

Solving Heterogeneous Agent Models with Non-convex Optimization Problems: Linearization and Beyond

Michael Reiter, Institute for Advanced Studies, Vienna and NYU Abu Dhabi

November 25, 2018

PRELIMINARY AND INCOMPLETE!

Abstract

This paper presents a solution method for heterogeneous agent models where agents solve non-convex optimization problems. It builds on the linearization approach with state reduction of Reiter (2010). State reduction is used as a basis for global nonlinear solutions in medium-dimensional state space. The method is applied to a model of heterogeneous households with indivisible labor supply, and to a model of heterogeneous firms with time-varying uncertainty.

JEL classification:

Address of the author:

Michael Reiter
Institute for Advanced Studies
Josefstädter Strasse 39
A-1080 Vienna
Austria
e-mail: michael.reiter@ihs.ac.at

1 Introduction

The solution of heterogeneous agent models with incomplete markets poses important technical challenges, mainly because of the dimension of the underlying state space. Krusell and Smith (1998) were using a very low-dimensional approximation to the aggregate law of motion, which was highly successful in their specific application, and this method has been the workhorse in this field for 20 years. Reiter (2009) uses the opposite approach: the solution is linearized in aggregate variables (while fully nonlinear in the solution of the individual problem), which allows to keep track of very high-dimensional approximation of the cross-sectional distribution. The approach of Reiter (2010a) pushes this even further: optimal state-reduction increases the number of state variables that can be handled by linearization. This approach has been applied in McKay and Reis (2016) and Reiter, Sveen, and Weinke (2013), and Ahn, Kaplan, Moll, Winberry, and Wolf (2018) developed a similar technique for continuous time models. The approach also has similarities to Kubler and Scheidegger (2018), who reduce the state space to a very small dimension with a bounded-rationality interpretation, while Reiter (2010a) uses a higher-dimensional approximation with the aim to provide almost-exact aggregation. Boppart, Krusell, and Mitman (2018) notice that there is an alternative, and in a certain sense simpler approach to compute linearized solution, noting that the simulation of a linearized model is just a linear superposition of impulse responses to a MIT shock.

While the use of the linearization method has grown in recent years, it is clear that certain classes of models require solution methods that are nonlinear also in aggregate variables. The most obvious case are models of portfolio choice. Another example are models with stochastic volatility, such as Bloom, Floetotto, Jaimovich, Eksten, and Terry (2018), where it is important to investigate to what extent aggregate uncertainty generates precautionary behavior, which again needs nonlinear solution techniques. On the other hand, models where aggregate uncertainty affects behavior in an essential way, need a solution technique that goes beyond linearization. The method of Krusell and Smith (1998) is a hybrid between linear and nonlinear solution. The household value function is a nonlinear function of both individual and aggregate states, but households assume a (log-)linear law of motion for the aggregate state. This rules out, for example, asymmetric responses to positive and negative shocks. The nonlinear method that I develop below avoids the parameterization of the aggregate law of motion. Aggregate transitions are computed in equilibrium, separately for each point on an aggregate grid.

In the present paper I make two contributions. First, I complete the theory of model reduction introduced in my working paper Reiter (2010a), by adding optimal value function reduction. Second, I test whether this theory is useful for state-of-the art heterogeneous agent models, which very often involve non-convexities in the agents' optimization problem. Non-convexities lead to discontinuous policy functions and irregular cross-sectional distribu-

tions. These make it difficult to compute extensive margin effects in models with continuous and discrete choice. This as two parts. First I investigate whether linearization and model reduction give results that are similar in precision to results presented in the literature with different methods. Second, I show how to build on the information from the linearized model to construct different types of nonlinear solutions.

The resulting method is complicated, but most of it is automated in a toolkit that is programmed in Julia, and which uses a special syntax for heterogeneous agent models, similar to what Dynare does for representative agent models.

2 Example Models

2.1 Example 1: the Model of Chang and Kim (2007)

This model is very similar to a standard model of the Krusell and Smith (1998)-type, except for introducing indivisible labor.

Ex-ante identical households face shocks to idiosyncratic labor productivity, which follows a Markov process. Labor markets work frictionlessly, but labor is indivisible: a household can work either zero hours or a fixed number of hours. The objective of the paper is to study the observed "labor wedge", defined as the discrepancy between wage and observed MRS between consumption and leisure:

$$wedge = w - \frac{U_L}{U_C} \quad (1)$$

where we use a standard utility function $U(C, L)$. Why are there systematic fluctuations in the wedge over the business cycle? The claim is that this is caused by the combination of indivisible labor in combination with aggregation over heterogeneous households.

Technology is standard. There is a representative firm with production function:

$$Y_t = F(L_t, K_t, \lambda_t) = \lambda_t L_t^\alpha K_t^{1-\alpha} \quad (2)$$

where λ_t is Markov with transition probability distribution π_λ . The household value function is given by

$$V(a, x; \lambda, \phi) = \max_{a' \in \mathcal{A}, h \in \{0, \bar{h}\}} \{ u(c, h) + \beta \mathbb{E} V(a', x'; \lambda', \phi') \} \quad (3)$$

s.t.

$$c = w(\lambda, \phi)xh + (1 + r(\lambda, \phi))a - a' \quad (4)$$

$$a' \geq \bar{a} \quad (5)$$

$$\phi' = T(\lambda, \phi) \quad (6)$$

where

- x is exogenous individual productivity process

- λ is aggregate TFP process
- ϕ is cross-sectional distribution of agents over (a, x) .

In this model, the aggregate number of hours is determined exclusively by the extensive margin, namely the fraction of households who decide to work in a given period. The effective labor supply also depends on the productivity level of the working households. The labor supply reaction is determined by two things:

- the change in the threshold level of capital where households switch from working to non-working
- the mass of households close to this threshold.

A main computational challenge is to pin down the labor supply response. This is made difficult by the irregular shape of the cross-sectional distribution of capital, which is caused by the discontinuities in households' savings function at the participation threshold.

2.2 Example 2: Business cycles with investment uncertainty

The model is a combination of Khan and Thomas (2008) and the time varying uncertainty of Bloom, Floetotto, Jaimovich, Eksten, and Terry (2018).

[DETAILS TO BE FILLED IN.]

2.3 A general model

For notational simplicity, we assume there is only one type of heterogeneous agent (the households in the Chang/Kim model, or the firms in the uncertainty model). Generalizing this to the case of several ex-ante heterogeneous agents is conceptually straightforward. The problem of this agent is described as a dynamic programming problem. The rest of the model consists of a finite set of dynamic equations, just as in a standard DSGE model.

The decision problem of the heterogeneous agent is assumed to be of the following form. For the agent there are four types of state variables:

1. One continuous individual endogenous state variable (such as capital), denoted by k .
2. One discrete individual endogenous state variable (such as employment states, size of owned house out of a finite choice of house sizes, etc.), denoted by e .
3. One discrete individual exogenous state variable (such as labor productivity) denoted by z .
4. The set of aggregate states, denoted by Ω , which consists of the predetermined states \mathbf{D} and the current states Z . \mathbf{D} contains the cross-sectional distribution of individual states at the end of the last period, and potentially some other predetermined states.

Z typically consists of the current values of the exogenous driving processes. The agent takes the transition law of the aggregate state Ω as given.

The agent has two control variables:

1. The continuous control variable k' , which for convenience we assume is equal to the end-of-period continuous state.
2. The discrete control variable d . We assume that this variable determines the end-of-period discrete state, but it also influences current utility. For example, in the Chang/Kim model there is no endogenous discrete state, but the discrete choice (working or not working) affects current utility, conditional on the end-of-period capital of the household.

The transition law of the individual exogenous state is of course independent of the actions of the individual agent, but can depend on the current aggregate state. We assume it is characterized by a finite Markov chain with transition probabilities $\pi(z, z'; \Omega)$. At this state we are not making any assumptions about the transition law of the aggregate state Ω .

This setup is somewhat more general than it may appear. The agent may have more continuous control variables, which one can handle if the other continuous controls are related to the end-of-period state by static optimality conditions, for example an optimal labor supply condition conditional on saving. A discrete state variable is often a discrete approximation to a continuous state variable, where we assume it is not necessary to have a very fine approximation. Several discrete variables can of course always be combined into one discrete variable by forming a Cartesian product. We assume for notational simplicity that the end-of-period endogenous states are also the beginning-of-period states in the following period. One could easily generalize this by introducing some further exogenous noise into the individual transition laws.

Discrete choice renders the decision problem non-convex. At points in the state space where the discrete choice changes, the continuous choice generally also jumps. In the numerical solution, we will use a discrete grid also for the continuous state, but we will identify the value of the continuous choice where the discrete choice jumps (the "switch points"). The change of the switch points in reaction to changes in the aggregate environment defines the extensive-margin reaction of the agents.

At an aggregate state Ω , we define the individual policy function $\mathcal{A}(\Omega)$ as comprising the two decision functions $k'(k, e, z; \Omega)$ and $d(k, e, z; \Omega)$.

Apart from the dynamic optimization problem of agents, the model is characterized by a transition equation for the cross-sectional distribution:

$$\mathbf{D}' = \mathcal{T}(\mathbf{D}, Z, p, \mathcal{A}) \tag{7}$$

and equilibrium conditions for endogenous aggregate variables as a function of state and decisions:

$$0 = \mathcal{E}(p, \mathbf{D}, Z, \mathcal{A}) \quad (8)$$

3 Linear Approximation and State Reduction

The linearized solution requires the following steps:

1. A finite-dimensional approximation of the model, cf. Section 3.1.
2. Model reduction (Sections 3.2–3.5).
3. Simulating the model (Section 3.6).
4. Accuracy check (Section 3.7).
5. Balanced reduction (Section 3.8).

3.1 Finite approximation of HA model

The finite approximation (discretization) of a HA model was explained in Reiter (2009) (cf. also Costain and Nakov (2011)) and is now standard. Here I will mainly explain the elements that refer to the value function approximation, because this is the part that comes in because of the non-convexity of the agents' decision problem.

3.1.1 Discretization of the value function

The problem of the heterogeneous agents is solved as a discrete dynamic programming problem. For this, choose a grid G_k of the continuous individual state variable with n_k points, denoted by \bar{k}_j for $j = 1, \dots, n_k$. From this, form a grid of individual states G_x of $n_k \cdot n_e \cdot n_z$ points by a Cartesian product between the capital grid and the sets of discrete individual states:

$$G_x = G_k \times \{e_1, \dots, e_{n_e}\} \times \{z_1, \dots, z_{n_z}\} \quad (9)$$

Denote the elements of this grid as \bar{x}_i for $i = 1, \dots, n_k \cdot n_e \cdot n_z$. Between grid points, the value function will be interpolated in k by a quadratic interpolation method that preserves the contraction property of the Bellman iteration (for details, cf. Appendix A).

This grid will be used in both the linearized and the nonlinear solution.

3.1.2 Finite approximation of the cross-sectional distribution

The cross-sectional distribution of individual states will also be approximated on a finite grid. There are two simple ways to approximate continuous distribution on a finite grid. The first one is to model them as a combination of point masses, more concretely as the fraction

of agents at each point of a fixed grid, which was used for example in Young (2010). The second one is to model the distribution as a histogram, assuming a constant density within a histogram bin, which means *between* two grid points (Reiter 2009). For a problem with both continuous and discrete decisions, the histogram approach is the appropriate one, for the following reason. The model solution determines threshold points, where the discrete choice changes. The extensive margin effect in a model will be determined by the change in the threshold points, as well as the density of agents at the threshold. Since the thresholds will generally not lie on a predetermined grid, the approximation as point masses cannot capture this effect. The histogram approach can handle this situation.

