

Alternative Methods for Solving Heterogeneous Firm Models

Stephen J. Terry*

April 9, 2014

Abstract

This paper implements and compares four alternative techniques for the solution of heterogeneous agents business cycle models within the lumpy capital adjustment framework. The widespread Krusell Smith algorithm consistently delivers high accuracy and economic implications quantitatively similar to other bounded rationality, projection-based approaches, but it does so at the cost of high computational intensity. The Parametrized Distributions and Explicit Aggregation methods yield important speed gains but reduced accuracy. The conceptually distinct Projection plus Perturbation method implies qualitatively similar economic results, even more dramatic reductions in computational cost, as well as an important scalability of the aggregate state space. A code package implementing each solution method is available online.

Keywords: Heterogeneous Agents, Computational Methods, Lumpy Investment

JEL: C63, E22, E32

*stephenjamesterry@gmail.com, Stanford University, Department of Economics, 579 Serra Mall, Stanford, CA 94305. All of the code used to produce the results in this paper can be found at Stephen Terry's website: <https://sites.google.com/site/stephenjamesterry/>. This research was supported by a Bradley dissertation fellowship from the Stanford Institute for Economic Policy Research. This paper was improved by comments from Brent Bundick, Itay Saporta-Eksten, and Nick Bloom. A portion of this research was completed as a visitor at the Federal Reserve Bank of Richmond, but this paper does not necessarily reflect the views of the Federal Reserve Bank of Richmond or the Federal Reserve System.

1 Introduction

Heterogeneous agent business cycle models offer the attractive possibility of combining a fully fledged business cycle structure with rich, testable implications for the cross-section of consumer or firm behavior. However, such frameworks, with seminal examples given by the incomplete markets model of Krusell and Smith (1998) and the heterogeneous firms model of Khan and Thomas (2008), pose several practical challenges for researchers. First, their solution and simulation are computationally intensive. Second, the traditional solution techniques used for these models, such as the Krusell Smith (KS) algorithm, rely on bounded rationality and aggregation assumptions and must be evaluated ex-post for the internal consistency of these assumptions. To help guide researchers around these issues in the practical solution of the incomplete markets model, many papers provide alternative solution techniques and computational strategies.¹ These advances profitably improve the speed and accuracy of solutions of the incomplete markets model, but the literature lacks a comprehensive analysis of their applicability to the heterogeneous firms context, which encompasses a fundamentally different economic and computational environment.²

This paper seeks to provide such a comparison of solution techniques specifically targeted towards the solution of the heterogeneous firms model. The Khan and Thomas (2008) model is a natural framework on which to base such a comparison because of the large number of papers using a similar underlying structure.³ The heterogeneous firms framework here combines aggregate uncertainty in the form of aggregate productivity shocks together with lumpy capital adjustment costs and a rich cross-sectional distribution of idiosyncratic productivity shocks and capital holdings. I adapt four existing algorithms to the heterogeneous firms structure, implementing each solution technique and comparing them along multiple dimensions: their simulated business cycle moments, cross-sectional investment rate distributions, impulse response functions, internal accuracy, as well as the computational burden posed by each algorithm. I implement the following four algorithms:

1. the traditional KS approach as first adapted by Khan and Thomas (2008),
2. the Parametrized Distributions (PARAM) algorithm due to Algan et al. (2008, 2010a),
3. the Explicit Aggregation (XPA) method of Den Haan and Rendahl (2010) as first adapted by Sunakawa (2012),

¹See, among others, Algan et al. (2010b), Algan et al. (2008, 2010a), Den Haan and Rendahl (2010), Den Haan (2010b), Maliar et al. (2010), Reiter (2010c), and Young (2010).

²In particular, the heterogeneous firms model requires a discrete investment choice by agents and use of Bellman equations rather than Euler equations (as in the incomplete markets model) to characterize optimal policies. The incomplete markets model can also be solved very efficiently at the microeconomic level through use of the Endogenous Grid Points method of Carroll (2006), while such a technique is unavailable for the lumpy investment problem of firms in the heterogeneous firms model.

³See, among others, Gourio and Kashyap (2007), Bloom et al. (2012), Khan and Thomas (2013), Bachmann and Bayer (2013), Bachmann and Ma (2012), Bachmann et al. (2013), with earlier work in Khan and Thomas (2003) and Thomas (2002). A complementary literature in heterogeneous agent state-dependent pricing models, typically dependent on the KS algorithm for solutions, includes papers by Vavra (2014), Klenow and Willis (2007), Knotek (2010), Knotek and Terry (2008), and many others.

4. and the Projection plus Perturbation (REITER) solution technique of Reiter (2009).

A quick word about the choice of these four techniques for comparison is in order. The KS approach is a natural and important choice because of its wide use in the heterogeneous firms literature to date. The PARAM algorithm is attractive both because it has been studied comprehensively in the context of the incomplete markets model but also because it bears conceptual similarity to another approach, the Backward Induction algorithm of Reiter (2010c). The XPA approach has been studied previously as a solution method for the heterogeneous firms model in Sunakawa (2012), and for comparability I rely on that paper’s adaptation of the original Den Haan and Rendahl (2010) technique.⁴ Finally, the REITER approach is an important addition to the algorithms considered here because it is conceptually distinct. The REITER method relies upon a linear perturbation approach in aggregates together with a rich, discretized projection problem at the microeconomic level to solve an approximation to the rational expectations equilibrium. For each solution method, code is readily available online.⁵

Three main conclusions can be drawn from the comparisons in this paper. First, the KS algorithm compares very favorably with the other solution techniques considered here. The KS routine offers a high degree of internal accuracy and delivers economic implications well in line with those of the other bounded rationality projection-based solution methods. This result should be interpreted as a favorable robustness check to the large number of papers relying on the KS algorithm in the heterogeneous firms context, including the original work by Khan and Thomas (2008). However, these advantages do come at the cost of fairly high computational intensity within both model solution and simulation steps.

Second, the other two bounded rationality solution techniques based on an approximation of the aggregate state space and projection methods, PARAM and XPA, both afford important speed gains relative to KS by avoiding simulation within the solution routine itself. As pointed out by Sunakawa (2012), these speed gains can open up the possibility of new applications for heterogeneous firms models, and the business cycle moments, micro-level investment rate distribution, and simulated impulse responses to productivity shocks are closely in line with those delivered by the KS algorithm. However, reduced computational burden comes at the cost of less internal accuracy for both the PARAM and XPA approaches. Furthermore, although model solution time is reduced by use of these algorithms, simulation for calculation of business cycle moments or impulse response analysis remains costly.

Third, the REITER approach stands apart from the KS, PARAM, and XPA algorithms for several reasons. The method is based on an alternative equilibrium concept, i.e. a perturba-

⁴There are several differences between my work and Sunakawa (2012), however. While that paper usefully draws out a potential application of speed gains from the XPA algorithm in structural estimation of the aggregate technology process, I focus here instead on a comparison of the solution techniques themselves. Also, my implementation of the XPA approach uses a bias-correction step within the forecasting structure based on a steady-state solution of the model which is different from the approach in Sunakawa (2012).

⁵See <https://sites.google.com/site/stephenjamesterry/>.

tion approximation to the full rational expectations equilibrium.⁶ The resulting solution to the heterogeneous firms model delivers qualitatively similar economic results as the traditional KS algorithm, although the aggregate dynamics do systematically differ in some ways discussed in more detail below. In practical terms, the use of perturbation methods delivers speed gains far above those achieved by either the PARAM or XPA algorithms. Perhaps more significantly, however, the REITER approach offers scalability of the macroeconomic complexity of heterogeneous agents models by providing a means for inclusion of a richer aggregate state space than would be currently tractable using a projection-based approach subject to the curse of dimensionality.

Section 2 lays out the model and calibration, a direct simplification of Khan and Thomas (2008). Section 3 provides a brief overview of each of the four solution techniques implemented in the paper. In Section 4, I compare the resulting simulations, impulse responses, accuracy, and time requirements of each solution method. In an example in Section 5, I extend the baseline model to include two additional aggregate state variables using the REITER method. Section 6 concludes. Appendix A contains detailed explanations of the solution algorithms as applied to the heterogeneous firms model, Appendix B contains some practical details on their numerical implementation as well as additional accuracy checks and figures, Appendix C discusses the simulation technique used to generate nonlinear impulse responses, and Appendix D discusses an extension to the aggregate state space of the model solved using the REITER approach.

2 Model and Calibration

The model solved in this code is a simplified version of the model presented and solved in Khan and Thomas (2008). The simplification involves only the removal of maintenance investment and trend investment growth, but crucially maintains aggregate uncertainty, the discrete nature of the investment decision, and idiosyncratic productivity shocks at the microeconomic level. Interested readers can find much more detail on the assumptions underlying this economic structure in Khan and Thomas (2008).

2.1 Households

A unit mass continuum of identical households trade a complete set of state-contingent claims, own a unit mass distribution of firms, and have flow utility given by $U(C, 1-N) = \log(C) + \phi(1-N)$, $\phi > 0$. C represents aggregate consumption, and N represents aggregate labor supply. For our purposes, there are two implications of the household problem from Khan and Thomas (2008) of importance for the solution of the model. First, firm value maximization is equivalent to maximization of

⁶Interestingly, in a projection-based solution technique distinct from the projection plus perturbation approach by Reiter (2009, 2010a,b), Gordon (2011) also discusses a means of solving an approximation to the full rational expectations equilibrium of the incomplete markets model with the cross-sectional distribution as a state variable, using a dramatically reduced Smolyak grid for the projection routine. Note also that although Reiter (2009) and this paper's analysis use a linear perturbation, higher-order approximations of the same equilibrium system are possible, as suggested by Reiter (2010b).

dividends weighted by a marginal utility price p . Second, household labor supply optimality and linear disutility of labor imply a trivial relationship between the wage and price p :

$$p = \frac{1}{C(A, \mu)}, \quad w(A, \mu) = \frac{\phi}{p(A, \mu)}.$$

Above, prices and wages are written in terms of an aggregate state (A, μ) including aggregate productivity A and a cross-sectional distribution μ of capital and productivity, both of which are discussed in more detail below.

2.2 Firms

In each period there is a distribution of firms $\mu(z, k)$ over idiosyncratic productivity and capital levels z and k .⁷ Individual firms are subject to both idiosyncratic and aggregate productivity shocks, which are exogenous and are assumed to follow independent AR(1) processes in logs:

$$\log(A') = \rho_A \log(A) + \sigma_A \varepsilon_A, \quad \log(z') = \rho_z \log(z) + \sigma_z \varepsilon_z$$

where innovations to both processes are iid $N(0, 1)$. The state vector for an individual firm is given by $(z, k; A, \mu)$, which contains both the idiosyncratic states for that firm as well as an aggregate state including productivity and all distributional information. Firms also receive a random draw of fixed capital adjustment costs in each period, discussed below. Conditional upon idiosyncratic productivity and capital (z, k) , a firm that chooses labor input n produces output given by the decreasing returns to scale technology $y(z, k, n, A) = zAk^\alpha n^\nu$, where $\alpha + \nu < 1$.

In a rational expectations equilibrium there is a known transition mapping Γ_μ tracking the evolution of the cross-sectional distribution, as well as a mapping Γ_p from the aggregate state to the marginal utility of the representative household-owner p :

$$\mu' = \Gamma_\mu(\mu, A), \quad p = \Gamma_p(\mu, A).$$

Recall that the wage is a function of the household marginal utility given linear labor disutility, so that these two aggregate mappings fully characterize the aggregate structure of the economy from the perspective of an individual firm. Then, in each period, a firm receives a stochastic draw of a fixed capital adjustment cost ξ , given in units of labor. The firm value function V , adjusted by the marginal utility of the representative households, is therefore given by

$$V(z, k; A, \mu) = \mathbb{E}_\xi \tilde{V}(z, k, \xi; A, \mu).$$

Once a firm receives a draw of a stochastic adjustment cost $\xi \sim G(\xi)$, the firm faces a choice between paying the capital adjustment cost or not adjusting the capital stock

$$\tilde{V}(z, k, \xi; A, \mu) = \max \{ -\xi p(A, \mu) w(A, \mu) + V^A(z, k; A, \mu), V^{NA}(z, k; A, \mu) \},$$

⁷Note that although I use the term “firm” throughout the paper for simplicity, such models are typically disciplined by the use of *establishment* data at the microeconomic level, treating individual establishments as separately operating business units for the purposes of the model. However, for a recent treatment of a heterogeneous firms structure centering on the distinction between establishments and firms, see Kehrig and Vincent (2013).

where the value upon adjustment V^A is given by optimization over investment and labor

$$V^A(z, k; A, \mu) = \max_{k', n} \{p(A, \mu) (zAk^\alpha n^\nu - k' + (1 - \delta)k - w(A, \mu)n) + \beta \mathbb{E}_{\Gamma_\mu, z', A'} V(z', k'; A', \mu')\}.$$

If a firm chooses not to adjust its capital stock, then it must face a dynamic return V^{NA} which involves optimization of only the labor input n holding capital levels fixed at the (depreciated) level from last period:

$$V^{NA}(z, k; A, \mu) = \max_n \{p(A, \mu) (zAk^\alpha n^\nu - w(A, \mu)n) + \beta \mathbb{E}_{\Gamma_\mu, z', A'} V(z', (1 - \delta)k; A', \mu')\}.$$

The trivial nature of the discrete choice problem leads to a cutoff rule for capital investment, such that firms adjust their capital stock if and only if the adjustment cost draw ξ is less than a cutoff level

$$\xi^*(z, k; A, \mu) = \frac{V^A(z, k; A, \mu) - V^{NA}(z, k; A, \mu)}{\phi},$$

where the numerator reflects the gains from capital adjustment relative to inaction and the denominator's adjustment by labor disutility ϕ is required to convert from marginal-utility to labor units. Further the distribution of lumpy capital adjustment costs is assumed to be given by $G(\xi) = U(0, \bar{\xi})$, where $\bar{\xi} > 0$ indexes the level of the adjustment friction in the economy.

2.3 Equilibrium

An equilibrium represents a set of firm value functions $\tilde{V}, V, V^A, V^{NA}$, firm policies and adjustment thresholds k', n, ξ^* , prices $p(A, \mu), w(A, \mu)$, and mappings Γ_μ, Γ_p such that

- Firm capital adjustment choices and policies conditional upon adjustment satisfy the Bellman equations defining V, V^A, V^{NA} above, and therefore firm capital transitions are given by

$$k'(z, k, \xi; A, \mu) = \begin{cases} k'(z, k; A, \mu), & \xi < \xi^*(z, k; A, \mu) \\ (1 - \delta)k, & \xi \geq \xi^*(z, k; A, \mu) \end{cases}.$$

- The distributional transition rule used in the calculation of expectations above by firms is consistent with the aggregate evolution of the distributional state

$$\Gamma_\mu(z', k') = \int \int \int I_A(z, k) d\mu(z, k) dG(\xi) d\Phi(\varepsilon_z)$$

$$A(z', k', \xi, \varepsilon_z; A, \mu) = \{(z, k) | k'(z, k, \xi; A, \mu) = k', z' = \rho_z z + \sigma_z \varepsilon_z\}, \Phi(x) = \mathbb{P}(\varepsilon_z \leq x)$$

- Aggregate output, investment, and labor are consistent with the current distribution μ and firm policies:

$$Y(A, \mu) = \int \int zAk^\alpha n(z, k, \xi; A, \mu)^\nu d\mu(z, k) dG(\xi)$$

$$I(A, \mu) = \int \int (k'(z, k, \xi; A, \mu) - (1 - \delta)k) d\mu(z, k) dG(\xi)$$

$$N(A, \mu) = \int \int n(z, k, \xi; A, \mu) d\mu(z, k) dG(\xi) + \int \int_0^{\xi^*(z, k; A, \mu)} dG(\xi) d\mu(z, k)$$

- Aggregate consumption satisfies the resource constraint

$$C(A, \mu) = Y(A, \mu) - I(A, \mu).$$

- The households are on their optimality schedules for savings and labor supply decisions, i.e. the first order conditions defining marginal utility and wages hold, and the price mapping is consistent

$$p(A, \mu) = \Gamma_p(A, \mu) = \frac{1}{C(A, \mu)}, \quad w(A, \mu) = \frac{\phi}{p(A, \mu)}.$$

- Aggregate productivity follows the assumed AR(1) process in logs.

2.4 Calibration

The parameter choices used in the solution method comparison below are those chosen by Khan and Thomas (2008). The parameter choices reflect an annual frequency and positive levels of capital adjustment costs at the firm level, as summarized in Table 1. Given that this paper is concerned with the comparison of numerical solution techniques, and that the model is a simplified version of the original structure, these parameter choices should be taken as purely illustrative.

Table 1: Model Calibration

Parameter	Role	Value	Parameter	Role	Value
α	Capital elasticity	0.256	ν	Labor elasticity	0.640
β	Discount rate	0.977	ϕ	Labor disutility	2.4
δ	Capital depreciation	0.069	ρ_A	Aggregate persistence	0.859
σ_A	Aggregate volatility	0.014	ρ_z	Idiosyncratic persistence	0.859
σ_z	Idiosyncratic volatility	0.022	$\bar{\xi}$	Capital adjustment costs	0.0083

Note: The calibration above is based on Khan and Thomas (2008), Table I, reflecting an annual calibration of the heterogeneous firms model with nonconvex capital adjustment costs.

3 Solution Methods Overview

3.1 Krusell Smith Algorithm

Khan and Thomas (2008) in the original exposition of the heterogeneous firms model use the first algorithm considered here, the KS approach. Their algorithm extends the one proposed in Krusell and Smith (1998) for use in the incomplete markets model and bases the general equilibrium components of the solution on a bounded rationality approach.

When solving their dynamic problem, firms approximate the intractable distribution $\mu(z, k)$ over idiosyncratic productivity and capital with some moments m . In practice, m is chosen to simply be the mean aggregate level of capital K . Given this approximation, two sets of forecast rules provide expectations for firms of both the aggregate level of consumption and the evolution of aggregate capital itself. Therefore, the intractable state vector $(z, k; A, \mu)$ for the firm problem

discussed above is replaced by $(z, k; A, m)$, and the transition and price mappings are replaced by forecast rules $\hat{m}' = \hat{\Gamma}_m$ and $\hat{p} = \hat{\Gamma}_p$. In practice, the forecast rules are assumed to take a loglinear form conditional upon aggregate productivity, although the algorithm allows for more flexibility in theory.

Solution of the model involves repeated simulation to obtain a fixed point on the forecast mappings for firms. First, a particular set forecast rules is assumed, allowing for the creation of value functions for the idiosyncratic firm problems using the simplified state space $(z, k; A, m)$. Then, given the idiosyncratic firm value functions, the model is simulated. Throughout this paper unless otherwise noted, aggregate and productivity shocks in the KS method, as well as the PARAM and XPA techniques, are discretized using the Markov chain approximation process of Tauchen (1986). Also, unless otherwise noted, simulation of the cross-sectional distribution of productivity and idiosyncratic capital makes use of the nonstochastic or histogram-based approach in Young (2010) rather than relying on simulation of individual firms. This histogram-based simulation technique avoids the sampling error associated with individual firm simulation and in practice is less computationally burdensome. In each period, market-clearing consumption must be found by repeated reoptimization of firm policies given a guessed price level, the currently simulated histogram of firm states, and continuation values and expectations as dictated by the firm values computed using the current forecast rule guess. Finally, after simulation is complete, the forecast rules are updated on the simulated aggregates. The entire process repeats until a fixed point is achieved in the forecast rules. Further details on the KS solution algorithm, as well the practical choices surrounding the numerical solution of the model can be found in Appendixes A and B.

