

Endogenous grid method*

Giulio Fella[†]

March 29, 2025

Abstract

The endogenous gridpoint method (EGM), pioneered by Carroll (2006), offers an efficient alternative to traditional numerical approaches for solving dynamic economic models with realistic heterogeneity, nonlinearities, occasionally binding constraints, and discrete-continuous choices. Unlike exogenous grid methods (EXGM), EGM reverses the standard solution logic by fixing the value of post-decision states and solving for pre-decision states using the Euler equation. The resulting problem involves finding the zero of a known, and often linear, function as opposed to one which is known only at a finite set of points. This significantly reduces interpolation errors and computational costs. EGM is especially advantageous in models with borrowing constraints, discrete choices, or non-concave objectives, offering significant computational savings and improved accuracy. This survey provides a comprehensive overview of EGM, detailing its core principles, implementation, and extensions to more complex settings, including problems with multiples continuous state variables and discrete-continuous choices. By serving as both an introduction for newcomers and a roadmap for advanced applications, this survey underscores the accuracy, computational efficiency and versatility of EGM which make it a valuable tool for solving high-dimensional dynamic optimization problems in economics and finance.

JEL classification: C63

Keywords: Endogenous grid method, dynamic programming, computational economics, Euler equation, nonconcavity

[†]Giulio Fella: Queen Mary University of London, University of Bologna, CFM and IFS, g.fella@qmul.ac.uk

1 Introduction

Dynamic economic models are fundamental tools for analyzing a wide range of intertemporal decisions, from household consumption-savings problems to firm investment choices and general equilibrium models. In the presence of realistic heterogeneity, occasionally binding constraints, adjustment costs or discrete choices the resulting value and policy functions are characterized by substantial nonlinearities and possibly nonconcavities and non-differentiabilities which make them unsuitable to local solution methods.

Traditional global methods to solve such problems involve choosing an exogenous, time-invariant, grid for the beginning-of-period (or “pre-decision”) state vector, solving forward for the end-of-period (or “post-decision”) optimal choice for the endogenous state variables and iterating on the value function and/or policy functions. While conceptually straightforward, this “exogenous grid” approach (EXGM) requires numerical root-finding and repeated interpolation of conditional expectations over large sets of points, ultimately increasing the computational burden and compounding the well known *curse of dimensionality* (Bellman, 1961). Moreover, in some models, such as those with occasionally binding constraints, exogenous grids can lead to significant interpolation errors near kinks of the policy functions.

The endogenous gridpoint method (EGM) pioneered by Carroll (2006) addresses these limitations by reversing the usual logic: instead of solving for the choice that corresponds to a beginning-of-period state, it chooses a grid for the end-of-period’s endogenous state variables and uses the Euler equation to solve endogenously for the associated value of the beginning-of-period endogenous state variables. The intuition is possibly best conveyed in the case of an optimal saving problem. Rather than solving for the next-period wealth $a'(a)$ on a grid for current wealth a , it solves for its inverse $a = a'^{-1}$ on an exogenous grid for a' , taking advantage of the fact that the Euler equation is often linear in, or a closed-form function of, a but not a' . This eschews numerical root-finding and the associated repeated evaluation of the Euler equation and the associated expectation. An

additional advantage in the case of models with exogenous borrowing constraints is that the location of the constraint on a' is known and the corresponding value of initial wealth can be computed exactly.

This survey presents a unified exposition of the endogenous gridpoint method, highlighting its core ideas, implementation details, and extensions that have been proposed to handle more complex environments than the optimal saving problem. We begin with the foundational single-asset framework, comparing EGM's computational performance and accuracy to those of conventional EXGM or time iteration. We then show how EGM can be extended to problems involving multiple continuous state variables, as well as to non-concave and non-differentiable problems involving a mixture of continuous and discrete choices, such as housing or default decisions. Throughout, we emphasize both the key theoretical insights—why and when EGM works well—as well as the practical steps needed to implement it efficiently in applications.

Economic applications of EGM. EGM has gained substantial popularity in solving high-dimensional dynamic problems in economics. The following provides only an incomplete survey of its applications. Carroll's (2006) seminal contribution introduced EGM in the context of optimal saving problems with a single state and control variables. Barillas and Fernández-Villaverde (2007) extended it to an environment, the stochastic growth model with endogenous labour supply, with more than one control. Hintermaier and Koeniger (2010) introduced a hybrid EGM method that can accommodate an additional continuous state variable. Guerrieri and Lorenzoni (2017) apply it to study the implications of a tightening of borrowing limits in a model with both durable and non-durable goods. Bayer, Luetticke, Pham-Dao and Tjaden (2019) use it to study the aggregate implications of surprise changes in aggregate income uncertainty in an environment with liquid and illiquid assets. White (2015) and Ludwig and Schön (2018) both proposed non-hybrid extensions of EGM to problem with multiple continuous state variables and how to address the complications of interpolating over the resulting multi-dimensional, non-regular endogenous grid. The need to speed up the estimation of the rich life-cycle