For the histogram of the cross-sectional distribution one can use the same grid as for the value function approximation, but in may in general be useful to allow for a different choice of grid points, to model the cross-sectional dynamics with better precision. We therefore introduce a different grid G_k^D of the continuous individual state variable with n_D points, denoted by \bar{k}_j^D for $j = 1, \dots, n_D$. From this, we form a grid of individual states G_x^D of $n_D \cdot n_e \cdot n_z$ points.

3.1.3 The discretized model

The model contains three types of variables and equations.

1. The value function at each of the points in the grid G_k . This gives $n_k \cdot n_e \cdot n_z$ points in each period.

The corresponding equation for each of these variables is the Bellman equation (10) at this grid point.

2. The mass of agents in each histogram bin defined by the grid G_k^D . This gives $(n_k - 1) \cdot n_e \cdot n_z$ points in each period.

The corresponding equation is the transition dynamics at this point, which depends on the individual policy function. Details are given in Appendix B.

3. Aggregate variables such as GDP, TFP etc., and their corresponding equations. This is the same as in standard DSGE models.

All these variables at time t are collected into the vector θ_t . Since this vector can be huge, some form of model reduction may be needed, which is the topic of Sections 3.2–3.5. Notice that the optimal decisions k and e are not part of θ_t . Given current states and next period's value function, they are implicitly given by the optimality conditions.

3.1.4 Differentiating the value function

The solution of the agent problem is characterized by the Bellman equation

$$V(k, e, z, \Omega, Z, p) = \max_{d, k'} \tilde{V}(k, e, z, \Omega, Z, d, k', p) \quad (10)$$

where the current value function conditional on current equilibrium variables p and on current actions d, k' is defined as

$$\tilde{V}(k, e, z, \Omega, Z, d, k', p) \equiv U(k, e, z, d, k'; p, \Omega) + \beta \sum_{z'} \pi(z, z'; \Omega) \mathbb{E}_{\Omega'} \{V(k', T(e, d), z', \Omega')\} \quad (11)$$

Writing $k'()$ and $d()$ for the optimal decision, we can write the Bellman equation as

$$V(k, e, z, \Omega, Z) = U(k, e, z, d(), k'(); p, \Omega) + \beta \sum_{z'} \pi(z, z'; \Omega) \mathbb{E}_{\Omega'} \{V(k'(), T(e, d()), z', \Omega')\} \quad (12)$$

For the linearized solution, we have to differentiate (12) with respect to all the variables of the model. Define by ω any of the variables with respect to which we differentiate. Then we get from (12) that

$$\begin{aligned} \frac{\partial V(k, e, z, \Omega, Z)}{\partial \omega} &= \frac{\partial U(k, e, z, d(), k'(); p, \Omega)}{\partial k'} \frac{\partial k'()}{\partial \omega} + \frac{\partial U(k, e, z, d(), k'(); p, \Omega)}{\partial \omega} \\ &+ \beta \left[\frac{\partial \sum_{z'} \pi(z, z'; \Omega)}{\partial \omega} \mathbb{E}_{\Omega'} \{V(k'(), T(e, d()), z', \Omega')\} \right. \\ &\left. + \sum_{z'} \pi(z, z'; \Omega) \mathbb{E}_{\Omega'} \left\{ \frac{\partial V(k'(), T(e, d()), z', \Omega')}{\partial \omega} + \frac{\partial V(k'(), T(e, d()), z', \Omega')}{\partial k'} \frac{\partial k'}{\partial \omega} \right\} \right] \quad (13) \end{aligned}$$

The problem is that $\frac{\partial k'}{\partial \omega}$ is not known. However, since we assume that k' is bound constrained, there are two possible cases. Either the optimal k' is interior, in which case the envelope theorem applies and (13) simplifies to

$$\begin{aligned} \frac{\partial V(k, e, z, \Omega, Z)}{\partial \omega} &= \frac{\partial U(k, e, z, d(), k'(); p, \Omega)}{\partial \omega} \\ &+ \beta \left[\frac{\partial \sum_{z'} \pi(z, z'; \Omega)}{\partial \omega} \mathbb{E}_{\Omega'} \{V(k'(), T(e, d()), z', \Omega')\} \right. \\ &\left. + \sum_{z'} \pi(z, z'; \Omega) \mathbb{E}_{\Omega'} \left\{ \frac{\partial V(k'(), T(e, d()), z', \Omega')}{\partial \omega} \right\} \right] \quad (14) \end{aligned}$$

Or k' is at a constraint, and which case $\frac{\partial k'}{\partial \omega}$ is simply the derivative of the constraint.

For given equilibrium vector p and discrete states e, z a threshold point k^* where the continuous choice switches from k'_- to k'_+ is characterized by

$$\begin{aligned} \max_{k'_-} U(k^*, e, z, d_-, k'_-; p, \Omega) + \beta \sum_{z'} \pi(z, z'; \Omega) \mathbb{E}_{\Omega'} \{V(k'_-, T(e, d_-), z', \Omega')\} = \\ \max_{k'_+} U(k^*, e, z, d_+, k'_+; p, \Omega) + \beta \sum_{z'} \pi(z, z'; \Omega) \mathbb{E}_{\Omega'} \{V(k'_+, T(e, d_-), z', \Omega')\} \quad (15) \end{aligned}$$

The extensive margin effect of any shock or change in parameters depends largely on how it affects this thresholds, and therefore it is essential that we compute the derivative of the

threshold w.r.t. any other variables. It is given by

$$\frac{\partial k^*}{\partial \omega} = - \frac{\frac{\partial \tilde{V}(X_+)}{\partial \Omega} \frac{\partial \Omega}{\partial \omega} + \frac{\partial \tilde{V}(X_+)}{\partial p} \frac{\partial p}{\partial \omega} + \frac{\partial \tilde{V}(X_+)}{\partial k'_+} \frac{\partial k'_+}{\partial \omega} - \frac{\partial \tilde{V}(X_-)}{\partial \Omega} \frac{\partial \Omega}{\partial \omega} - \frac{\partial \tilde{V}(X_-)}{\partial p} \frac{\partial p}{\partial \omega} - \frac{\partial \tilde{V}(X_-)}{\partial k'_-} \frac{\partial k'_-}{\partial \omega}}{\frac{\partial \tilde{V}(X_+)}{\partial k} - \frac{\partial \tilde{V}(X_-)}{\partial k}} \quad (16)$$

where $X_- \equiv (k, e, z, \Omega, d_-, k'_-, p)$ and $X_+ \equiv (k, e, z, \Omega, d_+, k'_+, p)$. Following the discussion of how to differentiate the value function, we set $\frac{\partial k'_-}{\partial \omega}$ and $\frac{\partial k'_+}{\partial \omega}$ in (16) as zero in case they are an interior solution, or equal to the derivative of the relevant bound constraint in case they are constrained.

3.2 Model reduction: general outline

Differentiating the equations of the model outlined in Section 3.1.3 we obtain a system of linear rational expectations equations

$$\Lambda \theta_{t-1} + \Gamma \theta_t + E_t \Phi \theta_{t+1} + \Psi \epsilon_t = 0 \quad (17)$$

with a very large vector of variables θ . We want to reduce the dimension of the model without any significant loss in accuracy. The engineering literature (Antoulas 2005) shows how to do this if the model is already given in the VAR form $\theta_t = A \theta_{t-1} + B \epsilon_t$ (or in state space form). We cannot apply this directly because we first have to solve the model in order to know the dynamics.

The first step is to partition the equation system (17) as

$$\begin{bmatrix} \Lambda_{ss} & 0 & 0 \\ \Lambda_{ys} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} s_{t-1} \\ y_{t-1} \\ v_{t-1} \end{bmatrix} + \begin{bmatrix} \Gamma_{ss} & \Gamma_{sy} & \Gamma_{sv} \\ \Gamma_{ys} & \Gamma_{yy} & \Gamma_{yv} \\ 0 & \Gamma_{vy} & \Gamma_{vv} \end{bmatrix} \begin{bmatrix} s_t \\ y_t \\ v_t \end{bmatrix} + E_t \begin{bmatrix} 0 & 0 & \Phi_{sv} \\ 0 & 0 & \Phi_{yv} \\ 0 & 0 & \Phi_{vv} \end{bmatrix} \begin{bmatrix} s_{t+1} \\ y_{t+1} \\ v_{t+1} \end{bmatrix} + \begin{bmatrix} \Psi_s \\ \Psi_y \\ 0 \end{bmatrix} \epsilon_t = 0 \quad (18)$$

The variable vector θ partitioned as (s, y, v) such that only v appears with time index $t + 1$, only s appears with time index $t - 1$, and only y , not s enters equations for v . We assume that Γ_{ss} , Γ_{yy} , and Γ_{vv} are regular. We further assume that Γ_{ss} and Γ_{vv} are very sparse (often just the identity matrix) and therefore easy to invert.

In practice, we take as s all the variables that appear with time index $t - 1$, and as v all the variables that appear with time index $t + 1$. These two groups must not overlap. y are all the other variables. We split the equation systems into equations so as to satisfy the constraints implicit in (18).

With the above assumptions, we can rewrite (18) as

$$\begin{bmatrix} \Gamma_{ss}^{-1} \Lambda_{ss} \\ \Lambda_{ys} \\ 0 \end{bmatrix} s_{t-1} + \begin{bmatrix} I & \Gamma_{ss}^{-1} \Gamma_{sy} & \Gamma_{ss}^{-1} \Gamma_{sv} \\ \Gamma_{ys} & \Gamma_{yy} & \Gamma_{yv} \\ 0 & \Gamma_{vv}^{-1} \Gamma_{vy} & I \end{bmatrix} \begin{bmatrix} s_t \\ y_t \\ v_t \end{bmatrix} + \mathbb{E}_t \begin{bmatrix} \Gamma_{ss}^{-1} \Phi_{sv} \\ \Phi_{yv} \\ \Gamma_{vv}^{-1} \Phi_{vv} \end{bmatrix} v_{t+1} + \begin{bmatrix} \Gamma_{ss}^{-1} \Psi_s \\ \Psi_y \\ 0 \end{bmatrix} \epsilon_t = 0 \quad (19)$$

Because both s and v are very large vectors, the task is to reduce the dimension of the model. This has two components:

1. State reduction: choose an $n_m \times n_s$ matrix \bar{M} with $n_m < n_s$ and define

$$m_t = \bar{M} s_t \quad (20)$$

Interpretation: m_t denotes the statistics ("m" is a memo of "moments") of the cross-sectional distribution that agents based their decision on (bounded rationality). We assume that those statistics are linear functions of the distribution.

The matrix must be such that there exist $\tilde{\Lambda}_{ys}$ and $\tilde{\Gamma}_{ys}$ with

$$\Lambda_{ys} = \tilde{\Lambda}_{ys} \bar{M}, \quad \Gamma_{ys} = \tilde{\Gamma}_{ys} \bar{M} \quad (21)$$

(21) is satisfied if the rows of Λ_{ys} and Γ_{ys} are spanned by the rows of \bar{M} :

$$\begin{bmatrix} \Lambda'_{ys} & \Gamma'_{ys} \end{bmatrix} \in \text{span}(\bar{M}') \quad (22)$$

2. Value reduction: choose a matrix \bar{V} that spans the space in which the value function is assumed to live:

$$v_t = \bar{V} f_t \quad (23)$$

$\dimen(f) \ll \dimen(v)$. W.l.o.g. we can choose \bar{V} as orthonormal so that $\bar{V}'\bar{V} = I$.

