

Solving General Incomplete Market Models with Substantial Heterogeneity

Thomas M. Mertens* Kenneth L. Judd†

June 7, 2011

PRELIMINARY DRAFT

Abstract

We propose a simple but general solution method for models with incomplete markets and finitely but arbitrarily many heterogeneous agents. Our method can handle many state and choice variables for each agent and thus an extremely high-dimensional state space. The solution technique is based on perturbation methods that build an approximation around a point at which the solution is known. At this point, agents are either identical or their policy functions are known. The novel idea is to exploit the symmetry of the problem to avoid the curse of dimensionality. Our method underlies an approximation theory that specifies the speed and radius of convergence as well as the class of models to which our method applies. As a result, we study the interaction between uninsurable idiosyncratic labor income risk and asset price dynamics within a standard macroeconomic model. We show that not only the variability of individual wealth but also its comovement with other agents' wealth plays an important role.

***New York University**, Stern School of Business; Postal Address: 44 W Fourth Street, Suite 9-73, New York, NY 10012, USA; E-mail: mertens@stern.nyu.edu.

†**Hoover Institution**, Stanford University; Postal Address: 434 Galvez Mall, Stanford, CA 94305, USA; E-mail: judd@hoover.stanford.edu

1 Introduction

A large body of literature in finance and macroeconomics makes the simplifying assumption that aggregate variables are determined by the behavior of a representative agent. In reality, different people earn different incomes, have different talents, and hold different expectations. For this heterogeneity to be reflected in economic outcomes, incompleteness of asset markets is essential. In reality, substantial amounts of idiosyncratic risk can only be partially insured. Unemployment risk serves as one of the prime examples. Economic models built on the assumption of a representative agent thus do not reflect a crucial aspect of our society. Furthermore, modeling heterogeneous agents permits a more stringent test of our current economic theory since we can use information about the entire distribution of economic variables across the population.

This paper provides a solution method for a broad range of incomplete markets models with substantial heterogeneity, i.e. an economy with finitely but arbitrarily many different agents. The goal is to develop a simple but general solution method for economies in which the state space is very large. A multitude of state variables arises from heterogeneity but might increase if there are several variables for each individual.

To demonstrate the performance of our solution method, we focus on one particular economy. We solve a dynamic stochastic general equilibrium model with aggregate risk in production and an endogenous capital stock. A firm produces a single consumption good which households consume according to a Cobb-Douglas production function. Future total factor productivity is uncertain due to aggregate shocks. Households maximize expected discounted utility given by a utility function featuring constant relative risk aversion in consumption. We add a cost of deviating from a target level of capital. This cost serves two purposes. First, it makes borrowing costly and thus serves as an endogenous borrowing constraint. Second, it facilitates the solution method as it determines the distribution of capital in the deterministic steady-state. Given their utility function and budget constraint, each household decides how much to consume and save each period.

We add idiosyncratic shocks to labor income which agents cannot insure against. Households only trade claims to capital which renders markets incomplete. Agents thus cannot fully insure against their idiosyncratic labor income shocks. As a result, equilibrium outcomes feature idiosyncratic components. Households hold different levels of capital which translates into inequality of wealth and consumption.

The analysis of this model presents a difficult problem. Ultimately, we want to be able to study the interaction of choices and prices with the distribution of capital. In particular, we

need a solution method that solves for individual behavior and aggregate variables including asset prices as a function of the entire distribution of economic conditions. But, in turn, this distribution is affected by all individuals' behavior.

To solve this problem, we develop a solution technique for models with many heterogeneous agents and incomplete markets based on perturbation methods. We focus on a finite number of agents where the overall number can be arbitrarily large, i.e. there is substantial heterogeneity. The entire distribution of individual state variables enters the analysis.

Perturbation methods build an approximation of the optimal policies as functions of the state variables based on Taylor expansions. The first step is to find a special case of the model in which the solution is known. This reference point could be a deterministic steady-state of the model or a closed-form solution of a simplified version. Around this solution, we build an approximation to all optimal policies of our true model.

Our model possesses a well-defined deterministic steady-state around which we expand optimal policies with respect to all state variables. A model without a cost of deviating from a target level of capital would not have this feature. A unit-root problem arises in this case. Only the aggregate level of capital and not its distribution matters in the deterministic steady-state. Thus, *any* distribution of capital with the average capital being the steady-state level would result in a feasible steady-state. However, the dependence on the target level of capital can be made vanishingly small once we move from the deterministic to the stochastic economy.

Having pinned down the deterministic steady-state, we build an expansion with respect to all state variables. We know that equilibrium outcomes are functions of the state space. Thus we expand the deterministic economy in all state variables. But since we allow for arbitrarily many agents, we might also have arbitrarily many state variables.

The novel idea lies in exploiting the symmetry of decision rules across agents. If two agents are identical in their objectives, they should respond identically in the same economic situation. For example, starting out from a case where both agents live under the same economic condition, a marginal increase in agent one's wealth will impact the decision of agent two the same way that agent one's decision would have been impacted by the same change in agent two's wealth. Exploiting the symmetry, we solve for the decision rules of all agents as a function of the entire distribution of individual states. In a second step, we also incorporate differences in individuals' objectives.

The last step lies in moving from the deterministic to the stochastic economy. Since shocks are part of the state space, the previous expansion delivers equilibrium reactions to known,

deterministic changes in the shocks. For example, the previous expansion would compute the asset price reaction if next period's productivity was above its steady-state level. To move to the stochastic economy, we integrate over all possible realizations of the shocks and weight them by their probability. From this logic it follows immediately that we need a higher-order expansion. If we were to resort to a first-order approximation, integrating over the first-order approximation would not affect equilibrium behavior since a linear solution is certainty-equivalent. Higher-order expansions bring in the effects of uncertainty. A second-order approximation reflects the effect of the variance of shocks, a cubic approximation additionally takes the skewness into account, and so on.

Our solution method is asymptotically valid and converges to the true solution. By adding higher moments, we can eventually recover the true policy function. In practice, of course, convergence is not complete. However, we have a means of testing the accuracy of our solution. We plug our approximation into the equilibrium conditions to check its optimality.

This solution method satisfies five important criteria. First, it is generally applicable. It handles general competitive equilibrium and dynamic programming problems for partial and general equilibrium problems whenever the distribution of state variables matters. Second, it solves for the distribution of all economic variables. In our solution, all economic variables are functions of the entire distribution of state variables. Furthermore, we can compute the distribution of all individual choices. Third, the user can specify any accuracy that the method in turn delivers. The approximation is asymptotically valid and the method can thus produce results with any desired precision. Fourth, the method comes with a way to check for its accuracy. And lastly, there exists an approximation theory underlying the solution method. These features require further explanation.

The solution method is applicable to a wide range of applications. It applies whenever equilibrium or optimality conditions for a competitive equilibrium or dynamic programming problems imply that the choice variables are smooth functions of state variables. The dynamic programming problem or competitive equilibrium can feature arbitrarily many state variables and is thus interesting for a large set of economic applications. The implementation of constraints invalidates the smoothness of choices. However, for our method to apply, there are two ways to fix it. Either, one can smooth out the constraint such that there is no kink, or one can implement an endogenous borrowing constraint. Our economy provides an example for the latter.

The approximation returns equilibrium choices and prices as functions of the entire distribution of state variables. Thus, our technique solves for the true functions without having to approximate the state space. Since it is based on perturbation methods, the user can choose

the degree to which equilibrium variables are approximated. Each degree requires adding new “moments” (or, better, statistics) of the distribution of states. These are the statistics that are most useful in reducing the approximation error.