3.2 Parametrized Distributions Algorithm

The PARAM algorithm is based on the work of Algan et al. (2008, 2010a), which was done in the context of the incomplete markets model, and the solution technique bears heavy resemblance to the Backward Induction algorithm of Reiter (2010c). To my knowledge, this paper represents the first application of the PARAM algorithm to a version of the Khan and Thomas (2008) model. The PARAM approach, like the KS method, relies upon a bounded rationality assumption, the replacement of the cross-sectional distribution $\mu(z, k)$ with a set of moments m in the dynamic problem of an individual firm.

However, and contrasting with the KS assumption of forecast rules $\hat{\Gamma}_m$ and $\hat{\Gamma}_p$ for the aggregate moments and prices, the PARAM approach instead relies upon a set of “reference moments” m^{ref} , equal to the higher-order centered moments of the cross-sectional distribution of firm capital, conditional upon idiosyncratic productivity. The moments m included in the approximate state for firm dynamic problems are either a subset of or implied by the reference moments m^{ref} , and they can be drawn from a steady-state solution of the model with no aggregate uncertainty if solution of the model without simulation is desired.

Solution of the model involves value function iteration with the simplified state space of $(z, k; A, m)$.

Given a guess for the firm value function which can be used in construction of the continuation value in the firm Bellman equations, optimization and calculation of the next iteration of the value function requires calculation of two objects: market-clearing price $p(A, m)$ for construction of current-period returns, and next-period moments m' for input into continuation values. Both p and m' can be computed within the value function iteration step quite naturally by using fixed point iteration. After guessing values for (p, m') , firm policies are computable, and implied aggregates can be obtained by integrating over the cross-sectional distribution of firm-level productivity and capital (z, k) . Such integration is the key step within the PARAM algorithm and is performed numerically using flexible exponential functional forms for the density of the model which exactly match the aggregate moments m together with the higher-order reference moments in the cross-section. Iteration on prices and next-period moments continues until a fixed point is achieved, at which point the next value function iteration step is taken. Once the value function converges, the model is solved.

Note that crucially the PARAM approach does not require simulation and therefore leads to large time savings relative to the KS algorithm's solution. However, if desired, new values for reference moments can be computed from simulation and updated until an outside fixed point is achieved, similar to the KS technique. In either case, however, simulation in each period requires a fixed-point iteration routine over market-clearing prices and next-period moments, similar to the process within the model solution step and involving integration over parametrized cross-sectional densities. See Appendix A for further details on the PARAM algorithm, as well as the functional forms used for the assumed cross-sectional densities.⁸

3.3 Explicit Aggregation Algorithm

The XPA solution method relies upon the techniques suggested by Den Haan and Rendahl (2010), as first adapted and applied to the heterogeneous firms model by Sunakawa (2012). The algorithm is essentially identical to the KS method, also making use of a bounded rationality assumption replacing the aggregate state space (A, μ) with an approximation based on moments (A, m) and using forecasting rules $\hat{\Gamma}_m$ and $\hat{\Gamma}_p$ for the aggregate moments and prices.

However, there is one main difference between the two techniques. XPA replaces the simulation step of the KS routine with an aggregation across a fixed cross-sectional distribution which is made feasible through the substitution of aggregate states into idiosyncratic policies. In other words, once value functions and policies are obtained based on a simplified state space of $(z, k; A, m)$ and the posited forecast rules, market-clearing prices are obtained by integrating policies over the constant exogenous ergodic distribution of z and ignoring heterogeneity in idiosyncratic capital k . Afterwards, the forecast rules can be updated from the moments and prices generated in this manner until a fixed point is achieved.

⁸The algorithm laid out in Appendix A, as well as the code posted online, allows for use of fixed steady-state reference moments and alternatively for updating of these moments through simulation. I use only the former in this paper, because doing so by itself already yields economic implications similar to the KS and XPA techniques.

As the original work by Den Haan and Rendahl (2010) noted, substitution of aggregate states into idiosyncratic policies creates a Jensen’s inequality-type bias in the forecast system (and hence the resulting aggregate solution) which can be ameliorated in a straightforward way by use of information from the steady-state solution.⁹ Importantly, however, avoiding simulation within the solution step also allows for large time savings, as emphasized and put to interesting practical use for structural estimation of technology shocks by Sunakawa (2012).

3.4 Projection plus Perturbation Algorithm

The REITER algorithm is based on the work of Reiter (2009) in the context of the incomplete markets model. The REITER approach has gained traction recently in the analysis of models with nonconvex costs and heterogeneity, being used in Costain and Nakov (2011) and Reiter et al. (2009). To my knowledge, this paper is the first application of the approach to a version of the Khan and Thomas (2008) model by itself, although it should be noted that Reiter et al. (2009) analyzes a similar sticky-capital environment with the addition of a New Keynesian sticky-price structure.

The REITER solution method departs in two important ways from the KS, PARAM, and XPA approaches. First, and importantly, the algorithm solves an approximation to the full, rational expectations equilibrium of the model rather than relying upon a bounded rationality assumption to reduce the state space using a set of moments. Therefore, the REITER method can be used as a check on the bounded rationality assumptions of the KS, PARAM, and XPA algorithms. Second, the REITER method relies upon linear perturbation of the model around the steady-state of a model with no aggregate uncertainty, although it still preserves idiosyncratic nonlinearity through a discretization of the firm-level problem. By contrast, the methods considered so far have relied upon global solution techniques, conditional on their bounded rationality assumptions. Importantly, such a perturbation approach leads to drastically reduced computational requirements, although of course the global nature of the other solutions must serve as a useful cross-check for large simulated fluctuations away from the steady-state solution.

The REITER approach relies upon three steps. The first step imposes almost trivial computational cost: the solution of a steady-state model with no aggregate uncertainty but maintaining micro-level nonlinearity, using a discretization or histogram for idiosyncratic states (z, k) . Then, the second step writes the full, discretized rational expectations equilibrium as the solution to a system of nonlinear equations F . The system is a function of current and lagged values of a large endogenous vector X_t , as well as some exogenous aggregate shocks ϵ_t . In the application to the heterogeneous firms model, the endogenous vector includes aggregate productivity, the cross-sectional histogram weights on each idiosyncratic point, firm values at a set of discrete points, as well as some implied model aggregates including price, output, investment, and labor. Therefore, the system F

⁹In particular, the forecast rules can be adjusted so that they are consistent with a fixed point at the steady-state levels of capital and price.

must take into account Bellman equations, distributional transitions, and equilibrium conditions. The third step involves the application of standard techniques for the solution of dynamic linear rational expectations systems, such as the method of Sims (2002) or Christiano (2002), to the solution of the heterogeneous firms model. Through numerical differentiation, the system F can be written as a linear approximation around the steady-state solution of the model, and then the standard methods for the solution of linear models may be applied. Further discussions of the details of the REITER solution method can be found in Appendix A.

4 Comparing Solutions

This section compares the four alternative solutions to the heterogeneous firms model along multiple dimensions: reported business cycle aggregate series and moments, impulse response functions to an aggregate productivity shock, microeconomic moments of investment rates, internal accuracy and diagnostic statistics, and finally a comparison of computational time. Unless otherwise noted, comparisons across methods are conducted with comparable levels of computational intensity, i.e. the projection grid ranges and densities do not vary across methods, similar interpolation and optimization techniques are used when solving Bellman equations, and of course random exogenous shocks are held constant during simulation across methods.¹⁰ Specific details about the numerical choices made are available in Appendix B. Appendix C contains details on the procedure used to simulate impulse responses in the nonlinear discretized models.

4.1 Unconditional Simulation: Business Cycle Moments and Micro Investment

To begin the comparison, Figure 1 plots a representative 100-period portion of larger 2000-period simulation of the solutions from each technique, displaying log aggregate output, investment, labor input, exogenous productivity, and market-clearing price. For strict comparability we must maintain identical driving processes across the discretized aggregate productivity series of the KS, PARAM, and XPA methods and the REITER linearized solution, which by contrast takes continuous local shocks to aggregate productivity. To accomplish this, a set of continuous productivity shocks duplicating the discretized aggregate productivity process were trivially computed and input into the REITER solution to produce Figure 1. Figure 2, by contrast, displays business cycle aggregates from a representative portion of a distinct unconditional simulation, with identical length, of the REITER solution based on continuous exogenous productivity shocks drawn from $N(0, \sigma_A^2)$, which is labelled REITER-OWN in discussion of business cycle moments below.

¹⁰The exception to this rule is a higher density of the discretized cross-sectional distribution of capital used when solving the REITER method. The linearization of the equilibrium centers crucially around this distribution and is slightly more sensitive to grid size than the projection-based KS, PARAM, or XPA methods. Since accuracy statistics are not directly compared between the REITER method and other methods, and since notwithstanding the higher grid size the REITER method simulation times are still substantially *lower* than for other techniques, as discussed below, the difference does not affect the qualitative conclusions of the paper. More details are available in Appendix B.

Table 2: Unconditional Simulation Differences

Method	Output	Investment	Labor	Price
Mean % Difference from KS				
PARAM	-0.2760	-0.6892	-0.1144	0.1751
XPA	0.0053	0.2849	0.0079	0.0118
REITER	-0.1736	0.0710	-0.1219	0.0527

Note: Mean percentage differences between the business cycle simulations for the PARAM, XPA, and REITER solutions over a 2000-period unconditional simulation, relative to an analogous KS simulation are reported in this table, with columns representing the average value of $100(\log(X_t^{method}) - \log(X_t^{KS}))$ for each solution *method* and series X_t . The exogenous aggregate productivity process reflects a Markov chain discretized using the Tauchen (1986) procedure and is held constant across solution methods during the simulation. To achieve this, an identical simulated discretized productivity process is input directly into the KS, PARAM, and XPA solutions, while a series of continuous aggregate shocks exactly replicating the discretized productivity process is input into the REITER solution. For each method, the full 2000 period simulation for each solution method begins after 500 periods, with an initial burn-in period discarded to avoid the influence of initial conditions on the simulated aggregates.

Although the simulated fluctuations are quite similar across solution methods in Figure 1, a few patterns are immediately visible to the naked eye. First, mean differences between the simulated aggregates appear small. This result is confirmed by Table 2, which reports average differences between each series and a KS method baseline of well under 0.7% in magnitude throughout the simulation. However, and especially for the price and labor series, there appear to be some differences in volatility between the projection-based solutions KS, PARAM, and XPA and the linearized REITER method. In particular, the REITER-simulated price (labor) series are slightly less (more) volatile than the other three simulations, a fact that is reflected in the business cycle moments available in Tables 3 and 4.

Table 3: HP-Filtered Volatilities

Method	Output	Investment	Labor	Productivity	Price
KS	(0.02526)	3.9313	0.6937	0.5740	0.3780
PARAM	(0.023778)	3.7041	0.6518	0.6098	0.4371
XPA	(0.023928)	3.7098	0.6538	0.6060	0.4380
REITER	(0.026462)	3.7757	0.8774	0.5479	0.2372
REITER - OWN	(0.024431)	3.7909	0.8748	0.5499	0.2448

Note: The business cycle volatility moments of the KS, PARAM, XPA, and REITER solutions over a 2000-period unconditional simulation are reported in this table in the first four rows. The first column reports, in parentheses, the standard deviation of the HP-filtered log aggregate output series. The second through fifth columns report the ratio of the standard deviation of the aggregate HP-filtered log investment, labor, exogenous productivity, and price series to the output standard deviation in the first row. Following Khan and Thomas (2008), the HP-filter uses smoothing parameter 100. In the first four rows, the exogenous aggregate productivity process reflects a Markov chain discretized using the Tauchen (1986) procedure and is held constant across solution methods during the simulation. To achieve this, an identical simulated discretized productivity process is input directly into the KS, PARAM, and XPA solutions, while a series of continuous aggregate shocks exactly replicating the discretized productivity process is input into the REITER solution. In the fifth row, analogous moments from a continuous shock simulation for the REITER solution is reported, labelled “REITER-OWN.” For each case, the full 2000 period simulation for each solution method begins after 500 periods, with an initial burn-in period discarded to avoid the influence of initial conditions on the simulated aggregates.

The standard set of HP-filtered business cycle volatilities and ratios in Table 3 reveals broadly similar and familiar variability of aggregate output, as well as relative volatilities of productivity, across all four solution methods.¹¹ Furthermore, for all series, flow investment is at least several times more volatile than output, with the volatility of consumption (equal to that of price, given the log transformation of $p = \frac{1}{c}$ used here) smallest due to the smoothing effects of general equilibrium. Importantly, however, as noted above, the REITER solution delivers higher variability of labor as well as lower variability of prices or consumption than the projection-based techniques. This phenomenon is not due to the use of discrete shocks with the linearized REITER solution, as the same conclusion can be drawn from the moments of the REITER-OWN simulation based on continuous shocks. Rather, the moderate price-dampening appears to be an implication of the use of local perturbation techniques rather than projection in the aggregate state space.

Business cycle correlations with output reported in Table 4 reveal a similar message: broad similarities among the KS, PARAM, and XPA methods, with a bit less (more) comovement of price (labor) with aggregate output in the REITER simulations. As I will argue below, the perturbation approach of the REITER technique offers attractive scalability of the complexity of the aggregate state space, but the potential for differences in macroeconomic volatilities relative to the KS method in the heterogeneous firms context should be kept in mind by researchers using the technique.

Table 4: HP-Filtered Correlation with Output

Method	Output	Investment	Labor	Productivity	Price
KS	1.0000	0.9779	0.9660	0.9999	-0.8806
PARAM	1.0000	0.9734	0.9481	0.9974	-0.8811
XPA	1.0000	0.9702	0.9434	0.9962	-0.8709
REITER	1.0000	0.9845	0.9769	0.9952	-0.5921
REITER-OWN	1.0000	0.9845	0.9751	0.9953	-0.5919

Note: The business cycle correlations for the KS, PARAM, XPA, and REITER solutions over a 2000-period unconditional simulation are reported in this table in the first four rows. The first column reports the correlation of the HP-filtered log aggregate output series with itself, 1. The second through fifth columns report the correlation of the aggregate HP-filtered log investment, labor, exogenous productivity, and price series with the filtered output series. Following Khan and Thomas (2008), the HP-filter is performed with smoothing parameter 100. In the first four rows, the exogenous aggregate productivity process reflects the Markov chain discretized using the Tauchen (1986) procedure and is held constant across solution methods during the simulation. To achieve this, an identical simulated discretized productivity process is input directly into the KS, PARAM, and XPA solutions, while a series of continuous aggregate shocks exactly replicating the discretized productivity process is input into the REITER solution. In the fifth row, analogous moments from a continuous shock simulation for the REITER solution is reported, labelled “REITER-OWN.” For each case, the full 2000 period simulation for each solution method begins after 500 periods, with an initial burn-in period discarded to avoid the influence of initial conditions on the simulated aggregates.

One important advantage of the heterogeneous agents approach is the possibility of disciplining the parametrization of a model with evidence from the micro-level distribution of investment

¹¹Comparison of Tables 3 and 4 here to Table IV of Khan and Thomas (2008) will reveal some moderate quantitative differences between the business cycle moments of this paper’s KS simulation and their “GE lumpy” model. Given the removal of maintenance investment for expositional clarity, and the distinct discretizations of aggregate productivity, we would not expect identical results. Therefore, it should be emphasized that the correct analysis of the results in this paper is based on relative comparisons across methods, rather than direct comparison with Khan and Thomas (2008).

rates. A number of cross-sectional moments can be computed for each period of the unconditional simulation of the model within each solution method, based directly on the Young (2010)-style simulated histograms for the KS and XPA solutions, the parametrized cross-sectional densities and coefficients of the PARAM method, and the endogenous discretized vector of histogram weights available directly from the REITER simulation. Here, we focus on the moments analyzed by Khan and Thomas (2008), including the prevalence of inaction, large positive or negative investment “spikes,” and the overall likelihood of positive or negative investment. Table 5 reports the mean levels of these statistics, together with the first and second moments of the investment-rate distribution. Comfortingly, the methods deliver broadly similar implications for the cross-section of investment, although the cross-sectional dispersion of investment rates is lower for the REITER solution.¹² Around three-quarters of firms are inactive in each period, with around one-fifth of firms exhibiting both positive investment spikes or positive investment overall. Much smaller proportions of observations see negative investment spikes or rates.¹³

Table 5: Microeconomic Investment-Rate Moments

	KS	PARAM	XPA	REITER	REITER-OWN
	Simulated Mean Value				
$\frac{i}{k}$	0.0891	0.0887	0.0903	0.0797	0.0364
$\sigma\left(\frac{i}{k}\right)$	0.2224	0.2247	0.2290	0.0939	0.0624
$\mathbb{P}\left(\frac{i}{k} = 0\right)$	0.7348	0.7411	0.7405	0.7706	0.7169
$\mathbb{P}\left(\frac{i}{k} \geq 0.2\right)$	0.1904	0.1982	0.1906	0.1628	0.0904
$\mathbb{P}\left(\frac{i}{k} \leq -0.2\right)$	0.0234	0.0255	0.0245	0.0101	0.0472
$\mathbb{P}\left(\frac{i}{k} > 0\right)$	0.2209	0.2167	0.2152	0.1925	0.1196
$\mathbb{P}\left(\frac{i}{k} < 0\right)$	0.0443	0.0423	0.0443	0.0397	0.0755

Note: The rows of the table above report the mean value, across periods, of the indicated microeconomic moment of the cross-sectional distribution of investment rates $\frac{i}{k}$ from a 2000-period unconditional simulation of the KS, PARAM, XPA, and REITER methods. The first row reports the level of investment rates, the second row the cross-sectional standard deviation of investment rates, the third column the probability of investment inaction, the fourth (fifth) columns the probability of positive (negative) investment spikes larger in magnitude than 20%, and the sixth (seventh) columns the probability of strictly positive (negative) investment rates. The first four columns report values from simulations based on identical discretized aggregate productivity processes, while the fifth column, labelled REITER-OWN reports a distinct simulation with continuous aggregate productivity shocks input into the REITER solution. In all columns, the full 2000 period simulation for each solution method begins after 500 periods, with an initial burn-in period discarded to avoid the influence of initial conditions on the simulated aggregates.

¹²The one exception to this pattern is the REITER-OWN simulation, which is subject to a distinct simulation of continuously varying aggregate productivity shocks, and is not directly comparable to the discretized aggregate productivity simulation used to produce the other microeconomic moments in the first four columns of the table.

¹³It’s important to note that, as opposed to the model in Khan and Thomas (2008) which allows for costless maintenance investment, the simplified structure here results in higher levels of investment inaction and lower levels of negative investment, because firms do not have a low-cost option for engaging in small adjustments to their capital stock. Consistently across solution methods, the environment here therefore delivers moments broadly similar to those of the “Traditional Model” of Table II in Khan and Thomas (2008).

4.2 Impulse Response Functions

Now we turn to a comparison of the heterogeneous firms solutions based on conditional responses, or impulse response analysis, rather than unconditional simulations. At this point, some concrete decisions must inevitably be made about the manner in which to simulate the underlying object of interest, i.e. the average change in the forecast of a given series in response to a shock to aggregate productivity of a certain size. Two considerations will always face a researcher working with nonlinear discretized models like those considered here. First, given the nonlinear structure of the KS, PARAM, and XPA solutions, the average conditional response to a shock will depend both upon initial conditions and upon the size of the shock. Second, we may wish to consider a shock scaled to a certain average size, such as the calibrated standard deviation of the underlying true aggregate productivity process, but a discrete Markov chain only admits discrete innovations in the aggregate productivity series. Neither challenge is present with a linearized solution such as that available for the REITER method, since in that case a classical impulse response is computable directly from the coefficients defining the model solution, and the local linearity guarantees that for small perturbations the impulse response scales directly with shock size.