Using (20), (21) and (23), and premultiplying the first block of equations in (19) by \bar{M} , (19) can be written as and the second block by \bar{V}' , (19) can be written as

$$\begin{bmatrix} \bar{M} \Gamma_{ss}^{-1} \Lambda_{ss} \\ \tilde{\Lambda}_{ys} \bar{M} \\ 0 \end{bmatrix} s_{t-1} + \begin{bmatrix} I & \bar{M} \Gamma_{ss}^{-1} \Gamma_{sy} & \bar{M} \Gamma_{ss}^{-1} \Gamma_{sv} \bar{V} \\ \tilde{\Gamma}_{ys} & \Gamma_{yy} & \Gamma_{yv} \bar{V} \\ 0 & \bar{V}' \Gamma_{vv}^{-1} \Gamma_{vy} & I \end{bmatrix} \begin{bmatrix} m_t \\ y_t \\ f_t \end{bmatrix} + \mathbb{E}_t \begin{bmatrix} \bar{M} \Gamma_{ss}^{-1} \Phi_{sv} \bar{V} \\ \Phi_{yv} \bar{V} \\ \bar{V}' \Gamma_{vv}^{-1} \Phi_{vv} \bar{V} \end{bmatrix} f_{t+1} + \begin{bmatrix} \bar{M} \Gamma_{ss}^{-1} \Psi_s \\ \Psi_y \\ 0 \end{bmatrix} \epsilon_t = 0 \quad (24)$$

Notice the asymmetry between state reduction versus value reduction: we have assumed we know a good approximation to the subspace in which the value function lives. I will show in Section 3.3 how to compute such a subspace by iterating forward on the Bellman equation. For the cross-sectional distribution, in contrast, we only assume that the statistics m_t contain the most relevant information about the distribution. There is no easy way to specify an approximate subspace, in particular because the cross-sectional distribution can take a quite erratic shape when individual policy functions are discontinuous, as they usually are with non-convex decision problems. Because of this problem, the equation system (24) still contains the full state vector s . The problem is the term $\bar{M}\Gamma_{ss}^{-1}\Lambda_{ss}s_{t-1}$. One can approach this problem in two ways. The simpler approach, which can be applied for any matrix \bar{M} , is to make some plausible choice of this subspace. Concretely, we set $s_t = \bar{P}m_{t-1}$ for some matrix \bar{P} so that $\bar{M}\Gamma_{ss}^{-1}\Lambda_{ss}s_{t-1}$ in (24) is replaced by $\bar{M}\Gamma_{ss}^{-1}\Lambda_{ss}\bar{P}m_{t-1}$. The logic behind this approach is the following. If it is true that $\bar{M}s_t$ contains the relevant information about s_t , the exact choice of s_t should not matter much. This should at least be good enough for a first solution of the model. One can then simulate the model to find a more appropriate PD and solve the model again. This is explained in Section

The alternative approach, which can lead to an high-precision solution, is to find a suitable matrix \bar{M} such that there exists a matrix \hat{A} with $\bar{M}\Gamma_{ss}^{-1}\Lambda_{ss} = \hat{A}\bar{M}$. In this case, $\bar{M}\Gamma_{ss}^{-1}\Lambda_{ss}s_{t-1}$ in (24) becomes $\hat{A}m_{t-1}$. This is explained in Section .

3.3 Almost Exact Value Function Reduction

The first step is to find lower-dimensional basis \bar{V} of the space in which the value function v must lie. This can be done with almost zero loss of accuracy in the following way:¹ Setting $\tilde{\Gamma}_{vy} = -\Gamma_{vv}^{-1}\Gamma_{vy}$ and $\tilde{\Phi}_{vv}v_{t+1} = -\Gamma_{vv}^{-1}\Phi_{vv}$, we can write the third line of equations in (19) as

$$v_t = \mathbf{E}_t \left(\tilde{\Gamma}_{vy}y_t + \tilde{\Phi}_{vv}v_{t+1} \right) \quad (25)$$

Iterating gives

$$\begin{aligned} v_t &= \mathbf{E}_t \left[\tilde{\Gamma}_{vy}y_t + \tilde{\Phi}_{vv}\tilde{\Gamma}_{vy}y_{t+1} + \tilde{\Phi}_{vv}^y\tilde{\Gamma}_{vy}y_{t+y} + \dots \right] \\ &= \mathbf{E}_t \sum_{i=0}^{\infty} \tilde{\Phi}_{vv}^i \tilde{\Gamma}_{vy}y_{t+i} \end{aligned} \quad (26)$$

At this stage, we do not know $\mathbf{E}_t y_{t+i}$, but (26) implies that v_t is spanned by the columns of the $\tilde{\Phi}_{vv}^i \tilde{\Gamma}_{vy}$. This information is useless if the rank of the $\tilde{\Phi}_{vv}^i \tilde{\Gamma}_{vy}$ taken together equals the dimension of v . It turns out, however, that the numerical rank of the $\tilde{\Phi}_{vv}^i \tilde{\Gamma}_{vy}$ as determined by the finite machine precision is much smaller than the dimension of v . An essential condition

¹In Reiter (2010a) I was using an iterative algorithm to determine \bar{V} , but the procedure described here is much better.

for this is that Γ_{vy} has small rank, which says again that there is only a small set of equilibrium variables y that enter the agents' utility functions.

For the basis \bar{V} of the space in which the value function v must lie, we take an orthonormal basis of $\cup_{i=0}^k \tilde{\Phi}_{vv}^i \tilde{\Gamma}_{vy}$ for some finite k . Truncating k implies no significant loss of accuracy because of discounting.

3.4 State aggregation: proxy distributions

We choose a matrix \bar{P} which selects to any vector of statistics m_t one specific distribution (called "proxy distribution", (Reiter 2010b)) $s_t^{PD} = \bar{P}m_t$ which has those statistics. This means we require

$$\bar{M}\bar{P} = I_{n_m} \quad (27)$$

The interpretation of this approach is that the state equations in (18) are satisfied not at all possible state vectors s , but only at those that are "typical" distributions in the sense $s = \bar{P}m$ for any m . One general way to choose the proxy distribution is to use steady state information: given any moment matrix \bar{M} , one can choose the corresponding proxy distributions as the distributions that are closest to steady subject to the moment constraint:

$$\max_{\tilde{s}} \frac{1}{2} \tilde{s}' \Omega^{-1} \tilde{s} \quad s.t. \quad \bar{M}\tilde{s} = \tilde{m} \quad (28)$$

The solution to (28) is given by $\tilde{s} = \Omega \bar{M}' (\bar{M} \Omega \bar{M}')^{-1} \tilde{m}$, which means

$$\bar{P} = \Omega \bar{M}' (\bar{M} \Omega \bar{M}')^{-1} \quad (29)$$

For an illustrative example (which does not satisfy (21)), choose \bar{M} such that it selects adjacent bins:

$$\bar{M} = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & \dots & 0 & 0 & 0 \\ \dots & & & \dots & \dots & \dots & \dots & & & \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 1 & 1 & 1 \end{pmatrix}$$

With $\Omega = I$, (29) gives

$$\bar{P} = \begin{pmatrix} 1/3 & 1/3 & 1/3 & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ 0 & 0 & 0 & 1/3 & 1/3 & 1/3 & \dots & 0 & 0 & 0 \\ \dots & & & \dots & \dots & \dots & \dots & & & \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 1/3 & 1/3 & 1/3 \end{pmatrix}$$

The proxy distribution assumes equal distribution within bins. Alternatively, if Ω is the diagonal matrix with elements equal to the steady state distribution s^* , the proxy distribution

is proportional to the steady state distribution within bins:

$$\bar{P} = \begin{pmatrix} \frac{s_1^*}{\sum_{i=1}^3 s_i^*} & \frac{s_2^*}{\sum_{i=1}^3 s_i^*} & \frac{s_3^*}{\sum_{i=1}^3 s_i^*} & \dots & 0 & 0 & 0 & & & & \\ 0 & 0 & 0 & \frac{s_4^*}{\sum_{i=4}^6 s_i^*} & \frac{s_5^*}{\sum_{i=4}^6 s_i^*} & \frac{s_6^*}{\sum_{i=4}^6 s_i^*} & \dots & 0 & 0 & 0 & \\ & \dots & & & \dots & & \dots & \dots & \dots & & \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & \frac{s_{n-2}^*}{\sum_{i=n-2}^n s_i^*} & \frac{s_{n-1}^*}{\sum_{i=n-2}^n s_i^*} & \frac{s_n^*}{\sum_{i=n-2}^n s_i^*} & \end{pmatrix}$$

3.5 Almost-exact state aggregation

As we have discussed at the end of Section 3.2, our aim is to find a selection matrix \bar{M} such that (21) is satisfied, and there exists a matrix \hat{A} with $\bar{M}\Gamma_{ss}^{-1}\Lambda_{ss} = \hat{A}\bar{M}$. To do this, we follow what in Reiter (2010a) is called the "Conditional Expectations Approach". So assume we want to predict the endogenous variables y_{t+i} for $i = 0, \dots, \infty$. in a linear model where y is related to the states s by

$$y_t = Cs_t \quad (30)$$

and we have the state transition equation

$$E_t s_{t+i} = A^i s_t \quad (31)$$

where the system matrix A has dimension $n \times n$, and C has dimension $m \times n$. Obviously, we need the following linear combinations of s :

$$Cs_t, \quad CA s_t, \quad CA^2 s_t, \quad \dots \quad (32)$$

Stack them into

$$Q = \begin{bmatrix} C \\ CA \\ CA^2 \\ \dots \\ CA^{n-1} \end{bmatrix} \quad (33)$$

The $m \cdot n \times n$ matrix Q is called "observability matrix". This approach is only useful if the rank of Q is substantially lower than n . Define $k \leq n$ as the rank of Q . The SVD of Q can be written as

$$Q = \begin{bmatrix} U_1 & U_2 \end{bmatrix} \begin{bmatrix} S & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} V_1' \\ V_2' \end{bmatrix} = U_1 S V_1' \quad (34)$$

$$S \equiv \text{diag}(\sigma_1, \dots, \sigma_k) \quad (35)$$

where U_1 and V_1 have dimension $m \cdot n \times k$, $U_1' U_1 = V_1' V_1 = I_k$.

From the Cayley-Hamilton theorem, there exists a Λ such that

$$QA = \Lambda Q \quad (36)$$

Using (34) in (36) we get $U_1 S V_1' A = \Lambda U_1 S V_1'$. Premultiplying by $S^{-1} U_1'$ we get $V_1' A = S^{-1} U_1' \Lambda U_1 S V_1'$ (notice that S is invertible). Setting $\bar{M} = V_1'$ we get

1. $\bar{M}\bar{M}' = I_k$

2. \bar{M} can be interchanged with A :

$$\bar{M}A = \hat{A}\bar{M} \quad (37)$$

with $\hat{A} = S^{-1}U_1'\Lambda U_1S$ being a $k \times k$ -matrix. From 1. it follows that $\hat{A} = \bar{M}A\bar{M}'$.

