

Görtz, Christoph; Mirza, Afrasiab

Working Paper

On the Applicability of Global Approximation Methods for Models with Jump Discontinuities in Policy Functions

CESifo Working Paper, No. 4837

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Görtz, Christoph; Mirza, Afrasiab (2014) : On the Applicability of Global Approximation Methods for Models with Jump Discontinuities in Policy Functions, CESifo Working Paper, No. 4837, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/102199>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Working Papers

www.cesifo.org/wp

On the Applicability of Global Approximation Methods for Models with Jump Discontinuities in Policy Functions

Christoph Görtz
Afrasiab Mirza

CESIFO WORKING PAPER NO. 4837
CATEGORY 6: FISCAL POLICY, MACROECONOMICS AND GROWTH
JUNE 2014

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

On the Applicability of Global Approximation Methods for Models with Jump Discontinuities in Policy Functions

Abstract

We show that the standard Value Function Iteration (VFI) algorithm has difficulties approximating models with jump discontinuities in policy functions. We find that VFI fails to accurately identify the location and size of jump discontinuities while other methods - such as the Endogenous Grid Method (EGM) and a Finite Element Method (FEM) - are much better at approximating this class of models. We illustrate differences across methods using a standard plant-level investment model with both variable and fixed capital adjustment costs. We find that the policy functions generated by VFI are quite different from those generated by EGM and FEM. Importantly, these differences are economically significant: for our baseline parameterization VFI generates investment spikes that are 5-8% larger in comparison to the other two methods. The choice between EGM and FEM depends on the context. While EGM is faster than FEM, it is much more difficult to implement. For larger models, the modifications necessary to apply EGM can lead to high code complexity. On the other hand, FEM can accommodate larger models with minimal implementation differences and its high scalability can reduce computation time significantly.

JEL-Code: C630, C680, E370.

Keywords: dynamic equilibrium economies, non-convex capital adjustment costs, computational methods, nonlinear solution methods.

Christoph Görtz
University of Birmingham
Department of Economics
J.G. Smith Building, Edgbaston
UK – Birmingham, B15 2TT
c.g.gortz@bham.ac.uk

Afrasiab Mirza
University of Birmingham
Department of Economics
J.G. Smith Building, Edgbaston
UK – Birmingham, B15 2TT
a.mirza@bham.ac.uk

This version: May 2014

We thank Christian Bayer, Andrew Clausen, Wouter den Haan, Giulio Fella, John Fender, Jesus Fernandez-Villaverde, Morten Ravn, Carlo Strub, Håkon Tretvoll, John Tsoukalas and participants at the Society of Computational Economics 2013 Conference and the Royal Economic Society Annual Meeting 2014 for useful comments and suggestions. All remaining errors are our own.

1 Introduction

We examine differences in the answers produced by global approximation methods for solving dynamic economies where agents face non-concave problems (i.e. non-convex choice sets). Non-concave problems can result from the inclusion of fixed adjustment costs that are empirically relevant in many circumstances.¹ For such problems, agents make decisions by comparing the option values associated with different adjustments. In the presence of fixed adjustment costs, a kink arises in the value function at the intersection of these option values and implies a jump discontinuity in the policy function. While differences across approximation methods have been extensively studied for dynamic economies where policy functions are continuous (e.g. McGrattan (1996), Santos (2000), Aruoba et al. (2006), Santos and Peralta-Alva (2012)), the literature provides little guidance about applicability of computational methods when policy functions exhibit jump discontinuities.

We document that the exact intersection of the option values — and thereby the jump discontinuity in the policy function — is difficult to determine using Value Function Iteration (VFI). The algorithm's restriction to a finite grid on state and control variables limits VFI to approximating the option values as step functions. This results in multiple intersections of these values and subsequently leads to an imprecise determination of the jump discontinuity. Sufficient mitigation of this problem requires very fine grids that are infeasible in many applications due to the curse of dimensionality.

To our knowledge this problem has not been documented in the literature. We explore its implications and show that two alternative global approximation methods are more suitable for solving models with jump discontinuities in policy functions. Specifically, we solve the model using a Finite Element Method (FEM), and an adaptation of the

¹The relevance of fixed adjustment cost is highlighted for example in studies of investment (e.g. Caballero et al. (1995), Doms and Dunne (1998), Power (1998), Cooper et al. (1999), Nilsen and Schiantarelli (2003) and Cooper and Haltiwanger (2006), Whited (2006), Bayer (2006), Khan and Thomas (2008), Bloom (2009), Wang and Wen (2012)), consumer-durables choice (e.g. Jose Luengo-Prado (2006), Bajari et al. (2013)), portfolio choice models with transaction costs and asset prices (e.g. Vayanos (1998)), costly technology adoption (e.g. Khan and Ravikumar (2002)) and optimal dynamic capital structure choice (e.g. Hennessy and Whited (2005)).

Endogenous Grid Method (EGM).² These methods overcome the limitations of VFI as they both approximate the option values using piece-wise linear functions — effectively approximating these values using an infinite set of points — that lead to a single crossing and therefore a unique determination of the jump discontinuity in the policy function.

We illustrate differences across approximation methods for non-concave problems using a partial equilibrium model of a plant where investment is subject to both variable and fixed capital adjustment costs and investment is irreversible. This model is well established in the literature and nests the standard Q-theory model as a special case. In the presence of fixed costs the plant determines its investment strategy each period by comparing the option value of remaining inactive (not investing) with the option value of becoming active (investing). The optimal investment strategy follows an (S, s) adjustment process whereby the plant does not make *any* investment until capital depreciates below a threshold level at which point the plant makes a *substantial* investment to rebuild its capital stock. The threshold is determined by the intersection of the plant's option values. To correctly capture the dynamics of investment it is crucial to determine this threshold accurately.

While we find that the approximation of the threshold in the policy function for investment is nearly identical for EGM and FEM, it is substantially different for VFI. Through simulations we show that this difference is economically significant by comparing key moments. For example, in our baseline setting VFI implies the size of the investment spike is 8% (5%) larger than when approximating the model with EGM (FEM). VFI also underestimates the average level of capital by approximately 1.4% relative to both EGM and FEM. Differences between VFI and the other methods persist even for very fine capital grids and are robust also for alternative parameterizations.

While EGM and FEM can overcome the limitations of VFI, there are tradeoffs when choosing between the two methods to solve problems with jump discontinuities in policy functions. While EGM is by far the fastest of the three solution methods, it is also the

²Given that we consider non-concave problems, we focus on piece-wise linear approximations and do not implement methods that involve higher order polynomial approximations.

most complex to implement as it requires a number of adaptations to be applicable to our model. The original EGM algorithm introduced by Carroll (2006) is limited to models with at most one control and one endogenous state variable. The literature has proposed numerous extensions to accommodate more complex models as the applicability of EGM is context dependent.³ Importantly, Fella (2014) shows how to extend EGM to non-smooth and non-concave problems such as ours. We adapt the algorithm – exemplified by Fella (2014) for a consumption model that involves fixed adjustment costs for durable goods – to our model of a plant with fixed capital adjustment costs.

FEM is far less complex to implement for our model than EGM.⁴ FEM is a general purpose method that requires only minimal changes to handle more complex models. It has been introduced into economics by McGrattan (1996) and Aruoba et al. (2006) but has since been used only for a limited number of applications. A drawback of this method is that it is even more expensive in terms of computation time than VFI. But the performance of FEM can be improved dramatically by making use of parallel computing. As the computation time of FEM increases only linearly with the number of grid points, it is particularly useful for approximating larger models with jump discontinuities in policy functions.

Overall, our results suggest that the optimal choice of numerical method for problems with jump discontinuities in policy functions depends on the context: EGM is applicable in a limited class of smaller models where one can benefit from the algorithms’ speed at cost of tolerable implementation complexity. FEM is best-suited to overcome the drawbacks of VFI for any other class of models and through its scalability it is especially applicable to models with a large number of state and control variables.

³These extensions often combine EGM with VFI. Barillas and Fernandez-Villaverde (2007) show how to introduce additional control variables; Hintermaier and Koeniger (2010) demonstrate how to introduce additional endogenous state variables in a durable goods model and Ludwig and Schön (2013) show how to accommodate additional endogenous state variables in a human capital model.

⁴Our FEM code approximates the value function using piece-wise linear functions with weights updated via iteration on the Bellman operator rather than minimization of the Galerkin residual as in McGrattan (1996) and Aruoba et al. (2006). The latter approach has been exemplified for smooth problems and is faster but more complex to implement. However, for our problem with jump discontinuities in the policy function we find that this approach is problematic as results are highly dependent on the algorithm’s start values.

