Universität
Augsburg
University

Solving Life Cycle Models, Optimal Age-Dependent Unemployment Insurance, and Adaptive Beliefs in a Real Business Cycle Model

Kumulative Dissertation

der Wirtschaftswissenschaftlichen Fakultät

der Universität Augsburg

zur Erlangung des Grades eines

Doktors der Wirtschaftswissenschaften

(Dr. rer. pol.)

Vorgelegt von

Sijmen Duineveld

Prof. Dr. Burkhard Heer (Erstgutachter)

Prof. Dr. Alfred Maußner (Zweitgutachter)

Prof. Dr. Robert Nuscheler (Vorsitzender der Disputation)

Tag der mündlichen Prüfung: 24. Januar 2019

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This dissertation consists of three essays[1], which all rely on dynamic optimization in discrete time. Dynamic optimization differs from static optimization, because it deals with optimal choices involving multiple periods. We can distinguish between finite and infinite horizon problems, which require different techniques to find the optimal solution. The first two essays deal with finite horizons problems, and the third essay with an infinite horizon problem. Although dynamic optimization is at the core of modern macroeconomics, and also common in the public finance literature, the most standard solution methods were not sufficient to solve the models described in this dissertation. We therefore had to rely on less used solution techniques.

The first essay is called 'Solving Life Cycle Models' and deals with the methodology of solving life-cycle models. This essay is inspired by practical problems that arose working on optimal unemployment insurance, which is the subject of the second essay. Life-cycle models have a finite horizon, which allows us to find the optimal solution using backward iteration. Although standard textbooks on numerical methods for economists cover methods for solving life-cycle models with backward iteration, the coverage is usually limited. For example Adda et al. (2003), Heer and Maussner (2009) and Miranda and Fackler (2004) do not discuss the efficiency of different techniques for solving life-cycle models, even though computation times

[1] I am the single author of all three essays.

can be relatively long in models with many periods. As our model on optimal unemployment insurance in the second essay has many periods we were forced to find a computationally efficient method to solve that model. The first essay is the result of that investigation.

Another issue we encountered is that the literature on solving life-cycle models pays little attention to the preservation of the convexity of the optimization problem, although this is a necessary condition to guarantee a unique solution. The exception is Heer and Maussner (2009), who explicitly preserve the convexity of the problem. However, their solution method is not the most efficient, since it does not use derivatives and relies on linear interpolation, resulting in relative long computation times and low accuracy. The preservation of the convexity is especially important in the model of the second essay, since it has many periods and errors might accumulate over time.

The second essay of this dissertation is called 'Optimal Age-Dependent Unemployment Insurance' and uses one of the methods developed in the first essay. The second essay optimizes unemployment insurance in a life-cycle model in partial equilibrium with wealth accumulation, and job search. More specifically, we consider the optimal constant replacement rate as is standard in the unemployment insurance literature. In addition, we study optimal replacement rates which are allowed to depend linearly on age, which has to our knowledge not been investigated before. We find that the optimal replacement rate is high early in life and drops sharply to reach a level of zero before the age of 45. The welfare gain of making replacement rates dependent on age is somewhat limited though at 5 to 15% of the Net Present Value of all unemployment benefits.

The third essay is called 'Adaptive Beliefs in a Real Business Cycle Model' and investigates how the business cycle properties of a standard macroeconomic model change when agents extrapolate recent changes in productivity into the future. The model in which these adaptive beliefs are applied fits in the Dynamic Stochastic General Equilibrium (DSGE) framework, which is the basis for modern macroeconomics. DSGE models involve dynamic optimization, usually in a model with an infinite horizon[2]. Solving DSGE models results in optimal behavioral rules for different states of the economy, where the behavioral rules usually depend on expectations about the future. The economy of a DSGE model is simulated with a stochastic process, for example shocks to productivity.

Regarding the expectations in macroeconomic models the Rational Expectations (RE) paradigm has dominated the field in the last 4 decades. Rational Expectations mean that agents have full knowledge of the economy, and as such form expectations about the future that are consistent with the (average) actual realizations of the model. Our model deviates from RE and makes use of Bounded Rationality, where agents extrapolate recent changes in productivity into the future. These kind of expectations are so-called 'animal spirits'. The best example of such extrapolative behavior comes from the late nineties where productivity growth was expected to continue in the new millennium (Hirshleifer et al., 2015). There was also data that supported these expectations (Edge et al., 2007), and the overly optimistic expectations may therefore not be attributed entirely to 'animal spirits'.

To analyze the role of adaptive beliefs in the macroeconomy we implement these beliefs in a Real Business Cycle (RBC) model, which is the workhorse model for the study of real variables, and has a labor and capital market. To solve this RBC model with adaptive beliefs we had to rely on a non-standard solution method. The standard solution method for DSGE models is an approximation of the economy around the equilibrium that would occur in the absence of any shocks. This approximation method is a so-called local solution method. In our model agents expect productivity to increase substantially over time, so we have to evaluate the economy far away from the steady state, which warrants the use of a so-called global solution method. We chose to use the Galerkin projection method, which is faster than the more widely used Value Function iteration, at least when good starting values are available.

We calibrate the adaptive beliefs in our RBC model to a measure of consumer

[2]Agents within the model might have a finite life though, such that we get an overlapping generations model.

confidence as both reflect ‘animal spirits’ in the economy. We show that the model with adaptive beliefs can improve the business cycle statistics significantly compared to the Rational Expectations version, but only when the variance of the economic sentiment is high.

1.A Bibliography

Adda, J., R. Cooper, and R. W. Cooper (2003). *Dynamic economics: quantitative methods and applications.* MIT press.

Edge, R. M., T. Laubach, and J. C. Williams (2007). Learning and shifts in long-run productivity growth. *Journal of Monetary Economics 54*(8), 2421–2438.

Heer, B. and A. Maussner (2009). *Dynamic general equilibrium modeling: computational methods and applications.* Springer Science & Business Media.

Hirshleifer, D., J. Li, and J. Yu (2015). Asset pricing in production economies with extrapolative expectations. *Journal of Monetary Economics 76*, 87–106.

Miranda, M. J. and P. L. Fackler (2004). *Applied computational economics and finance.* MIT press.

Chapter 2

Solving Life Cycle Models

2.1 Introduction

This essay shows how to solve a simple life-cycle model with savings. The emphasis is on the different numerical methods that can be employed. Each method will be evaluated in terms of accuracy and computational speed. The accuracy in the approximation of the Value Function is especially important when there are many periods, because errors can accumulate over time, and eventually make the approximation of the Value Function non-concave.

The numerical methods are evaluated in a partial equilibrium model without risk, such that an analytical solution can be obtained. This allows us to calculate the absolute accuracy of the numerical solutions. In addition, the interest rate R is set at the inverse of the time preference ($R = \frac{1}{\beta}$), which implies that a constant level of consumption is optimal. This makes it very easy to calculate the analytical solution. It should be noted that in overlapping generation models with General Equilibrium we usually have $r \neq \frac{1}{\beta} - 1$, so consumption will not be constant over the life-cycle, even without any risk or borrowing constraints.

The described methods are general, so they can also be used in general equilibrium models and models with idiosyncratic risk, like income shocks or mortality risk. We describe four methods, all of which use an interpolation scheme that preserves the monotonicity and concavity of the Value Function. The first three methods only use

Lagrangian data, which means no information on the slope of the Value Function is used for the interpolation. The first method, Golden Section search, is easy to program, since it does not require any derivatives. The second method is Constrained Optimization with Lagrangian data, and the third method is Solving First Order Conditions. Both the second and third method use constraints and derivatives to find the optimum. Finally, we describe Constrained Optimization with Hermite data, which means also the slope of the Value Function is approximated and used for the interpolation. This approximation is obtained by adding an extra constraint to the optimization problem (Cai and Judd, 2012, 2015).

2.2 Model

The model is a simple life-cycle model where an agent gets income S_t in period t. The agent maximizes his life-time utility:

$$\max \sum_{t=1}^{T+D} \beta^t u\left(c_t\right) \tag{2.1}$$

where T and D are the number of period's in working life and retirement, respectively, β is the discount factor, c_t is period t consumption, which gives utility $u\left(c_t\right)$. The utility function is assumed to be monotonically increasing and concave to ensure a unique solution. In our Matlab program we allow for both CARA and CRRA utility functions. With CARA the utility function is $u\left(c_t\right) = \frac{-\exp(-\nu c_t)}{\nu}$ and with CRRA it is $u\left(c_t\right) = \frac{(c_t+\epsilon)^{1-\nu}}{1-\nu}$, where ϵ is a small number (set to 0.01) to ensure the marginal utility of consumption does not go to infinity when $c_t = 0$, which could result in large errors in the interpolation. The parameter ν is the degree of risk aversion.

The constraints for this optimization problem are:

$$Rk_t + S_t \geq c_t + k_{t+1} \tag{2.2}$$

$$k_{t+1} \geq 0 \tag{2.3}$$

$$c_t \geq 0 \tag{2.4}$$

which are the budget constraint, borrowing constraint and non-negativity constraint on consumption, respectively. The interest rate is R, which is paid over wealth at the beginning of the period (k_t). The initial wealth in period 1 is given, with $k_1 \geq 0$. During working life ($t \leq T$) the agent's income is the wage ($S_t = w$) and during retirement ($t > T$) the agent gets a pension ($S_t = p$). Both the wage and the pension are known with certainty, and we impose that the wage is at least as high as the pension, $w \geq p$, such that neither the borrowing constraint nor the non-negativity constraint on consumption ever binds. The lack of uncertainty and absence of binding constraints allows us to compute the analytical solution, and evaluate the absolute accuracy of the numerical solutions.

2.3 Analytical solution

To calculate the accuracy of the numerical methods the analytical solution is computed at the same gridpoints as in the numerical solutions. For every period t and gridpoint i we calculate the constant consumption flow $\bar{c}_{i,t}$ that the agent can afford given his wealth $k_{i,t}$, where i is the indicator of the gridpoints with $i = 1, ..., I$. The Net Present Value of his wealth and income in the remaining periods of his life is:

$$Q_{i,t} = k_{i,t} + \sum_{s=t}^{T+D} \frac{1}{R^{T+D-s+1}} S_t \tag{2.5}$$

which allows him to buy the constant consumption flow:

$$\bar{c}_{i,t} = Q_{i,t} / \sum_{s=t}^{T+D} \frac{1}{R^{T+D-s+1}} \tag{2.6}$$

His total utility $V(k_{i,t})$ from period t onwards can be calculated as:

$$V(k_{i,t}) = \sum_{s=t}^{T+D} \beta^{T+D-s} u(\bar{c}_{i,t}) \tag{2.7}$$

Given this Value Function we can calculate the slope of the Value Function analytically:

$$\frac{\partial V}{\partial k_t} = \frac{\partial V}{\partial \overline{c}}\frac{\partial \overline{c}}{\partial k_t} = \sum_{s=t}^{T+D} \beta^{T+D-s} u'\left(\overline{c}\right) / \sum_{s=t}^{T+D} \frac{1}{R^{T+D-s+1}} \tag{2.8}$$

which is equal to $Ru'\left(\overline{c}\right)$ if $R = \frac{1}{\beta}$ as is the case in our parameter setting.

2.4 Standard numerical solution

The basic method for solving a life-cycle model with savings is as follows (part of terminology taken from Cai and Judd, 2015):

- Start in the last period alive, $t = T + D$, and work backwards to the first period of the model.
- For every period we **define a grid** for wealth with I grid points: $k_{i,t} \in [k_{min,t}, k_{max,t}]$ for $i = 1, ..., I$.
- **Maximization Step**: choose the capital stock in the next period (or consumption in this period) that maximizes (expected) utility, given current capital $k_{i,t}$. For this we transform the optimization problem to a Bellman equation:

$$V\left(k_{i,t}\right) = \max_{k_{t+1}} u\left(Rk_{i,t} + S_t - k_{t+1}\right) + \beta V\left(k_{t+1}\right) \tag{2.9}$$

  where $V\left(k_t\right)$ is the Value Function that gives the maximum expected life-time utility given wealth k_t. The Value Function is approximated in the next step.

- **Fitting Step**: since there is usually no analytical solution for the Value Function it is approximated by the solution at the gridpoints i, using interpolation to get the Value Function between those gridpoints. The interpolation is usually done with a spline or polynomial, where it is crucial that the interpolation method preserves the monotonicity and concavity of the Value Function. Shape preservation is crucial, since the problem should be convex to ensure a unique solution. If the shape of the Value Function is not preserved multiple solutions might exist and large errors might occur.

Defining a grid When no analytical solution can be obtained we have to make a guess for the size of the grid. The size of the grid should be such that we never go outside the boundaries, $[k_{min,t}, k_{max,t}]$, during either the Maximization Step or the simulation of our economy, because we are not able to preserve the shape of the Value Function when we extrapolate outside the grid. Getting a good guess for the grid increases the efficiency in terms of accuracy and computation time, since a narrower grid makes the solution more accurate or we can use less gridpoints for the same level of accuracy. We describe two methods for obtaining a grid that is neither too wide nor too narrow.

The first method is to start with a grid that is very wide, solve the model as described above, run a simulation and then shrink the grid based on the maximum level of the state variable(s) that occurred during the simulation. When several variants of the model have to be calculated, for example for the optimization of tax policy, it can be useful to keep the grid the same for all variants. In that case a sufficient margin should be added to the boundaries such that all variants stay within the grid during simulations.

The second method is to guess a maximum of the state variables based on a simpler version of the model that (1) has an analytical or fast numerical solution, and (2) has a path of the state variables that is higher than our economy. The path of the state variables from this solution is then our first guess for the upper bound of the grid. In our model we could for example calculate the wealth of an agent that has (a) a slightly higher income during working age, and (b) zero income during retirement. The wealth path of that agent will certainly be above the wealth path of an agent in our economy and should thus be a robust guess for the upper bound of the grid.

For both methods of obtaining a grid we recommend to use a small growth factor to the upper bound, such that we stay strictly within the boundaries of the grid, and avoid extrapolating outside the grid. For example, if the upper bound from either of the two methods is defined as $k_{max,t}$ we could adjust this boundary to $k^{adj}_{max,t} = (1 + m)^t \, k_{max,t}$ with m larger than zero, but small.

Methods Maximization Step We will review the following solution methods for the maximization step:

- Golden Section Search, which uses Lagrangian data
- Constrained Optimization with Lagrangian data
- Solve First Order Conditions, which uses Lagrangian data
- Constrained Optimization with Hermite data

Methods Fitting Step We describe two methods for the Fitting Step:

- Lagrangian data: use only the level of the Value Function at the gridpoints for the interpolation
- Hermite data: use both the level and the slope of the Value Function at the gridpoints for the interpolation

2.5 Numerical solution: Lagrange Value Function Iteration (L-VFI)

We have borrowed heavily from the description of Lagrangian Value Function Iteration from Cai and Judd (2015). Value Function iteration consists of three steps. In the Initialization Step we choose a grid and a functional form to approximate the Value Function. In the Maximization Step we solve the optimization problem. Finally in the Fitting Step we approximate the Value Function using the chosen interpolation method. With Lagrangian Value Function Iteration we only use data on the level of the Value Function at each gridpoint to approximate the Value Function. This contrasts with Hermite Value Function Iteration where we also use information on the slope of the Value Function at the gridpoints to approximate the Value Function (see Section 2.6).

2.5.1 Initialization Step

In the initialization step we choose approximation nodes $\mathbb{K}_t = \{k_{i,t} : 1 \leq i \leq I\} \subset \mathbb{R}$ for every $t \leq T + D$, and we choose a functional form for $\hat{V}(k_t; \mathbf{b}_t)$ where $\mathbf{b}$ is a vector of parameters (of a spline or polynomial for example).

2.5.2 Standard Lagrange Maximization Step

We solve for the Bellman equation (2.9) at each gridpoint for every period, starting in the last period and working backwards in time. For each gridpoint $k_{i,t}$ we have reformulated the problem slightly by inserting the budget constraint directly into both the utility function and the non-negativity constraint on consumption:

$$v_i = \max_{k_{t+1}} u\left(Rk_{i,t} + S_t - k_{t+1}\right) + \beta\hat{V}\left(k_{t+1}; \mathbf{b}_{t+1}\right) \tag{2.10}$$

subject to:

$$k_{t+1} \geq 0 \tag{2.11}$$

$$Rk_{i,t} + S_t - k_{t+1} \geq 0 \tag{2.12}$$

with Lagrangian multipliers μ_t and χ_t, respectively. The latter constraint is the non-negativity constraint on consumption. These constraints define a lower and upper bound for k_{t+1}. The First Order Condition (FOC) with respect to k_{t+1} is then:

$$u'\left(Rk_{i,t} + S_t - k_{t+1}\right) = \beta\frac{\partial\hat{V}\left(k_{t+1}; \mathbf{b}_{t+1}\right)}{\partial k_{t+1}} + \mu_t - \chi_t \tag{2.13}$$

where $\chi_t = 0$ since the upper bound (equation 2.12) should not bind as $c_t = 0$ is never optimal, and $\mu_t = 0$ as the non-negativity constraint (equation 2.11) will not bind either, unless we are at lower bound ($k_{i,t} = 0$) and income is at it's lowest level (when we are in retirement in our model). In the latter case it is optimal to consume total income: $c_t = S_t$.