3. We have $C' \in \text{span}(\bar{M}')$:

$$C = \begin{bmatrix} I & 0 & \dots & 0 \end{bmatrix} Q = \left(\begin{bmatrix} I & 0 & \dots & 0 \end{bmatrix} U_1S \right) \bar{M} \quad (38)$$

We can summarize the above discussion in the following

Proposition 1. *Assume that Q has rank $k \ll n$ when choosing $A = \Gamma_{ss}^{-1}\Lambda_{ss}$ and C such that*

$$\begin{bmatrix} \Lambda'_{ys} & \Gamma'_{ys} \end{bmatrix} \in \text{span}(C) \quad (39)$$

Then there is a $k \times n$ -matrix \bar{M} such that

- $\Lambda_{ys} = \tilde{\Lambda}_{ys}\bar{M}$, $\Gamma_{ys} = \tilde{\Gamma}_{ys}\bar{M}$
- There exists an $n_m \times n_m$ -matrix \hat{A} such that

$$\bar{M}\Gamma_{ss}^{-1}\Lambda_{ss} = \hat{A}\bar{M} \quad (40)$$

Using (40), we can replace $\bar{M}\Gamma_{ss}^{-1}\Lambda_{ss}s_t$ in (24) by $\hat{A}m_t$ and obtain

$$\begin{bmatrix} \hat{A} \\ \tilde{\Lambda}_{ys} \\ 0 \end{bmatrix} m_{t-1} + \begin{bmatrix} I & \bar{M}\Gamma_{ss}^{-1}\Gamma_{sy} & \bar{M}\Gamma_{ss}^{-1}\Gamma_{sv}\bar{V} \\ \tilde{\Gamma}_{ys} & \Gamma_{yy} & \Gamma_{yv}\bar{V} \\ 0 & \bar{V}'\Gamma_{vv}^{-1}\Gamma_{vy} & I \end{bmatrix} \begin{bmatrix} m_t \\ y_t \\ f_t \end{bmatrix} + \text{E}_t \begin{bmatrix} \bar{M}\Gamma_{ss}^{-1}\Phi_{sv}\bar{V} \\ \Phi_{yv}\bar{V} \\ \bar{V}'\Gamma_{vv}^{-1}\Phi_{vv}\bar{V} \end{bmatrix} f_{t+1} + \begin{bmatrix} \bar{M}\Gamma_{ss}^{-1}\Psi_s \\ \Psi_y \\ 0 \end{bmatrix} \epsilon_t = 0 \quad (41)$$

The commutation property (40) says that, knowing the statistics m_t , the exact distribution does not matter for the solution. It does not say that the distributions live in the space spanned by \hat{A} . In the terminology of Section 3.4, it does not necessarily give a good proxy distribution.

3.6 Simulation and Error Analysis

The reduced model (41) can be solved and simulated for (m_t, y_t, f_t) by standard methods. Having a simulation (m_t, y_t, f_t) , can we recover the full state vector s_t and the value function

v_t ? Since we have established in Section 3.3 that v_t lives in the space spanned by \bar{V} and parameterized by f_t , we get v_t directly by

$$v_t = \bar{V} f_t \tag{42}$$

Similarly, we could obtain the distribution by $s_t = \bar{P}m_t$, but this cannot be expected to yield a precise solution, because the proxy distribution \bar{P} given by (29). is a rather arbitrary approximation. And as we have mentioned above, the property (37) does not imply that $\bar{P} = \bar{M}'$ is a good proxy distribution. It only says that $(\bar{M}'m_t)$ is just as good for predicting the future as is s_t .² To avoid the arbitrariness of the proxy distribution, we do not use $s_t = \bar{P}m_t$, but rather use the following, computationally more involved procedure. In any period t of the simulation

1. obtain y_t from the reduced simulation
2. obtain $v_t = \bar{V} f_t$ from the reduced simulation
3. obtain $E_t v_{t+1} = \bar{V} E_t f_{t+1}$ from the reduced simulation
4. obtain s_t by solving the first block of equations in (19).

To obtain a simulated path for the s_t , we obviously have to specify an initial state s_0 .

3.7 Measuring accuracy

We have to deal with three types of approximation error:

1. The error from discretization. This arises already in steady state. We can check for it by varying the number of grid points etc.
2. The error from linearization. We will analyze it later, after obtaining nonlinear solutions. Alternatively, one can check for it by computing perfect foresight solutions after shocks of different size (Boppart, Krusell, and Mitman 2018).
3. The error from aggregation. Minimizing this error is the purpose of the aggregation procedure of Sections 3.3–3.5. By simulating the model as described in (cf. Section 3.6), we can compute the aggregation error by analyzing the residuals along any impulse response function. For any given initial state s_{t-1} , if $\epsilon_t = \epsilon_{t+1} = 0$ it follows from (17) and linearity that

$$\Lambda\theta_{t-1} + \Gamma\theta_t + \Phi\theta_{t+1} = 0 \tag{43}$$

We compute an IR function for $t = 1 : T$, starting from any s_0 . For $t = 2 : T - 1$, we compute the residual

$$Res(t) = \Lambda\theta_{t-1} + \Gamma\theta_t + \Phi\theta_{t+1} \tag{44}$$

²Numerically, (37) is not very precisely satisfied anyway.

To interpret the residual, we must scale it properly. Define

$$RScal(i, t) \equiv \frac{Res(i, t)}{\sum_j (|\Lambda_{i,j}\theta_{j,t-1}| + |\Gamma_{i,j}\theta_{j,t}| + |\Phi_{i,j}\theta_{j,t+1}|)} \quad (45)$$

as the residual of equation i normalized by the sum of the absolute entries in this equation.

With optimal state and value function reduction, the aggregation error turns out to be extremely small. For example, in the case of the Chang/Kim model, the maximum scaled error is of the order 10^{-10} , and the mean scaled error is of the order 10^{-11} .

3.8 Reducing even further: Balanced Reduction

Having solved the reduced model (41), we get the solution in the form

$$\begin{aligned} m_t &= Am_{t-1} + B\epsilon_t \\ y_t &= Cm_t \end{aligned} \quad (46)$$

This can be further analyzed and reduced by the methods developed in the engineering literature, called "balanced reduction" (Antoulas 2005). Define the matrices R , \mathcal{P} , Q , \mathcal{Q} , U , S , V , \tilde{H} as follows:

$$RR' = \mathcal{P} \equiv \mathcal{L}(A, B, \Sigma_\epsilon) \quad (47)$$

$$QQ' = \mathcal{Q} \equiv \mathcal{L}(A', C') \quad (48)$$

$$USV' = R'Q \quad (49)$$

$$\tilde{H} = S^{-1/2}V'Q' \quad (50)$$

where $\mathcal{L}(A, B, \Sigma_\epsilon)$ is defined as the matrix Σ that solves $\Sigma = A\Sigma A' + B\Sigma_\epsilon$. (47) says that R is the Cholesky factor of the covariance matrix \mathcal{P} , and Q in (48) is the Cholesky factor of the observability Gramian \mathcal{Q} , while U , S and V are the SVD of the matrix $R'Q$ with $U'U = I$, $V'V = I$ and S diagonal with decreasing entries. We take S as the square matrix containing only the non-zero singular values (and drop the columns of U and rows of V corresponding to the zero singular values), so that S is invertible by construction.

Now consider the variable transformation $\hat{m} = \tilde{H}m$, $\hat{A} = \tilde{H}A\tilde{H}^{-1}$, $\hat{B} = \tilde{H}B$, $\hat{C} = C\tilde{H}^{-1}$. Using that $\tilde{H}^{-1} = RUS^{-1/2}$, straightforward algebra shows that

$$\mathcal{L}(\hat{A}, \hat{B}, \Sigma_\epsilon) = \tilde{H}\mathcal{L}(A, B, \Sigma_\epsilon)\tilde{H}' = \mathcal{L}(\hat{A}', \hat{C}') = (\tilde{H}')^{-1}\mathcal{L}(A', C')\tilde{H}^{-1} = S. \quad (51)$$

Equ. (51) is a remarkable result. It shows that in the new vector \hat{m} the variables are ordered such that \hat{m}_i has both the i -th highest variance, and makes the i -th highest contribution to future values of y . For the reduced model, we pick the first k components of \hat{m} , or the first k rows of \tilde{H} , such that the diagonal elements $S_{i,i}$ are negligible for $i > k$:

$$H = \tilde{H}_{1:k,:} \quad (52)$$

Properties of balanced reduction

Is the (doubly) reduced model

$$\begin{aligned}\hat{m}_t &= \hat{A}\hat{m}_{t-1} + \hat{B}\epsilon_t \\ \hat{y}_t &= \hat{C}\hat{m}_t\end{aligned}\tag{53}$$

an optimal approximation to the (reduced) model (53) in any sense? With the choices of H that we have discussed, it is not a strictly optimal. Nevertheless, balanced reduction has a strong performance guarantee (cf. Antoulas (2005, Theorem 7.10), Antoulas (1999, Section 2.6)):

$$\text{distance}(\text{ExactModel}, \text{ReducedModel}) \leq 2(\sigma_{k+1} + \dots + \sigma_n)\tag{54}$$

Here, the σ 's are the singular values in (49) (called ‘‘Hankel singular values’’) that were omitted in the construction of H in (52). The distance measure in (54) is the Hankel norm, which is defined as the maximum distance in the future response

$$\sqrt{\sum_{i=0}^{\infty} \|y_{t+i} - \hat{y}_{t+i}\|^2}\tag{55}$$

to any sequence of past shocks ϵ_{t-i} with unit length:

$$\sqrt{\sum_{i=0}^{\infty} \|\epsilon_{t-i}\|^2} = 1\tag{56}$$

In particular, the difference in the usual impulses responses between exact and the reduced models cannot be bigger than the bound (54). This explains why balanced reduction is the standard aggregation method in the control literature.

There exist even better, but more complicated approximations than balanced reduction. The theoretical lower bound on the distance between the two models is σ_{k+1} . This bound can actually be attained (Antoulas 1999, Sections 2.6.1,3.2). For us, it seems not worthwhile to investigate more complicated methods for the linearized model, because our main concern is whether the reduced state from the linear model is still suitable for the nonlinear model.

4 Computing the Nonlinear Solution

4.1 Approximating the value function

The household value function is a function of the individual state (k, e, z) and the aggregate state $\Omega = (\mathbf{D}, Z)$. To make the computation feasible, we assume that the value function can be well approximated as a function of the reduced information set M rather than the full state vector Ω :

$$V(k, e, z, \Omega, Z) \approx \hat{V}(k, e, z, M, Z)\tag{57}$$

where

$$M = H\mathbf{D} \quad (58)$$

for some given matrix H . For ease of notation, I assume that the full Z is in the information set, which is typically the case.

At each point $(k, e, z)_i$ in the individual grid G_x , with $i = 1, \dots, n_k \cdot n_e \cdot n_z$, we approximate this function as a linear combination of n_p known basis functions $\phi_j(M, Z)$,

$$\hat{V}((k, e, z)_i, M, Z) = \sum_{j=1}^{n_p} \phi_j(M, Z) \gamma_{j,i} \quad (59)$$

Notice that we have separate coefficients $\gamma_{j,i}$ for each individual state $(k, e, z)_i$. Obviously, we can write the approximated value function as a function of the original state by $V((k, e, z)_i, \mathbf{D}, Z) = \hat{V}((k, e, z)_i, H\mathbf{D}, Z)$.

Having computed \hat{V} at each grid point $(k, e, z)_i$, we can again interpolate quadratically in the individual state k . In this way, we can compute an interpolated value function at any state (k, e, z, \mathbf{D}, Z) . Denote the interpolated value function as $V^I(k, e, z, \mathbf{D}, Z; \gamma)$. It is parameterized by the vector of coefficients γ .

For the approximated value function, we can define expected continuation values based on two different information sets. The expected continuation value conditional on *end-of-period* information V^{EC} is given by

$$V^{EC}(k', e', z, \Omega, Z; M') = \sum_{z'} \pi(z, z'; \Omega) E_{Z'} V^I(k', e', z', M'Z'; \gamma) \quad (60)$$

If we make a finite approximation of the law of motion of aggregate Z , we can write it as

$$V^{EC}(k', e', z, \Omega, Z; M') = \sum_{z'} \pi(z, z'; \Omega) \sum_{Z'} \pi_Z(Z, Z') V^I(k', e', z', M'Z'; \gamma) \quad (61)$$

Notice that the continuation value depends on end-of-period moments M' , which are determined in equilibrium. Since \mathbf{D}' is predetermined (does not depend on next period's shocks), we only need to know the finite number of statistics $M' = H\mathbf{D}'$ in order to compute the expected continuation value.