model with discrete education and crime choices in Fella and Gallipoli (2014) lead to Fella’s (2014) extension of EGM to non-concave and non-differentiable environments with both a discrete and continuous state (DC-EGM) variables. Iskhakov, Jørgensen, Rust and Schjerning (2017) derived independently a very similar DC-EGM algorithm. In addition to proposing a more efficient way to discard local, but not global, maxima, they show how extrinsic taste shocks can be used to control the propagation of kinks inherent to discrete-continuous problems. Iskhakov and Keane (2016) and De Nardi, Fella and Paz-Pardo (2024) use DC-EGM to study, respectively, the effect of social security and optimal welfare policies in life-cycle models with wealth and human capital accumulation and a discrete labour supply choice. Ameriks, Briggs, Caplin, Shapiro and Tonetti (2020) study the interaction between precautionary saving and bequest motive in a life-cycle model with a binary choice between private and public provision of long-term care. Yao, Fagereng and Natvik (2021) study how housing and mortgage debt influence the marginal propensity to consume in a life-cycle model with a discrete housing choice and non-convex transaction costs. Finally, Druedahl and Jørgensen (2017) further generalised EGM to environment with multiple continuous and discrete state variables and proposes a simpler interpolation technique compared to White (2015) and Ludwig and Schön (2018). Druedahl and Jørgensen (2018) use it to study the coexistence of positive gross credit-card debt and positive gross wealth, while Ejrnaes and Jørgensen (2020) apply it to study fertility and abortion choices in a life-cycle model.

Goals and outline. This survey aims to serve two purposes. First, it acts as a tutorial for researchers who wish to implement EGM in relatively standard models but may not have been fully aware of the associated computational savings and accuracy improvements. Second, it provides a roadmap for tackling challenging classes of dynamic problems—from multidimensional concave problems to discrete-continuous choice models—where EGM and its variants can be usefully applied. Ultimately, the goal is to underscore the versatility of EGM and inspire further application of its core insight into contemporary quantitative economics research.

The survey is organized as follows. Section 2 introduces EGM in Carroll’s (2006) original set up with a single continuous state variable. The simple set up is used to bring out the differences from standard exogenous grid methods, such as value or policy function iteration, that account for the substantial gains in computational costs. The section also studies a numerical example that quantifies the relative contribution of each of these difference to the overall improvement.¹ Section 3 discusses extensions of EGM to environments with multiple continuous state variable that satisfy concavity and differentiability of the objective. It is well known that discrete choices and fixed costs introduce nonconcavity and non-differentiabilities even with respect to the continuous state variables of a problem. Section 4 discusses extensions of EGM to problems with both discrete and continuous state variable for which the advantages, in terms of execution time and accuracy, of EGM are starkest. Section 5 concludes

2 EGM with a single, continuous state variable

This section introduces EGM and discusses its advantages and tradeoffs in a slightly simplified version of the income fluctuation problem in Carroll’s (2006) seminal paper.

Consider an infinitely lived consumer whose goal is to maximize lifetime utility

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c(t))$$

where $c(t)$ is consumption at time t and $\beta \in (0,1)$ is the discount factor. Time is discrete. The felicity function $u(\cdot)$ is twice differentiable, strictly increasing and concave and satisfies the Inada conditions. In each period, households draw stochastic labor income y taking non-negative values in the ordered grid $G_y = \{y_1, y_2, \dots, y_m\}$ according to a stationary, first-order Markov chain. The consumer can save at the gross risk-free rate R but cannot borrow. For the moment, we are also going to assume that the income

¹The Julia code used for the numerical example is available at https://github.com/gfell/egm_oreef.

process implies that the borrowing constraint is never binding along an optimal path.²

The consumer problem can be written in recursive form as

$$V(a, y) = \max_{a' \geq 0} u(y + Ra - a') + \beta EV(a', y'). \quad (1)$$

Under our assumptions, a solution to the Bellman equation exists and is unique, the value function V is differentiable, strictly increasing and strictly concave. The optimal policy satisfies the first-order condition (FOC)

$$u'(y + Ra - a') = \beta EV_a(a', y'), \quad (2)$$

where V_a , the partial derivative of V with respect to a' , satisfies the envelope condition

$$V_a(a, y) = Ru'(y + Ra - a'). \quad (3)$$

Conventional exogenous grid methods (EXGM), such as value function or time iteration, solve the problem characterized by equations (1) and (2) where (a, y) is the state at the beginning of the period or, equivalently, before the choice of a' (*pre-decision*). More specifically, if one denotes by $G_a = \{\alpha_1, \alpha_2, \dots, \alpha_n\}$ the discrete grid for wealth and by $G_{ay} = G_a \times G_y$ the grid for the state vector (a, y) , conventional EXGM finds the value and policy functions by the following iterative procedure.