Golden Section Search

Golden Section Search maximizes a univariate objective function by shrinking the interval in which the optimum solution is known to lie, until the interval is considered small enough. The algorithm for Golden Section Search (as described by Heer and Maussner, 2009) is very easy to program, since we do not need to calculate derivatives. The algorithm is relatively fast given it's simplicity. Especially if the optimization toolbox of Matlab is not available this algorithm is good alternative. For good computation times it is crucial to exploit that the optimal solution for k_{t+1} is monotonically increasing in current capital, k_t.

Golden Section Search requires a lower and upper bound as inputs, which makes the method less suitable when these bounds can not be determined analytically. For example, if the labor supply is endogenous then the maximum consumption level depends on the labor supply. We would thus first have to solve for the maximum possible level of consumption before we can apply the Golden Section Search. This decreases the computational efficiency. For such optimization problems a Constrained Optimization algorithm that can handle non-linear constraints is more suitable.

The idea of Golden Section Search is to find the minimum (or maximum) of a function $f(x)$ on some interval $[A, D]$ and make the interval smaller and smaller until the difference between A and D is considered small enough. To shrink the interval we choose two additional points B and C that divide the interval such that we maximize the shrinkage of the interval in each iteration. Given points A and D we calculate $B = pA + (1-p)D$ and $C = (1-p)A + pD$, where p is the Golden Section ratio: $p = \frac{\sqrt{5}}{2} - \frac{1}{2} \approx 0.618$ (see Heer and Maussner, 2009, for the derivation). We have slightly adjusted the algorithm from Heer and Maussner (2009) such that it returns the boundary solution if the boundary is the maximum. This adjustment allows us to check whether the interval was set wide enough, since we should get an interior solution in our model, except at the lower bound during retirement. The Golden Section algorithm is described in Section 2.A.

We loop over gridpoints $k_{i,t}$ and find the optimal capital stock in the next period, using (2.10) as the objective. We implement the constraints as lower and upper boundaries for k_{t+1} given gridpoint $k_{i,t}$, which are denoted as $\underline{k}_{i,t+1}$ and $\overline{k}_{i,t+1}$, respectively. Since we know that the optimal solution k^*_{t+1} is increasing with k_t we adjust the lower bound upward as we loop over $k_{i,t}$. We set the lower bounds as follows:

$$\underline{k}_{i,t+1} = \begin{cases} k_{min,t+1} & \text{if } i = 1 \\ k^*_{i-1,t+1} & \text{if } i > 1 \end{cases} \tag{2.14}$$

The upper bound is the capital stock in the next period if consumption is null, so $\overline{k}_{i,t+1} = Rk_{i,t} + S_t$.

Constrained Optimization

The second method to maximize the objective is Constrained Optimization for which we use Matlab's '*fmincon*'. In our case the constraints are a simple lower and upper bound for k_{t+1}. Since our constrained optimizer uses (approximations) of the derivatives we need to approximate the Value Function with a continuously differentiable function ($\mathbb{C}^1$) (see Subsection 2.5.3 for more details). The derivatives can be calculated analytically or with finite differences [1]. We use analytical derivatives as they are much faster, and they are easy to calculate in our model. For more complex models the use of finite difference approximations is usually more convenient. The analytical derivatives with respect to k_{t+1} are given in equation (2.13). Since we approximate the Value Function with a continuously differentiable function the term $\frac{\partial \hat{V}(k_{t+1};\mathbf{b}_{t+1})}{\partial k_{t+1}}$ can be calculated either with finite differences or by taking the derivative of the approximation of the Value Function. We use a quadratic spline to approximate the Value Function, so we can easily calculate the slope of the Value Function by differentiating this spline[2].

---

[1]Note that central difference approximation is much more accurate than the default forward difference approximation used by Matlab, but central differences are more expensive to calculate.

[2]Matlab has the '*fnder*' procedure in the '*Curve Fitting Toolbox*' to get the slope of piecewise polynomials, which are created using '*mkpp* and can be evaluated with '*ppval*'.

Constrained Optimization is not the most efficient method given the simplicity of our model, but Constrained Optimization will be the most efficient option when the constraints are non-linear. For example, with endogenous labor supply the maximum consumption level will depend on the labor supply. Such a constraint is easily be implemented in and efficiently handled by constrained optimizers, like Matlab's '*fmincon*'.

Solving First Order Conditions

Since the constraints are a simple lower and upper bound an alternative would be to solve directly for the First Order Condition (FOC) of the objective function, using a constrained solver that can handle boundaries[3], like Matlab's '*lsqnonlin*'[4]. In this case we solve directly for:

$$u'\left(Rk_{i,t} + S_t - k_{t+1}\right) = \beta \frac{\partial \hat{V}\left(k_{t+1}; \mathbf{b}_{t+1}\right)}{\partial k_{t+1}} \tag{2.15}$$

where the solver handles the upper and lower bound given by equations (2.11) and (2.12). The left hand side can be calculated analytically, and the right hand side is calculated as the derivative of the quadratic spline that we use to approximate the Value Function (see the previous paragraph on 'Constrained-Optimization').

2.5.3 Fitting Step

For all algorithms we need an interpolation method to obtain the Value Function $V(k_t)$ between gridpoints. For this we fit a (piece-wise) spline or polynomial given the solutions at the gridpoints $V(k_{i,t})$. The interpolation method should preserve the monotonicity and concavity of the value function. For the Constrained Optimization

[3] As mentioned before the lower bound for k_{t+1} will only bind when we are on the lower bound of k_t and in retirement. The upper bound will never bind.

[4] With Matlab's '*lsqnonlin*' we can use "vectorization", which means we solve for all gridpoints in period t simultaneously and greatly improves speed. For this we use the '*JacobPattern*' option to set the Jacobian matrix to a sparse matrix with ones on the diagonal, and zeros everywhere else.

and Solving the First Order Conditions we need the function approximation of $V(k_t)$ to be a continuously differentiable function ($\mathbb{C}^1$) as well.

A well-known univariate interpolation algorithm is the Schumaker (1983) shape preserving quadratic spline, which uses knot insertion. However, this algorithm with Lagrangian data (i.e. no information on slope) does not always preserve monotonicity. For this reason we use the improved version by Lam (1990), which always preserves both monotonicity and concavity, and is continuously differentiable as well[5].

2.6 Numerical solution: Hermite Value Function Iteration (H-VFI)

With Lagrangian methods described above we only use information on the level of the Value Function at the gridpoints to approximate the Value Function. With Hermite interpolation we also calculate the slope of the Value Function to get a more accurate approximation of the Value Function. To obtain the slope we use the method described in Cai and Judd (2015). This method requires a constrained optimization algorithm which reports the Lagrangian multiplier on the constraints, for which we used Matlab's '*fmincon*' .

2.6.1 Hermite Maximization Step

As with the Lagrangian Maximization Step we solve the Bellman equation (2.9) at each gridpoint, starting in the last period and working backwards in time. This time we change the optimization problem such that the state variable $k_{i,t}$ is replaced with y_i in the objective function, and we impose the equality constraint $y_i = k_{i,t}$ such

[5] The shape preserving spline of Lam (1990) has the same shape as the interpolation algorithm described McAllister and Roulier (1981), but the algorithm used by Lam seems simpler to implement. This interpolation scheme is added in the the toolbox as '*lam_mcp_spl*'.

that the solution is not affected. Our optimization problem is thus:

$$v_i = \max_{y,k_{t+1}} u\left(Ry + S_t - k_{t+1}\right) + \beta\hat{V}\left(k_{t+1}; \mathbf{b}_{t+1}\right) \tag{2.16}$$

subject to:

$$\begin{aligned} k_{t+1} &\geq 0 \\ Ry + S_t - k_{t+1} &\geq 0 \\ k_{i,t} - y &= 0 \end{aligned} \tag{2.17}$$

with Lagrangian multipliers μ_t, χ_t and ψ_t, respectively. The Lagrangian multiplier on the equality constraint, ψ_t, will be equal to the slope of the Value Function at each gridpoint, $\frac{\partial V(k_{i,t})}{\partial k_{i,t}}$. This Lagrangian multiplier is reported by most algorithms that solve constrained optimization problems, such as Matlab's '*fmincon*'. Knowing the value of $\frac{\partial V(k_{i,t})}{\partial k_{i,t}}$ allows us to use Hermite data in the Fitting Step, which should increase the accuracy of the approximation of the Value Function.

The extra costs in computational time will be limited if adding the extra constraint (equation (2.17)) does not affect computational efficiency much. For our model we have to change the algorithm from the fast '*trust-region-reflective*' to the slower '*sqp*' algorithm when we implement the Hermite method, since the '*trust-region-reflective*' algorithm can not handle linear constraints and boundaries at the same time. As the '*sqp*' algorithm is much slower the Hermite method results in considerably longer computation times than the Lagrangian method, so the use of Hermite data would have to improve the accuracy significantly to make up for the longer computation time. However, in a model with a non-linear constraint, for example due to endogenous labor supply, we already have to rely on a slower algorithm that can handle non-linear constraints and boundaries. In that case adding the extra constraint (2.17) will hardly affect computation times, and makes the use of Hermite Value Function iteration more appealing.

2.6.2 Fitting Step

To approximate $\hat{V}(k_t; \mathbf{b}_t)$ we use both the levels, v_i, and the slopes, $s_i = \frac{\partial V(k_{i,t})}{\partial k_{i,t}}$, at each gridpoint and fit a spline or polynomial. The spline or polynomial should use a continuously differentiable function ($\mathbb{C}^1$) that preserves the monotonicity and concavity of the Value Function. Once the slopes at the gridpoints are known the method by Schumaker (1983) is a shape preserving quadratic spline. We use the algorithm of Lam (1990), which is the same as Schumaker's algorithm once the slopes are determined. Another possibility would be the use of a rational spline for which Cai and Judd (2012) describe an algorithm[6].

2.7 Performance Results

In this section we report the accuracy and computation times of the following methods:

- Lagrangian data using the three types of solution methods: Golden Section Search (GSS), Solving FOCs (FOC), and Constrained Optimization (CON);
- Hermite data using Constrained Optimization (HER).

We use the parameters as listed in Table 2.1. We have used two types of grids, one with equidistant gridpoints, and one with more gridpoints close to the lower bound, called "Transformed". The latter is especially recommended when there is a binding constraint at the lower bound. This will for example occur when there is income risk in combination with a borrowing constraint. The Value Function will then be steep close to the lower bound as the marginal utility of consumption is high, so increasing the number of gridpoints there should increase the accuracy.

In Tables 2.2 and 2.3 we report the computation time and the accuracy. The accuracy measure is the maximum absolute deviation of consumption from the analytical

[6] Note that calculating the derivative of a rational spline is a bit more involved than calculating the derivative of a quadratic spline as we did.

Table 2.1: Parameters

	Value
Risk aversion (ν)	2.0
ϵ in utility function	0.01
Years working (T)	480
Years retirement (D)	180
Wage (w)	1.00
Pension (p)	0.10
Initial wealth (k_0)	0.00

solution for consumption, at the gridpoints in any period. To lower computation times for the derivative based methods (all methods except Golden Section Search) we have used analytical gradients and "vectorization". Using analytical gradients is more complicated to implement, but is more accurate and can decrease computation times significantly.

With "vectorization" in Matlab we mean that we solve the maximization problem for all gridpoints in any period simultaneously, which can also reduce computation times significantly. With vectorization the objective function is the sum of the Value Function at all gridpoints (in period t). Since we used an analytical gradient and vectorization we can not decrease computation times using parallel computing in our derivative based methods. When we estimate the gradient using numerical approximation (finite differences) parallel computing could improve computation times.

The use of parallel computing prevents us from exploiting the convexity of our problem, so it is not *a priori* clear if we can improve computation times with parallel computing, even when our maximization problem is not vectorized. Exploiting the convexity of the problem means that we know that the optimal solution of k_{t+1} is increasing with k_t, so when we loop over $k_{i,t}$ we can revise the lower bound or initial guess for k_{t+1} up as described in Subsection 2.5.2 under Golden Section Search. When we solve for each gridpoints consecutively the upward revision of the initial guess or lower bound will reduce computation times. This is not possible with parallel computing as we start the optimization for all gridpoints simultaneously. For our model and setup with 2 computational cores it turns out that Golden Section Search with a loop over the grid and exploiting the convexity results in lower computation times than the use of parallel computing.

2.7.1 Computation times

The first thing to note is that Solving First Order Conditions (FOC) will be very fast when we only have a lower and upper bound as constraints, and we can use an analytical gradient. Constrained Optimization (CON) will already be significantly slower, and will only be the better choice when there are more constraints than a simple lower and upper bound. With non-linear constraints Constrained Optimization (CON) is our preferred method, since it will handle those constraints efficiently. Golden Section Search is relatively fast, and is easy to program, so this is a good alternative in case the lower and upper bound can be determined. It should be noted that computation times with Golden Section search will depend more on the error tolerance (ε in Section 2.A) than the derivative based algorithms, because the number of iterations with Golden Section Search will increase directly when we lower the tolerance. With the derivative based methods the number of iterations does not necessarily increase when we lower the tolerance.

The Hermite Value Function iteration (HER) is considerably slower than the other alternatives, but the difference with Constrained Optimization is mostly due to the choice of the algorithm. With the Hermite method Matlab's '*sqp*' algorithm will be the fastest, since we can not use the '*trust-region-reflective*' algorithm, and the '*interior-point*' will be slow. The '*interior-point*' algorithm will be slow in combination with the Hermite method, because the extra constraint on k_t, equation (2.17), will always bind and the '*interior-point*' algorithm will converge relatively slow to a solution with a binding constraint. If we can not use '*trust-region-reflective*' for the Constrained Optimization then the Hermite method will only be marginally slower. For example, with Lagrangian Constrained Optimization and the '*interior-point*' algorithm (the second fastest method) the computation time is already roughly equal to the computation time with Hermite Constrained Optimization.

Table 2.2: Accuracy and computation time, yearly model, 100 nods

	Equidistant nodes		Transformed	
	Max. error C_t ($\cdot 10^{-3}$)	Time (s.)	Max. error C_t ($\cdot 10^{-3}$)	Time (s.)
GSS	0.0488	3.28	0.3208	3.05
FOC	0.0490	0.46	0.3212	0.47
CON	0.3896	3.59	0.3212	3.71
HER	0.0716	13.98	0.1633	20.52

Table 2.3: Accuracy and computation time, monthly model, 25 nods

	Equidistant nodes		Transformed	
	Max. error C_t ($\cdot 10^{-3}$)	Time (s.)	Max. error C_t ($\cdot 10^{-3}$)	Time (s.)
GSS	2.2033	13.11	3.8588	12.31
FOC	2.2031	3.98	3.8577	3.99
CON	2.2043	9.09	3.8577	8.72
HER	0.4331	13.94	1.1716	16.70

2.7.2 Accuracy of results

The Tables 2.2 and 2.3 show that all Lagrangian methods have similar errors. This implies that the size of the errors depends mostly on the interpolation scheme, and not the algorithm. This indicates that the accuracy of the algorithms was set sufficiently high. The only exception is the Constrained Optimization in the yearly model with equidistant nodes. The relatively large errors are caused by Matlab's '*trust-region-reflective*' algorithm, since tests with other algorithms for Constrained Optimization result in errors of similar magnitude as the GSS and FOC methods.

The Hermite method will in general be more accurate, and would allow the use of less gridpoints to get accuracy levels similar to that of the other methods. In Table 2.2 we see that the Hermite solution results in a lower accuracy compared to the GSS and FOC methods in the yearly model with equidistant nodes, although the error is still relatively small.

To analyze the cause of this relatively large error we have looked at the solution in the one but last period of life, $t = T + D - 1$. When we use the analytical derivative for $\frac{\partial V(k_{t+1})}{\partial k_{t+1}}$ instead of the slope of the approximated Value Function $\frac{\partial \hat{V}(k_{t+1};\mathbf{b}_{t+1})}{\partial k_{t+1}}$ the errors are very close to zero. This indicates that the errors originate from the interpolation method, and not the optimization. We may thus conclude that Lagrangian interpolation accidentally does better than Hermite interpolation with the equidistant grid in the yearly model.