The rest of the paper is organized as follows. The next section presents the model we use to illustrate our results. We then provide descriptions of the three approximation methods we use to solve the model. Section 4 discusses the parameterization which we use for the analysis of solutions generated by the approximation methods in Section 5. The final section concludes.

2 The Model

We consider a general class of models where in every period the agent makes both a continuous and discrete choice (c', d') based on the state variable (c, d) consisting of previous period's choices. The set of possible states is denoted by Ω . The agent's choice set is constrained as follows:

$$(c, c', d, d') \in \Gamma.^5$$

Importantly, this specification of the constraints includes the case where c or d are subject to non-convex adjustment costs. The agent solves the following dynamic programming problem:

$$V(c, d, A) = \sup_{(c', d') \in \Gamma(c, \cdot; d, \cdot; A)} u(c, c'; d, d'; A) + \beta \sum_{A' \in \mathcal{A}} \pi(A'|A) V(c', d', A')$$

where \mathcal{A} is the set of all possible shock realizations $A \in \mathcal{A}$, π is the corresponding transition matrix, the domain of V is $\Omega \times \mathcal{A}$, the per-period utility function of the agent is u , and the discount-factor is β . We assume that $u(\cdot, c'; d, d'; A)$ and $u(c, \cdot; d, d'; A)$ are differentiable on $\text{int}(\Gamma(\cdot, c'; d, d'; A))$ and $\text{int}(\Gamma(c, \cdot; d, d'; A))$, respectively. Importantly, the value function V is non-concave in the presence of non-convex adjustment costs to c or d . As a result, the agent compares the *option values* associated with choices of c' and d' . A kink in the value function arises at the point of indifference between these options and implies a jump discontinuity in the policy functions (see e.g. Clausen and

⁵We define particular subsets of Γ as follows: $\Gamma(c, \cdot; d, \cdot) = \{(c', d') : (c, c', d, d') \in \Gamma\}$, $\Gamma(c, \cdot; d, d') = \{c' : (c, c', d, d') \in \Gamma\}$, $\Gamma(\cdot, c'; d, d') = \{c : (c, c', d, d') \in \Gamma\}$.

Strub (2012)).

The general framework described above nests a number of important models with jump discontinuities used in the literature. This includes models with costly technology adoption (e.g. Khan and Ravikumar (2002)), durable consumption goods (e.g. Bajari et al. (2013)), and firm-level investment (e.g. Cooper and Haltiwanger (2006), Wang and Wen (2012)). We illustrate the applicability of different approximation methods using a model that captures key elements of models in the firm-investment literature. Specifically, we employ a partial equilibrium model of a plant in which capital adjustment is subject to both fixed and variable adjustment costs that we describe in detail below.

The plant produces output Y_t via the production function

$$Y_t = A_t K_t^\alpha, \quad 0 < \alpha < 1, \quad (1)$$

where K_t denotes capital and productivity A_t evolves according to the AR(1) process

$$\log A_{t+1} = \rho \log A_t + \varepsilon_t, \quad 0 < \rho < 1, \quad (2)$$

where $\varepsilon_t \sim N(0, \sigma_\varepsilon)$. The plant's capital stock evolves according to the law of motion

$$K_{t+1} = (1 - \delta)K_t + I_t, \quad 0 < \delta < 1, \quad (3)$$

where I_t is investment. When the plant chooses to invest, it has to pay a price p_I per investment good as well as adjustment costs $C(K_t, I_t)$. These are given by

$$C(K_t, I_t) = \frac{\gamma}{2} \left(\frac{I_t}{K_t} \right)^2 K_t + F K_t, \quad \gamma \geq 0, F \geq 0.$$

where the first term denotes convex variable investment adjustment costs and the latter term the non-convex fixed costs. These are proportional to the capital stock to eliminate any size effects. Further, investment is completely irreversible as we assume for simplicity

that capital cannot be resold on a secondary market. Formally, we impose $I_t \geq 0 \forall t$.

Note that the model includes the standard Q-theory model of investment, in which the value function is proportional to the stock of capital, as a special case.⁶ The plant's problem consists of choosing a sequence of investments $\{I_t\}_{t=0}^{\infty}$ to maximize discounted life-time profits:

$$V(K, A) = \max_{\{I_t \geq 0\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t \left[AK_t^\alpha - p_I I_t - F \mathcal{I}_{(I_t > 0)} K_t - \frac{\gamma K_t}{2} \left(\frac{I_t}{K_t} \right)^2 \right] \quad (4)$$

subject to equations (1) and (2) and the constraint $I_t \geq 0$, given an initial level of capital, K_0 , and productivity, A_0 . $\mathcal{I}_{(I_t > 0)}$ is an indicator function that equals 1 if investment is positive and zero otherwise. The constraint $I_t \geq 0$ may bind in equilibrium when capital is too costly relative to the increase in future profits from additional plant capacity.

Dropping time indices, we can write the problem recursively as:

$$V(K, A) = \max\{V^a(K, A), V^i(K, A)\},$$

where $V^a(K, A)$ and $V^i(K, A)$ are the values to the plant to exercising its option to either remain inactive (i.e. not invest) or active (invest). We can characterize the value of the option to invest as follows:

$$V^a(K, A) = \max_{I > 0} \left[AK^\alpha - p_I I - FK - \frac{\gamma K}{2} \left(\frac{I}{K} \right)^2 + \beta E_{A'|A} V(K(1 - \delta) + I, A') \right], \quad (5)$$

where $K' > K(1 - \delta)$. Similarly, we can characterize the value of the option to not invest as

$$V^i(K, A) = AK^\alpha + \beta E_{A'|A} V(K(1 - \delta), A'),$$

⁶This is the case in our setup if the profits are proportional to the capital stock which is guaranteed if the plant's profit function is homogeneous of degree one ($\alpha = 1$) and the adjustment cost function is convex ($F = 0$).

where $K' = K(1 - \delta)$ because $I = 0$. In each period, the plant computes the value of these two options and chooses its investment strategy accordingly.

In the presence of fixed costs ($F > 0$), it is optimal for the plant to follow an (S, s) adjustment strategy for investment. In other words, investment will be zero for all periods in which the capital stock exceeds a threshold level $\hat{K}(A)$. When capital has depreciated below the threshold level the plant will make a substantial investment to build capital up again. Hence, there is a jump discontinuity in the policy function for investment at the threshold $\hat{K}(A)$.

The intuition behind the plant's choice is the following: for capital stock levels below $\hat{K}(A)$ the value of investing will be higher than the value of not investing: $V^a(K, A) > V^i(K, A)$. That is, the benefit from having a larger capital stock in the future exceeds the costs of investing today. For capital stocks above $\hat{K}(A)$ the opposite is true: the benefit from having an even larger capital stock tomorrow diminishes (due to decreasing returns to scale in production) and is smaller than the costs of investment. In this case $V^a(K, A) < V^i(K, A)$ and the plant will not invest.

The convexity of the adjustment costs in investment and the monotonicity of the value function $V(K, A)$ in capital entail that $V^i(K, A)$ and $V^a(K, A)$ cross exactly at one point for a given productivity, namely at $\hat{K}(A)$.⁷ This implies that the value function $V(K, A)$ exhibits a kink at $\hat{K}(A)$ and is globally non-concave which is illustrated in Figure 1 (see also Clausen and Strub (2012)). As there is no closed form solution for the value function, we need to approximate the solution numerically.

3 Solution Methods

We solve the model using three global approximation methods and provide brief descriptions of these in this section. Additional details on the implementation of the methods are provided in the Appendix. For all solution methods we approximate the AR(1)

⁷Monotonicity of the value function in capital follows from the monotonicity of the Bellman operator.