Table 6: IRF Simulation Differences

Method	Output	Investment	Labor	Price
Mean % Difference from KS				
PARAM	-0.1093	-0.3188	-0.0853	0.0248
XPA	-0.1035	-0.2742	-0.0550	0.0533
REITER	-0.0413	0.3062	0.0905	0.1354

Note: Mean percentage differences between the impulse response simulations for the PARAM, XPA, and REITER solutions, relative to an analogous KS impulse response simulation are reported in this table, with entries representing the average percentage value of $\hat{x}_t^{method} - \hat{x}_t^{KS}$ for each solution method *method* and impulse response \hat{x}_t^{method} , averaged over the post-shock period. The exogenous aggregate productivity impulse response is simulated as suggested by Koop et al. (1996), discussed in more detail in Appendix C, and involves the computing the average log difference between 2000 independent simulations of 50-period length, with and without positive productivity shocks.

To create an approximation to the average conditional response in our context, we therefore simulate the “generalized nonlinear impulse responses” of Koop et al. (1996), although for simplicity we refer to these simply as “impulse responses.” The details of this approach are contained in Appendix C, but a brief overview is in order here. The approach relies upon a large number of pairs of simulations, with one “shock” simulation and one “no shock simulation.” Within each pair, the two simulations are run under identical exogenous shock processes, with one difference. At a designated period after some initialization range, we impose a positive shock to aggregate productivity in the shock simulation, allowing the aggregates to evolve as normal afterwards. The average percentage difference, across simulation pairs, between the shocked and no shock simulations provides a reasonable approximation to the average innovation to a given series in response to a productivity shock.

To generate a flexibly-sized aggregate shock using discretized productivity, we simply convexify the shock arrival within each simulation pair described above, imposing a shock only with a probability calculated to generate any desired average change in aggregate productivity. The details of this additional modification, as well as Figure C1 comparing the virtually identical linearized and simulated impulse responses for the REITER method are again available in Appendix C.

Figure 3 plots the impulse response to a one-standard deviation (1.4%) average positive aggregate productivity shock for aggregate output, investment, labor, and price. The responses are qualitatively identical: an increase in aggregate productivity leads immediately to a jump in output, labor, and investment, together with an increase in consumption (reduction in price) that dissipates more gradually than the other series. Unsurprisingly, investment responds most strongly, with more moderate responses from other aggregates. Table 6 reports mean percentage differences between the conditional responses in the KS case and the PARAM, XPA, and REITER simulations, revealing little difference on average, all under a third of a percent in magnitude. Unsurprisingly, given the unconditional business cycle patterns considered above, the price (labor) response is smaller (stronger) for the REITER method than for the projection based KS, PARAM, and XPA methods.

4.3 Accuracy Statistics

As described above, firms making investment choices in the bounded rationality KS, PARAM, and XPA solution techniques must rely upon the simplified aggregate state space (A, K) to form expectations both about market-clearing prices today p , as well as the aggregate capital level in the next period K' . In the case of the KS and XPA solutions, firms make use of explicit log linear forecast rules for both quantities, conditional upon the discretized aggregate productivity state today. The PARAM method does not make available an explicit forecast rule, but PARAM does endogenously generate within the solution step a mapping over some projection grid on (A, K) to clearing levels of (p, K') , as discussed in more detail in Appendix A. By linearly interpolating this mapping we can generate a forecast system for price and aggregate capital from the PARAM solution.

There are several metrics commonly used in the heterogeneous agents literature to evaluate the internal accuracy of the expectations embedded in bounded rationality forecasting rules. For each forecast system, it is natural to compare several horizons of forecasts to gauge their internal accuracy, and the one-step ahead forecast accuracy is commonly evaluated using the R^2 of the regressions, as in Khan and Thomas (2008) or the original Krusell and Smith (1998) analysis. Similarly, the root mean squared error (RMSE) in percentage terms, can be computed from the one-step ahead capital and current-period price forecast rules. Both the R^2 and $RMSE$ statistics are reported in Appendix B in Table B2, conditional upon the realized discretized level of aggregate productivity A and based on the same common-across-methods 2000-period unconditional simulation of the aggregate productivity process used above for the comparison of business-cycle and micro-level investment moments. Appendix Figure B2 provides a comparison plot.

Table 7: Accuracy Statistics for Forecast Rules

	Max DH, K'	Mean DH, K'	Max DH p	Mean DH p
KS	0.6199	0.1909	0.2155	0.0411
PARAM	2.5986	0.4657	1.6889	0.2851
XPA	3.7564	0.7417	1.5376	0.3008

Note: The table above reports internal accuracy statistics based on unconditional simulations for the three solution methods with explicit forecast mappings from the approximate aggregate state (A, K) to realized next-period capital K' (the first two columns) and to market-clearing prices p (the final two columns). The first two columns report the maximum and mean Den Haan (2010a) statistic for aggregate capital K' , i.e. the maximum and mean values, in 100 times log differences, between a capital series based purely on iteration forward of the forecast system and the realized simulated values. The third and fourth columns reflect analogous statistics for the market clearing price p . The exogenous aggregate productivity process reflects a Markov chain discretized using the Tauchen (1986) procedure and is held constant across solution methods during the simulation. To achieve this, an identical simulated discretized productivity process is input directly into the KS, PARAM, and XPA solutions, while a series of continuous aggregate shocks exactly replicating the discretized productivity process is input into the REITER solution. For each method, the full 2000 period simulation for each solution method begins after 500 periods, with an initial burn-in period discarded to avoid the influence of initial conditions on the simulated aggregates. The exogenous productivity simulation used for calculation of the DH statistics are distinct from the simulation used within the model-solution step of any algorithm.

The R^2 and $RMSE$ measures are widespread and therefore important forecast accuracy statistics. However, Den Haan (2010a) emphasizes that forecast series which “correct” errors within the forecast system by substitution of realized aggregate capital series into the regression each period can be a weak tool for identification of the accumulation of forecast errors over time. Therefore, Table 7 reports the maximum Den Haan (DH) statistic and mean DH statistic over the simulations. These statistics rely after initialization only on forward iteration of the forecasts themselves through the forecasting system and are equal to the maximum and mean differences between realizations and the iterated forecast values, in percentage terms.

Two messages clearly emerge from the results in Table 7. First, the KS method delivers the most consistently accurate results by any metric, with maximum DH errors around half a percent or lower for both capital and price at an extended horizon, as well as even smaller average DH statistics for both series. As can be seen in Appendix Table B2, R^2 measures very near 1 also result from the KS solution, a result which was emphasized by Khan and Thomas (2008). Clearly, the absolute levels of these errors are dependent upon the exact density of the aggregate state space projection grid chosen in this paper (see details on that in Appendix B), and we would like internal accuracy to deliver errors far smaller than these KS results. However, given a fixed projection grid, as considered here, we can conclude that the *relative* accuracy of the KS method is substantially higher than that of the PARAM or XPA solutions. A second result also becomes apparent from Table 7: overall, although the accuracy of both the PARAM and XPA methods is comparable, the PARAM method appears to deliver slightly more accurate internal expectations. Maximum and mean DH statistics reflect generally superior capital forecasting performance from PARAM, with about the same level of accuracy for price forecasts. As can be seen in Appendix Table B2, the RMSE and R^2 metrics do also indicate reduced forecast accuracy for the XPA approach relative to the PARAM solution, especially for extreme values of aggregate productivity.

A final note is in order concerning the REITER solution method. Since it is based upon an approximation to the rational expectations equilibrium of the model, there is no directly comparable notion of forecast accuracy for that solution technique. However, as suggested by Reiter (2009), we can increase the density of the underlying discretization of the cross-sectional distribution substantially (by one-third in the case considered here), and compare the maximum and mean simulated difference between market-clearing aggregate price series in the baseline REITER discretization and the higher density approximation. Those statistics, only very roughly analogous to the price DH statistics reported in Table 7, are 0.1830(0.0466%) for the max (mean) differences. The baseline REITER price simulation, together with the same series generated using the denser grid, are plotted in Appendix Figure B1.

4.4 Computational Time

A final explicit comparison of the solution techniques involves computational time. Although run-time comparisons inevitably depend upon the efficiency and choices made when coding the solution methods, as well as the specific language or software used and the details of the numerical approach, a few considerations help to allay those concerns in this case. The projection-based solution KS, PARAM, and XPA solutions techniques used here, as well as the code solving the model with no aggregate uncertainty, were all written in Fortran by the same researcher, liberally parallelizing when possible and executing the programs on the same hardware. Then, using the model with no aggregate uncertainty from Fortran as an input, the REITER method was implemented by numerically differentiating the equilibrium system described in Appendix A around the steady-state and solving the resulting linear system in MATLAB using the standard method of Sims (2002).¹⁴ We should therefore feel comfortable in at least relative comparisons of computational time across methods, which are reported in Table 8.

Table 8: Computational Time

Time (in seconds)	Solution	Unconditional Simulation	IRF Simulation
	KS	518	44950
	PARAM	1090	77994
	XPA	471	42236
	REITER	0.5	33.9

Note: The quantities above refer to the runtime on a six-core 3.33GHz Dell Precision T3500 with 24 GB of RAM. of model solution and simulation code in parallelized Fortran (KS, PARAM, and XPA methods) as well as Fortran and MATLAB (REITER method). All of the code used to produced these results can be found on Stephen Terry’s website. “Solution” refers to the time required for calculation of a forecast rule fixed point (KS and XPA), completion of value function iteration (PARAM), or calculation of the linearized model solution representation (REITER). Unconditional simulation refers to the time required for unconditional simulation of a 2000-period economy with identical exogenous aggregate shocks across solution methods and an initial simulation range before 500 periods discarded. IRF simulation refers to the time required for simulation of a conditional impulse response function to an aggregate productivity shock using the method of Koop et al. (1996), with a simulation length of 50 years, 2000 replications, and shocks held

¹⁴In particular, the linear solution step uses the `gensys` software from Chris Sims’ website.

constant across solution methods. All models are solved using comparable idiosyncratic and aggregate grids, identical Bellman equation or policy iteration tolerances, and identical forecast rule initial conditions, with the exception of the REITER method, which is solved using a denser cross-sectional grid. See Appendix B for details.

Within the model solution step, the KS algorithm takes approximately 20 times as long as the PARAM or XPA techniques, due to the necessity of repeated model simulation to find a forecasting system fixed point. By avoiding simulation, each of those two alternative approaches reduces time within the model solution step substantially. Although the steady-state solution with no aggregate uncertainty, an initial input into the REITER method, can be solved within a couple of seconds, the numerical differentiation and solution of the resulting linear system take a bit more time. Overall, for the numerical choices made here, the REITER solution takes approximately two and a half times as long as the XPA method.

Simulation speeds fall into two distinct groups. The bounded rationality projection-based KS, PARAM, and XPA approaches are costly to simulate but take around the same amount of time. Each approach requires iteration on either the market-clearing price (KS, XPA), or the price, next period's capital stock, and the approximating coefficients of a simulated cross-sectional density (PARAM). Although the PARAM technique takes the most time within this group for simulation, costs are roughly comparable. By contrast, once a linear representation of the equilibrium is obtained in the REITER solution step, simulation is virtually costless, about 1000 times faster than the next-quickest XPA approach for the unconditional simulation. A similar ratio is evident for the much longer repeated simulations of impulse response analysis: the REITER method dominates in simulation time by about three orders of magnitude.¹⁵ The increased simulation speed, as well as the availability of linearized impulse responses, opens the door for expansion of the aggregate state space within the REITER approach, considered next.

5 Extending the Aggregate State Space: A Simple Example

Perhaps the most significant conceptual contribution of the REITER approach is the scalability it offers. By relying on perturbations of the model in aggregates, it is possible to simultaneously simulate fully nonlinear micro-level behavior together with a rich aggregate structure. By avoiding the curse of dimensionality inherent in the KS, PARAM, and XPA methods, it is also possible at low cost to add a rich set of aggregate shocks and states which can offer a connection between the traditional representative agent approach within macroeconomics and the micro-founded heterogeneous agents literature. A recent and important example of such a connection is McKay and Reis (2013), which integrates the incomplete markets model, a New Keynesian structure, and nontrivial fiscal policy dynamics to study the implications of fiscal transfers for output dynamics. Similarly, Reiter et al. (2009) have applied their method to a structure with sticky capital and sticky investment choices for firms within a New Keynesian environment, and Costain and Nakov (2011) have

¹⁵As emphasized by McKay (2013), the easy and fast manipulability of the linearized equilibrium from the REITER solution opens the door to the application of full-information estimation techniques for heterogeneous agent models, an intriguing possibility.

studied the dynamics of a rich state-dependent pricing model using the REITER method.¹⁶

As a concrete example within the heterogeneous firms model, with details deferred to Appendix D, I illustrate how at essentially no additional computational cost or complexity it is possible to add two demand or preference shocks, one to the rate of time-preference and one to labor disutility, within the simplified Khan and Thomas (2008) framework. It is clear that adding two additional aggregate states is computationally quite burdensome within the bounded rationality projection-based KS, PARAM, or XPA approaches due to the curse of dimensionality.¹⁷ Adding the full richness of some large representative agent models in terms of shocks and equilibrium interactions would become even more infeasible. In Figure 5 I plot the linearized impulse response of this extended model to a positive time-preference or demand shock, somewhat arbitrarily scaled to equal the magnitude of the aggregate productivity shock at 1.4%. Unsurprisingly, the demand shock delivers increased consumption but reductions in investment, output, and labor. While such dynamics leading to a lack of comovement are not necessarily empirically plausible, the REITER method is obviously capable of generating nontrivially richer aggregate dynamics with essentially the same computational burden. Impulse responses to the other two aggregate shocks are deferred to Appendix D.

6 Conclusion

By comparing the KS, PARAM, XPA, and REITER solution methods along the dimension of business cycle and micro-level investment moments, conditional impulse responses, internal accuracy, and runtime, we are able to draw a more complete picture of the tradeoffs among solution techniques available for the heterogeneous firms model. Overall, the KS algorithm is time consuming but internally extremely accurate and robust. The related XPA and PARAM algorithms deliver less internal accuracy, quantitatively similar economic implications to the KS approach, but large within-solution step speed gains by avoiding the KS algorithm’s dependence upon simulation within the model solution step. Quantitatively and conceptually, these three algorithms based on bounded rationality, approximation of the aggregate state space, and global projection techniques generally deliver similar conclusions, and this paper can be interpreted as a favorable robustness check for the long literature using the KS method in heterogeneous firms contexts.

By contrast, with a conceptually different rational expectations equilibrium concept, yet still qualitatively similar economic conclusions, even more dramatic time savings, and an important scalability, the REITER method offers an alternative to the other three methods considered that can potentially serve as a useful link between the representative agent and heterogeneous agent

¹⁶Note that when the discretization of the cross-sectional distribution is too dense, or the aggregate dynamics too rich, to handle with standard linear solution techniques, Reiter (2010a) provides an overview of model reduction techniques which can be used to solve a smaller system of linear equations with dynamics similar to those of the original, larger system. Such model reduction is implemented by McKay and Reis (2013).

¹⁷Of course, although quite costly, such expanded analysis may still be feasible. See Bloom et al. (2012), Khan and Thomas (2013), and Bachmann and Ma (2012), among others, for KS method approaches with richer aggregate state spaces.

business cycle literatures by allowing for an extremely rich aggregate state space within the context of a fully specified nonlinear microeconomic set of distributions and policies.

An interesting set of recently proposed solution techniques, omitted from this paper’s comparison but potentially extremely useful in future as a means of efficiently solving a global projection-based approximation to a full rational expectations heterogeneous agents equilibrium are presented in Gordon (2011) and Judd et al. (2012). In Gordon (2011), the use of sparse Smolyak projection grids allows for the solution of models with the full cross-sectional distribution within the state space but still potentially subject to extreme calibrations or shocks. In Judd et al. (2012) a simulation-based reduction of the size of a projection grid is proposed that naturally might allow for the incorporation of a full discretized distribution within the aggregate state space of a heterogeneous agents model. Possible application of the Gordon (2011) method to the heterogeneous firms context and the general application of the Judd et al. (2012) projection grid simplification technique to heterogeneous agents frameworks like the incomplete markets model are the subject of ongoing work.

References

- Algan, Yann, Olivier Allais, and Wouter J Den Haan (2008), “Solving heterogeneous-agent models with parameterized cross-sectional distributions.” *Journal of Economic Dynamics and Control*, 32, 875–908.
- Algan, Yann, Olivier Allais, and Wouter J Den Haan (2010a), “Solving the incomplete markets model with aggregate uncertainty using parameterized cross-sectional distributions.” *Journal of Economic Dynamics and Control*, 34, 59–68.
- Algan, Yann, Olivier Allais, Wouter J Den Haan, and Pontus Rendahl (2010b), “Solving and simulating models with heterogeneous agents and aggregate uncertainty.” *Handbook of Computational Economics*.
- Bachmann, Rüdiger and Christian Bayer (2013), “Wait-and-see business cycles?” *Journal of Monetary Economics*, 60, 704–719.
- Bachmann, Rüdiger, Ricardo J Caballero, and Eduardo MRA Engel (2013), “Aggregate implications of lumpy investment: new evidence and a dsge model.” *American Economic Journal: Macroeconomics*, 5, 29–67.
- Bachmann, Rüdiger and Lin Ma (2012), “Lumpy investment, lumpy inventories.” Working paper.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J Terry (2012), “Really uncertain business cycles.” Working paper.
- Carroll, Christopher D (2006), “The method of endogenous gridpoints for solving dynamic stochastic optimization problems.” *Economics letters*, 91, 312–320.
- Christiano, Lawrence J (2002), “Solving dynamic equilibrium models by a method of undetermined coefficients.” *Computational Economics*, 20, 21–55.

- Costain, James and Anton Nakov (2011), “Distributional dynamics under smoothly state-dependent pricing.” *Journal of Monetary Economics*, 58, 646–665.
- Den Haan, Wouter J (2010a), “Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents.” *Journal of Economic Dynamics and Control*, 34, 79–99.
- Den Haan, Wouter J (2010b), “Comparison of solutions to the incomplete markets model with aggregate uncertainty.” *Journal of Economic Dynamics and Control*, 34, 4–27.
- Den Haan, Wouter J and Pontus Rendahl (2010), “Solving the incomplete markets model with aggregate uncertainty using explicit aggregation.” *Journal of Economic Dynamics and Control*, 34, 69–78.
- Gordon, Grey (2011), “Computing dynamic heterogeneous-agent economies: Tracking the distribution.” Working paper.
- Gourio, François and Anil K Kashyap (2007), “Investment spikes: New facts and a general equilibrium exploration.” *Journal of Monetary Economics*, 54, 1–22.
- Judd, Kenneth L, Lilia Maliar, and Serguei Maliar (2012), “Merging simulation and projection approaches to solve high-dimensional problems.” Working paper.
- Kehrig, Matthias and Nicolas Vincent (2013), “Financial frictions and investment dynamics in multi-plant firms.” *US Census Bureau Center for Economic Studies Paper No. CES-WP-56*.
- Khan, Aubhik and Julia K Thomas (2003), “Nonconvex factor adjustments in equilibrium business cycle models: do nonlinearities matter?” *Journal of monetary economics*, 50, 331–360.
- Khan, Aubhik and Julia K Thomas (2008), “Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics.” *Econometrica*, 76, 395–436.
- Khan, Aubhik and Julia K Thomas (2013), “Credit shocks and aggregate fluctuations in an economy with production heterogeneity.” *Journal of Political Economy*, 121, 1055–1107.
- Klenow, Peter J and Jonathan L Willis (2007), “Sticky information and sticky prices.” *Journal of Monetary Economics*, 54, 79–99.
- Knotek, Edward (2010), “A tale of two rigidities: Sticky prices in a sticky-information environment.” *Journal of Money, Credit and Banking*, 42, 1543–1564.
- Knotek, Edward and Stephen J Terry (2008), “Alternative methods of solving state-dependent pricing models.” Working paper.
- Koop, Gary, M Hashem Pesaran, and Simon M Potter (1996), “Impulse response analysis in non-linear multivariate models.” *Journal of Econometrics*, 74, 119–147.
- Krusell, Per and Anthony A Smith, Jr (1998), “Income and wealth heterogeneity in the macroeconomy.” *Journal of Political Economy*, 106, 867–896.
- Maliar, Lilia, Serguei Maliar, and Fernando Valli (2010), “Solving the incomplete markets model with aggregate uncertainty using the krusell–smith algorithm.” *Journal of Economic Dynamics and Control*, 34, 42–49.