It should be noted that the Hermite Value Function iteration could gain in accuracy compared to the Lagrangian methods when the solution is less regular than in our simple model. For example, when there is income uncertainty and the borrowing constraint binds the slope of the Value Function will be steep close to the lower bound. Since the Lagrangian method will estimate the slope using a 'harmonic mean', an irregular shape of the Value Function might result in a relatively lower accuracy for the Lagrangian methods (compared to the accuracy differences between Lagrangian and Hermite methods reported for our model).

2.8 Concluding Remarks

We have tested the accuracy and speed of four methods to solve a one-dimensional life-cycle model, for which we are also able to calculate the analytical solution. For a simple problem like ours, where the only constraints are boundaries, solving for the First Order Conditions under these constraints is by far the fastest method, without compromising on accuracy. If computation times are not important Golden Section Search is a simple, but robust alternative, since we do not need to approximate derivatives. In fact, we can use Golden Section Search in combination with linear interpolation and still preserve the shape in terms of concavity and monotonicity (see Heer and Maussner, 2009, for example). For problems with non-linear constraints Constrained Optimization will be the most efficient method. In these cases the use of Hermite interpolation will not add much to the computation time, while it should improve the accuracy.

Binding constraints Since the borrowing constraint did not bind (except in retirement) the Value Function has a regular shape, which makes it easier to approximate the Value Function using interpolation. When constraints bind, for example at the lower bound in case of uncertainty, the slope of the Value Function will be much steeper close to the boundary. For this reason it will be recommended to use a transformation of the grid, such that we have more gridpoints closer to the lower bound (see also Cai and Judd, 2015).

In addition, having a binding constraint might also increase the accuracy of the Hermite interpolation, since it will directly calculate the slope of the Value Function at the lower bound. With the other solution methods the slope of the Value Function will be determined by the gridpoints next to the bound, which may result in large inaccuracies if the irregular shape of the Value Function is poorly approximated without information on the slope.

Higher dimensional problems Our model had only one state variable, for which we used I gridpoints. When solving higher dimensional problems the number of gridpoints increases rapidly with the number of dimensions. If we use I gridpoints in k dimensions then the total number of gridpoints is I^k, which we call the "curse of dimensionality". This curse makes the Hermite interpolation method more attractive, since we should be able to get a more accurate approximation of the Value Function with less gridpoints in each dimension. Cai and Judd (2015) describe Hermite interpolation in a higher dimensional problem using Chebyshev polynomials to approximate the Value Function.

Some caution is recommended though with the use of their method. First of all, their method does not preserve the shape of the Value Function perfectly, so the problem might become locally non-convex, especially when the errors accumulate in a model with many periods. Second, the polynomial will not go through all data points, which also leads to inaccuracies that might accumulate over time, especially close to the boundaries of the grid. Unfortunately, to our knowledge there is no interpolation method for higher dimensional problems that preserves the shape and goes through all gridpoints. For that reason the Chebyshev regression that Cai and Judd (2015) employ with high order polynomials (say order 40) seems the only feasible method for multi-dimensional problems.

2.A Golden Section Search algorithm

The algorithm as described by Heer and Maussner (2009) is slightly adjusted. We add two variables which track if the initial lower and upper bound of 'Step 1' have been adjusted. If the lower bound is adjusted at least once we set $\Phi_{lb} = 1$ and if the upper bound is adjusted at least once we set $\Phi_{ub} = 1$. If the lower (upper) bound was never adjusted we add the original lower (upper) bound to the points from which we select the minimum, such that the algorithm returns the lower (upper) bound if that is the optimum.

The adjusted algorithm to find the minimum of $f(x)$ in interval $x \in [\underline{x}, \overline{x}]$ is:

Step 1: Initialization

$A = \underline{x}$, $D = \overline{x}$

$$B = pA + (1-p)D$$
$$C = (1-p)A + pD$$

and set $fB = f(B)$ and $fC = f(C)$.

Set $\Phi_{lb} = 0$, $\Phi_{ub} = 0$.

Step 2: Adjust interval if $|D - A| \geq \varepsilon$

If $fB < fC$:

$$\begin{aligned} C &= B \\ fC &= fB \\ D &= C \\ fD &= fC \\ B &= pC + (1-p)\,A \\ fB &= f\,(B) \\ \varPhi_{ub} &= 1 \end{aligned}$$

Else:

$$\begin{aligned} A &= B \\ fA &= fB \\ B &= C \\ fB &= fC \\ C &= pB + (1-p)\,D \\ fC &= f\,(C) \\ \varPhi_{lb} &= 1 \end{aligned}$$

Step 3: Choose minimum if $|D - A| < \varepsilon$

- If $\Phi_{lb} = 0$: set $fA = f\,(A)$
- If $\Phi_{ub} = 0$: set $fD = f\,(D)$

Choose the minimum from fA, fB, fC and fD, and return A, B, C or D as the optimum.

2.B Bibliography

Cai, Y. and K. L. Judd (2012). Dynamic programming with shape-preserving rational spline hermite interpolation. *Economics Letters 117*(1), 161–164.

Cai, Y. and K. L. Judd (2015). Dynamic programming with hermite approximation. *Mathematical Methods of Operations Research 81*(3), 245–267.

Heer, B. and A. Maussner (2009). *Dynamic general equilibrium modeling: computational methods and applications.* Springer Science & Business Media.

Lam, M. (1990). Monotone and convex quadratic spline interpolation. *Virginia Journal of Science 41*(1), 3–13.

McAllister, D. F. and J. A. Roulier (1981). An algorithm for computing a shape-preserving osculatory quadratic spline. *ACM Transactions on Mathematical Software (TOMS) 7*(3), 331–347.

Schumaker, L. I. (1983). On shape preserving quadratic spline interpolation. *SIAM Journal on Numerical Analysis 20*(4), 854–864.

Chapter 3

Optimal Age-Dependent Unemployment Insurance

3.1 Introduction

This research sets out to determine the optimal level of Unemployment Insurance (UI) benefits, where the level of benefits are allowed to vary with age. Providing Unemployment Insurance has the traditional public policy trade-off between equity and efficiency. On the one hand higher unemployment benefits increase equity as it provides agents with insurance against a (temporary) fall in income. On the other hand, it reduces efficiency because it weakens the incentives to take up a job when unemployed.

There are empirical results that indicate that an increase in unemployment benefits increases unemployment duration (see Krueger and Meyer, 2002). As Chetty (2008) notes this has traditionally been interpreted as an increase in moral hazard as higher unemployment benefits reduce the relative benefits of having a job, distorting search effort and resulting in higher unemployment rates. Chetty (2008) then continues that the increase in unemployment duration as a result of higher UI benefits may not be entirely attributed to moral hazard. When agents have little liquid assets their search effort may be too high from a social point of view. Providing the unemployed with liquidity to smooth consumption can increase social welfare, even though unemployment duration increases. This is very relevant for the design of optimal Unemployment Insurance in the U.S. as almost half of the unemployed report zero liquid wealth (Chetty, 2008).

This trade-off between insurance and giving the right incentives is likely to change over the life-cycle for various reasons. In the first place, young individuals have less assets and are less able to self-insure against income shocks. Empirical research has indeed found that unemployment duration increases with wealth (see Lentz, 2009; Lentz and Tranaes, 2005; Bloemen and Stancanelli, 2001; Algan et al., 2003). Second, wages are known to be age-dependent (see for example Hanse, 1993; Rupert and Zanella, 2015) which will influence the benefits of finding a job, and thus search effort of the unemployed. That wages influence search behavior has been shown by Lentz (2009) for Denmark. Furthermore, the closer the unemployed get to their retirement age, the less long they will benefit from re-employment. This will reduce search effort. Finally, older individuals might have more difficulties finding a job, because employers might prefer younger employees.

Such age dependent factors imply that age-dependent unemployment benefits might be welfare improving. To test this we develop a life-cycle model with wealth accumulation and job search. All three factors are important to include in the model.

The life-cycle perspective is crucial to capture age-specific effects. Another difference between a life-cycle model and a model with infinitely living agents is that liquidity insurance through loans is more easy to implement with infinitely lived agents as they have infinite time to repay their debts. In fact, models with infinitely-lived agents and an interest rate equal to or higher than one divided by the discount factor have the odd feature that agents will accumulate infinite assets as a means of self-insurance (Chamberlain and Wilson, 2000). In addition, the life-cycle approach is important because discounted life-time utility is likely to vary more than in a model with infinitely lived agents. As (discounted) income varies more between agents in a life-cycle model redistribution is more important.

The second aspect, search effort, is important as Lentz (2009) argues. Empirical evidence from Devine and Kiefer (1991) shows that rejection or acceptance of job offers plays a small role in the observed unemployment hazard rate variation and indicates that search intensity is the most important driver for unemployment hazard. In our model search effort will depend on the difference in (discounted) utility between being employed and being unemployed. This difference in utility is likely to depend strongly on the phase of the life-cycle.

The third feature of the model, the accumulation of assets, is important to allow agents to self-insure against income shocks through pre-cautionary savings as insurance against unemployment is imperfect. In our model agents will save for two reasons. The first is retirement, since agents are not provided with sufficient pension benefits. The second motive for savings is precaution as the agent can use his assets to smooth consumption over his life.

Regarding savings the level of the interest rate is very important as Lentz (2009) found that the ratio between the interest rate and the discount factor of agents has strong effects on the optimal level of unemployment benefits. To be able to set the interest rate freely we choose the partial equilibrium approach, where both wages and interest rates are given and independent of labor supply and wealth accumulation, respectively.

Despite the importance of a life-cycle approach, asset accumulation and search effort, the combination of these three aspects in one model is rare, and has to our knowledge only been used in Costain (1997) and Heer (2003). Our model extends their approach with a more extensive calibration, and in addition considers optimal age-dependent UI benefit levels. The extensive calibration includes three measures that are non-standard. These are the elasticity of unemployment duration with respect to the level of unemployment benefits, the average consumption drop when an agent becomes unemployed and the average duration of unemployment, with a linear trend in age. The combination of these factors should give the model a good fit with respect to the trade-off between insurance and the negative impact of insurance on search effort.

The extra parameters that are calibrated for are the disutility of working, adjustment costs to private wealth when the change in private wealth is negative, and the matching probability given search effort, where the parameter that determines the matching probability is a linear function of age.

The unemployment insurance instrument that we evaluate is the net replacement rate, meaning the ratio of UI benefits to net wages as is standard in the literature (see for example Wang and Williamson, 2002; Davidson and Woodbury, 1997). We use two variants for this. One variant where the replacement rate is fixed over the life-cycle and one where the replacement rate is allowed to linearly depend on age. In both cases the maximum duration is fixed at 6 months. It turns out that making replacement rates dependent on age can improve welfare, with welfare gains of between 5 to 15% of the Net Present Value of UI contributions.

3.2 Model

The model is a life-cycle model with search effort and wealth accumulation, in partial equilibrium. The wage is allowed to be age specific, and the labor supply when employed is fixed at 1 unit. The UI policies that will be analyzed in this paper consist of unemployment benefits with a maximum duration of 6 months, after which the agent gets (close to) nothing as is standard in the literature.

Agents maximize utility:

$$E_0 \left\{ \sum_{t=1}^{T+D} \beta^{t-1} \left[u\left(c_t\right) - d_t - a_t \right] \right\} \tag{3.1}$$

where E_0 is the expected value before the start of the first period, $u\left(c_t\right)$ is the period t utility from consumption c_t, T is the number of periods in working age, and D is the number of periods in retirement, β is the discount factor, d is the (fixed) discomfort of working and a is the job search intensity with the latter two being zero during retirement.

We abstract from uncertainty about the time of death and we also abstract from a bequest motive. In addition we will set the gross interest rate equal to one over the discount factor: $R = 1/\beta$. The combination of certainty about the time of death, lack of bequest motives and $R = 1/\beta$ ensures that agents prefer a constant consumption level. In models with stochastic survival probabilities the interest rate is usually set higher than one over the discount factor ($R > 1/\beta$), which results in higher savings compared to our model with $R = 1/\beta$.

3.2.1 Retired individuals

The retired individuals optimize (3.1) for $t = [T+1, ..., T+D]$ subject to the following two constraints:

$$Rk_t + n = c_t + k_{t+1} \tag{3.2}$$

$$k_t \geq B_{t,lim} \tag{3.3}$$

The first constraint is the budget constraint with private wealth denoted by k, and the gross interest rate by R. The public pension is denoted by n. We set $n = 0$, since we want the consumption level during retirement to depend on income during working life as in the current U.S. pension system. Endowing all agents with the same pension would result in a counterfactual redistribution of income through pensions, given that agents are ex ante the same (there is no heterogeneity in ability or disutility of working) but differ in earned income. The second constraint is the borrowing constraint. We have set $B_{t,lim} = 0$ as the borrowing limit, since agents do not get a public pension, and hence have a natural borrowing limit of 0.

Writing the optimization problem as a Lagrangian yields:

$$\mathcal{L} = E\left\{\sum_{t=T+1}^{T+D} \beta^{t-1}\left[u\left(c_t\right) + \delta_t\left(Rk_t + n - c_t - k_{t+1}\right) + \gamma_t\left(k_t - B_{t,lim}\right)\right]\right\} \tag{3.4}$$

Combining the First Order Conditions (FOC) with respect to c_t and k_{t+1} results in the Euler equation:

$$u'\left(c_t\right) = \beta R E\left\{u'\left(c_{t+1}\right) + \gamma_{t+1}\right\} \tag{3.5}$$

which implies consumption is constant if $R = 1/\beta$ (as in this paper) and wealth is positive at the beginning of retirement (when the non-negativity constraint on private capital k does not bind). Both these conditions hold in this paper.

3.2.2 Working age: employed individuals

Employed individuals can only choose their consumption level. The Bellman equation that workers optimize in periods $t = [1, ..., T]$ is:

$$\begin{aligned} V(k_t, z_t) = &\max_{c_t, k_{t+1}} u(c_t) - d + \\ &\beta E\{(1-\lambda) V(k_{t+1}, z_{t+1} = 0) + \lambda V(k_{t+1}, z_{t+1} = 1)\} \end{aligned} \tag{3.6}$$

where $V(k_t, z_t)$ is the value function, and k_t is private savings. Symbol z is the status, where $z = 0$ means employed and $z = 1$ means the first period of unemployment. The disutility of working is d and the exogenous probability of loosing a job is λ.

The constraints of the optimization are:

$$Rk_t + (1 - \tau - \mu) w_t \geq c_t + Q(Rk_t - k_{t+1}) + k_{t+1} \tag{3.7}$$

$$k_{t+1} \geq B_{t,lim} \tag{3.8}$$

where τ is the tax rate and μ is the UI contribution rate, both as a percentage of the gross wage. The UI contributions μw_t will be added to a public UI fund. Taxes are used to fund a budget surplus. The adjustment costs to wealth, if the change in capital is negative, are denoted by Q, and defined as:

$$Q = \begin{cases} \varrho (Rk_t - k_{t+1})^2 & \text{if } Rk_t - k_{t+1} < 0 \\ 0 & \text{otherwise} \end{cases} \tag{3.9}$$

These adjustment costs allow us to match the consumption drop when becoming unemployed of the model with empirical data, while using only one type of wealth[1].

[1] An alternative would be to introduce two types of wealth, illiquid and liquid wealth, where illiquid

The costs are quadratic, which makes it more costly to withdraw larger amounts. This reflects that the costs of obtaining liquid funds is likely to be increasing with the amount required. For example, the costs of transforming illiquid assets to liquid funds is likely to be increasing in the amount involved.

The First Order Condition with respect to consumption is:

$$u'(c_t) = \delta_t \tag{3.10}$$

where δ_t is the multiplier of the budget constraint (3.7). The FOC with respect to k_{t+1} is:

$$\beta E\left\{[1-\lambda]\frac{\partial V(k_{t+1}, z_{t+1}=0)}{\partial k_{t+1}} + \lambda\frac{\partial V(k_{t+1}, z_{t+1}=1)}{\partial k_{t+1}}\right\} + \gamma_t = \delta_t\left(1 - \frac{dQ}{dk_{t+1}}\right) \tag{3.11}$$

where δ_t is multiplier on the budget constraint (3.7) and γ_t is multiplier on the non-negativity constraint of private capital (3.8).

3.2.3 Working age: unemployed individuals

Unemployed individuals also have to decide on their consumption and savings, but they have one extra control variable, which is search effort. Search effort determines the chance of finding a job, but results in disutility.

