Karlsruher Institut für Technologie

Department of Economics and Management
Institute of Economics (ECON)
Prof. Dr. Johannes Brumm

Department of Informatics
Institute of Theoretical Informatics (ITI)
T.T.-Prof. Dr. Thomas Bläsius

Solving Dynamic Macroeconomic Models with an Entrepreneurial Sector

Bachelor Thesis

Submission Date: November 8, 2021

Henriette Kissling
Matr.-Nr. 2055203
henriette.kissling@student.kit.edu

Statement of Authorship

Ich versichere wahrheitsgemäß, die Arbeit selbstständig verfasst, alle benutzten Hilfsmittel vollständig und genau angegeben und alles kenntlich gemacht zu haben, was aus Arbeiten anderer unverändert oder mit Abänderungen entnommen wurde sowie die Satzung des KIT zur Sicherung guter wissenschaftlicher Praxis in der jeweils gültigen Fassung beachtet zu haben.

Karlsruhe, November 8, 2021

Solving Dynamic Macroeconomic Models with an Entrepreneurial Sector

Henriette Kissling

Abstract

In this thesis, we make two contributions: first, we transfer an overlapping-generations model of intertemporal savings and investment decisions used to evaluate taxation systems to an equivalent infinite-horizon Aiyagari-style model. We compare the results of the two model types and investigate the mechanisms at play. Our work is the first step of an ongoing research program with the overall objective of evaluating wealth taxation in the context of heterogeneous returns. We show that the results generated by the transferred model are fairly close to the empirical data on the U.S. earnings, income and wealth distribution. Crucially, the model yields almost the same extent of inequality in terms of Gini coefficients.

Second, we explore the technical boundaries of solving complex dynamic macroeconomic models with common methods and extend these methods in order to obtain a good compromise between accuracy and computational feasibility. We show that our approach manages to speedup the calculation by factor 9 and allows us to impose configurable standards on the minimum precision.

Contents

1 Introduction

In this thesis, we investigate a heterogeneous agent model of intertemporal savings decisions in the context of wealth taxes. Evaluating the economic potential of wealth taxes is interesting, as center-left parties in various economies have recently picked up on their (re-)introduction. Advocators of wealth taxation refer to the large wealth inequality and the relatively high tax burden of the working population in contrast to capital taxation. The debate has been further accelerated by the question whether the high expenses during the COVID-19 pandemic should at least partly be financed from a one-time conscription of wealth.[1]

Consequently, Guvenen et al. (2019) introduce a model to investigate wealth taxation in the presence of return heterogeneity of investments. They show that wealth taxes, in contrast to capital income taxes, shift the tax base from working to idle capital and therefore reduce the extent of capital misallocation. Furthermore, optimally designed wealth taxation can raise aggregate welfare as it subsequently allows for lower income tax rates and thus shifts the tax burden from the bulk of the working population to few individuals at the top of the wealth distribution.

These findings, however, are built on the assumption of perfect substitutability between entrepreneurial capital and private assets. In reality, entrepreneurs are bond to the decisions they made in the past, since investments cannot be withdrawn on short notice. This raises the interesting question, whether the results by Brumm and Scheidegger (2017) can be reproduced if the assumption of perfect substitutability is dropped. One conceivable approach would be to impose adjustment costs that penalize fluctuating investment decisions. As a consequence, it would be necessary to add one continuous dimension to the state space in order to distinguish between entrepreneurial capital and private assets. However, note that complex macroeconomic models like this one suffer from the *curse of dimensionality*, implying that the size of the state space grows exponentially in the number of dimensions. Thus, each additional dimension poses a substantial computational challenge.

As a potential remedy, we could switch to Adaptive Sparse Grids (ASGs) instead of conventional grids. The main idea of ASGs, as introduced by Guvenen et al. (2019), is to alleviate the *curse of dimensionality* by solving dynamic economic models on a sparse grid and adaptively refine the grid in regions with high function curvature, e.g., as in models with occasionally binding constraints. Yet, we face another drawback when attempting to solve the model with ASGs: In order to depict the lifecycle of entrepreneurs who typically start from high productivity but low capital endowment and tend to revert this relationship by the end of their life, Guvenen et al. (2019) choose an Overlapping-Generations

[1] Bundesministerium der Finanzen, 03/2021, `https://www.bundesfinanzministerium.de/Content/DE/Downloads/Ministerium/Wissenschaftlicher-Beirat/Gutachten/Vermoegensabgabe-Corona.pdf?__blob=publicationFile&v=3`

(OLG) structure. OLG models usually involve interpolation of value functions during the solution process. However, ASGs with hierarchical basis functions, as proposed by Brumm and Scheidegger (2017) do not preserve concavity. This may cause problems regarding the maximization step in the value function iteration. Thus, we circumvent the problem by transferring the OLG model into an infinite-horizon model.

We start by transferring the original overlapping-generations (OLG) structure into an Aiyagari-style model with infinite time horizon. The Aiyagari model is the obvious starting point for our analysis as it captures individual savings decisions subject to fluctuating income levels. We extend the Aiyagari model mainly by two mechanisms.

First, we add entrepreneurial engagement as a third source of income next to labor and capital income. Precisely, agents run their own businesses that turn capital into tradeable goods, and periodically decide how much capital to invest, subject to their idiosyncratic predisposition for entrepreneurial success. Adding this entrepreneurial process to the Aiyagari model adds a full layer of complexity with an associated optimization problem that needs to be solved separately in each period.

Second, we simulate the life cycle component of the OLG model and ensure constant redeployment of agents by imposing a mortality risk and replacing old agents with newborns. We provide detailed insights into the necessary adjustments to the classical solution approaches in order to incorporate the mortality process.

We then evaluate if our model can explain (i.e., reproduce) the characteristics of the U.S. economy. The model builds on the same set of assumptions about the underlying economy as Guvenen et al. (2019), but differs in the way these assumptions are implemented. In this light, we compare the results and discuss the impact of the different approaches. Furthermore, we perform a sensitivity analysis on our choice of parameters, since it is subject to high degrees of freedom.

This thesis is contributing to the research program motivated above by conducting the transformation and investigating the robustness of the results. Furthermore, it goes into the details of the necessary adjustments to the model as well as considerations regarding the efficiency of the solution process.

The scope of our work, however, goes beyond the presented results. We develop the model and choose methods to solve it such that it allows for later extensions. Specifically, we discuss the accuracy of our solution both in absolute terms and in comparison to the required computational effort. We observe that established solution approaches reach their limits when applied to our model.

Thus, we introduce a solution approach that exhausts the common practices by dynamically placing grid points at areas of interest and interpolating in between. While this does not automatically contribute to faster convergence, at least it saves us a lot of computing-intensive steps and dramatically improves the overall runtime. Crucially, our algorithm guarantees a configurable upper bound for the deviation of *all* grid points from

the real solution and thus allows us to predetermine the accuracy of the solution. We introduce a figure that allows us to quantify the loss of precision we need to cope with in exchange for the reduced complexity. Furthermore, we obtain the full-blown solutions for two selected configurations of the model and evaluate the extent to which they deviate from the approximate results.

The remainder of this thesis is organized as follows: In Chapter 2, we present the full Aiyagari-style model with infinite time horizon, extended by a mortality risk, a government and a complete additional sector of entrepreneurial business. In Chapter 3, we introduce common numerical solution approaches for dynamic optimization problems that cannot be solved analytically. In our particular case, we need to obtain the macroeconomic equilibrium as well as the associated optimal decisions and the resulting distribution of earnings, income and wealth. Additionally, we present the necessary extensions in relation to the mortality risk. Chapter 4 discusses our choice of parameters and the corresponding results of the model. Chapter 5 shifts the focus from the economic to an algorithmic perspective and examines in detail how we can obtain sufficient solutions in a computationally feasible way. We conclude with Chapter 6 where we outline the main results and sketch out further analyses conducted with our model. Furthermore, we propose ideas to extend the model in order to challenge selected assumptions.

2 The Heterogeneous Agent Model

In this section we introduce the dynamic stochastic heterogenous agent model the remainder our work is based on. We start from a model of intertemporal savings decisions made by private households which was introduced by Aiyagari (1994).

In its simplest form, the model describes households who face stochastic variation in their labor income in each period and try to hedge against this risk by accumulating savings over their lifetime. Each period t, a household aims to maximize their expected utility $u(c)$ over lifetime by choosing how much to consume and how much to save for the next period, given today's assets a_t, interest rate r and income $\varepsilon_t \cdot \bar{w}$, where ε is the stochastic labor productivity and $\bar{w}$ is the equilibrium wage per efficiency unit. This leads to the following optimization problem:

$$\max_{c_t} \mathbb{E}_0\Big[\sum_{t=0}^{\infty} \beta^t \cdot u(c_t)\Big], \tag{1}$$

subject to

$$c_t + a_{t+1} = (1+r)\ a_t + \varepsilon_t \cdot \bar{w} \tag{2}$$

$$a_{t+1} \geq -b, \tag{3}$$

where β is the discount factor on future assets, $u(c_t)$ is the household's utility function depending on the current level of consumption c_t and $-b$ is the lower bound on borrowings. The right-hand side of (2) describes the household's current income from capital and labor which must equal their total expenses on consumption and savings on the left-hand side in each period. Additionally, (3) states that households are constrained in their savings decision, which must not be lower than the borrowing constraint $-b$. The interest rate r and the wage per efficiency unit $\bar{w}$ are determined outside the optimization problem and are assumed to be constant over all periods.

The stochastic labor productivity ε is drawn from the discrete space $\mathbf{E} = \{\varepsilon_1, \ldots, \varepsilon_{n_\varepsilon}\}$. Let $\mathbf{\Pi}(\varepsilon'|\varepsilon)$ denote the probability to move from labor productivity state ε in the current period to labor productivity state ε' in the next. We assume that $\mathbf{\Pi}(\varepsilon'|\varepsilon)$ follows a Markov process which is exogenously determined outside the model and is therefore not affected by any decisions made by the agents. Note that in contrast, next period's assets a' directly result from the savings decision made in the current period. The consumption decision and the decision how much to save for next period are just two sides of the same coin, as every unit not spent on consumption will be saved for next period. In the following, we will focus on the optimal savings decision $a_{t+1}(a_t, \varepsilon_t)$ for each asset state a_t and labor

productivity ε_t given r and $\bar{w}$. This function is called the *savings policy* function. We can directly retrieve the *consumption policy* from the savings policy using (2).

2.1 Solving the Model

We can solve the optimization problem (1) with constraints (2) and (3) by setting up the Lagrangian and deriving the first-order conditions. This yields the optimality condition we will refer to as the *Euler equation* in the following:

$$\frac{\partial u(c_t)}{\partial c_t} = \beta \cdot \mathbb{E}_{t+1}\left[(1+r) \cdot \frac{\partial u(c_{t+1})}{\partial c_{t+1}}\right]. \tag{4}$$

Note that (4) breaks down the infinite-horizon problem (1) to a two-period optimality condition which must hold for each period t. Thus, we impose the condition that decisions do not change, i.e., $c_t = c_{t+1}$, and solve for the root of the Euler equation. The root yields the optimal choice of consumption with the central feature being that the marginal utility of consumption today equals the discounted expected marginal utility of consumption tomorrow times the reward on having foregone one unit of consumption today $(1 + r)$. We can insert $c_t = (1 + r) \cdot a_t + \varepsilon_t \cdot \bar{w} - a_{t+1}$ from (3) in order to find the optimal savings decision $a_{t+1}(a_t, \varepsilon_t)$ for each asset state a_t and labor productivity ε_t.

2.2 Extensions of the Ayiagari Model

In this section, we present two common extensions of the Aiyagari model introduced above. First, we make the labor decision "endogenous" by adding leisure to the utility function and letting households decide how much to work in each period subject to the current prices. Second, we add a government that levies taxes on consumption, labor and capital income on all agents.

2.2.1 Endogenous Labor Decision

So far, the only decision households make is how much to consume and how much to save in each period, subject to idiosyncratic shocks of their labor productivity. However, a more realistic approach is that households adapt their working behavior depending on their current productivity. If their current labor productivity is low, households might decide to live on their savings and reduce their hours worked, depending on their preferences. On the contrary, households might use a productivity boost to accumulate more assets and therefore temporarily decide to work more.

Since a household's labor income now depends on their choice of hours worked, the budget constraint (3) extends to

$$c_t + a_{t+1} = (1+r)a_t + \varepsilon_t \cdot \bar{w} \cdot \ell_t \tag{5}$$

$$a_{t+1} \geq -b, \tag{6}$$

where $\ell_t \in [0,1]$ denotes this period's labor choice as a fraction of total time, thus, $(1-\ell_t)$ denotes the share of free time entering the utility function. The trade-off between free time and income that can be used on additional consumption is reflected by the extended utility function $u\Big(c_t, (1-\ell_t)\Big)$. The objective function including endogenous labor choice becomes:

$$\max_{c_t, \ell_t} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \cdot u\Big(c_t, (1-\ell_t)\Big) \right]. \tag{7}$$

Introducing endogenous labor to the model adds another layer of complexity. We thereby abstract from the labor decision for now but develop the model such that we can easily adjust it for endogenous labor later. For the remainder of this thesis, we assume that agents spend all of their time working and do not gain any utility from leisure.

2.2.2 Government and Taxes on Consumption, Labor and Capital Income

Next, we add a government to the model. The government levies taxes to finance (exogenous) governmental expenses G. Let τ_{c} denote the tax rate on consumption, τ_ℓ the tax rate on labor income and τ_{cap} the tax rate on capital. In view of our later analysis, τ_{cap} can either be a tax τ_{k} on *capital income* or a tax τ_{a} on *wealth*, where the latter targets total assets and not only the returns on those assets. This yields the governmental budget constraint:

$$G = \begin{cases} \tau_{\mathrm{c}} \cdot C + \tau_\ell \cdot \bar{w} \cdot N + \tau_{\mathrm{cap}} \cdot r \cdot A, & \text{if } \tau_{\mathrm{cap}} = \tau_{\mathrm{k}} \\ \tau_{\mathrm{c}} \cdot C + \tau_\ell \cdot \bar{w} \cdot N + \tau_{\mathrm{cap}} \cdot (1+r) \cdot A, & \text{if } \tau_{\mathrm{cap}} = \tau_{\mathrm{a}}. \end{cases} \tag{8}$$

The right-hand side of the equation yields the total tax revenue, where C denotes aggregate consumption, N denotes aggregate productivity hours and A denotes aggregate capital.

At the individual level, the taxes affect each household's budget constraint, which now takes into account the after-tax values of consumption, labor and capital income. The

budget constraint including taxes depends on whether a capital income tax or a wealth tax is levied:

$$c_t \cdot (1-\tau_{\mathrm{c}}) + a_{t+1} = \begin{cases} a_t + r \cdot a_t \cdot (1-\tau_{\mathrm{cap}}) + \varepsilon_t \cdot \bar{w} \cdot \ell_t (1-\tau_\ell), & \text{if } \tau_{\mathrm{cap}} = \tau_{\mathrm{k}} \\ [a_t + r \cdot a_t] \cdot (1-\tau_{\mathrm{cap}}) + \varepsilon_t \cdot \bar{w} \cdot \ell_t (1-\tau_\ell), & \text{if } \tau_{\mathrm{cap}} = \tau_{\mathrm{a}}. \end{cases} \tag{9}$$

2.3 Entrepreneurial Business Decision

Having discussed two possible extensions of the Ayiagari model, we now add an additional source of income that is associated with a much more complex process than the labor income $\varepsilon \cdot \bar{w}$. The extension is taken from Guvenen et al. (2019), who introduce a stochastic heterogeneous agent model where agents run their own entrepreneurial businesses and produce intermediate goods x. We will refer to agents as entrepreneurs in an interchangeable manner. The key feature of the model is heterogeneity of returns among agents. Precisely, agents face different predispositions for entrepreneurial success, which we call the *entrepreneurial productivity* z. The entrepreneurial productivity is composed from a fixed inherent component and a lifecycle component and determines each agent's returns on their invested capital.

At the beginning of life, each agent is assigned an inherent *entrepreneurial ability* $\bar{z}$. However, this inherent potential is not constant over lifetime. Some agents enter life in the *fast lane* Λ_{f} with probability p_{fl}, where their ability $\bar{z}$ is further amplified by a factor $\lambda > 1$. Each period, they risk losing their place in the fast lane and ending up in the *normal lane* Λ_{n} with probability p_{n}, where their ability is no longer amplified. Agents who have not been born into the fast lane enter life in the normal lane. Both of these groups face another risk p_{r} of losing their entrepreneurial productivity completely. These agents end up in the *retirement lane* Λ_{r}, where they do not gain any returns on their investments anymore and therefore must inevitably decide to drop out of entrepreneurial engagement. The entrepreneurial productivity $z(\bar{z}, \Lambda)$ depends on inherent ability and current lane status and can then be described by:

$$z(\bar{z}, \Lambda) = \begin{cases} \bar{z}^{\lambda} & \text{if } \Lambda = \Lambda_{\mathrm{f}} \\ \bar{z} & \text{if } \Lambda = \Lambda_{\mathrm{n}} \\ 0 & \text{if } \Lambda = \Lambda_{\mathrm{r}}. \end{cases} \tag{10}$$

Figure 1 illustrates the possible transitions between lanes with the respective probabilities.

Note that in this setting, agents can never go back to a higher lane. The model by Guvenen et al. (2019) is a so-called Overlapping Generations (OLG) Model, where agents only live up to a certain number of years before they die. It is designed such that the probability to end up in the retirement state grows over lifetime, but there is still a chance

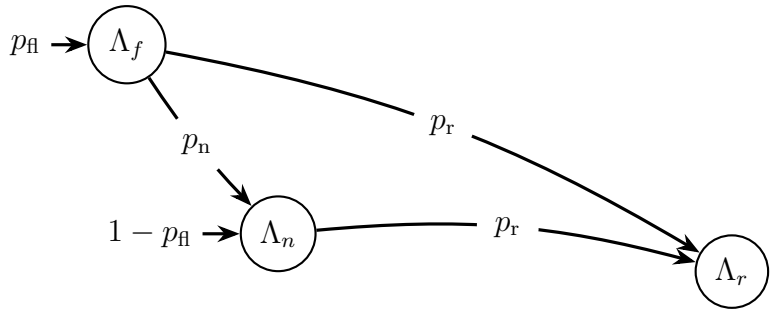

Figure 1: Possible transitions between lanes with respective probabilities. The horizontal position indicates the lifecycle, the vertical position indicates the level of entrepreneurial productivity in the respective lane.

for agents to keep up their entrepreneurial engagement until they die. With an infinite time horizon as in our model, however, all agents will inevitably end up in the retirement state. This is compensated by the mortality risk, we introduce in the following section.

2.3.1 Adding a Mortality Risk to Match the OLG Structure

In the OLG model by Guvenen et al. (2019), the transition graph (1) for entrepreneurial lanes is designed such that agents can never go back to a higher lane. However, the OLG structure ensures that each period, a fixed share of retired agents is replaced by newborns having a positive entrepreneurial productivity. In our infinite-horizon model, we capture this OLG feature by adding a lane-dependent mortality risk: Each period, we take a fixed share of agents from the retirement lane and replace them with newborns in the fast and normal lane. Thus, we impose a positive mortality risk only on agents in the retirement lane. The conditional mortality risk is:

$$\eta(\Lambda) = \begin{cases} \eta, & \text{if } \Lambda = \Lambda_{\text{r}} \\ 0 & \text{otherwise.} \end{cases} \tag{11}$$

Figure 2 illustrates the effect of the mortality risk on the possible lane transitions. Note that these additional transitions do not apply to a single agent (as they cannot be newborn) but only on a macroeconomic level in order to ensure a constant share of agents in the fast and the normal lane.

Newborns enter life without any assets. We do not take assets leftover by dead households into account. These assets accrue to the government but are not further considered.

2.3.2 Investment Decision

Subject to the equilibrium input factor price p and their current productivity z, each entrepreneur decides how much to produce of intermediate good x in the current period.

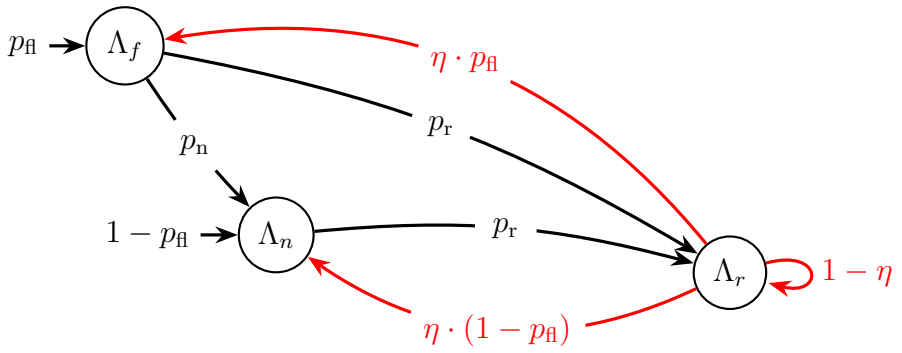

Figure 2: Lane transition graph adjusted for the mortality risk. Each period, a share η of all retired agents is redistributed from the retirement lane over the other two lanes. The fraction of agents born into the fast lane is denoted by p_{fl}.

Specifically, the amount of x is determined by the amount of capital k the entrepreneur chooses to invest in their business:

$$x = z \cdot k. \tag{12}$$

Crucially, producing the intermediate goods does *not* require any labor input. The investment k is collateralized by the entrepreneur's own assets. Entrepreneurs can borrow capital for investment in the bond market, but only up to a certain factor ϑ of their own assets a, which depends on their entrepreneurial productivity z. This yields the collateral constraint:

$$k \leq \vartheta \cdot a. \tag{13}$$

Note that (13) depends on the entrepreneurial productivity z_i, not the inherent ability $\bar{z}_i$. If $\vartheta = 1$, entrepreneurs do not have access to the bond market but can only invest their own assets in their businesses. It is reasonable that ϑ is a monotonically increasing function in z, since more productive entrepreneurs gain higher returns on their investment and are thus more likely to repay their debt. Likewise, entrepreneurs in the retirement lane face entrepreneurial productivity $z = 0$ and do not have access to the bond market, even if their inherent ability $\bar{z}$ is high.

2.4 The Complete Optimization Problem

We extend the Aiyagari model presented at the beginning of this section by the entrepreneur process from Section 2.3 in order to approach the model by Guvenen et al. (2019), with the main difference being the infinite time horizon in our model instead of the OLG structure in theirs. In this section, we focus on the conditions that must hold for the optimal solution. We will talk about how this solution can be obtained practically in Section 3. For simplicity, we suppress the subscripts i and t in the following.

Given the interest rate r, the wage per efficiency unit $\bar{w}$ and the price on intermediate goods $p(x)$, we determine the optimal savings decision of each agent, depending on their asset level a, labor productivity ε, inherent entrepreneurial ability $\bar{z}$ and current lane Λ. The entrepreneurial process we just introduced generates an additional source of income which does not depend on the savings decision and can thus be determined beforehand. We will refer to this first optimization problem as the *Static Entrepreneurial Investment Decision Problem*:

Each period, an entrepreneur chooses the amount of capital to invest in their business. As shown in (12), Guvenen et al. (2019) assume there is no labor effort necessary to produce the intermediate good x. This allows the Static Entrepreneurial Investment Decision Problem to be completely decoupled from the household's savings decision: since p, a and z are predetermined at the beginning of the period, the investment decision over k can be simply made from the first order conditions and does not need to be solved dynamically. Precisely, each entrepreneur aims to maximize their profits π

$$\pi(a, z) = \max_{k \leq \vartheta \cdot a} \left\{ p(z \cdot k) \cdot z \cdot k - (r + \delta) \cdot k \right\}, \tag{14}$$

where r is the current interest rate at which they can borrow capital in the bond market and δ is the capital depreciation rate. Section 2.4.1 describes the way to obtain the optimal investment decision and the corresponding return on investment from (14) in detail.

Having solved the Static Entrepreneurial Investment Decision Problem, we can insert this additional kind of earnings into the *Household Dynamic Programming Problem* in order to obtain the optimal savings policy.

Let $\omega(a, z; \tau_{\text{cap}})$ denote post-production, after-tax wealth. This is the sum of current assets a plus return on these assets $r \cdot a$ plus return on investment from entrepreneurial business $\pi(a, z)$ minus the tax burden which depends on the type of capital tax levied.

$$\omega(a, z; \tau_{\text{cap}}) = \begin{cases} a + \Big[ra + \pi(a, z)\Big] \cdot (1 - \tau_{\text{cap}}), & \text{if } \tau_{\text{cap}} = \tau_{\text{k}} \\ \Big[a + ra + \pi(a, z)\Big] \cdot (1 - \tau_{\text{cap}}), & \text{if } \tau_{\text{cap}} = \tau_{\text{a}}. \end{cases} \tag{15}$$

The individual dynamic programming problem is then given by:

$$\max_{c_t,(1-\ell_t)} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \cdot (1-\eta_i)^t \cdot u(c_t, 1-\ell_t) \right]. \tag{16}$$

subject to

$$c_t \cdot (1+\tau_c) + a_{t+1} = \omega(a_t, z_t; \tau_{\text{cap}}) + \varepsilon_t \cdot \bar{w} \cdot \ell_t \cdot (1-\tau_\ell) \tag{17}$$

$$a_{t+1} \geq -b. \tag{18}$$

Section 2.4.2 provides the associated Lagrangian and all first-order conditions necessary to solve (16) with constraints (17) and (18) for the optimal consumption and labor decisions.