We do not make any simplifying assumptions on the state space and the solution is asymptotically valid, i.e. the approximation error decreases to zero in the limit. The higher the degree of accuracy, the better the solution will be within the radius of convergence.

Thus, we can implement a stopping criterion that requires the solution to be exact to any given accuracy. After every step of adding another degree, we can compute errors in the equilibrium conditions. Once these errors are driven below a certain stopping criterion, we can terminate the algorithm.

Importantly, our solution underlies an approximation theory. Hence we know that the approximation converges to the true solution within the radius of convergence, how big that radius is, and at which speed convergence occurs.

As a result, our solution method delivers to which order moments of the wealth distribution affect equilibrium outcomes. We find that the average capital stock has a first-order effect on equilibrium outcomes.

A first-order approximation delivers the coefficients on two values that affect the deterministic steady-state: one coefficient captures the reaction of an agent’s choice as a function of her own capital holdings, the other coefficient the reaction to any other agent’s capital holdings. The symmetry plays again a crucial role. The reaction to changes in capital holdings by any other agent are identical. Individual choices are thus, to a first-order approximation, a function of individual capital holdings and the aggregate capital of others. However, the fact that the first-order effect is a function of individual and aggregate capital does not imply that heterogeneity does not play a role.

A second-order approximation entails four coefficients. There is the double change in individual capital holdings that will, together with the double change in other agents’ capital, lead the average choice to be a function of the variance of the distribution of capital. However, there are also coefficients on the change of individual capital interacted with the change in another agent’s capital and the interaction of another agent’s capital with a third agent’s capital. In other words, the variance of the distribution of capital is not the only statistic that plays a role.

We confirm previous research in finding that the impact of heterogeneity has an effect on the steady-state level of capital. Since agents face idiosyncratic risk, they respond by building up a buffer stock of precautionary savings. With aggregate risk this channel is enforced. Due

the utility specification featuring constant relative risk aversion, agents increase their capital holdings further due to uncertain returns to capital.

Furthermore, we find that the comovement of an individual's capital with the average capital of others plays as much of a role as the variance of the wealth distribution. Depending on the specific model, its quantitative importance might dominate that of the variance. From an asset pricing perspective, the importance might not come as a surprise. The covariance of the return with marginal utility is the main driver of asset prices.

Our methodology gives us two new ways of testing the impact of heterogeneity. First, have a way of testing whether heterogeneity matters quantitatively. We solve the model and check the effect of a varying numbers of individuals on equilibrium outcomes. Second, our technique allows us to check which features of the model influence the impact of heterogeneity. In other words, we can check which modeling assumptions need to be altered in order to change the quantitative implications of heterogeneous agents.

This paper contributes to a growing literature on introducing heterogeneity into economic models. Therefore, we relate to several strands of research. One strand of the literature introduces idiosyncratic risk for individuals with incomplete markets but abstracts from aggregate risk. The most prominent examples in this literature are Bewley (1977) and Aiyagari (1994). In contrast, however, we want to focus on models with aggregate risk to study portfolio choice and financial market equilibrium.

Another strand of the literature has focused on idiosyncratic risk with aggregate shocks. First, in special cases one might be able to find closed-form solutions as in Heathcote, Storesletten and Violante (2009) and Moll (2009). Another promising idea is to use a multiplier approach to characterize features of the distribution of state variables across the population as in Chien, Cole and Lustig (2010) and Chien and Lustig (2010). Other papers make simplifying assumptions on the number of agents and the number of possible shocks, as in Dumas and Lyasoff (2010). Special cases with closed-form solution can be used as a starting point for the expansion.

Most of the literature, however, is concerned with approximations. One idea is to replace the distribution of wealth by aggregate wealth only when calculating the equation of motion for aggregate variables. This method was developed in Krusell and Anthony A. Smith (1998) and inspired methods in the subsequent literature, for example in Storesletten, Telmer and Yaron (2007), Gomes and Schmid (2010), and Favilukis, Ludvigson and van Nieuwerburgh (2011). Alternatively, one might work with a limited history of shocks as in Lustig and van Nieuwerburgh (2010). Since we are particularly interested in the effect of distributions on

equity prices and the effect on new financial securities, this approximation method is not appropriate for this research project.

Recently, alternative solution method for models with heterogeneous agents have been developed in Haan, Judd and Juillard (2010), Haan (2010), Judd, Maliar and Maliar (2009), and Judd, Maliar and Maliar (2010). Haan, Judd and Juillard (2010) and Haan (2010) parameterize the distribution of state variables. Feng, Miao, Peralta-Alva, and Santos (2009) approximate the equilibrium on a lower-dimensional space. This paper develops a technique that does not require the specification of a class of distributions. Compared to Judd, Maliar and Maliar (2009) and Judd, Maliar and Maliar (2010), the method in this paper has the advantages that we can study as many agents as desired whereas the number of agents is limited in their method. Furthermore, our method applies to models with many state variables and choices for each individual. Furthermore, the usual differences between perturbation and projection apply: our method returns quasi-analytical expressions, allows to prove theorems, get intuition for the impact of parameter changes, and is fast and simple.

Our method builds on perturbation methods. These methods have been used in Jin and Judd (2002), Judd and Guu (2000), Judd (2002), Hassan and Mertens (2010), Mertens (2009), Fernández-Villaverde and Rubio-Ramírez (2006), Garlappi and Skoulakis (2010). Most recently, Mertens (2009) uses perturbation methods to study heterogeneity induced by private information. This solution method was also applied in Hassan and Mertens (2010).

This paper, however, is not the first paper that attempts to use perturbation methods to analyze general equilibrium models with substantial heterogeneity. An alternative idea to the one in this paper has been explored by Preston and Roca (2007) and Kim, Kollmann and Kim (2010). This work starts by restricting the state space from the outset. Only certain moments of the wealth distribution enter the state space. However, it is easy to construct examples where a limited set of moments lead to poor approximations. Instead, this paper is the first to recognize the symmetry of the problem and build a solution method that exploits it. No limitations on the state space are required.

2 An economy with heterogeneous agents

This section sets up a standard RBC model except for the fact that households are subject to idiosyncratic risk. Each household receives shocks to their individual labor income each period for which no tradable asset exists. We solve two versions of the model. In the first setup, the only tradable asset is a claim to capital which is risky due to aggregate shocks

to total factor productivity. The second version augments the first setup with an additional risk-free savings technology which exists in zero net supply. Agents then solve a portfolio choice problem on top of the consumption-savings and labor supply decision.

2.1 Households

A finite number I of households lives for an infinite number of periods. Households are each endowed with one unit of time which they devote towards labor inelastically. While they are identical in their preferences, households differ in their productivity. Each period, they receive an idiosyncratic shock to their productivity and thus their labor income. Agents can only partially insure against this shock by holding saving to buffer the shocks. A tradable contract consists of claims to capital which is risky due to aggregate productivity shocks. A second tradable contract is a bond and thus a claim to a risk-free payoff. Households derive utility from consuming a single consumption good which they pay for from their income and savings.

Each household builds rational expectations and chooses streams of consumption, labor supply, and capital holdings so as to maximize the expected utility function

$$\max_{c_t^i, k_t^i} \sum_{t=0}^{\infty} \beta^t E_0 [u_c(c_t^i) - u_k(k_t^i, b_t^i)] \quad i = 1, \dots, I \quad (1)$$

where β is the time discount factor, c_t^i is household i 's consumption choice in period t , k_t^i are household i 's capital and b_t^i bond holdings.

