and a trivial structure to connect the other variables:

$$\boldsymbol{a} = \begin{pmatrix} -a(\tau_1)/\tau_1 \\ \vdots \\ -a(\tau_m)/\tau_m \\ 0 \\ 0 \\ 0 \end{pmatrix}, \qquad \boldsymbol{B} = \begin{pmatrix} -\boldsymbol{b}(\tau_1)'/\tau_1 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ -\boldsymbol{b}(\tau_1)'/\tau_1 & 0 & 0 & 0 \\ 0 \quad 0 & 1 & 0 & 0 \\ 0 \quad 0 & 0 & 1 & 0 \\ 0 \quad 0 & 0 & 0 & 1 \end{pmatrix}.$$

The model without housing data omits the last row of $\boldsymbol{a}$ and the last row and column of $\boldsymbol{B}$. The term structure coefficients follow the ordinary differential equations

$$\begin{aligned} a'(\tau) &= -\boldsymbol{b}(\tau)'\tilde{\boldsymbol{\lambda}}_0 + \frac{1}{2}\boldsymbol{b}(\tau)'\boldsymbol{b}(\tau)' - \delta_{0R}, \\ \boldsymbol{b}'(\tau) &= -(\boldsymbol{K}' + \tilde{\boldsymbol{\Lambda}}_1')\boldsymbol{b}(\tau) - \boldsymbol{\delta}_{1R}, \end{aligned}$$

with solutions

$$\begin{aligned} \boldsymbol{b}(\tau) &= (\boldsymbol{K}' + \tilde{\boldsymbol{\Lambda}}_1')^{-1}[\exp(-(\boldsymbol{K}' + \tilde{\boldsymbol{\Lambda}}_1')\tau) - \boldsymbol{I}_2]\boldsymbol{\delta}_{1R} \\ a(\tau) &= \int_0^\tau a'(s)ds, \end{aligned}$$

and $a(0) = 0$, $\boldsymbol{b}(0) = \mathbf{0}$. Alternatively, Muns (2015) derives closed form solutions for $a(\tau)$ that don't require numerical integration and are hence computationally more efficient. We adopt that approach and refer to Muns (2015) for the full derivation.

The second part of the state space formulation is the transition equation, which describes how the state vector evolves over time. In continuous time, this can be written as the multivariate Ornstein-Uhlenbeck process:

$$\mathrm{d}\tilde{\boldsymbol{x}}_t = \left[\begin{pmatrix} \mathbf{0}_{2\times1} \\ \delta_{0\pi} - \frac{1}{2}\boldsymbol{\sigma}_\pi'\boldsymbol{\sigma}_\pi \\ \delta_{0R} + \eta_S - \frac{1}{2}\boldsymbol{\sigma}_S'\boldsymbol{\sigma}_S \\ \delta_{0R} + \eta_Q - r^{imp} - \frac{1}{2}\boldsymbol{\sigma}_Q'\boldsymbol{\sigma}_Q \end{pmatrix} + \begin{pmatrix} -\boldsymbol{K} & \mathbf{O}_{2\times3} \\ \boldsymbol{\delta}_{1\pi}' & \mathbf{0}_{1\times3} \\ \boldsymbol{\delta}_{1R}' & \mathbf{0}_{1\times3} \\ \boldsymbol{\delta}_{1R}' & \mathbf{0}_{1\times3} \end{pmatrix} \tilde{\boldsymbol{x}}_t \right] \mathrm{d}t + \begin{pmatrix} [\boldsymbol{I}_{2\times2}\ \mathbf{O}_{2\times3}] \\ \boldsymbol{\sigma}_\pi' \\ \boldsymbol{\sigma}_S' \\ \boldsymbol{\sigma}_Q' \end{pmatrix} \mathrm{d}\tilde{\boldsymbol{z}}_t.$$

However, we require a discrete version as a transition equation. We consider the exact discretisation from $\mathrm{d}\tilde{\boldsymbol{x}}_t = (\boldsymbol{\theta}_0 + \boldsymbol{\Theta}_1\tilde{\boldsymbol{x}}_t)\,\mathrm{d}t + \boldsymbol{\Sigma}_X\,\mathrm{d}\tilde{\boldsymbol{z}}_t$ to

$$\tilde{\boldsymbol{x}}_t = \boldsymbol{\phi} + \boldsymbol{\Phi}\tilde{\boldsymbol{x}}_{t-h} + \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \sim \mathrm{N}(\mathbf{0}, \boldsymbol{Q}).$$