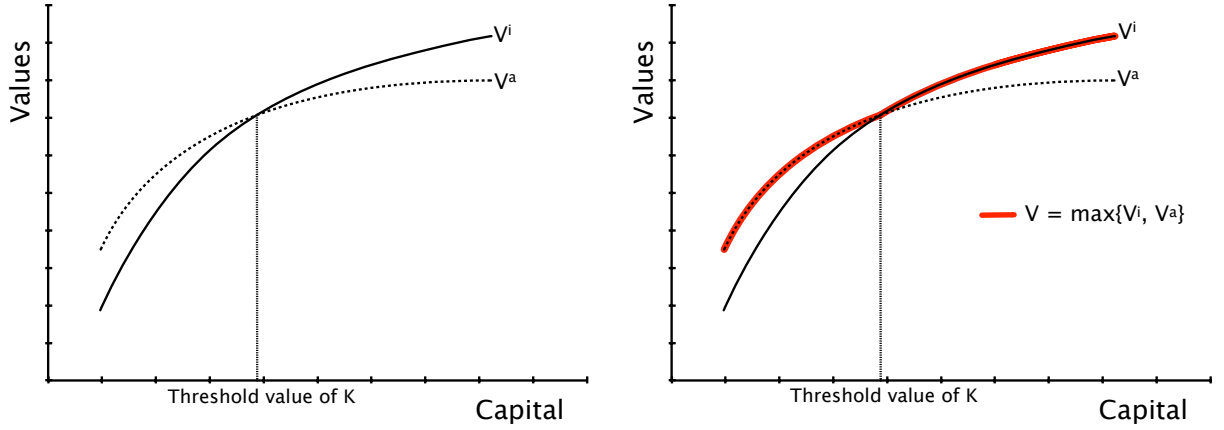


Figure 1: The left diagram shows the option values of active (V^a , dotted line) and inactive (V^i , solid line) investment for a given level of productivity. The right diagram shows the value function for the plant's problem (V , in red) that results from choosing the maximum value of the options V^a and V^i for each level of capital.

process for productivity using a discrete Markov chain as in Rouwenhorst (1995).

3.1 Value Function Iteration

We solve the model first using Value Function Iteration (VFI). To implement VFI, we first discretize the state space by assigning a grid to the state and control variables. We then iterate on the Bellman operator and compute an approximation of the value function over the grid, i.e. for every grid point combination of the endogenous and exogenous state variable. The Bellman operator in our case is:

$$TV(K, A) = \max \left\{ AK^\alpha + \beta E_{A'|A} V(A', (1 - \delta)K), \max_{K' \geq (1 - \delta)K} AK^\alpha - pI - FK - \frac{\gamma K}{2} \left(\frac{I}{K} \right)^2 + \beta E_{A'|A} V(K', A') \right\} \quad (6)$$

$$K' = (1 - \delta)K + I$$

$$\log A' = \log A + \varepsilon$$

Note that this method requires the explicit computation of $V^i(K, A)$ and $V^a(K, A)$ at each grid point combination of the state variables. The value function is then updated

according to $V(K, A) = \max\{V^i(K, A), V^a(K, A)\}$. The repeated application of the max operator over the grid leads to an exponential increase in computing time with finer grids.

3.2 Finite Element Method

The main idea behind FEM is to approximate a function of interest using a number of much simpler basis functions. Each of these basis functions are typically non-zero only on a small part of the state space, or equivalently on a small number of elements. This sparsity allows a large number of elements to be handled and the algorithm is well suited for parallel computing.

Our FEM algorithm approximates the value function using a piece-wise linear approximation. We partition the state space into rectangles of the form $[K_i, K_{i+1}] \times [A_j, A_{j+1}]$ using an equally spaced grid for capital where K_i (A_j) is the i^{th} (j^{th}) grid point for capital (the shock). We then approximate the value function over the state space using a piece-wise linear function over the grid points of the partition. Given an initial guess for the value function $V^0(K, A)$ at each grid point in the state space, we approximate the value function as $\hat{V}(K, A) = \sum_{ij} \hat{V}_{ij}(K, A)$ where

$$\hat{V}_{ij}(K, A) = \begin{cases} V_{ij}^0(K, A) + \frac{V_{i+1j}^0 - V_{ij}^0}{K_{i+1} - K_i} (K - K_i) & \text{if } K \in [K_i, K_{i+1}] \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

so that we effectively use a piece-wise linear approximation for each value of productivity. We then apply the Bellman operator (6) using $\hat{V}(K', A')$ as our guess for tomorrow's value function and update our initial guess $V^0(K, A)$ on the grid points. Finally, we iterate to convergence on $\hat{V}(K, A)$.

The key difference between FEM and VFI is that with FEM tomorrow's value function can be evaluated at any point in the state space. Crucially, this implies that the optimal choice of tomorrow's capital is not restricted to be on the exogenous grid $[K_1, \dots, K_n]$. Therefore, the optimization step in the Bellman operator can be carried out using a stan-

standard constrained optimization routine that enforces the irreversible investment constraint $I \geq 0$. Hence, FEM permits an additional degree of freedom above VFI but it comes at a cost as we are forced to employ the computationally expensive constrained optimization routine repeatedly. Note that the policy function generated by this procedure is also a piece-wise linear function akin to (7). Our algorithm of FEM is no more difficult to implement than VFI given that we do not rely on Galerkin weighting and given the built-in optimization routines in MATLAB.⁸

3.3 Endogenous Grid Method

EGM as introduced by Carroll (2006) suggests assigning an exogenous grid over the control variable K' rather than the state variable K . Then, the following first-order condition allows us to determine an endogenous grid over K , given the exogenous grid K' and the derivative of the value function with respect to K' , $V_{K'}(K', A')$,⁹

$$p_I + \gamma \frac{K' - (1 - \delta)K}{K} = \beta E_{A'|A} V_{K'}(K', A'). \quad (8)$$

Interpolating over the endogenous state space, and evaluating it at the exogenous grid points, we can obtain a set of optimal control and state pairs that can then be used to approximate the value function.

Crucially, this procedure requires a unique solution to the first-order condition when solving for the endogenous grid over K . As shown in Figure 1, in our case fixed costs introduce kink(s) in the value function resulting in jump discontinuities in the (otherwise smooth and decreasing) slope of the value function, $V_{K'}(K', A')$. As a result, the first-order condition (8) does not imply a unique endogenous grid over K .¹⁰ Therefore, EGM as introduced by Carroll (2006) is not directly applicable.

⁸For our case of jump discontinuities in the policy function we find that Galerkin weighting is not suitable as it leads to results that are highly dependent on start values for the algorithm.

⁹The derivation of the first-order condition is shown in the Appendix.

¹⁰Clausen and Strub (2012) show that at the optimum the first-order condition holds and the envelope condition is valid.

Instead, we employ a modification of EGM proposed by Fella (2014) and implement the following steps for our case with fixed capital adjustment costs:

1. We begin by assigning an (exogenous) grid on K' and an initial guess for $V_{K'}(K', A')$.
2. We then proceed to generate an endogenous grid for K using the first-order condition (8).¹¹
3. We split our endogenous state space into two regions: one where the value function is concave ($V_{K'}(K', A')$ is smooth) and another where the value function is not concave ($V_{K'}(K', A')$ exhibits jump discontinuities).
 - (a) We apply the algorithm by Carroll (2006) in the region where the value function is concave.
 - (b) We apply VFI in the region where the value function is non-concave to identify and retain only the global optimum.
4. We then proceed to interpolate over the endogenous state space and construct optimal (K, K') pairs in both the active and inactive cases.
5. We use these pairs to construct an approximation of the values to being active and inactive and thereby the overall value function.
6. We use the slope of this value function to construct the endogenous grid as in step 2. Steps 2-5 are repeated until the value function is deemed to have converged.

The computationally most demanding task of this algorithm is the interpolation step which is far less expensive than the maximization/optimization steps in VFI or FEM. Note however that the applicability of EGM is context dependent, for example it cannot necessarily accommodate additional variables. The reason is that EGM's applicability rests on finding a unique solution of the first-order conditions for the endogenous grids. The literature shows how to accommodate additional variables for specific classes of

¹¹Note that this implies that the endogenous grid changes in every iteration.

models by often combining EGM with VFI steps (see for example Barillas and Fernandez-Villaverde (2007), Hintermaier and Koeniger (2010), Fella (2014), and Ludwig and Schön (2013)).

4 Parameterization

Our choice of the model parameters is based on estimates by Cooper and Haltiwanger (2006). Using annual plant level data of the Longitudinal Research Database, they estimate the above model with convex and non-convex capital adjustment costs and find that a combination of these fits the data well.¹² The estimates of Cooper and Haltiwanger (2006), that we use to calibrate the model, are summarized in Table 1. Importantly, they find evidence for substantial fixed adjustment costs of roughly 4% of the average plant-level capital stock.

We approximate the stochastic process for productivity by a four-state Markov chain. Assigning the values $[e^{-0.1116}, e^{-0.0372}, e^{0.0372}, e^{0.1116}]$ for the states of the shock implies a standard deviation $\sigma_\epsilon = 0.03$ consistent with the estimate by Cooper and Haltiwanger (2006). We match their estimate for the shock's persistence, 0.885, by using the method proposed by Rouwenhorst (1995).