- McKay, Alisdair (2013), “Idiosyncratic risk, insurance, and aggregate consumption dynamics: a likelihood perspective.” Working paper.
- McKay, Alisdair and Ricardo Reis (2013), “The role of automatic stabilizers in the us business cycle.” Working paper.
- Reiter, Michael (2009), “Solving heterogeneous-agent models by projection and perturbation.” *Journal of Economic Dynamics and Control*, 33, 649–665.
- Reiter, Michael (2010a), “Approximate and almost-exact aggregation in dynamic stochastic heterogeneous-agent models.” Working paper.
- Reiter, Michael (2010b), “Nonlinear solution of heterogeneous agent models by approximate aggregation.” Working paper.
- Reiter, Michael (2010c), “Solving the incomplete markets model with aggregate uncertainty by backward induction.” *Journal of Economic Dynamics and Control*, 34, 28–35.
- Reiter, Michael, Tommy Sveen, and Lutz Weinke (2009), “Lumpy investment and state-dependent pricing in general equilibrium.” Working paper.
- Sims, Christopher A (2002), “Solving linear rational expectations models.” *Computational economics*, 20, 1–20.
- Sunakawa, Takeki (2012), “Applying the explicit aggregation algorithm to discrete choice economies: With an application to estimating the aggregate technology shock process.” Working paper.
- Tauchen, George (1986), “Finite state markov-chain approximations to univariate and vector autoregressions.” *Economics letters*, 20, 177–181.
- Thomas, Julia K (2002), “Is lumpy investment relevant for the business cycle?” *Journal of Political Economy*, 110, 508–534.
- Vavra, Joseph (2014), “Inflation dynamics and time-varying volatility: New evidence and an ss interpretation.” *The Quarterly Journal of Economics*, 129, 215–258.
- Young, Eric 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, 36–41.

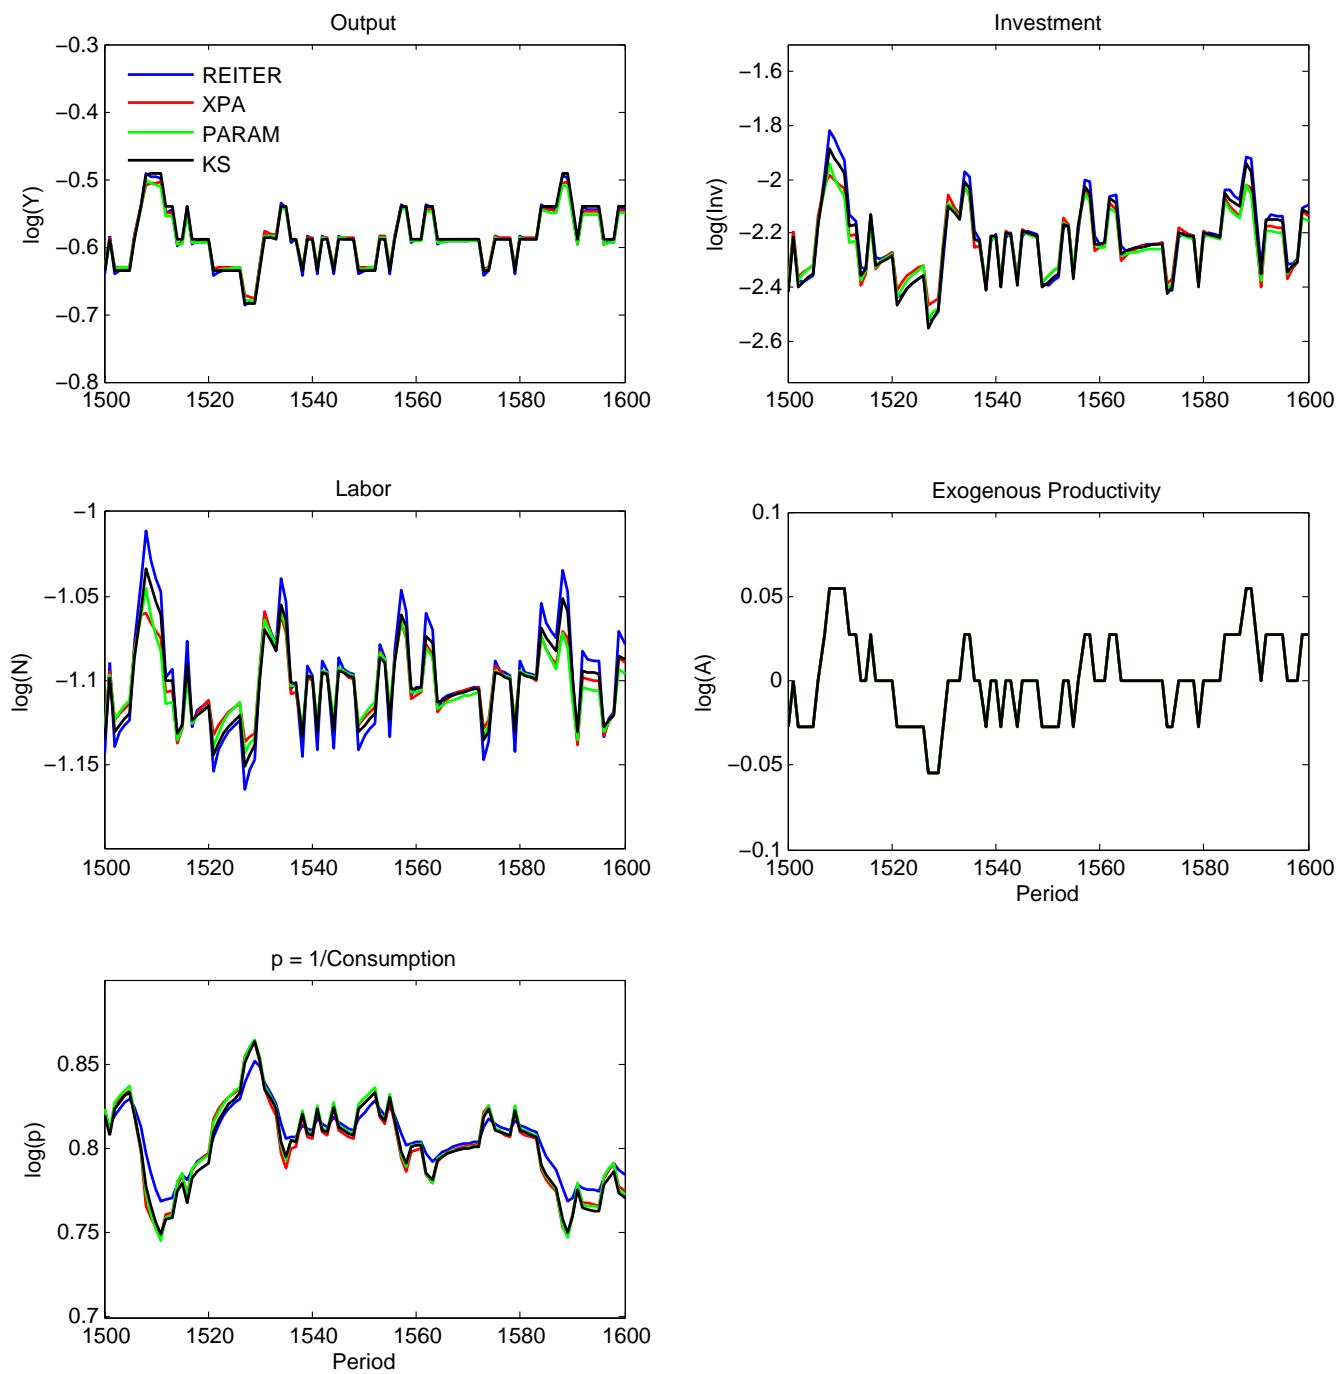


Figure 1: Unconditional Simulation Comparison

Note: The figure above plots a representative 100-period portion of a 2000-period unconditional simulation of the KS, PARAM, XPA, and REITER solutions. The KS solution is in black, the XPA in red, the PARAM in green, and the REITER in blue. The exogenous aggregate productivity process reflects a Markov chain discretized using the Tauchen (1986) procedure and is held constant across solution methods during the simulation. To achieve this, an identical simulated discretized productivity process is input directly into the KS, PARAM, and XPA solutions, while a series of continuous aggregate shocks exactly replicating the discretized productivity process is input into the REITER solution. The full 2000 period simulation for each solution method begins after 500 periods, with an initial burn-in period discarded to avoid the influence of initial conditions on the simulated aggregates.

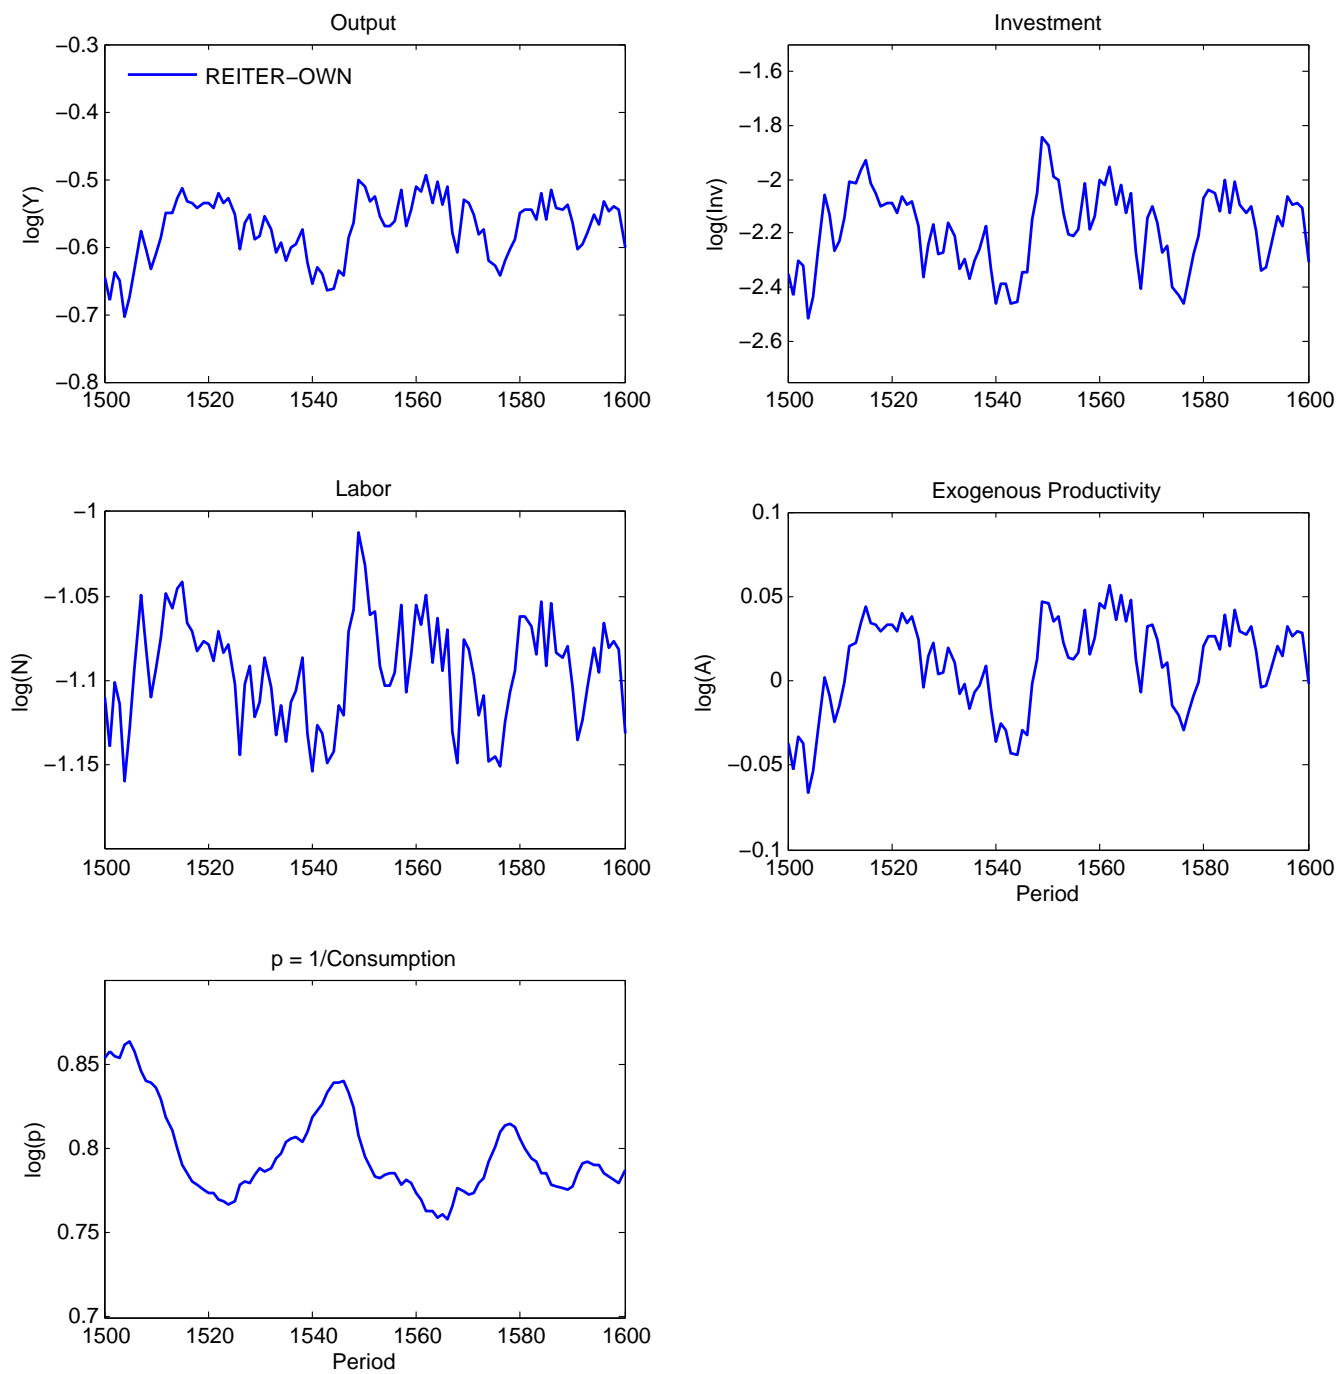


Figure 2: Unconditional Continuous Simulation of REITER-OWN

Note: The figure above plots a representative 100-period portion of a 2000-period unconditional simulation of the REITER solution, labelled REITER-OWN in business cycle and investment-rate moment tables in the main text. The shock process to aggregate productivity is drawn from a continuous distribution $N(0, \sigma_A^2)$, and because of the continuity of the aggregate productivity process the simulation differs from the discretized simulations possible for the KS, PARAM, and XPA solutions. The full 2000 period simulation for each solution method begins after 500 periods, with an initial burn-in period discarded to avoid the influence of initial conditions on the simulated aggregates.

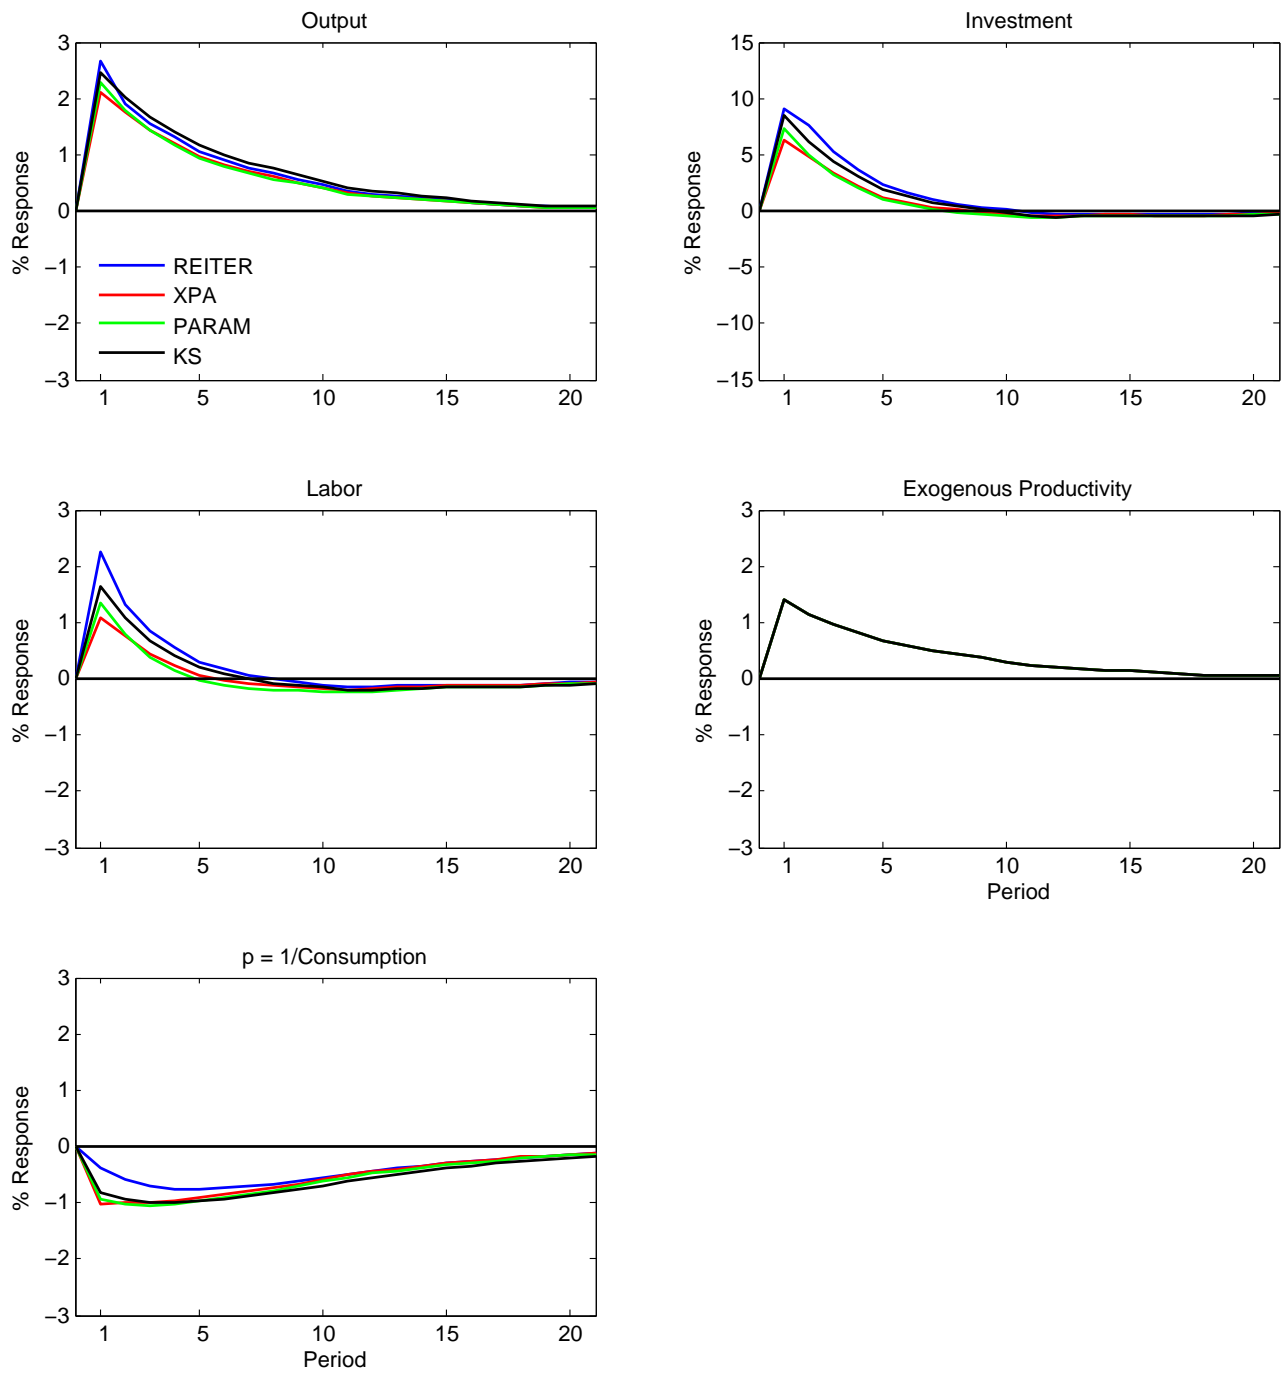


Figure 3: Impulse Response, Productivity Shock

Note: Simulated impulse responses to an aggregate productivity shocks for the KS, PARAM, XPA, and REITER series are plotted above, in percentages. The KS solution is in black, the XPA in red, the PARAM in green, and the REITER in blue. The exogenous aggregate productivity impulse response is simulated as suggested by Koop et al. (1996) and discussed in more detail in Appendix C. The simulation consists of 2000 independent simulations of 50-period length, with and without imposed productivity shocks which occur during a period labelled 1 above. In one simulation, a positive shock to aggregate productivity with mean equal to one-standard deviation of the aggregate productivity shock process occurs, while in another simulation with otherwise identical exogenous series, no aggregate productivity innovation occurs. The reported series \hat{x}_t^{method} is the mean of $100(\log(X_t^{shock}) - \log(X_t^{no shock}))$ across simulations, for each series X_t and solution method $method$.