Finally, we determine the corresponding market prices r, $\bar{w}$ and p by introducing the *final good producer*. Having introduced the dynamics of entrepreneurial engagement on the individual level, we shed light on the question how these individual decisions are put together. So far, the interest rate r, the wage per efficiency unit $\bar{w}$ and the price $p(x)$ of intermediate goods were given. We want to understand how these prices are determined on a macroeconomic level.

We use the concept of an anonymous final good producer who collects all input factors and turns them into a final good that determines the economy's aggregate output. For the aggregation, we assume an infinite continuum of agents i with total measure $\int i \, \mathrm{d}i = 1$, where the contribution of a single individual to the aggregate is negligibly small.

We assume a Cobb-Douglas production function with input factors L (labor supply) and aggregate input goods Q. We call Q the *quality-adjusted capital stock*, which is composed from the intermediate goods x_i supplied by each entrepreneur i:

$$Q = \left(\int x_i^\mu di \right)^{\frac{1}{\mu}}. \tag{19}$$

The underlying assumption is monopolistic competition among entrepreneurs who supply their intermediate goods to the final good producer. The parameter μ determines the markup over marginal costs: With perfect competition among suppliers, the equilibrium price of a supplied unit must equal the marginal costs of producing this unit. In a setting with monopolistic competition, however, each supplier has at least some market power, causing the equilibrium price to exceed the marginal costs. This difference is expressed in $1-\mu$, for example, if $\mu = 0.9$, the markup is 10%.

Labor supply and quality-adjusted capital yield aggregate output Y:

$$Y = Q^\alpha \cdot L^{1-\alpha}, \tag{20}$$

where α denotes the relative share of capital in production. The Final Good Producer chooses the amount of input factors $\int x_i \, \mathrm{d}i$ and $L = \int \varepsilon_i \cdot \ell_i \, \mathrm{d}i$ such that the total revenue is maximized subject to the individual price of the intermediate good $p(x_i)$ and wage per efficiency unit $\bar{w}$:

$$\max_{\{x_i\},L} \left(\int x_i^{\mu} \, \mathrm{d}i \right)^{\frac{\alpha}{\mu}} \cdot L^{1-\alpha} - \int p_i \cdot x_i \, \mathrm{d}i - \bar{w} \cdot L. \tag{21}$$

Note that the markup μ appears in the aggregation of the output but not in the aggregation of input factors. The first-order conditions yield:

$$p(x_i) = \alpha \cdot Q^{\alpha-\mu} \cdot L^{1-\alpha} \cdot x_i^{\mu-1} =: \mathcal{R} \cdot x_i^{\mu-1} \tag{22}$$

$$\bar{w} = (1-\alpha) \, Q^{\alpha} \cdot L^{-\alpha}. \tag{23}$$

Due to the markup μ, the equilibrium price $p(x_i)$ is calculated from the *quality-adjusted* capital stock Q (19) instead of the *unadjusted* capital stock $K = \int k_i \, \mathrm{d}i$. The absence of perfect competition implies that the equilibrium price of capital, commonly known as the interest rate, can no longer be derived from the firm's optimality condition. In general, the input factor price $p(x_i)$ does *not* equal the interest rate. We will discuss ways to determine the equilibrium interest rate in Section 3.

2.4.1 Optimality Conditions for the Static Entrepreneurial Investment Decision Problem

We insert the equilibrium price $p(x)$ of the intermediate good (22) into (25):

$$k(a,z) = \min \left\{ \left(\frac{\mu \cdot \mathcal{R} \cdot z^{\mu}}{r+\delta} \right)^{\frac{1}{1-\mu}}, \vartheta \cdot a \right\}. \tag{24}$$

The entrepreneurial profit $\pi(a,z)$ is then determined by

$$\pi(a,z) = \begin{cases} \mathcal{R}\left(z\vartheta \cdot a\right) - (r+\delta)\vartheta \cdot a, & \text{if } \; k(a,z) = \vartheta \cdot a \\ (1-\mu)\mathcal{R} z^{\mu} \left(\dfrac{\mu \mathcal{R} z^{\mu}}{r+\delta} \right)^{\frac{\mu}{1-\mu}}, & \text{if } \; k(a,z) < \vartheta \cdot a. \end{cases} \tag{25}$$

2.4.2 Optimality Conditions for the Dynamic Household Optimization Problem

Setting up the Lagrangian from (16) yields:

$$\begin{aligned}\mathcal{L} = \mathbb{E}_0\Bigg[\sum_{t=0}^{\infty}\beta^t(1-\eta_i)^t\Big\{&u(c_t, 1-\ell_t)\\ &+\lambda_t\cdot\Big[\omega(a,z;\tau_{\text{cap}})+\varepsilon_t\cdot\bar{w}\cdot\ell_t\cdot(1-\tau_\ell)-c_t\cdot(1+\tau_{\text{c}})-a_{t+1}\Big]\\ &+\mu_t\cdot a_{t+1},\Big\}\Bigg],\end{aligned} \tag{26}$$

where λ is the Lagrangian multiplier of the budget constraint (17) in period t and μ is the Lagrangian multiplier of the borrowing constraint (18). Note that for each period t, the asset choice a_t affects tomorrow's decision via a_{t+1}, but no other periods. Thus, the optimization problem can be broken down to the following two-period optimality conditions. For simplicity, we denote tomorrow's state variables with a prime, e.g. a' denotes tomorrow's assets, instead of using the subscript t. The first-order conditions are given by:

$$\frac{\partial\mathcal{L}}{\partial c} = \frac{\partial u}{\partial c} - \lambda\cdot(1+\tau_{\text{c}}) = 0 \qquad \Leftrightarrow \lambda = \frac{\partial u}{\partial c}\cdot\frac{1}{1+\tau_{\text{c}}} \tag{27}$$

$$\frac{\partial\mathcal{L}}{\partial(1-\ell)} = -\frac{\partial u}{\partial(1-\ell)} + \lambda\cdot(1-\tau_\ell)\cdot\bar{w}\cdot\varepsilon = 0 \quad \Leftrightarrow \frac{\partial u}{\partial(1-\ell)} = \frac{\partial u}{\partial c}\cdot\frac{1-\tau_\ell}{1+\tau_{\text{c}}}\cdot\bar{w}\cdot\varepsilon. \tag{28}$$

Equation (27) states that the Lagrangian multiplier of the budget constraint must equal after-tax marginal utility of consumption today. Equation (28) states that marginal utility of leisure today must equal after-tax marginal utility of consumption today times after-tax labor income per hour worked. We point out that these two conditions must hold for each period but do not depend on the other periods, so there is no need to calculate the expected value.

For the third FOC, we get:

$$\begin{aligned}&\frac{\partial\mathcal{L}}{\partial a'} = -\lambda+\mu+\beta\cdot(1-\eta_i)\cdot\mathbb{E}_{\varepsilon',z'}\left[\lambda'\cdot\frac{\partial\omega(a',z';\tau_{\text{cap}})}{\partial a'}\right] = 0\\ \Leftrightarrow\quad &\frac{\partial u}{\partial c} = (1+\tau_{\text{c}})\cdot\beta\cdot(1-\eta_i)\cdot\mathbb{E}_{\varepsilon',z'}\left[\lambda'\cdot\frac{\partial\omega(a',z';\tau_{\text{cap}})}{\partial a'}\right]+\mu.\end{aligned} \tag{29}$$

Inserting (27) into (29) yields the final Euler Equation of the complete model:

$$\frac{\partial u}{\partial c} = \beta \cdot (1 - \eta_i) \cdot \mathbb{E}_{\varepsilon', z'} \left[\frac{\partial u(c')}{\partial c'} \cdot \frac{\partial \omega(a', z'; \tau_{\text{cap}})}{\partial a'} \right] + \mu \tag{30}$$

with

$$\frac{\partial \omega(a', z'; \tau_{\text{cap}})}{\partial a'} = \begin{cases} 1 + \left(r + \dfrac{\partial \pi(a', z')}{\partial a'} \right) \cdot (1 - \tau_{\text{k}}), & \text{if } \tau_{\text{cap}} = \tau_{\text{k}} \\ \left(1 + r + \dfrac{\partial \pi(a', z')}{\partial a'} \right) \cdot (1 - \tau_{\text{k}}), & \text{if } \tau_{\text{cap}} = \tau_{\text{a}}. \end{cases} \tag{31}$$

The derivative of the entrepreneurial revenue function with respect to a' depends on whether the collateral constraint is binding or not. As long as the constraint is binding, each additional unit of assets today allows for additional investments in the amount of ϑ. Once the investment decision (24) falls below the collateral constraint, the marginal surplus of an additional unit of a is zero:

$$\frac{\partial \pi(a, z)}{\partial a} = \begin{cases} \mathcal{R} \cdot \mu \cdot (z \cdot \vartheta \cdot a)^{\mu - 1} \cdot z \cdot \vartheta - (r + \delta) \cdot \vartheta, & \text{if } k(a, z) = \vartheta \cdot a \\ 0, & \text{if } k(a, z) < \vartheta \cdot a. \end{cases} \tag{32}$$

3 Solving the Model

In this chapter, we introduce common solution approaches for the model presented in Chapter 2. The complete algorithm is composed from several nested subroutines that are dedicated to obtain locally optimal partial solutions. We start by breaking down the model to a four-dimensional state space in order to establish a shared wording for the remainder of this chapter. Afterwards, we work our way from the outer algorithm that puts the partial solutions obtained by the subroutines together to a detailed description of each subroutine. At the end of the chapter, we provide detailed information about the adjustments that arise from introducing a mortality risk to the model.

3.1 Discretization of the State Space

The state space of the optimization problem is four-dimensional.

$$\begin{aligned}
\mathbf{S} &\coloneqq \mathbf{E} \times \bar{\mathbf{Z}} \times \mathbf{L}, \\
\mathbf{E} &\coloneqq \{\varepsilon_1, \ldots, \varepsilon_{n_\varepsilon}\}, \\
\bar{\mathbf{Z}} &\coloneqq \{\bar{z}_1, \ldots, \bar{z}_{n_{\bar{z}}}\}, \\
\mathbf{L} &\coloneqq \{\Lambda_\mathrm{r}, \Lambda_\mathrm{n}, \Lambda_\mathrm{f}\}
\end{aligned}$$

An element $\mathbf{s} \in \mathbf{S}$ is a distinct combination of the three exogenous states $\varepsilon \in \mathbf{E}$, $\bar{z} \in \bar{\mathbf{Z}}$, and $\Lambda \in \mathbf{L}$. The transition matrix $\mathbf{\Pi}(\mathbf{s}'|\mathbf{s})$ holds the probability to switch from any tuple over exogenous states $\mathbf{s}$ today to any tuple $\mathbf{s}'$ tomorrow. It is calculated from the Kronecker product of the single transition matrices Π_ε, $\Pi_{\bar{z}}$ and Π_Λ.

The space over labor productivity $\mathbf{E}$, the space over entrepreneurial ability $\bar{\mathbf{Z}}$ and the lane states $\mathbf{L}$ are inherently discrete, however the asset space $\mathbf{A}$ is continuous. Since we cannot calculate the policy function and the stationary distribution analytically, we need to discretize $\mathbf{A}$. We choose a non-equidistant monotonically increasing grid $\mathbf{A} = \{a_1, \ldots, a_{n_a}\}$ of length n_a, where $a_1 = a_\mathrm{min}$ and $a_{n_a} = a_\mathrm{max}$.

The asset states are not equally spaced but concentrated at the bottom of the asset distribution in order to locate the kink at the borrowing constraint as precisely as possible. More precisely, the grid points are calculated from

$$a_i = \left(\frac{i-1}{n_a - 1}\right)^{\vartheta_a} \cdot (a_\mathrm{max} - a_\mathrm{min}) + a_\mathrm{min} \quad \forall i \in \{1, \ldots, n_a\}, \tag{33}$$

where $\vartheta_a = 2.0$.

Adding the three discrete states $\mathbf{E}$, $\bar{\mathbf{Z}}$ and $\mathbf{L}$ yields a discrete grid over states $(a, \varepsilon, \bar{z}, \Lambda)$ of shape $(n_a \times n_\varepsilon \times n_{\bar{z}} \times n_\Lambda)$.

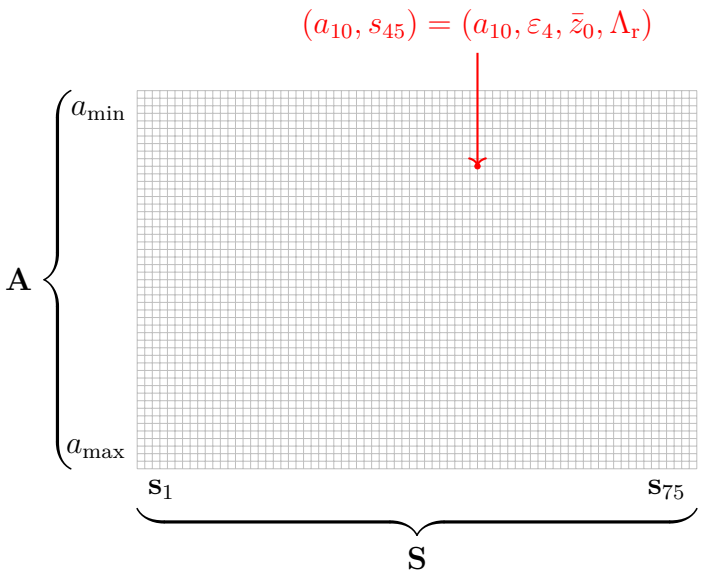

Figure 3: Projection of the four-dimensional state space onto a two-dimensional grid $\mathbf{A} \times \mathbf{S}$. The grid point (a_{10}, s_{45}) is exemplary marked.

For illustration purposes, we set $n_g \coloneqq n_\varepsilon \cdot n_{\bar{z}} \cdot n_\Lambda$ and project the four-dimensional state space $\mathbf{A} \times \mathbf{E} \times \bar{\mathbf{Z}} \times \mathbf{L}$ onto a two-dimensional grid $\mathbf{A} \times \mathbf{S}$ of shape $n_a \times n_g$, as shown in Figure 3. We will refer to the tuple $(a_i, \varepsilon_k, \bar{z}_l, \Lambda_m)$ as grid point $(a_i, \mathbf{s}_j)$, where a_i corresponds the i^{th} row of the grid and $\mathbf{s}_j = (\varepsilon_k, \bar{z}_l, \Lambda_m)$ corresponds to the j^{th} column of the grid. For each grid point $(a, \mathbf{s})$, we calculate the policy function and the stationary distribution and interpolate in between.

3.2 Finding the Stationary Equilibrium

We want to find the stationary equilibrium of the economy where prices, decisions and the distribution of agents over income, earnings and wealth remain constant. The economy described by the model is a closed economy, which implies that there are no exports and imports of goods or capital. Therefore, the total sum of capital K invested into entrepreneurial business must equal the total sum of assets A held by all individuals:

$$A \coloneqq \int a \, \mathrm{d}\Phi(a, \varepsilon, \bar{z}, \Lambda) = \int k(a, z) \, \mathrm{d}\Phi(a, \varepsilon, \bar{z}, \Lambda) =: K, \tag{34}$$

where $\Phi(a, \varepsilon, \bar{z}, \Lambda)$ is the distribution of agents over assets a, labor productivity ε, entrepreneurial ability $\bar{z}$ and lane Λ and $k(a, z)$ is the investment decision which depends on a and $\bar{z}$. In order to find the macroeconomic equilibrium, we need to identify for each market the corresponding prices where supply equals demand. These are the equilibrium interest rate r^* for the bond market, the equilibrium wage per efficiency unit $\bar{w}^*$ for the labor market and the price of intermediate goods p^* for the market of intermediate en-

trepreneurial goods. We can calculate $\bar{w}$ and p from the firm's first-order conditions (22) and (23) if we know the quality-adjusted capital stock Q and the labor supply L. In a situation with exogenous labor supply, L is constant. However, Q directly depends on the investment decisions entrepreneurs make subject to the prices *which are in turn affected by* Q. Therefore, we need an iterative process to determine the equilibrium quality-adjusted capital stock Q^*. We start from a guess $\hat{Q}$ for Q, calculate the policies for the corresponding prices $\bar{w}(\hat{Q})$ and $p(\hat{Q})$, determine the implied Q_{impl}, adjust the guess and iterate until $\hat{Q} = Q_{\text{impl}}$.

However, the market-clearing interest rate r^* cannot be obtained from the firm's first-order optimality conditions, since the quality-adjusted capital stock Q does not equal the sum of investments K. Thus, we need to make another guess $\hat{r}$ for r, calculate A and K as in (34) and, as long as these values differ, update the guess and go to the next iteration. Since r denotes the price of borrowing and lending capital respectively, aggregate savings A is a monotonically increasing function in r, while aggregate investments K is a monotonically decreasing function in r. Therefore, if aggregate savings exceed aggregate investments, we lower the interest rate and similarly, if investments exceed savings, we raise it. How do we do that?

The naive approach is to increment and decrement r by a constant value, depending of the sign of excess capital $K - A$. However, in order to prevent us from running around the equilibrium interest rate, we need to adjust the step size as the iteration converges. In order to speed up the process, consider looking at the ratio K/A instead of the difference $K - A$: if K exceeds A, aggregate investments exceed aggregate savings, which implies that the current interest rate is too low. If we multiply r with $K/A > 1$, r is increased and the change in r reflects the extent of capital shortage. Likewise, if K/A is smaller than 1, r is decreased. As the algorithm converges towards the equilibrium, the change in r is thus slowly moderated.

However, the adjustment of r does not inevitably yield the equilibrium interest rate r^* but can cause divergence if our guess for the quality-adjusted capital stock Q is not sufficiently stable. Figure 4 shows the evolution of the guessed interest rate in a situation where the solution does not converge.

Adding a smoothing factor to the K/A ratio reduces the chance of divergence, since the actual deviation is not one-to-one translated into the interest rate adjustment. For example, if we choose a smoothing factor of 0.5 and savings are twice as high as investments, the new guess is 0.75 times the old guess. Of course, this may come at the cost of additional iterations of the outer algorithm, since the adjustment is slower. However, to the best of our knowledge, this prevents the algorithm from diverging.

Algorithm 1 provides the complete algorithm that solves the model for the equilibrium interest rate r^* and quality-adjusted capital stock Q^*. Note that we use several subroutines to determine the optimal investment decision and the optimal savings policy for each

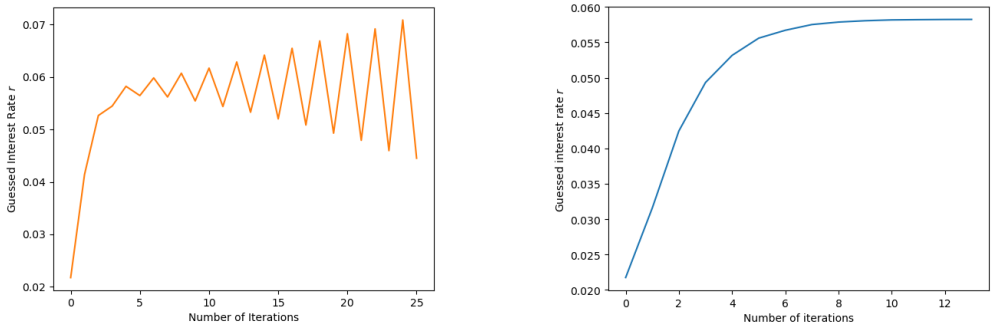

Figure 4: Divergence of the guessed interest rate r. As r and Q influence each other, the adjustment of r does not inevitably yield an equilibrium (left-hand plot). Adding a smoothing parameter of 0.5 mitigates the adjustment of the interest rate and yields convergence for the same configuration.

agent. We refer to the former as the *Static Entrepreneurial Optimization Problem*, which will be discussed in detail in Section 3.3, and to the latter as the *Dynamic Household Optimization Problem* (Section 3.4). As both optimization problems depend on the current prices r, $\bar{w}$ and p, we need to solve them in each iteration of the outer algorithm. Subsequently, we need to determine the PDF using *CDF Iteration* (Section 3.6) in order to calculate the aggregate variables Q_{impl}, K and A. The following sections provide detailed information about the subroutines.

3.3 Solving the Static Entrepreneurial Optimization Problem

As described in Section 2.4.1, the entrepreneurial investment decision problem is completely decoupled from the household's savings decision and can thus be solved beforehand. Knowing the current period's price level $p(x)$, we can calculate the optimal investment decision $k(a, z)$ from (24) and the return on investment $\pi(a, z)$ from (25). Note that $\pi(a, z)$ generates an additional source of income in the dynamic households optimization problem.

1. Choose a smoothing parameter $\rho \in [0, 1]$ and a tolerance $\varphi > 0$
2. Guess $\hat{Q}$ and $r\hat{r}$
3. Calculate $\bar{w}$ and p from the firm's FOC (22) and (23)
4. Solve the Static Entrepreneurial Optimization Problem (Section 3.3) to obtain the investment decisions $k(a, z)$ and the return on investment $\pi(a, z)$ for all asset states $a \in \mathbf{A}$ and entrepreneurial productivity states $z \in \bar{\mathbf{Z}}$
5. Solve the Dynamic Household Optimization Problem (Section 3.4) using *Time Iteration*, i.e., find the policy function $a'(a, \mathbf{s})$ for all asset states $a \in \mathbf{A}$ and exogenous states $\mathbf{s} \in \mathbf{S}$
6. Use *CDF Iteration* to determine the distribution $\Phi(a, \varepsilon, \bar{z}, \Lambda)$ (Section 3.6) from the inverse policy function $a'^{-1}(a', s)$
7. Calculate

$$Q_{\text{impl}} = \left(\int \left(k(a, z) \cdot z \right)^{\mu} \mathrm{d}\Phi(a, \varepsilon, \bar{z}, \Lambda) \right)^{1/\mu},$$
$$K = \int k(a, z) \, \mathrm{d}\Phi(a, \varepsilon, \bar{z}, \Lambda),$$
$$A = \int a \, \mathrm{d}\Phi(a, \varepsilon, \bar{z}, \Lambda)$$

8. **If** $\max\left\{ \left| K - A \right|, \left| Q_{\text{impl}} - \hat{Q} \right| \right\} < \varphi$: **break**
9. **Else:** Set

$$\hat{Q} \leftarrow \rho \cdot \hat{Q} + (1 - \rho) \cdot Q_{\text{impl}}$$
$$\hat{r} \leftarrow \hat{r} \cdot \frac{K}{A}$$

and go to 3

Algorithm 1: Outer algorithm to determine the equilibrium interest rate and capital stock

3.4 Solving the Dynamic Household Optimization Problem

We use the *Time Iteration* (TI) algorithm to determine the policy functions for savings $a'(a, s)$ and consumption $c(a, s)$. The main idea of Time Iteration (TI) is to project the infinite horizon savings decision onto a two-period decision model. Each period, a household maximizes their expected lifetime utility by choosing how much to save for next period (and thus, how much to put aside for current consumption), given their assets and net income from labor, capital and entrepreneurial business. This decision is captured by the policy function $a'(a, \mathbf{s})$, which maps the current asset level a and exogenous state $\mathbf{s} = (\varepsilon, \bar{z}, \Lambda)$ to next period's assets a', subject to the factor prices $\bar{w}$, r and p.

Given: wage $\bar{w}$ and interest rate r, profit function from the solution of the Static Entrepreneurial Optimization Problem $\pi(a, z; \tau_{\text{cap}})$ (Section 3.3)

1. Choose a smoothing parameter ρ_{TI} and tolerance value φ_{TI}
2. Guess $a' \leftarrow \hat{a}'$
3. **Loop:** For each asset state $a \in \mathbf{A}$ and $\mathbf{s} \in \mathbf{S}$, $\mathbf{s} = (\varepsilon, \bar{z}, \Lambda)$:
   (a) Calculate
   $$\hat{a}''(a, \mathbf{s}) = \sum_{\mathbf{s}' \in \mathbf{S}} \Pi(\mathbf{s}'|\mathbf{s}) \cdot \hat{a}'(\hat{a}'(a, \mathbf{s}), \mathbf{s}')$$
   (b) Use a Newton solver to obtain the root $a'(a, \mathbf{s})$ of the Euler equation (30), assuming $a'' = \hat{a}''$.
4. **If** $|a' - \hat{a}'| < \varphi_{\text{TI}}$: **break**
5. **Else:** Adjust the guess: $\hat{a}' \leftarrow \rho_{\text{TI}} \cdot \hat{a}' + (1 - \rho_{\text{TI}}) \cdot a'$

Algorithm 2: Time Iteration algorithm to determine the optimal policy functions for all asset states a and exogenous states $\mathbf{s}$ for given prices.

The complete algorithm is given by Algorithm 2. In order to determine the current policy, we solve for the root of the Euler equation (30) where the marginal utility of consumption today equals the expected discounted marginal utility of consumption tomorrow. In order to solve the Euler equation, we need to know the current after-tax wealth $\omega(a, \mathbf{s}; \tau_{\text{cap}})$ as well as the current labor income $\bar{w} \cdot \varepsilon \cdot \ell$ and, crucially, next period's savings decision $a''(a', \mathbf{s}')$, which in turn depends on the result we get for a'. This is where the idea of TI comes in: We need to *guess* the policy a' first in order to solve the Euler equation for the root, using a fixed guess $a'' = \mathbb{E}_{\mathbf{s}'}\, a'(a', \mathbf{s}')$. Note that we found the correct solution exactly when we get the same value for the root as our guess. Consequently, for any given state $\mathbf{s}$, today's policy is the same as tomorrow's. As long as the root finding yields another result than the a' we assumed in order to guess a'', we adjust the guess and

solve for the root of the Euler equation again. Once the distance between the guess and the real policy falls below a certain threshold, we consider the found root to be optimal and terminate.