The utility function is comprised of two additively separable parts. The first part is a standard utility function defined over individual consumption. To pin down the functional form, we impose constant relative risk aversion over consumption. The latter part is non-standard and thus merit some more attention.

The second term in the utility specification is a cost of deviating from a target level of capital and bond holdings. It serves three purposes. First, it pins down the distribution of capital and bond holdings in the deterministic steady-state by penalizing deviations from a target level of capital and bonds which we define later.¹ Second, it imposes an endogenous limit to borrowing. And third, our penalty function rules out Ponzi schemes. However, it does not impose a no-Ponzi condition ex ante. Rather, agents are permitted to run a Ponzi scheme but will not choose them in equilibrium due to the penalty function. It is thus closely related

¹The same goal can be achieved by endogenizing the discount factor.

to endogenous borrowing constraints. An alternative approach would be to fix the budget constraint through an exogenous borrowing constraint.

The functional form we impose on the second term reflects both the endogenous borrowing constraint and the determination of the distributions of capital and portfolio choice in the deterministic steady-state

$$u_k(k_t^i, b_t^i) = \nu_1 \frac{1}{(k_t^i + b_t^i - \bar{k})^2} + \nu_2 (k_t^i - \bar{k})^2 + \nu_3 (b_t^i)^2 + \nu_4 (k_t^i + b_t^i) \quad (2)$$

where $\bar{k}$ denotes the target level of capital and ν parameters for the penalty function.

This specification takes on positive values whenever capital is not at its target level but particularly high values when agents want to borrow against future endowments. Deviations from the target level will this impose a cost to agents. We impose the restriction $\nu_4 = \nu_1 \frac{2}{(\bar{k} - \underline{k})^3}$ to ensure that the penalty function has its global minimum at $\bar{k}$. In particular, the derivative with respect to capital vanishes at its target level $\bar{k}$.

The first part of the penalty function imposes the asymmetry between borrowing and saving. Borrowing is heavily penalized in order to build in an endogenous borrowing constraint. Note that the penalty function only takes effect when going to at least a second-order expansion. The asymmetry will arise only for higher orders. The second part is responsible for fixing the steady-state distribution of capital while the third part determines portfolio choice at the steady-state. The last term is merely a technical matter. It ensures that the choice between capital and bonds is well defined in the deterministic steady-state. Once we move to a stochastic economy, capital and bonds are no longer perfect substitutes and we set the parameter ν_3 to zero. In that sense, the third term reflects a mere technicality and is not present in our solution of the stochastic model.

Households maximize utility subject to their budget constraint

$$c_t^i + k_{t+1}^i + b_{t+1}^i = (1 + r_t^k)k_t + (1 + r_t^b)b_t + w_t(1 - \psi_t^i). \quad (3)$$

The rates of return are denoted by r_t^k for capital and r_t^b for bonds, and wages by w_t .

There is a shock to individual productivity denoted by ψ_t^i which is independent and identically distributed across households. Agents cannot directly insure against this shock since there is no tradable asset contingent on individual labor productivity. Since we work with a finite number of households, we cannot apply a law of large number to ensure that the individual labor productivity shock ψ_t^i is purely idiosyncratic. However, we can decompose the shock into an aggregate component which is due to the finite number of agents and an idiosyncratic

component. The standard deviation of the aggregate component is a decreasing function in the number of agents.

To keep a concise notation, we introduce capital case letters for aggregate quantities of consumption, capital, bonds, and labor productivity

$$C_t = \sum_{i=1}^I c_t^i \quad K_t = \sum_{i=1}^I k_t^i \quad B_t = \sum_{i=1}^I b_t^i \quad \Psi_t = \sum_{i=1}^I \psi_t^i.$$

2.2 Technology

Aggregate capital and effective labor enter the production process for the single consumption good. A Cobb-Douglas function describes the production process that uses capital and labor as its inputs. Aggregate productivity is stochastic. The parameter for the production function is given by α which leads to a functional form for output given by $Y = f(K, L, z) = e^z K^\alpha L^{1-\alpha}$ where e^z denotes the shock to total factor productivity, K aggregate capital and L aggregate effective labor demand. Total factor productivity is log-normally distributed.

The logarithm of total factor productivity follows an AR(1) process

$$z_{t+1} = \phi_z z_t + \varepsilon_t \tag{4}$$

where the parameter ϕ_z determines the degree of mean reversion in total factor productivity.

Given the constant returns to scale of the production function, wages and dividends pay their marginal product

$$\begin{aligned} r_t^k &= \alpha e^{z_t} K_t^{\alpha-1} L_t^{1-\alpha} \\ w_t &= (1 - \alpha) e^{z_t} K_t^\alpha L_t^{-\alpha} \end{aligned} \tag{5}$$

Due to the shocks to total factor productivity, the returns to capital and wages are risky. Only one shock is driving the uncertainty of proceeds for both factor inputs. As a result, labor income and returns to capital are positively correlated.

2.3 Definition of equilibrium

This section lays out the definition of equilibrium for two versions of the economy. The first economy allows for trading only one asset, claims to capital. The second economy introduces a risk-free bond as an additional asset in zero net supply.

Both versions of the economy introduce heterogeneity in the simplest form. Household are affected by an individual shock to their labor productivity. In an economy with complete markets, the idiosyncratic part of the risk would not affect equilibrium since it could be traded away perfectly by households. Idiosyncratic risk would thus neither affect aggregate variables nor consumption choices.

In our setup, however, asset markets are restricted to a limited set of assets. In particular, there are no assets available that pay contingent on individual labor productivity. Agents thus cannot directly insure against their idiosyncratic risk. They can only partially insure by building up precautionary savings.

The second economy features heterogeneity in the same manner but introduces an additional risk-free security. Hence we have to deal with portfolio choice in this case. Asset trading takes place in stock and bond trading where bonds are available in zero net supply.

In either economy, the aggregate resource constraint is given by the following equation

$$C_t + K_{t+1} - (1 - \delta)K_t = Y_t \quad (6)$$

which shows how current output and depreciated capital can be used for consumption or next period's capital stock. The derivation follows from the households' budget constraints, the first-order conditions of the firm, and a zero net aggregate supply of bonds.

2.3.1 Version I

For the first version of our economy, we impose zero bond holdings by all agents in the economy. The utility function is thus independent of bond holdings. Furthermore, there is only one consumption savings decision by households and no portfolio choice decision needs to take place.

The individual productivity shock is independent and identically distributed and follows a stochastic process of the form

$$\psi_{t+1} = \phi + \phi_\psi \psi_t + \phi_\theta(\psi_t)\theta_{t+1}. \quad (7)$$

The parameter ϕ_ψ governs the degree of persistence in the evolution of the shock. ϕ is a parameter and ϕ_θ governs the standard deviation of the shock. θ is white noise with unit variance.