Conditional on M' , which is exogenous to the individual agent, the optimization problem can be compactly written as

$$\max_{d, k'} U(k, e, z, d, k'; p, \Omega) + \beta V^{EC}(k', T(e, d), z', \Omega, Z; M') \quad (62)$$

4.2 Temporary Equilibrium

Given any continuation value function $V^{EC}(k', e', z, \Omega, Z; M')$, a temporary equilibrium (p, m', \mathcal{A}) at any aggregate state \mathbf{D}, Z is defined as consisting of

- a set of equilibrium values p

- a set of end-of-period statistics m'
- a policy function \mathcal{A}

such that

- policy functions are optimal, which means that they satisfy (62).
- prices satisfy the equilibrium conditions

$$\mathcal{E}(p, \mathbf{D}, Z, \mathcal{A}) = 0 \tag{63}$$

- end-of-period statistics satisfy the transition law

$$m' = H_D \mathcal{T}(\mathbf{D}, Z, p, \mathcal{A}) \tag{64}$$

This concept of temporary equilibrium requires to solve for the set of statistics M' , which might be impractical if this set is large. We therefore also define a *partial temporary equilibrium*, where we replace the consistency condition (64) by $m' = \mathcal{T}_0(\mathbf{D}', Z)$ with an arbitrarily given transition law $\mathcal{T}_0(\mathbf{D}', Z)$. Below, we will use the transition law obtained from the linearized solution. In this case, the fixed point problem is reduced to finding the equilibrium variables p .

4.3 Backward iteration algorithm

I now describe an algorithm to solve for a global nonlinear approximation of the solution. The task is to find the parameters $\gamma_{j,i}$ that define the value function at each grid point as a function of the reduced aggregate state. This requires that we have chosen a set of states M and a set of basis functions $\phi_j(M, Z)$, as described above. The backward iteration is carried out on a finite set of aggregate grid points Ω_l for $l = 1 \dots, n_A$. For the results reported in Section 5), I have used a grid obtained from an earlier simulation of the model, similar to the approach in (Judd, Maliar, and Maliar 2012).

In the following, I describe here a "conceptual" version of the algorithm, which solves for temporary equilibrium in each step and at each grid point. This is very time consuming, but should be the best guarantee for convergence. A practical algorithm makes choices about what to update in which step, so as to increase speed while still achieving convergence, but this is an implementation detail.

1. Denote the maximum iteration count by T . Initialize the value function $V_{T+1}(\bar{x}_i; \bar{\Omega}_j)$ for $i = 1, \dots, n_k \cdot n_e \cdot n_z$ and $j = 1, \dots, n_A$ by the value function obtained in an earlier solution (for example the linearized solution).
2. For $t = T : -1 : 1$ do

- (a) Compute the coefficients of the polynomial approximation separately for each $k, e, z_j \in G_x$ by the linear projection $\gamma_{i,t+1} = B \backslash \vec{V}_{i,t+1}$ where

$$B = \begin{pmatrix} \phi_1(\bar{M}_1) & \cdots & \phi_{n_p}(\bar{M}_1) \\ \vdots & \vdots & \vdots \\ \phi_1(\bar{M}_{n_M}) & \cdots & \phi_{n_p}(\bar{M}_{n_M}) \end{pmatrix}, \quad \vec{V}_{i,t+1} = \begin{pmatrix} V_{t+1}(\bar{x}_i; \bar{\Omega}_1) \\ \vdots \\ V_{t+1}(\bar{x}_i; \bar{\Omega}_{n_A}) \end{pmatrix} \quad (65)$$

- (b) For each aggregate grid point $M_j, j = 1, \dots, n_A$ do

- i. Guess (p, m')
- ii. Define an expected continuation value at grid points as

$$\tilde{V}(k_i, e_i, z_i) = \sum_{z'} \pi(z_i, z'; \Omega) \sum_{Z'} \pi_Z(Z, Z') \sum_{j=1}^{n_p} \phi_j(m') \gamma_{j,t+1}(k_i, e_i, z_i) \quad (66)$$

- iii. Compute policy \mathcal{A} that solves (62) with continuation value function

$$[d(x_i), k'(x_i)] = \operatorname{argmax}_{d, k'} \left\{ U(x_i, d, k'; p, M_j) + \beta \tilde{V}_{\mathcal{I}}(k'; T(e_i, d), z_i) \right\} \quad (67)$$

where x_i stands for k_i, e_i, z_i .

- iv. Check whether (p, m', \mathcal{A}) satisfy the conditions for a temporary equilibrium (62), (63) and (64). If not, update (p, m') until convergence is achieved.

Finding the equilibrium (p, m') describes an $n_e + n_M$ -dimensional fixed point problem which can be solved, for example, by quasi-Newton methods such as Broyden's algorithm.

- (c) Update the value function for all $i = 1, \dots, n_k n_e n_z$ by

$$V_t(x_i, M_j) = U(x_i, d(x_i), k'(x_i), p^*) + \beta \tilde{V}_{\mathcal{I}}(k'(x_i); (\cdot, T(e_i, d), z_i)) \quad (68)$$

After convergence, the value function should approximately satisfy the recursive relationship

$$V_t(x_i, M_j) = U(x_i, d(x_i, M_j), k'(x_i, M_j), p(M_j)) + \beta \tilde{V}_{\mathcal{I}}(k'(x_i, M_j); (\cdot, T(e_i, d(x_i, M_j)), z_i; m'(M_j))) \quad (69)$$

4.4 A sequence of solution methods

The methods described above can be used to obtain a variety of different solutions. I classify them into five groups. Computing such a sequence of solutions helps to understand the different aspects of nonlinearity in the model.

1. **Lin** : Linearized solution

2. **NonL1** : Partial temporary equilibrium (cf. Section 4.2) with expected continuation value obtained from linearized solution. This shows the effect of the distribution on switch points, and captures nonlinear extensive margin effect.

At any point in the simulation, the partial temporary equilibrium reflects the current cross-sectional distribution. Asymmetries in the impulse responses due to changes in the distribution are therefore reflected in this solution. However, the agents do not yet take into account the effect of these nonlinearities on next period's state and on their continuation value.

3. **NonL2** : Temporary equilibrium (cf. Section 4.2) with expected continuation value obtained from linearized solution. This is like **NonL1** , but now the agents compute their continuation value at the equilibrium distribution at the end of the period. This captures another aspect of nonlinearity.
4. **NonL3** : Temporary equilibrium with global linear value function. The value function at each point of the individual grid is still linear in the aggregate states, but this linear function is obtained by a regression over an aggregate grid covering the relevant parts of the state space. This does not yet incorporate precautionary effects w.r.t. aggregate variables.
5. **NonL4** : Temporary equilibrium with a nonlinear value function

Computing all these solutions involves the following steps.

1. Solve the model by linearization.
2. Compute the continuation value function (61) from the linearized model.
Solve the model as a sequence of temporary equilibria (method **NonL1**) and check for deviations from linearity.
3. Simulate a long time series of the model with method **NonL1** or **NonL2** and choose a set of aggregate grid points from this simulation (Judd, Maliar, and Maliar 2012).
4. Approximate the value function as a polynomial in a very small set of aggregate variables, such as aggregate capital and the exogenous shock processes (with a first order polynomial, this is method **NonL3** ; with a higher order polynomial, it is method **NonL4**). Notice that we are using here basically the same value function approximation as in a Krusell/Smith solution. The difference is that we are not assuming a function form for the aggregate law of motion, but rather solve for the aggregate transition as an equilibrium outcome separately at each aggregate state.
5. Simulate a new time series of the model with method **NonL3** or 4. Compute the value function and the Bellman equation residuals ($\hat{V} - V$ in the notation of Section 4.5)

along the simulation path. Project the residuals along the simulation path on a set of potential state variables such as higher moments, states from balanced reduction etc. to see which states are helpful to predict the value function.

6. Approximate the value function as a polynomial in a larger set of aggregate states and repeat Step 4.

4.5 Measuring accuracy of a nonlinear solution

One can measure accuracy by looking at the Bellman equation residual along a simulation path. At any aggregate state that we visit in the simulation, we first solve for the temporary equilibrium, and then compute the value function at this state by Equ. (68), as we do in the update step of the backward iteration algorithm. Denote this value function by V_t . Then we compute the approximated value function \hat{V}_t given by (59). The difference $\tilde{V}_t = \hat{V}_t - V_t$ is the Bellman equation error, similar to the Euler residual in models where the policy function is approximated. We need to scale this error so that it can be interpreted. What matters is not really the level but the slope of the value function. Denote by $V_t[i]$ the value at point i of the value function grid (for ease of notation, I suppress the arguments of the grid of discrete variables). We scale the error as

$$err_t[i] = \left| \frac{(\hat{V}_t[i+1] - V_t[i+1]) - (\hat{V}_t[i] - V_t[i])}{V_t[i+1] - V_t[i]} \right| \quad (70)$$

With log utility, this has the interpretation of a relative Euler approximation error in consumption. We measure the average and the maximum error over the whole simulation.

4.6 Comparison with Krusell/Smith method

In both methods, the value function is approximated as a function of a reduced set of aggregate state variables. There are two fundamental differences. A first one is that Krusell/Smith choose these reduced states as smooth functions of the underlying large state spaces, such as the first or higher moments of the cross-sectional distribution. I supplement these variables by a set of statistics that are the result of balanced reduction of the linearized model. The second fundamental difference is that Krusell/Smith impose a smooth functional form on the aggregate law of motion in these variable when they compute the value function of the economic agent. This method therefore parameterizes both the value function and the aggregate law of motion. In contrast, I solve for the temporary equilibrium at each point of the aggregate grid, which involves to find next period's realization of the cross-sectional distribution (at least of the reduced state that characterizes this distribution). No a priori assumption is made about the aggregate transition law. The only parameterization is the value function. The exact solution will then depend on the grid of aggregate states on which the solution is computed. This grid is obtained from a simulation of an linearized solution

of the model. This is one of the ways in which linearization helps to prepare the nonlinear solution. By not imposing smoothness on the aggregate transition, the method becomes more general, but at the same time it becomes more vulnerable to instability.

5 Numerical Results

5.1 The model of Chang and Kim

I have solved this model using the same parameter values as in the original paper. For the numerical approximation, I have approximated the stochastic process of individual productivity by a 17-state Markov chain. I have chosen a grid of 1000 histogram bins for the capital distribution, and a grid of 400 points for the value function, for each level of productivity. With these choices, the linearized model has somewhat more than 23800 variables. Optimal state reduction reduces the number of states from around 17000 to 389. The value function parameters were reduced from 6800 to 253. As mentioned in Section 3.7, the maximum aggregation error is of the order 10^{-10} .

Although the model is easy to write down, the numerical solution still poses substantial problems. Takahashi (2014) argues that the numerical solution presented in the original paper Chang and Kim (2007) is numerically imprecise, and presents a corrected version of the solution. Table 1 presents key results of the model. The first column shows the original solution of Chang and Kim, the second column shows the corrected solution of Takahashi, the third column shows my solution obtained from the linearized model. The remaining columns show results from several nonlinear approximations, of the different types outlined in Section 4.4.