Figure 4: Den Haan Fundamental Accuracy Plot

Note: The figures above plot a representative 100-period portion of a larger 2000-period unconditional simulation of the three solution methods with explicit forecast mappings from the approximate aggregate state space (A, K) to realized prices and next-period productivity (p, K') , i.e. the KS (in black), PARAM (in green), and XPA (in red) methods. The first column plots the realized market-clearing price p (solid line) and the forecast value of price (dotted line), computed by iterating forward the forecasting system as suggested by Den Haan (2010a) rather than substituting realizations of aggregate capital at each point. Similarly, the second column plots the realized capital series against the iterated forecasts for aggregate capital made on the basis of the forecast system alone. During the simulation, the exogenous aggregate productivity process reflects a Markov chain discretized using the Tauchen (1986) procedure and is held constant across solution methods during the simulation. To achieve this, an identical simulated discretized productivity process is input directly into the KS, PARAM, and XPA solutions, while a series of continuous aggregate shocks exactly replicating the discretized productivity process is input into the REITER solution. For each method, the full 2000 period simulation for each solution method begins after 500 periods, with an initial burn-in period discarded to avoid the influence of initial conditions on the simulated aggregates. The exogenous productivity simulation used for calculation of the DH forecast series is distinct from the simulation used within the model-solution step of any algorithm.

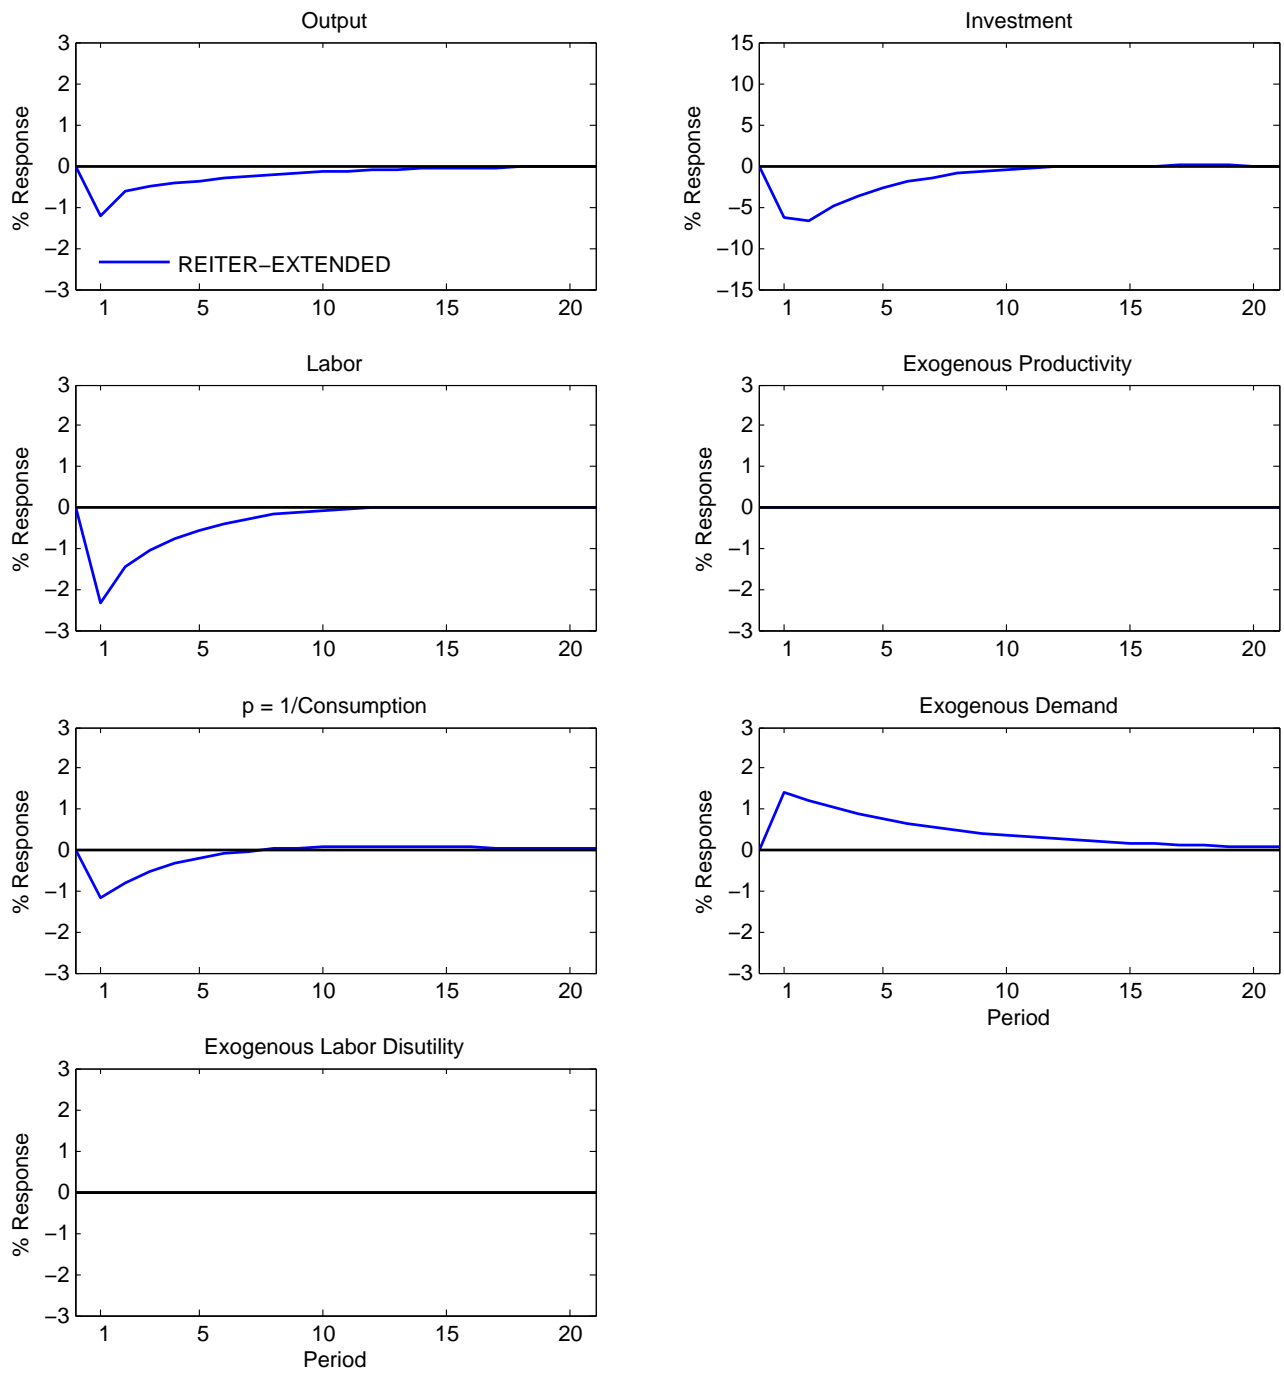


Figure 5: Impulse Response, Demand Shock in Extended Model

Note: The figure above plots the linearized impulse response function to a positive shock to the aggregate demand series D_t in the extension of the heterogeneous firms model augmented with aggregate demand and labor preference shocks. These impulse responses are computed after solving the extended model using the REITER technique and are therefore labelled REITER-EXTENDED. Given the endogenous variables X_t of the linearized system describing the economy, a solution $X_t - X^{SS} = A(X_{t-1} - X^{SS}) + B\epsilon_t$ of the model trivially yields impulse responses $A^{t-1}B$ representing the economy's local response to each aggregate shock in the vector ϵ_t . The impulse response here is scaled by the assumed standard deviation of innovations to the aggregate demand series, with the exogenous shocked series itself plotted in the right hand side of the third row.

A Solution Methods

This section describes the equilibrium of the model with no aggregate uncertainty, as well as the details of the KS, PARAM, XPA, and REITER solution methods. To maintain generality, the bulk of the section focuses on describing the algorithms or equilibrium concepts themselves. Therefore, some practical numerical issues involved in the solution techniques are mentioned in passing, but most numerical details (on grid sizes, optimization algorithms, etc) are deferred to a listing in Appendix B.

A.1 No Aggregate Uncertainty Model

The equilibrium definition of the steady-state model or model with no aggregate uncertainty is identical to the equilibrium with aggregate uncertainty discussed in the main text with constant aggregate productivity A . Unless otherwise specified, the steady-state model will be solved with $A = 1$. With a constant aggregate state space, individual firm states are given by the far smaller state space (z, k) , and solution of the no aggregate uncertainty model simply involves repeatedly guessing values of the market-clearing price or consumption, computing an ergodic cross-sectional distribution $\mu(z, k)$ based on the price-implied policies and adjustment thresholds, and then checking consistency with the guessed price level. In the code available for this paper, the price clearing is performed using bisection, and calculation of an ergodic cross-sectional distribution given a price level follows the nonstochastic or histogram-based approach of Young (2010). In turn, this approach requires a projection grid for value functions, a denser simulation grid for idiosyncratic capital, and a discretized productivity process at the idiosyncratic level. For completeness, the equilibrium equations are listed below:

$$\begin{aligned}
 V^A(z, k) &= \max_{k', n} \left\{ p \left(Azk^\alpha n^\nu - k' + (1 - \delta)k - \frac{\phi}{p}n \right) + \beta \mathbb{E}_{z'} V(z', k') \right\} \\
 V^{NA}(z, k) &= \max_n \left\{ p \left(Azk^\alpha n^\nu - \frac{\phi}{p}n \right) + \beta \mathbb{E}_{z'} V(z', (1 - \delta)k) \right\} \\
 V(z, k) &= -\phi \int_0^{\xi^*(z, k)} \xi dG(\xi) + G(\xi^*(z, k))V^A(z, k) + [1 - G(\xi^*(z, k))]V^{NA}(z, k) \\
 \xi^*(z, k) &= \frac{V^A(z, k) - V^{NA}(z, k)}{\phi} \\
 \frac{1}{p} &= \int \int Azk^\alpha n(z, k, \xi)^\nu d\mu(z, k) dG(\xi) - \int \int (k'(z, k, \xi) - (1 - \delta)k) d\mu(z, k) dG(\xi),
 \end{aligned}$$

where we have that policies for labor have a closed-form and capital policies follow a threshold rule, given optimal policies $k^*(z, k)$ from $V^A(z, k)$:

$$k'(z, k, \xi) = \begin{cases} k^*(z, k), & \xi \leq \xi^*(z, k) \\ (1 - \delta)k, & \xi > \xi^*(z, k) \end{cases},$$

and that $\mu(z, k)$ is the ergodic distribution implied by a discretization as in Young (2010) together with application of the capital policies, adjustment cost thresholds, and idiosyncratic productivity transitions. Once the equilibrium is obtained, it is trivial to compute aggregate capital as $K = \int k d\mu(z, k)$.

A.2 Krusell Smith (KS)

The KS solution method used in the paper, following Khan and Thomas (2008) assumes that the set of approximating moments m in the aggregate state space of the model is equal to the mean or aggregate capital level K . Further, the solution method discretizes idiosyncratic and aggregate productivity processes following Tauchen (1986). Conditional upon a discretized level of aggregate productivity A , the forecast rules for price and next period's aggregate capital level take a loglinear form:

$$\log(\hat{p}) = \alpha_p(A) + \beta_p(A) \log(K)$$

$$\log(\hat{K}') = \alpha_K(A) + \beta_K(A) \log(K).$$

Recall that from the household problem this yields a forecast wage level $\hat{w} = \frac{\phi}{\bar{p}}$. Given these choices, the solution algorithm works as follows. First, guess an initial forecast rule system $\hat{\Gamma}_{(1)} = (\alpha_p^{(1)}, \beta_p^{(1)}, \alpha_K^{(1)}, \beta_K^{(1)})$. Then, before the solution iteration takes place, draw a large number T of exogenous aggregate productivity values based on the discretized Markov chain for A . Finally, in iteration $s = 1, 2, \dots$ do the following

1. Given forecast rule system $\hat{\Gamma}_{(s)}$, solve the following system of equations via projection on some grid of states $(z, k; A, K)$

$$V_{(s)}^A(z, k; A, K) = \max_{k', n} \left\{ \hat{p}_{(s)}(A, K) \left(z A k^\alpha n^\nu - k' + (1 - \delta)k - \hat{w}_{(s)}(A, K)n \right) + \beta \mathbb{E}_{z', A'} V_{(s)}(z', k'; A', \hat{K}'_{(s)}) \right\}.$$

$$V_{(s)}^{NA}(z, k; A, K) = \max_n \left\{ \hat{p}_{(s)}(A, K) \left(z A k^\alpha n^\nu - \hat{w}_{(s)}(A, K)n \right) + \beta \mathbb{E}_{z', A'} V_{(s)}(z', (1 - \delta)k; A', \hat{K}'_{(s)}) \right\}$$

$$\xi_{(s)}^*(z, k; A, K) = \frac{V_{(s)}^A(z, k; A, K) - V_{(s)}^{NA}(z, k; A, K)}{\phi},$$

$$V_{(s)}(z, k; A, K) = -\phi \int_0^{\xi_{(s)}^*(z, k; A, K)} \xi dG(\xi) + G\left(\xi_{(s)}^*(z, k; A, K)\right) V_{(s)}^A(z, k; A, K) + \left(1 - G\left(\xi_{(s)}^*(z, k; A, K)\right)\right) V_{(s)}^{NA}(z, k; A, K)$$

which determines value function $V_{(s)}(z, k; A, K)$. Note that expectations are computed via summation over the appropriate portion of the discretized transition matrices Π_z and Π_A for idiosyncratic and aggregate productivity, respectively, and that the solution to the Bellman equations can be achieved with policy iteration, Howard acceleration of the Bellman equations, and continuous univariate optimization techniques in next period's capital level k' , such as Brent optimization or golden section search.

2. Given the value function solution from step 1, simulate the model. Simulation follows the nonstochastic approach of Young (2010), and requires initialization of the cross-sectional distribution of capital and productivity $\mu(z, k)$ over a dense, discrete grid of (z_i, k_j) with n_d points for discretized capital. In each period t , market-clearing prices p_t must be determined without reference to the forecast price level $\hat{p}_{(s)}$ (which appear only through embedded expectations in the continuation value). The price clearing algorithm, for a given guessed price value \tilde{p} requires calculation of the excess demand or error function, which is done as follows:

- Compute new capital and adjustment cutoff thresholds at each point on the dense simulation grid (z_i, k_j) through

$$\tilde{V}^A(z_i, k_j) = \max_{k', n} \left\{ \tilde{p} \left(z_i A k_j^\alpha n^\nu - k' + (1 - \delta)k_j - \frac{\phi}{\tilde{p}}n \right) + \beta \mathbb{E}_{z', A'} V_{(s)}(z', k'; A', \hat{K}'_{(s)}) \right\}$$

$$\tilde{V}^{NA}(z_i, k_j) = \max_n \left\{ \tilde{p} \left(z_i A k_j^\alpha n^\nu - \frac{\phi}{\tilde{p}}n \right) + \beta \mathbb{E}_{z', A'} V_{(s)}(z', (1 - \delta)k_j; A', \hat{K}'_{(s)}) \right\}$$

$$\tilde{\xi}^*(z_i, k_j) = \frac{\tilde{V}^A(z_i, k_j) - \tilde{V}^{NA}(z_i, k_j)}{\phi}$$

where the optimization over labor is a static problem which yields an analytical reduced-form for the right hand side expressions which can be optimized in capital k' alone. Let optimal labor and capital policies conditional upon adjustment be given by $n^*(z_i, k_j)$ and $k'^*(z_i, k_j)$. We have suppressed dependence on the aggregate states (A, K) for the moment, which is constant through the market-clearing process.