Given the setup of the economy, each agent has to determine the solution to her utility

maximization problem by picking the optimal amount of consumption. Optimal choices are the solution to the following optimality condition:

$$u'_c(c_t^i) = \beta E_t \left[(1 + r_{t+1}^k) u'_c(c_{t+1}^i) - u_k^{(1)}(k_{t+1}^i, b_{t+1}^i) \right] \quad (8)$$

If we set all parameters ν to zero in the utility function, this optimality condition results in the standard Euler equation. With the parameters ν being non-zero, we impose an endogenous borrowing constraint. In the Euler equation (8), there will be additional terms which stem from the effect of consumption this period on the capital stock tomorrow. Since the penalty function, which imposes the borrowing constraint in the following period, depends on next period's capital stock, a marginal unit of consumption today marginally increases the expected penalty. It thus has a distortionary effect on the consumption savings decision in the current period which are governed by the extra terms in the Euler equation when plugging in the functional form of utility.

To define the state space in a concise manner, we introduce the following notational convention. We denote vectors by a small case bold letter and a matrix by an upper case bold font letter.

The state space of this economy consists of the set of individual capital holdings of each of the I households as well as their productivity $\mathbb{X}^1 = \mathbb{R}^{I \times 2}$ as well as the set of possible states of aggregate productivity $\mathbb{Z}^1 = \mathbb{R}$. Furthermore, we clarify whether the state space belongs to a stochastic or deterministic economy by making each function dependent on $\sigma \in [0, 1]$, the standard deviation of shocks. When the standard deviation of shocks σ equals one, we refer to the stochastic economy. The deterministic counterpart is denoted by $\sigma = 0$. An element of the state space is denoted by $(\mathbf{X}, \mathbf{z})$.

Definition 1 (Definition of Equilibrium for Economy I) *An equilibrium of the first version of the economy is a function $c(\mathbf{X}, \mathbf{z}, \sigma)$ and prices $r(\mathbf{X}, \mathbf{z}, \sigma)$ and $w(\mathbf{X}, \mathbf{z}, \sigma)$ such that (i) each household solve his optimization problem (1) with bond holdings being restricted to zero subject to the budget constraint (3), (ii) prices r and w are competitive, and consumption choices c are consistent with the aggregate law of motion.*

This definition shows that the number of equations defining the equilibrium is huge. Each household must solve the individual optimization problem. Thus I first-order conditions along with their budget constraints are necessary to pin down the equilibrium where wages and rates of return are equated to the corresponding marginal product.

Definition 2 (Deterministic Steady-State) *A deterministic steady-state is a point in the*

state space $(\mathbf{X}_0, \mathbf{z}_0, 0)$ such that each household's first-order conditions are satisfied, consumption is constant, and capital does not change.

2.3.2 Version II

The second version of the economy makes the same assumptions as the first version but introduces an additional risk-free asset on top of the existing claims to capital. Agents have thus one more choice to make. Besides the consumption-savings decision, they also have to decide on the portfolio share that determines which fraction of their savings is allocated towards capital.

We follow the specifications of the first version of the economy to keep the solution as comparable as possible. The shock to individual labor productivity thus follows the same stochastic process.

The difference lies thus merely in the existence of a tradable risk-free asset. An additional optimality condition determines the choice of bond holdings for each agent

$$u'_c(c_t^i) = \beta E_t \left[(1 + r_{t+1}^b) u'_c(c_{t+1}^i) - u_k^{(2)}(k_{t+1}^i, b_{t+1}^i) \right] \quad (9)$$

The state space of this economy consists of the set of individual capital holdings of each of the I households, their bond holdings, and the level of their individual productivity, $\mathbb{X}^2 = \mathbb{R}^{I \times 3}$. Furthermore, we need to keep track of aggregate productivity $\mathbb{Z}^2 = \mathbb{R}$. The standard deviation of shocks follows the same notation as before

Definition 3 (Definition of Equilibrium for Economy II) *An equilibrium of the second version of the economy is a consumption function $c(\mathbf{X}, \mathbf{z}, \sigma)$, portfolio choice decision determined by $b(\mathbf{X}, \mathbf{z}, \sigma)$, and prices $r(\mathbf{X}, \mathbf{z}, \sigma)$ and $w(\mathbf{X}, \mathbf{z}, \sigma)$ such that (i) households' optimality conditions (8) and (8) are satisfied along with the budget constraint (3), (ii) prices r and w are competitive, and choices for consumption c and bond holdings b are consistent with the aggregate law of motion.*

3 Numerical method

This section describes a new solution method to solve models with incomplete information and substantial heterogeneity. The model of the previous section serves as an example of this class. We lay out the general structure of the setup and apply perturbation methods to it.

We derive a higher-order approximation to the true solution. As we explain in this section, a simple linearization of the model does not provide a desirable outcome since the resulting approximation is certainty equivalent. Heterogeneity would not impact aggregate variables and prices.

3.1 Overview over the solution method

Before getting into details, we provide a summary of the solution method. The rest of section is concerned with laying out the specifics of each of the steps mentioned here. The summary serves the purpose of giving an overview to the reader who only wants to get the gist of our method.

The solution method starts with listing equilibrium conditions which entail optimality conditions, budget constraints, and equations of motion. In a first step, we shut down all uncertainty and find the deterministic steady-state. The economic model has to be such that a unique steady-state exists. We discuss ways to ensure a proper point of expansion. Importantly, all economic agents should be identical at this point in the state space.

Then we expand the system around this deterministic steady-state in all state variables which include all individual states of all I agents. Since we put particular emphasis on the number of agents I being very large, we would have to expand with respect to an extremely large number of variables. The key step to avoid this infeasible effort is to recognize the symmetry of the problem. Expanding an optimal choice of, say, agent 1 with respect to agent 2's state variables delivers just the same coefficients as the expansion with respect to any other agent's state variables. At the same time, due to the symmetry at the deterministic steady-state, the policy function for agent 1 is just the same as that of any other agent.

We build a higher-order expansion of the system in all state variables $\mathbf{X}$ and $\mathbf{z}$ before moving from the deterministic to the stochastic economy. A higher-order expansion is crucial for the study of heterogeneity as we explain below. Uncertainty manifests itself in a change in the steady-state level and response to state variables.

From the distribution of state variables, we can compute the distribution of all choices and economic variables through a nonlinear change of variables.

3.2 General Setup

This paper's solution method handles competitive equilibria as well as dynamic programming problems. To demonstrate the generality of the solution method, we define a matrix of

individual state variables $\mathbf{X} \in \mathbb{R}^{C \times I}$ and a vector of aggregate shocks $\mathbf{z} \in \mathbb{R}^Z$ where S denotes the number of individual state variables and S the number of aggregate shocks. The state space is thus extremely high dimensional. We write the equilibrium conditions (in the case of a competitive equilibrium) or the Bellman equation (in the case of a dynamic programming problem) along with the equation of motions, market clearing conditions, and budget constraints in the general form

$$\begin{aligned} E_t [F(\mathbf{X}_t, \mathbf{z}_t, \mathbf{P}_t, \mathbf{X}_{t+1}, \mathbf{z}_{t+1}, \mathbf{P}_{t+1})] &= 0 \\ \mathbf{X}_{t+1} &= G(\mathbf{X}_t, \mathbf{z}_t, \mathbf{P}_t). \end{aligned} \tag{10}$$

$\mathbf{P}_t \in \mathbb{R}^{S \times I}$ denotes the matrix of choices for all agents in period t . F and G are functionals operating on the high-dimensional space spanned by state variables and choices.

We impose additional structure on the functionals F and G to ensure symmetry at the deterministic steady-state. We require the equations in F and G that determine the choices $\mathbf{p}^i$ to be identical for all agents i . This way, we guarantee the existence of a deterministic steady-state and that the policy functions are identical for all agents at this point.

