

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 edged 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 probably 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.

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 rms 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 rms model in Sunakawa (2012), and for comparability I rely on that paper's adaptation of the original Den Haan and Rendahl (2010) technique⁴. 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 online⁵.

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

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

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 :

$$p = \frac{1}{C(A; \theta)}; \quad w(A; \theta) = \frac{1}{p(A; \theta)}:$$

Above, prices and wages are written in terms of an aggregate state $A_t(\theta)$ including aggregate productivity A and a cross-sectional distribution θ 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 $(z; k)$ over idiosyncratic productivity and capital levels z and k .

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

$$V^A(z; k; A; \theta) = \max_{k^{\theta}; n} p(A; \theta) z A k^{\theta} n^{\theta} + (1 - \delta) k^{\theta} w(A; \theta) n + E_{z^{\theta}; A^{\theta}} V(z^{\theta}; k^{\theta}; A^{\theta}; \theta) :$$

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:

$$V^{NA}(z; k; A; \theta) = \max_n p(A; \theta) (z A k^{\theta} n^{\theta} - w(A; \theta) n) + E_{z^{\theta}; A^{\theta}} V(z^{\theta}; (1 - \delta) k^{\theta}; A^{\theta}; \theta) :$$

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 θ is less than a cutoff level

$$\theta(z; k; A; \theta) = \frac{V^A(z; k; A; \theta) - V^{NA}(z; k; A; \theta)}{p(A; \theta)}$$

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

2.3 Equilibrium

An equilibrium represents a set of firm value functions $\{V; V^A; V^{NA}\}$, firm policies and adjustment thresholds $k^{\theta}; n; \theta$, prices $p(A; \theta); w(A; \theta)$, and mappings $\theta; \rho$ 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

$$k^{\theta}(z; k; \theta; A; \theta) = k^{\theta}(z; k; A; \theta); \quad \theta < \theta$$

Aggregate consumption satisfies the resource constraint

$$C(A;) = Y(A;) - I(A;):$$

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

$$p(A;) = p(A;) = \frac{1}{C(A;)}; \quad w(A;) = \frac{1}{p(A;)}:$$

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

discussed above is replaced by $z(k; A; m)$, and the transition and price mappings are replaced by forecast rules $m^0 = \hat{m}$ and $p = \hat{p}$. 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 of 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

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^j for input into continuation values. Both p and m^j can be computed within the value function iteration step quite naturally by using fixed point iteration. After guessing values for $(p; m^j)$, 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 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.

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 $A(\cdot)$ with an approximation based on moments (

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 rms 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 rms 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 methods¹⁰. 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

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 $100(\log(X_t^{\text{method}}) - \log(X_t^{\text{KS}}))$ 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

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	PARAM
	Mean % Difference from KS			
PARAM				

ARAM

To generate a exibly-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%) 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,

Table 7: Accuracy Statistics for Forecast Rules

	Max DH, K^0	Mean DH, K^0	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^0 (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^0 , i.e. the maximum and mean values,

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,

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

studied the dynamics of a rich state-dependent pricing model using the REITER method¹⁶

As a concrete example within the heterogeneous rms 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 14%. 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 rms 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 rms 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*

- 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 2: Unconditional Continuous Simulation of REITER-OWN

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 x_t^{method} is the mean of $100(\log X_t^{\text{shock}} - \log(X_t^{\text{noshock}}))$ 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 $(p; K^0)$, 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.

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 $(z; k)$ 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:

$$V^A(z; k) = \max_{k^0, n} p A z k^n - k^0 + (1 - \delta) k - \frac{p}{p} n + E_{z^0} V(z^0; k^0)$$

$$V^{NA}(z; k) = \max_n p A z k^n - \frac{p}{p} n + E_{z^0} V(z^0; (1 - \delta) k)$$



$$\log(\hat{K}^0) = \kappa(A) + \kappa(A) \log(K)$$

Recall that from the household problem this yields a forecast wage level $\bar{w} \triangleq \bar{p}$. Given these choices, the solution algorithm works as follows. First, guess an initial forecast rule system $\hat{\pi}_{(1)} = (\overset{(1)}{p}; \overset{(1)}{p}; \overset{(1)}{K}; \overset{(1)}{K})$. Then, before the solution iteration takes place, draw a large number \bar{p}

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 $(z; k)$ 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 $\sim z$ 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 $(z_i; k)$ at discretized levels of z_i :

$$P(k; z_i) = \frac{z_i}{z_0} \exp \left[\frac{z_i}{1} (k - m_1^{z_i}) + \frac{z_i}{2} (k - m_1^{z_i})^2 - m_2^{z_i} \right]$$

+ : levels of

Loop over the aggregate states $(A; K)$. For each $(A; K)$, use any nonlinear equation system solver in $p; K^0$ to obtain $p(A; K)$ and $K^0(A; K)$. The method used in this paper is dampened fixed-point iteration in the pair $p; K^0$.