Conventional (pre-decision state) EXGM

Given a guess of values $V_a^n(a, y)$ with $(a, y) \in G_{ay}$, iterate until convergence on:

1. For each point $(a, y) \in G_{ay}$ find $a'(a, y)$ satisfying the FOC (2).
 2. Replace for a' in (1) or (3) to compute $V_a^{n+1}(a, y)$.
-

Step 1 in conventional EXGM is costly because the FOC (2) cannot be solved analytically for a' . Therefore, for each (a, y) the FOC has to be evaluated at each element

²This would be the case if the the lower bound on the labor income support $y_1 = 0$.

a'_k of the sequence of guesses of the numerical root finding routine. Each such evaluation involves (a) interpolating V_a^n at points (a'_k, y') ; and (b) integrating over the y' to compute the approximate conditional expectation $EV_a^n(a'_k, y')$. The computation of the conditional expectation for *each* candidate zero a'_k , is actually the costlier part of Step 1.³

The endogenous grid method (EGM) of Carroll (2006) provides an alternative solution method that not only avoids the repeated computation of the expectation, but also the costly numerical root finding. Rather than solving the FOC (2) for the end-of-period endogenous state variable a' given (a, y) , it treats a' as fixed and solves for the beginning-of-period state variable a given (a', y) . In other words, EGM defines exogenously a grid G_a for a' and, for each income state y , finds *endogenously* the value of beginning-of-periods assets a^{end} for which $a' = a$ is the optimal saving choice. To the extent that the marginal utility of consumption has an analytic inverse the approach eschews root finding as the FOC (2) can be rewritten as

$$y + Ra = a' + u'^{-1}[\beta EV_a(a', y')]. \quad (4)$$

In other words, EGM exploits the fact that the FOC is analytic, in fact linear, in a although not in a' .

Formally, the EGM algorithm is as follows.

EGM

Given a guess of values $V_a^n(a, y)$ with $(a, y) \in G_{ay}$, iterate until convergence on:

1. For each point $(a', y) = (\alpha_i, y_k) \in G_{ay}$ find $a_{ik}^{end} = a$ satisfying the FOC (4).
This defines the value of the policy function $a'(s_{ik}) = \alpha_i$ in state $s_{ik} = (a_{ik}^{end}, y_k)$.
 2. For each y_j interpolate $\{(s_{ik}, a'(s_{ik}))\}_{i=1}^n$ on G_a to obtain $a'(\cdot)$ on the, iteration-invariant, exogenous grid G_{ay} .
 3. Replace for a' in (1) or (3) to compute $V_a^{n+1}(a, y)$.
-

³Step 2 is meant to encompass both value function and time iteration. The former uses the Bellman equation (1) to update the value function and then computes its partial derivative. The latter directly updates the partial derivative using the envelope condition (3).

Step 1 highlights the big advantage of EGM compared to conventional EXGM. The conditional expectation on the right of the FOC is a function of y and the end-of-period (*post-decision*) endogenous state variable a' . By making a' fall on a fixed grid, EGM evaluates the expectation (and the FOC in general) only on the $(n \times m)$ grid points of G_{ay} . Under conventional EXGM, the number of evaluations of the expectation (and the FOC in general) equals the number of iterations of the root finding routine for *each* of the $(n \times m)$ grid points for the beginning-of-period (pre-decision) state (a, y) .

Step 2, though, entails a cost which is absent from EXGM. In order to check convergence of the solution consecutive iterates of the value function or its derivative have to be evaluated on the same set of points. This requires interpolating the policy function computed on the endogenous grid on a time-invariant, exogenous grid.⁴ This trade-off is present only in infinite horizon models. In finite horizon models one can let the grid for a' at time t coincide with the endogenous grid for a at time $t + 1$, as in Iskhakov (2015) and Iskhakov et al. (2017).

The above discussion of the main advantage of EGM compared to EXGM is not fully fair as conventional (pre-decision) EXGM, though the most common, is not the most efficient implementation of EXGM. Consider the alternative formulation of our problem where the choice of state is (a', y) , the state at the end of the period or, equivalently, after the choice of a' (post-decision).⁵ Let $W(\cdot, \cdot)$ denote the associated value function which implies the alternative formulation of the Bellman equation (1)

$$W(a, y_{-1}) = E[\max_{a'} u(y + Ra - a') + \beta W(a', y)], \quad (5)$$

⁴One could invert the ordering of step 2 and 3 and interpolate the value function or its partial derivative on the endogenous grid.