The notation quickly becomes cumbersome in our setting. In our economy of the previous section, the optimality condition requires to specify marginal utility. To see the structure of marginal utility of consumption in period t expressed in more detail, we first recognize that marginal utility $u'_c(c_t^i) = \tilde{u}(\mathbf{X}, \mathbf{z}, \mathbf{p}^i)$ is a function of the state of the economy and choice variables $\mathbf{p}^i$. We could use consumption as one of the choice variables but the more general case is discussed here. For example, capital and bond holdings could be the state variables in which case consumption would have to be computed via the budget constraint from these choices. Now we can demonstrate more clearly how much information we need to compute marginal utility.

$$\tilde{u} \left(\begin{pmatrix} x_1^1 & x_1^2 & \dots & x_1^N \\ x_I^1 & x_I^2 & \dots & x_I^N \end{pmatrix}, \begin{pmatrix} z_1 \\ \vdots \\ z_Z \end{pmatrix}, \begin{pmatrix} \mathbf{p}_1^i \left(\begin{pmatrix} x_1^1 & x_1^2 & \dots & x_1^N \\ x_I^1 & x_I^2 & \dots & x_I^N \end{pmatrix}, \begin{pmatrix} z_1 \\ \vdots \\ z_Z \end{pmatrix}, \sigma \right) \\ \vdots \\ \mathbf{p}_P^i \left(\begin{pmatrix} x_1^1 & x_1^2 & \dots & x_1^N \\ x_I^1 & x_I^2 & \dots & x_I^N \end{pmatrix}, \begin{pmatrix} z_1 \\ \vdots \\ z_Z \end{pmatrix}, \sigma \right) \end{pmatrix} \right) \tag{11}$$

In our setup, the exogenous state variables in the next period do not depend on last period's choice. However, the extension to the more general case is straightforward.

3.3 Range of Applications

As the previous subsection shows, the setup is quite general. Here, we clarify the two main requirements for the models we can solve.

First, we require the model to imply smooth policy functions. We apply perturbation methods to the problem which build a Taylor series expansion of the optimal policies around a deterministic steady-state. For the Taylor series to converge to the true solution, the operator defining the optimality conditions and the optimal policy need to be smooth which means they need to possess many derivatives. The more derivatives exist, the better the approximation can be. In most economic problems, optimal policies are indeed analytic which means they possess all derivatives.

If the derivatives of the solution exist within some range, the Taylor series converges within this radius. In particular, if the solution is analytic in some range, we obtain a global approximation within that region.

Second, we require the existence of a deterministic steady-state for which we can solve. This steady-state can either be obtained by having agents being symmetric at this point or a closed-form solution with heterogeneity exists for this special case. Imposing identical agents in that steady-state does not mean that we cannot allow for heterogeneity. We can expand in the dimension in which agents are heterogeneous, for example with respect to their risk aversion.

For most purposes, we impose complete symmetry on the functional F . For example, in our competitive equilibrium all agents' first-order conditions are the same. If we denote $\mathbf{X}^{i \leftrightarrow j}$ the matrix of state variables where we exchange the state variables of agent i with agent j and vice versa and the same for policy functions $P^{i \leftrightarrow j}$, then we can express the symmetry requirement as

$$F_i(\mathbf{X}_t, \mathbf{z}_t, \mathbf{P}_t, \mathbf{X}_{t+1}, \mathbf{z}_{t+1}, \mathbf{P}_{t+1}) = F_j(\mathbf{X}_t^{i \leftrightarrow j}, \mathbf{z}_t, \mathbf{P}_t^{i \leftrightarrow j}, \mathbf{X}_{t+1}^{i \leftrightarrow j}, \mathbf{z}_{t+1}, \mathbf{P}_{t+1}^{i \leftrightarrow j}). \quad (12)$$

As a consequence, we can restrict ourselves to building an expansion for only one agent and thus expanding only the optimality condition for only one agent.

3.4 Deterministic steady-state

To solve for the deterministic steady-state, we shut down all shocks. Since agents are heterogeneous only with respect to their idiosyncratic labor income shocks in our model, the

deterministic steady-state features identical agents and no heterogeneity.

Two difficulties arise with respect to the deterministic steady-state by setting their standard deviation to zero. For this deterministic economy, we solve the equilibrium conditions for their steady-state variables and choices. Even when there is only one state variable, capital, for each agent in our model, the deterministic steady-state is not well defined. To see this, imagine a deterministic version of the economy in the previous section. The condition that ensures a steady-state is given by

$$1 - \delta + f_K(K_t, L_t, 0) = \frac{1}{\beta}$$

since the Euler equation then implies constant consumption. This equation, however, does not depend on the distribution of wealth. *Any* distribution satisfying the above steady-state restriction is thus a deterministic steady-state.

In our setup, agents obtain disutility from deviating from a target level of capital. This penalty function implies that there is a unique distribution of capital in the deterministic steady-state, irrespective of how small the penalty from deviating from the target is. We choose the target level of capital to match the level in the deterministic steady-state. Our starting point for the expansion is thus an equal distribution of capital across the population.

In the economy with portfolio choice, a second problem arises. In the deterministic steady-state of the problem, capital and bonds become perfect substitutes. Both assets pay a riskless return and, by a no arbitrage condition, the two have to pay the same return. The perfect substitutability implies an indeterminacy of the portfolio in the deterministic steady-state.

Therefore, we introduce an analogous trick to the one we used for the deterministic steady-state distribution of capital. We introduce a small penalty for deviating from the optimal portfolio. Any long or short positions in bonds will lead to a reduction in utility.

When introducing the penalty function in our program, we keep the coefficients ν as free parameters. After moving from the deterministic to the stochastic economy, we can let parameters that eliminate the unit root in the deterministic steady-state go to zero. In this sense, fixing the steady-state distribution is a mere technicality that is necessary to solve for the proper expansions of the deterministic economy.

3.5 Higher-order expansion

Computing a higher-order Taylor series for the equilibrium policy functions, quantities, and prices is essential to our solution method. There are two reasons for it. First, heterogeneity

manifests its impact only in higher-order terms and second, so does stochasticity.

A first-order approximation implements standard linearization. Due to this linearity, heterogeneity does not affect equilibrium outcomes because under these rules, the average choice is the choice of the average person. The impact of heterogeneity on equilibrium outcomes thus only enters through higher-order terms, starting with a second-order approximation.

For the same reason, stochasticity impacts equilibrium only through higher-order terms. The first-order approximation is certainly equivalent while higher-order terms add the effects of variance, skewness, and higher moments. Mathematically, one can get this result by building an expansion with respect to the standard deviation of shocks as we do in the next section. There it shows that the first-order term in the expansion does not impact equilibrium outcomes.

A Taylor series expansion of high order serves as a good approximation to equilibrium outcomes. Since we impose conditions to ensure an analytic solution, our algorithm is asymptotically valid. For analytic functions, the approximation will converge with higher-order terms within the radius of convergence. An infinite-order Taylor series fully recovers the true solution. Thus we know that eventually we get very close to the true solution. In practice, of course, we have to truncate the Taylor series. But the stage at which we stop can be endogenous to the accuracy of the solution. To compute high-order derivatives, a high precision arithmetic might be necessary as Swanson, Anderson and Levin (2005) points out.