Conditional on our choice of model parameters, we approximate the value function for all solution methods over the same state space for capital. The state space is chosen so that capital does not hit any boundaries during our simulations. Convergence is evaluated by considering the largest absolute distance between corresponding points of the value function of two consecutive iterations. If this absolute distance falls below 10^{-4} the algorithm is deemed to have converged.

As a baseline scenario we use 420 (VFI) and 60 (EGM and FEM) capital grid points,

¹²For the sake of simplicity of exposition we do not include the possibility of selling capital considered by Cooper and Haltiwanger (2006). Selling plant's capital stock at a price smaller than p_I would introduce an additional kink in the value function. The solution methods can be adjusted to accommodate the additional choice, but as our findings can be generalized to these additional kinks we assume irreversibility of capital for ease of exposition.

which is representative for many practical applications. This choice generates comparable average log-absolute Euler equation errors across methods (see Table 2).¹³ As conventional in studies which consider the performance of different approximation methods we use an equally spaced grid for capital.

Table 1: Model Parameters (based on Cooper and Haltiwanger (2006))

β	0.95	discount factor
δ	0.069	capital depreciation rate
p_I	1	price to buy capital
α	0.592	returns of capital
ρ	0.885	persistence of plant specific shock
σ_ϵ	0.03	standard deviation of plant specific shock
γ	0.049	convex adjustment costs
F	0.039	fixed adjustment costs

5 Results

As noted in Section 2, theory predicts that the policy function exhibits a jump discontinuity at the threshold separating the active (positive investment) and inactive (no investment) regions. Theory also predicts additional jump discontinuities in the policy function in the active region due to the interaction between fixed and convex variable adjustment costs (see e.g. Clausen and Strub (2012)). The variable costs penalize the plant for making large adjustments while fixed costs penalize the plant for making small and frequent investments. The result is that the active region of the policy function consists of concave parts that are separated by jump discontinuities.

Figure 2 shows the policy functions for tomorrow’s capital generated by VFI, EGM and FEM at a common level of accuracy as measured by average log-absolute Euler equation errors (baseline scenario). This figure highlights that while EGM and FEM deliver nearly identical policy functions, these differ substantially from the one approximated by VFI

¹³As the Euler equation is a necessary but not a sufficient condition in our setup, Euler equation error statistics are calculated across policy functions for all shocks (using equation (A.1)) in the area of the state space in which all approximation methods imply positive investment.

in two important dimensions. First, the threshold separating the active and inactive regions is not uniquely determined by VFI. Second, the concave parts and the jump discontinuities in the active region are inaccurately approximated.

To clearly see the problems that arise in the determination of the threshold with VFI, we show in the top panel of Figure 3 the values to the plant of being active and inactive for increasingly finer capital grids. The intersection of these values determines the capital threshold, $\hat{K}(A)$, below which the plant is active and above which the plant is inactive. While theory predicts a single intersection of these functions, however VFI generates multiple intersections as a result of approximating these values using step functions. The reason for these steps is that only a finite set of points can be used to approximate the values of being active and inactive because VFI limits the choices for both the values of the endogenous state and the control variable to a fixed grid.¹⁴

The inaccurate determination of the threshold shows up directly in the corresponding policy functions for tomorrow's capital which are shown in the bottom panel of Figure 3. For instance, for the capital grid consisting of 200 points, VFI predicts the threshold to be between 30.5-36.9 capital units. As we make the grid finer, the prediction clearly improves: however, while the baseline capital grid of 420 points implies a threshold between 33.28-34.96, even with a grid of 2050 points the threshold cannot be approximated to within one decimal place. We can only say that it lies in range of 32.94-34.96. As one can see from Figure 3, this problem can only be mitigated through very fine grids that are typically infeasible in most applications due to the curse of dimensionality.

For similar reasons, VFI is also unable to correctly approximate the jump discontinuities and concave parts of the policy function in the active region. However, while the threshold is crucial for the dynamics of the model, the poor approximation of the active region is only of larger importance when the persistence of the technology shock is low. The problems of VFI in approximating policy functions with jump discontinuities can be

¹⁴Such approximations are particularly prone to error when the slope of the underlying function is steep. In our problem, the slope of the value to being inactive is much larger than the slope of the value to being active. Hence, as shown in Figure 3, the approximation of the value of being inactive is much worse than the approximation of the value of being active.

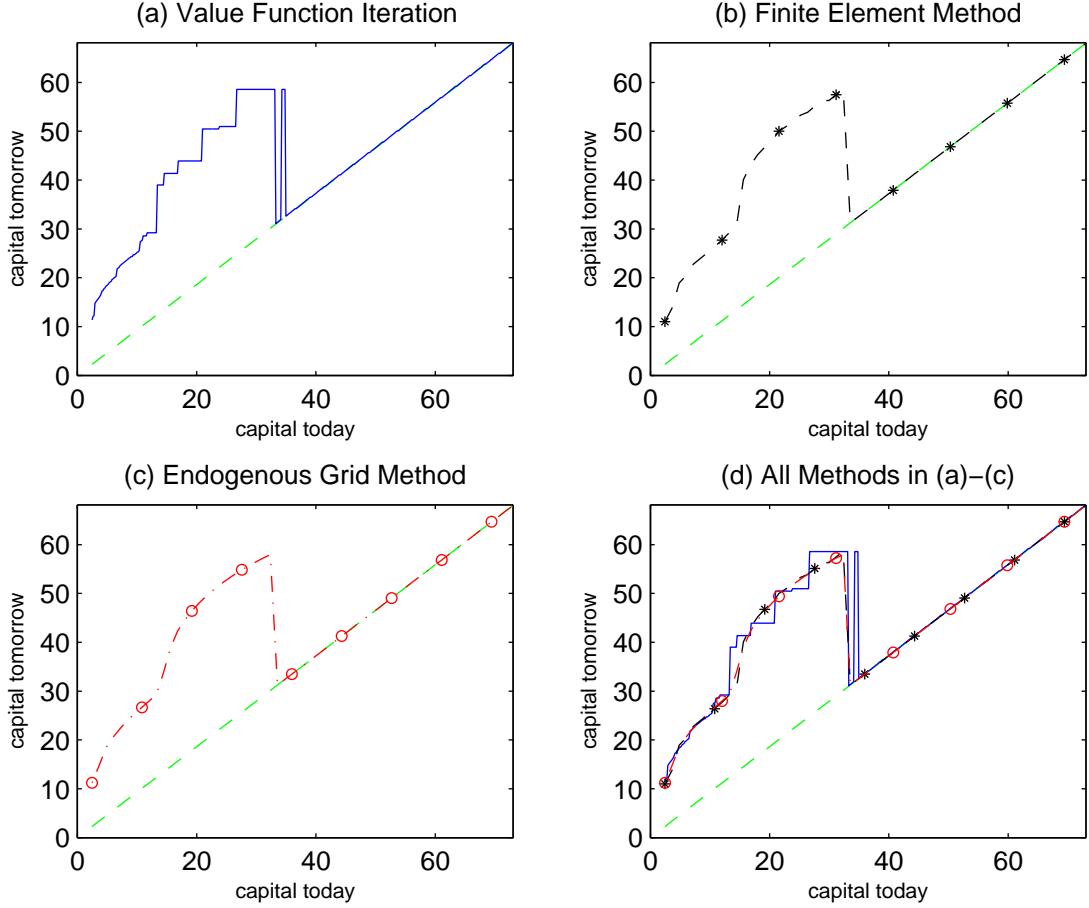


Figure 2: Policy Function for capital (given $A = e^{0.0372}$) implied by different approximation methods. The green dashed line indicates the no investment decision $(1 - \delta)K$. The baseline capital grid (VFI: 420, FEM: 60, EGM: 60) generates comparable average log-absolute Euler equation errors across methods.