- Compute implied output, investment, consumption, and labor from

$$\begin{aligned}\tilde{Y} &= \sum_{z_i, k_j} \mu_t(z_i, k_j) y(z_i, k_j), \quad y(z_i, k_j) = z_i A k_j^\alpha n^*(z_i, k_j)^\nu \\ \tilde{I} &= \sum_{z_i, k_j} \mu_t(z_i, k_j) \left[G(\tilde{\xi}^*(z_i, k_j)) (k'^*(z_i, k_j) - (1 - \delta)k_j) \right] \\ \tilde{C} &= \tilde{Y} - \tilde{I} \\ \tilde{N} &= \sum_{z_i, k_j} \mu_t(z_i, k_j) \left(n^*(z_i, k_j) + \int_0^{\xi^*(z_i, k_j)} \xi dG(\xi) \right)\end{aligned}$$

- Define the excess demand function or clearing error as $\frac{1}{\tilde{p}} - \tilde{C}$
- If excess demand is suitably small, the clearing algorithm stops. If not, repeat with an updated value of \tilde{p} , using any preferred method such as bisection.
- With market clearing price p_t in hand (as well as Y_t, I_t, C_t, N_t), update the discretized distribution $\mu_{t+1}(z_i, k_j)$ for the next period. Based on Young (2010), distributional transitions are calculated through linear interpolation of capital policies and use of the idiosyncratic productivity transition matrix

$$\begin{aligned}\mu_{t+1}(z_{i'}, k_{j'}) &= \sum_{z_i} \sum_{k_j} \Pi_{z_{i'i}} \mu_t(z_i, k_j) \left[\begin{array}{c} \omega^a(i, j, j') \mathbb{I}(k'^*(z_i, k_j) \in [k_{j'-1}, k_{j'+1}]) \\ + \omega^{na}(i, j, j') \mathbb{I}(k_j \in [k_{j'-1}, k_{j'+1}]) \end{array} \right] \\ \omega^a(i, j, j') &= \begin{cases} G(\xi^*(z_i, k_j)) \left(\frac{k'^*(z_i, k_j) - k_{j'}}{k_{j'+1} - k_{j'}} \right) & k'^*(z_i, k_j) \in [k_{j'}, k_{j'+1}], 1 < j' < n_d \\ G(\xi^*(z_i, k_j)) & k'^*(z_i, k_j) \geq k_{n_d}, j' = n_d \\ G(\xi^*(z_i, k_j)) & k'^*(z_i, k_j) \leq k_1, j' = 1 \end{cases} \\ \omega^{na}(i, j, j') &= \begin{cases} [1 - G(\xi^*(z_i, k_j))] \left(\frac{\tilde{k}_j - k_{j'-1}}{k_{j'} - k_{j'-1}} \right) & \tilde{k}_j \in [k_{j'-1}, k_{j'}], 1 < j' < n_d \\ [1 - G(\xi^*(z_i, k_j))] & \tilde{k}_j \geq k_{n_d}, j' = n_d \\ [1 - G(\xi^*(z_i, k_j))] & \tilde{k}_j \leq k_1, j' = 1 \end{cases}\end{aligned}$$

where above $\tilde{k}_j = \max((1 - \delta)k_j, k_1)$ is the non-adjustment next-period capital stock. Also, with μ_{t+1} in hand, next period's simulated capital stock is computable as

$$K_{t+1} = \sum_{z_i, k_j} \mu_{t+1}(z_i, k_j) k_j.$$

- Continue to simulate period $t + 1$.
3. With the simulation from periods 1 to T completed given forecast rules $\hat{\Gamma}_{(s)}$, discard some number T_{erg} of periods as initialization, and update the forecast rules based on OLS regressions. Run the loglinear OLS regressions of *realized* prices p_t on aggregate capital stocks K_t and realized next-period capital stocks \hat{K}_{t+1} on current stocks K_t , where new coefficients $(\hat{\alpha}_p, \hat{\beta}_p, \hat{\alpha}_K, \hat{\beta}_K)$ are obtained separately for each level of aggregate productivity, and compute forecast rule errors as the maximum absolute difference between assumed coefficients $\hat{\Gamma}_{(s)}$ and the new estimated values. If the coefficients have converged to an acceptable tolerance, the model is solved. If not, then update $\hat{\Gamma}_{(s+1)}$ using dampened fixed-point iteration (i.e. set $\hat{\Gamma}_{(s+1)}$ equal to a weighted average of $\hat{\Gamma}_{(s)}$ and the newly estimated coefficients), then continue to solution iteration $s + 1$.

A.3 Parametrized Distributions (PARAM)

This is a discussion of the computational algorithm due to Algan et al. (2008, 2010a) (henceforth AADH) which relies upon higher-order reference moments, as well as assumed functional forms for the cross-sectional distribution of idiosyncratic capital and productivity $\mu(z, k)$ in the solution of the model.

Just as in the KS algorithm, we first must discretize the aggregate and idiosyncratic productivity processes following Tauchen (1986). Let the number of idiosyncratic (aggregate) productivity points be given by $n_z(n_A)$. Then, prior to the solution of the model, we first determine a set of aggregate moments m to be included in the aggregate state space. Here, this will be the singleton of aggregate capital, the cross-sectional mean K . Together with aggregate productivity, (A, K) therefore forms the aggregate state. Then, determine a set of reference moments m^{ref} used to help pin down the shape of the cross-sectional distribution of idiosyncratic capital and productivity. Here, this will be the first $n_M n_z$ centered moments of the capital distribution conditional upon each value of idiosyncratic productivity. Also, compute the exogenous ergodic distribution of idiosyncratic productivity $\tilde{\pi}_z$ for future use.

The reference moments are needed in this algorithm because, together with the aggregate capital state, they jointly determine the coefficients of the flexible exponential function form for the approximation to $\mu(z_i, k)$ at discretized levels of z_i :

$$P(k, \rho^{z_i}) = \rho_0^{z_i} \exp \left\{ \begin{array}{l} \rho_1^{z_i} (k - m_1^{z_i}) + \rho_2^{z_i} [(k - m_1^{z_i})^2 - m_2^{z_i}] \\ + \dots + \rho_{n_M}^{z_i} [(k - m_1^{z_i})^{n_M} - m_{n_M}^{z_i}] \end{array} \right\}.$$

As discussed in AADH, a well-behaved convex minimization problem which is laid out below yields a mapping between the $n_z n_M$ moments of the cross-sectional distribution and coefficient vectors ρ^{z_i} , $i = 1, \dots, n_z$, because the first-order conditions of this problem are identical to the moment conditions setting distributional and reference moments equal to each other. We assume when manipulating or integrating this distribution that idiosyncratic capital lies in the bounds $[\underline{k}, \bar{k}]$.

To solve the model, start with an initial constant guess for these reference moments, from the distribution in the steady-state model with no aggregate uncertainty. If desired, simulation can later provide a set of reference moments which vary with the level of aggregate productivity, or these steady-state reference moments themselves can be held fixed. Although we will outline both the repeated simulation and simulation-free approaches below, the results reported in the paper hold the reference moments at their steady-state levels. Finally, another option is also to use a regression forecast system from $(A, K) \rightarrow m^{ref}$, although following AADH we forego that approach here.

1. Iterate over the reference moments or the reference moment forecasting system, which yields a constant or variable mapping $(A, K) \rightarrow m^{ref}$.
 - (a) Loop over the aggregate states (A, K) , on the discretized grid for A and some projection grid for K .
 - For each (A, K) and implied $m^{ref}(A, K)$, choose distributional coefficients $\rho_{A,K,1}^{z_i}, \dots, \rho_{A,K,n_M}^{z_i}$, $i = 1, \dots, n_z$ to solve

$$\min_{\rho_{A,K,i=1,\dots,n_z}^{z_i}} \sum_{i=1}^{n_z} \int_{\underline{k}}^{\bar{k}} P(k, \rho_{A,K}^{z_i}) dk,$$

noting that the constants $\rho_{a,K,0}^{z_i}$ are irrelevant for the minimization and are simply chosen to ensure integration to 1. This is a well-behaved convex minimization problem and the integral can be computed via any standard quadrature rule. Here and throughout, integrals are computed via Simpson's rule. It is important to note that any rule with fixed nodes and weights is preferable to adaptive methods because fixed rules contribute to stability in the minimization problem above. The minimization is implemented in this paper using a robust and quick quasi-Newton routine with symmetric rank-one (SR1) updating of the approximation to the inverse Hessian.

- (b) Iterate over $V_{(s)}(z, k; a, K)$ to convergence.

- Loop over the aggregate states (A, K) . For each (A, K) , use any nonlinear equation system solver in p, K' to obtain $p(A, K)$ and $K'(A, K)$. The method used in this paper is dampened fixed-point iteration in the pair p, K' .

- For each value of p, K' , evaluate on the discretized grid for z and some spline projection grid for k the following equations:

$$V_{(s+1)}^A(z, k; A, K) = \max_{k', n} \left\{ p \left(zAk^\alpha n^\nu - k' + (1 - \delta)k - \frac{\phi}{p}n \right) + \beta \mathbb{E}_{z', A'} V_{(s)}(z', k'; A', K') \right\}$$

$$V_{(s+1)}^{NA}(z, k; A, K) = \max_n \left\{ p \left(zAk^\alpha n^\nu - \frac{\phi}{p}n \right) + \beta \mathbb{E}_{z', A'} V_{(s)}(z', (1 - \delta)k; A', K') \right\}$$

$$\xi^*(z, k; A, K) = \frac{V_{(s+1)}^A(z, k; A, K) - V_{(s+1)}^{NA}(z, k; A, K)}{\phi}$$

$$V_{(s+1)}(z, k; A, K) = -\phi \int_0^{\xi^*(z, k; A, K)} \xi dG(\xi) + G(\xi^*(z, k; A, K))V_{(s+1)}^A(z, k; A, K) + (1 - G(\xi^*(z, k; A, K)))V_{(s+1)}^{NA}(z, k; A, K)$$

- Then, compute the errors to the system of equations in p and K' given by

$$\frac{1}{p} = \sum_{i=1}^{n_z} \tilde{\pi}_{zi} \int_{\underline{k}}^{\bar{k}} \left[\begin{array}{c} zAk^\alpha n(z_i, k; A, K)^\nu \\ -G(\xi^*(z_i, k; A, K))(k'(z_i, k; A, K) - (1 - \delta)k) \end{array} \right] P(k, \rho_{A, K}^{z_i}) dk$$

$$K' = \sum_{i=1}^{n_z} \tilde{\pi}_{zi} \int_{\underline{k}}^{\bar{k}} \left[\begin{array}{c} G(\xi^*(z_i, k; A, K))k'(z_i, k; A, K) \\ +(1 - G(\xi^*(z_i, k; A, K)))(1 - \delta)k \end{array} \right] P(k, \rho_{A, K}^{z_i}) dk,$$

where above the calculation requires the ability to compute $n(z_i, k; A, K)$, $k'(z_i, k; A, K)$, and $\xi^*(z_i, k; A, K)$ at a set of Simpson integration nodes in k . These can be computed by recalculation of the right hand sides of the Bellman equations above at the quadrature nodes in k , and this task is simplified by noting that capital policies depend only on idiosyncratic productivity and that the labor demand optimization is static and can be obtained in closed-form. Also, the densities $P(k, \rho_{A, K}^{z_i})$ must be evaluated using moments which reflect both the reference moments for higher order terms but also the current mean level of aggregate capital. Therefore, the first moments of capital used in the construction of $P(k, \rho_{A, K}^{z_i})$ are linearly shifted to deliver consistency with K as the current aggregate shape.

- Error on the value function iteration is given by $\|V_{(s+1)} - V_{(s)}\|$, which can be defined as desired. The solution used in this paper is based on the max absolute percentage difference. If the value function has converged, exit the value function iteration process. If not, go iteration $s + 1$.
- (c) At this point, a decision must be made. If the reference moments used in the solution are to be held fixed at their steady-state values, the model is now solved. However, if the reference moments are to be updated with simulation, then simulation must be performed as part of the solution itself. In either case, simulation proceeds as follows, noting that a set of exogenous productivity draws A_t for T periods have been made outside of the solution loop. Given converged values V , simulate the economy for a large number of periods, in each period imposing market clearing to obtain p, K' . For each period, do the following, taking advantage of the functional forms assumed for the cross-sectional density.
- Start period t with a set of coefficients $\rho_t^{z_i}$ and moments m_t (the first $n_z n_M$ moments of the idiosyncratic capital density conditional upon idiosyncratic productivity) which jointly pin down the cross-sectional density in (z, K) for the period. Then using the value function from above and the same fixed point iteration approach, compute the equilibrium p, K' consistent with both the value function and the cross-sectional distributions for the current period.

- Compute the value of all reference moments for the next period $t + 1$. These are higher-order centered moments of the cross-sectional distribution of capital next period, and can be computed directly via quadrature, given the policies and cross-sectional distribution coefficients of period t .
 - Now that the reference moments are set for period $t + 1$, along with the aggregate capital state, compute the coefficients of the cross-sectional distribution associated with period $t + 1$ using the exact same minimization step as above.
- (d) After simulation is completed for all T periods, and a certain number T_{erg} of initial periods are discarded, you have two options. If the reference moments are held constant at their steady-state values, you simply have an unconditional simulation of the model. If a fixed-point on the reference moments is desired, then update the reference moments in the outer loop now. The appropriate method depends on your assumptions for the reference moments. If you have assumed one unconditional constant set of reference moments not varying with aggregates, compute the unconditional average of each reference moment over the simulation. If you have assumed reference moments which vary with aggregate productivity, then compute the conditional average of each reference moment on the value of aggregate productivity. If you have assumed a reference moment forecasting system, update this system with OLS. Return to step (a) if the reference moments have not converged based on some criterion, say max absolute difference.

A.4 Explicit Aggregation (XPA)

First discretize the aggregate productivity process A and solve the steady-state model for each aggregate productivity value A_k , $k = 1, \dots, n_A$. For each level, save values of equilibrium price and capital stocks $p^{SS}(A_k), K^{SS}(A_k)$. Also, discretize the idiosyncratic productivity process and store the exogenous ergodic distribution $\tilde{\pi}_z$ of discretized idiosyncratic productivity for future use.

Set up the aggregate state space to include (A, K) and posit forecast rules for market-clearing prices and next-period capital identical to the KS case. Then, solve the model *exactly* as in the case of the KS algorithm, with two modifications:

- (2') Replace KS simulation step (2) with an “explicit aggregation” step. In particular, loop over aggregate states (A, K) , where A varies over its discretized grid and K varies over the same projection grid used to compute the value functions.
- At each point (A, K) , compute market-clearing prices via a solution routine like bisection over price p , with implied output, investment, consumption, and next-period aggregate capital given a guess \tilde{p} defined by

$$\begin{aligned}\tilde{Y} &= \sum_{z_i} \tilde{\pi}_i y(z_i, K), \quad y(z_i, K) = z_i A K^\alpha n^*(z_i, K)^\nu \\ \tilde{I} &= \sum_{z_i} \tilde{\pi}_i \left[G(\tilde{\xi}^*(z_i, K)) (k'^*(z_i, K) - (1 - \delta)K) \right] \\ \tilde{C} &= \tilde{Y} - \tilde{I} \\ \tilde{K}' &= \sum_{z_i} \tilde{\pi}_i \left[G(\tilde{\xi}^*(z_i, K)) k'^*(z_i, K) + (1 - [G(\tilde{\xi}^*(z_i, K))]) (1 - \delta)K \right]\end{aligned}$$

with clearing error given as before by $\frac{1}{\tilde{p}} - \tilde{C}$.

- Note that the end of this iteration generates a dataset $(A, K) \rightarrow (p, K')$ as (A, K) varies over the discretization and projection grids.
- (3') Replace KS update step (3) with a forecast update rule step using the dataset defined by explicit aggregation step (2').

- First, update the current vector of forecast rule coefficients $\hat{\Gamma}_{(s)}$ by estimating $(\hat{\alpha}_p(A), \hat{\beta}_p(A), \hat{\alpha}_K(A), \hat{\beta}_K(A))$ with OLS on the explicit aggregation dataset, segmented by discretized value of A .
- Then, define a vector of n_A bias correction terms

$$x_p^{Bias}(A) = \hat{\alpha}_p(A) + \hat{\beta}_p(A) \log(K^{SS}(A)) - \log(p^{SS}(A))$$

$$x_K^{Bias}(A) = \hat{\alpha}_K(A) + \hat{\beta}_K(A) \log(K^{SS}(A)) - \log(K^{SS}(A))$$

and adjust the new forecast rule coefficients' constant terms with

$$\hat{\alpha}_p(A) \leftarrow \hat{\alpha}_p(A) - x_p^{Bias}(A)$$

$$\hat{\alpha}_K(A) \leftarrow \hat{\alpha}_K(A) - x_K^{Bias}(A).$$

- Then, check the estimated coefficients against the old coefficients $\hat{\Gamma}_{(s)}$. If they are within some tolerance according to max absolute deviations, the model is solved and exit the routine. If the forecast rules have not converged, use dampened fixed-point iteration to update the forecast rule system $\hat{\Gamma}_{(s+1)}$ based on rule (s) and the newly estimated system.

Note that the differencing off of $x^{Bias}(A)$ is an attempt to correct for the Jensen's inequality bias induced by substitution of aggregate states into idiosyncratic policies. The bias results from lack of variation in the cross-section of idiosyncratic capital when recovering market-clearing prices and next-period capital stocks. However, the steady-state model prices and capital stocks do incorporate cross-sectional integration over a distribution of idiosyncratic capital presumably similar to the distributions within the model with aggregate uncertainty. The modification by the term $x^{Bias}(A)$ requires that the estimated forecast system be able to exactly reproduce as a fixed point the steady-state prices and aggregate capital stocks $p^{SS}(A)$ and $K^{SS}(A)$, conditional upon aggregate productivity.

Note also that after the model is solved, simulation is completed exactly as in the KS algorithm, using the Young (2010) nonstochastic or histogram-based approach, and requiring market-clearing in each period with integration over the full cross-sectional distribution of idiosyncratic capital.

A.5 Projection Plus Perturbation (REITER)

The REITER solution method is based on three steps, and provides a perturbation approximation to the full rational expectations equilibrium. The first step is to solve the steady-state version of the model, with no aggregate uncertainty and aggregate productivity held fixed at a value of $A = 1$. The steady-state solution is identical to the one used, for example, as an input into the PARAM solution. The second step is to set up a system of nonlinear equations defining the model's equilibrium, which is covered in the first subsection below. The final step is to linearize and solve the system using standard numerical differentiation and solution techniques, covered in the second subsection below.