We denote a vector of indices $\mathbf{i} = (i_1, i_2, \dots, i_N)$ with a bold letter. Furthermore, we refer to the sum of all the entries of such a multi-index by $|\mathbf{i}| = \sum_{\iota=1}^N i_\iota$ and to the product of all entries as $\mathbf{i}! = \prod_{\iota=1}^I i_\iota$. For a collection of multi-indices, we let the sum of all entries be denoted by the sum of its component vectors $\|\mathbf{I}\| = \|(\mathbf{i}^1, \mathbf{i}^2, \dots, \mathbf{i}^I)\| = \sum_{\iota=1}^I |\mathbf{i}^\iota|$ and the product of all entries by $\mathbf{I}! = \prod_{\iota=1}^I \mathbf{i}^\iota!$.

Let $\mathbf{P}^{(\mathbf{I}, \mathbf{j}, \mathbf{k})}$ denote the derivatives of the choices $\mathbf{P}$ of the respective orders. Order $\mathbf{i}^1 = (i_1^1, \dots, i_I^1)$ describes the derivatives in the directions $x_{i_1^1}^1$ to $x_{i_I^1}^1$. We also introduce a convenient notation for exponentials of the collection of state variables

$$\mathbf{X}^{\mathbf{I}} = \begin{pmatrix} (X_1^1)^{i_1^1} & \dots & (X_S^1)^{i_S^1} \\ \vdots & & \vdots \\ (X_1^I)^{i_1^I} & \dots & (X_S^I)^{i_S^I} \end{pmatrix} \quad \text{and} \quad \mathbf{z}^{\mathbf{j}} = \begin{pmatrix} (z_1)^{j_1} \\ \vdots \\ (z_Z)^{j_Z} \end{pmatrix}$$

Once we know the derivatives at a specific point, we can recover the choice variable of the

ϖ 's choice of agent i as

$$(\mathbf{P}_{\varpi}^i)^{(\mathbf{l}, \mathbf{j}, \mathbf{k})} = \sum_{o=1}^{\infty} \sum_{\|\mathbf{l}\| + \|\mathbf{j}\| + \|\mathbf{k}\| = o} \frac{1}{\mathbf{l}! \cdot \mathbf{j}! \cdot \mathbf{k}!} \left. \frac{\partial \mathbf{P}_{\varpi}^i}{\partial \mathbf{X} \partial \mathbf{z} \partial \sigma} \right|_{(\mathbf{X}^0, \mathbf{z}^0, 0)} (\mathbf{X} - \mathbf{X}^0)^{\mathbf{l}!} \cdot (\mathbf{z} - \mathbf{z}^0)^{\mathbf{j}!} \cdot \sigma^{\mathbf{k}}. \quad (13)$$

Perturbation methods tell us how to compute these derivatives at the deterministic steady-state. We apply the implicit function theorem and take derivatives of equilibrium conditions (10). By the chain rule, we get equations for the derivatives of the policy function at the deterministic steady-state.

$$\left. \frac{dF_i(\mathbf{X}_t, \mathbf{z}_t, \mathbf{P}_t, \mathbf{X}_{t+1}, \mathbf{z}_{t+1}, \mathbf{P}_{t+1})}{dx_1^1} \right|_{(\mathbf{X}^0, \mathbf{z}^0, 0)} = \frac{\partial F_i}{\partial x_1^1} + \frac{\partial F_i}{\partial \mathbf{P}_t} \frac{\partial \mathbf{P}_t}{\partial x_1^1} + \frac{\partial F_i}{\partial \mathbf{X}_{t+1}} \frac{\partial \mathbf{X}_{t+1}}{\partial x_1^1} + \frac{\partial F_i}{\partial \mathbf{z}_{t+1}} \frac{\partial \mathbf{z}_{t+1}}{\partial x_1^1} + \frac{\partial F_i}{\partial \mathbf{P}_{t+1}} \frac{\partial \mathbf{P}_{t+1}}{\partial x_1^1} \quad (14)$$

Now we use the fact that all the partial derivatives of F_i are known. They are simply the derivatives of the equilibrium conditions. When evaluating them at the deterministic steady-state, the only remaining variables in the differentiated equilibrium conditions are the derivatives of the optimal policies $\mathbf{P}$ at the deterministic steady-state. These are the coefficients in the Taylor series of the optimal policies. We differentiate the first-order conditions of each agent with respect to each state variable and thus obtain as many equations as we have state variables. We need one coefficient for each state variable in the first-order term. Thus we have as many equations as unknowns. Once we solve for them, we obtain an approximation to the optimal policy from (13).

The key innovation of this paper lies in recognizing of the symmetry of the problem. In principle, we would have to start with agent one, differentiate his first-order conditions with respect to each agent's state variables, move to agent two and so on. However, we do not have to go through this entire process. There are two crucial ways in which the problem is symmetric.

First, we only have to take the derivative with respect to state variables of agents that are different from the view point of that household whose first-order condition we differentiate. For example, when differentiating agent one's first-order condition, there are only two different agents in the first order expansion. There is agent one himself and there is any other agent. Since all other agents are identical, the derivative with respect to agent two's state variable is the same as the derivative with respect to agent three's state variable and so on. Thus it suffices to solve for the reaction of agent one's policy with respect to *any other* state variable. As a result, our expansion with respect to the first set of state variables $\mathbf{X}_1$, we only have to

keep track of two numbers.

Second, we only have to take all derivatives for one agent. It suffices to differentiate the first-order conditions for agent one. Agent two's first-order conditions look identical and thus lead to the same result. The symmetry here is that agent one's response to a marginal increase in agent two's state variable is the same as agent two's response to a marginal increase in agent one's state variable. This carries over to all derivatives.

The first-order term corresponds to standard linearization. Exploiting this symmetry, there are two coefficients to be computed for the first-order term for each state variable. One coefficient returns the change in policy of an agent in response to a change in her own wealth. The second coefficient asks for this agent's reaction in response to a change in the state variable by somebody else.

For the second order term, the system becomes slightly more complex. There we aim at solving for the quadratic terms in the expansion. For each state variable, we have to compute four values: an agent's change in response to a change in her state variable to a change in response to her wealth, an agent's change in response to a change in her state variable to a change in response to somebody else's state variable, an agent's change in response to a change in somebody else's state variable to a change in response to this person's state variable, and an agent's change in response to a change in somebody else's state variable to a change in response to a third person's state variable. Although the economic interpretation is tedious, solving the system of equations is straightforward. From the second order on, the system of equations for the unknown coefficients is linear.

As a direct consequence, our equation of motion is not only a function of an aggregate statistic of state variables but the entire distribution as the theory would tell us. In this respect we deviate from some of the previously proposed solution techniques. Another idea mentioned in the literature suggests to include further moments of the distribution to approximate the equation of motion with a small state space. Our solution method does something similar but with an important difference.

Our solution method includes the moments of the distribution of state variables with every order which have the most impact. In this sense, the solution method proposes a set of approximating statistics with which to approximate policy functions. As an additional feature, a better approximation simply adds moments to the previous approximation without the necessity to recompute previous approximations.

So far, the perturbation method we discussed solves for a solution of the deterministic economy (although equation (13) allows for stochastic terms). The next section discusses the

effect of stochasticity on the approximation.

3.6 Uncertainty

Having obtained a high order approximation of the deterministic economy, we move towards its stochastic counterpart. We accomplish the transition by varying the standard deviation of shocks.

Taking a first-order expansion with respect to the standard deviation of shocks produces coefficients that are all zero. The reason lies in the fact that the first-order expansion of the standard deviation introduces shocks only into the linearized economy. This first-order approximation of the system, however, is certainty equivalent and the effect of uncertainty does not play a role.