3.5 Determining the Stationary Distribution

The stationary distribution $\lambda^*(\mathbf{S})$ over all exogenous states $\mathbf{s} \in \mathbf{S}$ can be calculated from the transition matrix $\Pi(\mathbf{s}'|\mathbf{s})$. By definition, the stationary distribution is a vector $v \in [0,1]^{n_g}$ that does not change when multiplied with the transition matrix Π, i.e., $\Pi(\mathbf{s}|\mathbf{s}') \cdot v = v$. We use this mathematical relationship and determine λ^* by calculating the eigenvector corresponding to the eigenvalue 1 of $\Pi(\mathbf{s}'|\mathbf{s})$ and normalize it to $\|v\|_1 = 1$ in order to obtain the stationary distribution $\lambda^*(\mathbf{s}) : \mathbf{S} \to [0,1]$. Let $\mathcal{S} \subseteq \mathbf{S} \coloneqq \mathbf{E} \times \bar{\mathbf{Z}} \times \mathbf{L}$ denote a subset of $\mathbf{S}$. We define $\mathcal{S}_{\text{retired}}$, $\mathcal{S}_{\text{normal}}$ and $\mathcal{S}_{\text{fast}}$ as the subsets of agents in the retirement (normal/ fast) lane. The corresponding measure of agents in the retirement (normal/ fast) lane is given by

$$\lambda^*(\mathcal{S}_{\text{retire}}) \coloneqq \sum_{s \in \mathcal{S}_{\text{retire}}} \lambda^*(s) \tag{35}$$

and similar for $\lambda^*(\mathcal{S}_{\text{normal}})$ and $\lambda^*(\mathcal{S}_{\text{fast}})$ respectively.

When we calculate the distribution of agents over the grid in the next step, the probability mass per column $\mathbf{s}_i$ summarized over all rows must be equal to the corresponding stationary distribution $\lambda^*(\mathbf{s}_i)$. Put differently, we need to ensure that the CDF for the highest asset state yields $\lambda^*(\mathbf{s}_i)$ for all $\mathbf{s}_i \in \mathbf{S}$.

3.6 Aggregation

Having determined the policy functions for a given interest rate r and quality-adjusted capital stock Q, we need to calculate the distribution over the grid points, i.e., the percentage of agents for any combination of assets and exogenous states. Agents "move" between grid points depending on their savings choices and the (stochastic) variation in their labor productivity ε and entrepreneurial lane Λ. Note that the inherent entrepreneurial ability $\bar{z}$ is not subject to changes over the lifetime. The projection of the four-dimensional state space onto a two-dimensional grid is helpful here: each vertical movement, i.e., change in assets, is fully captured by the policy function $a'(a, \mathbf{s})$, whereas each horizontal movement, i.e., change in exogenous states, is captured by the transition matrix $\Pi(a, \mathbf{s})$. We already know the stationary distribution over exogenous states from Section 3.5. We need to expand the stationary distribution to the asset dimension in order to find the distribution $\Phi(a, \mathbf{s})$ corresponding to the savings decisions $a'(a, \mathbf{s})$.

To calculate $\Phi(a, \mathbf{s})$, we iterate over the cumulative distribution function (CDF). Like TI, CDF Iteration is an iterative process where we initially guess the result and iterate until the solution converges. Again, we search for an equilibrium where the macroeconomic circumstances as well as the individual decisions do not change anymore. In this situation, the distribution over agents is constant.

The exact algorithm is taken from Heer and Maussner (2009). We use an asset grid $\{a_1^\circ, \ldots, a_{n_{a_f}}^\circ\}$, $a_1^\circ = a_{\min}$, $a_{n_{a_f}}^\circ = a_{\max}$ finer than the one we calculated the policy for, i.e., $n_{a_f} \gg n_a$. We start by guessing the CDF $F(a^\circ, \mathbf{s})$ over the fine grid. First, we determine the movement along the vertical dimension of the grid, i.e., the change in the asset distribution, from the *inverse policy function*: Given an asset state a' tomorrow and an exogenous state $\mathbf{s}_i$ today, the inverse policy function $a'^{-1}(a', \mathbf{s}_i)$ yields the asset level a that would make the savings choice a' optimal. We interpolate F and insert a'^{-1} to obtain the CDF that corresponds to the policy functions. However, due to the fact that assets must not exceed the grid throughout the whole iteration, we need to apply some rules:

- If the inverse policy function for $(a'_j, \mathbf{s}_i)$ yields $a'^{-1}(a'_j, \mathbf{s}_i) \geq a_{\max}$, the CDF for $\mathbf{s}_i$ reaches its maximum, thus, we set $F(a, \mathbf{s}_i) = \lambda^*(\mathbf{s}_i)$ for all $a \geq a_j$.
- If the inverse policy function for $(a'_k, \mathbf{s}_i)$ yields $a'^{-1}(a'_k, \mathbf{s}_i) \leq a_{\min}$, the CDF must be zero for this and all lower asset states, i.e., $F(a, \mathbf{s}_i) = 0$ for all $a \leq a_k$.

Applying these rules returns the CDF $F(a, \mathbf{s})$ today. Since exogenous states change between periods, we multiply $F(a, s)$ with the transition matrix $\Pi(\mathbf{s}'|\mathbf{s})$ in order to obtain the CDF $\hat{F}(a', s')$ over *tomorrow's* exogenous states $s' \in \mathbf{S}$. Note that this step is equivalent to an in-row redistribution over the grid, as it solely changes the distribution of agents over the exogenous states but not over assets. Section 3 provides the CDF Iteration algorithm. Note that in step 5(a)iv, we ensure that the CDF at the highest asset state equals the stationary distribution.

3.6.1 Aggregating with Mortality Risk

As described in Section 2.3.1, we impose a mortality risk on retired agents in order to redistribute them to the normal and the fast lane. The effect of adding a mortality risk to the model is threefold:
First, it effectively reduces the discount factor for agents who have a positive mortality risk and thus their willingness to accumulate assets. This is reflected by the Euler equation (30) but not in the transition matrix $\Pi(\mathbf{s}'|\mathbf{s})$, as the single individual does not take the possibility of being newborn into account.

Second, the mortality risk allows us to restrict aggregate savings in the model. In the simplest case, newborns enter life without any assets. With this approach we guarantee

Given: Policy functions $a'(a, s)$, transition matrix $\Pi(\mathbf{s}'|\mathbf{s})$

1. Choose a tolerance criterion φ_{CDF}
2. Choose a grid over assets $a^\circ = \{a_1^\circ, \ldots, a_{n_{a_f}}^\circ\}$ that is finer than the asset grid the policy was calculated for, i.e., $n_{a_f} \gg n_a$.
3. Calculate the (interpolated) inverse policy function $a'^{-1}(a', \mathbf{s})$ over the fine grid a°
4. Guess the piecewise cumulative distribution function (CDF) $F(a, \mathbf{s})$ over the fine grid a°
5. **Loop**:
   (a) $\forall a' \in \mathbf{A},\ \mathbf{s} \in \mathbf{S}$:
      i. Calculate $F\Big(a'^{-1}(a', \mathbf{s}), \mathbf{s}\Big)$ by interpolating $F(a, \mathbf{s})$
      ii. Set $F\Big(a'^{-1}(a', \mathbf{s}), \mathbf{s}\Big) \leftarrow 0$, if $a'^{-1}(a', \mathbf{s}) \leq a_{\min}$
      iii. Set $F\Big(a'^{-1}(a', \mathbf{s}), s\Big) \leftarrow \lambda^*(\mathbf{s})$, if $a'^{-1}(a', \mathbf{s}) \geq a_{\max}$
      iv. Set $F(a_{\max}, \mathbf{s}) \leftarrow \lambda^*(\mathbf{s})\ \ \forall \mathbf{s} \in \mathbf{S}$
   (b) For all $a' \in \mathbf{A},\ \mathbf{s}' \in \mathbf{S}$ calculate

$$\hat{F}(a', \mathbf{s}') = \sum_s \Pi(\mathbf{s}'|\mathbf{s}) \cdot F\left(a'^{-1}(a', \mathbf{s}), \mathbf{s}\right)$$

   (c) **If** $\left\|\hat{F} - F\right\|_{\max} < \varphi_{\text{CDF}}$: **break**
   (d) **Else**: Update the guess by setting $F = \hat{F}$
6. Calculate the PDF $\Phi(a, \mathbf{s})$ from the CDF

Algorithm 3: CDF Iteration algorithm to determine the PDF/ CDF over all grid points

to always have a minimum proportion of agents at the lowest asset level. The effective mortality risk, i.e., the total share of agents who are redistributed in each period, is determined by the measure of retirement states in the stationary distribution:

$$\eta_{\text{eff}} = \eta \cdot \lambda^*(\mathcal{S}_{\text{retired}}) \tag{36}$$

Third, we do not only redistribute agents over the asset grid but also over the exogenous states. From a macroeconomic perspective, agents who die simply lose their assets and switch from the retirement lane to the fast or normal lane. Let p_{fl} denote the fixed share of agents born into the fast lane and $(1 - p_{\text{fl}})$ the fixed share of agents born into the normal lane. We need to ensure a minimum share of $\eta_{\text{eff}} \cdot p_{\text{fl}}$ in the lowest asset state of the fast lane and of $\eta_{\text{eff}} \cdot (1 - p_{\text{fl}})$ in the lowest asset state of the normal lane respectively.

Whereas a single individual can, theoretically, live infinitely long in our model but cannot be newborn once having died, this *is* possible at the macroeconomic level. Effectively, introducing a mortality risk reduces the probability to remain in the retirement lane by η and increases the chance to switch from the retirement lane to the normal lane by $\eta \cdot (1 - p_{\text{fl}})$ and to the fast lane by $\eta \cdot p_{\text{fl}}$ respectively. The corresponding lane transition graph is shown by Figure 2, p. 9.

Specifically, we need to consider the modified lane transition graph from Figure 2 instead of Figure 1 (p. 1) in order to obtain the stationary distribution λ^*. However, in the CDF Iteration, using 2 will not serve our purpose, as the process of death and rebirth is inextricably linked to the complete loss of assets. Instead, we manually reduce the probability mass at each retirement lane grid point by η and redistribute η_{eff} over the fast and normal lane states with respect to p_{fl} and the stationary distribution λ^*.

Algorithm 4 provides the additional steps (bb) and (cc) we need to insert into the CDF Iteration algorithm proposed by Heer and Maussner (2009). Note that we use the "naive" transition matrix $\Pi(\mathbf{s}'|\mathbf{s})$ corresponding to Figure (1) in step (bb) and manually replace dead agents with newborns. Crucially, in step (cc), newborn agents draw their exogenous states from the stationary distribution and not from the last state at which their predecessor left the model. This procedure is consistent with the determination of the policy functions for the households.

5. **Loop**:

(bb) Take a fixed share η of agents from each grid point in the retirement lane, i.e., $\forall \mathbf{s} \in \mathcal{S}_{\text{retire}}$:

$$F\Big(a'^{-1}(a', \mathbf{s}), \mathbf{s}\Big) \leftarrow F\Big(a'^{-1}(a', \mathbf{s}), \mathbf{s}\Big) \cdot (1 - \eta)$$

(c) For all $a' \in \mathbf{A},\ \mathbf{s}' \in \mathbf{S}$ calculate

$$\hat{F}(a', \mathbf{s}') = \sum_{s} \Pi(\mathbf{s}'|\mathbf{s}) \cdot F\left(a'^{-1}(a', \mathbf{s}), \mathbf{s}\right)$$

(cc) Redistribute the share of dead agents η_{eff} to the normal and the fast lane in the respective measures of the stationary distribution, i.e., $\forall a \in \mathbf{A}$:

$$F(a, \mathbf{s}_i) = F(a, \mathbf{s}_i) + \begin{cases} p_{\text{fl}} \cdot \eta_{\text{eff}} \cdot \dfrac{\lambda^*(\mathbf{s}_i)}{\lambda^*(\mathcal{S}_{\text{fast}})}, & \text{if } \ \mathbf{s}_i \in \mathcal{S}_{\text{fast}} \\ (1 - p_{\text{fl}}) \cdot \eta_{\text{eff}} \cdot \dfrac{\lambda^*(\mathbf{s}_i)}{\lambda^*(\mathcal{S}_{\text{normal}})}, & \text{if } \mathbf{s}_i \in \mathcal{S}_{\text{normal}} \end{cases}$$

Algorithm 4: Extension of the CDF Iteration algorithm for the mortality risk.

4 Policy Analysis

In Chapter 3, we described the algorithms to iteratively determine the macroeconomic equilibrium. However, whether there is an equilibrium at all, i.e., whether we find a (positive) interest rate such that the market clearing condition (34) (p. 16) holds, highly depends on the choice of economic parameters. Thus, our goal is to find a parametrization that yields a realistic equilibrium. We assess the extent by which our model represents the real world by targetting the U.S. distribution over earnings, income and wealth found by Kuhn, Rios-Rull, et al. (2016).[2]

In the following section, we propose a configuration that yields an equilibrium and is as close to the empirical U.S. distribution data as we could get. We refer to this configuration as the *benchmark configuration* in the following. Subsequently, we investigate the shape of the policy functions as well as the macroeconomic parameters and the distribution over earnings, income and wealth and evaluate possible channels to explain our observations. Furthermore, we compare our results to the ones by Guvenen et al. (2019) our model builds on.

Finally, we conduct sensitivity checks by systematically changing selected parameter values and analyzing their effect on the model. This helps us to develop an intuition for the dynamics of the model and to understand the influence of the choice of parameters.

4.1 Parametrization of the Benchmark Model

In the following, we present our choice of economic parameters the remainder of this (and the next) chapter builds on. We refer to this configuration as the *benchmark configuration*.

Choice of Grid Points and Discretization of the Continuous Asset Space We need to choose the maximum asset level $a_{\max}$ carefully such that optimal savings do not exceed the grid. Crucially, if the share of agents in the highest asset state of the stationary distribution is unreasonably high, this indicates that $a_{\max}$ is too small and we need to extend the grid. For now, we set $a_{\max} = 100$ but we monitor the corresponding share in the PDF closely throughout our investigation.

We set the number of asset states to $n_a = 51$, the number of labor productivity states to $n_\varepsilon = 5$ and the number of entrepreneurial ability states to $n_{\bar{z}} = 5$. The number of lanes is $n_\Lambda = 3$ by construction. This yields a grid of $51 \times (5 \cdot 5 \cdot 3) = 51 \times 75$ entries. The grid concentration parameter ϑ_a (see Section 3.1) is set to 2. For the CDF Iteration we need a finer grid over assets than the one we calculated the policies for. We set the number of asset states for the fine grid $n_{a_f} = 1001$ and the concentration parameter to $\vartheta_{a_f} = 4.0$.

[2] In conformity with Kuhn, Rios-Rull, et al. (2016), we use the term "(labor) earnings" on labor income and "income" on total income from labor, returns on capital and returns on entrepreneurial investment.

Private Households The utility function of the household is a Cobb-Douglas function:

$$u(c,(1-\ell)) = \frac{\left(c^{\gamma}(1-\ell)^{1-\gamma}\right)^{1-\sigma} - 1}{1-\sigma}, \tag{37}$$

where σ is the risk aversion parameter and γ is the share of consumption in utility.

Note that we assume exogenous labor, i.e., households do not decide over their amount of hours worked. Thus, the optimization problem becomes much simpler.[3] We set $\gamma = 1$, which implies that the household gains all its utility from consumption and none from leisure and make sure the choice of labor hours ℓ is always 1. The simple utility function that does not take leisure into account is

$$u(c) = \frac{c^{1-\sigma} - 1}{1-\sigma}. \tag{38}$$

We set the risk aversion parameters σ to 2, which is at the low side and implies a rather small preference of agents to smooth their consumption over lifetime. In particular, the higher σ the less volatile the savings subject to changing circumstances.

We set $\beta = 0.92$ which implies that one unit of consumption tomorrow yields only 92% of the utility it would yield today.

The log labor productivity $\log \varepsilon$ is generated from an AR1 process with persistence $\rho_\varepsilon = 0.965$ and standard deviation $\sigma_\varepsilon = 0.25$. This yields $\varepsilon_{\text{min}} = \varepsilon_0 = 0.06$, $\varepsilon_{\text{median}} = \varepsilon_2 = 1$ and $\varepsilon_{\text{max}} = \varepsilon_4 = 17.46$. The stationary distribution is provided by Table 1. The bulk of agents (44%) has a median labor productivity.

Table 1: Measure of each labor productivity state ε in the stationary distribution.

	ε_0	ε_1	ε_2	ε_3	ε_4
Measure	0.04	0.24	0.44	0.24	0.04

Entrepreneurs The evolution of inherent entrepreneurial ability $\bar{z}$ is simpler in our model than in the one by Guvenen et al. (2019), since we do not consider inheritance between parents and their children. Instead, the logarithmized ability states $\log \bar{z}$ are drawn from an equally spaced discrete grid of length $n_{\bar{z}} = 5$, where $\log \bar{z}_{\text{min}} = -\sigma_{\bar{z}}$ and $\log \bar{z}_{\text{max}} = 3 \cdot \sigma_{\text{z}}$. Following Guvenen et al. (2019), we set $\sigma_{\bar{z}} = 0.12$ such that $\bar{z}_{\text{min}} \approx 0.89$ and $\bar{z}_{\text{max}} \approx 1.43$. We choose the distribution over $\bar{z}$ such that only a small elite of agents (5%) is highly productive, as shown by Table 2.

The probabilities of leaving the fast lane $p_1 = 0.05$ and retiring from entrepreneurial engagement $p_2 = 0.03$ are taken from Guvenen et al. (2019). Table 3 shows the corresponding probabilities to end up in the retirement, normal or fast lane respectively at different ages for agents born into the fast lane. Agents born into the normal lane skip

[3] See Appendix A for the complete optimization problem with endogenous labor

Table 2: Measure of each inherent entrepreneurial ability state $\bar{z}$ in the stationary distribution.

	$\bar{z}_1$	$\bar{z}_2$	$\bar{z}_3$	$\bar{z}_4$	$\bar{z}_5$
Measure	0.35	0.3	0.2	0.1	0.05

the fast lane but face the same probability to retire at any age. After 80 periods, agents born into the fast or the normal lane will have retired with a probability of more than 90 percent. This probability implies a mean duration of 12.5 years in the fast lane and another 18 years until retirement.

Table 3: Guvenen: Probability to be in the retirement lane, normal lane or fastlane at a certain age for agents who start life in the fastlane.

age	p_{retire}	p_{normal}	p_{fast}
1	0.03	0.05	0.92
10	0.26	0.30	0.44
20	0.46	0.35	0.19
30	0.5	0.35	0.15
40	0.60	0.32	0.08
50	0.70	0.26	0.04
60	0.78	0.20	0.015
70	0.84	0.15	0.007
80	0.88	0.116	0.003
90	0.91	0.086	0.001

As described in Section 2.3.1, we impose a mortality risk η on agents in order to match the OLG structure. We set $\eta = 0.2$ for now, which implies an unreasonably short mean duration of 5 years in the retirement lane. However, remember that we use the mortality risk as a tool to limit savings and ensure a minimum fraction of η_{eff} at the lowest asset states. We will try to set a more realistic value for the mortality risk later. For the share of agents born into the fast lane, we choose $p_{\text{fl}} = 1.0$. This yields the stationary distribution over lanes presented by Table 4 as well as an effective mortality risk $\eta_{\text{eff}} = 2.61\%$.

Table 4: Measure of each entrepreneurial lane in the stationary distribution.

	Λ_{r}	Λ_{n}	Λ_{f}
Measure	0.13	0.54	0.33

Each period, agents decide how much they want to invest in their businesses. As proposed in Section 3.3, the optimal investment choice function $k(a, \mathbf{s})$ given exogenous state tuple $\mathbf{s} \in \mathbf{S}$ and collateral constraint ϑ is:

$$k(a, \mathbf{s}) = \min \left\{ \left(\frac{\mu \mathcal{R} z^{\mu}}{r + \delta} \right)^{\frac{1}{1-\mu}}, \vartheta a \right\}, \tag{39}$$

where z denotes the entrepreneurial productivity depending on inherent ability $\bar{z}$ and current lane Λ. Very productive agents with high values of z (i.e., agents in the fast lane or agents in the normal lane with very high inherent ability $\bar{z}$) are subject to a collateral constraint ϑ, as the first entry of the minimum function grows rapidly in z. Thus, investments are highly sensitive to the choice of ϑ. Like Guvenen et al. (2019), we define five groups $i = 1, \ldots, 5$ that correspond to the five inherent ability states $\bar{z}_i$ and define the collateral constraint $\vartheta(i) = 1 + 1.5 \cdot (i-1)/(n_{\bar{z}} - 1)$ for each group. Note that ϑ meets the condition that retired agents must not have access to the bond market anymore (since they are assigned to the lowest group $i = 1$) we imposed in Section 2.3.2. Furthermore, for agents with $\bar{z} > \bar{z}_{\text{median}} = 1$ the collateral constraint is loosened as long as they experience positive amplification of their ability in the fast lane.[4]

The Final Good Producer We set the share of capital in production to $\alpha = 0.4$ and the markup over marginal costs to $\mu = 0.9$ as in the model by Guvenen et al. (2019). The capital depreciation rate is $\delta = 0.04$.

Government We choose the same tax rates as Guvenen et al. (2019), setting the labor tax rate to $\tau_\ell = 0.224$, the consumption tax rate to $\tau_\text{c} = 0.075$ and the capital income tax rate to $\tau_\text{cap} = 0.25$. For now, we do not require the government budget to be balanced but let the government expenses scale up and down with the total tax income.

Table 5 gives an overview over all chosen parameters.

[4]If the $\bar{z}$ values are chosen as indicated above, an amplification by $\lambda = 2$ implies a step to the next higher i group and thus access to additional borrowings of $0.5 \cdot a$ for the $\bar{z}_3$ and $\bar{z}_4$ group. With $\lambda = 3$ and greater, all three $\bar{z}$ groups above the median have maximum access to the bond market. For the highest $\bar{z}$ group, ϑ is at maximum in the normal lane already.

Table 5: Choice of economic parameters.

Economic Parameters		
Parameter	Symbol	Value
Lower bond on (private) borrowings	-b	-0.0
Risk aversion	σ	2.0
Share of consumption in utility	γ	1.0
Discount factor	β	0.92
Persistence of the labor productivity process	ρ_ε	0.965
Standard deviation of the labor productivity process	σ_ε	0.25
Standard deviation of log entrepreneurial ability	$\sigma_{\tilde{z}}$	0.12
Probability of leaving the fast lane	p_1	0.05
Probability of retirement	p_2	0.03
Conditional mortality risk when in the retirement lane	η	0.2
Share of agents born into the fast lane	p_{fl}	1.0
Share of capital in production	α	0.4
Markup over marginal costs	μ	0.9
Labor income tax rate	τ_{l}	0.224
Consumption tax rate	τ_{c}	0.075
Capital income tax rate	τ_{cap}	0.25
Technical Parameters		
Parameter	Symbol	Value
Minimum asset level	$a_{\min}$	0.0
Maximum asset level	$a_{\max}$	100.0
Number of asset states (normal grid)	n_a	51
Number of asset states (fine grid)	n_{a_f}	1001
Parameter of asset state concentration	ϑ_a	2.0
Parameter of fine asset state concentration	ϑ_{a_f}	4.0
Number of labor productivity states	n_ε	5
Number of entrepreneurial ability states	$n_{\tilde{z}}$	5
Number of entrepreneurial lanes	n_Λ	3
Smoothing parameter for Q adjustment	ρ_Q	0.5
Smoothing parameter for r adjustment	ρ_r	0
Tolerance criterion for outer algorithm	φ	0.01
Tolerance criterion for TI	φ_{TI}	0.001
Tolerance criterion for CDF Iteration	φ_{CDF}	10^{-6}

4.2 Macroeconomic and Distributional Results

With the parametrization proposed in the last Section, we find the equilibrium at an interest rate of 7.11%. The ratio of the quality-adjusted capital stock Q to output is 6 and the ratio of Q to aggregate labor supply L is 14. The total tax revenue is 18% of total GDP, of whom 53% account for labor income taxation, 25% for consumption taxation and 22% for capital taxation. Finally, the ratio of aggregate input goods to aggregate investments is around 7.6.

The PDF over asset states yields a share of 0.6% at the highest asset state, of whom 3% account for the retired agents, 11% account for agents in the normal lane and the remaining 86% for agents in the fast lane. The measure of agents at the lowest asset state is 9.5%. Recall that the effective mortality risk, i.e., the share of agents who are redistributed in each period, is 2.61%, indicating that the model naturally generates a substantial additional share of agents at the bottom asset level. The remaining 6.89% at the lowest asset state are solely retirees. The corresponding PDF and CDF are shown by Figure 5. There is a small outlier at the highest asset state indicating that the maximum asset state is exceeded. We will check this when conducting the sensitivity analyses.