The Bellman equation for the unemployed is:

$$V(k_t, z_t > 0) = \max_{c_t, a_t, k_{t+1}} u(c_t) - a_t + \beta\{p(a_t)V(k_{t+1}, z_{t+1}=0) + [1-p(a_t)]V(k_{t+1}, z_{t+1}>0)\} \tag{3.12}$$

---

wealth would earn a higher return. An extra state variable would increase computation times significantly though, and make optimization of the policies too time-consuming.

where $z_{t+1} > 0$ means $z_{t+1} = \min(z_t + 1, z_{max})$. This optimization is subject to constraints:

$$Rk_t + b_{z,t} \geq c_t + Q(Rk_t - k_{t+1}) + k_{t+1} \tag{3.13}$$

$$k_{t+1} \geq B_{t,lim} \tag{3.14}$$

$$a_t \geq 0 \tag{3.15}$$

where a_t is search effort, $p(a_t)$ is the chance of finding a job given effort a_t (see equation 3.18 for the functional form), $b_{z,t}$ are the benefits for the unemployed ($z > 0$) with status z. With a maximum duration of UI benefits we have $z_{max} = ui_{max}+1$, meaning the maximum duration of unemployment benefits plus one period. This is the period when the agent no longer receives unemployment benefits but gets Social Assistance, which is set equal to (almost) zero throughout this paper. The adjustment costs of private wealth Q are the same as for the employed and described in equation (3.9).

The FOCs for c_t and k_{t+1} can be combined to yield:

$$u'(c_t)\left(1 - \frac{dQ}{dk_{t+1}}\right) = \beta\left\{p(a_t)\frac{\partial V(k_{t+1}, z_{t+1} = 0)}{\partial k_{t+1}} + [1 - p(a_t)]\frac{\partial V(k_{t+1}, z_{t+1} > 0)}{\partial k_{t+1}}\right\} + \gamma_t \tag{3.16}$$

where γ_t is the multiplier of the non-negativity constraint on private wealth (3.14). The FOC for search effort a_t is:

$$1 = \beta p'(a_t)\left\{V(k_{t+1}, z_{t+1} = 0) - V(k_{t+1}, z_{t+1} > 0)\right\} + \varphi_t \tag{3.17}$$

where φ_t is the multiplier on the non-negativity constraint on search effort (3.15). The model is numerically solved backwards so at time t the values for $V(k_{t+1}, z_{t+1})$ are known. Search effort therefore only depends on k_{t+1} and the optimization problem can be reduced to a single variable (k_{t+1}) problem.

Figure 3.1: Wage over the life-cycle

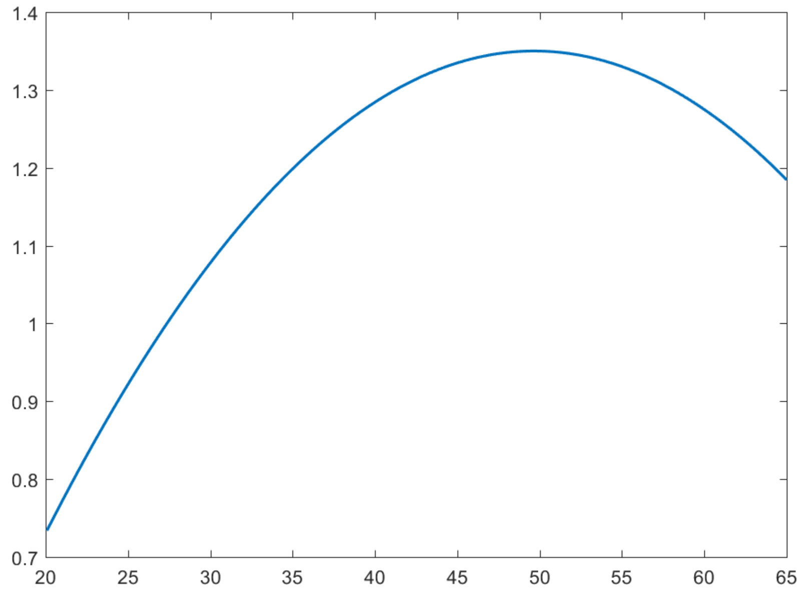

3.3 Calibration

The model is calibrated using U.S. data, where one period is a month. The unemployment benefits in the calibration are fixed at 50% of the net wage as in Wang and Williamson (2002) with a maximum duration of 6 months. If someone is unemployed for longer he gets (almost) zero income[2]. The net average wage over the life-cycle is set at 1.0 and the average gross wage at 1.2. We use productivity profiles of Hanse (1993) as the age-dependent wage, which is plotted in Figure 3.1.

The utility of consumption is of the CRRA-type, $\frac{(Ac+\epsilon)^{1-\alpha}}{1-\alpha}$, with $\alpha = 2$ in the benchmark calibration, $A = 0.25$ and $\epsilon = 0.01$ where the latter two values ensure that the marginal utility at $c_t = 0$ is limited[3]. The interest rate is set at $R = 1/\beta$. Pensions are set at (almost) zero, because agents can save for their pension without much distortion, and we do not want to redistribute income through pensions (see

[2]In practice we set the Pension and Social Assistance to 0.0001 to make it slightly easier to solve the optimization problem numerically, since this allows some difference between the upper and the lower bound in the optimization so we can always use a constraint optimization routine.

[3]A high marginal utility of consumption increases the error in the approximation of the Value Function, especially when the borrowing constraint is binding.

subsection 3.2.1 on Retired individuals).

The model is calibrated to match the following estimates of Kroft and Notowidigdo (2016):

- Unemployment rate with an average of 6.2% in the U.S. over 1968-2007;
- Duration of unemployment with an average of 4.25 months in the U.S. between ages 20 to 62, and a linear increase of 0.032 months per year of age;
- Elasticity of average unemployment duration with respect to the level of UI benefits of 0.63. This estimate of Kroft and Notowidigdo (2016) is close to the middle of the range of estimates reported by Krueger and Meyer (2002)
- Consumption drop (in logs) when people become unemployed of 0.069 based on the PSID data from 1968 to 1997, which only measures food consumption.

Kroft and Notowidigdo (2016) used the Survey of Income and Program Participation (SIPP) data over the period 1985-2000 (except for the estimation of the consumption drop). The sample is restricted to 1) prime age males, 2) job-search reporters, 3) not laid off temporarily, 4) at least three months of work history, and 5) claimed UI benefits. The data is right-censored at 50 weeks of unemployment.

The parameters to be calibrated are the job loss probability λ, the search efficiency in the matching function with a linear age-dependent trend r_t, the disutility of working d, and the parameter for the adjustment costs in private wealth ϱ. The results of the calibration are listed in Table 3.1.

The job loss probability λ is not taken from the data as unemployment spells shorter than one month are not possible in the model and hence should be omitted. The results from the calibration are a job loss probability of 1.56% per month. This is significantly lower than the 5.32% reported by Shimer (2005), but that number includes many very short term unemployment spells, which are ignored in this model except for their influence on the average duration.

One of the key assumptions of the model is the age-dependent job-finding probability, which is assumed to be lower for older ages compared to younger ages, because of

Table 3.1: Parameters

	Value
UI contr. (μ in %)	2.37
Tax (τ in %)	14.30
Net repl. rate (% of net wage)	50.00
Net soc. ass.	0.0001
Net public pension	0.0001
Max. UI dur. (months)	6
Years in working age (T)	45
Years in retirement (D)	15
Population	1000000
Discount factor (β)	0.9966
Risk aversion (α)	2.00
Job loss prob. (λ in %)	1.56
Avg. matching efficiency (r_0)	0.22
Age trend in match. eff. (r_1)	0.0057
Adjustment Costs (ϱ)	0.035
Disutility working (d)	2.991
C.V. disutility [a]	*0.620*

[a] *C.V. is calculated here as* $\frac{d}{u'(\hat{c})}$ *for simplicity, where* $\hat{c}$ *is average consumption*

differences in job-matching efficiency. That older unemployed have more difficulties finding a job is consistent with earlier findings of Chan and Stevens (2001). Hence, the duration of unemployment is, on average, higher among older unemployed than younger unemployed. Since this is an important feature of the model we allow for an age-specific matching probability (for given search effort), and calibrate it as follows.

The matching function is:

$$p\left(a_t\right) = \frac{1}{1-p_0}\left(r_t a_t + \epsilon_p\right)^{\nu} / \left[1 + \left(r_t a_t + \epsilon_p\right)^{\nu}\right] - \frac{p_0}{1-p_0} \tag{3.18}$$

$$p_0 = \frac{\epsilon_p^{\nu}}{1+\epsilon_p^{\nu}} \tag{3.19}$$

which results in $p\left(0\right) = 0$, $p'\left(0\right) \gg 0$ and $p\left(\infty\right) = 1$. The value ν is chosen arbitrary ($\nu = 0.7$), and $\epsilon_p = 0.05$ is added to ensure that $p'\left(0\right) \neq \infty$, which would make it difficult to approximate the matching function for a_t close to zero. The parameter

Figure 3.2: Unemployment duration over life-cycle (model results)

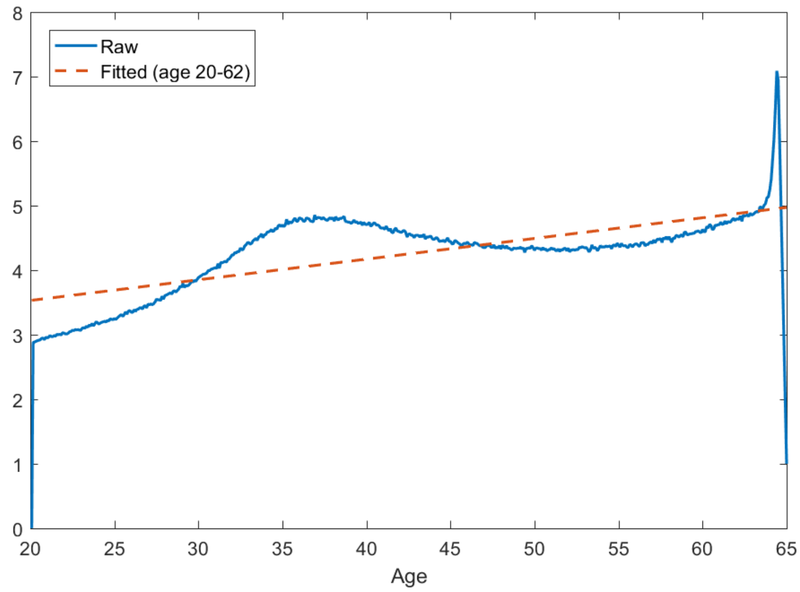

r_t is made linearly dependent on age:

$$r_t = \max\left[0, r_0 + r_1 \left(age_t - \overline{age}\right)\right] \tag{3.20}$$

where $\overline{age}$ is the average age, and r_0 and r_1 are calibrated to match the empirical values of 1) the average unemployment duration, and 2) the linear trend in the average unemployment duration as a function of age.

For this purpose we calculate how the average duration is linearly dependent on age using the same SIPP sample as Kroft and Notowidigdo (2016). Using a simple regression we obtain an average duration of 4.25 months, and an increase of 0.032 months per year of age. We match the average duration and linear trend in duration with age by adjusting r_0 and r_1 in equation (3.20). The average duration per age in the simulation of the model is plotted in Figure 3.2 and is labeled 'Raw'. The line 'Fitted' is the linear trend fitted to these points and corresponds to the calibration targets of the average duration and the linear trend of duration with age.

As reported in Table 3.1 the disutility of working is set at $d \approx 3$. The Compensating Variation of the disutility of working is about 60% of average consumption, which intuitively seems realistic. The parameter of the adjustment costs ϱ in equation (3.9)

Table 3.2: Main results for 6 months of benefits

	Benchmark	Opt. Constant UI	Opt. UI-age
Net repl. rate (t=1)	0.500	0.374	1.000
UI contributions (%)	2.37	1.62	1.91
Unemployment rate (%)	6.20	5.50	5.09
Compensating Variation	1.114	0.947	0.000
as % of NPV UI in Benchm.[a]	18.13	15.41	0.00

[a] *The Net Present Value of UI contributions in Benchmark is 6.15.*

is set to match the consumption drop (in logs) when agents become unemployed.

3.4 Main Results

The main results consist of the welfare effects of 3 variants: the 'Benchmark', which is the current UI system in the U.S., 'Optimal Constant UI' which is the optimized constant replacement rate, and 'Optimal UI-age' which is the optimal system when UI benefits are allowed to depend linearly on age, with an upper bound of the replacement rate at 100%. The welfare effects are measured by the Compensating Variation relative to the variant with the highest welfare (the age-dependent UI system). For all systems the maximum UI duration is fixed at 6 months.

The results are reported in Table 3.2, which shows that the welfare gain of an Optimal Constant UI replacement rate is relatively small compared to the current system ('Benchmark'), with a difference in the Compensating Variation of just under 3% of the Net Present Value (NPV) of all UI contributions in the Benchmark. The results indicate that the replacement rate in the current UI system is too high, since the optimal constant replacement rate is about 12% lower.

The optimal constant level of UI benefits is somewhat lower than typically reported in the literature. For example Davidson and Woodbury (1997) and Wang and Williamson (2002) report optimal replacement rates of 1.3 and 0.56 for a UI duration of 6 months. One reason that the optimal replacement rates are higher in these papers is that agents have no or very little savings, meaning they have no buffers to smooth their consumption, which increases the value of higher benefit levels.

Figure 3.3: Replacement rate over the life-cycle

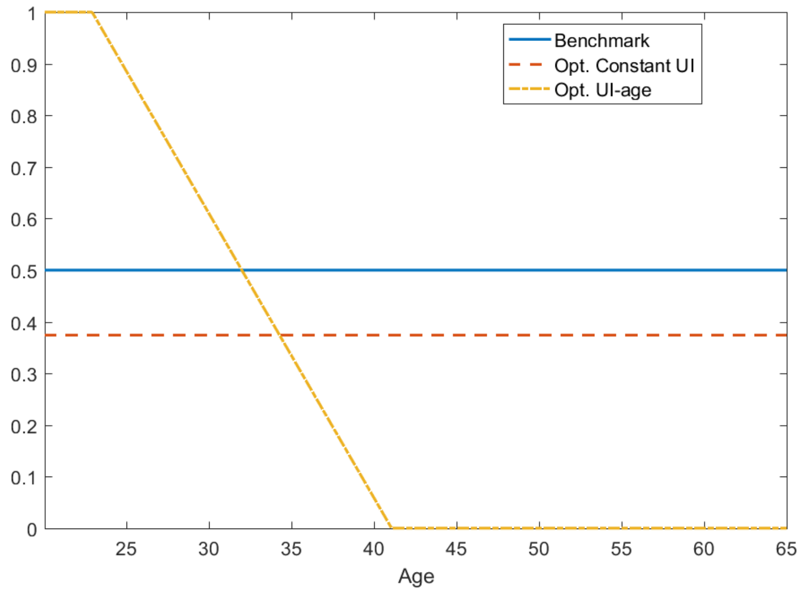

The papers of Davidson and Woodbury (1997) and Wang and Williamson (2002) are also not calibrated as extensive as ours. For example, these authors do not calibrate the elasticity of the duration to the level of UI benefits, or the consumption drop when becoming unemployed. Furthermore, the agents in their models live infinitely, which most likely results in higher replacement rates compared to a life-cycle model. The reason is that the benefits of job search decline as the retirement age is approached, since the time that job seekers will benefit from employment becomes shorter.

The lower constant replacement rate compared to the 'Benchmark' results in a larger consumption drop when becoming unemployed as shown in Figure 3.4. This is directly the result of lower income insurance. The slightly higher welfare in the 'Optimal Constant UI' variant compared to the 'Benchmark' can therefore be attributed completely to a reduction in moral hazard.

This contrasts with the result of the Optimal age-dependent UI benefits ('Opt. UI-age'), where the increase in welfare stems from a better targeted insurance. The resulting welfare difference is relatively large, with an increase in welfare of over 18% of the Net Present Value of all UI contributions in the Benchmark case as reported in Table 3.2. The optimal age-dependent replacement rate is very high

Figure 3.4: Consumption change (logs) when becoming unemployed

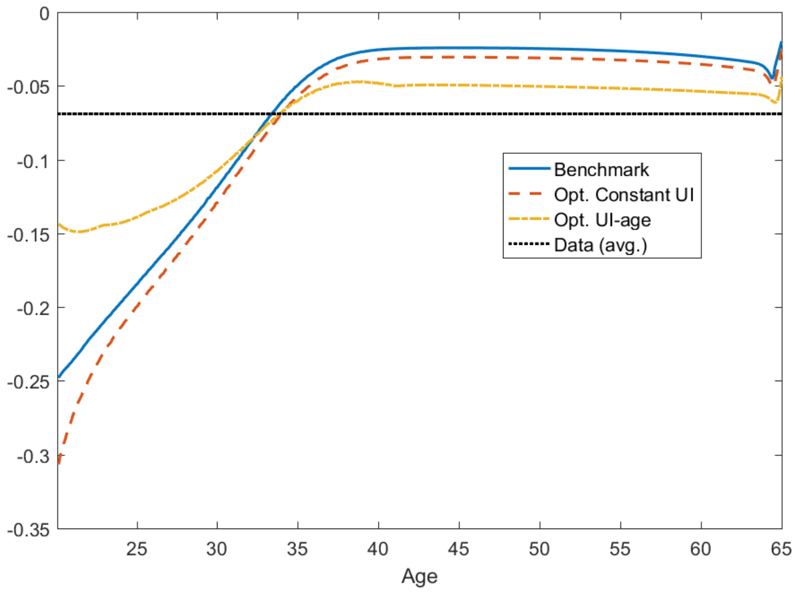

at age 20, reaching the upper bound of 100%. The optimal replacement rate then drops sharply to reach 0 at an age of just over 40 as shown in Figure 3.3.