	Chang/Kim	Takahashi	Lin	NonL1	NonL3	NonL4	NonL4
σ_H	0.76	0.57	0.56	0.56	0.44	0.44	0.44
σ_{wedge}	0.76	0.24	0.23	0.22	0.15	0.14	0.15
σ_Y	1.28	1.30	1.20	1.23	1.14	1.14	1.14
σ_C	0.39	0.33	0.33	0.33	0.42	0.42	0.42
σ_I	3.06	3.08	3.12	3.13	2.75	2.75	2.76
σ_L	0.50	0.41	0.41	0.42	0.34	0.33	0.34
$\sigma_{Y/H}$	0.50	0.49	0.49	0.48	0.57	0.58	0.57
$\rho(H, wedge)$	0.87	0.95	0.95	0.97	0.96	0.96	0.96

Table 1: Results CK model

It turns out that the second moments obtained from the linearized solution are very close to the results in Takahashi (2014). On the one hand, this is not surprising, because the linearized solution does not suffer from the numerical problems pointed out by Takahashi. On the other hand, it is remarkable that two solution methods that are based on very different

simplifying assumptions in the end come to almost identical conclusions. I will argue below that this is due to one commonality between the methods: in both cases, agents solve their optimization problem under the assumption of a linear aggregate law of motion (log-linear versus linear plays very little role in this model).

Figure 1 illustrates why linearity in aggregate variables is a dubious assumption in this model. The upper panel illustrates the labor supply response to a one percent increase in TFP. The blue line shows the increase in the threshold point following the technology shock. The green line shows the density of households at the threshold point, and the red line shows the response of hours, which is the product of the two earlier lines. It was divided by the total number of hours in steady state, so the sum over all the points of the red line is the aggregate elasticity of hours to TFP. One can see that the bulk of the labor supply response comes with household at the intermediate productivity levels (level 7-11 out of 17 possible levels). The lower panel shows some snippets of the cross-sectional distribution of assets in the steady state, for these five levels of individual labor productivity. The x-axis indicates the histogram bin in the capital distribution, relative to the bin in which the labor market participation threshold lies. The y-axis draws the fraction of households in this bin. This means, for each level of productivity, the value at $x=0$ gives the number of households in the capital bin in which the participation threshold lies. As is clear from the upper part of the graph, the density in these bins is crucial for determining aggregate labor supply. Since a linearized solution always gives a scaled up version of the response to very small shocks, only the number of agents in the bin of the threshold matters, not the agents in neighboring bins. This is different in a nonlinear solution, where the threshold does not stay within the same bin of the capital grid. We see that for all productivity levels the mass of agents to the left of the threshold is much larger than to the right of the threshold. This is because households that switch from working to non-working also reduce their saving level. At least in a partial equilibrium settings, this would mean that a lowering of the threshold changes the behavior of many more households than an increase in the threshold, which implies asymmetry in the response to aggregate shocks. In general equilibrium, endogenous price reactions will probably dampen but not eliminate the asymmetry,

To see how this asymmetry affects the solution, Figure 2 draws impulse responses obtained from the first type of nonlinear solution, **NonL1**. For each variable, three lines are drawn. The green line gives the response to a positive shock of two standard deviations. The red line gives the response to a negative shock of two standard deviations, multiplied by -1. If the response was symmetric, the two lines would exactly coincide. The blue line gives the linear impulse response as a benchmark. The shows the asymmetric response that one would expect from Figure 1. In response to a positive shock, labor supply increases by less and wages increase by more than in the linear benchmark, because the participation threshold moves up, and the cross-sectional density function is sharply decreasing to the right of the threshold. In response to a negative shock, labor reacts more and wages react less than in

the linear case. The deviations from symmetry appear quantitatively small. This is mainly because labor supply reacts very strongly to the current wage, for a *given* continuation value. Since there are no labor market frictions, and most households hold enough assets to be well insured against TFP fluctuations, they react very elastically to transitory changes in the real wage. The fourth column in Table 1 presents the results obtained by simulating the temporary equilibrium in this way. The slight asymmetries detected in Figure 2 have almost no effect on second moments.

Figure 3 draws the same impulse responses, this time calculated with a solution of type **NonL3**. For this graph, households are approximating their value function as a cubic function of aggregate capital and TFP only. Results are almost identical if we use a linear function of aggregate capital and TFP, or as a function of capital, TFP and several other moments, obtained from balanced reduction. What matters is not the states that households use for approximation, but the type of approximation: we solve for temporary equilibrium with a *global* approximation of the value function. The result is surprising: for both positive and negative shocks, labor supply reacts less and the real wage reacts more than in the linearized solution. This is mainly driven by the fact that in a stochastic simulation, the distribution around the threshold points is on average different from what it is in the steady state. A further interesting aspect of Figure 3 versus Figure 2 is that the global approximation reduces the asymmetry that was observed with the linearized continuation value function. Temporary equilibrium with market clearing wage leads to less asymmetry than one would expect from Figure 1. Equilibrium with globally approximated value functions reduces it even further.

These results are reflected in columns 4–7 of Table 1. If a global approximation is used, the variance of hours, effective labor and of the labor wedge all go further down. It makes almost no difference whether the approximation is globally linear (column 5), globally nonlinear with only one moment of the distribution (column 6), or globally nonlinear with 5 moments of the distribution (column 7).

To check the accuracy of the solution, Figure 4 plots the Bellman equation errors at some randomly chosen steps in the simulation process. In each case, the errors for three levels of individual productivity are chosen: the 5th, 9th and 13th out of 17 productivity levels. We see the same pattern at both points: errors are extremely small for most part of the asset grid, but for each productivity level there is a spike in a narrow range, which are around the threshold level of assets where households of this productivity switch to non-working. The errors were computed in the solution that approximates the value function as a function of aggregate capital and current TFP only. In this simple model, this already gives a reasonably precise solution, and it turns out to be very difficult to improve on it. The main problem is that aggregate labor supply effectively depends on a few household types that are located around the thresholds from the intermediate productivity levels. This makes aggregate responses rather irregular, in a way that cannot be captured by smooth functions of the aggregate states. To increase accuracy of the solution, it is probably necessary to increase

the number of individual productivity states.

5.2 The model of investment under uncertainty

[TO BE FILLED IN.]

6 Conclusions

Linearization can be successfully applied in heterogeneous agent models with non-convex optimization problems. For a number of baseline models, the linearized solution is very close to the solution obtained in the original papers (for example Takahashi (2014) and Khan and Thomas (2008)) by different methods.

I have developed a method of state and value function reduction that allows to solve rather large heterogeneous agent models by linearization. This reduction process is fully automatized, and independent of the specific structure of the model such as the number of individual states. A key condition for the applicability of reduction is that the cross sectional distribution affects the utility or budget constraint of agents only through a small set of equilibrium quantities such as prices.

Linearization with state reduction is the basis for a variety of nonlinear solution methods. Simulating the model as a sequence of temporary equilibria where the expected continuation value function comes from the linearized solution can uncover asymmetries in the impulse responses that are typical for models with non-convex optimization problems. We can also compute solutions where the agents take these asymmetries into account in their continuation value function. This is achieved by backward iteration on an aggregate grid that was obtained from simulation. This solution can significantly deviate from the linearized solution.

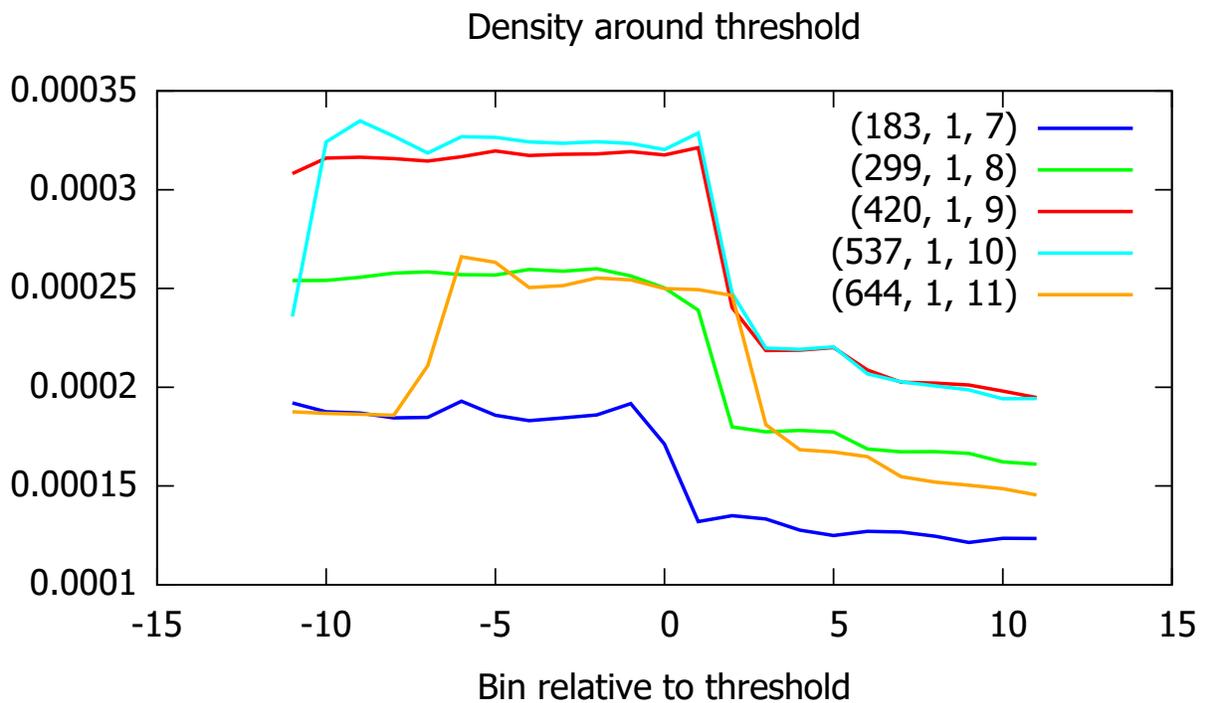
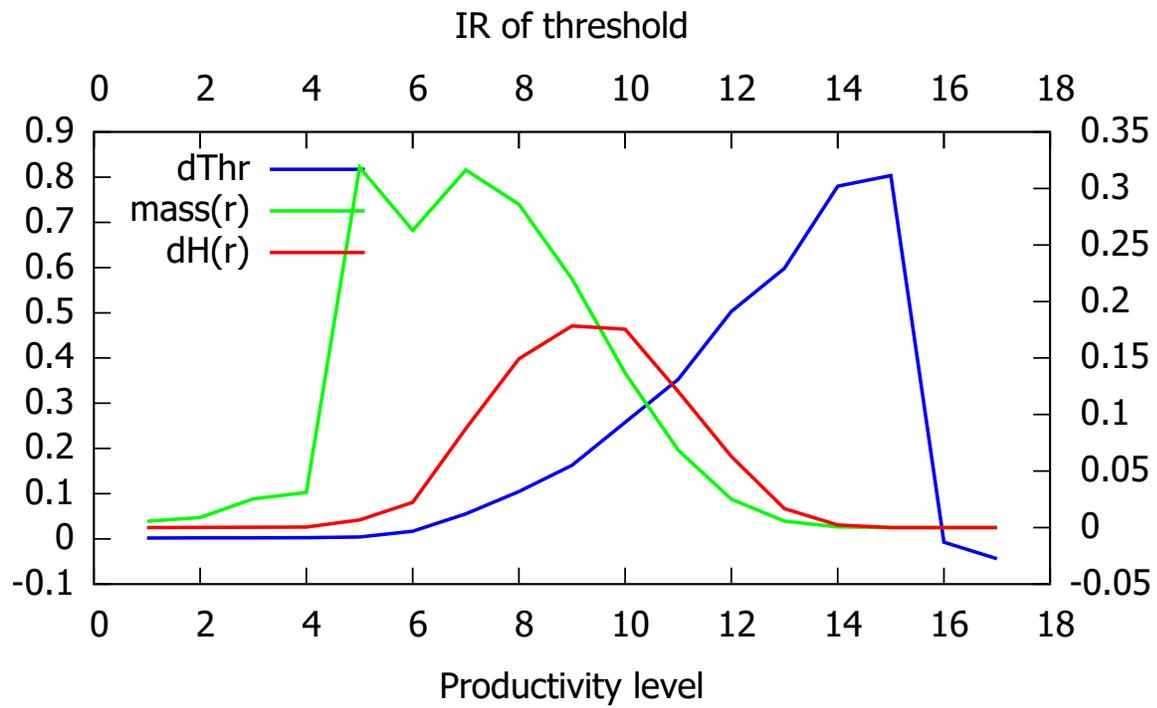