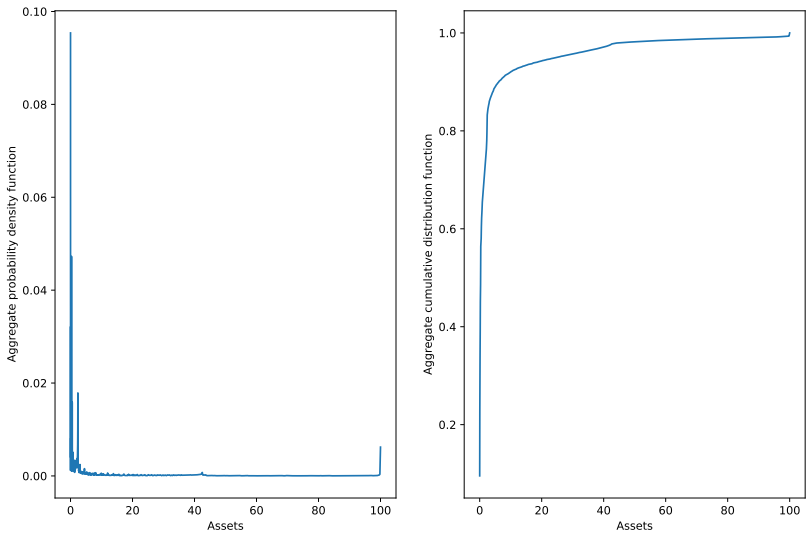

Figure 5: PDF (left) and CDF (right) over asset states for the benchmark configuration.

Next, we investigate the distribution of aggregate savings over the three entrepreneurial lanes. We find that 7% of all savings are made by retirees whose measure in the stationary distribution is almost twice as large at 13%. The bulk of savings comes from agents in the normal lane with a share of 54%, exactly matching their respective measure in the stationary distribution. Agents in the fast lane make savings above average, as their

contribution to the aggregate capital stock is 39% compared to measure 33% in the stationary distribution.

In order to assess the extent of inequality generated from the benchmark configuration, we compare the distribution over earnings, income and wealth to the empirical data found by Kuhn, Rios-Rull, et al. (2016). Specifically, we partition the data by earnings, income and wealth and determine the quintiles as well as the 90-95%, 95-99% and 99-100% percentiles. Subsequently, we compute the share of earnings, income and wealth of each percentile group and evaluate how well our model fits the empirical data.

Figures 6 to 8 compare the empirical data with the distribution generated from the model. Looking at the quintiles, the model does a fairly good job at matching the real data with a maximum deviation of 3.7 percentage points for the earnings distribution, 2.5 percentage points for the income distribution and 4 percentage points for the wealth distribution. However, the 90-95%, the 95-99% and the 99-100% percentile need further refinement. If our model is to be aligned with the findings by Guvenen et al. (2019), we need to put emphasis on matching the top 10% and 1% closely. Figure 8 shows that the upper 10% of the wealth distribution in our model hold 83% of the total wealth, which overstates the empirical share of 75% from the U.S. and also the share of 66% generated by Guvenen et al. (2019). The shares of the 95-99% quintile generated by our model tend to be too large in all cases, whereas the top 1% shares tend to be too small.

The respective Gini coefficients are 0.54 (0.67) for the earnings partition, 0.52 (0.58) for the income partition and 0.75 (0.85) for the wealth partition (bracketed values denote the target Gini coefficients from the empirical data), which is already promising but also emphasises the need to adjust the configuration. We will evaluate if other configurations fit the empirical data better in the next section.

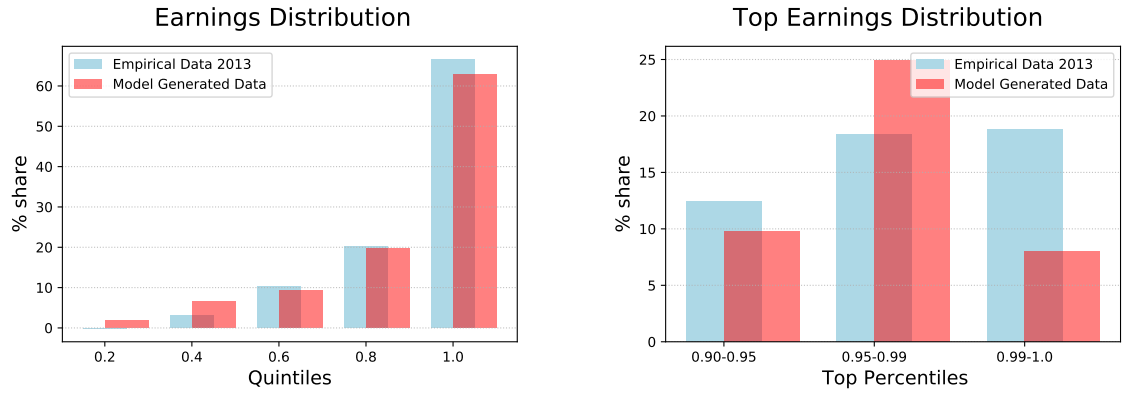

Figure 6: Share of the quintiles (left) and upper percentiles (right) on total earnings. The blue bars show the empirically determined data while the red bars show the data generated by the model. All percentiles are calculated over the earnings partition. Empirical data from Kuhn, Rios-Rull, et al. (2016).

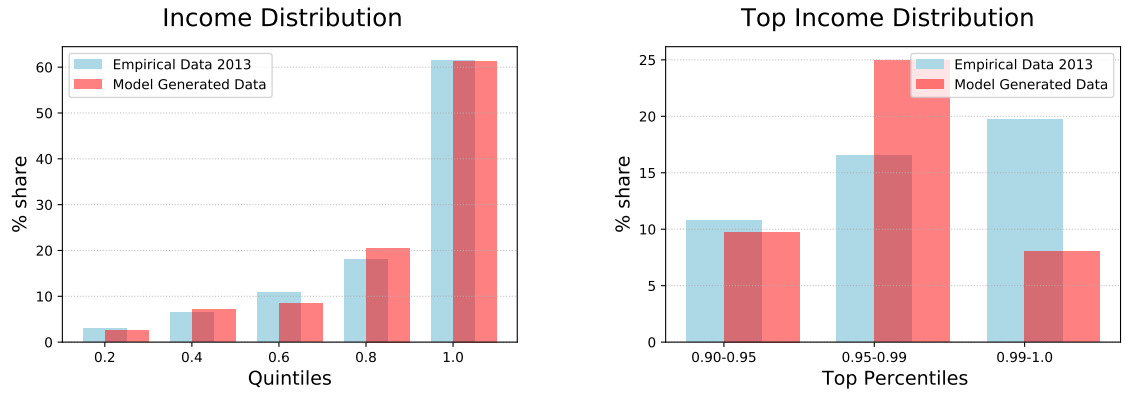

Figure 7: Share of the quintiles (left) and upper percentiles (right) on total income (labor earnings, capital income and entrepreneurial profits). The blue bars show the empirically determined data while the red bars show the data generated by the model. All percentiles are calculated over the income partition. Empirical data from Kuhn, Rios-Rull, et al. (2016).

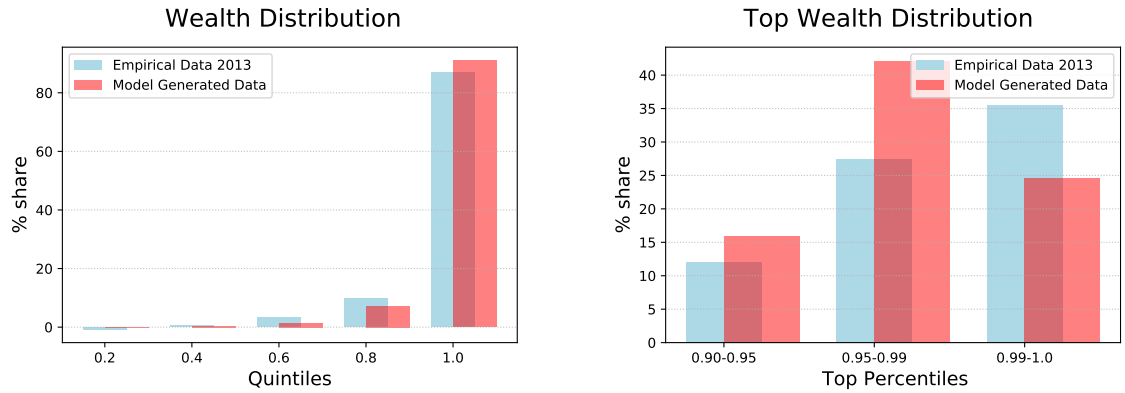

Figure 8: Share of the quintiles (left) and upper percentiles (right) on total wealth. The blue bars show the empirically determined data while the red bars show the data generated by the model. All percentiles are calculated over the wealth partition. Empirical data from Kuhn, Rios-Rull, et al. (2016).

Table 6 outlines the main findings in comparison to the figures found by Guvenen et al. (2019). Our model manages to match some of their targeted moments better than others. There are two possible explanations that are related to the specific implementations of the models, besides the differences in the parametrization. First, our model features less variability of the inherent entrepreneurial ability $\bar{z}$. As a result, the most productive entrepreneurs are gathered in the fast lane state of the highest ability $\bar{z}_4$ that still accounts for 1.6% of the stationary distribution. In order to consider few *billionaires* with substantially higher productivity, as Guvenen et al. (2019) do, we would need a finer resolution for the ability states.

Furthermore, as discussed in Section 2.3.1, when agents die in our model, their assets accrue to the government and are not further considered. Crucially, we do not model any parent-child relationship that could serve as a channel for bequests. Therefore, we do not allow for billionaires by inheritance, which makes it harder for agents to accumulate extremely large asset levels. If at all, our model allows solely for "self-made billionaires".

Table 6: Macroeconomic Measures found by Guvenen et al. (2019) in comparison to the benchmark model and sensitivity checks. Abbreviations: Total Tax Revenue (TTR), Gross Domestic Product (GDP), Capital Income Tax Revenue (CITR)

	U.S. Data	Guvenen Model	Benchmark Model
Capital-Output-Ratio	3	3	5.96
Top 10% Wealth Share [%]	75	66	83
Top 1% Wealth Share [%]	36	36	25
TTR / GDP [%]	29.5	25	18
CITR / GDP [%]	28	25	22
Bequest / Wealth [%]	1.29	0.99	1.35

4.3 Policy Evaluation

Figure 9 shows the savings policy functions $a'(a, \mathbf{s})$ over the asset grid $\mathbf{A}$. We plot the policy for each lane Λ, holding labor productivity $\varepsilon = \varepsilon_4$ and entrepreneurial ability $\bar{z} = \bar{z}_4$ fixed at their highest level each. The grey line indicates the savings level at which the asset level is kept constant.

All policy functions are strictly monotonically increasing in the asset level a, thus, higher assets correspond with higher savings. Furthermore, we observe that the savings level rises in effective entrepreneurial productivity, implying that fast lane agents save the most and their net savings are positive. On the contrary, agents in the retirement lane live on their savings, i.e., save less than their current asset level. However, when we evaluate the policy functions for the benchmark configuration, several aspects leap to the eye. We identify two occurences of nonlinearities from the plots: First, the retirement lane policy shows a kink at the bottom of the asset range. Second, although not manifested in a

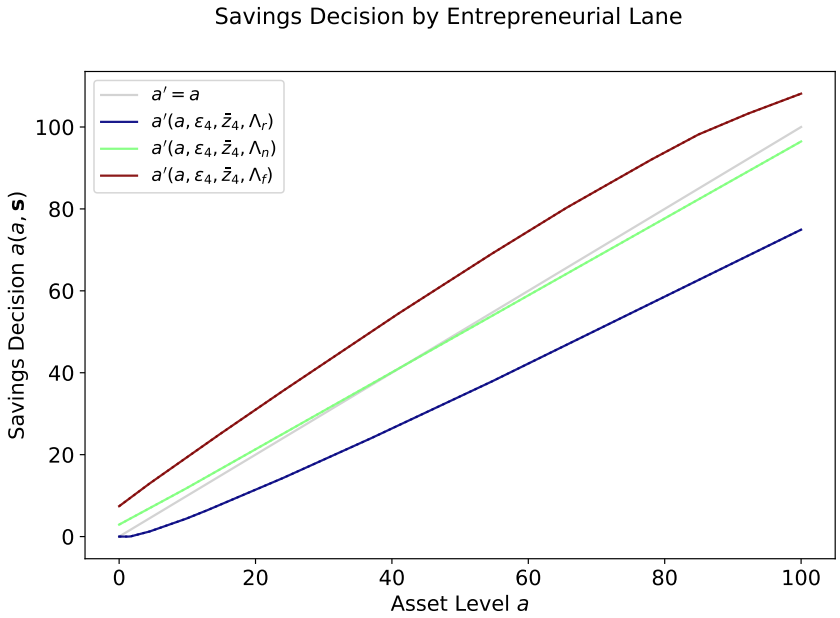

Figure 9: Savings policy functions for fixed labor productivity ε_4 and entrepreneurial ability $\bar{z}_4$ over all three lanes. The grey line indicates the savings level where agents keep their asset levels constant.

distinct kink, the slope of the fast lane policy function changes within the upper quarter of the asset range and correspondingly, the fast and the normal lane policy approach each other beyond this point.

In the following section, we evaluate the different sources of income and future prospects agents base their decisions on in order to explain the observed behavior of the policy functions.

4.3.1 Sources of Income

In order to better understand the effect of the different channels, let us recap the three possible sources of income in our model: First, labor earnings, which, due to the exogenous labor choice, solely depend on the wage $\bar{w}$ per efficiency unit and the current labor productivity ε but *not* on the asset level a. Second, returns of investment in entrepreneurial activities $\pi(a, z; \tau_{\text{cap}})$ as in (25) (p. 12), which are growing in a as long as the entrepreneur is constrained in their investment decision and are constant otherwise. Third, returns on lending capital at the bond market, which grow linearly in the asset level a. Thus, we expect the impact of the different labor productivity levels to decrease with rising asset levels, as other sources of income gain relative importance.

Table 7 and Table 8 provide data about the maximum possible income by source for each group of the $\bar{z}$ distribution. Note that the top-level income from entrepreneurial investment $\pi(a_{\max}, \bar{z}_4, \Lambda_{\text{f}})$ is comparable to almost two additional top-level salaries $\bar{w} \cdot \varepsilon_4$. Contrary to that, in the normal lane, maximum profits from entrepreneurial engagement

are only little above the lowest earnings level $\bar{w} \cdot \varepsilon_0$. The maximum possible return on capital is $r \cdot a_{\max} = 7.11$.

Table 7: Maximum possible return on investments π for the normal lane Λ_{n} and the fast lane Λ_{f} at the highest asset level $a = a_{\max}$ by ability $\bar{z}$.

	$\bar{z}_0$	$\bar{z}_1$	$\bar{z}_2$	$\bar{z}_3$	$\bar{z}_4$
$\pi(a_{\max}, \bar{z}, \Lambda_{\mathrm{n}})$	0.002	0.006	0.02	0.05	0.16
$\pi(a_{\max}, \bar{z}, \Lambda_{\mathrm{f}})$	0.00003	0.006	1.39	21.13	56.15

Table 8: Labor earnings for each labor productivity group ε.

	ε_0	ε_1	ε_2	ε_3	ε_4
$\bar{w} \cdot \varepsilon$	0.1	0.41	1.73	7.21	30.13

Driven by this insight, we ask which source of income is most important, depending on labor productivity, entrepreneurial ability and lane. Keep in mind that at the lowest asset level $a_{\min}$, agents solely rely on their labor earnings, since they do not own any capital to invest or lend. Table 9 depicts the most important source of income for each combination of labor productivity ε and ability $\bar{z}$ at the maximum asset level. The data emphasizes that the significance of entrepreneurial engagement vanishes for agents apart from the fast lane.

Table 9: Predominant income source for all ε and $\bar{z}$ groups at the highest asset level $a = a_{\max}$, where RoC = return on capital, RoI = return on (entrepreneurial) investment, L = labor income. Bracketed values denote the predominant source for the fast lane where they deviate from the other two lanes.

	$\bar{z}_0$	$\bar{z}_1$	$\bar{z}_2$	$\bar{z}_3$	$\bar{z}_4$
ε_0	RoC	RoC	RoC	RoC (RoI)	RoC (RoI)
ε_1	RoC	RoC	RoC	RoC (RoI)	RoC (RoI)
ε_2	RoC	RoC	RoC	RoC (RoI)	RoC (RoI)
ε_3	L	L	L	L (RoI)	L (RoI)
ε_4	L	L	L	L	L (RoI)

4.3.2 Retirement Lane

Agents in the retirement lane have an entrepreneurial productivity of zero by construction, regardless of their inherent ability $\bar{z}$. Since they can never go back to a higher lane, these agents differ from each other only in their current and expected labor income $\bar{w} \cdot \varepsilon$ with the corresponding levels shown by Table 8. Thus, the policies for agents with the same labor productivity are exactly the same. We will exploit this feature in Chapter 5 in order to reduce the number of grid points we need to evaluate the policies for. All retired agents face a substantially high mortality risk of 20% in each period. Consequentially, all retired

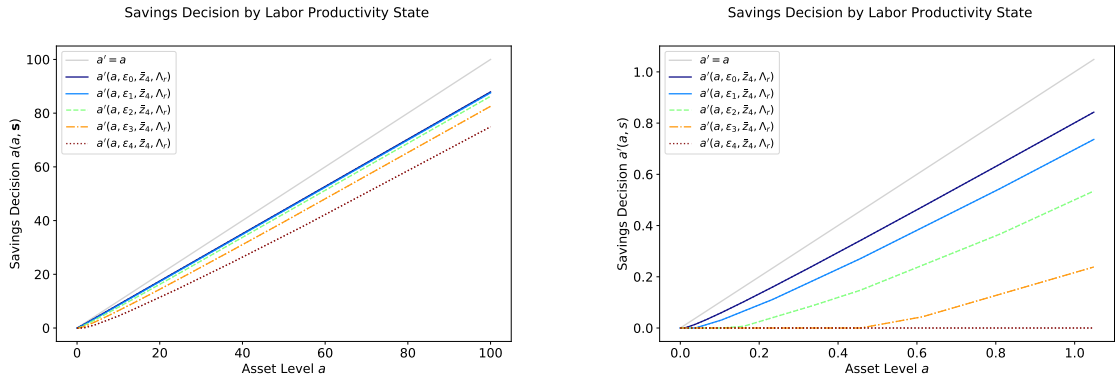

Figure 10: Savings policies for agents in the retirement lane, by labor productivity ε over the whole asset range (left) and zoomed in (right).

agents live on their savings in that they save less than would be necessary to keep their asset level constant.

Table 10: Percentage share of return in capital (RoC) on total income at the top asset level $a_{\max}$ by labor productivity for agents in the retirement lane.

	ε_0	ε_1	ε_2	ε_3	ε_4
Share of RoC [%]	98.6	94.5	81.5	49.7	19.1

In the retirement lane, agents rely on the two remaining sources of income: their labor earnings and the return on capital, which is determined by the interest rate r. Table 10 shows that for agents at the maximum asset level $a_{\max}$ in the lower three labor productivity groups, return on capital is the most important source of income. Thus, each unit of capital saved today provides a huge utility gain from consumption tomorrow. As a consequence, these agents are very likely to accumulate high savings. We identify two contrary effects. First, agents with high labor productivity gain a substantial share of income from labor and are thus less dependent on the interest payments, which reduces their propensity to accumulate savings. Second, the relative importance of labor income on total income decreases in the asset level for *all* agents, thus, at high asset levels, the differences in labor productivity should hardly be reflected in the savings behavior. However, Figure 10 shows that the former effect outweighs the latter. If hold ability $\bar{z}$ and lane Λ_{r} fixed and plot the different savings levels by labor productivity state ε, we observe a negative relationship between labor earnings and savings, and this difference grows in the asset level. Furthermore, zooming in reveals that the higher the labor productivity the higher the maximum asset level where the borrowing constraint is still binding.

Figure 11 shows the consumption policies for retired agents by labor productivity. We observe that the higher the labor productivity level the higher the consumption level, as agents with high labor productivity can simply afford to consume more (remember that the top before-tax salary is around 30, whereas the lowest is around 0.1). Furthermore, we can see the kink of the ε_4 consumption policy that corresponds to the kink of the

Figure 11: Consumption policy for agents in the retirement lane at all labor productivity levels.

savings policy. Up to this kink, the borrowing constraint for the ε_4 group is binding, thus, agents spend all their after-tax income on consumption (and would prefer to consume even more). The borrowing constraint has a continuing effect up to an asset level of roughly 25, which is reflected in a decreasing slope of the consumption policy.

Remember that we set the risk aversion parameter σ comparably low. This implies that the preference of agents to smooth their consumption over lifetime is rather small. Simultaneously, because the persistence ρ_ε of the AR1 process that generates the labor productivity is at the high side, the probability to remain in the same state is about 99.5%. Thus, agents with high labor productivity are well-advised to assume the same amount of earnings tomorrow as today and consequently rely less on savings to finance their consumption. If we set the risk aversion parameter σ to 4 or lower the persistence ρ_ε to 0.9 while holding all other circumstances fixed, all agents reduce their consumption levels as shown by the left and middle plot of Figure 12. Crucially, the borrowing constraint is no longer binding as the valuation of the future (for higher σ) and the expected value of the future (for lower ρ_ε) decrease, and the consumption policies are flattened. Note that these effects are solely due to changes in the decision rules, as the current income is neither affected by σ nor by ρ_ε.

The same happens if the mortality risk is reduced from 20% to 0%. Remember that the mortality risk effectively decreases the discount factor for agents in the retirement lane. We can see these mechanisms at play if we plot the consumption policies for a given labor productivity and ability over the three lanes in Figure 13. Retired agents spend substantially more on consumption than agents in the normal and even most agents in

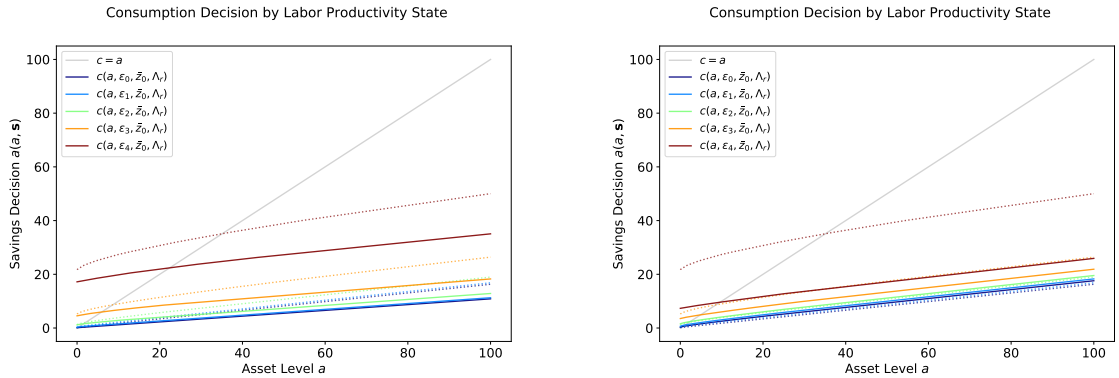

Figure 12: Consumption policy of retired agents for changed parameters $\sigma = 4$ (left plot) and $\rho_\varepsilon = 0.9$ (right plot), holding all other factors fixed at their equilibrium benchmark level. The dotted lines indicate the benchmark consumption policy.

the fast lane. If we eliminate the mortality risk, holding all other circumstances fixed, the consumption of retired agents drops to the normal lane level, as these two lanes almost face the same income. Note that the lower mortality risk also slightly affects the consumption policies of the normal and the fast lane but rather indirectly via the future prospects.

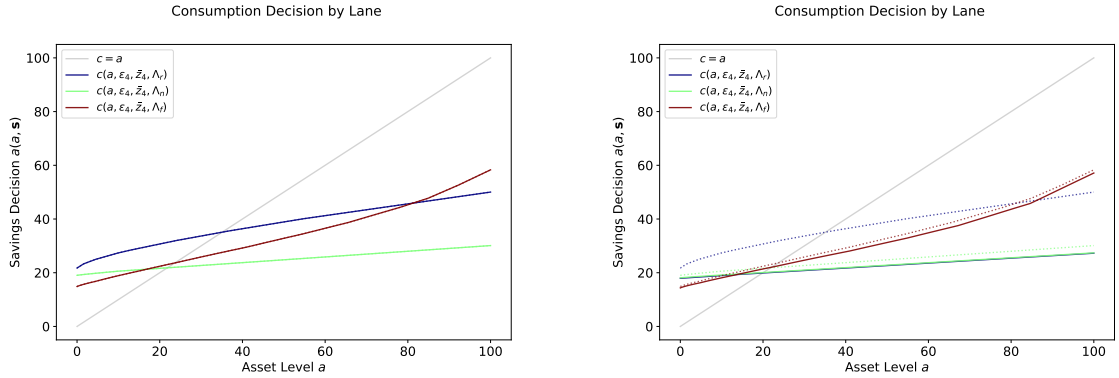

Figure 13: Consumption policy by lane with 20% mortality risk in the retirement lane (left plot) and 0% mortality risk (right plot). The dotted lines in the left plot show the benchmark policy.

4.3.3 Normal Lane

Due to the high interest rate, agents in the normal lane face only little or no incentives to invest much in their business. Figure 14 shows the optimal investment policies and the corresponding returns on investment for the normal lane. We observe that all agents but the ones with very low assets are restricted by the first entry of the minimum function (39). Consequently, for normal lane agents, returns on entrepreneurial investment account for only a small share of their total income. Even for agents with the highest ability $\bar{z}_4$, this share is no higher than 2% at maximum.