To obtain the transition parameters we consider the eigenvalue decomposition $\boldsymbol{\Theta}_1 = \boldsymbol{U}\boldsymbol{D}\boldsymbol{U}^{-1}$. Given that $\boldsymbol{K}$ is a lower triangular matrix, the eigenvalues of $\boldsymbol{\Theta}_1$ are k_1, k_2 and 0 with multiplicity 3. Then the transition parameters are as follows.

$$\boldsymbol{\Phi} = \boldsymbol{U}\exp(\boldsymbol{D}h)\boldsymbol{U}^{-1}$$

$$\boldsymbol{\phi} = \boldsymbol{U}\boldsymbol{F}\boldsymbol{U}^{-1}\boldsymbol{\theta}_0,$$

where $\boldsymbol{F}$ is a diagonal matrix with elements

$$(\boldsymbol{F})_{ii} = h\alpha((\boldsymbol{D})_{ii}h), \qquad \alpha(x) = \frac{\exp(x) - 1}{x},$$

with $\alpha(0) = 1$. The derivation for $\boldsymbol{Q}$ gives

$$\boldsymbol{Q} = \boldsymbol{U}\boldsymbol{V}\boldsymbol{U}^{-1},$$

$$(\boldsymbol{V})_{ij} = [\boldsymbol{U}^{-1}\boldsymbol{\Sigma}_X\boldsymbol{\Sigma}_X'(\boldsymbol{U}^{-1})']_{ij}h\alpha([(\boldsymbol{D})_{ii} + (\boldsymbol{D})_{jj}]h)$$

B Kalman filter

The measurement and transition equation allow us to estimate the model with a Kalman filter. We follow Pelsser (2019) in setting the initial estimates equal to their unconditional expected values:

$$\hat{\boldsymbol{x}}_0 = \lim_{t\to\infty} \mathrm{E}[\tilde{\boldsymbol{x}}_t] = \mathbf{0},$$

$$\boldsymbol{P}_0 = \lim_{t\to\infty} \mathrm{var}[\tilde{\boldsymbol{x}}], \qquad \mathrm{vec}(\boldsymbol{P}_0) = (\boldsymbol{I} - \boldsymbol{\Phi} \otimes \boldsymbol{\Phi})^{-1}\, \mathrm{vec}(\boldsymbol{Q}).$$

The prediction equations are

$$\hat{\boldsymbol{x}}_{t|t-h} = \boldsymbol{\phi} + \boldsymbol{\Phi}\hat{\boldsymbol{x}}_{t-h},$$

$$\boldsymbol{P}_{t|t-h} = \boldsymbol{\Phi}\boldsymbol{P}_{t-h}\boldsymbol{\Phi}' + \boldsymbol{Q}.$$

The likelihood contribution is

$$\hat{\boldsymbol{y}}_{t|t-h} = \boldsymbol{a} + \boldsymbol{B}\hat{\boldsymbol{x}}_{t|t-h},$$

$$\boldsymbol{u}_t = \boldsymbol{y}_t - \hat{\boldsymbol{y}}_{t|t-h}$$

$$\boldsymbol{V}_t = \boldsymbol{B}\boldsymbol{P}_{t|t-h}\boldsymbol{B}' + \boldsymbol{H},$$

$$\log L_t = -\frac{1}{2}\ln|\boldsymbol{V}_t| - \frac{1}{2}\boldsymbol{u}_t'\boldsymbol{V}_t^{-1}\boldsymbol{u}_t.$$

Finally, the updating step is

$$\boldsymbol{K}_t = \boldsymbol{P}_{t|t-h}\boldsymbol{B}'\boldsymbol{V}_t^{-1},$$

$$\boldsymbol{L}_t = \boldsymbol{I} - \boldsymbol{K}_t\boldsymbol{B},$$

$$\hat{\boldsymbol{x}}_t = \hat{\boldsymbol{x}}_{t|t-h} + \boldsymbol{K}_t\boldsymbol{u}_t,$$

$$\boldsymbol{P}_t = \boldsymbol{L}_t\boldsymbol{P}_{t|t-h}.$$

The algorithm maximises the log likelihood function

$$\log L = \sum_{t=2}^{T} \log L_t.$$