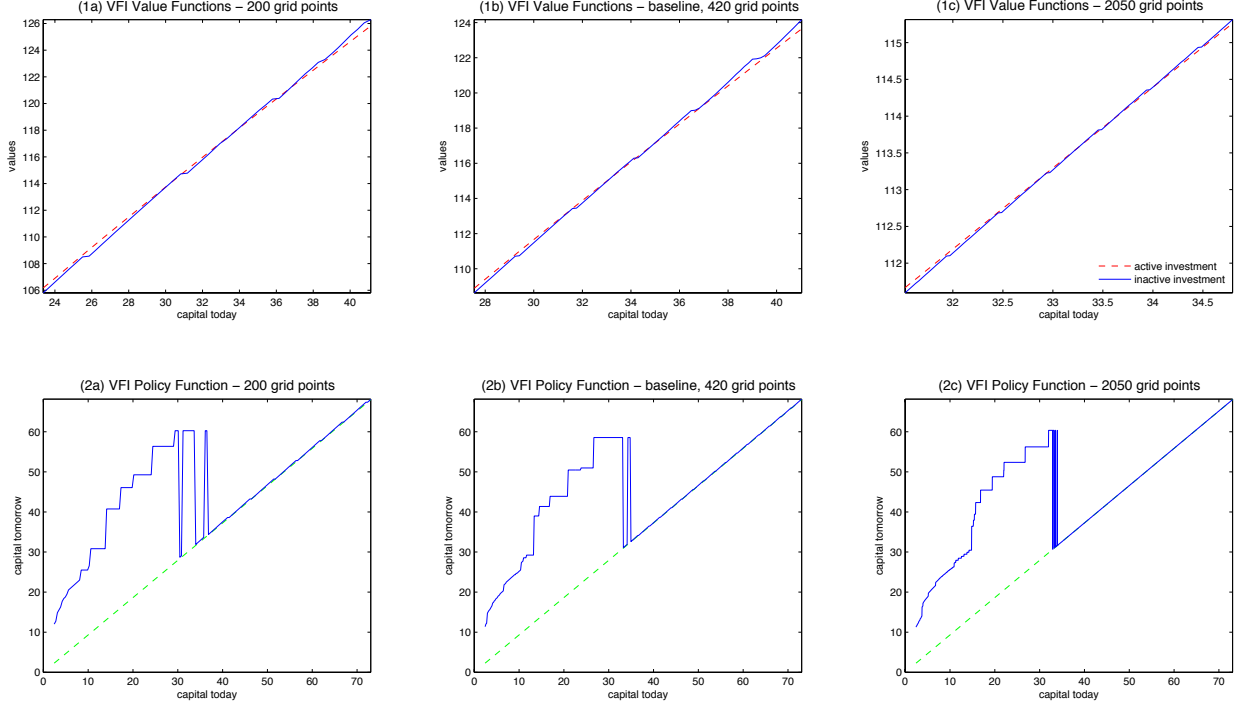


Figure 3: Approximation using Value Function Iteration. Top panel: Option values to the plant of being active (red dashed) and inactive (blue solid) for different capital grids (zoomed in to show multiple intersections). Bottom panel: Policy functions for capital for different capital grids. In all sub-plots the shock is fixed at $A = e^{0.0372}$.

overcome by using FEM or EGM. Figure 2 shows that FEM and EGM are able to very precisely determine the threshold value of capital. Both methods yield an estimate of 33.55 capital units for the threshold (while VFI estimates it be in the range of 33.28-34.96 capital units). Also, the approximation of the active (and inactive) regions by EGM and FEM are almost identical.

To understand the economic relevance of differences between approximated policy functions we compare key moments based on simulations. Table 2 shows the spike rate of investment (i.e. the share of periods in which investment exceeds 20% of the capital stock), the size of the investment spikes and the mean of capital in a deterministic setting.¹⁵ This experiment allows us to focus on the impact of the position of the threshold $\hat{K}(A)$ and the size of the jump discontinuity as these directly impact both the spike rate

¹⁵We fix productivity at $A = e^{-0.116}$. All moments are calculated from the dynamics implied by this shock value to avoid that differences in moments average out across shock values. However, other values of A generate differences across methods comparable to the ones shown in Table 2.

and the size of the investment spike. These two statistics are often used to calibrate models with (S, s) adjustment of capital to the data. The mean of capital is a popular measure for firm size.

For our baseline scenario (see Figure 2 for the policy functions), we report that moments differ substantially between VFI and the other two approximation methods. Table 2 shows that the spike rate delivered by VFI (9%) is substantially lower than approximations with EGM and FEM (both 10%). Further, VFI overestimates the size of the investment spikes (23.11) by 8% relative to EGM (21.47) and by 5% relative to FEM (21.92) and underestimates the mean of capital by 1.4% relative to the other two methods. Hence, EGM and FEM determine a very similar position and size of the threshold, while VFI implies economically very different values.

Table 2 also shows that even for a small number of grid points, both EGM and FEM very quickly converge to nearly identical moments. These moments change only very little for even much finer grids, reflecting the precise determination of the threshold. On the other hand, VFI requires very fine grids to produce a comparable spike rate. However, this measure is only of limited value for comparison as its variation is bounded by the number of simulation periods. Even for very fine grids (VFI: 2050, EGM: 500, FEM: 2000) that deliver very small Euler equation errors (absolute average log-error approximately -7.20), the investment spike size and the mean of capital produced by VFI still differ from the other two methods. It is important to understand that in many practical applications (e.g. with more endogenous state variables or more shock grid points) such fine capital grids for VFI are infeasible due to the curse of dimensionality. However, for our purpose these rather accurate solutions offer useful guidance for comparison as the model – and typically non-concave models in the literature – cannot be solved analytically.

While we use a specific model to exemplify the problems of VFI to accurately identify a jump discontinuity in policy functions, note from the exposition above that these problems will be present in any application in which the location of discontinuities are determined by the intersection of option values. The economic significance of the dif-

ferences between VFI and the other two methods clearly depends on specific model and the parameterization. Note that our parameterization is relatively conservative. Differences between VFI and the other two methods would be even more pronounced for other, widely used, parameter values in the literature. For example, our value for the parameter determining the returns of capital, $\alpha = 0.592$, is at the upper bound of used values. Lower values for α emphasize the problems of VFI to identify the threshold as it leads to flatter option values V^i and V^a .¹⁶ Appendix A.5 shows that the problems of VFI to accurately identify jump discontinuities are economically relevant for various parameterizations.

Trade-offs between EGM and FEM The discussion above suggests that both, EGM and FEM can overcome VFI's problems in approximating policy functions with jump discontinuities. However, there are pros and cons between EGM and FEM in terms of computation time and implementation complexity. We ran all programs on an Intel i7-3770 (3.4 GHz) Processor with 4 active cores and 16 GB of memory running Windows 7. As we implemented all methods using Matlab, we can directly compare running time across methods and the number of code lines can serve as a measure of implementation complexity.

Table 2 summarizes the running time for our baseline scenario. EGM solves the model much faster than FEM (and VFI) as it does not rely on expensive root finding operations but instead on repeated linear interpolation which is much less expensive. The performance of FEM can be improved dramatically as the root-finding operations can benefit from parallel computing.¹⁷ For example, FEM is highly scalable and with 4 processing units we can reduce the running time in the baseline case by up to 60%. We also find that computing time for FEM increases nearly linearly with additional capital grid points (see Table 2), while it increases exponentially with VFI and EGM.

While the adapted EGM method is by far the fastest, it is also the most complex

¹⁶Commonly used values are between 0.30 and 0.42, see for example Gomes (2001), Görtz and Tsoukalas (2013) and King et al. (1988).

¹⁷This does not lead to any improvements in speed for EGM.

Table 2: Statistics across different approximation methods

	capital grid points	spike rate (in %)	investment spike size	capital mean	Euler equation error		CPU (seconds)
					average	maximum	
VFI	200	8.40	24.63	32.77	-6.05	-4.69	6.36
	300	9.00	22.32	31.56	-6.43	-4.87	10.58
	350	9.10	24.50	33.95	-6.47	-4.83	12.58
	400	9.10	24.28	33.65	-6.53	-4.87	15.29
baseline:	420	9.00	23.11	32.02	-6.59	-4.88	16.00
	500	11.10	19.69	33.15	-6.63	-4.84	21.66
	600	9.00	23.34	31.86	-6.69	-4.89	26.56
	1000	10.00	22.50	33.28	-6.94	-4.92	54.18
	1500	10.00	21.03	30.92	-7.12	-4.94	99.40
	2050	10.00	22.02	32.23	-7.19	-4.94	168.01
EGM	21	14.30	21.38	34.22	-6.09	-4.79	0.89
	30	11.10	21.58	31.67	-6.29	-4.95	0.96
	40	9.00	21.56	31.58	-6.33	-5.06	1.33
	50	9.10	21.56	31.13	-6.42	-5.08	1.36
baseline:	60	10.00	21.47	32.48	-6.59	-5.07	1.72
	80	10.00	21.48	31.93	-6.66	-5.05	2.38
	100	10.00	21.57	31.25	-6.67	-5.04	3.53
	200	10.00	21.64	31.50	-6.92	-5.04	12.96
	400	10.00	21.63	31.48	-7.13	-5.03	56.51
	500	10.00	21.67	31.35	-7.19	-5.03	77.71
FEM	25	12.50	22.19	33.33	-6.08	-4.68	35.26
	30	11.10	20.99	31.67	-6.37	-5.01	39.71
	40	9.00	21.46	31.58	-6.41	-5.06	48.76
	50	9.10	21.74	31.13	-6.46	-5.06	58.49
baseline:	60	10.00	21.92	32.48	-6.57	-5.04	66.76
	100	10.00	21.52	31.25	-6.66	-5.05	102.14
	150	10.00	21.46	31.30	-6.72	-5.01	143.78
	300	10.00	21.48	31.27	-6.93	-5.01	265.18
	800	10.00	21.68	31.26	-7.13	-5.01	651.83
	2000	10.00	21.66	31.37	-7.20	-5.00	1573.6