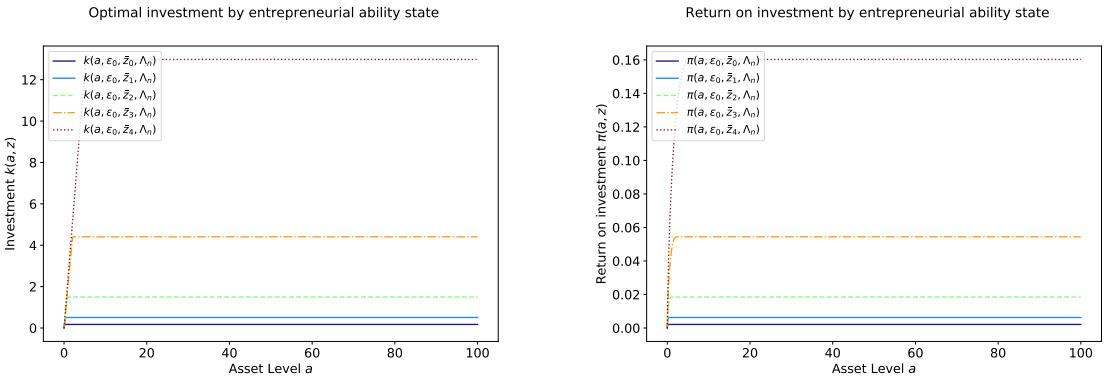

Figure 14: Optimal investment policy and corresponding return on investments for agents in the normal lane. Since optimal investment does not depend on labor earnings, the investment policies are the same for all ε states.

Figure 15: Optimal savings policy by labor productivity for agents in the normal lane.

Since the significance of entrepreneurial engagement almost vanishes for normal lane agents, the policy functions differ only in the labor productivity. In Figure 15, we plot the policies for fixed ability and lane. We observe that the savings policies for the different labor productivities approach and even cross each other with rising asset levels. This leads to the initial ordering almost being reversed at the highest asset state compared to the lowest, i.e., those agents with the lowest earnings levels tend to save the most in order to keep their consumption level high.

Crucially, we observe large differences in the savings behavior of normal lane agents compared to retirees as shown in Figure 10. The explanation for the reversed order is fairly similar, however, the big difference is the mortality risk we imposed on retirees but not on the other lanes. As discussed in Section 5.1.1, if we drop the mortality risk, the

savings and consumption policies for both lanes are almost the same, since entrepreneurial engagement if of no importance for any of the two groups.

4.3.4 Fast Lane

Being in the fast lane amplifies the differences in inherent ability $\bar{z}$, as agents in the bottom ability group with $\bar{z}_0 < 1$ are worse off when in the fast lane while agents with $\bar{z} > \bar{z}_{\text{median}}$ gain substantially. Agents in the second lowest ability group $\bar{z} = \bar{z}_1$ neither profit nor lose when in the fast lane. Consequently, their same savings policies for the normal and the fast lane are the same.

Figure 16 shows the optimal investment policy and the corresponding return on investment for fast lane agents. We observe that the two top ability groups are constantly constrained in their investment decision. Remarkably, we find that for the most productive group, returns on entrepreneurial investment amplify their total wealth by factor 1.5 in a single period.

As a result, profits from entrepreneurial business are the predominant source of income (up to 89%) for the two highest ability groups. Table 11 depicts the share of entrepreneurial profits in total income for entrepreneurs in the fast lane and all possible combinations of labor productivity ε and ability $\bar{z}$ at the top asset level $a_{\max}$. The substantial endowment of highly productive agents explains the higher savings levels of these groups compared to less gifted agents as shown by Figure 17.

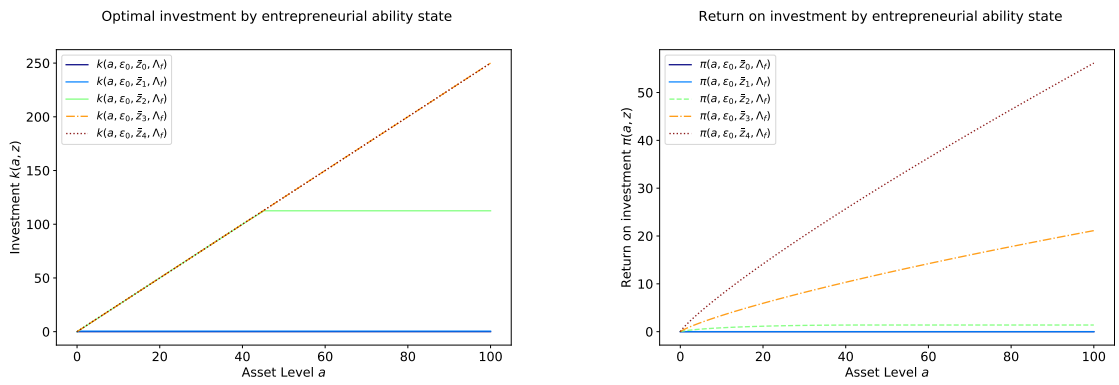

Figure 16: Optimal investment policy and corresponding return on investments for agents in the fast lane. Since optimal investment does not depend on labor earnings, the investment policies are the same for all ε states. The upper two ability groups are constantly constrained in their investments (thus, optimal investment grows linearly in the asset level), while the middle ability group $\bar{z}_2$ is constrained at least up to an asset level around 40.

However, we observe a kink in their policy functions, located in the upper quarter of the asset range, which cannot be explained solely by the level of endowment. As shown by Figure 18, the kink appears for the two top ability groups of entrepreneurs, regardless of their labor productivity. It indicates a threshold above which the slope of the policy

Table 11: Percentage share of fast lane entrepreneurs' profits in total income at the top asset level $a_{\max}$.

	$\bar{z}_0$	$\bar{z}_1$	$\bar{z}_2$	$\bar{z}_3$	$\bar{z}_4$
ε_0	0.0004	0.09	16.15	74.55	88.62
ε_1	0.0004	0.08	15.58	73.73	88.18
ε_2	0.0003	0.07	13.58	70.5	86.4
ε_3	0.0002	0.044	8.84	59.59	79.67
ε_4	0.0008	0.02	3.6	36.2	60.12

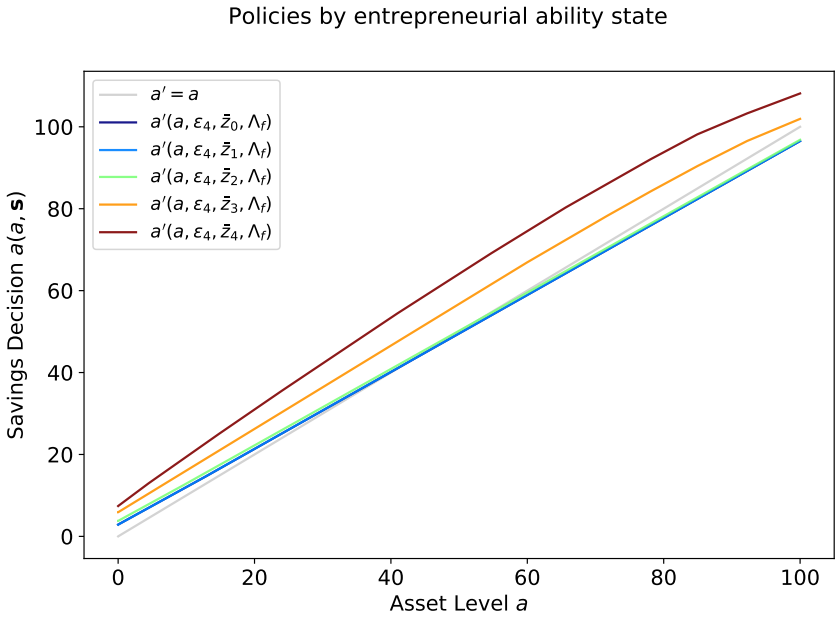

Figure 17: Savings policy functions for fast lane agents at the highest labor productivity state ε_4 over all entrepreneurial ability states.

Figure 18: Savings policy functions for fast lane agents at the highest ability state over all labor productivity states.

function is smaller than before. The most most productive fast lane groups differ from the others in two aspects: First, their current income from entrepreneurial engagement increases in the asset level due to the constantly binding constraint ϑ on investments.

Second, just because of this constraint, these agents gain an additional reward on each additional unit of savings, since the first derivative of the investment function from (39) is nonzero. Due to the collateral constraint, highly productive agents cannot invest as much in their businesses as would be optimal. Constrained agents gain very much from each additional unit of assets tomorrow as their optimal investment grows linearly in their asset level. Precisely, each unit of consumption let go today implies 2.5 additional units of investment tomorrow, due to the choice of the collateral constraint ϑ. Remember that the Euler equation (30) includes a reward on each additional unit of savings expressed by:

$$\frac{\partial \omega(a', z'; \tau_{\text{cap}})}{\partial a'} = 1 + \left(r + \frac{\partial \pi(a', z')}{\partial a'} \right) \cdot (1 - \tau_{\text{k}}), \quad \text{if} \quad \tau_{\text{cap}} = \tau_{\text{k}}$$

For the most productive entrepreneurs, the derivative of the investment choice function $\partial \pi(a', z')/\partial a'$ is 3.33 at maximum, which is almost 50 times higher than the 0.07 reward gained from the interest rate. With growing asset levels, the derivative approaches 0.5, which is still about seven times higher than the interest rate.[5] For all non-constrained agents, the derivative is zero. Crucially, because prices are assumed to be constant in equi-

[5] We left out $a_0 = a_{\min} = 0$ as $\partial \pi(a', z'; \tau_{\text{cap}})/\partial a'$ approaches ∞ for $a \to 0$.

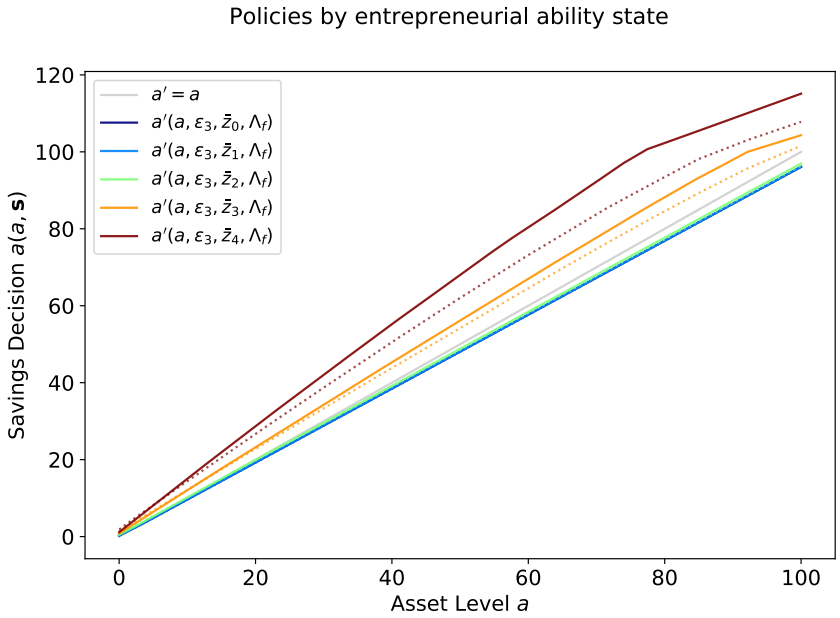

Figure 19: Savings policy functions by ability state for agents in the fast lane with labor productivity ε_3. If the probability to descend to the fast lane p_{normal} is increased from 0.05 to 0.5, the savings levels rise and the kinks of the policy functions are shifted to lower asset states.

librium and each agent's productivity can only deteriorate, agents who are unconstrained today will be unconstrained tomorrow.

The reward on additional savings partly explains our observation related to Figure 9, where the distance in savings between the fast and the normal lane grows up to the location of the kink and is continually reduced afterwards. However, this is only the case if agents remain in the fast lane in the next period, and the corresponding probability is 92% due to our choice of the probabilities to descend $p_{\text{normal}} = 0.05$ and $p_{\text{retire}} = 0.03$. Thus, if we increase the risk of descending to the normal lane to $p_{\text{normal}} = 0.5$, holding all other factors fixed, we expect the kink to be shifted towards lower asset states. This is confirmed by Figure 19, where we increased p_{normal} to 0.5, holding all other factors fixed at their benchmark equilibrium level. Note that, again, all income sources remain the same, thus, our observation is solely due to a change in the decision rules.

Besides the risk of descending to lower lanes, the choice of the collateral constraint has a huge impact on the model, as mentioned in Section 4.1. If we relax the collateral constraint ϑ while keeping all prices fixed at their benchmark equilibrium level, we observe that the kink is more pronounced and moves towards lower asset states. The results are shown by Figure 20. In the most extreme case, we disable the collateral constraint completely by setting $\vartheta = \infty$. This allows highly productive agents with low assets to borrow heavily and gain from large profits. The effect of the higher income implies a jump

Table 12: Macroeconomic moments from the U.S. data compared to the results found by Guvenen et al. (2019), our benchmark model and various sensitivity checks. In order to establish the same basis of comparison, the U.S. data is also taken from Guvenen et al. (2019).

	U.S. Data	Guvenen Model	Benchmark Model	$p_{\text{fl}} = 0.5$	$\beta = 0.9475$ $\delta = 0.05$	$\eta = 0.05$	Identical $\eta > 0$	$\sigma = 4$
Capital-Output-Ratio	3	3	5.96	4.6	6.6	5.5	4.49	8.4
Top 10% Wealth Share [%]	75	66	83	78	75	83	90	53
Top 1% Wealth Share [%]	36	36	25	18	24	33	33	21
TTR / GDP [%]	29.5	25	18	20	17	19	19	15
CITR / GDP [%]	28	25	22	20	21	21	15	30
Bequest / Wealth [%]	1.29	0.99	1.35	1.31	1.31	3.77	5.81	1.89

of the savings policy from the lowest to the second lowest asset state.[6] However, since no agent is constrained in their investments, the first derivative of the profit function with respect to the asset level drops to zero for all agents. Consequently, the kink is eliminated.

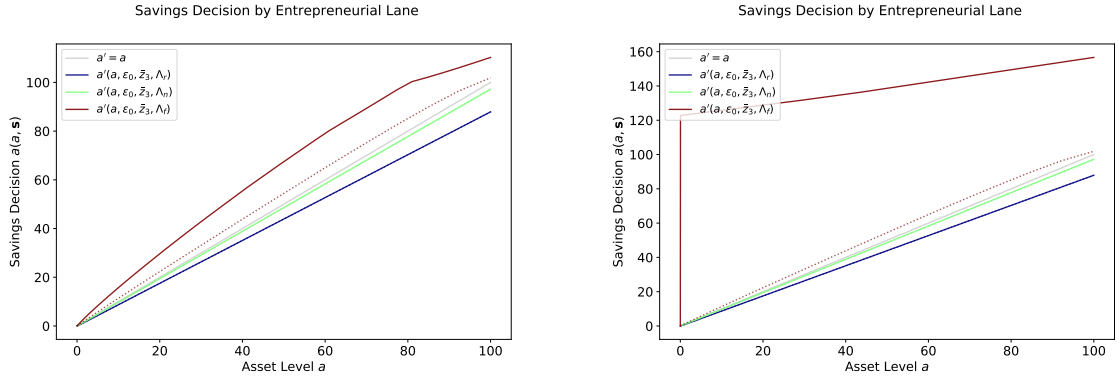

Figure 20: Savings policy for fixed lowest labor productivity state ε_0 and the second highest ability state $\bar{z}_3$ over all three lanes. The left plot depicts the policy if ϑ is raised by factor 50 compared to the benchmark value. The right plot depicts the policy for $\vartheta = \infty$.

4.4 Sensitivity Analysis

We conduct sensitivity checks varying selected economic parameters. The corresponding macroeconomic outcomes are provided by Table 12. Apparently, changing only one or two parameters can have a substantial impact on the resulting equlibria. In contrast to the last section, where we analyzed few changes in the parametrization holding all other circumstances fixed, we now turn to the comparison of the final equilibria.

Lowering the Share of Fast Lane Agents We cut the share of agents born into the fast lane from 1.0 to 0.5. This does not have any direct effect on the policy functions

[6] Due to the design of the bond market that requires participants to own at least one infinitesimal small unit of assets, agents at the lowest asset state $a = 0$ do not have access to the bond market.

but solely on the distribution of agents and thus on the macroeconomic outcomes. The measure of retirees remains the same, however, the measure of fast lane agents is cut in half by 16 percentage points and these accrue to the normal lane, which now makes up the bulk of the population with measure 71%.

Since there are fewer highly productive agents, the quality-adjusted capital stock is cut by one third compared to the benchmark level, which allows the price on input goods to scale up by 25%. As a result, returns on investment rise for all active entrepreneurs. This drives up the incentives to accumulate savings, especially for agents in the fast lane.

Looking at the top 10% and top 1% wealth share provided by Table 13, we observe that this configuration generates slightly less inequality compared to the benchmark configuration.

Table 13: Macroeconomic moments from the U.S. data in comparison to the results of the model by Guvenen et al. (2019), the benchmark model and a lower share of fastlane agents.

	U.S. Data	Guvenen Model	Benchmark Model	$p_{\mathfrak{fl}} = 0.5$
Capital-Output-Ratio	3	3	5.96	4.6
Top 10% Wealth Share [%]	75	66	83	78
Top 1% Wealth Share [%]	36	36	25	18
TTR / GDP [%]	29.5	25	18	20
CITR / GDP [%]	28	25	22	20
Bequest / Wealth [%]	1.29	0.99	1.35	1.31

Approaching the Guvenen Configuration We increase the discount factor for future utility β from 0.92 to 0.9475 and the capital depreciation rate δ from 0.04 to 0.05 in order to match the parametrization by Guvenen et al. (2019). The configuration manages to exactly match the empirical top 10% wealth share as shown by Table 14. The impact on the other targeted parameters compared to the benchmark configuration is rather moderate, except for a much higher capital-output-ratio.

The equilibrium savings and investment levels increase by 37% and the quality-adjusted capital stock rises by 18,6%. Consequently, the interest rate falls by 41% to 4.2%. While the earnings and the income distribution remain almost constant, the share of the top wealth quintile decreases by 4 percentage points, as shown by Figure 21. The policy functions stay roughly the same.

Reducing the Mortality Risk We lower the mortality risk of retired agents from 20% to 5%. This is far nearer to the OLG structure by Guvenen et al. (2019), since it implies an average duration of 12.5 in the fast, 18 years in the normal and 20 years in the retirement lane. Consequently, the measures of lanes in the stationary distribution

Table 14: Macroeconomic moments from the U.S. data in comparison to the results of the model by Guvenen et al. (2019), the benchmark model and larger discount factor β and capital depreciation rate δ.

	U.S. Data	Guvenen Model	Benchmark Model	$\beta = 0.9475$ $\delta = 0.05$
Capital-Output-Ratio	3	3	5.96	6.6
Top 10% Wealth Share [%]	75	66	83	75
Top 1% Wealth Share [%]	36	36	25	24
TTR / GDP [%]	29.5	25	18	17
CITR / GDP [%]	28	25	22	21
Bequest / Wealth [%]	1.29	0.99	1.35	1.31

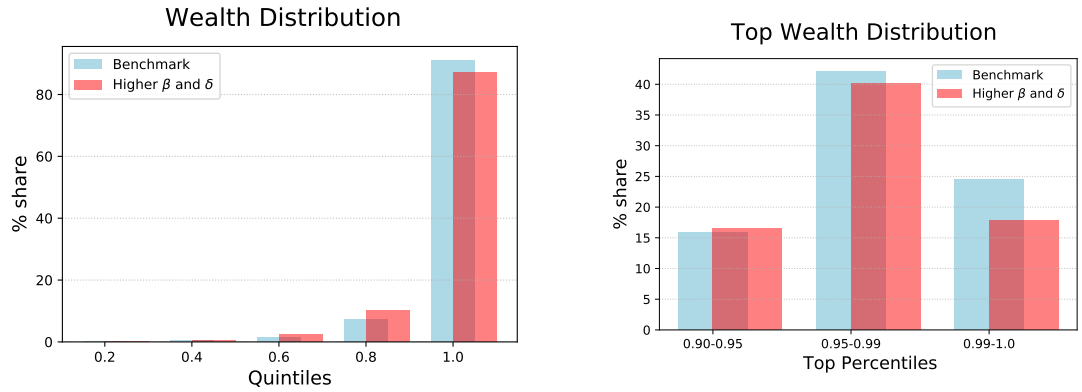

Figure 21: Wealth Distribution for increased discount factor β and capital depreciation rate δ.

change from 13% to 38% for the retirement lane, 54% to 39% for the normal lane and 33% to 23% for the fast lane. The effective mortality risk is lowered from 2.61% to 1.87%. As a result, we get 30% less agents at the lowest asset state. The share of savings held by retirees jumps by 11 percentage points from 6.84% to 18.9%.

Table 15: Macroeconomic moments from the U.S. data in comparison to the results of the model by Guvenen et al. (2019), the benchmark model and lower mortality risk η.

	U.S. Data	Guvenen Model	Benchmark Model	$\eta = 0.05$
Capital-Output-Ratio	3	3	5.96	5.5
Top 10% Wealth Share [%]	75	66	83	83
Top 1% Wealth Share [%]	36	36	25	33
TTR / GDP [%]	29.5	25	18	19
CITR / GDP [%]	28	25	22	21
Bequest / Wealth [%]	1.29	0.99	1.35	3.77

The interest rate is lower by 7% and the quality-adjusted capital stock is lower by 13% compared to the benchmark level. However, there is almost no impact on the distribution of earnings, income and wealth, except the share of total wealth for the upper quintile

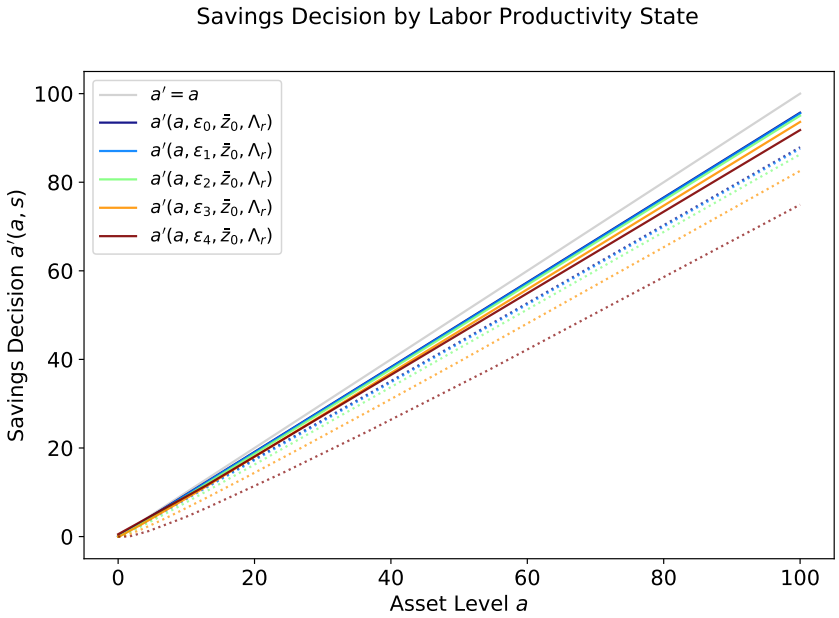

Figure 22: Savings policy functions for the retirement lane. Dotted lines indicate the respective policies for the benchmark configuration.

rises by 1.5 percentage points to 92.5%. This corresponds with a slightly higher Gini coefficient of 0.765 compared to 0.755 before.

The effect on the policy functions for retired agents is remarkable. Figure 22 shows that savings levels are much higher and the differences by labor productivity vanish. Furthermore, no retired agent is affected by the borrowing constraint. The corresponding consumption policy functions are almost linear and much lower than at the benchmark level as shown by Figure 12 in Section 4.3.2.

Making the Mortality Risk Independent of Lane We impose a mortality risk on all agents, regardless of the lane, and set it to 2.6% in order to obtain the same effective mortality risk as the benchmark configuration. This dramatically reduces the share of active entrepreneurs in the stationary distribution, as the measure of retired agents jumps from 13% to 54% and the measures of normal and fast lane agents are roughly cut in half. As a result, aggregate input goods fall by one third, which allows for higher prices on input goods and thus higher returns on investments for all active entrepreneurs.

For retirees, whose mortality risk is now 2.6% compared to 20% before, the effective discount on future consumption decreases. In response, they accumulate more savings, increasing their share in aggregate wealth from 6.84% to 29%. Note that this capital is idle in that it is not invested in production.

The effect on the wealth distribution is large. The share of the upper quintile is raised by almost 7 percentage points to 98%, which is almost 11 percentage points above the empirical value. The top 1% wealth share is 33% which is fairly near to the 36% share

Table 16: Macroeconomic moments from the U.S. data in comparison to the results of the model by Guvenen et al. (2019), the benchmark model and identical mortality risk $\eta = 0.026$ imposed on all agents.

	U.S. Data	Guvenen Model	Benchmark Model	Identical $\eta = 0.026$
Capital-Output-Ratio	3	3	5.96	4.49
Top 10% Wealth Share [%]	75	66	83	90
Top 1% Wealth Share [%]	36	36	25	33
TTR / GDP [%]	29.5	25	18	19
CITR / GDP [%]	28	25	22	15
Bequest / Wealth [%]	1.29	0.99	1.35	5.81

targeted by Guvenen et al. (2019). Looking at the PDF and CDF over assets presented by Figure 23 reveals that in this case, the high top 1% wealth share is indeed generated by the model and not due to grid exceedance. However, the model overstates the 90-95% and the 95-99% percentiles as shown by Figure 24.

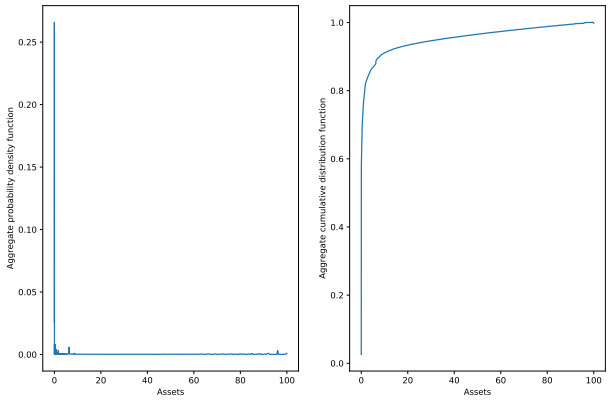

Figure 23: PDF (right) and CDF (left) of the configuration with lane-independent mortality risk.