How these targeted benefits increase welfare is most clear from Figure 3.4, which shows the average consumption drop (in logs) over the life-cycle, where we have also included the average of -0.069 as reported by Kroft and Notowidigdo (2016). The figure shows that the consumption drop in all variants is largest early in the life-cycle. Targeting Unemployment Insurance at this period reduces the consumption drop strongly (the 'Opt. UI-age' variant) at young ages. After the age of around 35 the replacement rate in the age-dependent system is lower than in the other variants, which results in the largest consumption drop after that age. The lower insurance should on the one hand reduce welfare due to higher income risk, but on the other hand reduce moral hazard due to stronger incentives to search for a job. In fact, in the age dependent variant those over 40 get no UI benefits, which completely eliminates moral hazard effects later in life.

3.5 Sensitivity analysis: lower risk aversion

To evaluate to what extend our results in the previous section depend on the choice of the parameters we conduct a sensitivity analysis. Since the risk aversion might

Table 3.3: Parameters, $\alpha = 1$

	Value
UI contr. (μ in %)	2.35
Tax (τ in %)	14.31
Risk aversion (α)	1.00
Job loss prob. (λ in %)	1.56
Avg. matching efficiency (r_0)	1.13
Age trend in match. eff. (r_1)	-0.0279
Adjustment Costs (ϱ)	0.014
Disutility working (d)	0.655
C.V. disutility [a]	*0.596*

[a] *C.V. is calculated here as* $\frac{d}{u'(\hat{c})}$ *for simplicity, where* $\hat{c}$ *is average consumption*

Table 3.4: Main results for 6 months of benefits, $\alpha = 1$

	Benchmark	Opt. Constant UI	Opt. UI-age
Net repl. rate (t=1)	0.500	0.188	0.551
UI contributions (%)	2.35	0.70	0.54
Unemployment rate (%)	6.20	4.81	4.60
Compensating Variation	1.055	0.296	0.000
as % of NPV UI in Benchm. [a]	17.29	4.85	0.00

[a] *The Net Present Value of UI contributions in Benchmark is 6.10.*

play an important role in the welfare gain of our optimized policies we carry out the same analysis as in the previous section, but with risk aversion $\alpha = 1$. The calibration targets the same objectives and results in the parameters listed in Table 3.3, where we have omitted the unaffected parameters.

The main results with the lower risk aversion are listed in Table 3.4. The welfare difference between Optimal Constant UI and Optimal age-dependent UI is reduced from over 15% with a high risk aversion ($\alpha = 2$) to just under 5% with a low risk aversion ($\alpha = 1$). The welfare difference between the Benchmark and the age-dependent system stays roughly the same at 17-18% of the Net Present Value of UI contributions in the Benchmark case.

The main difference between Table 3.2 and Table 3.4 is the much lower level of insurance in the two optimized variants in Table 3.4, which is also reflected in the differences between Figure 3.3 and Figure 3.5. This despite the calibration of the

Figure 3.5: Replacement rate over the life-cycle, $\alpha = 1$

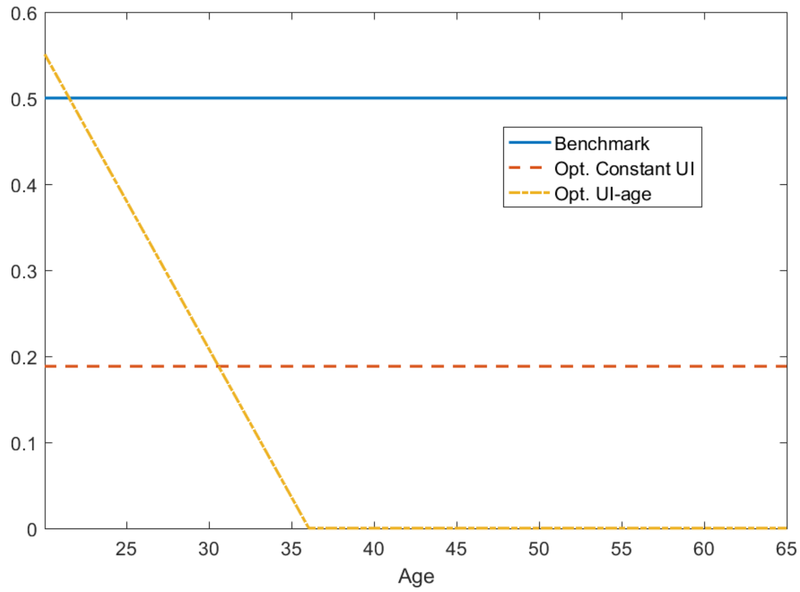

model on both the consumption drop when becoming unemployed and the elasticity of unemployment duration with respect to the level of UI benefits. Apparently, fitting the adjustment costs of wealth ϱ and the disutility of working d to these two objectives does not make the model insensitive to the choice of the risk aversion parameter.

Interesting is that the 'Optimal Constant UI' variant results in a welfare loss of just under 5% compared to the age-dependent variant. This despite a constant replacement rate of less than 20% resulting in large consumption drop when becoming unemployed as shown in Figure 3.6. Despite a much larger consumption drop early in life in the 'Optimal Constant UI' system the welfare is still higher than in the 'Benchmark'. This indicates that moral hazard in the Benchmark system is quite strong as lower insurance and less moral hazard ('Optimal Constant UI') results in higher welfare than high insurance with high moral hazard ('Benchmark').

In the Optimal age-dependent UI system the level of UI benefits drops fast, reaching a level of zero just over 35 years. With the low risk aversion the level of insurance at the beginning of the life-cycle is much lower (replacement rate at $t = 1$ of 55%) compared to $\alpha = 2$ (replacement rate at $t = 1$ of 100%). The welfare benefits of age-dependent benefits comes most of all from reducing moral hazard, since insurance

Figure 3.6: Consumption change (logs) when becoming unemployed, $\alpha = 1$

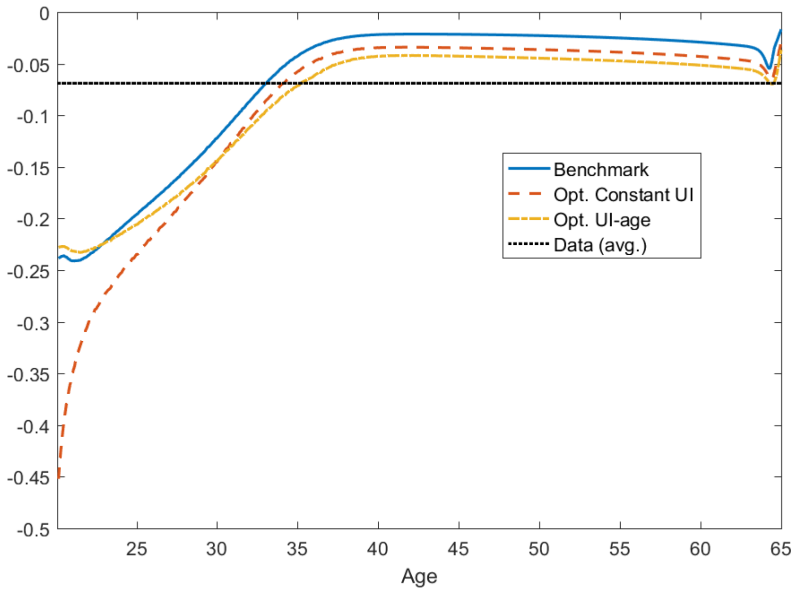

levels are higher than in the 'Benchmark' only in the first few years of the life-cycle.

3.6 Conclusion

We have analyzed the optimal unemployment insurance benefits over the life-cycle in an extensively calibrated model. Our results indicate that the welfare benefits of targeting unemployment insurance at young people can be relatively large, with welfare gains roughly between 5 and 15% of the Net Present Value of UI contributions in the current system. These results are mainly driven by the lack of self-insurance for the young as they have accumulated few assets. The young have less assets, because their income is low compared to the average earnings over the life-cycle. This makes them credit constrained as they could increase their welfare by borrowing against future, higher income, and so raise their consumption when young. In addition, they have had less time to accumulate assets for precautionary motives.

Regarding the insurance of older people our model results indicate that moral hazard effects are strong later in life, which reduces optimal replacement rates for the old. In fact, when using replacement rates linear in age, the optimal replacement rate of the old is zero. This implies that the welfare gains of a higher replacement rate early in the life-cycle are bigger than the welfare gains of raising the replacement rate above zero for the old.

3.A Bibliography

Algan, Y., A. Chéron, J.-O. Hairault, and F. Langot (2003). Wealth effect on labor market transitions. *Review of Economic Dynamics 6*(1), 156–178.

Bloemen, H. G. and E. G. Stancanelli (2001). Individual wealth, reservation wages, and transitions into employment. *Journal of Labor Economics 19*(2), 400–439.

Chamberlain, G. and C. A. Wilson (2000). Optimal intertemporal consumption under uncertainty. *Review of Economic dynamics 3*(3), 365–395.

Chan, S. and A. H. Stevens (2001). Job loss and employment patterns of older workers. *Journal of Labor Economics 19*(2), 484–521.

Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of political Economy 116*(2), 173–234.

Costain, J. S. (1997). Unemployment insurance with endogenous search intensity and precautionary saving.

Davidson, C. and S. A. Woodbury (1997). Optimal unemployment insurance. *Journal of Public Economics 64*(3), 359–387.

Devine, T. J. and N. M. Kiefer (1991). *Empirical labor economics: the search approach.* Oxford University Press on Demand.

Hanse, G. D. (1993). The cyclical and secular behaviour of the labour input: Comparing efficiency units and hours worked. *Journal of Applied Econometrics 8*(1), 71–80.

Heer, B. (2003). Employment and welfare effects of a two-tier unemployment compensation system. *International Tax and Public Finance 10*(2), 147–168.

Kroft, K. and M. J. Notowidigdo (2016). Should unemployment insurance vary with the unemployment rate? theory and evidence. *The Review of Economic Studies 83*(3), 1092–1124.

Krueger, A. B. and B. D. Meyer (2002). Labor supply effects of social insurance. *Handbook of public economics 4*, 2327–2392.

Lentz, R. (2009). Optimal unemployment insurance in an estimated job search model with savings. *Review of Economic Dynamics 12*(1), 37–57.

Lentz, R. and T. Tranaes (2005). Job search and savings: Wealth effects and duration dependence. *Journal of Labor Economics 23*(3), 467–489.

Rupert, P. and G. Zanella (2015). Revisiting wage, earnings, and hours profiles. *Journal of Monetary Economics 72*, 114–130.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American economic review 95*(1), 25–49.

Wang, C. and S. D. Williamson (2002). Moral hazard, optimal unemployment insurance, and experience rating. *Journal of monetary Economics 49*(7), 1337–1371.

Chapter 4

Adaptive Beliefs in a RBC Model

4.1 Introduction

The macroeconomics paradigm of Rational Expectations has dictated the field over the last four decades, but more recently attention is shifting towards models where agents are boundedly rational in some way (Hommes, 2011). Especially after the financial crisis there are increasing calls to replace the rigidity of Rational Expectations (Blanchard, 2018). A wide range of different types of expectations that deviate from Rational Expectations have been developed. Our main interest is the role that 'animal spirits' can play in business cycle fluctuations, where 'animal spirits' are defined as waves of optimism and pessimism about future economic prosperity (De Grauwe, 2011).

Other papers that have looked into 'animal spirits' as a source of business cycle fluctuations are De Grauwe (2011), Massaro (2013) and Hirshleifer et al. (2015). The first two implemented extrapolative biases in a New Keynesian framework without capital accumulation. In their framework extrapolative biases can result in instability. The paper by Hirshleifer et al. (2015) features an economy that is more close to our economy although their labor supply is fixed, and their main interest lies with the effect of 'animal spirits' on asset prices.

There are several other approaches in the literature to investigate expectations as a source of business cycle fluctuations. Two examples are exogenous shocks to expectations (see for example Milani, 2011) or adaptive learning (see Evans and Honkapohja, 2001, for an extensive overview). Within the latter literature there are several variants such as signal extraction problems (Edge et al., 2007), misspecified models (Lorenzoni, 2009; Branch and McGough, 2011), and incomplete information models (Eusepi and Preston, 2011).

In our model 'animal spirits' are confined to exogenous productivity without any feedback loop that endogenously reinforces positive or negative shocks. Agents perfectly observe all variables including the productivity level, but use heuristics to predict future levels of productivity. To be more precise, agents extrapolate recent changes in the productivity level. The agents expect future productivity levels to be a weighted average of two extreme forecasting rules, an optimistic and a pessimistic forecasting rule. The weights of each rule depend on the recent performance of each rule as in an Adaptive Beliefs System (Brock and Hommes, 1997, 1998). When there have recently been more positive (negative) shocks, then the agents expect a positive bias (negative) in future productivity shocks. As in Massaro (2013) and De Grauwe (2011) we aggregate the expectations of heterogeneous agents into a representative agent framework[1].

The intuition behind such an extrapolative bias is best explained with the uncertainty about productivity that we have seen for example during the dotcom bubble. At the height of the dotcom bubble there was the belief that new technologies were able to support a higher growth rate of productivity than before (Hirshleifer et al., 2015). Apparently there was also data that supported this claim as even professionals and policy makers seemed to expect higher growth rates of productivity. In effect they extrapolated the higher growth rate of the late nineties into the new millennium (Edge et al., 2007).

The extrapolative bias in our model does something similar: recent changes in productivity levels are projected into the future. There is empirical evidence that in real life people rely on heuristics to form expectations about the future (see Hommes,

[1]Although Massaro (2013) and De Grauwe (2011) start from a heterogeneous agent framework, they both take the average expectation at some point.

2011, for references). These heuristics include trend extrapolation and mean reversion as forecasting strategies. There is also evidence that the fraction of agents using a particular forecasting strategy changes over time (Hommes, 2011). This evidence justifies the use of adaptive beliefs as a micro-foundation for macroeconomic models.

One critique of models that deviate from Rational Expectations is that it allows the modeler too many degrees of freedom, or the 'wilderness' of bounded rationality as Sims (1980) called it. A common approach to address this issue is the calibration of the forecasting error of the model to the forecasting errors of professionals. This is for example the approach of Eusepi and Preston (2011). However, in general this modeling assumption results in correlations between output and consumption, and output and hours worked that are too high, as in Real Business Cycle (RBC) models with Rational Expectations. In contrast we calibrate the changes in beliefs in our model to a measure of consumer confidence. In particular, the standard deviation and autocorrelation of the weight of the optimistic forecasting rule is set such that it matches the standard deviation and autocorrelation of consumer confidence.

Another issue that arises in models that deviate from Rational Expectations is the length of the forecasting horizon of agents in the model. Several models use agents that only make one-period-ahead forecasts, which makes their decisions suboptimal given their maintained beliefs (Eusepi and Preston, 2011). In this paper agents assume they will not update their expectations in the future. This means agents make infinite horizon forecasts under the assumption of maintained beliefs as in Eusepi and Preston (2011), which implies that agents in our model expect the steady state to change.

The model that we use is a RBC model, which is calibrated on the second moments of the US economy. We use the preferences of Jaimovich and Rebelo (2009), which allow us to adjust the strength of the wealth effect in response to expected income shocks. This feature is crucial in our model, because in a standard RBC model higher expected productivity generates a wealth effect that results in higher consumption and more leisure, thus reducing labor supply therewith resulting in a recession (Beaudry and Portier, 2007; Jaimovich and Rebelo, 2009; Den Haan and

Kaltenbrunner, 2009). With adaptive expectations a productivity shocks can result in a much higher expected productivity, which will have a large wealth effect and reduce labor supply strongly. With the preferences from Jaimovich and Rebelo (2009) we can manage the strength of this wealth effect.

The paper is organized as follows. The next section describes the model, including the adaptive beliefs and the solution method. Section 3 describes the Calibration. Section 4 gives the Results, and Section 5 is a sensitivity analysis with respect to the standard deviation of the 'animal spirits'. Finally the paper is concluded in Section 6.

4.2 Model

The model is a standard RBC model with the same preferences as in Jaimovich and Rebelo (2009). These preferences allow us to control the wealth effect. The preferences include the form used by Greenwood et al. (1988) (when $\phi = 0$) and the form chosen by King et al. (1988) (when $\phi = 1$) as special cases. The parameter ϕ thus controls the strength of the short run wealth-effect on labor supply.