VFI: Value Function Iteration, EGM: Endogenous Grid Method, FEM: Finite Element Method. The spike rate is calculated as the share of periods in which investment exceeds 20% of the capital stock. Moments are calculated from 50 simulations over 1050 periods where the first 50 periods are discarded. Euler Equation error statistics are calculated across policy functions for all shocks in the area of the state space in which the Euler equation holds. CPU time for FEM is reported utilising four processing units.

to implement. For our implementation, both VFI and FEM were of comparable complexity, each requiring 100-150 lines of code. On the other hand, EGM was by far the most complex method requiring nearly 300 lines of code. Moreover, additional state and control variables would make this complexity gap between EGM and the other methods even larger. This is because VFI and FEM require minimal changes to handle more complex models whereas EGM requires a number of intricate extensions that are often problem specific (see for example Barillas and Fernandez-Villaverde (2007), Hintermaier and Koeniger (2010), Fella (2014) and Ludwig and Schön (2013)).

Based on these observations, VFI's problems in approximating policy functions with jump discontinuities can be overcome by either EGM or FEM, however choosing one over the other is problem dependent. EGM is best suited to specific cases with a limited number of control and state variables. FEM is better suited to approximate larger and more complex models.

6 Conclusion

We evaluate the applicability of global approximation methods for models with jump discontinuities in policy functions. We find that using Value Function Iteration (VFI) can be problematic as it fails to accurately identify the location and size of jump discontinuities. Other methods such as the Endogenous Grid Method (EGM) or the Finite Element Method (FEM) are far better at accurately approximating this class of models. As a result, policy functions for FEM and EGM are nearly identical but differ considerably from those generated by VFI. Using a standard manufacturing plant model with variable and fixed investment adjustment costs, we find these differences are also economically significant as evidenced through key moments of the model. Our results suggest that the choice of numerical method is context dependent: While EGM is considerably faster than FEM, it is more complicated to implement. For larger models the modifications necessary to apply EGM can lead to high code complexity. On the other hand, the implementation

of FEM even for larger models is relatively straightforward and model alternations can be accommodated with minimal implementation differences.

References

- Aruoba, S. B., Fernandez-Villaverde, J., and Rubio-Ramirez, J. F. (2006). Comparing solution methods for dynamic equilibrium economies. *Journal of Economic Dynamics and Control*, 30(12):2477–2508.
- Bajari, P., Chan, P., Krueger, D., and Miller, D. (2013). A dynamic model of housing demand: Estimation and policy implications. *International Economic Review*, 54(2):409–442.
- Barillas, F. and Fernandez-Villaverde, J. (2007). A generalization of the endogenous grid method. *Journal of Economic Dynamics and Control*, 31(8):2698 – 2712.
- Bayer, C. (2006). Investment dynamics with fixed capital adjustment cost and capital market imperfections. *Journal of Monetary Economics*, 53(8):1909–1947.
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.
- Caballero, R. J., Engel, E. M. R. A., Haltiwanger, J. C., Woodford, M., and Hall, R. E. (1995). Plant-level adjustment and aggregate investment dynamics. *Brookings Papers on Economic Activity*, 1995(2):1–54.
- Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320.
- Clausen, A. and Strub, C. (2012). Envelope theorems for non-smooth and non-concave optimization. *Department of Economics - University of Zurich: ECON - Working Papers*, (062).
- Cooper, R., Haltiwanger, J., and Power, L. (1999). Machine replacement and the business cycle: Lumps and bumps. *American Economic Review*, 89(4):921–946.
- Cooper, R. W. and Haltiwanger, J. (2006). On the nature of capital adjustment costs. *Review of Economic Studies*, 73(3):611–633.

- Doms, M. E. and Dunne, T. (1998). Capital adjustment patterns in manufacturing plants. *Review of Economic Dynamics*, 1(2):409–429.
- Fella, G. (2014). A generalized endogenous grid method for non-smooth and non-concave problems. *Review of Economic Dynamics*, 17(2):329–344.
- Gomes, J. F. (2001). Financing Investment. *American Economic Review*, 91(5):1263–1285.
- Görtz, C. and Tsoukalas, J. D. (2013). Learning, Capital Embodied Technology and Aggregate Fluctuations. *Review of Economic Dynamics*, 16(4):708–723.
- Hennessey, C. and Whited, T. (2005). Debt dynamics. *Journal of Finance*, 60(3):1129–1165.
- Hintermaier, T. and Koeniger, W. (2010). The method of endogenous gridpoints with occasionally binding constraints among endogenous variables. *Journal of Economic Dynamics and Control*, 34(10):2074–2088.
- Jose Luengo-Prado, M. (2006). Durables, nondurables, down payments and consumption excesses. *Journal of Monetary Economics*, 53(7):1509–1539.
- Khan, A. and Ravikumar, B. (2002). Costly technology adoption and capital accumulation. *Review of Economic Dynamics*, 5(2):489 – 502.
- Khan, A. and Thomas, J. K. (2008). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica*, 76(2):395–436.
- King, R. G., Plosser, C. I., and Rebelo, S. T. (1988). Production, growth and business cycles : I. The basic neoclassical model. *Journal of Monetary Economics*, 21(2-3):195–232.
- Ludwig, A. and Schön, M. (2013). Endogenous grids in higher dimensions: Delaunay interpolation and hybrid methods. MEA discussion paper series 13274, Munich Center

for the Economics of Aging (MEA) at the Max Planck Institute for Social Law and Social Policy.

- McGrattan, E. R. (1996). Solving the stochastic growth model with a finite element method. *Journal of Economic Dynamics and Control*, 20(1-3):19–42.
- Nilsen, O. A. and Schiantarelli, F. (2003). Zeros and lumps in investment: Empirical evidence on irreversibilities and nonconvexities. *The Review of Economics and Statistics*, 85(4):1021–1037.
- Power, L. (1998). The missing link: Technology, investment, and productivity. *The Review of Economics and Statistics*, 80(2):300–313.
- Rouwenhorst, K. (1995). Asset pricing implications of equilibrium business cycle models. In Cooley, T., editor, *Frontiers of Business Cycle Research*, pages 294–330.
- Santos, M. S. (2000). Accuracy of numerical solutions using the euler equation residuals. *Econometrica*, 68(6):1377–1402.
- Santos, M. S. and Peralta-Alva, A. (2012). Analysis of numerical errors. Working Papers 2012-6, University of Miami, Department of Economics.
- Vayanos, D. (1998). Transaction costs and asset prices: A dynamic equilibrium analysis. *Review of Financial Studies*, 11(1):1–58.
- Wang, P. and Wen, Y. (2012). Hayashi meets kiyotaki and moore: A theory of capital adjustment costs. *Review of Economic Dynamics*, 15(2):207 – 225.
- Whited, T. M. (2006). External finance constraints and the intertemporal pattern of intermittent investment. *Journal of Financial Economics*, 81(3):467–502.

A Technical Appendix

In this section, we provide additional details on the implementation of the three approximation methods and derive the Euler equation when the plant is active.

A.1 Euler equation when plant is active

When the plant is active ($I > 0$), the optimal investment strategy can be characterized by an Euler equation. Using equations (3) and (5) the plant's problem in this case can be formulated as

$$V(K, A) = \max_{K'} AK^\alpha - p_I(K' - (1 - \delta)K) - FK - \frac{\gamma}{2}K \left(\frac{K' - (1 - \delta)K}{K} \right)^2 + \beta E_{A'|A} V(K', A').$$

Following Proposition 1 of Clausen and Strub (2012), at the optimal choice of capital tomorrow the following first-order condition holds:

$$p_I + \gamma \frac{K' - (1 - \delta)K}{K} = \beta E_{A'|A} V_{K'}(K', A')$$

where $V_{K'}(\cdot)$ denotes the function's derivative with respect to K' . The following Euler equation characterizes investment dynamics when the plant is active

$$p + \gamma \frac{I}{K} = \beta E_{A'|A} \left(\alpha A'(K')^{\alpha-1} + p(1 - \delta) - F + \frac{\gamma}{2} \left(\frac{I'}{K'} \right)^2 + \gamma(1 - \delta) \frac{I'}{K'} \right), \quad (\text{A.1})$$

where $K'' = (1 - \delta)K' + I'$. Given that the plant is not active for all possible values of the state variables, the above equation holds only when investment is strictly positive.