Only through second- and higher-order terms do we recover the solution to the stochastic system. The second-order term introduces shocks into the quadratic economy. This approximation is no longer certainty equivalent and uncertainty takes effect. To be more precise, the second-order term introduces the effect of the variance of shocks, the third-order term recovers the reaction to skewness, and so on.

We can interpret the way uncertainty enters the equilibrium as effectively altering the coefficients in the Taylor series. Building the expansion with respect to the standard deviation of shocks effectively alters the coefficients to the Taylor series of the deterministic system. This can be most easily seen if we look at a univariate problem. If we expand the deterministic version of a system $h(y(x, \sigma = 0), x, \sigma = 0) = 0$, we get the solution

$$y(x, 0) \approx \sum_{i=0}^I \frac{1}{i!} y^{(i,0)}(\bar{x}, 0) (x - \bar{x})^i. \quad (15)$$

The stochastic version has an expansion

$$\begin{aligned} y(x, \sigma) &\approx \sum_{i=0}^I \sum_{j=0}^J \frac{1}{(i+j)!} h^{(i,j)}(\bar{x}, 0) (x - \bar{x})^i \sigma^j \\ &= \sum_{i=1}^I \frac{1}{i!} \left(\sum_{j=0}^J \frac{i!}{(i+j)!} h^{(i,j)}(\bar{x}, 0) \sigma^j \right) (x - \bar{x})^i. \end{aligned} \quad (16)$$

The rearrangement demonstrates that the expansion of the stochastic system looks just like the deterministic system except that the coefficients contain a “correction term” for the

stochasticity of the function.

We can see this term graphically as depicted in figure 1. In the second-order, the function shifts while the third-order term would also tilt the function while even higher orders change its curvature.




Figure 1: Perturbation methods build an approximation in state variables around the deterministic steady-state (thick solid line). The expansion with respect to the standard deviation shifts (second order) and tilts this line (third order).

3.7 Distribution of Equilibrium Variables

Given our approximation method, we can compute the distribution of any equilibrium outcomes or nonlinear functions thereof. We therefore combine our perturbation method with a non-linear change of variables.² For example, from capital and bond holdings we could compute portfolio weights or Sharpe ratios of individual portfolios.

Suppose we have some economic variable of interest which is a nonlinear function $H(\mathbf{X}, \mathbf{z}, \mathbf{P})$ of the state variables and choices, then we can approximate it with a Taylor series. The coefficients can be computed as follows

$$\frac{dH}{dx_s^i} = \frac{\partial H}{\partial \mathbf{X}} \frac{\partial \mathbf{X}}{\partial x_t^i} + \frac{\partial \mathbf{P}}{\partial \mathbf{X}} \frac{\partial \mathbf{X}}{\partial x_t^i} \quad (17)$$

and analogously for other state variables.

²Judd (2002), Fernández-Villaverde and Rubio-Ramírez (2006), and Mertens (2009) explain nonlinear changes of variables in conjunction with perturbation methods.

The computation of the coefficient is trivial once we make the observation that the first term is given by the derivative of H (which is given) and the second one has already been computed in the previous approximation. Hence there is nothing left to do except for assembling the Taylor series expansion.

Thus, computing the distribution of any variable of interest within the economy is not more intricate than computing the distribution of capital.

3.8 Accuracy

The solution method comes with a natural way to check for its accuracy. The equilibrium conditions are satisfied when our functional F satisfies the optimality condition. For any approximation, we can check the deviation from zero and thus the error in the optimality condition. Since we have asymptotic validity of the solution method, we specify a tolerance as a threshold for the error. Once the error is below the tolerance in some norm, we terminate the approximation process.

To get a meaningful measurement for the error, it makes sense to normalize the optimality conditions such that they are unit-free. For example, we should rewrite the Euler equation (8) in the form

$$u'_c(c_t^i) = \beta E_t \left[(1 + r_{t+1}^k) u'_c(c_{t+1}^i) - u_k^{(1)}(k_{t+1}^i, b_{t+1}^i) \right] \quad (18)$$

to avoid the error being in units of marginal utility.

This measurement provides a way to check for accuracy after adding an order of approximation. Thus one can decide at each step whether the approximation suffices the criteria or not.

As an additional benefit, there is no need to recompute previous orders after each step. The approximation method keeps previous coefficients unaltered when refining the solution.

4 Results

This section summarizes the findings for the particular economy of section 2. We first discuss the choice of functional forms and parameter values. Then we show the accuracy of the solution method before discussing findings with respect to the two versions of the economy.

4.1 Calibration

Most of the parameters and functional forms are standard in the literature. In large parts, there is a consensus on how to calibrate a real business cycle model. And the introduction of heterogeneity has predecessors in the literature. For comparability of our results, we aim at matching the same parameter combination.

We choose constant relative risk aversion as functional form for our utility specification defined over consumption. We set the coefficient of relative risk aversion to 2. The time preference factor is chosen to be 0.95.

We implement the penalty function for deviations from steady-state capital and bond holdings that imposes an endogenous borrowing constraint. Therefore, we set the parameter $\nu_1 = 3$ to ensure that the borrowing constraint receives a lot of weight. The parameter $\nu_2 = 0.01$ is set to a small number since it merely ensures that the steady-state is defined for the deterministic economy. Finally, we set $\nu_3 = 0.00001$ for the deterministic economy and to zero once we move to the stochastic version. To implement the borrowing constraint, we set $\underline{k} = -0.1$. As a consequence, for almost all cases, borrowing becomes prohibitively costly. The parameter $\bar{k}$ corresponds to the steady-state value of capital which is set such that the return on capital equals the reciprocal of the time preference factor in the deterministic steady-state.

The parameters governing the macroeconomic considerations do not vary much across different works in the literature. We set the capital share of output to α to 1/3 and the parameter for depreciation to 0.1.

For the shocks to aggregate productivity and shocks to individual labor productivity, we follow the calibration in Haan, Judd and Juillard (2010) via Kim, Kollmann and Kim (2010). We set the stochastic process for aggregate uncertainty to be

$$z_{t+1} = 0.25 + 0.75z_t + 0.00661\varepsilon_{t+1}.$$

Idiosyncratic shocks to labor income evolve according to

$$\psi_{t+1}^i = 0.4 + 0.55555\psi_t^i + (0.48989 - 0.28381\psi_t^i)\theta_{t+1}^i$$

The remaining question concerns the number of agents in the economy. As demonstrated when describing the solution method, the computing power required is the same for any number of individuals. To generate the results of this section, we set this number to 100.

4.2 Convergence

As mentioned in section 3, the numerical method leads to a natural check of the accuracy of the solution. We normalize the Euler equation by dividing by marginal utility on both sides as in equation (18).

Figure 2 plots the logarithm of the Euler equation error as a function of one agent's capital stock. This check for accuracy corresponds to the deterministic version of the economy.




Figure 2: Euler equation error for the deterministic version of our economy.

The deterministic steady-state satisfies the deterministic optimality conditions. Thus the Euler equation error is zero at this point. Thus the logarithm is at negative infinity.

Two observations stand out from this graph. First, we see convergence. The Euler equation error decreases for the interval. And second, the result approximates the solution not just locally but globally on a sizeable interval.