A.5.1 Nonlinear System of Equations in the Discretized Model

We first establish a grid of n_z idiosyncratic productivity points and a Markov transition matrix $\Pi_{zij} = \mathbb{P}(z_{t+1} = z_j | z_t = z_i)$ following Tauchen (1986). Then, we establish a grid of n_k idiosyncratic capital stock nodes k_i , which will function as knot points for cubic spline interpolation of the value functions. Finally, a denser simulation grid of n_d idiosyncratic points will be used to store the distribution. We then have the following equations, which are to be linearized around the zero-aggregate uncertainty steady-state of the model.

$$V_{t-1}^A(z_i, k_j) = \max_{k'} \left\{ p_{t-1} (f(A_{t-1}, z_i, k_j, w_{t-1}) - k' + (1 - \delta)k_j) + \beta \sum_{z_{i'}} \Pi_{zii'} V_t(z_{i'}, k') \right\} + \eta_{t,i,j}^A$$

$$\begin{aligned}
f(A_{t-1}, z_i, k_j, w_{t-1}) &= (1 - \nu) \left(\frac{\nu}{w_{t-1}} \right)^{\frac{\nu}{1-\nu}} (A_{t-1} z_i)^{\frac{1}{1-\nu}} k_j^{\frac{\alpha}{1-\nu}} \\
V_{t-1}^{NA}(z_i, k_j) &= p_{t-1} f(A_{t-1}, z_i, k_j, w_{t-1}) + \beta \sum_{z_{i'}} \Pi_{z_{i'}} V_t(z_{i'}, \tilde{k}_j) + \eta_{t,i,j}^{NA} \\
\tilde{k}_j &= \max((1 - \delta)k_j, k_1) \\
V_{t-1}(z_i, k_j) &= -\phi \int_0^{\xi_{t-1}^*(z_i, k_j)} \xi dG(\xi) + G(\xi_{t-1}^*(z_i, k_j)) V_{t-1}^A(z_i, k_j) \\
&\quad + [1 - G(\xi_{t-1}^*(z_i, k_j))] V_{t-1}^{NA}(z_i, k_j) \\
\xi_{t-1}^*(z_i, k_j) &= \frac{V_{t-1}^A(z_i, k_j) - V_{t-1}^{NA}(z_i, k_j)}{\phi} \\
\mu_t(z_{i'}, k_{j'}) &= \sum_{z_i} \sum_{k_j} \Pi_{z_{i'}} \mu_{t-1}(z_i, k_j) \left[\begin{array}{l} \omega^a(i, j, j') \mathbb{I}(k'(z_i, k_j) \in [k_{j'-1}, k_{j'+1}]) \\ + \omega^{na}(i, j, j') \mathbb{I}(\tilde{k}_j \in [k_{j'-1}, k_{j'+1}]) \end{array} \right] \\
\omega^a(i, j, j') &= \begin{cases} G(\xi_{t-1}^*(z_i, k_j)) \left(\frac{k'(z_i, k_j) - k_{j'}}{k_{j'+1} - k_{j'}} \right) & k'(z_i, k_j) \in [k_{j'}, k_{j'+1}], 1 < j' < n_d \\ G(\xi_{t-1}^*(z_i, k_j)) & k'(z_i, k_j) \geq k_{n_d}, j' = n_d \\ G(\xi_{t-1}^*(z_i, k_j)) & k'(z_i, k_j) \leq k_1, j' = 1 \end{cases} \\
\omega^{na}(i, j, j') &= \begin{cases} [1 - G(\xi_{t-1}^*(z_i, k_j))] \left(\frac{\tilde{k}_j - k_{j'-1}}{k_{j'} - k_{j'-1}} \right) & \tilde{k}_j \in [k_{j'-1}, k_{j'}], 1 < j' < n_d \\ [1 - G(\xi_{t-1}^*(z_i, k_j))] & \tilde{k}_j \geq k_{n_d}, j' = n_d \\ [1 - G(\xi_{t-1}^*(z_i, k_j))] & \tilde{k}_j \leq k_1, j' = 1 \end{cases} \\
\frac{1}{p_{t-1}} &= \sum_{z_i} \sum_{k_j} \mu_{t-1}(z_i, k_j) [y(A_{t-1}, z_i, k_j, w_{t-1}) - G(\xi_{t-1}^*(z_i, k_j)) (k'(z_i, k_j) - (1 - \delta)k_j)] \\
w_{t-1} &= \frac{\phi}{p_{t-1}} \\
Y_{t-1} &= \sum_{z_i} \sum_{k_j} \mu_{t-1}(z_i, k_j) y(A_{t-1}, z_i, k_j, w_{t-1}) \\
y(A_{t-1}, z_i, k_j, w_{t-1}) &= \left(\frac{\nu}{w_{t-1}} \right)^{\frac{\nu}{1-\nu}} (A_{t-1} z_i)^{\frac{1}{1-\nu}} k_j^{\frac{\alpha}{1-\nu}} \\
I_{t-1} &= \sum_{z_i} \sum_{k_j} \mu_{t-1}(z_i, k_j) G(\xi_{t-1}^*(z_i, k_j)) (k'(z_i, k_j) - (1 - \delta)k_j) \\
N_{t-1} &= \sum_{z_i} \sum_{k_j} \mu_{t-1}(z_i, k_j) \left[n(A_{t-1}, z_i, k_j, w_{t-1}) + \int_0^{\xi_{t-1}^*(z_i, k_j)} \xi dG(\xi) \right] \\
n(A_{t-1}, z_i, k_j, w_{t-1}) &= \left(\frac{\nu}{w_{t-1}} \right)^{\frac{1}{1-\nu}} (A_{t-1} z_i)^{\frac{1}{1-\nu}} k_j^{\frac{\alpha}{1-\nu}} \\
\log(A_t) &= \rho_A \log(A_{t-1}) + \varepsilon_t, \quad \varepsilon_t \sim iid N(0, \sigma_A^2)
\end{aligned}$$

Together, the above system represents a total of $n_s = 3n_z n_k + n_z n_d + 5$ equations in the $n_s \times 1$ vector of endogenous variables $X_t = ((V_t^A)', (V_t^{NA})', V_t', \mu_t', \log(p_t), \log(Y_t), \log(I_t), \log(N_t), \log(A_t))'$, together with the exogenous shock process ε_t . We write this nonlinear rational expectations system corresponding to the discretized model as $F(X_t, X_{t-1}, \eta_t, \varepsilon_t) = 0$, where $F(\cdot)$ is the left hand side minus the right hand side of each of the n_s equations above. Note that we know this equation is satisfied exactly at the steady-state value of $X_t = X_{t-1} = X^{SS}$ and $\eta_t = 0$, $\varepsilon_t = 0$ corresponding to no aggregate uncertainty and $A_t = 1$. The vectors X_t are $n_s \times 1$ while the vector η_t is $n_\eta \times 1$, where $n_\eta = 2n_z n_k$.

A few practical comments are in order. First, in general the approximation nodes used for interpolation of the value function in k will be different (and less dense) than the discrete values of k used to store the cross-sectional distribution μ_t above. The value ε_t is the exogenous shock to aggregate productivity. The vector η_t is the stacked set of expectational errors using Sims (2002) notation which must be applied to the expectations in the Bellman equations above, and these expectational errors depend upon aggregates only as the idiosyncratic uncertainty reflected in the discretization of idiosyncratic productivity is already taken into account through the summation with respect to the transition matrix Π_z .

Also, note that when the system above is evaluated, policies must be reoptimized to obtain the values of $k'(z_i, k_j)$ used in the distributional transitions above.¹⁸ In practice, this reoptimization to compute F is not particularly burdensome because the optimization must only be performed once, and because the capital policies conditional upon adjustment can be shown to depend only on the value of idiosyncratic productivity z . The functions f , y , and n above represent the reduced-form expressions for revenue net of labor costs, output, and labor input at firms, conditional upon the analytic solution to the static labor optimization problem.

A.5.2 Linearizing the System

We then numerically differentiate F with respect to each of its arguments, at the steady-state, to obtain the system of n_s equations below

$$F_1(X_t - X^{SS}) + F_2(X_{t-1} - X^{SS}) + F_4\eta_t + F_5\varepsilon_t = 0,$$

where $F_1 = \frac{\partial F}{\partial X_t}(n_s \times n_s)$, $F_2 = \frac{\partial F}{\partial X_{t-1}}(n_s \times n_s)$, $F_3 = \frac{\partial F}{\partial \eta_t}(n_s \times n_\eta)$, and $F_4 = \frac{\partial F}{\partial \varepsilon_t}(n_s \times 1)$ are the approximated derivative matrices evaluated at the steady-state of the model. This system of equations can be solved using a large variety of solution methods for linear rational expectations models, including those in Christiano (2002) or Sims (2002). In the calculations performed in this paper, the Sims (2002) algorithm as applied by the `gensys` software in MATLAB available on Chris Sim's website is used to solve the linear model, and the numerical differentiation is carried out by forward differentiation from the steady-state with relative step size $1e - 6$.

For concreteness, note that the linear solution to the model is simply a set of coefficient matrices A ($n_s \times n_s$) and B ($n_s \times 1$) such that locally around the steady-state the model satisfies

$$(X_t - X^{SS}) = A(X_{t-1} - X^{SS}) + B\varepsilon_t.$$

Immediately, the traditional local impulse responses $IRF_t, t = 1, \dots, T_{IRF}$ ($n_s \times 1$ vectors) to a shock to aggregate productivity can be computed as

$$IRF_t = A^{t-1}B,$$

and the model can be simulated by drawing $N(0, \sigma_A^2)$ shocks and substituting directly into the solution equation above.

At this point, it's important to note that by numerically differentiating the system F we are assuming that although nonlinearity and threshold decision rules exist and are preserved at the microeconomic level, the dependence of these micro-level decisions, as embedded in the value functions and distributional transition weights above, on aggregate shocks to productivity is smooth. In the context of the Khan and Thomas (2008) model with stochastic adjustment costs, the resultingly smooth value functions and policies make such an assumption sensible. However, models with discrete choices and fixed, nonstochastic adjustment costs, such as those in Bloom et al. (2012), which can not be expected to see policies vary smoothly with aggregate shocks would not allow for a reasonable application of the REITER approach. Finally, for the levels of discreteness chosen in this paper's solution, the linear system has approximately 900 equations, but for a denser discretization model reduction techniques can be applied which still allow for the model's solution. These reduction techniques are laid out in Reiter (2010a) and recently applied in McKay and Reis (2013).

¹⁸This required reoptimization when evaluating the system F differentiates the REITER solution method in Euler equation-based models, such as the incomplete markets model, and Bellman equation-based models such as the heterogeneous firms model studied here.

B General Numerical Choices and Additional Accuracy Statistics

To complement the general discussion of each solution algorithm, it is also useful to list some practical choices made in the projection, optimization, and discretization of the model, with further information available in Table B1.

Table B1: Some Choices for the Numerical Solutions

Object	Value	Type or Location Used
k nodes for V spline interpolation	10	Loglinear [0.1,8.0]
k points for KS, XPA distribution histogram	50	Loglinear [0.1,8.0]
k points for REITER distribution histogram	150	Loglinear [0.1,8.0]
K points for V linear interpolation	10	Linear [1.25,2.0]
A points in discretization	5	Loglinear [0.94,1.0562]
z points in discretization	5	Loglinear [0.9176,1.0897]
V or k' , ξ^* tolerance	0.0001	All methods
Forecast rule coefficient tolerance	0.001	KS, XPA methods
Howard accelerations	50	All methods but PARAM
Number of reference moments	4	PARAM method densities
Quadrature nodes in k	36	PARAM method integration
Nondiscarded simulation periods $T - T_{erg}$	2000	All methods
Number of IRF simulations	2000	All methods
IRF simulations length	50	All methods
IRF shock period	25	All methods

Note: The table above indicates specific methods, tolerances, grid sizes and densities, and simulation lengths used in the numerical implementation of the solution techniques discussed in the paper. The code used to produce the results in this paper is available on Stephen Terry's website.

- The mean or aggregate level of capital K is used as the approximating moment m for the cross-sectional distribution $\mu(z, k)$ for the three solution methods requiring such an approximation (KS, PARAM, and XPA).
- In all solution techniques, idiosyncratic and aggregate productivity processes z and A are discretized according to Tauchen (1986) and along grids spanning two standard deviations of the unconditional process standard deviation around the process steady-state.¹⁹
- Value functions are approximated as cubic splines in idiosyncratic capital k with a natural spline endpoint condition, and using linear interpolation in aggregate capital K . For the KS, XPA, and no-aggregate uncertainty models, the firm problem is solved using policy iteration with Howard acceleration of the value function, while in the PARAM method value function iteration is performed.
- The idiosyncratic capital policies within the Bellman equations are determined using Brent optimization (KS, XPA, PARAM, and no aggregate uncertainty), or golden section search (REITER differentiation).
- Price-clearing is performed during simulation using bisection (XPA, no aggregate uncertainty) or hybrid bisection/inverse quadratic interpolation (KS). Within the model solution step, joint search for clearing price and next period aggregate capital is performed using dampened fixed point iteration (PARAM).
- Minimization of the density-based objective determining distributional coefficients during simulation and solution of the PARAM model is performed using a standard quasi-Newton algorithm with symmetric rank-one (SR1) updates to the inverse Hessian approximation.

¹⁹For the REITER method, the aggregate productivity process A must take on continuous values within the solution process, but when simulating we use both the discretized process as well as separately consider the continuously simulated shocks.

- Within the PARAM solution, integration over the cross-sectional densities is performed using standard Simpson quadrature rules.

Table B2: Additional Accuracy Statistics for Forecast Rules

Statistic	K' , KS	K' , XPA	K' , PARAM	p , KS	p , XPA,	p , PARAM
RMSE, $A = A_1$	0.0182	0.5506	0.1257	0.0122	0.8911	0.0501
RMSE, $A = A_2$	0.0337	0.2081	0.0903	0.0247	0.1832	0.0607
RMSE, $A = A_3$	0.0486	0.3535	0.1007	0.0430	0.1230	0.0663
RMSE, $A = A_4$	0.0613	0.3804	0.1899	0.0604	0.2732	0.1211
RMSE, $A = A_5$	0.0792	0.4256	0.3090	0.0883	0.7054	0.2010
R^2 , $A = A_1$	0.9999	0.9350	0.9973	0.9999	-0.1915	0.9977
R^2 , $A = A_2$	0.9999	0.9910	0.9988	0.9997	0.9777	0.9981
R^2 , $A = A_3$	0.9998	0.9670	0.9981	0.9993	0.9916	0.9977
R^2 , $A = A_4$	0.9996	0.9604	0.9912	0.9984	0.9468	0.9910
R^2 , $A = A_5$	0.9992	0.9575	0.9692	0.9959	0.3850	0.9696

Note: The table above reports internal accuracy statistics based on unconditional simulations for the three solution methods with explicit forecast mappings from the approximate aggregate state (A, K) to realized next-period capital K' (the first three columns) and to market-clearing prices p (the final three columns). The first five rows represent the root mean squared error, in 100 times logs, between the forecasts of one-period ahead capital and current price and realized values, conditional upon the current discretized value of aggregate productivity A_1, \dots, A_5 . The final five rows represent the R^2 of each forecast regression, conditional upon the current value of aggregate productivity. The exogenous aggregate productivity process reflects a Markov chain discretized using the Tauchen (1986) procedure and is held constant across solution methods during the simulation. To achieve this, an identical simulated discretized productivity process is input directly into the KS, PARAM, and XPA solutions, while a series of continuous aggregate shocks exactly replicating the discretized productivity process is input into the REITER solution. For each method, the full 2000 period simulation for each solution method begins after 500 periods, with an initial burn-in period discarded to avoid the influence of initial conditions on the simulated aggregates. The exogenous productivity simulation used for calculation of the accuracy statistics are distinct from the simulation used within the model-solution step of any algorithm.

C Simulated Impulse Responses with Nonlinear Models

As noted in the main text, nonlinear impulse response analysis must take into account variation in the initial conditions and size of the imposed shocks, and following Koop et al. (1996) we take the following approach to compute the average conditional response to a one-standard deviation aggregate productivity shock:

1. Fix a large number N of simulations, a per-simulation length T_{IRF} , and a shock-period T_{shock} .
2. Independently draw exogenous $u_{it} \sim U(0, 1)$ shocks for each period t of each simulation i , as well as one “shock occurrence” $s_i \sim U(0, 1)$ draw for each simulation i .
3. For each simulation i , based on comparison of the shocks u_{it} with the entries of the cumulated transition matrix of discretized aggregate productivity, create two versions of simulated aggregate productivity series A_{it}^{shock} and $A_{it}^{noshocks}$, which are identical until the shock-period for $t < T_{shock}$. In the shock period T_{shock} for each simulation, compare the shock occurrence draw s_i for simulation i with the cutoff thresholds \bar{s} defined below. If $s_i \leq \bar{s}$, set $A_{it}^{shock} = A_{n_A}$, where n_A is the highest level of the discretized aggregate productivity process, and then allow A_{it}^{shock} to transit normally based on u_{it} for $t > T_{shock}$. If $s_i > \bar{s}$, allow A_{it}^{shock} to transit normally based on u_{it} for all periods $t \geq T_{shock}$. In either case, allow $A_{it}^{noshocks}$ to transit normally based on u_{it} for all $t \geq T_{shock}$.
4. Simulate any aggregate series of interest X twice, using both A_{it}^{shock} and $A_{it}^{noshock}$ exogenous processes. The impulse responses \hat{x}_t of series X in period t to an aggregate productivity shock in period T_{shock} is given by

$$\hat{x}_t = 100 \frac{1}{N} \sum_{i=1}^N \log \left(\frac{X_{it}^{shock}}{X_{it}^{noshock}} \right).$$

To obtain an average percentage innovation in aggregate productivity which equals σ_A exactly, we choose the shock threshold \bar{s} to solve

$$\sigma_A = \bar{s} \sum_{k=1}^{n_A} \tilde{\pi}_{Ak} (\log(A_{n_A}) - \log(A_k)),$$

where $\tilde{\pi}_A$ is the ergodic distribution of the discretized aggregate productivity process A .

As noted in Appendix B, to compute the impulse responses plotted in the main text, we set $T_{IRF} = 50$, $T_{shock} = 25$, and $N = 2000$, and we hold exogenous draws u_{it}, s_i constant across simulation methods. We also set $n_A = 5$.

One final comment is in order regarding the REITER solution method. Because the REITER approach yields a linearized solution, the simulation-based analysis of Koop et al. (1996) is unnecessary. Although for completeness and comparability we perform the simulation-based impulse response with the REITER method, a much simpler alternative, invariant to shock scaling or initial conditions, is available. In particular, when writing the REITER solution as $X_t = AX_{t-1} + B\varepsilon_{At}$, where X_t is the endogenous vector defined in Appendix A and ε_{At} is a continuous shock to aggregate productivity $\varepsilon_{At} \sim N(0, \sigma_A^2)$, the linear impulse response is given by

$$\hat{x}_t^{REITER-LIN} = A^{t-1}B,$$

which is a vector of responses of each endogenous variable to an innovation in ε_{At} . Figure C1 plots the simulated impulse response and the linearized version immediately computable from the solution matrices, unsurprisingly confirming that they are virtually identical.

D Adding Aggregate Complexity with the REITER Method

To add a persistent demand or preference shock D_t as well as a labor disutility process D_t^N to the baseline heterogeneous firms context considered here, replace the household objective based on maximization of

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t, 1 - N_t)$$

with maximization of the alternative

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t D_t U^*(C_t, N_t, D_t^N),$$

$$U^*(C_t, N_t, D_t^N) = \log(C_t) - \phi D_t^N N_t$$

$$\log(D_t) = \rho_D \log(D_{t-1}) + \sigma_D \varepsilon_{Dt}$$

$$\log(D_t^N) = \rho_N \log(D_{t-1}^N) + \sigma_N \varepsilon_{Nt}.$$

First order conditions for the modified problem immediately reveal that firms owned by households subject to the shocks D_t and D_t^N will value dividend flows at the rate $p_t D_t$ where $p_t = \frac{1}{C_t}$, and equilibrium wages are given by $w_t = \frac{\phi D_t^N}{p_t}$. By replacing price p_t and wage w_t with the formulas above in the system F from the REITER discussion of Appendix A, and adding the evolution equations for D_t and D_t^N , we immediately obtain a system with two more endogenous variables and two more exogenous shock processes, which can be linearized and solved using numerical differentiation and the solution technique of Sims (2002), exactly as in the baseline REITER method.

The impulse responses from this model are computed using the somewhat arbitrary choices of $\rho_D = \rho_N = 0.859$, identical to the persistence of aggregate productivity, scaled by the standard deviations $\sigma_D = \sigma_N = 0.014$, also equal to baseline aggregate productivity volatility. The main text displays the response to a demand shock D_t , while Appendix Figures D1 and D2 plot local impulse responses to aggregate productivity and labor disutility shocks, respectively.




Figure B1: REITER Price Dynamics, Baseline and Denser Discretization

Note: The figure above plots a representative 100-period portion of a larger 2000-period unconditional simulation of the market clearing price series p_t obtained by applying the REITER solution method with two histogram grid densities: 150 grid points (solid line) and 200 grid points (dotted line). For each method, the full 2000 period simulation for each solution method begins after 500 periods, with an initial burn-in period discarded to avoid the influence of initial conditions on the simulated aggregates.




Figure B2: Realizations vs. One-Period Ahead Forecasts

Note: The figures above plot a representative 100-period portion of a larger 2000-period unconditional simulation of the three solution methods with explicit forecast mappings from the approximate aggregate state space (A, K) to realized prices and next-period productivity (p, K') , i.e. the KS (in black), PARAM (in green), and XPA (in red) methods. The first column plots the realized market-clearing price p (solid line) and the forecast value of price (dotted line), given the currently realized value of (A, K) . The second column plots the realized capital series (solid line) against the one-period ahead forecasts (dotted line) made on the basis of (A, K) . During the simulation, the exogenous aggregate productivity process reflects a Markov chain discretized using the Tauchen (1986) procedure and is held constant across solution methods during the simulation. To achieve this, an identical simulated discretized productivity process is input directly into the KS, PARAM, and XPA solutions, while a series of continuous aggregate shocks exactly replicating the discretized productivity process is input into the REITER solution. For each method, the full 2000 period simulation for each solution method begins after 500 periods, with an initial burn-in period discarded to avoid the influence of initial conditions on the simulated aggregates.