Increasing the Risk Aversion Finally, increasing the risk aversion σ to 4 as in Guvenen et al. (2019) inflates the entrepreneurial business and yields *less* inequality, as shown by Table 17. Savings and investments almost triple compared to the benchmark configuration and the quality-adjusted capital stock is inflated to almost 180% of its benchmark value. As a result, the interest rate is cut by 75% compared to the benchmark configuration and the price on input goods falls by 25%.

The resulting distrubution over earnings, income and wealth is far from the empirical data as shown by Figure 25. The top 10% wealth share is 10 percentage points lower than the targeted empirical share and 14 percentage points lower than the benchmark model.

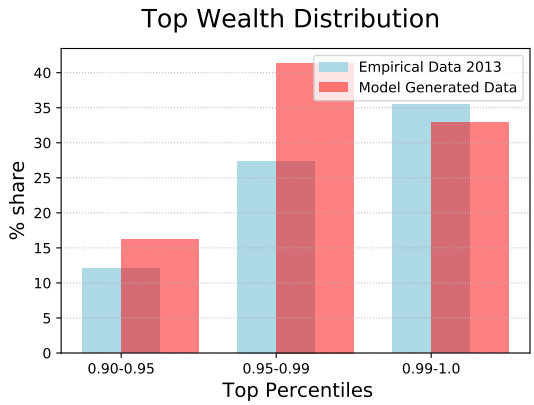

Figure 24: Top 10% percentiles for the wealth distribution where all agents face the same mortality risk of 2.6% in comparison to the empirical data.

Table 17: Macroeconomic moments from the U.S. data in comparison to the results of the model by Guvenen et al. (2019), the benchmark model and higher risk aversion σ.

	U.S. Data	Guvenen Model	Benchmark Model	$\sigma = 4$
Capital-Output-Ratio	3	3	5.96	8.4
Top 10% Wealth Share [%]	75	66	83	53
Top 1% Wealth Share [%]	36	36	25	21
TTR / GDP [%]	29.5	25	18	15
CITR / GDP [%]	28	25	22	30
Bequest / Wealth [%]	1.29	0.99	1.35	1.89

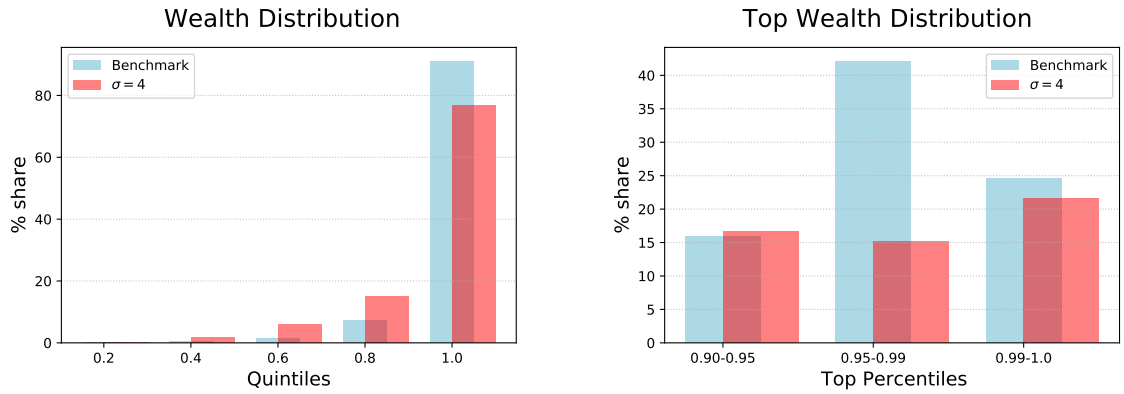

Figure 25: Wealth Distribution for higher risk aversion $\sigma = 4$.

5 Improving the Efficiency of the Solution

As we have seen in the last chapter, we need to run the model over and over in order to find a choice of parameters that yields a realistic equilibrium. Thus, pushing down the runtime to a computationally feasible minimum is crucial. This applies to our model introduced in Chapter 2, but even more to future extensions of the model which are likely to further drive up the computational effort.

We start by comparing the complete Guvenen model including mortality risk, entrepreneurial productivity process, taxes and exogenous labor supply to the simple Aiyagari model from the beginning of Chapter 2. In order to establish comparability between the two models, we inflate the number of exogenous states of the latter to 75. The additional complexity from the entrepreneur process drives up the runtime to about 55 minutes, which is roughly four times higher than for the simple Aiyagari model.

Note that we cannot perform isolated runtime experiments holding all circumstances fixed, thus, the observed runtime is subject to natural variations. We collect the runtime data on a private notebook with an Intel(R)Core(TM) i7-10510U CDP @ 1.8GHz 2.3GHz processor and 16 GB RAM. Due to the low validity of the runtime data, we focus on the number of computing-intensive steps in the following, which we can control and reproduce.

It is obvious, however, that we need to dramatically reduce the runtime of the complex model. Remember that we need to perform several repetitions of the TI as the guesses for the macroeconomic variables adjust (see Algorithm 1, p. 19). Thus, assuming a number of 30 iterations of Algorithm 1, an expected runtime of 55 minutes per iteration yields a total time effort of 27 hours until convergence every time we change a single parameter, which is highly unfeasible.

We use line profiling tools in order to identify the policy determination and especially the root finding (3b of Algorithm 2, p. 20) as the most time-consuming step. Note that one iteration of Algorithm 1 consists of one execution of the TI algorithm and one execution of the CDF Iteration algorithm (Algorithm 3, p. 23). Of the 55 minutes total runtime, only half a minute accounts for the CDF Iteration.

We find that one subiteration (assuming a savings policy for tomorrow and calculating the corresponding root of the Euler equation) of the TI algorithm needs roughly 63.9 seconds on average. These 63.9 seconds are solely due to the root finding steps (we need to find the root for each of the 3825 grid points). However, the root finding is executed by an external solver and thus beyond our control. Therefore, our best chance is to reduce the number of grid points we need to calculate the root for and interpolate in between. We use the total number of root finding steps in order to assess the efficiency improvements in the following.

In this chapter, we focus on similarities of the policy functions along the asset range. We start by exploiting our preliminary knowledge about the policy functions gained from

Chapter 4 in Section 5.1. Afterwards, in Section 5.2, we develop an approach that does not make use of a priori knowledge but is dynamically tailored to the model and its configuration. In order to do this, we need to introduce a metric to approximate the deviation of the interpolated solution from the optimal solution: the *Euler error*. Section 5.3 discusses our choice of a suitable tolerance criterion. Finally, we evaluate our results in Section 5.4. We give a brief outlook on exploiting possible similarities in other dimensions of the state space at the end of this chapter in Section 5.5.

5.1 Using A Priori Knowledge to Reduce the Number of Grid Points

As we have seen in Chapter 4, for each combination of exogenous states held fixed, the policy functions are monotonically increasing in the asset level and look almost linear in most parts of the asset space. We identified two main areas of interest: First, the lowest asset levels where households have very small current income levels but are likely to be better off in the future. These households prefer to make negative savings, i.e., they would like to borrow additional capital to smooth their lifetime consumption, but are restricted by the borrowing constraint.

Second, policy functions for agents in the fast lane with high levels of entrepreneurial ability show a kink in the upper half of the asset range. This indicates a threshold above which agents decide to save less and consume more. We have already discussed possible explanations for the appearance of the kink in Chapter 4. In this section, we use this a priori knowledge to reduce the number of grid points beforehand.

5.1.1 Reduce Grid Points in the Retirement Lane

Once they end up in the retirement lane, all entrepreneurs, regardless of their inherent ability $\bar{z}$, face the same entrepreneurial productivity $z = 0$. None of these households will ever engage in entrepreneurial business again until they die. Thus, for each labor productivity state ε, current income as well as the transition probability for all $n_{\bar{z}}$ ability states are exactly the same. By calculating the policy only once for each labor productivity state ε, we reduce the number of retirement lane states from 25 to 5 and the total number of grid points from $n_{\mathrm{a}} \cdot (25 + 25 + 25) = 3825$ to $n_{\mathrm{a}} \cdot (25 + 25 + 5) = 2805$.

We save our choice of grid points in the sparse matrix $\mathcal{M} \in \{0, 1\}^{n_a \times n_s}$, where n_a is the number of asset states and $n_s = n_\varepsilon \cdot n_{\bar{z}} \cdot n_\Lambda$ is the number of exogenous state tuples. This matrix has the same shape as the policy grid and marks the grid points where the policy was calculated exactly. Each row denotes an asset state, while each column denotes a distinct combination of exogenous states. The binary-coded matrix $\mathcal{M}$ for the reduced number of retirement states is shown by Figure 26.

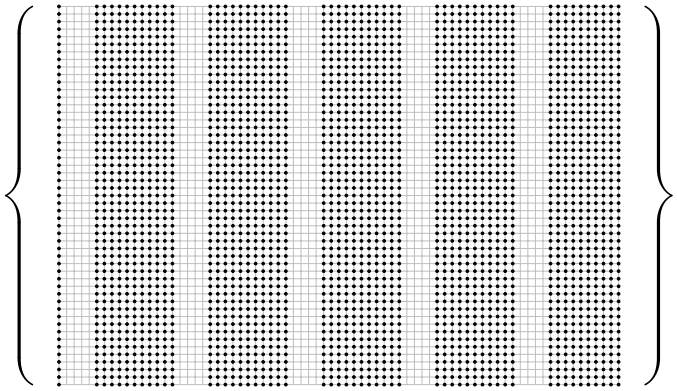

Figure 26: Binary-coded sparse matrix $\mathcal{M}$ with reduced number of retirement states.

We calculate the policy for each marked entry of $\mathcal{M}$. Afterwards, for each labor productivity state ε, we copy the obtained policy to all other columns with the same ε but different ability $\bar{z}$. Note that this kind of redundancy arises from having vectorized the three-dimensional exogenous state space $\mathbf{E} \times \bar{\mathbf{Z}} \times \mathbf{L}$ to the quasi one-dimensional space $\mathbf{S}$ which is associated with building the cross product of the three substates. It would be sufficient to define only one accumulative retirement state per labor productivity state in the first place.

5.1.2 Locate the Second Kink

Having reduced the number of marked columns in $\mathcal{M}$, we make use of our insights from Chapter 4 in order to reduce the number of marked grid points per column. Since each row of $\mathcal{M}$ corresponds to an asset state a, we exploit the almost linear behavior of the policy functions for agents with low entrepreneurial skills. Our naive approach to reduce the number of grid points along the asset range is to assume that the second kink appears only for non-retired entrepreneurs with ability $\bar{z}$ equal to or above the median $\bar{z}_{\text{median}} = 1$. Note that these are precisely the entrepreneurs who are not worse off when in the fast lane. In contrast, entrepreneurs with $\bar{z} < \bar{z}_{\text{median}}$ have an even smaller productivity when in the fast lane compared to the normal lane, since the amplifier λ is greater than 1. Moreover, in the retirement lane, agents drop out of entrepreneurial business completely, thus, we do not assume any second kinks here.

We assume sufficiently linear policies for $\bar{z}_0$ as well as for all retirement states and content ourselves with simply interpolating between the lowest and the highest asset state. For the remaining ability states in the normal and fast lane, we calculate the policy for every second asset state exactly and interpolate in between, which implies $\lceil n_a/2 \rceil$ grid points per column. The corresponding sparse matrix $\mathcal{M}$ can be fully initialized beforehand and is shown by Figure 27.

Table 18 provides the total number of grid points for each variant discussed above. Using a priori knowledge about the policy functions saves us about 72% of all root finding

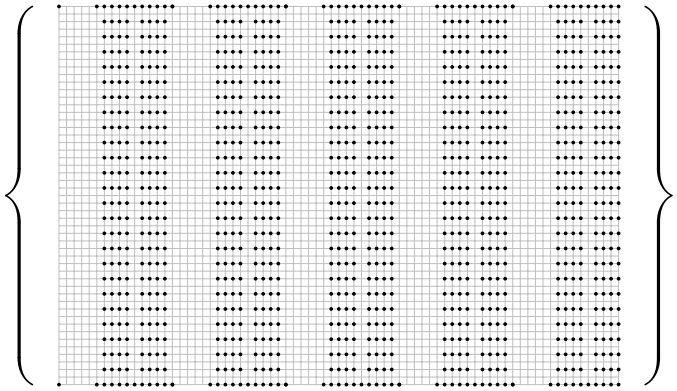

Figure 27: Sparse matrix $\mathcal{M}$ assuming linear policies for lower entrepreneurial productivity states.

Table 18: Total number of grid points for each initialization of $\mathcal{M}$.

Ability States	all		$\bar{z}_0$		all		Total number of grid points
Lanes	Λ_r		$\Lambda_\mathrm{n}, \Lambda_\mathrm{f}$		$\Lambda_\mathrm{n}, \Lambda_\mathrm{f}$		
	nrow	ncol	nrow	ncol	nrow	ncol	
Unadjusted complete model	51	25	51	10	51	40	3825
Reduce columns (Figure 26)	51	5	51	10	51	40	2805
Reduce rows (Figure 27)	2	5	2	10	$\lceil 25/2 \rceil$	40	1070

steps. Expressed in total runtime of the subiterations, reducing the number of columns with the corresponding initialization of $\mathcal{M}$ shown by Figure 26 decreases the runtime from 64 seconds to 39 seconds. If we additionally reduce the number of rows (Figure 27), the runtime drops to 9 seconds on average.

However, there are some drawbacks: First, calculating only every second grid point for those states which presumably show a kink implies a loss of precision. It is also inefficient in that we do not isolate the kink but instead search over the whole asset range.

Second, we used our limited knowledge derived from few iterations of the TI algorithm to assume a relationship between low entrepreneurial productivity and linear policies. How do we ensure these policies never show the second kink, subject to changing macroeconomic circumstances such as interest rate and prices?

Third, the efficiency gain from the reduction of grid points heavily depends on the configuration of the model, especially the entrepreneurial ability $\bar{z}$. If we change the parameters of the entrepreneurial ability process or add additional states above the median, this approach does not bring us any advantage.

5.2 Using a Bisection-based Approach to Determine the Grid Points of Interest

Having discussed the several disadvantages of using preliminary knowledge and statically determining the choice of columns in advance, we now focus on reducing the grid points along the asset range. This implies a reduction of marked entries per column of $\mathcal{M}$. We introduce a flexible approach that is not based on a priori knowledge but dynamically chooses grid points of interest. In the easiest case, where the policy function for a given combination of exogenous states is completely linear, we need to calculate only the first and the last grid point of each column and obtain a sufficient result by interpolating in between. Therefore, all approaches we introduce in the following start by initializing $\mathcal{M}$ with ones in the first and the last row as shown in Figure 28.

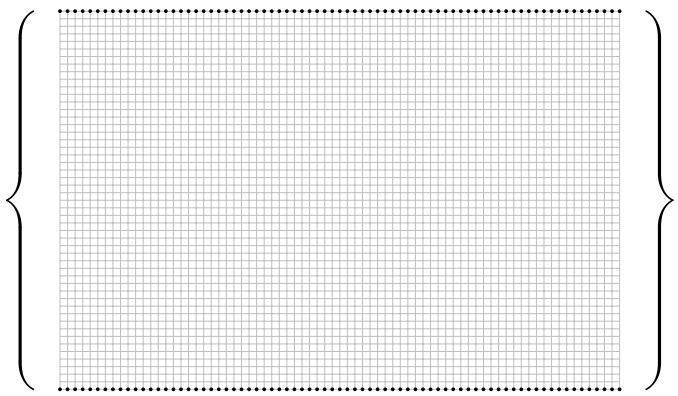

Figure 28: Initialization of $\mathcal{M}$. Black dots indicate ones.

We use a bisection-based approach in order to dynamically insert data points as long as the interpolation error is too large. For each combination of exogenous states (i.e., each column in $\mathcal{M}$), we calculate the root for the lowest and the highest asset state and interpolate in between. Subsequently, we check for each asset level whether the interpolated policy value is close enough to the real solution. We can quantify the interpolation error using the concept of Euler errors we introduce below in order to determine the grid points of interest. If the Euler error for a given asset level is above a chosen tolerance value, we calculate the root for this point exactly and interpolate again between all points we have already observed. The order in which we observe the grid points is taken from the bisection algorithm, i.e., we look at the middle asset level first, then divide the interval in two, look at the resulting two middle points again and continue until we have evaluated all points. Figure 29 provides an example for the dynamic choice of grid points. Note that unlike the example suggests, we neither know the true policy function nor the exact deviations e_{A} and e_{B}.

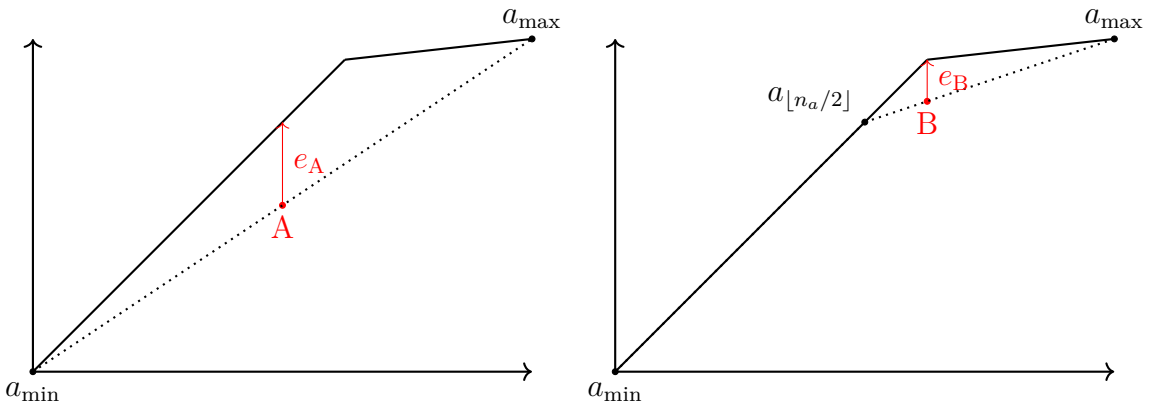

Figure 29: Determination of the piecewise linear policy function with initial data points $a_{\min}$ and $a_{\max}$. The solid line shows the true functional form (which is unknown), while the dotted line shows the interpolated function. Point $\mathrm{A} = a_{\lfloor n_a/2 \rfloor}$ is evaluated first. Since the deviation e_{A} is above the tolerance value, A is calculated exactly and marked in $\mathcal{M}$. We interpolate again over the data points $a_{\min}$, A and $a_{\max}$. In the proceeding of the algorithm, each asset state to the left of A will be accepted, while the asset states to the right need further adjustment. When point B is evaluated, the deviation e_{B} from the true value is beyond the tolerance criterion. The algorithm will calculate the root for B exactly and mark B in $\mathcal{M}$.

5.2.1 The Euler Error

For each grid point, we calculate the Euler error in order to determine the deviation of the interpolated policy from the exact policy we would obtain by root finding. For a given asset level a and state $\mathbf{s}$, the Euler error yields the percentual deviation of the estimated consumption policy from the optimal one. To compute the Euler error, we use the Euler equation (30) from Section 2. First, we calculate the consumption choice $\hat{c}$ today implied by the savings level a' from the budget constraint (17):

$$\hat{c} = \frac{1}{1+\tau_c}\Big(\omega(a, z; \tau_{\text{cap}}) + \varepsilon \cdot \bar{w} \cdot \ell \cdot (1-\tau_\ell) - a'\Big). \tag{40}$$

Likewise, we calculate tomorrow's approximate consumption $\hat{c}'$ for tomorrow's assets a' and all exogenous states $\mathbf{s}'$, using the same function we assume for a' to derive the expected policy a'' for the period after tomorrow. We then use the Euler equation to obtain the optimal consumption choice c^* given the approximate consumption choice $\hat{c}'$ tomorrow:

$$\begin{aligned} \frac{\partial u}{\partial c^*} &= \beta \cdot (1-\eta_\Lambda) \cdot \mathbb{E}\left[\frac{\partial u}{\partial \hat{c}'} \cdot \frac{\partial \omega(a', z'; \tau_{\text{cap}})}{\partial a'}\right] + \mu \\ \Leftrightarrow c^* &= u_c^{-1}\left(\beta \cdot (1-\eta_\Lambda) \cdot \mathbb{E}\left[\frac{\partial u}{\partial \hat{c}'} \cdot \frac{\partial \omega(a', z'; \tau_{\text{cap}})}{\partial a'}\right] + \mu\right). \end{aligned} \tag{41}$$

How do we obtain the Lagrangian multiplier μ? In the full optimization problem (26), μ is nonzero only if the borrowing constraint is binding, i.e., $a' \geq -b$. This is the case for low-income households with low productivity. For these households, the multiplier is the residual of the Euler equation when solved for $a' = 0$ instead of the optimal value $a' < 0$:

$$\mu(a' = 0) = \frac{\partial u}{\partial c} - \beta \cdot (1 - \eta_\Lambda) \cdot \mathbb{E}\left[\frac{\partial u}{\partial c'} \cdot \lim_{a' \to 0} \frac{\partial \omega(a', z'; \tau_{\text{cap}})}{\partial a'}\right]. \tag{42}$$

For constrained households, the optimal savings decision is negative, because the marginal utility of consumption today is greater than the expected discounted marginal utility of consumption tomorrow weighted by the return on saving an additional unit. This is precisely the case where the residual of (42) is positive. For all non-constrained households, the residual is zero.

We are finally equipped to calculate the Euler error as the fraction of consumption $\hat{c}$ implied by the policy from (40) and the optimal consumption c^* from (41):

$$e = \left| \frac{\hat{c}}{c^*} - 1 \right| \tag{43}$$

Note that we need the multiplier μ from (42) to "cancel out" the fact that we manipulate the policy for the borrowing-constrained households. If we calculate the policy for the borrowing-constrained households exactly, we do not make any interpolation error. However, since we manually set the policy to zero, the residual of the Euler equation for $a' = 0$ is nonzero which would yield a positive Euler error. By adding μ to (41), we ensure that the Euler error solely reflects the interpolation error made and is not affected by the manual adjustment at the borrowing constraint.

However, if we do not calculate the policy for the borrowing-constrained households exactly, we cannot disentangle the effect of the interpolation error from the effect of the binding constraint. In this case, the multiplier μ incorporates both influences, implying an Euler error of zero for *all* grid points where the constraint is binding, whether we calculated them exactly or not.

For this reason, we distinguish between the *interpolation error*, which reflects solely the error due to non-exact policy determination and determines our choice of grid points throughout the calculation, and the *a posteriori Euler error* we calculate after the TI has converged for the resulting policy. Why do we need to distinguish between these two concepts?

First, as described above, the a posteriori Euler error fails to properly reflect the interpolation error made at the very low asset states where the borrowing constraint is binding. However, we can estimate the interpolation error before we manually adjust negative pol-

icy values. This allows us to decide whether the interpolation error is reasonably small or if we should make additional efforts to determine the location of the borrowing constraint exactly. We will go into the details in Section 5.4.5.

Second, our algorithm is designed such that it ensures an upper bound on the Euler error over all points of the original grid. However, remember that the grid over assets $\mathbf{A}$ we defined in Section 4.1 is itself only a discrete approximation of the continuous asset space $[a_{\min}, a_{\max}]$. This implies, that we rely on interpolation if we want to evaluate the policy function at any asset level off the grid. As a result, the Euler errors between the grid points might far exceed the maximum Euler error our algorithm can ensure. We choose a finer grid for the a posteriori calculation of the Euler errors with $n_{a_f} = 1001$ and an even higher concentration of grid points at the bottom asset levels.

5.2.2 Dynamically Determine the Sequence of Grid Points

The Euler error itself is of course only an approximation of the deviation from the real value as it depends crucially on the guessed policy $\hat{a}'$ and $\hat{c}'$ respectively. Nevertheless, we assume that it serves as an approximation for the real deviation (which implies a monotonic relationship between the Euler error and the real deviation). We can exploit this feature in order to make the choice of grid points even more efficient than the bisection-based approach does. The exact procedure is described by Algorithm 5. Precisely, we let go of the fixed order and evaluate the grid points with the highest Euler error first. If this is already below the tolerance value, we are done. Otherwise, we calculate the root for this grid point exactly, interpolate again and recalculate the associated Euler errors. With the added information, we expect the maximum Euler error to be smaller than before. Moreover, the Euler error for the grid points we already calculated exactly is zero. We repeat until the maximum Euler error falls below the tolerance criterion.

Figure 30 illustrates the development of the Euler error for a given state $\mathbf{s}$ over the asset grid. With only six grid points calculated exactly, the maximum Euler error is still at 0.06 which is larger than the allowed tolerance of 0.05. When the grid point with the currently highest Euler error is calculated exactly in the next iteration, the number of marked grid points increases by one and the Euler error is now below the tolerance value for all remaining asset states. We accept the found policy interpolated over 7 states and go to the next exogenous state.