A.2 Value Function Iteration

We then implement value function iteration via the following algorithm:

1. Create an equally spaced grid for capital over an interval $[K^{\min}, K^{\max}]$.
2. Approximate the AR(1) process for the log of productivity using a 4-state Markov chain. Denote the set of states by \mathcal{A} , and the transition matrix by π_4 . Note that this implies that productivity is drawn from the set $\mathbb{A} \equiv e^{\mathcal{A}}$.
3. Guess an initial value function $V^0(K_i, A_j) = A_j K_i^\alpha$ and policy function $K^0(K_i, A_j) = 0$ at each point $[K_i, A_j]$ on the grid where $K_i \in [K^{\min}, K^{\max}]$ and $A_j \in \mathbb{A}$.
4. Set the tolerance parameter $tol = 10^{-4}$. This parameter is used to determine if the value function has converged.
5. For each level of capital $K_i \in [K^{\min}, K^{\max}]$ and productivity A_j in \mathbb{A} :

- (a) compute the value of being inactive:

$$V^{ina}(K_i, A_j) \equiv A_j K_i^\alpha + \beta E_{A'|A_j} V^0(K_i(1-d), A')$$

- (b) compute the value of being active:

- i. for each possible $K' \geq (1-d)K_i \in [K^{\min}, K^{\max}]$ compute

$$\begin{aligned} \tilde{V}(K') \equiv & A_j K_i^\alpha - p(K' - K_i(1-\delta)) - F K_i - \frac{\gamma K_i}{2} \left(\frac{K' - K_i(1-\delta)}{K_i} \right)^2 \\ & + \beta E_{A'|A_j} V^0(K', A') \end{aligned}$$

- ii. the value of being active is $V^{act}(K_i, A_j) \equiv \max_{K'} \tilde{V}$

- (c) update the value and policy functions

- i. if $V^{ina}(K_i, A_j) \geq V^{act}(K_i, A_j)$ then $V^1(K_i, A_j) = V^{ina}(K_i, A_j)$, $K'^1(K_i, A_j) = (1-d)K_i$
ii. if $V^{ina}(K_i, A_j) < V^{act}(K_i, A_j)$ then $V^1(K_i, A_j) = V^{act}(K_i, A_j)$, $K'^1(K_i, A_j) = \arg \max \tilde{V}$

6. Check if $\|V^0 - V^1\|_\infty < tol$. If not set $V^0 = V^1$, $K'^0 = K'^1$ and repeat Step 5.

7. Verify that $K^{\min} < K'(K^{\min}, A_j) < K'(K^{\max}, A_j) < K^{\max}$ for all j . If not, enlarge grid and repeat Steps 1-6.

A.3 Finite Element Method

We implement a Finite Element Method approximation to the value function via the following algorithm:

1. Set the tolerance parameter $tol = 10^{-4}$.
2. Set the number of grid points for capital to n .
3. Fix the upper and lower bound for capital at $[K^{\min}, K^{\max}]$.
4. Generate an equally spaced capital grid $\{K_i\}_{i=1}^n$ on $[K^{\min}, K^{\max}]$.
5. Approximate the AR(1) process for the log of productivity using a 4-state Markov chain. Denote the set of states by \mathcal{A} , and the transition matrix by π_4 . Note that this implies that productivity is drawn from the set $\mathbb{A} \equiv e^{\mathcal{A}}$.
6. Guess an initial value function $\{\{V_{ij}^0\}_{i=1}^n\}_{j=1}^4$ at each point $[K_i, A_j]$ of the state space. We set $V_{ij}^0 = A_j K_i^\alpha$ for all i, j .
7. We approximate the value function $V(K, A)$ as $\hat{V}(K, A)$, a piece-wise linear interpolation through the points $\{\{V_{ij}^0\}_{i=1}^n\}_{j=1}^4$ where

$$\hat{V}_{ij}^0(K, A) = \begin{cases} V^0(K_i, A_j) + \frac{V_{i+1}^0 - V_i^0}{K_{i+1} - K_i} (K - K_i) & \text{if } K \in [K_i, K_{i+1}] \\ 0 & \text{otherwise} \end{cases}$$

8. For each point in the capital grid find the value of being inactive $V^{ina}(K, A)$ where

$$V^{ina}(K, A) = AK^\alpha + \beta E_{A'|A} \hat{V}^0(K(1-d), A')$$

9. Find the value of being active, $V^{act}(K_i, A_j)$:

- first find $K'(K_i, A_j) = \arg \max_{K' \geq K_i(1-d)} \tilde{V}(K, A) \equiv AK_i^\alpha - p(K' - K_i(1-d)) - FK_i - \frac{\gamma K_i}{2} \left(\frac{K' - K_i(1-d)}{K_i} \right)^2 + \beta E_{A'|A} \hat{V}^0(K', A')$
- the value of being active is then $V^{act}(K_i, A_j) = \tilde{V}(K'(K_i, A_j))$.

10. Update the value and policy functions

- (a) if $V^{inv}(K_i, A_j) \geq V^{act}(K_i, A_j)$ then $V_1(K_i, A_j) = V^{inv}(K_i, A_j)$, $K'(K_i, A_j) = (1-d)K_i$
- (b) if $V^{inv}(K_i, A_j) < V^{act}(K_i, A_j)$ then $V_1(K_i, A_j) = V^{act}(K_i, A_j)$, $K'(K_i, A_j) = \arg \max \tilde{V}$

11. Check if $\|V^0 - V^1\|_\infty < tol$. If not set $V^0 = V^1$ and repeat the steps above.
12. Verify that $K^{\min} < K'(K^{\min}, A_j) < K'(K^{\max}, A_j) < K^{\max}$ for all j . If not, enlarge grid and repeat Steps 1-12.

A.4 Endogenous Grid Method

We implement the Endogenous Grid Method as follows:

1. Set the tolerance parameter $tol = 10^{-4}$.
2. Set the number of grid points for capital to n .
3. Fix the upper and lower bound for capital at $[K^{\min}, K^{\max}]$.
4. Generate a capital grid $\{K'_i\}_{i=1}^n$ for tomorrow's capital.
5. Approximate the process for productivity using a 4-state Markov chain over \mathbb{A} .
6. Guess an initial value for $EV(K', A')$ at each point $[K_i, A_j]$ of the state space and construct a corresponding guess for $EV_{K'}(K', A')$.
7. Construct an endogenous grid of capital points $\{K_i^{end}\}_{i=1}^n$ using the Euler equation (8) to obtain n matching pairs $\{K_i^{end}, K'_i\}_{i=1}^n$.
8. Enforce irreversibility: if $K'/K_i^{end} < (1-d)$ then set $K_i^{end} = K'_i/(1-d)$.
9. Identify the non-concave region
 - (a) identify the set of jumps in $EV_{K'}(K', A')$
 - (b) find the minimum and maximum of values of $EV_{K'}(K', A')$ at these jumps: \underline{V}, \bar{V} .

- (c) the non-concave region for tomorrow's productivity A'_j consists of all pairs $\{K_i^{end}, K'_i\}$ where $EV_{K'}(K'_i, A'_j) \in [\underline{V}, \bar{V}]$.
- 10. For each pair $\{K_i^{end}, K'_i\}$ inside the non-concave region
 - (a) compute the value to being active for every $K'_j \neq K'_i$
 - (b) if the maximum does not occur at K'_i then discard the pair $\{K_i^{end}, K'_i\}$
- 11. Interpolate to recover new endogenous capital values for discarded values of K' .
- 12. Compute values to being active and inactive at each grid point and construct the new value function.
- 13. Update $EV(K', A')$ using the new value function and the transition matrix.
- 14. Check if value function has converged to within tol . If not repeat steps 7-14.