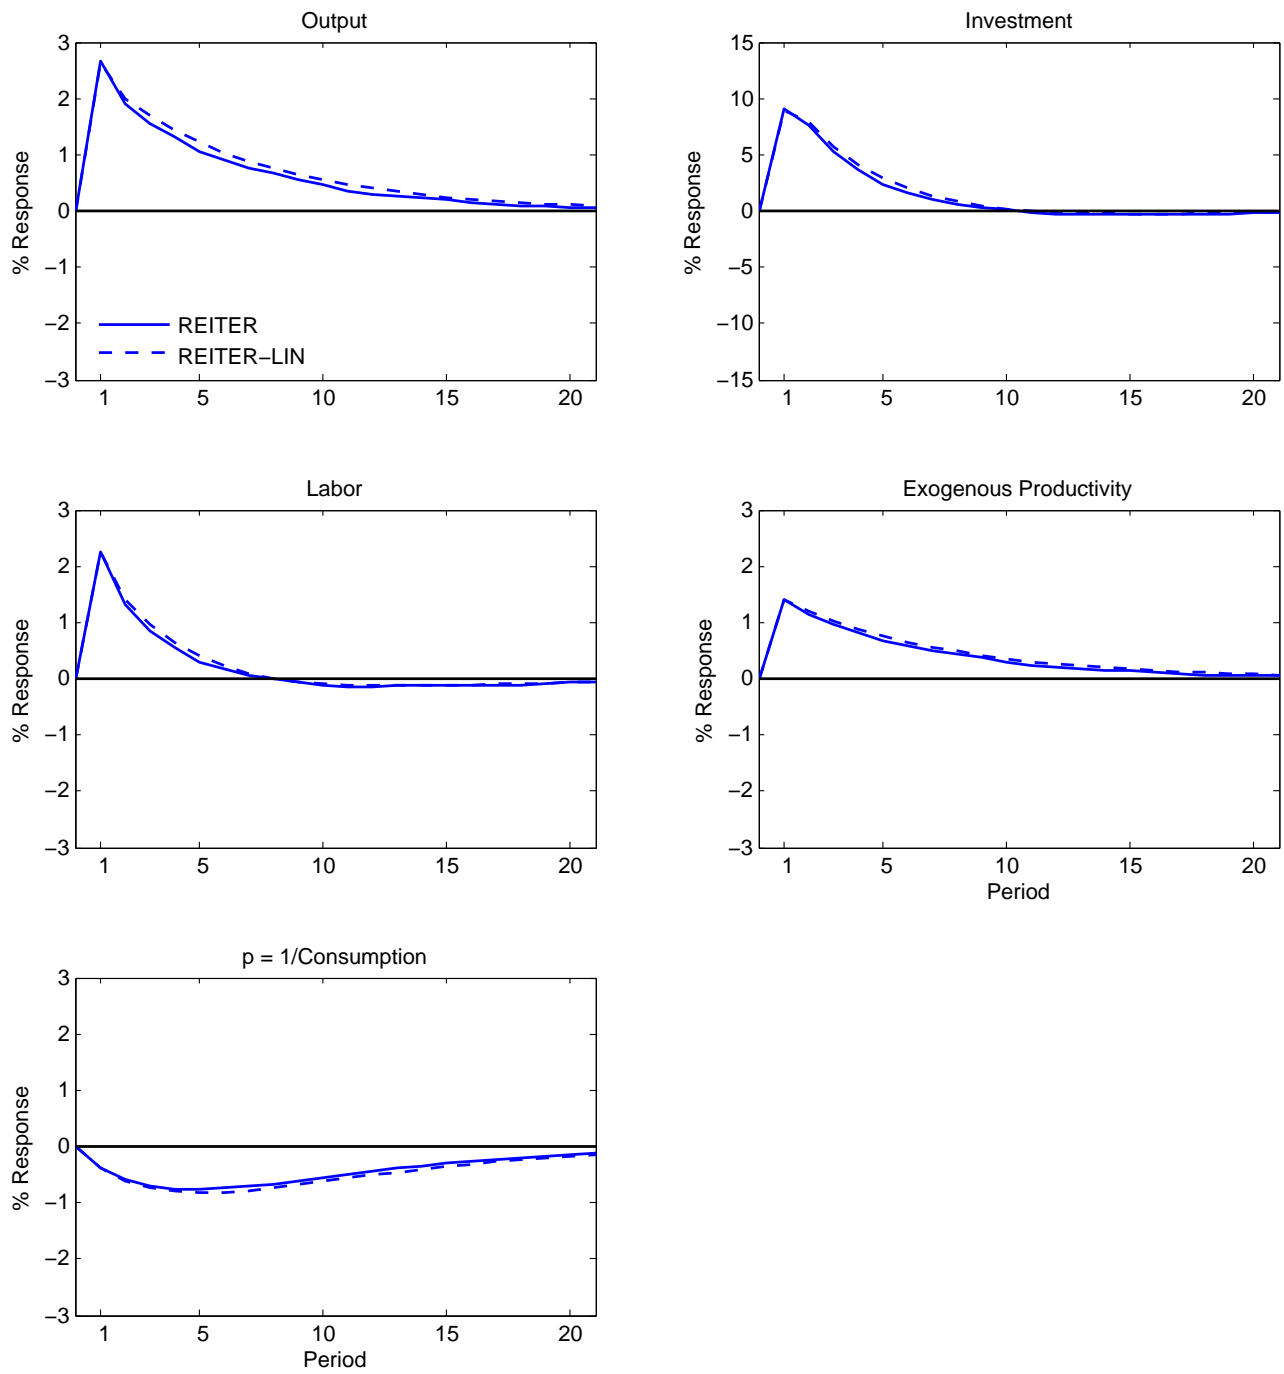


Figure C1: Impulse Response, Linear vs. Simulation-Based

Note: The figure above plots the linearized impulse response function (in dotted lines, labelled REITER-LIN) to a positive shock to the aggregate productivity series A_t in the baseline REITER solution of the heterogeneous firms model. Given the endogenous variables X_t of the linearized system describing the economy, a solution $X_t - X^{SS} = A(X_{t-1} - X^{SS}) + B\varepsilon_{A_t}$ of the model trivially yields impulse responses $A^{t-1}B$ representing the economy's local response to the aggregate productivity shock ε_{A_t} . The impulse response here is scaled by the assumed standard deviation of innovations to the aggregate productivity series, with the exogenous shocked series itself plotted in the right hand side of the second row. In solid lines labelled REITER, the figure plots analogous impulse responses computed as suggested by Koop et al. (1996), based on 2000 repeated simulations of a discretized aggregate productivity process of 50-period length each, with and without imposed productivity shocks at the period labelled 1 above. The REITER-LIN responses should be interpreted as 100 times log deviations from steady-states, while the REITER responses are equal to 100 times average log differences between the shocked and unshocked simulations.

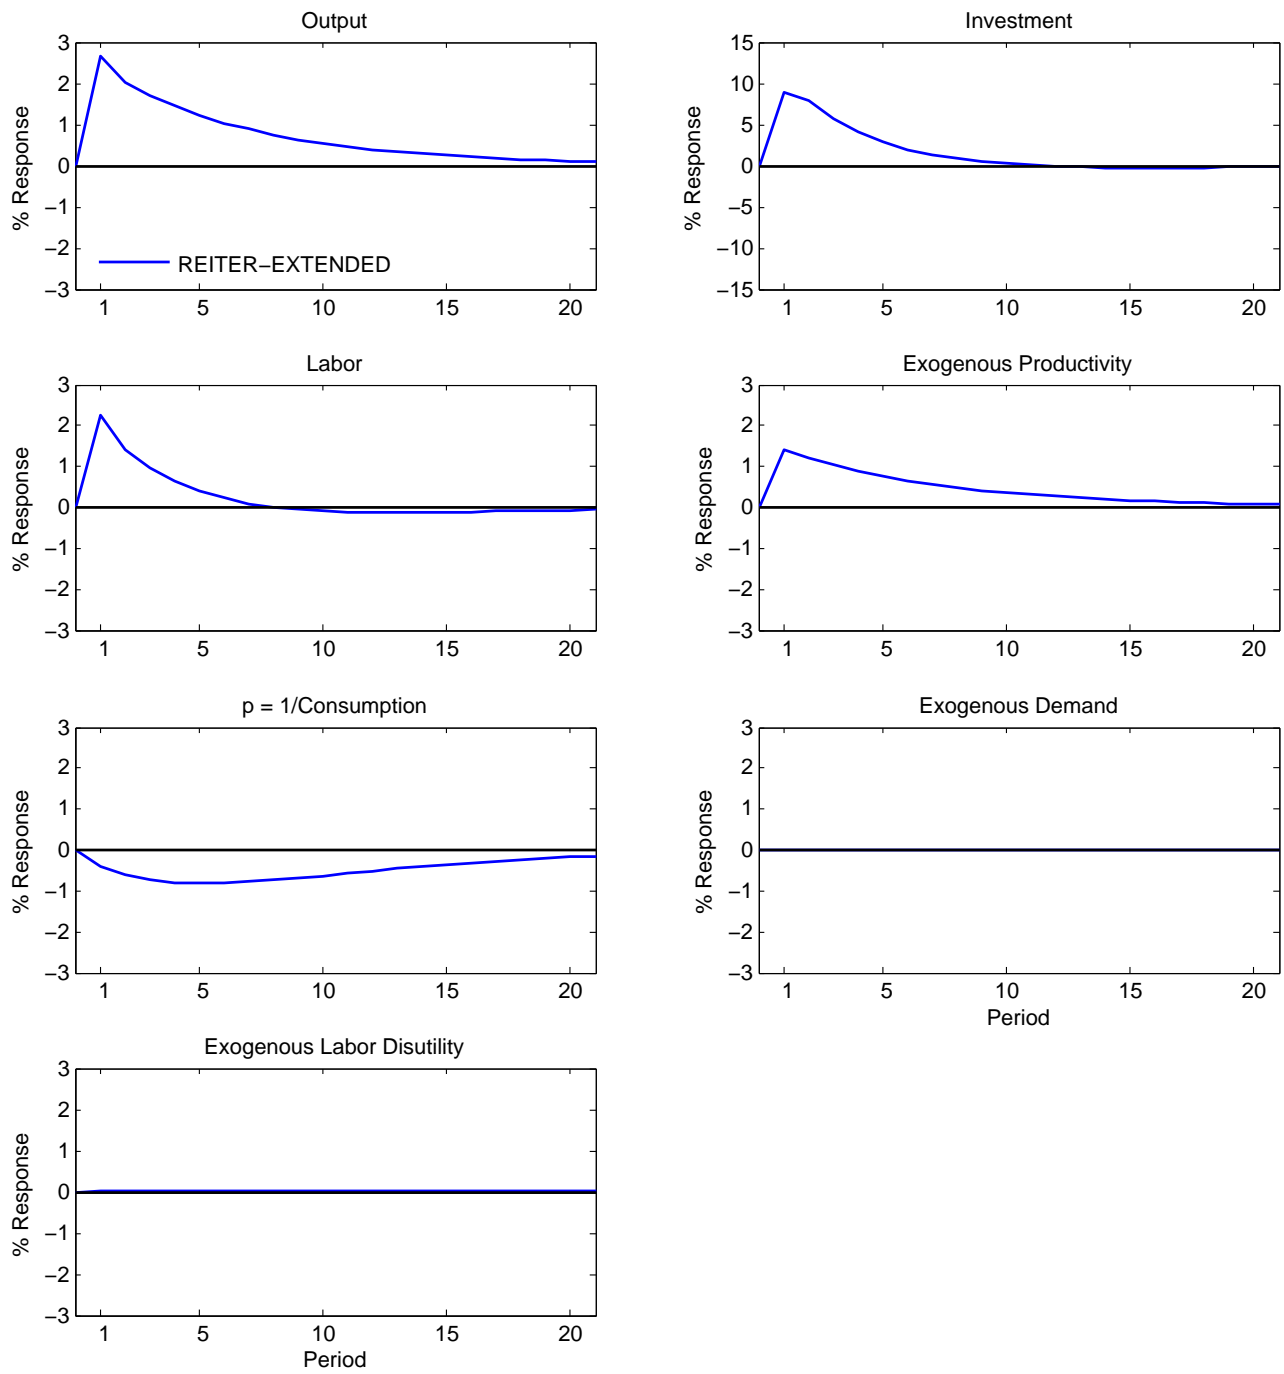


Figure D1: Impulse Response, Productivity Shock in Extended Model

Note: The figure above plots the linearized impulse response function to a positive shock to the aggregate productivity series A_t in the extension of the heterogeneous firms model augmented with aggregate demand and labor preference shocks. These impulse responses are computed after solving the extended model using the REITER technique and are therefore labelled REITER-EXTENDED. Given the endogenous variables X_t of the linearized system describing the economy, a solution $X_t - X^{SS} = A(X_{t-1} - X^{SS}) + B\epsilon_t$ of the model trivially yields impulse responses $A^{t-1}B$ representing the economy's local response to each aggregate shock in the vector ϵ_t . The impulse response here is scaled by the assumed standard deviation of innovations to the aggregate productivity series, with the exogenous shocked series itself plotted in the right hand side of the second row.

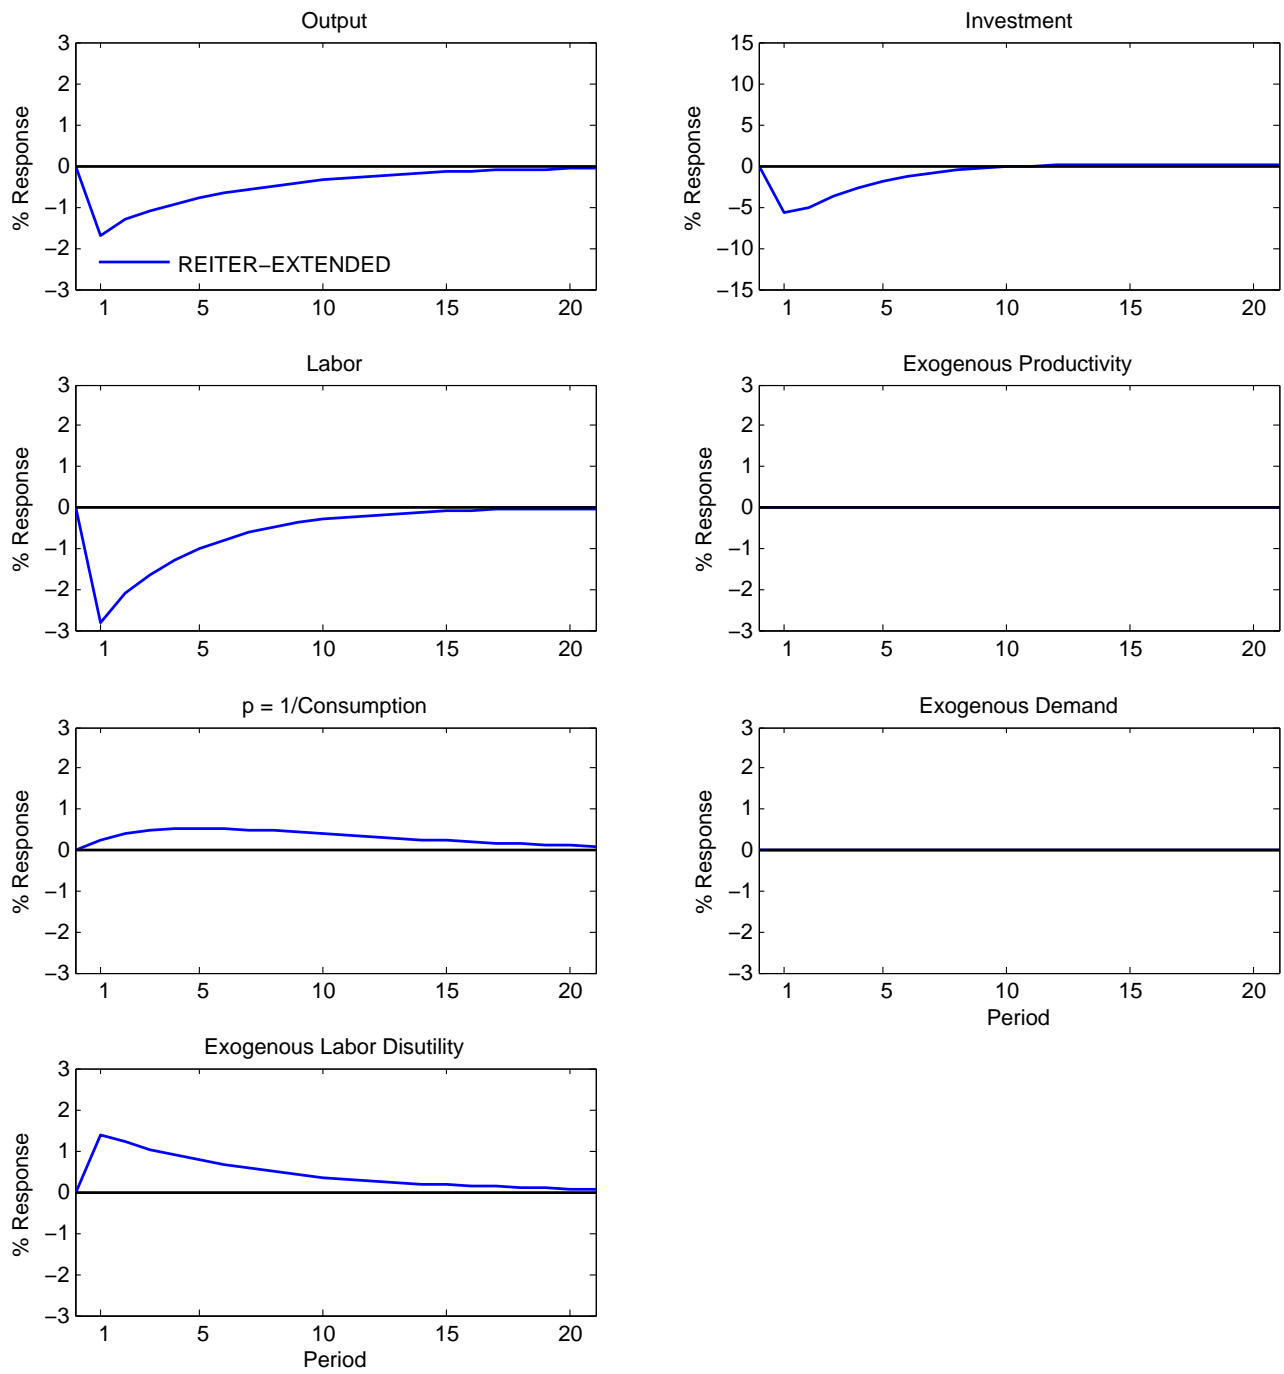


Figure D2: Impulse Response, Labor Preference Shock in Extended Model

Note: The figure above plots the linearized impulse response function to a positive shock to the labor disutility series D_t^N in the extension of the heterogeneous firms model augmented with aggregate demand and labor preference shocks. These impulse responses are computed after solving the extended model using the REITER technique and are therefore labelled REITER-EXTENDED. Given the endogenous variables X_t of the linearized system describing the economy, a solution $X_t - X^{SS} = A(X_{t-1} - X^{SS}) + B\epsilon_t$ of the model trivially yields impulse responses $A^{t-1}B$ representing the economy's local response to each aggregate shock in the vector ϵ_t . The impulse response here is scaled by the assumed standard deviation of innovations to the labor disutility series, with the exogenous shocked series itself plotted in the left hand side of the fourth row.