4.3 Impact of heterogeneity

Heterogeneous agent economies have very different dynamics from their representative agent counterparts. For example, consider the economy where, relative to the steady-state, agent one has a higher level of capital whereas all other agents are at their steady-state level. All agents but agent one will take the same decision today. This decision, however, will not be the steady-state choice. Next period, all agents' capital holdings will deviate from their steady-state level. Agent *one*, however, will still hold a different level of capital compared to all other agents. From this simple thought experiment, we see how the solution method helps us

in understanding the propagation of heterogeneity over time. This was a particularly simple example. But we can feed more realistic wealth distributions into the model and simulate their evolution.

The quantitative impact of heterogeneity across agents, however, is small due. Higher-order coefficients on the impact of heterogeneity, i.e. an agents reaction to a mean preserving spread in other agents' capital holdings, is very small.

Heterogeneity with aggregate risk increases the steady-state level of capital. This result is known from the previous literature. There are two reasons for it. First, idiosyncratic risk leads to precautionary savings on the part of households. Since households cannot trade claims contingent on their labor income, they try to partially insure against these shocks by building up a buffer stock of savings. Second, due to aggregate productivity shocks, holding capital is risky. There are two opposing forces. On the one hand, agents are risk averse and demand a higher risk premium for holding risky capital. Each unit should thus return a higher dividend which implies a higher marginal product of capital and thus a lower steady-state level of capital. On the other hand though, since returns to capital are risky, agents again respond by building up savings which implies a lower steady-state level of capital. With our utility specification of constant relative risk aversion, the latter effect dominates. Thus, heterogeneity with aggregate risk increases the steady-state level of capital.

We find an interesting result for the impact of heterogeneity. In a second order approximation, not only the impact of the variance of the distribution of capital determines equilibrium outcomes but also the comovement of an agent's capital with aggregate capital holdings of all other agents. This is akin to findings in the asset pricing literature where a price of the asset is determined by the covariance with individual marginal utility. Thus it comes as no surprise that this result plays an important role in determining the portfolio choice of each agent for the economy with portfolio choice.

5 Conclusion

In this paper, we presented a simple but nevertheless general solution method for incomplete market models with substantial heterogeneity. The method is built on perturbation methods which build a Taylor series expansion around a deterministic steady-state. This solution method is particularly useful for models with many state and choice variables. Generally, this idea can be employed not only to competitive equilibria but also to dynamic programming problems.

As an example, we solved a dynamic stochastic general equilibrium model with idiosyncratic shocks to labor income. We demonstrated the convergence properties for this particular example.

References

- Aiyagari, S. Rao**, “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 1994, *109*, 659–684.
- Bewley, Truman F.**, “The Permanent Income Hypothesis: A Theoretical Foundation,” *Journal of Economic Theory*, 1977, *16*, 252–292.
- Chien, YiLi and Hanno Lustig**, “The Market Price of Aggregate Risk and the Wealth Distribution,” *Review of Financial Studies*, 2010, *23(4)*, 1596–1650.
- , **Harold Cole, and Hanno Lustig**, “A Multiplier Approach to Understanding the Macro Implications of Household Finance,” *Review of Economic Studies*, 2010, *78*, 199–234.
- Dumas, Bernard and Andrew Lyasoff**, “Incomplete Market Equilibria Solved Recursively on an Event Tree,” *mimeo, Boston University School of Management*, 2010.
- Favilukis, Jack, Sydney Ludvigson, and Stijn van Nieuwerburgh**, “The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium,” *mimeo, NYU Stern School of Business*, 2011.
- Feng, Zhigang, Jianjun Miao, Adrian Peralta-Alva, , and Manuel S. Santos**, “Numerical Simulation of Nonoptimal Dynamic Equilibrium Models,” *mimeo, Boston University*, 2009.
- Fernández-Villaverde, Jesús and Juan F. Rubio-Ramírez**, “Solving DSGE Models with Perturbation Methods and a Change of Variables,” *Journal of Economic Dynamics and Control*, 2006, *30 (12)*, 2509–2531.
- Garlappi, Lorenzo and Georgios Skoulakis**, “Solving Consumption and Portfolio Choice Problems: The State Variable Decomposition Method,” *Review of Financial Studies*, 2010, *23(9)*, 3346–3400.
- Gomes, Joao F. and Lukas Schmid**, “Equilibrium Credit Spreads and the Macroeconomy,” *mimeo, Fuqua School of Business, Duke University*, 2010.
- Haan, Wouter J. Den**, “Comparison of Solutions to the Incomplete Markets Model with Aggregate Uncertainty,” *Journal of Economic Dynamics and Control*, 2010, *34*, 4–27.
- , **Kenneth L. Judd, and Michel Juillard**, “Computational Suite of Models with Heterogeneous Agents: Incomplete Markets and Aggregate Uncertainty,” *Journal of Economic Dynamics and Control*, 2010, *34*, 1–3.

- Hassan, Tarek A. and Thomas M. Mertens**, “The Social Cost of Near-Rational Investment,” *mimeo*, *NYU Stern*, 2010.
- Heathcote, Jonathan, Kjetil Storesletten, and Gianluca L. Violante**, “Consumption and Labor Supply with Partial Insurance: An Analytical Framework,” *mimeo*, *New York University*, 2009.
- Jin, Hehui and Kenneth L. Judd**, “Perturbation Methods for General Dynamic Stochastic Models,” *mimeo*, *Hoover Institution, Stanford University*, 2002.
- Judd, Kenneth L.**, “Perturbation Methods with Nonlinear Change of Variables,” *mimeo*, *Hoover Institution, Stanford University*, 2002.
- and **Sy-Ming Guu**, “The Economic Effects of New Assets: An Asymptotic Approach,” *mimeo*, *Hoover Institution*, 2000.
- , **Lilia Maliar, and Serguei Maliar**, “Numerically Stable Stochastic Simulation Approaches for Solving Dynamic Economic Models,” *NBER Working Paper 15296*, 2009.
- , —, and —, “A Cluster-Grid Projection Method: Solving Problems with High Dimensionality,” *NBER Working Paper 15965*, 2010.
- Kim, Sunghyun Henry, Robert Kollmann, and Jinill Kim**, “Solving the Incomplete Market Model with Aggregate Uncertainty Using a Perturbation Method,” *Journal of Economic Dynamics and Control*, 2010, *34*, 50–58.
- Krusell, Per and Jr. Anthony A. Smith**, “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 1998, *106 (5)*, 867–896.
- Lustig, Hanno and Stijn van Nieuwerburgh**, “How Much Does Household Collateral Constrain Regional Risk Sharing?,” *Review of Economic Dynamics*, 2010, *13 (2)*, 265–294.
- Mertens, Thomas M.**, “Excessively Volatile Stock Markets: Equilibrium Computation and Policy Analysis,” *mimeo*, *NYU Stern*, 2009.
- Moll, Benjamin**, “Creditor Rights, Inequality and Development in a Neoclassical Growth Model,” *mimeo*, *University of Chicago*, 2009.
- Preston, Bruce and Mauro Roca**, “Incomplete Markets, Heterogeneity and Macroeconomic Dynamics,” *mimeo*, *Columbia University*, 2007.

Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron, “Asset Pricing with Idiosyncratic Risk and Overlapping Generations,” *Review of Economic Dynamics*, 2007, *10*, 519–548.

Swanson, Eric, Gary Anderson, and Andrew Levin, “Higher-Order Perturbation Solutions to Dynamic, Discrete-Time Rational Expectations Models,” *mimeo*, *Federal Reserve Bank of San Francisco*, 2005.