A.5 Results for alternative Parameterizations

Tables 3 - 5 show moments generated by VFI, EGM and FEM for alternative parameterizations. For each of these alternatives we deviate from the parameterization shown in Table 1 by alternating one parameter. We evaluate commonly used values in the literature: a lower value of capital in the production, $\alpha = 0.4$, a higher capital depreciation rate, $\delta = 0.1$, and a higher parameter for the convex capital adjustment costs, $\gamma = 0.1$.

Table 3: Statistics across different approximation methods — alternative parameterization: $\alpha = 0.4$

	capital grid points	spike rate (in %)	investment spike size	capital mean	Euler equation error		CPU (seconds)
					average	maximum	
VFI	200	10.00	3.55	5.71	-4.26	-2.87	6.09
	300	10.00	3.65	5.69	-4.52	-2.91	9.61
	350	11.10	3.12	5.36	-4.66	-3.00	11.61
	400	11.10	3.45	5.85	-4.75	-2.99	13.81
baseline:	420	10.00	3.79	5.77	-4.73	-2.99	14.83
	500	12.50	2.89	5.51	-4.84	-2.90	18.52
	600	11.10	3.07	5.16	-4.82	-2.94	23.63
	1000	11.10	3.30	5.43	-5.14	-3.00	46.37
	1500	10.50	3.60	5.53	-5.31	-3.03	83.34
	2050	11.20	3.43	5.58	-5.42	-3.04	139.14
EGM	31	12.50	3.43	5.53	-4.48	-3.06	0.93
	41	10.00	3.37	5.58	-4.53	-3.17	0.98
	50	10.00	3.44	5.34	-4.63	-3.16	1.05
baseline:	60	11.10	3.41	5.46	-4.73	-3.15	1.26
	80	11.20	3.45	5.42	-4.89	-3.14	1.72
	110	10.20	3.40	5.52	-4.91	-3.14	2.46
	200	11.20	3.42	5.51	-5.22	-3.10	6.08
	400	11.20	3.41	5.52	-5.46	-3.12	19.85
FEM	25	14.20	3.57	5.42	-3.98	-2.88	31.02
	30	12.50	3.48	5.70	-4.25	-2.85	34.58
	40	9.00	3.55	5.55	-4.41	-3.03	42.56
	50	10.00	3.51	5.34	-4.47	-3.16	51.25
	60	11.10	3.42	5.46	-4.59	-3.14	59.09
baseline:	85	11.20	3.43	5.51	-4.74	-3.12	82.36
	100	10.50	3.48	5.44	-4.78	-3.12	92.57
	150	11.20	3.39	5.50	-4.94	-3.12	132.22
	300	11.10	3.44	5.52	-5.12	-3.11	249.09
	800	11.20	3.45	5.56	-5.41	-2.98	614.13
	2000	11.20	3.45	5.56	-5.43	-2.70	1559.57

VFI: Value Function Iteration, EGM: Endogenous Grid Method, FEM: Finite Element Method. The spike rate is calculated as the share of periods in which investment exceeds 20% of the capital stock. Moments are calculated from 50 simulations over 1050 periods where the first 50 periods are discarded. Euler Equation error statistics are calculated across policy functions for all shocks in the area of the state space in which the Euler equation holds. CPU time for FEM is reported utilising four processing units.

Table 4: Statistics across different approximation methods — alternative parameterization: $\delta = 0.1$

	capital grid points	spike rate (in %)	investment spike size	capital mean	Euler equation error		CPU (seconds)
					average	maximum	
VFI	200	12.5	14.96	20.13	-5.58	-4.13	6.18
	300	12.5	13.30	17.54	-5.79	-4.25	10.12
	350	14.3	12.48	18.48	-5.86	-4.31	12.34
baseline:	400	12.5	14.60	19.01	-5.97	-4.33	14.28
	500	14.3	12.61	18.56	-6.04	-4.33	19.38
	600	14.3	11.63	17.15	-6.10	-4.35	24.51
	1000	14.3	12.47	18.12	-6.27	-4.41	49.47
	1500	14.3	12.43	17.98	-6.43	-4.44	89.21
EGM	20	16.7	12.41	19.63	-5.42	-4.28	0.94
	30	14.3	12.53	17.57	-5.79	-4.47	1.02
	40	14.3	12.47	18.49	-5.86	-4.57	1.08
	50	15.3	12.55	18.53	-5.82	-4.54	1.17
baseline:	58	14.3	12.57	17.97	-5.99	-4.57	1.33
	80	14.3	12.49	18.12	-6.02	-4.56	1.71
	100	14.3	12.49	18.14	-6.12	-4.54	2.44
	200	14.3	12.55	18.03	-6.31	-4.53	7.70
	400	14.3	12.52	18.03	-6.53	-4.51	27.01
FEM	25	14.3	12.73	18.73	-5.55	-4.47	33.86
	30	14.3	12.30	17.57	-5.65	-4.53	38.48
	40	14.3	12.55	18.49	-5.70	-4.57	47.35
	50	15.3	12.57	18.53	-5.80	-4.40	56.81
	60	14.3	12.62	18.52	-5.91	-4.55	65.76
baseline:	75	14.3	12.40	17.83	-5.97	-4.56	79.83
	100	14.3	12.40	18.14	-6.19	-4.54	103.42
	150	14.3	12.65	18.15	-6.24	-4.53	147.35
	300	14.3	12.59	18.18	-6.27	-4.53	283.11
	800	14.3	12.63	18.15	-6.49	-4.47	722.43

VFI: Value Function Iteration, EGM: Endogenous Grid Method, FEM: Finite Element Method. The spike rate is calculated as the share of periods in which investment exceeds 20% of the capital stock. Moments are calculated from 50 simulations over 1050 periods where the first 50 periods are discarded. Euler Equation error statistics are calculated across policy functions for all shocks in the area of the state space in which the Euler equation holds. CPU time for FEM is reported utilising four processing units.

Table 5: Statistics across different approximation methods — alternative parameterization: $\gamma = 0.1$

	capital grid points	spike rate (in %)	investment spike size	capital mean	Euler equation error		CPU (seconds)
					average	maximum	
VFI	200	12.50	13.92	28.30	-6.19	-4.67	6.53
	300	10.00	19.40	30.26	-6.53	-4.88	10.38
	350	11.10	16.59	28.47	-6.68	-4.92	12.69
	400	11.10	18.35	31.08	-6.67	-4.90	15.14
baseline:	420	10.00	20.33	30.94	-6.75	-4.90	15.91
	500	11.10	19.60	32.99	-6.79	-4.91	19.80
	600	11.80	18.53	31.13	-6.85	-4.93	25.16
	1000	12.50	16.55	30.66	-7.19	-5.00	50.19
	1500	11.80	19.01	31.46	-7.38	-5.05	91.06
	1750	12.50	16.44	30.16	-7.43	-5.04	130.21
EGM	30	14.30	17.20	29.08	-6.15	-4.91	1.03
	40	10.00	17.23	30.36	-6.41	-5.02	1.11
	50	11.10	17.27	30.63	-6.47	-5.08	1.28
	60	12.50	17.25	31.16	-6.55	-5.13	1.37
baseline:	70	12.50	17.24	30.05	-6.73	-5.10	1.56
	80	11.80	17.27	30.63	-6.67	-5.11	1.83
	100	12.50	17.26	30.82	-6.82	-5.12	2.29
	200	12.50	17.26	30.42	-7.06	-5.11	7.61
	400	12.50	17.25	31.20	-7.37	-5.10	29.74
	450	12.20	17.27	30.46	-7.43	-5.10	44.82
FEM	25	16.70	16.84	30.25	-6.29	-4.79	33.16
	30	14.30	17.23	29.08	-6.38	-4.89	36.94
	40	10.00	17.52	30.36	-6.54	-5.14	47.65
	50	11.10	17.43	30.63	-6.60	-5.13	60.74
	60	12.50	17.13	31.16	-6.61	-5.11	68.25
baseline:	65	12.50	17.16	29.88	-6.74	-5.13	73.11
	100	12.50	17.26	30.82	-6.91	-5.11	108.25
	150	12.50	17.17	30.43	-7.11	-5.08	157.01
	300	12.50	17.16	30.93	-7.27	-5.08	305.97
	500	12.50	17.00	30.96	-7.42	-5.08	717.14

VFI: Value Function Iteration, EGM: Endogenous Grid Method, FEM: Finite Element Method. The spike rate is calculated as the share of periods in which investment exceeds 20% of the capital stock. Moments are calculated from 50 simulations over 1050 periods where the first 50 periods are discarded. Euler Equation error statistics are calculated across policy functions for all shocks in the area of the state space in which the Euler equation holds. CPU time for FEM is reported utilising four processing units.