The main difference with standard RBC models is that we introduce adaptive beliefs with respect to the exogenous productivity process. The agent can observe all variables including the productivity level, but will forecast future productivity levels with a bias. The bias results from putting more weight on recent realizations of the productivity level when forming expectations about the future level of productivity. The optimal decisions are made under the assumption that the agent will maintain his belief as in Eusepi and Preston (2011). Our framework has homogeneous beliefs, and does not have externalities, which allows us to use a representative agent framework.

4.2.1 Economy

The representative agent maximizes expected life-time utility:

$$U = \widetilde{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{\left(C_t - \psi N_t^{\theta} X_t\right)^{1-\nu} - 1}{1-\nu} \tag{4.1}$$

where C is consumption, N is hours worked, and X is a habit variable. The risk aversion is ν, the discount factor is β, and θ governs the labor supply elasticity. The expectation operator $\widetilde{E}_t$ indicates a (possibly) bounded rational expectation.

The habit variable X_t is given by:

$$X_t = C_t^{\phi} X_{t-1}^{1-\phi} \tag{4.2}$$

where ϕ governs the strength of the short run wealth-effect on labor supply. The habit formation makes preferences non-time-separable in consumption and hours.

Production is given by the Cobb-Douglas production function $Y_t = Z_t K_t^{1-\alpha} N_t^{\alpha}$. In this equation Y_t is output, K_t is the capital stock at the beginning of the period, and Z_t is the exogenous Total Factor Productivity (TFP). Output is used for consumption or investment: $Y_t = C_t + I_t$ and capital evolves according to $K_{t+1} = (1-\delta) K_t + I_t$ where δ is the depreciation rate. Combining these three equations yields the budget constraint:

$$K_{t+1} + C_t = Z_t K_t^{1-\alpha} N_t^{\alpha} + (1-\delta) K_t \tag{4.3}$$

The First Order Conditions with respect to C_t, X_t, N_t, and K_{t+1} are:

$$\lambda_t = \left(C_t - \psi N_t^{\theta} X_t\right)^{-\nu} - \mu_t \phi C_t^{\phi-1} X_{t-1}^{1-\phi} \tag{4.4}$$

$$\beta \widetilde{E}_t \left[\mu_{t+1} \left(1-\phi\right) C_{t+1}^{\phi} X_t^{-\phi}\right] = -\left(C_t - \psi N_t^{\theta} X_t\right)^{-\nu} \psi N_t^{\theta} + \mu \tag{4.5}$$

$$\lambda_t \alpha Z_t K_t^{1-\alpha} N_t^{\alpha-1} = \left(C_t - \psi N_t^{\theta} X_t\right)^{-\nu} \theta \psi N_t^{\theta-1} X_t \tag{4.6}$$

$$\lambda_t = \beta \widetilde{E}_t \lambda_{t+1} \left\{\left(1-\alpha\right) Z_{t+1} K_{t+1}^{-\alpha} N_{t+1}^{\alpha} + \left[1-\delta\right]\right\} \tag{4.7}$$

where μ_t and λ_t are the Lagrangian multipliers on the constraints (4.2) and (4.3) respectively. The description of the agent's beliefs follows in the next subsection.

The Actual Law of Motion for productivity Z_t is the following exogenous process:

$$z_t = \rho_z z_{t-1} + \epsilon_t \tag{4.8}$$

where $z_t = \log\left(Z_t\right)$. The autocorrelation coefficient is ρ_z, and shock ϵ has standard deviation σ_z and a zero mean.

4.2.2 Adaptive Beliefs

In our model the representative agent observes current period's variables, including the capital stock in the next period, K_{t+1}, which is determined in period t. The agent has Bounded Rational beliefs with respect to the law of motion of the exogenous productivity process described by equation (4.8), but is rational otherwise. The agent's Perceived Law of Motion for the productivity level is biased:

$$z_{t+1} = \rho_z z_t + \hat{b}_t + \epsilon_{t+1} \tag{4.9}$$

where the bias $\hat{b}_t$ depends on the current and past realizations of the productivity shocks, following the Adaptive Beliefs System (Brock and Hommes, 1997) as described hereunder.

The expected productivity level of the representative agent is a weighted average of an optimistic and a pessimistic forecasting rule. The expected productivity level for the pessimistic and optimistic rule are:

$$\widetilde{E}_t^{pes} z_{t+1} = \rho_z z_t - g \tag{4.10}$$
$$\widetilde{E}_t^{opt} z_{t+1} = \rho_z z_t + g \tag{4.11}$$

where $z_t = \log(Z_t)$. We set $g = \sigma_z$ meaning the bias in each forecasting rule is one standard deviation, and therefore also limits the maximum bias to one standard deviation. Note that the bias in each forecasting rule lasts indefinitely, so the pessimistic expectation for the steady state productivity level is: $\widetilde{E}^{pes}\bar{z} = \frac{-g}{1-\rho_z}$.

The weights of each forecasting rule are determined in the same way as Brock and Hommes (1998). The current period fitness measures are:

$$\tau_t^{pes} = \left(\widetilde{E}_{t-1}^{pes} z_t - z_t\right)^2 = (-g - \epsilon_t)^2 \tag{4.12}$$
$$\tau_t^{opt} = \left(\widetilde{E}_{t-1}^{opt} z_t - z_t\right)^2 = (g - \epsilon_t)^2 \tag{4.13}$$

where ϵ_t is the productivity shock in period t. The total fitness measure for each rule $i = [pes, opt]$ is:

$$M_t^i = \tau_{i,t} + \chi M_{t-1}^i \tag{4.14}$$

where χ is a measure of the memory strength. The weight of each forecasting rule is:

$$\omega_t^i = \frac{e^{-\gamma M_t^i}}{e^{-\gamma M_t^{pes}} + e^{-\gamma M_t^{opt}}} \tag{4.15}$$

where γ is a measure of the 'intensity of choice' (De Grauwe, 2011) as it determines how responsive the weights are to changes in the fitness measure. If $\gamma = 0$ each rule has a weight of 0.5, and if $\gamma = \infty$ the weight of each rule switches between 0 and 1.

The bias in the Perceived Law of Motion for Total factor Productivity can be written

Figure 4.1: Bias as function of shock ϵ_t

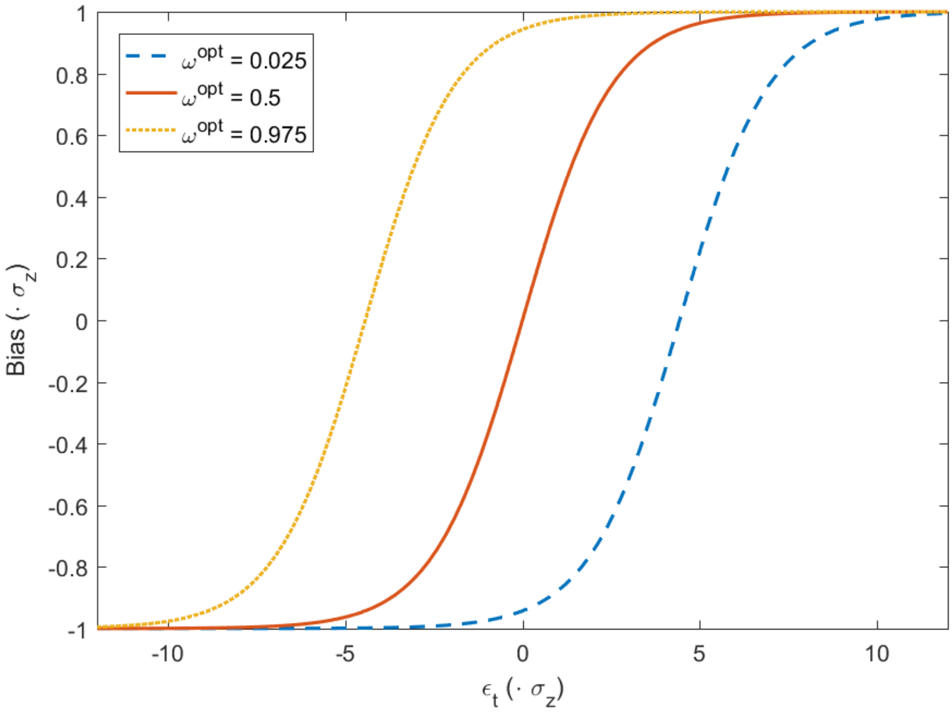

as:

$$\hat{b}_t = g\left(\omega_t^{opt} - \omega_t^{pes}\right) \tag{4.16}$$

This bias in the Perceived Law of Motion as described by equation (4.9) is a function of the shock in this period, ϵ_t, and the fitness measures of each forecasting rule in the previous period, M_{t-1}^{pes} and M_{t-1}^{opt}, so we can write: $\hat{b}_t = B\left(\epsilon_t, M_{t-1}^{pes}, M_{t-1}^{opt}\right)$.

A graph of the bias as a function of the shock is shown in Figure 4.1, where we have used the benchmark calibration (see Section 4.3) with $\chi = 0.9567$ and $\gamma = 2015$. In the figure it is clear that the marginal effect on expectations depends strongly on the size of the shock. When the sentiment is neutral ($\omega^{opt} = 0.5$) a marginal increase in the size of the shock around the zero mean affects the bias strongly. This contrasts with a small change in the bias as result of a marginal change in the shock if the level of the shock is already large in absolute value.

4.2.3 Optimal decision & numerical solution

The agent solves the First Order Conditions of equations (4.4) to (4.7), under the constraints (4.2) and (4.3), given his beliefs about the exogenous productivity process. These beliefs are described in equation (4.9), where the agent assumes that he will maintain his beliefs about the bias $\hat{b}_t$, similar to the maintained beliefs in Eusepi and Preston (2011). This means that the Bounded Rational agent's optimal decision rule coincides with the optimal decision rule that a Rational Expectations agent would have if the Actual Law of Motion for productivity was described by (4.9). Note that this implies that the agent expects the steady state of productivity, and therewith the steady state of the other variables, to change.

The decision rules of the agent are conditional on the bias, $\hat{b}_t = B\left(\epsilon_t, M_{t-1}^{pes}, M_{t-1}^{opt}\right)$. Given an agent's bias his decision rules are a function of the remaining state variables, $D_t = [K_t, Z_t, X_{t-1}]$. The agent only has to form expectations about the future values of these three state variables. The decision rules for consumption and hours[2] conditional on the bias can then be written as:

$$\hat{C}_t = C\left(K_t, Z_t, X_{t-1}|\hat{b}_t\right) \tag{4.17}$$

$$\hat{N}_t = N\left(K_t, Z_t, X_{t-1}|\hat{b}_t\right) \tag{4.18}$$

Since we use projection methods it would take a considerable amount of time to calculate the policy function in each period, as the bias $(\hat{b}_t)$ changes[3]. Instead we calculate the policy function for a limited number of gridpoints over the range of possible biases. We use gridpoints $j = 1, ..., J$, with $-g \leq b_j \leq g$. Interpolation between these gridpoints gives us the policy function for any possible level of the bias.

---

[2]We use projection methods, and have to solve for at least two choice variables, since there are two Euler equations, equation (4.5) and equation (4.7). We choose to solve for consumption and hours worked such that we can analytically solve for the other variables.

[3]De Grauwe (2011) for example recalculates the policy function in every period, but this is only feasible when the computation time of the decision rule is small.

We use the Galerkin projection method with a second order Chebyshev polynomial to obtain the policy functions[4]. During simulations we draw the shocks ϵ_t, and calculate the bias using equations (4.12) to (4.16). Given the bias and the value of the other state variables we can calculate consumption and hours work from our policy functions, which are conditional on the bias as described in equations (4.17) and (4.18). Together with (4.3) and (4.2) we can calculate all endogenous state variables in the next period.

4.3 Calibration

We calibrate a Rational Expectations (RE) and Bounded Rational (BR) version of the model. The calibration consists of choosing the intertemporal rate of substitution ν, the strength of the wealth effect ϕ, and the labor supply elasticity θ. The output elasticity of labor α and the discount factor β are set to the same values as in Jaimovich and Rebelo (2009, henceforth, JR). The steady state depreciation rate is set to the standard value of 0.025 per quarter (compared to 0.0125 in JR).

Ideally the persistence parameter ρ_z of productivity Z_t should be set to a value such that we could match the empirical value of the autocorrelation in output. However, the autocorrelation of output in our model is too low even for values of ρ_z close to 1. For this reason we fix the parameter ρ_z at a relatively standard value for quarterly models, namely $\rho_z = 0.95$. Since the RE and BR model will differ in the amplification of productivity shocks we adjust the parameter σ_z such that the standard deviation of output is 1.7%, which is the empirical value in a quarterly model after detrending.

The remaining three parameters (ν, ϕ, θ) are then set at values that minimize the sum of the normalized squared errors of three second moments. We restricted the long run labor supply elasticity to be below 10, meaning the parameter restriction $\theta \geq 1.1$. The three targeted second moments are the standard deviation of con-

[4]We have only used point estimates for productivity in the next period, $\widetilde{E}_t z_{t+1} = \rho z_t + \hat{b}_t$, instead of Gauss-Hermite nodes that would take full account of the distribution of shocks ϵ_{t+1}.

Table 4.1: Parameters

	RE	BR
Labor supply elasticity (θ)	1.1002[a]	1.1692
Risk aversion (ν)	0.6932	0.7352
Strength of wealth effect (ϕ)	0.0119	0.0097
Standard deviation shocks (σ_z)	0.0048	0.0096
Autocorrelation shocks (ρ_z)	0.9500	
Labor share in output (α)	0.6400	
Discount factor (β)	0.9850	
Depreciation rate (δ)	0.0250	
Scalar labour supply (ψ)	1.0000	

[a] *This is the lower bound.*

sumption relative to the standard deviation of output σ_c/σ_y, the relative standard deviation of labor supply σ_n/σ_y, and the relative standard deviation of investment σ_i/σ_y. We also report four contemporaneous correlations, namely the correlation between output and labor supply, output and consumption, output and Total Factor Productivity, and consumption and labor supply. The resulting parameter values of the calibration are given in Table 4.1.

Calibration of Bounded Rationality The parameters γ and χ of the adaptive beliefs system are set such that the volatility in the weights of each forecasting rule ω^i as in equation (4.15) roughly matches the standard deviation and autocorrelation coefficient in Consumer Confidence[5], which are 15% and 70%, respectively. The parameters γ and χ will only depend the parameters ρ_z and σ_z of the stochastic productivity process. The result of the calibration of the adaptive beliefs are $\chi = 0.9567$ and $\gamma = 2015$. This fixes the two additional parameters of the Bounded Rational model, so we have no extra degree of freedom in the Bounded Rational variant of the model.

[5]Mnemonic USCNFCONQ in Datastream, data from 1970-2006.

Table 4.2: Business Cycle statistics

	Data	RE	BR
St.d. Y (σ_y)	0.0170	0.0170	0.0170
Rel. st.d. C (σ_c/σ_y)	0.7500	0.6207	0.7423
Rel. st.d. N (σ_n/σ_y)	1.1200	0.9893	1.1268
Rel. st.d. I (σ_i/σ_y)	2.8500	2.3845	2.8695
Correlation Y and N	0.8600	0.9707	0.7359
Correlation Y and C	0.6600	0.9686	0.9353
Correlation Y and Z	0.7800	0.7977	0.7626
Correlation C and N	0.7300	0.8939	0.7572

Only the first four rows are targeted in the calibration.

4.4 Results

The results are obtained using Galerkin projection with second order Chebyshev polynomials as an approximation of the decision rules for consumption and labor supply, as in equation (4.17) and equation (4.18) respectively. All series are obtained after detrending with the Hodrick-Prescott filter with the smoothing parameter set to $\lambda = 1600$ as is standard for a quarterly model. We use 100 simulations with 1500 quarters each of which we omit the first 500 observations.

The results in terms of Business Cycle statistics are given in Table 4.2. This clearly shows that the Bounded Rational model almost perfectly matches the targeted moments, although the model has the same degrees of freedom as the Rational Expectations model. In addition, the Bounded Rational model matches the correlation between consumption and hours worked much better. Also interesting is that the Bounded Rational model results in a much lower correlation between output and hours worked, which is close to one in most RBC models with (close to) Rational Expectations (see for example Eusepi and Preston, 2011, Table 2 and Table 3).