This approach, however, has one drawback: it might prevent the time iteration algorithm from converging. Consider the case where the two largest Euler errors, e_{A} and e_{B} with the corresponding grid points A and B, are very close to each other but both above the tolerance value. In the first iteration, e_{A} is greater than e_{B}, thus, grid point A is calculated exactly. Now the Euler errors are recalculated and e_{B} is below the tolerance value, thus, we accept the interpolated solution for B. As we have seen in Section 3.4, the result of the current iteration enters the next iteration as the (possibly smoothed) new

Given a combination of exogenous states $s = (\varepsilon, \bar{z}, \Lambda)$:

1. For all data points $\{a_k : \mathcal{M}_{(a_k,s)} = 1\}$, calculate $a'(a_k, s)$ exactly
2. Interpolate linearly between $\{a_k : \mathcal{M}_{(a_k,s)} = 1\}$
3. Calculate the Euler error e_a over all asset states a
4. For all a_i with $\mathcal{M}_{(a_i,s)} = 0$:
   (a) Set $j \leftarrow \arg\max_{a_i} e_{a_i}$
   (b) **If** $e_j < \varphi$: **break**
   (c) **Else**:
      i. Calculate $a'(a_j, s)$ exactly by root finding
      ii. Set $\mathcal{M}_{(a_j,s)} = 1$
      iii. Interpolate again between all data points $\{a_k : \mathcal{M}_{(a_k,s)} = 1\}$ and recalculate the Euler errors

Algorithm 5: Dynamic approach to calculate the grid points in descending order of their Euler errors until all remaining points fall below the tolerance value.

policy guess. This time, e_{B} is greater than e_{A} and B is calculated exactly instead of A. In the worst case, we alternate between two solutions which yields a constant deviation between the guessed policy and the result.

We solve this problem by storing the grid points we have calculated exactly in former iterations in the sparse matrix $\mathcal{M}$. In each subsequent iteration, we calculate all grid points $(a, \mathbf{s})$ with $\mathcal{M}_{(a,\mathbf{s})} = 1$, regardless of their current Euler error, in order to avoid alternation. This approach ensures that precision can only improve with each iteration but might lead to slight redundancies. In the context of the example above, we likely end up calculating both grid points in each iteration. We will evaluate the additional effort on calculating these additional grid points in Section 5.4.

5.3 Choice of Tolerance Criterion

How do we decide if the interpolated value is close enough to the real solution? We need to define a tolerance criterion for the Euler error above which we calculate the root for the observed grid point exactly. The smaller the tolerance the more grid points we calculate exactly and the higher the calculation effort. However, if we set the tolerance criterion too loose, we risk making a high interpolation error and overlooking areas of interest.

As a compromise, we choose a very loose tolerance criterion at the beginning where we need to approximate the functional form. As the iteration converges, we simultaneously tighten the tolerance criterion such that the number of grid points calculated exactly

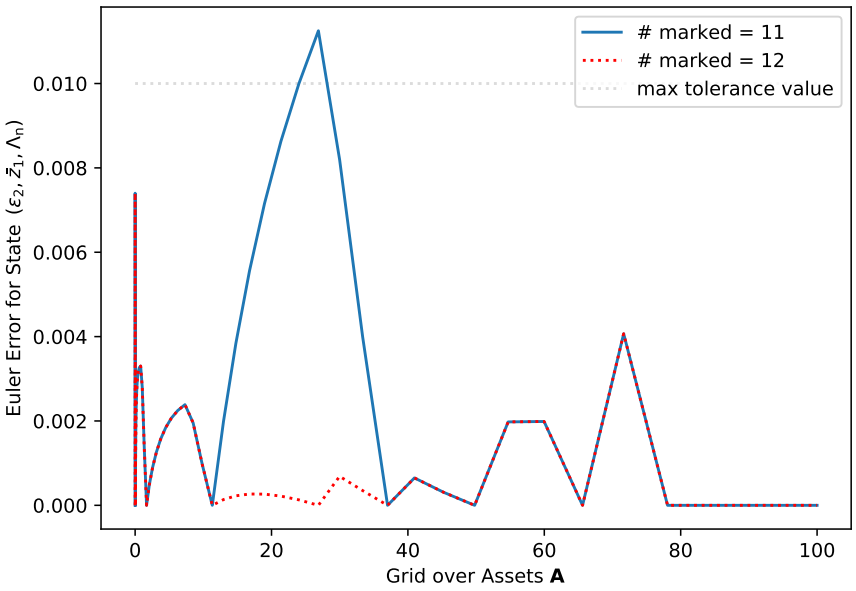

Figure 30: Development of the Euler error for an exogenous state $\mathbf{s}$ held fixed over all asset states. The blue line shows the initial Euler error for the policy interpolated over 11 data points. The dashed horizontal line marks the maximum tolerance value of 0.01. After the grid point with the highest Euler error has been adjusted, the number of marked grid points increments by one and the re-calculated Euler error (red dotted line) is below the tolerance value of 0.01 for all grid points.

increases throughout the calculation. Crucially, the tolerance criterion is reduced in *each* iteration, so we cannot run into a deadlock as described in Section 5.2.2.

We construct a function t for the tolerance criterion that becomes stricter the more iterations have been calculated. First, we choose a maximum tolerance $t_{\max}$. This is the maximum Euler error allowed, thus, it defines a lower bound for the accuracy. Second, each iteration we multiply the tolerance by a fixed factor $0 < f_t < 1$. The smaller f_t the faster the tolerance criterion is reduced. Third, we add an offset o if t should approach a value greater zero. Grid points with Euler errors below o will not be calculated. The formula to calculate the tolerance criterion is:

$$
\begin{aligned}
t_i &= \hat{t}_i + o, \\
\hat{t}_i &= \begin{cases} t_{\max}, & \text{if } i = 0 \\ t_{i-1} \cdot f_t, & \text{if } i > 0. \end{cases}
\end{aligned} \tag{44}
$$

5.3.1 Using a Constant Tolerance Value

Using a lose tolerance criterion in the beginning and tightening it as the convergence makes progress seems reasonable as it brings the algorithm quickly towards an approximate solution. However, this comes at the cost of more iterations since premature errors have

to be adjusted later. We compare two scenarios: First, we start at $t_{\max} = 0.8$ and decrease exponentially with factor $f_{\mathrm{t}} = 0.95$ down to an offset $o = 0.05$. Second, we keep the tolerance value constant at 0.05 throughout the whole iteration.

As expected, the mean number of grid points per iteration of the latter variant is larger by a factor of roughly 1.25 (176 compared to 138) and the maximum number of grid points is larger by a factor of roughly 1.1 (183 compared to 167). However, we observe a surprising reduction in the number of iterations from 105 for the decreasing tolerance value to 79 for the constant tolerance value. In terms of total root finding steps during the TI execution, the constant tolerance leaves us with 600 or 5% root finding steps *less* as well as a smaller mean Euler error of 1.02% compared to 1.17%.

Driven by this insight, we conduct several experiments with different tolerance criteria, varying $t_{\max}$, f_t and o, and compare the results with several constant tolerance values. Table 19 provides the data for different tolerance criteria, sorted by the mean Euler error in ascending order. We use the total number of root finding steps throughout the whole execution of the TI algorithm as an indicator for the computational effort.

First, the data shows a negative correlation between the number of root finding steps and the mean Euler error, implying that we need to make some effort in order to gain sufficiently accurate results. Notably, each change that speeds up the convergence of the tolerance value (i.e., choosing a smaller starting value or a larger adjustment factor) leads to a smaller mean Euler error, indicating that roughly approximating the policy initially does not bring us any advantage. Moreover, the data reveals that if the adjustment factor f_{t} is chosen too high and simultaneously the stopping criterion is chosen too lose, we risk premature termination. This is the case for run no. 12, where the calculation terminates already at $t = 38.56\%$. Thus, another advantage of using a constant tolerance criterion is that we can precisely quantify the maximum possible Euler error beforehand and do not risk early convergence.

While we observe large improvements concerning the mean Euler error, the *maximum* Euler error is rather invariant under different tolerance criteria. We explore different explanations for the high maximum Euler error in the next section.

We conclude that we need to choose the tolerance criterion carefully in order to obtain accurate results. Moreover, we need to take into account the termination condition for the TI algorithm and the smoothing parameter we calculate the new guess with. In any case, we face a tradeoff: Increasing the smoothing parameter or decreasing the termination condition leads to a higher number of iterations and therefore also to a smaller tolerance criterion. However, if we choose these parameters too losely, we risk an early termination and inaccurate results.

Table 19: Comparison of runtime data for different tolerance criteria. The list is sorted by the mean Euler error (EE) in ascending order. #root denotes the number of root finding steps. t_{final} denotes the minimum tolerance value at which the algorithm terminates. If the tolerance value is an exponentially decreasing function, t_{final} cannot be determined beforehand.

	$t_{\max}$	f_{t}	o	#root	mean EE [%]	max EE [%]	t_{final}[%]
1	0.001	0	0.001	97840	0.04	6.62	0.1
2	0.01	0	0.01	27208	0.27	6.32	1.0
3	0.02	0	0.02	20097	0.47	6.64	2.0
4	0.03	0	0.03	16511	0.72	6.59	3.0
5	0.04	0	0.04	14852	0.9	6.59	4.0
6	0.05	0	0.05	13904	1.02	6.62	5.0
7	0.8	0.9	0.05	11850	1.08	6.65	5.02
8	0.2	0.95	0.05	16116	1.09	6.67	5.08
9	0.5	0.95	0.05	13490	1.17	6.59	5.36
10	0.8	0.95	0.05	14490	1.17	6.58	5.36
11	1.0	0.95	0.05	14960	1.18	6.58	5.35
12	0.8	0.99	0.05	8991	2.88	29.94	38.56

5.4 Results

In view of the upcoming comparative analysis, we choose a benchmark configuration we consider to be sufficiently accurate while simultaneously feasible and conduct several experiments concerning the way grid points are chosen. Our experimental design needs to ensure constant circumstances as far as possible. Since the runtime of Algorithm 1 highly depends on the quality of the initial guesses for the interest rate r and the quality-adjusted capital stock Q, we concentrate on one single iteration and specifically on the execution of the TI algorithm. Furthermore, we try to eliminate any disturbances caused by erroneous guesses of r and Q by choosing the last iteration before Algorithm 1 converges. We evaluate the total TI runtime, the duration of the single subiterations as well as the respective Euler errors solely in the context of this last near-to-equilibrium iteration.

The goal of this analysis is to find the best compromise between maximum accuracy and minimum computational effort. We concentrate on the TI Algorithm and evaluate the choice of grid points and the total number of root finding steps necessary. Furthermore, we observe the Euler error over the fine and the interpolation error over the original grid in order to determine "blind spots". There are different reasons for high Euler errors that are associated with the discretization of the asset space, the choice of the tolerance criterion, the consideration of the borrowing constraint as well as the ordering of grid point evaluation.

5.4.1 Benchmark Configuration

We choose the same economic parameters as in Chapter 4 for the benchmark configuration and a constant tolerance value of 0.01. The last execution of the TI algorithm before the equilibrium is reached takes 5.3 minutes and 76 iterations in total. The mean duration is 4.2 seconds. The first iteration stands out with 7.9 seconds as we start from an initial all-zero policy guess.

Of these 4.2 seconds per subiteration, 92.8% account for the root finding and 6.7% account for the calculation of the interpolation errors. The total time spent on calculating the interpolation error roughly corresponds to 5 additional root finding steps.[7]

Compared to the static approach introduced in Section 5.1.1, where we calculated all grid points except for the redundant retirement lanes, the dynamic choice of grid points reduces the runtime per subiteration roughly by factor 9 (38.7 seconds compared to 4.2 seconds).

Turning to the number of grid points calculated exactly, we find a mean number of 358 grid points per subiteration which accounts for 9.4% of all possible grid points, or 12.8% if we do not take the redundant retirements into account as proposed in Section 5.1.1. Keep in mind that the number of marked grid points can only increase from one iteration to the next, since grid points that have been marked once will be calculated exactly until the TI converges. Figure 31 plots the number of marked grid points from the first to the last subiteration of the TI algorithm. Although the tolerance criterion is constant at 0.01, the number of marked grid points rises from 207 to 379 due to the redundancy. This implies that the share of marked grid points in all possible grid points rises from 5.4% to 9.9%, or from 7.4% to 13.5% without the redundant retirement lanes. However, the number does not grow constantly but reaches a plateau around iteration no. 30. This corresponds to the progress of the convergence, which we measure by the volatility of the mean interpolation error and illustrate in Figure 32

The total number of root finding steps over all subiterations is 27,228. Figure 33 shows the average number of exact calculations over all iterations for each grid point. Note that the grid has the same structure and the same initialization as in Figure 26. The figure acts as a good indicator for areas of interest in the policy functions. For example, there is a large coherent block of marked grid points in the lowest exogenous state $\mathbf{s}_0 = (\varepsilon_0, \bar{z}_0, \Lambda_\mathrm{r})$. In fact, as discussed in Chapter 4, this is precisely the area where the shape of the consumption policy function is rather curved than linear. This applies to all five retirement lane states which are marked by dashed lines in Figure 33.

Next, we evaluate the Euler errors in order to assess the accuracy of the found solution. The mean Euler error over the fine grid is 0.27% compared to a mean interpolation error of 0.23% for the original grid. The maximum Euler error of the fine grid is 6.31%.

[7] Based on an average number of 25 iterations over the system of equations per root finding step and an approximate duration of 0.0027 seconds for solving the system of equations.

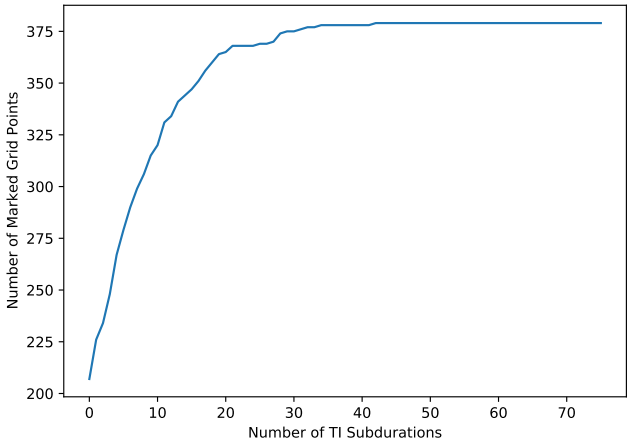

Figure 31: Number of grid points calculated exactly over the 76 subiterations of the TI algorithm.

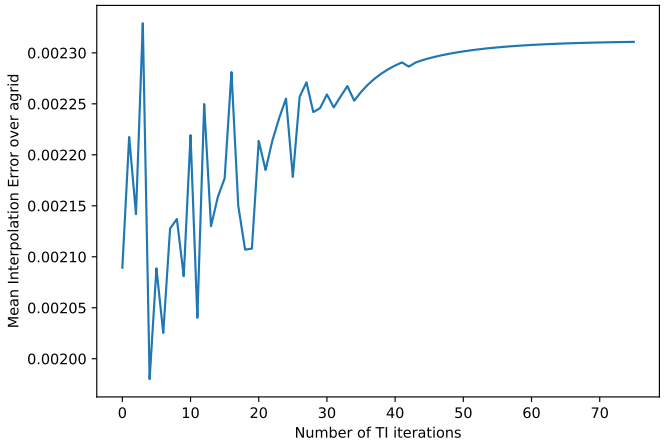

Figure 32: Mean interpolation error over the 76 subiterations of the TI algorithm.

In comparison to the static approach based on a priori knowledge we introduced in Section 5.1, the dynamic approach relies on only one third of the grid points but produces more accurate results. Precisely, we manage to reduce the mean Euler error by 75% and the maximum Euler error by 50% compared to the static approach.

However, the maximum Euler error of the fine grid is still more than six times larger than the allowed 1% interpolation error we specified by the tolerance value. Figure 34 reveals where the high Euler error comes from: At the middle asset levels, the mean Euler error is constantly small at around 1%. The high Euler errors appear at the very low and the very high asset levels. At the low asset levels, this is due to the borrowing constraint we did not identify exactly. At the high asset levels, the reason for the high Euler errors is entirely different. Remember that when discretizing the asset space $\mathbf{A}$, we placed many grid points at the bottom of the asset range in order to identify the first kink, and only

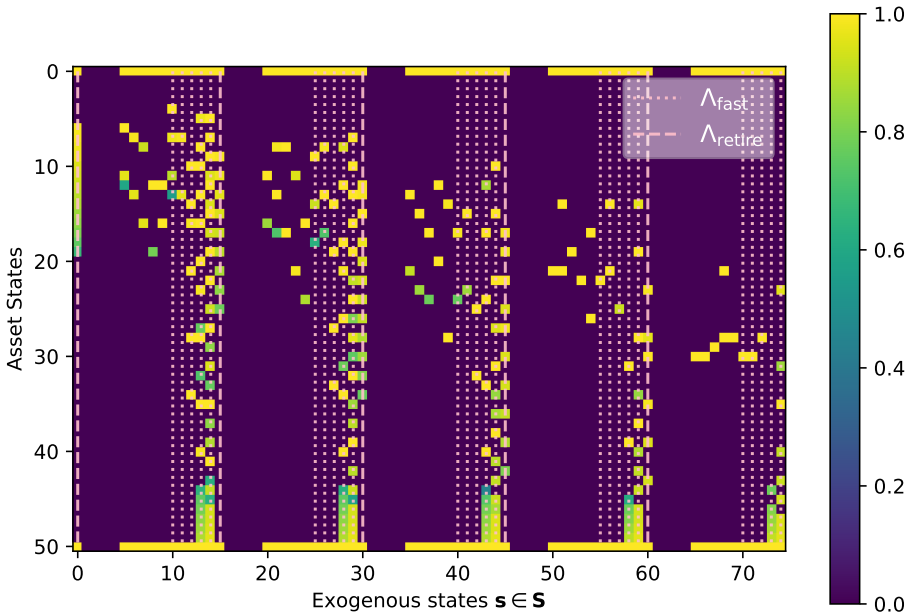

Figure 33: Mean binary-coded matrix $\mathcal{M}$ over all subiterations of the TI. The grid has the same shape as the grid for the policy calculation, thus, the horizontal dimension reflects the space over exogenous states and the vertical dimension the asset space. In each subiteration, $\mathcal{M}$ can only take the values 0 or 1. Thus, the color code ranges from purple (grid points are never marked) to yellow (grid points are marked in each iteration). The fast lane as well as the retirement lane states relevant for the calculation are marked by vertical dotted and dashed lines.

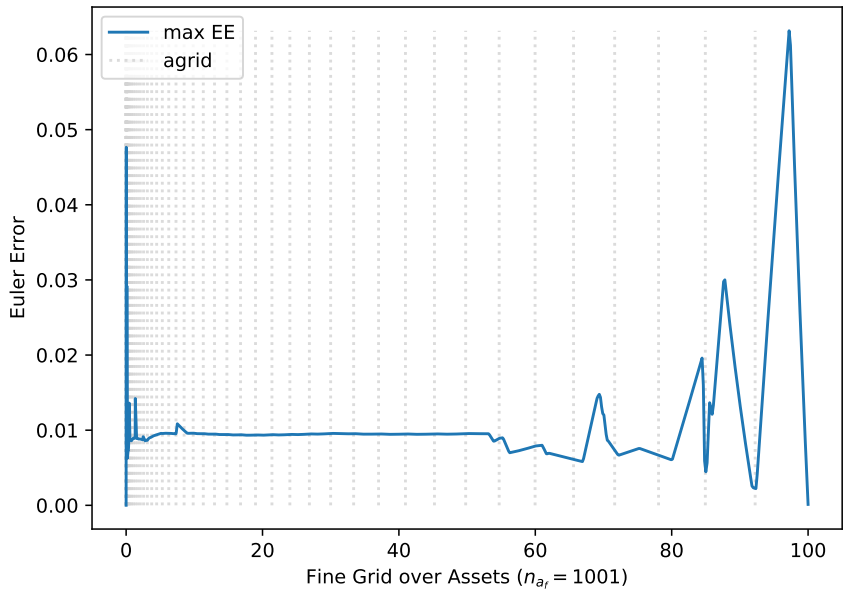

Figure 34: Maximum Euler error calculated over the fine asset grid. The vertical dotted lines mark the position of grid points the policy was calculated for.

few at the top. Figure 34 marks the location of grid points in the original grid the policy was initially calculated for. Clearly, the Euler error is low where grid points have been placed - by construction, it must be below 0.01 at any grid point. However, the distances between the grid points at the right half of the plot are comparably large. As this is exactly the area where the second kink is likely to appear, we get high Euler errors in the intervals between two grid points.

Our conjecture is further emphasized when we evaluate the maximum Euler error by exogenous state $\mathbf{s}$. Figure 35 shows the maximum Euler error by exogenous state over all asset states. We add vertical grey lines to the graph to distinguish between the different lanes. Most of the tips appear in the fast lane but we also observe high errors in the retirement lane. These are precisely the lanes where the second and the first kink appear. Again, we see that the maximum interpolation error for the original grid (orange plot) is moderate and nearly constant around 1%.

5.4.2 Exact Solution

We compare the results found by the adaptive approach to the full-blown solution where we calculate each grid point exactly. The duration per execution of the TI algorithm jumps up to 40 minutes and the average duration of one subiteration is 32 seconds. The macroeconomic results are slightly different: the interest rate obtained by the full solution is higher by 11% (or 0.0084 percentage points, respectively), aggregate savings and investment levels are 6.5% lower, and the quality-adjusted capital stock is 7% lower.

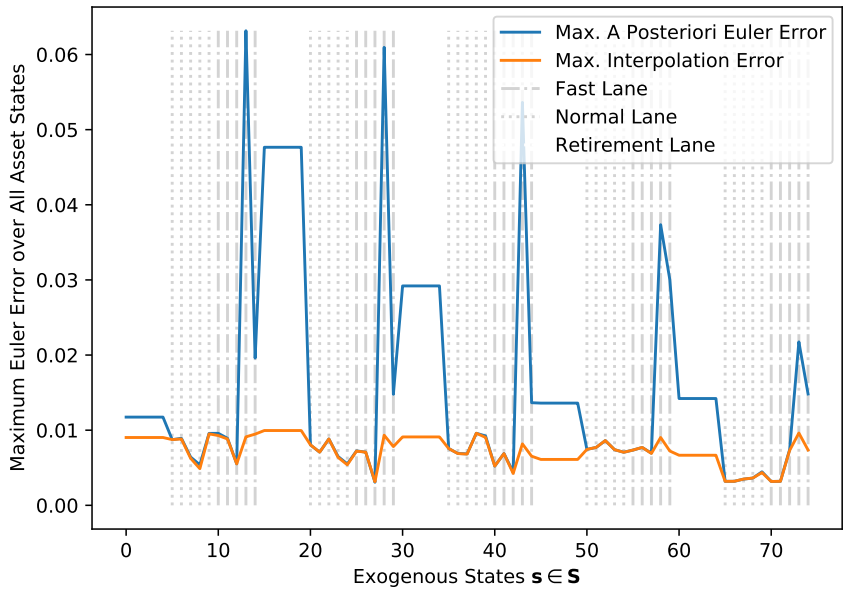

Figure 35: Maximum Euler error by Exogenous State. The blue plot shows the a posteriori Euler error over the fine asset grid, the orange plot shows the interpolation error over the original grid. The vertical grey dotted and dashed lines mark the position of the different lanes.

The capital-to-output ratio slightly drops from 5.96 to 5.7 and the share of the wealthiest 1% in total wealth rises by 1.17 percentage points to 25.75%. Overall, the resulting distribution is fairly similar with a maximum deviation of one percentage point for the earnings and income percentiles and 1.17 percentage points for the wealth distribution.

While these differences are not to be underestimated when evaluating the results of the model in detail, the overall qualitative results regarding the shape of the policy functions, the distribution of agents and the macroeconomic outcomes remain roughly the same. We recall that the motivation behind developing the adaptive approach in the first place was to improve the feasibility of working with the model. Specifically, the approximate approach dramatically decreases the runtime during the trial-and-error process of finding the correct parametrization. In this light, our implementation fully serves the purpose in that the found equilibrium is not expected to substantially deviate from the exact solution. Once a sufficient parametrization has been found, we suggest to recalculate the model exactly, with the initial guesses being already fairly near to the real outcome.

Furthermore, evaluating the Euler errors of the full solution provides an interesting insight: The mean Euler error of the exact solution is only 9% lower (0.244%) than for the adaptive approach (0.268%). However, Figure 36 reveals that the high mean Euler error is mainly due to the lowest asset states: the mean Euler error over all exogenous states is

almost four times as high for the exact solution compared to the approximate solution.[8] Therefore, we identify the same areas of high errors in both solutions, indicating that the problem is indeed caused by the discretization of the asset grid and the sensitivity of the model for small deviations.

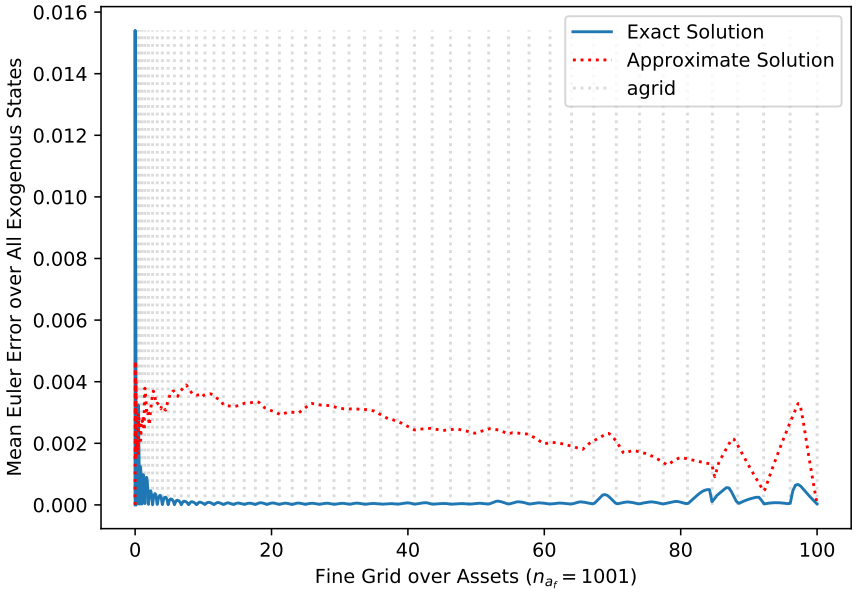

Figure 36: Mean a posteriori Euler error over the asset grid for the exact solution (solid blue line) and the approximate solution (dotted red line). The vertical grey dotted lines mark the position of the asset states.