{ For each value of $p; K^0$, evaluate on the discretized grid for z and some spline projection grid for k the following equations:

$$V_{(s+1)}^A(z; k; A; K) = \max_{k^0; n} \left(p z A k^{-n} \frac{k^0 + (1 - \beta)k}{\beta} + E_{z^0; A^0} V_{(s)}(z^0; k^0; A^0; K^0) \right)$$

$$V_{(s+1)}^{NA}(z; k; A; K) = \max_n \left(p z A k^{-n} + V \right) \text{down the cross-sectional density in}$$

Compute the value of all reference moments for the next period + 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 + 1, along with the aggregate capital state, compute the coefficients of the cross-sectional distribution associated with period + 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

First, update the current vector of forecast rule coefficients $\hat{\alpha}_{(s)}$ by estimating $(\hat{\alpha}_p(A); \hat{\alpha}_p(A); \hat{\alpha}_K(A); \hat{\alpha}_K(A))$ with OLS on the explicit aggregation dataset, segmented by discretized value of A .

Then, define a vector of n_A bias correction terms

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

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

and adjust the new forecast rule coefficients' constant terms with

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

$$\hat{\alpha}_K(A) - \hat{\alpha}_K(A) x_K^{\text{Bias}}(A):$$

Then, check the estimated coefficients against the old coefficients $\hat{\alpha}_{(s)}$. 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 $\hat{\alpha}_{(s+1)}$ based on rule ξ) and the newly estimated system.

Note that the differencing of $x^{\text{Bias}}(A)$ 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 $x^{\text{Bias}}(A)$ requires that the estimated forecast system be able to exactly reproduce as a fixed point the steady-state prices and aggregate capital stocks $p^{\text{SS}}(A)$ and $K^{\text{SS}}(A)$, 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 $z_{ij} = P(z_{t+1} = z_j | z_t = z_i)$ following Tauchen (1986). Then, we establish a grid of n_k idiosyncratic capital stock nodes k_i

$$f(A_{t-1}; z_i; k_j; w_{t-1})$$

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 z .

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^0 , KS | K^0

To obtain an average percentage innovation in aggregate productivity which equals $\bar{\Delta}_A$ exactly, we choose the shock thresholds to solve

$$\bar{\Delta}_A = \sum_{k=1}^K \bar{\Delta}_{Ak} (\log(A_{n_A}) - \log(A_k));$$

where $\bar{\Delta}_A$ is the ergodic distribution of the discretized aggregate productivity process Δ_A .

As noted in Appendix B, to compute the impulse responses plotted in the main text, we set $\bar{\Delta}_{IRF} = 50$, $T_{shock} = 25$, and $N = 2000$, and we hold exogenous draws $\epsilon_{it}; s_i$ 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 $X_t = AX_{t-1} + B\epsilon_{At}$, where X_t is the endogenous vector defined in Appendix A and ϵ_{At} is a continuous shock to aggregate productivity $\epsilon_{At} \sim N(0, \sigma_{At}^2)$. As a notation-based




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 $(p; K^0)$, 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 $X_t - X^{SS} = A(X_{t-1} - X^{SS}) + B(A_t - A^{SS})$ of the model trivially yields impulse responses $A^{-1}B$ representing the economy's local response to the aggregate productivity shock 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.

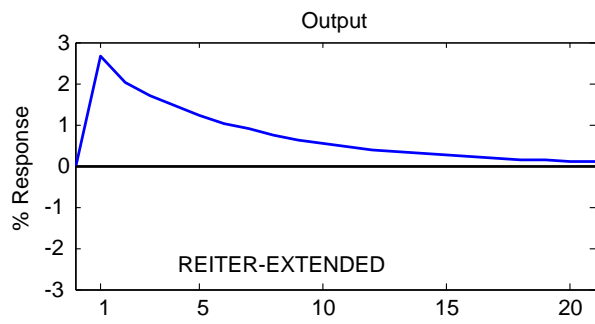


Figure D1: Impulse Response, Productivity Shock in Extended Model

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 $X_t - X^{SS} = A(X_{t-1} - X^{SS}) + B_t$ 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 $X_t - X^{SS} = A(X_{t-1} - X^{SS}) + B_t$ of the model trivially yields impulse responses $A^t B$ 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.