To explain the very good fit of the Bounded Rational model to the data we first consider the Impulse Response Functions (IRFs) to a shock in Total Factor Productivity of one standard deviation. These IRFs are are shown in Figure 4.2. The figure shows that we get a negative response of hours worked, output and consumption to a TFP shock starting at the steady state. The IRFs seem to imply a negative cor-

Figure 4.2: IRFs to productivity shock, starting at steady state

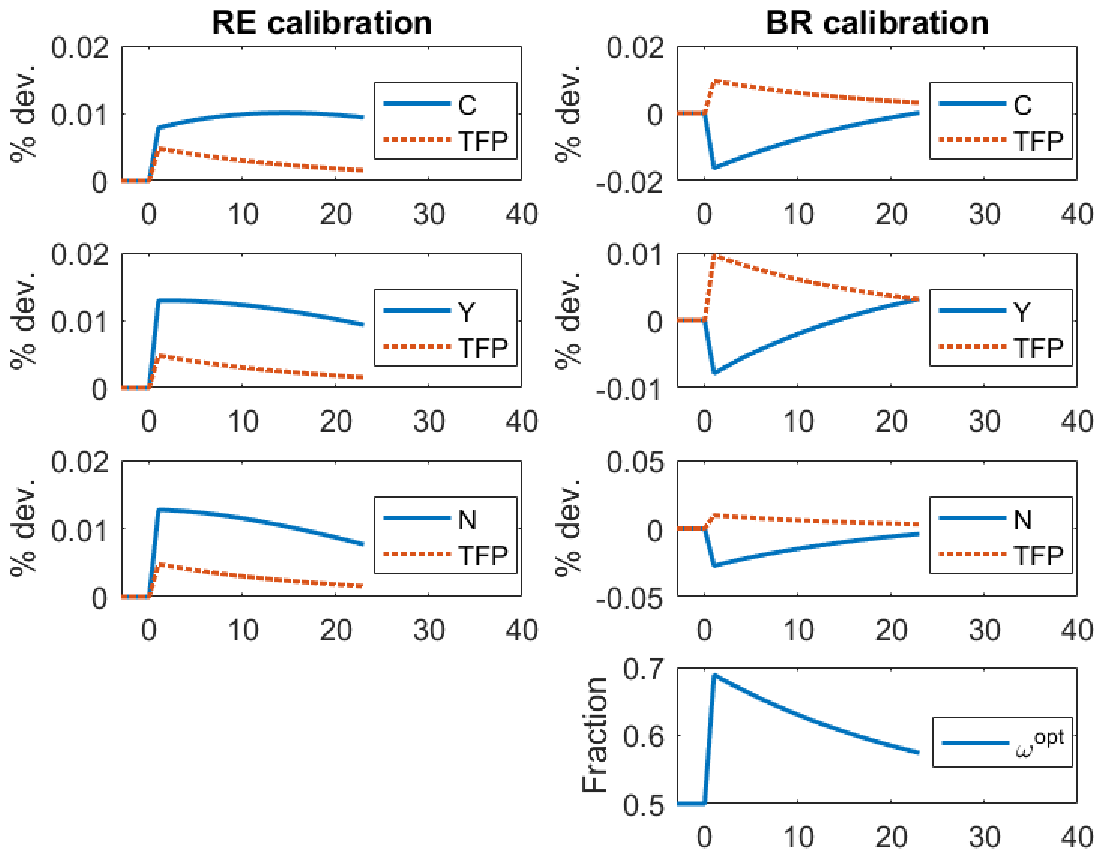

relation between Total Factor Productivity and output, which is at odds with the positive correlation as reported in Table 4.2. The explanation for both the positive correlation between TFP and output and the better fit of the data is combination of three factors.

The first factor that explains the good fit of the model is the strength of the wealth effect of expected future productivity levels. Figure 4.4 shows that the decision rule for hours worked is not monotonic in the TFP shock ϵ_t. For any of the Bounded Rational decision rules there is a section where an increase in the TFP shock reduces hours worked, because the wealth effect of higher expected productivity dominates the substitution effect. In other words at those points the effect of higher expected productivity, which reduces labor supply as agents expect to get richer in the future, is stronger than the increase in hours worked due to the immediate increase in the wage. For a neutral sentiment (weight of the optimistic rule $\omega^{opt} = 0.5$) the wealth effect dominates for shocks around the zero mean.

Figure 4.3: Histogram of weight of optimistic forecasting rule, ω^{opt}

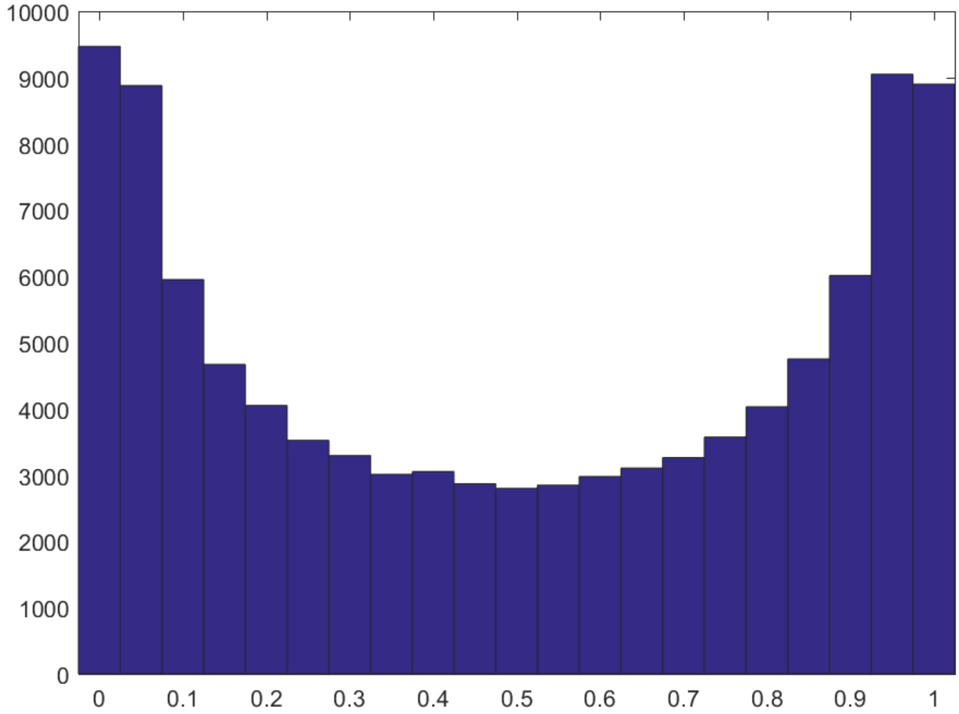

This brings us to the second factor that contributes to the near perfect fit, which is the high standard deviation in the sentiment, corresponding to the high standard deviation in consumer confidence. The high standard deviation in sentiment means that the agent is either very pessimistic or very optimistic most of the time as shown in Figure 4.3. As a consequence the marginal effect of a TFP shock on expected productivity is relatively weak for average shock sizes as shown in Figure 4.1. That figure shows that with a very pessimistic sentiment ($\omega^{opt} = 0.025$) the change in the bias is relatively small for any TFP shock smaller than roughly plus two standard deviations, meaning about 97.5% of all shocks. The same is true for a very optimistic sentiment: in that case the marginal effect of a TFP shock on expectations is small for any shock larger than minus two standard deviations. The marginal effect of TFP shocks on expectations will thus be small most of the time in our Bounded Rational economy. As a result agents are in the upward sloping part of the decision rules for shocks of average size as demonstrated in Figure 4.4. This means the substitution effect will dominate on average, which adds to the positive correlation between TFP and output as reported in Table 4.2.

Figure 4.4: Decision rule for hours worked as function of TFP shock

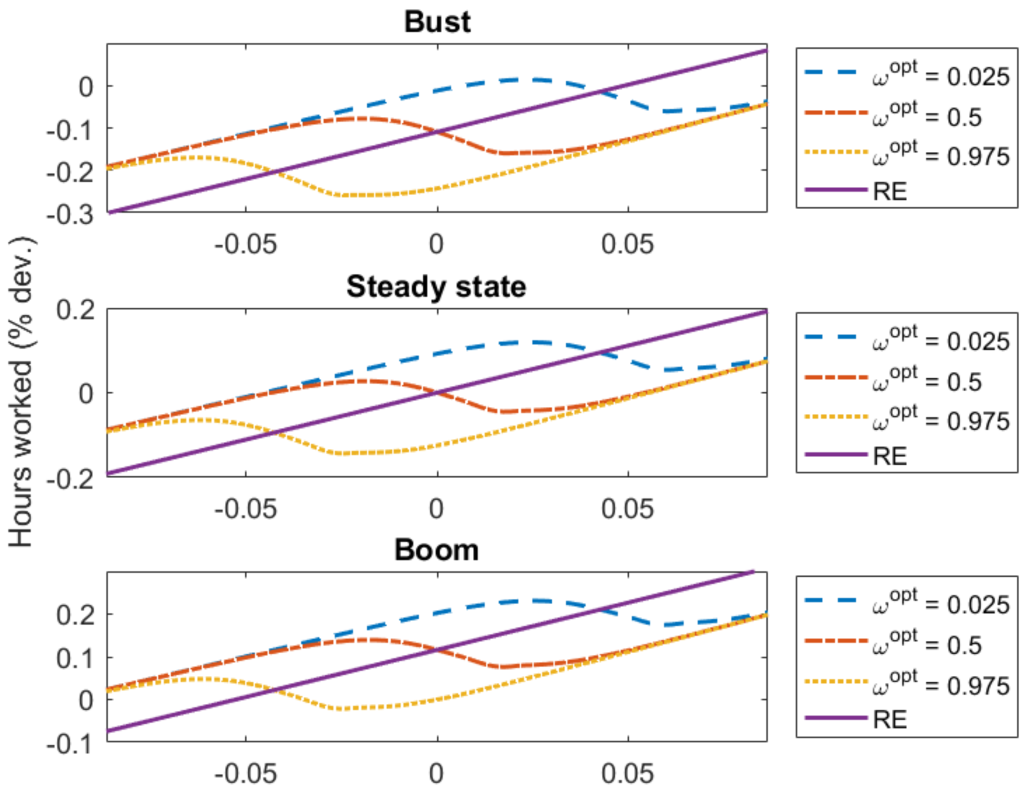

A third factor that contributes to the positive correlation is the substitution effect of Total Factor Productivity and capital. In pessimistic (optimistic) states the TFP and capital levels are below (above) the average, which shifts the policy function for hours worked down (up) as is shown in the upper (lower) panel of Figure 4.4. In the upper panel all policy functions are shifted down as a lower capital stock and lower productivity level[6] reduce the wage and therewith the labor supply[7]. This substitution effect shifts the labor supply function in the same direction as Total Factor Productivity and the capital stock, and adds to the positive average correlation between Total Factor Productivity and output. Together these factors result in a positive correlation between TFP and output and the good fit of the

[6]The starting values of the state variables in 'Bust' and 'Boom' are the averages (of the state variables) at the points that are closest to $\omega^{opt} = 0.025$ for the 'Bust' and $\omega^{opt} = 0.975$ for the 'Boom'. The state variables K, Z and X will all be below (above) their steady state levels in a 'Bust' ('Boom').

[7] The effect of the level of TFP and capital dominate the effect of habit variable X that works in the opposite direction.

Figure 4.5: Simulation Rational Expectations

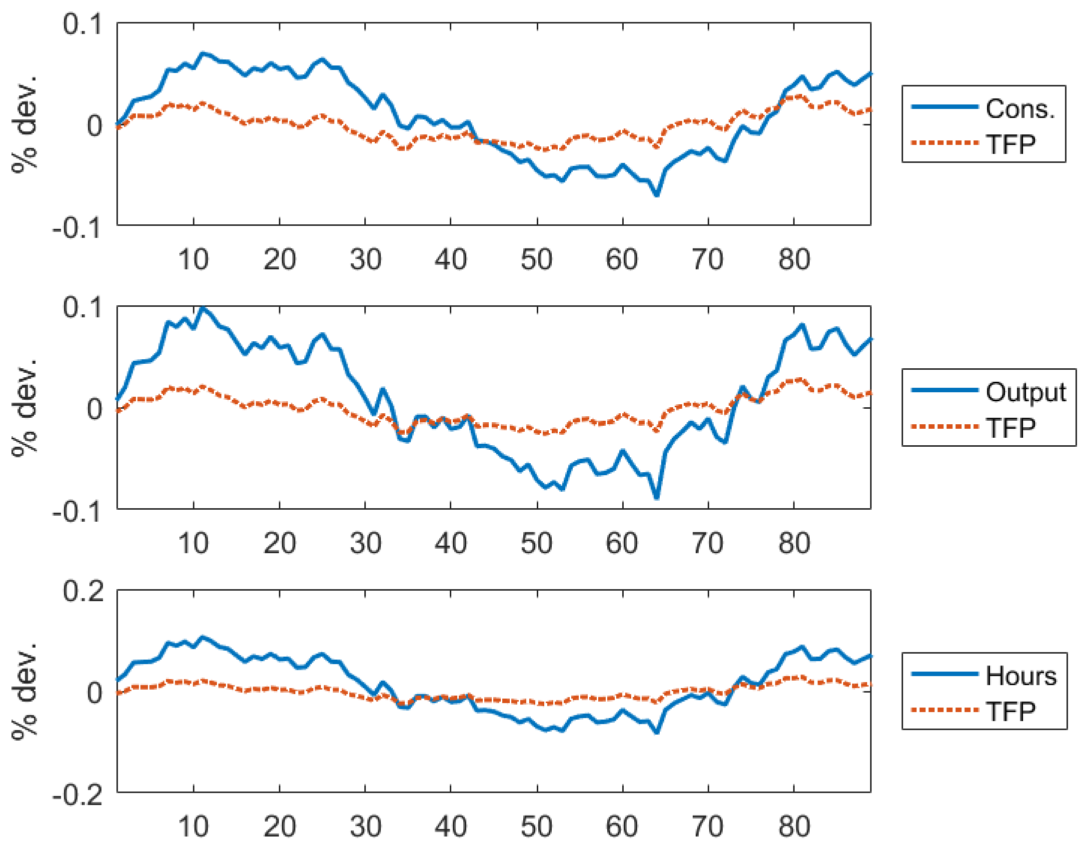

second moments.

To conclude this section we look at a simulation of the Rational Expectations model and the Bounded Rational model which are shown in Figure 4.5 and Figure 4.6, respectively. The first difference between the two models is that the amplification of TFP shocks is much stronger in the Rational Expectations model, which could also be concluded from the calibration, since the standard deviation of the productivity shocks of the RE model is only half the size of the shocks in the BR model. This difference in amplification is clearly visible in the second panel of Figure 4.5 and 4.6, where output tracks TFP quite closely in the Bounded Rational model, but output is much more volatile in the Rational Expectations model.

Interestingly the correlation between TFP and output is almost the same in the RE model (0.7977) and the Bounded Rational model (0.7626), although the behavior underlying these correlations is quite different. In the RE model the positive correlation seems to hold in every period as shown in Figure 4.5. However, in the

Figure 4.6: Simulation Bounded Rational expectations

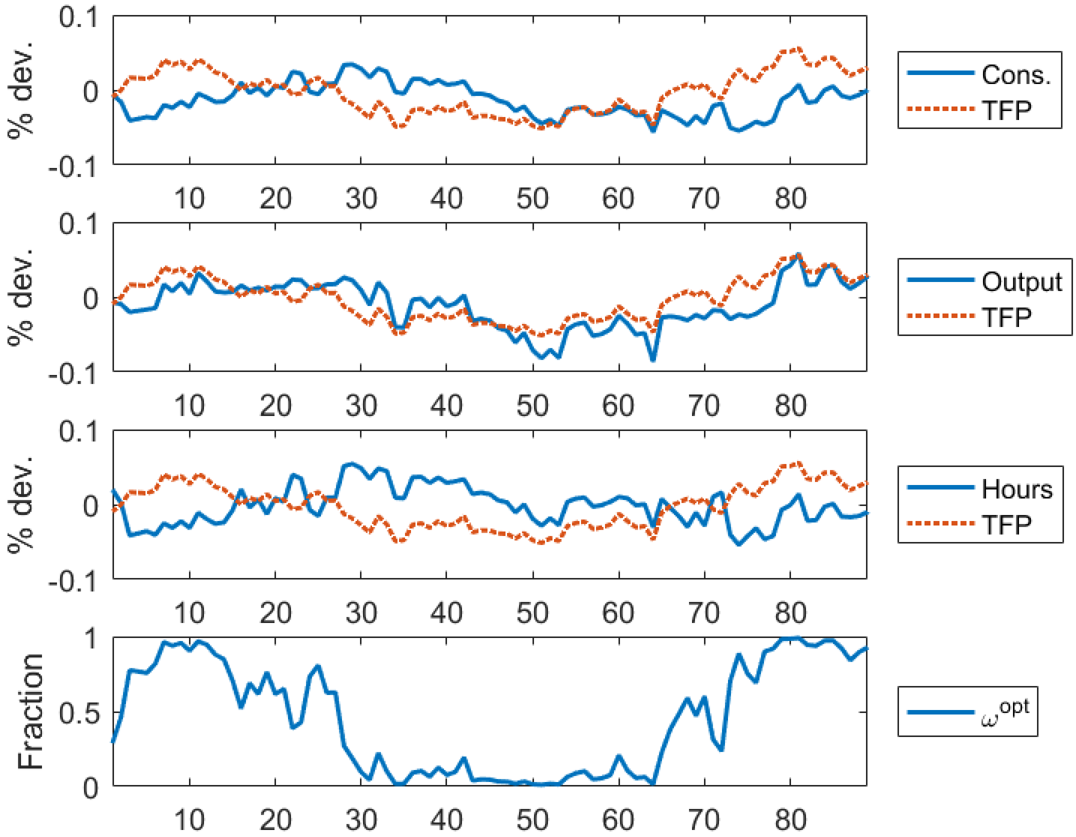

BR model the positive correlation does not hold in every period, although we see positive co-movement between output and TFP in most periods in Figure 4.6. The exceptions are the periods where sentiment is close to neutral. For example, we see output and TFP moving in different directions in the first periods, just after period 20, and around period 70. In all these periods the fraction of the optimistic rule was close to a half as shown in the last panel. This is in accordance with our above analysis of Figure 4.4: when we are in a boom or bust a positive correlation between TFP and hours worked is more likely, but around the steady state with neutral sentiment a negative correlation is more likely.