Figure 1: CK model, IR of work threshold

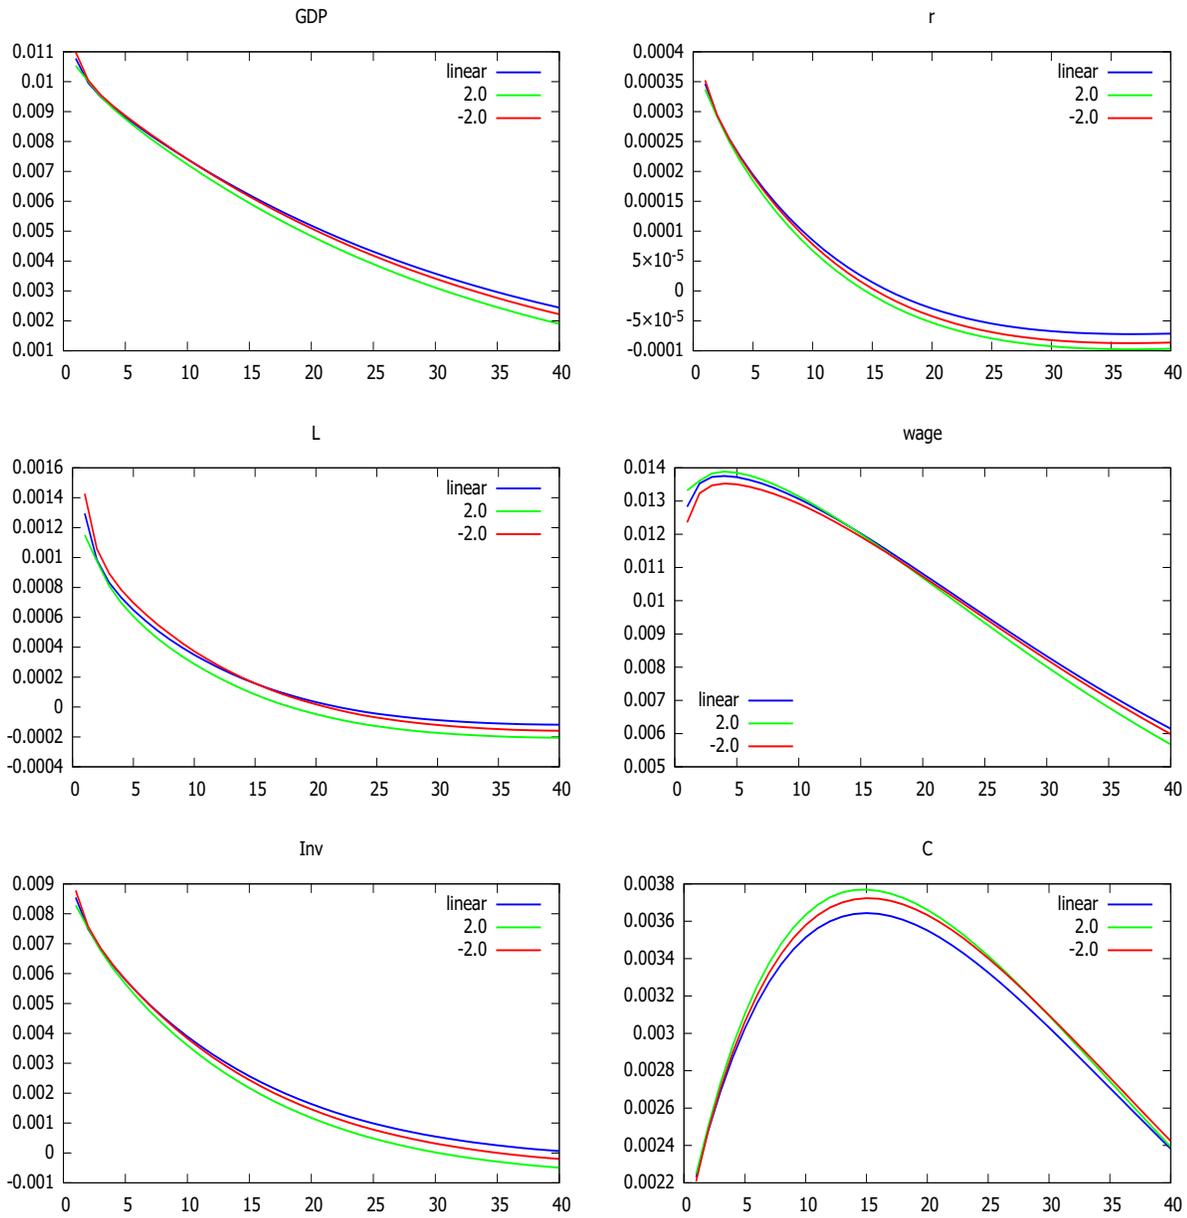


Figure 2: Impulse response to technology shock, **NonL1**

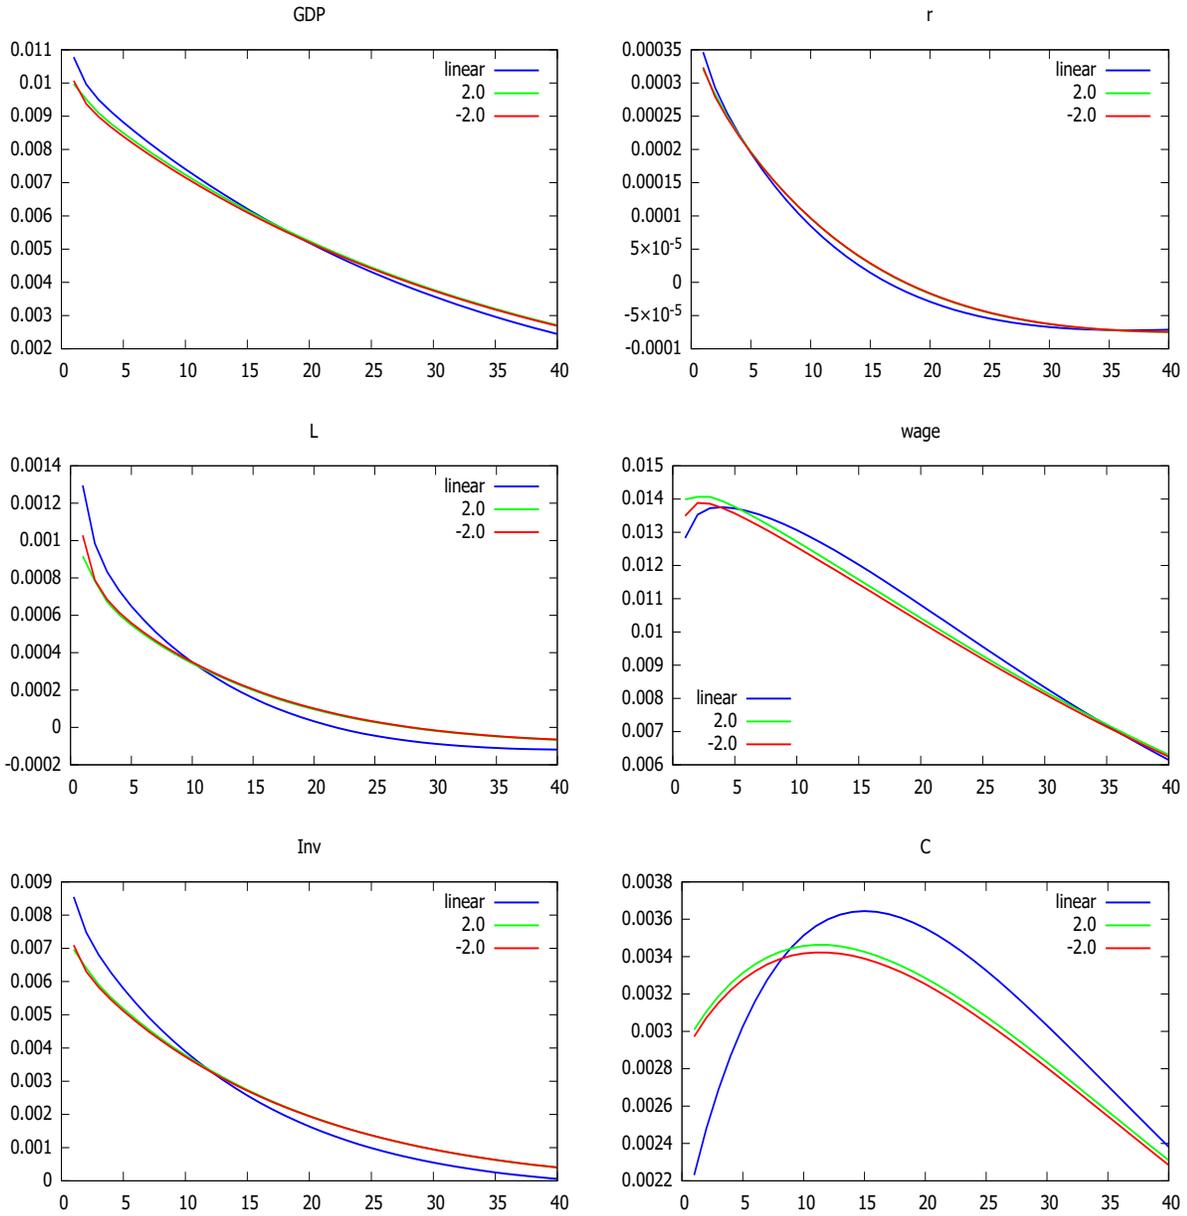


Figure 3: Impulse response to technology shock **NonL4**

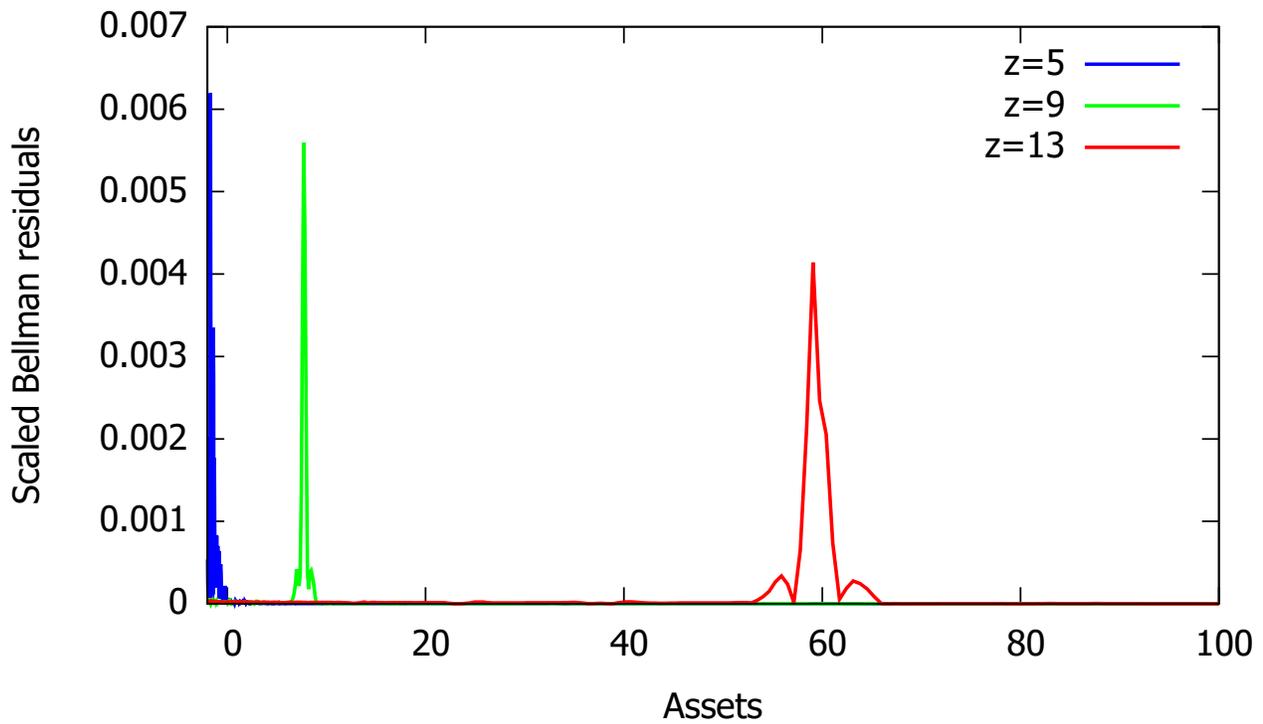
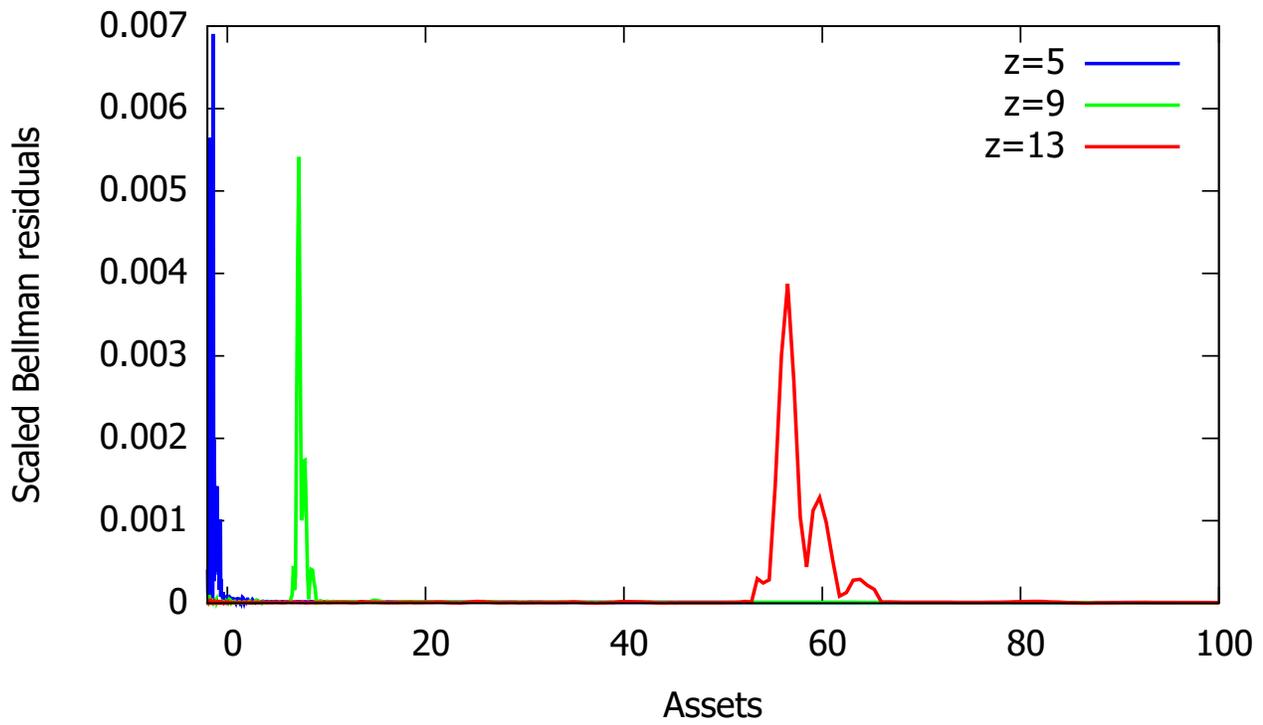


Figure 4: Bellman equation error in simulation

A Interpolating the value function

[TO BE FILLED IN.]

B Finite representation of the distribution

Denote by $\pi_k(i, j)$ the transition probability from histogram bin i to histogram bin j during period t , conditional on individual productivity being x_k at the beginning of t . We now have to approximate the transition probabilities. I assume that the discrete decision is the same at the lower and the upper end of the bin. If there is a threshold point within the bin, I treat this case simply as two separate bins. Denote by $a_0 = a(\bar{a}_i^D, x_k; \mu, \lambda)$ the continuous decision taken at the lower end of the bin, $a_1 = a(\bar{a}_{i+1}^D, x_k; \mu, \lambda)$ the continuous decision taken at the upper end. We make the assumption that the histogram bins are so small that the continuous decision over this range is well approximated by a linear function.

For simplicity of exposition we assume that $a_0 \leq a_1$; the change of formulas for the opposite case is straightforward. Denote by ι_0 and ι_1 the indices of the histogram bin in which a_0 and a_1 lie, respectively. We look for a transition law that preserves the expected value of a' , assuming that $a(\cdot)$ is linear on the bin, and that the density is constant on bins. Defining $\hat{a}_j \equiv \frac{\bar{a}_j^D + \bar{a}_{j+1}^D}{2}$, this can be written as

$$\hat{a} \equiv \frac{a_0 + a_1}{2} = \sum_{j=1}^{n_x} \pi_k(i, j) \hat{a}_j \quad (71)$$

We have to distinguish the following three cases.

1. $\iota_1 = \iota_0$. The image of bin i under the mapping $a(\cdot, x_k; \mu, \lambda)$ is contained in the bin ι_1 . To preserve expected value, we have to allow for a positive probability of going to an adjacent bin. We choose

$$\pi_k(i, \iota_0) = \begin{cases} \frac{\hat{a} - \hat{a}_{\iota_0+1}}{\hat{a}_{\iota_0} - \hat{a}_{\iota_0+1}}, & \text{if } \hat{a} \geq \hat{a}_{\iota_0} \\ \frac{\hat{a} - \hat{a}_{\iota_0-1}}{\hat{a}_{\iota_0} - \hat{a}_{\iota_0-1}}, & \text{otherwise} \end{cases} \quad (72)$$

with $\pi_k(i, \iota_0 + 1) = 1 - \pi_k(i, \iota_0)$ in the first of these case, and $\pi_k(i, \iota_0 - 1) = 1 - \pi_k(i, \iota_0)$ in the second one.

2. $\iota_1 = \iota_0 + 1$. The image of bin i is contained in two adjacent bins. We preserve expected values by choosing $\pi_k(i, \iota_0) = \frac{\hat{a} - \hat{a}_{\iota_1}}{\hat{a}_{\iota_0} - \hat{a}_{\iota_1}}$ and $\pi_k(i, \iota_0 + 1) = 1 - \pi_k(i, \iota_0)$.
3. $\iota_1 = \iota_0 + 2$. Now we split the probability mass between three intervals. This gives us the flexibility to match both the first and the second moments of the conditional distribution. We first try to do this by solving this as a linear problem. This can fail in the sense that the probability for one of the intervals is negative. This will never

be the middle interval, it must be either ι_0 or ι_1 . We set this probability to 0 and use the remaining two intervals so as to match the conditional mean. This is the choice that minimizes the cross-sectional variance, and comes as close as possible to the target variance.

4. $\iota_1 > \iota_0 + 2$. If the mean \hat{a} lies within one of the middle intervals, it is certainly possible to match both mean and variance, because one can then match the mean using only middle intervals, and this will certainly have a variance that is lower than $Var(a)$. One can then use the outer intervals to achieve a mean preserving spread, until the variance is matched. Since we have more than two free parameters, there is in general a continuum of choices that matches both moments. It would be natural to assign the interior bins probabilities that are proportional to their respective widths; in other words, to treat the interior bins just like one big bin. Since this will not always be possible, I proceed in the following iterative manner:

- (a) I treat the bins $\iota_0 + 1, \dots, \iota_1 - 1$ as one bin, and assign the probabilities as explain above in the case $\iota_0 + 2 = \iota_1$.
- (b) If one of the outer bins has negative probability, I drop this bin from the list. Then I treat again the inner ones of the remaining bins as one bin and go back to step 4.
- (c) If we are left with only three bins, we proceed with step 3.

This procedure will always lead to a solution that matches both moments as long as \hat{a} lies within the interior bins. If this is not the case, it is not in general possible to match both moments.

These rules define the transition matrix Π_k , Denote by $\pi_k(i)$ the mass of households with productivity x_k in histogram bin i , and by π_k the vector containing the $\pi_k(i)$'s. The total distribution is then characterized by stacking all the π_k 's into the big vector π . The transition matrix for the asset distribution from the beginning to the end of the period is then given by the block-diagonal matrix

$$\Pi = \begin{bmatrix} \Pi_1 & 0 & \dots & 0 & 0 \\ 0 & \Pi_2 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & \Pi_{n_x-1} & 0 \\ 0 & 0 & \dots & 0 & \Pi_{n_x} \end{bmatrix} \quad (73)$$

References

- Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf (2018). When Inequality Matters for Macro and Macro Matters for Inequality. *NBER Macroeconomics Annual* 32(1), 1–75.
- Antoulas, A. C. (1999). Approximation of linear dynamical systems. In J. Webster (Ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering, volume 11*. Wiley.
- Antoulas, A. C. (2005). *Approximation of Large-Scale Dynamical Systems*. SIAM.
- Bloom, N., M. Floetotto, N. Jaimovich, I. S. Eksten, and S. J. Terry (2018, May). Really Uncertain Business Cycles. *Econometrica* 86(3), 1031–1065.
- Boppart, T., P. Krusell, and K. Mitman (2018). Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control* 89(C), 68–92.
- Chang, Y. and S.-B. Kim (2007). Heterogeneity and aggregation: Implications for labor-market fluctuations. *American Economic Review* 97(5), 1939–1956.
- Costain, J. and A. Nakov (2011). Distributional dynamics under smoothly state-dependent pricing. *Journal of Monetary Economics* 58(6), 646–665.
- Judd, K. L., L. Maliar, and S. Maliar (2012). Merging simulation and projection approaches to solve high-dimensional problems. NBER Working Papers 18501, National Bureau of Economic Research, Inc.
- Khan, A. and J. K. Thomas (2008). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica* 76(2), 395–436.
- Krusell, P. and A. A. Smith (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy* 106(5), 867–96.
- Kubler, F. and S. Scheidegger (2018). Self-justified equilibria: Existence and computation. Manuscript, University of Zurich.
- McKay, A. and R. Reis (2016, January). The Role of Automatic Stabilizers in the U.S. Business Cycle. *Econometrica* 84, 141–194.
- Reiter, M. (2009). Solving heterogeneous agent models by projection and perturbation. *Journal of Economic Dynamics and Control* 33(3), 649–665.
- Reiter, M. (2010a). Approximate and almost-exact aggregation in dynamic stochastic heterogeneous-agent models. IHS Working Paper 258.
- Reiter, M. (2010b). Solving the incomplete markets model with aggregate uncertainty by backward induction. *Journal of Economic Dynamics and Control* 34(1), 28–35.
- Reiter, M., T. Sveen, and L. Weinke (2013). Lumpy investment and the monetary transmission mechanism. *Journal of Monetary Economics* 60(7), 821–834.

- Takahashi, S. (2014, April). Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations: Comment. *American Economic Review* 104(4), 1446–60.
- Young, E. R. (2010). Solving the incomplete markets model with aggregate uncertainty using the krusell-smith algorithm and non-stochastic simulations. *Journal of Economic Dynamics and Control* 34(1), 36–41.