5.4.3 Modifying the Discretization of the Asset Space

As we have seen in the last section, choosing only few grid points in the upper half of the asset range can cause high Euler errors. In this section, we modify two central parameters of the asset space discretization: first, by doubling the total number of grid points, and second, by making the asset grid equally-spaced.

A central feature of our algorithm is that it is designed to ignore the majority of the grid points and dynamically pick only few of interest. Thus, we expect that doubling the total number of asset states does not yield much additional effort (in contrast to the static a priori approach discussed in Section 5.1).

Indeed, if we double the number of asset states (and thus the total number of grid points), the mean number of marked grid points per iteration rises only moderately by 2%. Note that this effectively implies a reduction in the share of marked grid points in

[8] Remember, that the Euler error itself is only an approximation of the real deviation, as we do not know the true policy function. Put differently, the Euler errors for the two solutions assume different underlying policy functions.

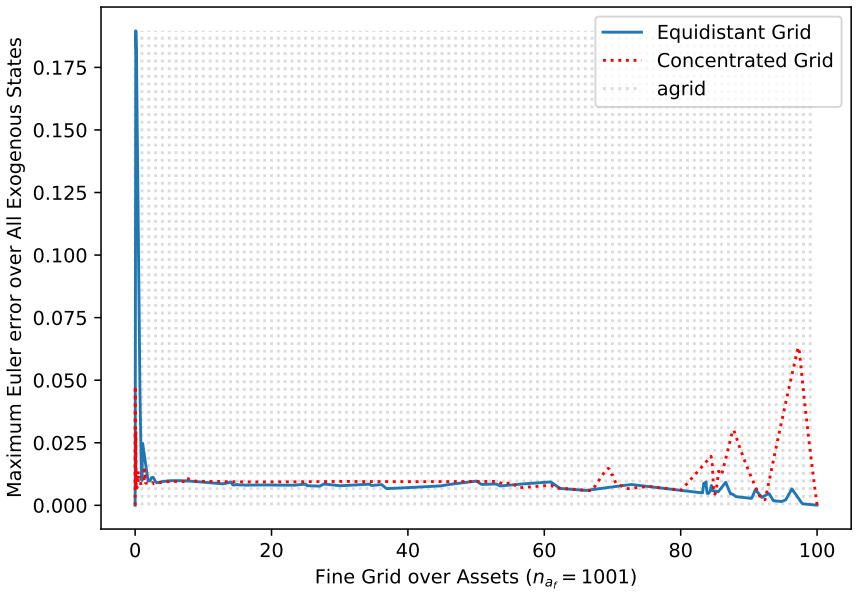

Figure 37: Maximum Euler error by asset states for the equidistant grid (solid line) vs. grid point concentration at the bottom of the asset range (dotted line). The vertical dotted lines mark the values of asset states for the equally-spaced case.

all possible grid points (leaving the redundant retirement states out) from 12.7% to 6.6% on average.

With an equally spaced grid of double precision, we can locate the second kink more precisely. In the most extreme case, the interval size between $a_{\max}$ and $a_{\max-1}$ is almost divided by four from 3.96 to 1 when the concentration parameter ϑ_a is changed from 2 to 1 and if, additionally, we double the number of asset states. However, this implies that the size of the intervals at the lowest asset states simultaneously grows by factor four.

The results of this approach are shown by Figure 37. The maximum Euler errors at the top of the asset range vanish thanks to the smaller interpolation intervals. However, this comes at the cost of precision loss at the bottom of the asset space which causes the maximum Euler error to triple compared to the benchmark level. The mean Euler error over the whole grid is reduced by 9% compared to the benchmark level.

We conclude that the chosen level of moderate grid point concentration is the best compromise we can get as it keeps both the mean and the maximum Euler error in reasonable bounds over the whole grid.

5.4.4 Bisection Approach

Next, we compare the results of the bisection-based approach from Section 5.2 where grid points are evaluated in a fixed order to the dynamic approach where grid points are evaluated in descending order by their Euler errors. Figure 38 plots the sparse matrices

$\mathcal{M}$ for the two variants. The overall placement of marked grid points is overall very similar. However, the left plot shows horizontal bars at the middle asset state, since it is the first and often the only grid point to be calculated exactly. In contrast, the adaptive approach places roughly the same number of grid points per column but precisely those where the deviation is the highest.

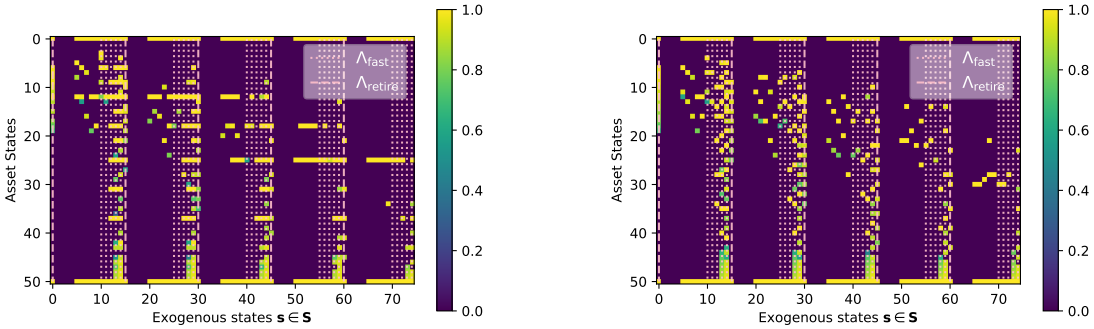

Figure 38: Mean binary-coded matrix $\mathcal{M}$ over all iterations for the bisection-based fixed-order approach (left side) vs. the adaptive approach that evaluates grid points in the descending order of their Euler errors (right side).

Surprisingly, both approaches yield almost exactly the same equilibrium interest rate, quality-adjusted capital stock and distributional data. For a constant tolerance value of 0.01, the equilibrium interest rate is 7.11% for the adaptive approach and 7.10% for the bisection-based approach and the quality-adjusted capital stocks differ by 0.04%. The technical data is exactly the same: both approaches yield a total number of 76 iterations in the near-equilibrium TI execution with 358 marked grid points on average and 376 and 379 grid points respectively at maximum. The mean Euler error is smaller by 0.07 percentage points for the adaptive approach compared to the bisection variant and the maximum Euler error is smaller by 0.055 percentage points. We conclude that at least at this level of precision, the adaptive approach does not save us any grid points, although it presumably targets the kinks more efficiently.

Nevertheless, these differences grow in the tolerance criterion, implying that the loser the final tolerance the larger the deviation. Repeating the analysis for a tolerance value that starts at 0.8 and decreases with factor 0.95 down to an offset of 0.05 yields remarkable differences between the two approaches. The adaptive approach relies on 105 iterations, 138 grid points on average and 167 at maximum. In contrast, the bisection-based variant requires 27 additional iterations and 143 (170) grid points on average (at maximum), which yields an increase in the number of root finding steps by 30%. While the former converges already at $t = 5.36\%$, the minimum tolerance for the latter is $t = 5.09\%$ which yields a slightly smaller mean Euler error (which, again, advocates a constant tolerance value).

Furthermore, the fixed order in which the grid points are evaluated can lead to divergence of the outer algorithm when combined with a constant tolerance criterion. This situation appears when we combine the bisection-approach with a constant tolerance value (whose advantages we have discussed in detail in Section 5.3). Remember that the order in which grid points are evaluated is preassigned by the bisection algorithm. In combination with a comparably lose tolerance value of 0.05 this leads to divergence of the outer algorithm, where every second iteration yields the same policy and Q and r alternate between two values. We therefore highly recommend to use the adaptive approach instead which ensures more variation as the algorithm dynamically decides which grid points to evaluate first.

5.4.5 Precisely Locate the First Kink

As we have seen in Chapter 4, the existence and exact location of the two kinks in the policy functions depend on the combination of assets and exogenous states $\mathbf{s}$ as well as the state transition probability and cannot be analytically determined. However, since the first kink applies to all households with negative optimal savings, we can relatively easily identify it "on the fly" during execution of the Time Iteration algorithm.

As proposed in Section 5.2, we initialize the binary-coded matrix $\mathcal{M}$ with ones in the first and the last row and calculate the interpolated policy for each exogenous state $\mathbf{s}$ separately. If the policy $a'(a_i, \mathbf{s})$ is negative or zero, we register the next grid point (a_{i+1}, s) for the calculation by setting the respective entry $\mathcal{M}_{(a_{i+1},\mathbf{s})} = 1$. As soon as the first positive policy value has been found, we can stop the exact calculation for this exogenous state $\mathbf{s}$ (as agents with sufficiently high skills are not likely to be constrained at all, the policy is positive already at the lowest asset state for most columns). Subsequently, we proceed with the dynamic choice of grid points subject to the Euler error.

This approach increases the number of total root finding steps by 15% when used with a constant tolerance value of 0.01, compared to the benchmark approach where the first kink is not targeted directly. The mean Euler error is slightly smaller by 0.01 percentage points. The results for the euqilibrium interest rate and the quality-adjusted capital stock are the same. The share of agents at the lowest asset state is about 0.1 percentage points larger, since the borrowing constraint can be located more precisely. However, zooming into the policy function at the constrained states reveals that the difference between the two approaches is rather negligible as shown in Figure 39. When zoomed out, the difference is not even visible anymore.

Figure 40 reveals that the maximum Euler error at the bottom of the asset grid is still relatively high at 0.03 compared to 0.05 before - indicating that the error does not arise from the points *at* the borrowing constraint but rather from the ones slightly above.

This conjecture is further strengthened when we evaluate at the interpolation error made during the calculation that does not search specifically for the borrowing constraint.

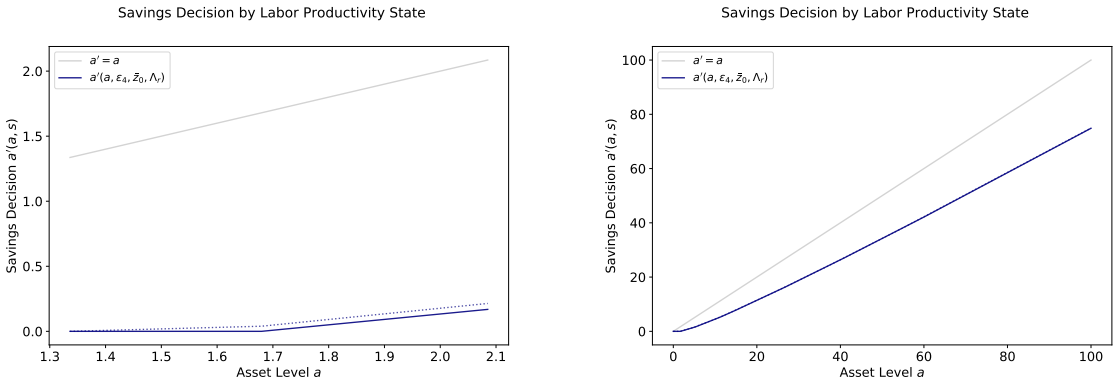

Figure 39: Policy function for the exact calculation of constrained asset states (solid line) vs. the interpolated approach (dotted line). The difference near the kink is almost negligible when zoomed in and vanishes when zoomed out to the whole asset grid.

Figure 41 plots the interpolation error by exogenous state for the five lowest asset states (except $a_{\min}$ which is by construction calculated exactly). The maximum asset level where the borrowing constraint is binding depends on the exogenous state, thus, we plot only those asset states where the constraint is binding for all exogenous states. We mark the states where the constraint is binding with vertical dotted lines. Clearly, the figure reveals that the interpolation error near the borrowing constraint is negligible compared to the interpolation error for the non-constrained states.

We conclude that exactly locating the first kink does not target the areas that produce the highest Euler errors. The additional effort of 4250 root finding steps is not justified by a systematic precision gain.

5.5 Further Ideas

While in the above section we focused on interpolation along the asset level, this approach can clearly be extended to other dimensions. For example, we could fix assets, entrepreneurial ability and lane and assume the same policy for all labor productivity states as long as the deviation from the real solution is sufficiently low. As we have seen in Chapter 4, if we hold two dimensions of exogenous states fixed and vary the third, the policy functions are often but not always monotonically increasing or decreasing. This further emphasizes the potential of a dynamic approach.

We already used this idea in Section 5.1.1, where all retired agents face the same entrepreneurial ability $z = 0$, regardless of their inherent ability $\bar{z}$. However, whereas this approach is based on a priori knowledge and statically embedded in the algorithm, we could also do it dynamically. Precisely, we could configure the algorithm to detect similarities of the policy functions along the different labor productivity, entrepreneurial ability or lane states. Clearly there is potential for further efficiency improvements and correspondingly a reduction in runtime of the TI algorithm.

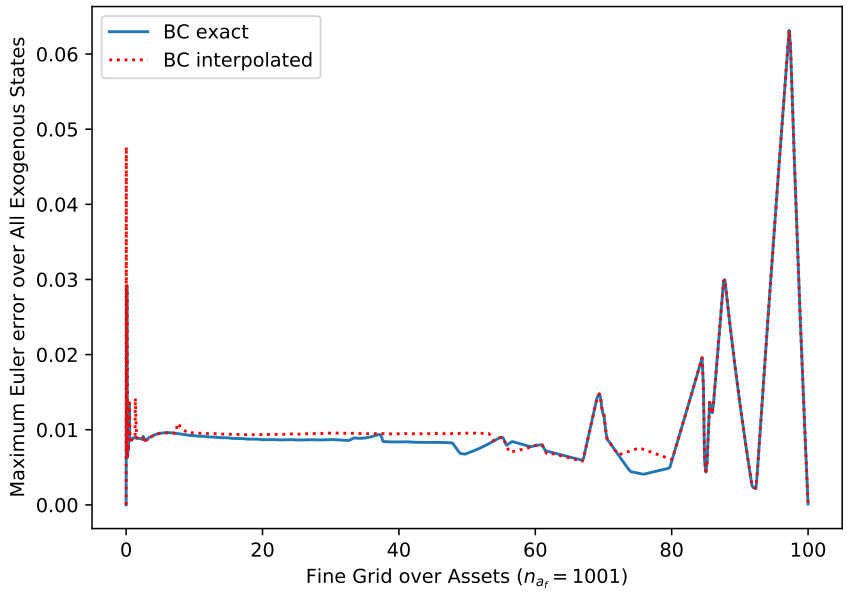

Figure 40: Maximum Euler error for the exact calculation of constrained asset states (solid line) vs. the interpolated approach (dotted line).

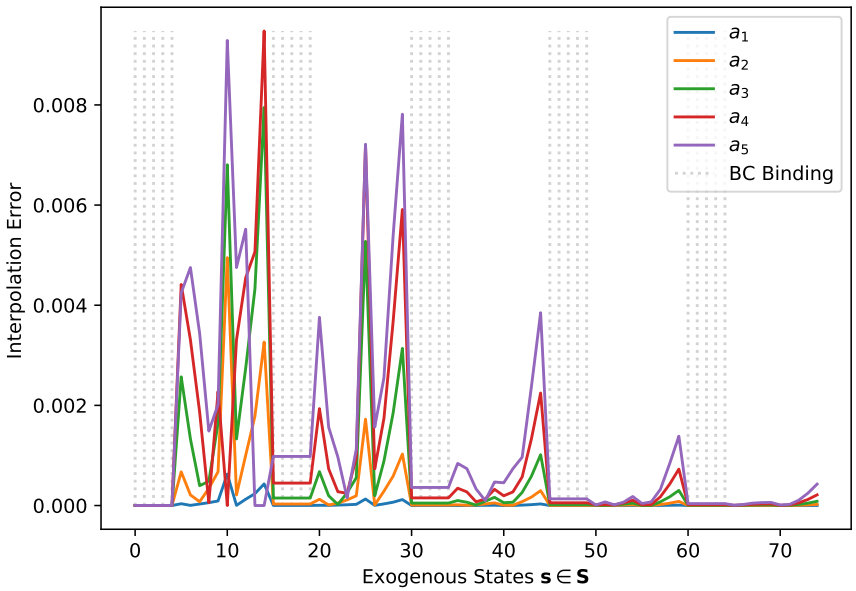

Figure 41: Interpolation error for the five lowest asset states (except $a_{\min}$). Whether a household is constrained depends on their assets as well as their exogenous state. The vertical dotted lines mark the exogenous states where the borrowing constraint is binding at all plotted asset states.

6 Conclusion and Outlook

In this thesis we translated an OLG model of intertemporal savings decisions made by entrepreneurs to an equivalent infinite-horizon Aiyagari-style model. We simplified the evolution of labor productivity and entrepreneurial ability in the model and reduced the respective number of states in each dimension. Remarkably, letting go of the OLG structure saved us an entire state variable, an agent's age, and allowed us to simplify the complex inheritance of skills between parents and children as modeled by Guvenen et al. (2019).

We discussed the differences between our model and the one by Guvenen et al. (2019) concerning our implementation as well as our choice of parameters and their impact on the results. Furthermore, we showed that despite these differences, our model is able to generate distributional results fairly close to the empirical data for the U.S. economy. Thus, it proves to be a good starting point for further tax analyses. However, there is some need for refinement of the parametrization in order to tightly match the top 1% share of the wealth distribution and lower. Further research is needed in order to evaluate whether this problem arises from a too small resolution of entrepreneurial states or if another choice of economic parameters can solve the issue.

In order to solve the model in a computationally feasible way, we designed an approach to reduce the computational effort by dynamically placing grid points at areas of interest and interpolating in between, where the deviation is below a configurable upper bound. We showed that our algorithm manages to reduce the number of grid points calculated in each iteration by 90.5%. The corresponding runtime per iteration of the TI algorithm is reduced from 55 to 4.2 minutes, which corresponds to a speedup of roughly factor 9 and yields the final solution within few hours, depending on the number of iterations. Furthermore, our algorithm allows us to predetermine the upper bound on the interpolation error the choice of grid points is based on. As a caveat, we pointed out that the a posteriori error might still exceed this lower bound since we inevitably need to discretize the continuous space over assets beforehand. However, this issue applies to the full-blown approach alike and is solely related to the sensitivity of the model in response to slight changes.

As the first subsequent step, we suggest to conduct similar tax experiments with the model as in Guvenen et al. (2019) and replace the capital income tax with a wealth tax. Section 2 already takes the necessary adjustments to the optimization problem into account. In order to obtain comparable results, we require the government budget to be balanced. Note that in our model, agents can only retire from entrepreneurial business but not from labor. Thus, the government expenses are not fed into the model in the form of pension payments and therefore did not have any practical relevance so far. However, when conducting tax experiments, we target the governmental expenses at

their equilibrium level from the capital income tax economy and adjust the wealth tax rate until the total tax revenue equals the expenses. This approach is equivalent to the "balanced budget" reform by Guvenen et al. (2019). We do not expect our results to be substantially different than theirs, since the key structure and, crucially, the assumption of heterogeneous returns remain the same.

As a second step we suggest to target several assumptions made by the authors of the original model and perform a robustness check on the results. For example, the original model features only one type of assets that are completely interchangeable between periods. However, in reality, entrepreneurs are often bound to the decisions they made in the past, since investments might not be withdrawable within a short time horizon. This yields the follow-up question whether the key findings hold if we distinguish between private assets and entrepreneurial capital and add frictions to the model, e.g., capital adjustment costs that penalize fluctuating investment decisions.

Since assumptions are made to keep the model (relatively) simple, each assumption let go poses a challenge on computational feasibility. In our case, in order to distinguish between different asset types, we need to add another continuous dimension to the state space. This is where ASGs can prove to be benefitial.

Brumm and Scheidegger (2017) state that regular grids suffer from the curse of dimensionality, i.e., the grid complexity grows exponentially in the number of dimensions. We already discussed the effect of adding states in Sec. 5.4.3. If we double the number of asset states, our dynamic choice of grid points manages to keep the computational effort almost constant. However, our approach is limited to the asset range and does not bring any advantage if we add exogenous states. In contrast, ASGs can handle highly complex models at reasonable computational effort in exchange for only a small accuracy loss. Brumm and Scheidegger (2017) state that for smooth policy functions, the number of computationally feasible dimensions can grow up to 100 when using ASGs and still up to 20 dimensions for non-smooth policy functions.

Another advantage of ASGs is that they are especially suited for economic models with occasionally binding constraints that imply kinks in the policy function. The adaptivity of the sparse grid allows us to place grid points precisely at those kinks, since they offer a substantial information gain. In our work, we make a similar attempt by introducing an algorithm in Sec. 5.2.2 to dynamically choose grid points where high Euler errors indicate movements in the policy functions. For example, our approach places many grid points at the low asset states where the consumption policy is rather curved than linear. Furthermore, we locate the kink associated with the borrowing constraint by working our way up from the lowest asset state until we found the first positive policy value. However, this approach is neither efficient nor is it applicable to the second kink in the upper half of the asset range. We further observe that the algorithm is limited to the choice of discrete asset states we make beforehand and with only little prior knowledge of the true

functional form, yielding high interpolation errors as a result. There are other approaches that are able to precisely locate the kink arising from the borrowing constraint and do not rely on root finding, like the Endogenous Grid Method (EGM). However, EGM does not work with more than one continuous state variable and is therefore not suited for extensions of the model.

In summary, the usage of ASGs yields promising gains both in terms of accuracy and computational feasibility. In view of the subsequent research endeavor building on this thesis, we are now fully equipped to take the next step, conduct tax experiments and challenge the assumption of perfect substitutability of capital. Our substantial insights from developing the dynamic approach and our detailed discussion of the Euler errors provide us with a solid intuition for the model and the mechanisms at play.

Acronyms

ASGs Adaptive Sparse Grids. 1, 2, 75, 76

CITR Capital Income Tax Revenue. 34

EGM Endogenous Grid Method. 76

GDP Gross Domestic Product. 34

OLG Overlapping-Generations. 1, 2, 7, 8, 10, 28, 46, 74

TI Time Iteration. 20, 51, 54, 57, 61–64, 66, 70, 72, 74

TTR Total Tax Revenue. 34

Bibliography

Aiyagari, S. Rao (). "Uninsured Idiosyncratic Risk and Aggregate Saving". *The Quarterly Journal of Economics 109* (3), 659–684.

Brumm, Johannes and Simon Scheidegger (). "Using adaptive sparse grids to solve high-dimensional dynamic models". *Econometrica 85* (5), 1575–1612.

Guvenen, Fatih, Gueorgui Kambourov, Burhanettin Kuruscu, Sergio Ocampo-Diaz, and Daphne Chen (). "Use it or lose it: Efficiency gains from wealth taxation". Tech. rep. National Bureau of Economic Research.

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

Kuhn, Moritz, Jose-Victor Rios-Rull, et al. (). "2013 Update on the US earnings, income, and wealth distributional facts: A View from Macroeconomics". *Federal Reserve Bank of Minneapolis Quarterly Review 37* (1), 2–73.

A Full Optimization Problem with Endogenous labor

The utility function of the household is a Cobb-Douglas function:

$$u(c,(1-\ell)) = \frac{(c^{\gamma}(1-\ell)^{1-\gamma})^{1-\sigma}-1}{1-\sigma}, \tag{45}$$

where γ is the share of consumption in utility and σ is the risk aversion parameter.

The derivatives of the utility function are given by:

$$\frac{\partial u(c,(1-\ell))}{\partial c} = \gamma \cdot \left(c^{\gamma}(1-\ell)^{1-\gamma}\right)^{-\sigma} \cdot \left(\frac{c}{1-\ell}\right)^{\gamma-1} \tag{46}$$

$$\frac{\partial u(c,(1-\ell))}{\partial (1-\ell)} = -(1-\gamma) \cdot \left(c^{\gamma}(1-\ell)^{1-\gamma}\right)^{-\sigma} \cdot \left(\frac{c}{1-\ell}\right)^{-\gamma} \tag{47}$$

Inserting (46) and (47) into (28) yields:

$$(1-\gamma) \cdot \left(c^{\gamma}(1-\ell)^{1-\gamma}\right)^{-\sigma} \cdot \left(\frac{c}{1-\ell}\right)^{-\gamma} = \gamma \cdot \left(c^{\gamma}(1-\ell)^{1-\gamma}\right)^{-\sigma} \cdot \left(\frac{c}{1-\ell}\right)^{\gamma-1} \cdot \frac{1-\tau_{\ell}}{1+\tau_c} \cdot \bar{w}\varepsilon \tag{48}$$

We can simplify (48) in order to obtain ℓ:

$$\begin{aligned} (1-\gamma) &= \gamma \cdot c \cdot (1-\ell) \cdot \frac{1-\tau_{\ell}}{1+\tau_c} \cdot \bar{w}\varepsilon \\ \Leftrightarrow \quad \ell &= 1 - \frac{1-\gamma}{\gamma} \cdot \frac{1+\tau_c}{1-\tau_{\ell}} \cdot \frac{c}{\bar{w}\varepsilon} \end{aligned} \tag{49}$$

To solve for c, we insert (49) into (18):

$$c = \frac{\gamma}{1+\tau_c}\Big[\omega(a,z;\tau_{\text{cap}}) - a' + \bar{w}\varepsilon(1-\tau_{\ell})\Big] \tag{50}$$