4.5 Sensitivity Analysis

The benchmark calibration above almost perfectly matched the targeted second moments, but also resulted in negative co-movement between TFP and output for

Table 4.3: Parameters with $\sigma_{cc} = 0.075$

	RE	BR
Labor supply elasticity (θ)	1.1002[a]	1.1128
Risk aversion (ν)	0.6932	0.7207
Strength of wealth effect (ϕ)	0.01191	0.00015
Standard deviation shocks (σ_a)	0.0048	0.0057
Autocorrelation shocks (ρ_z)	0.9500	
Labor share in output (α)	0.6400	
Discount factor (β)	0.9850	
Depreciation rate (δ)	0.0250	
Scalar labour supply (ψ)	1.0000	

[a] *This is the lower bound.*

Table 4.4: Business Cycle statistics with $\sigma_{cc} = 0.075$

	Data	RE	BR
St.d. Y (σ_y)	0.0170	0.0170	0.0170
Rel. st.d. C (σ_c/σ_y)	0.7500	0.6207	0.6252
Rel. st.d. N (σ_n/σ_y)	1.1200	0.9893	0.8866
Rel. st.d. I (σ_i/σ_y)	2.8500	2.3845	2.3563
Correlation Y and N	0.8600	0.9707	0.9997
Correlation Y and C	0.6600	0.9686	0.9899
Correlation Y and Z	0.7800	0.7977	0.6903
Correlation C and N	0.7300	0.8939	0.9893

Only the first four rows are targeted in the calibration.

certain states of the economy, even though the average correlation between TFP and output is very close to it's empirical value. Crucial for the good fit of the second moments and the correlation coefficients is that the economy spends most of it's time in either a pessimistic or optimistic state, which matches the strong swings in consumer confidence observed in the data. This section tests how well the Bounded Rational model can match the second moments when we fit the model to a lower standard deviation of consumer confidence.

We chose to set the standard deviation of Consumer Confidence to $\sigma_{cc} = 7.5\%$ instead of 15% in the benchmark calibration. This results in the same intensity of choice $\gamma = 2015$ as in the benchmark calibration, but a weaker memory strength with $\chi = 0.9155$ compared to $\chi = 0.9567$ in the benchmark calibration. A weaker memory strength means that expectations adjust faster to current shocks. The

Figure 4.7: IRFs to productivity shock with $\sigma_{cc} = 0.075$

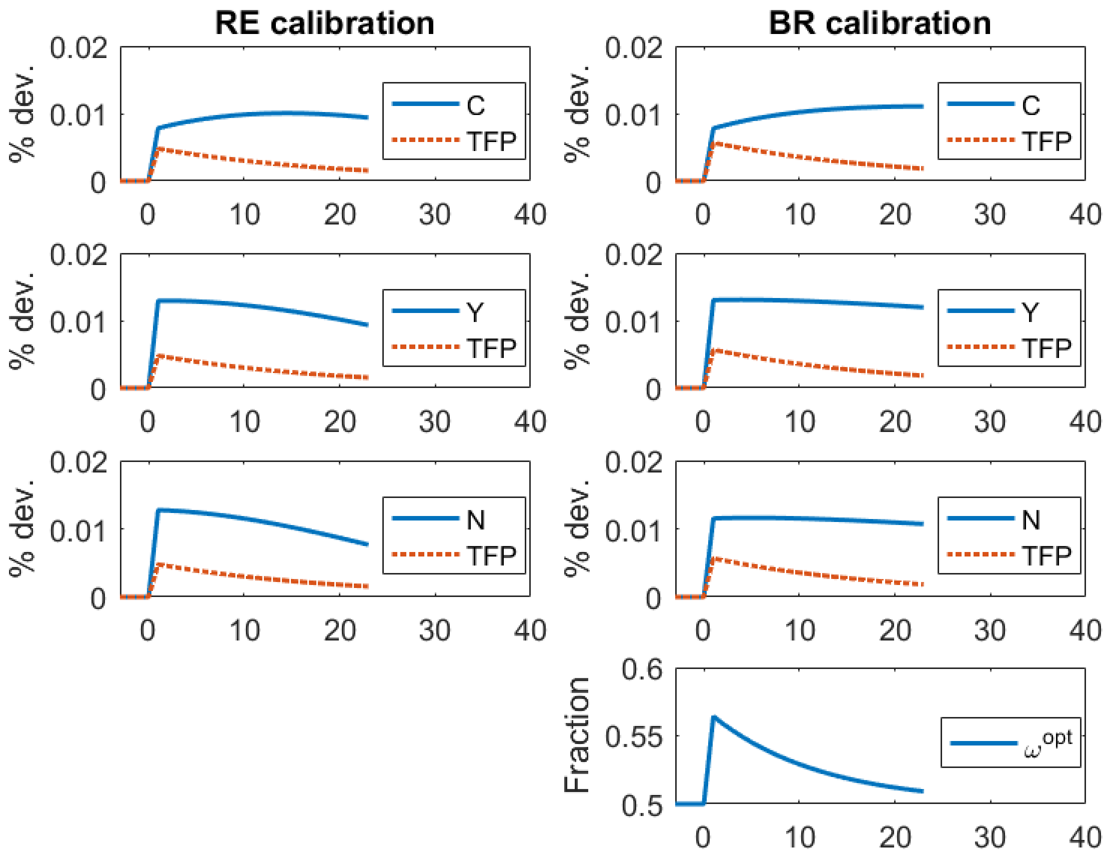

other parameters and results of the calibration are given in Table 4.3 and 4.4, which also repeats the parameters and statistics of the Rational Expectations variant for comparison.

The Business Cycle statistics with Bounded Rationality in this calibration do not improve much compared to the RE version. This alternative calibration results in ϕ close to 0, so preferences are close to GHH, meaning there is almost no wealth effect of changes in expected productivity. This is also reflected in the IRFs at the steady state in Figure 4.7, which show a positive correlation between productivity and output, hours and consumption. This positive correlation also holds for both more pessimistic and optimistic sentiments, since the wealth effect is muted. In addition, the lower standard deviation of Consumer Confidence ensures that the economy is more often in states with a relatively neutral bias as shown in Figure 4.8.

Figure 4.8: Histogram of weight of optimistic forecasting rule with $\sigma_{cc} = 0.075$

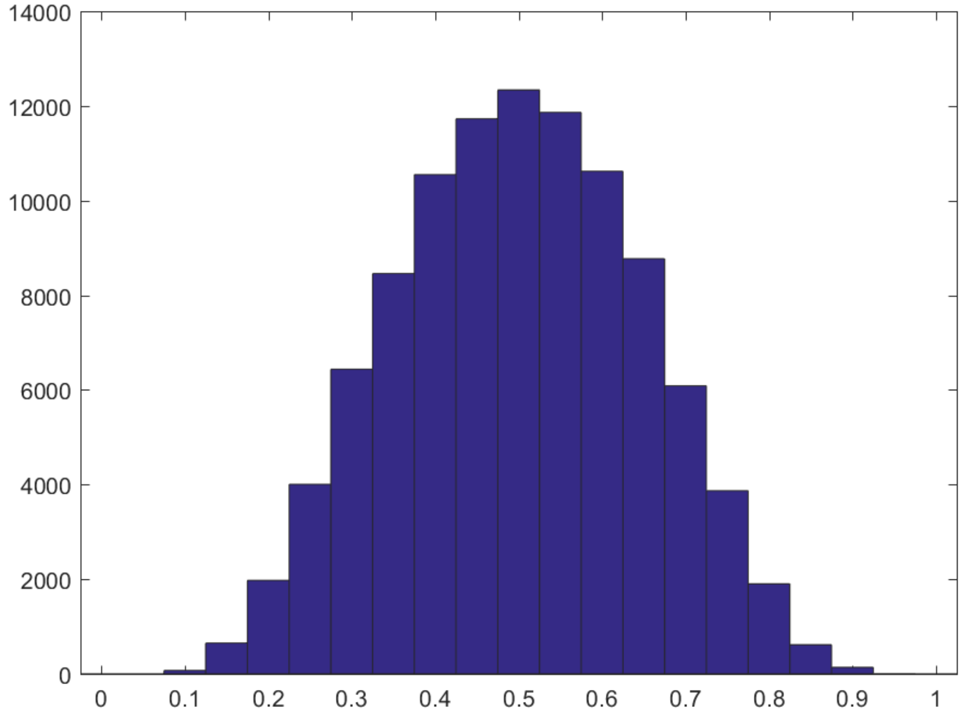

4.6 Conclusion

We have shown how an extrapolative bias with respect to the productivity shocks can improve business cycle statistics compared to a Rational Expectations model. We were able to match the relative standard deviation of hours, consumption and investment almost perfectly. At the same time we brought down the contemporaneous correlation between output and hours worked, which is close to one in standard RBC models with (close to) Rational Expectations, even when the model allows for variations in the strength of the wealth effect. To a lesser extent the Bounded Rational model also results in a contemporaneous correlation between consumption and output, and between consumption and hours that are closer to their empirical values.

Two factors are crucial for this near perfect fit of the second moments of the Bounded Rational model. The first factor is the non-monotonicity of the decision rules as a result of the strong wealth effect of changes in expected productivity. The second factor is that sentiment is either very pessimistic or very optimistic in most states of the economy, which means average sized productivity shocks do not shift expectations too much. Apparently the calibration was able to control the wealth effect to such an extend that it resulted in an almost perfectly match of the second moments. When the standard deviation of the sentiment is halved the Bounded Rational model performs about as good as the Rational Expectations version of the model.

4.A Bibliography

Beaudry, P. and F. Portier (2007). When can changes in expectations cause business cycle fluctuations in neo-classical settings? *Journal of Economic Theory 135*(1), 458–477.

Blanchard, O. (2018). On the future of macroeconomic models. *Oxford Review of Economic Policy 34*(1-2), 43–54.

Branch, W. A. and B. McGough (2011). Business cycle amplification with heterogeneous expectations. *Economic Theory 47*(2), 395–421.

Brock, W. A. and C. H. Hommes (1997). A rational route to randomness. *Econometrica: Journal of the Econometric Society*, 1059–1095.

Brock, W. A. and C. H. Hommes (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic dynamics and Control 22*(8), 1235–1274.

De Grauwe, P. (2011). Monetary policy and animal spirits. *Economic Theory 47*(2-3), 423–57.

Den Haan, W. J. and G. Kaltenbrunner (2009). Anticipated growth and business cycles in matching models. *Journal of Monetary Economics 56*(3), 309–327.

Edge, R. M., T. Laubach, and J. C. Williams (2007). Learning and shifts in long-run productivity growth. *Journal of Monetary Economics 54*(8), 2421–2438.

Eusepi, S. and B. Preston (2011). Expectations, learning, and business cycle fluctuations. *American Economic Review 101*(6), 2844–72.

Evans, G. W. and S. Honkapohja (2001). *Learning and expectations in macroeconomics*. Princeton University Press.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988). Investment, capacity utilization, and the real business cycle. *The American Economic Review*, 402–417.

Hirshleifer, D., J. Li, and J. Yu (2015). Asset pricing in production economies with extrapolative expectations. *Journal of Monetary Economics 76*, 87–106.

Hommes, C. (2011). The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic dynamics and control 35*(1), 1–24.

Jaimovich, N. and S. Rebelo (2009). Can news about the future drive the business cycle? *American Economic Review 99*(4), 1097–1118.

King, R. G., C. I. Plosser, and S. T. Rebelo (1988). Production, growth and business cycles: I. the basic neoclassical model. *Journal of monetary Economics 21*(2-3), 195–232.

Lorenzoni, G. (2009). A theory of demand shocks. *The American economic review 99*(5), 2050–2084.

Massaro, D. (2013). Heterogeneous expectations in monetary dsge models. *Journal of Economic Dynamics and Control 37*(3), 680–692.

Milani, F. (2011). Expectation shocks and learning as drivers of the business cycle. *The Economic Journal 121*(552), 379–401.

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica: Journal of the Econometric Society*, 1–48.

Chapter 5

Summary

The three essays of this thesis all make use of dynamic optimization techniques. The first two essays deal with life-cycle models, and are thus finite horizon problems which can be solved using backward iteration. The third essay concerns an infinite horizon model for which we have used projection methods as the solution method.

The first essay called 'Solving Life Cycle Models' analyzes four different numerical solution methods for solving life-cycle models with backward iteration. We use interpolation methods that explicitly preserve the monotonicity and concavity of the Value Function, such that the optimization problem has a unique solution. To be able to calculate the absolute errors of each numerical method we use a simple life-cycle model in partial equilibrium for which we can also calculate the analytical solution.

Three of the numerical methods that we evaluate use Lagrangian data for the interpolation of the Value Function. This means the interpolation method only uses information on the level of the Value Function at the gridpoints. The three methods that rely on Lagrangian data are Golden Section Search, Constrained Optimization, and Solving First Order Conditions.

Finally we also analyze a numerical method that relies on Hermite data, which means the interpolation method uses information on both the level and the slope of the Value Function. The information on the slope is obtained by adding an extra constraint to the optimization problem. The Lagrangian multiplier on this constraint will equal the slope of the Value Function.

The conclusion of the evaluation is that Solving First Order Conditions will be the most efficient method. A simple, but robust alternative is the Golden Section Search. However, both methods will only be efficient when the boundaries of the choice variable can be determined *ex ante*. When non-linear constraints apply Constrained Optimization will be more efficient. In that case the use of Hermite data can possibly add to the accuracy without increasing computation times much.

The second essay 'Optimal Age-Dependent Unemployment Insurance' evaluates what the optimal level of unemployment benefits is, when these benefits are allowed to depend linearly on age. The optimal unemployment benefits depend on the a trade off between insurance and moral hazard as higher unemployment benefits will on the one hand lower income risk, but on the other hand induce longer unemployment durations as people have less incentives to look for a job. To study how this trade off is affected by age we use a life-cycle model with job search and wealth accumulation.

The model is calibrated extensively on the following measures of the U.S. economy: the average unemployment rate, the average unemployment duration plus the increase in the average unemployment duration per year of age, the elasticity of unemployment duration with respect to the level of unemployment benefits, and finally the consumption drop when someone becomes unemployed. To match these targets we calibrate for the job loss probability, the matching efficiency of search effort (with a linear trend in age), the disutility of working, and finally the adjustment costs of wealth, which apply when an agent dissaves.

We evaluate three different unemployment systems, which are the current system with a replacement rate of 50%, an optimal constant replacement rate, and an optimal age-dependent system, where benefits depend linearly on age. The model results indicate that it is optimal to target unemployment benefits at young people. Compared to the system with an optimal constant replacement rate the increase in welfare with age-dependent benefits is about 5-15% of the Net Present Value of all unemployment insurance benefits.

In the third essay 'Adaptive Beliefs in a Real Business Cycle Model' we investigate how 'animal spirits' affect the business cycle statistics of a relatively standard real macroeconomic model. The 'animal spirits' are implemented as an extrapolative bias with respect to future productivity levels. This extrapolative bias means that recent positive changes in productivity will result in higher expected productivity in the future.

We calibrate the standard deviation and autocorrelation of these adaptive beliefs to the standard deviation and autocorrelation of a measure of consumer confidence. As a result the model with adaptive beliefs has the same degrees of freedom as the model with Rational Expectations. When we calibrate both model versions the standard Rational Expectations variant can not match the business cycle statistics well, but the version with adaptive beliefs can almost perfectly match the targeted business cycle statistics. This result, however, depends crucially on the high variance of the 'animal spirits'. When we lower the standard deviation of the adaptive beliefs the business cycle statistics are similar to those of the Rational Expectations version.