⁵Wright and Williams (1982) and Wright and Williams (1984) were the first to use a solution method based on a post-decision state. Judd (1998, p. 429) shows how to construct the end-of-period state Bellman equation. The method is extensively used in approximate dynamic programming; i.e. reinforcement learning (Van Roy, Bertsekas, Lee and Tsitsiklis, 1997; Powell, 2007).

where $E(\cdot)$ denotes the conditional expectation over y . The associated FOC is

$$u'(y + Ra - a') = \beta W_a(a', y) \quad (6)$$

and the envelope condition

$$W_a(a, y_{-1}) = REu'(y + Ra - a'). \quad (7)$$

It is straightforward to see that $W(a', y)$ in equation (5) equals $EV(a', y')$ in equation (1); the post-decision value function is the conditional expectation of the pre-decision one. Under such a formulation EXGM finds the value and policy functions by the following iterative procedure.

Post-decision state EXGM

Given a guess of values $W_a^n(a, y)$ with $(a, y) \in G_{ay}$, iterate until convergence on:

1. For each point $(a, y) \in G_{ay}$ find $a'(a, y)$ satisfying the FOC (6).
 2. Replace for a' in (5) or (7) to compute W_a^{n+1} .
-

If the value function is defined in terms of the end-of-period, or post-decision, endogenous state, finding a solution to the FOC in step 1 is a deterministic problem, although it still requires numerical root finding. As step 2 makes clear though the conditional expectation is now outside the maximum operator and is computed only once for each of the $n \times m$ points on the grid G_{ay} . As is the case with EGM, the post-decision state EXGM minimizes the number of numeric integrals to compute in solving the problem.⁶

Having established a fair EXGM benchmark we can finally list the relative merits of EGM. First, EGM avoids numerical root finding in Step 1, although in infinite horizon problems this has to be traded off against the additional interpolation in Step 2. Second

⁶Strictly speaking this is the case only within the class of non-parametric solutions for the value function. Judd, Maliar, Maliar and Tsener (2017) propose a method which further reduces the cost of computing the expectation provided one is willing to restrict attention to value functions in the (large) class of analytic functions which are separable in endogenous and exogenous state variables.

and related, EGM produces an exact solution to the Euler equation—the Euler errors exactly equal zero on the endogenous grid points—while for EXGM this is the case only up to an interpolation/approximation error.⁷ Third, in models with occasionally binding inequality constraints on the post-decision state variables, EGM determines exactly the level of beginning-of-period wealth where the constraints become binding. For example, suppose that the lowest income state $y_1 > 0$ and the ad-hoc borrowing constraint $a' \geq 0$ is binding with strictly positive probability. If the first grid point in G_a is $\alpha_1 = 0$, for each $y_j \in G_y$ EGM solves for the value of initial wealth a_{1j}^{end} for which the FOC holds with $a' = 0$. By monotonicity, the policy function is known in closed form— $a' = 0$ —for all $a \leq a_{1j}^{end}$. Therefore, unlike EXGM, EGM never interpolates across the level of initial wealth for which the borrowing constraint becomes binding. On the other hand, EGM is at a disadvantage when solving models featuring inequality constraints in terms of beginning-of-period wealth such as models with default risk.

2.1 Numerical performance comparison

This section compares the computational speed and accuracy of the three algorithms discussed above by solving a parameterized version of problem (1). It assumes log utility, $R = 1.025$ and a discretized (log-)income process which corresponds to an AR(1) with autoregressive coefficient $\rho = 0.97$ and conditional standard deviation $\sigma = 0.24$. The discretization follows Rouwenhorst with 11 grid points. The calibrated value of the discount factor is $\beta = 0.955$ to match an average wealth-income ratio of about 4.4. The parameterization implies that the borrowing constraint is binding with positive probability on the optimal path. Finally, the asset grid has 100 points. The solution is computed by piecewise-linear interpolation off nodes using Julia version 1.10 on an Apple[®] MacBook Pro[®] with M2 CPU and 16GB RAM. To solve non-linear equations we use NLSolve.

Table 1 illustrates the performance advantages, in terms of both of both execution

⁷Note that step 2 in the EGM algorithm is necessary only to establish convergence. Upon convergence, one can use the policy function on the endogenous grid in the model simulation.

Table 1: Speed and accuracy of EGM and time iteration.

	EGM		Time iteration	
	End. grid simulation (1)	Ex. grid simulation (2)	Pre-decision (3)	Post-decision (4)
CPU (s)	1.39 (2.43)	1.39 (2.43)	19.98	4.84
L_1	-3.89	-3.94	-4.02	-3.53
L_∞	-2.04	-1.39	-1.39	-1.26
W/Y ratio	4.42	4.42	4.43	4.20

Notes: CPU is the time (in seconds) necessary to compute the solution; L_1 and L_∞ are, respectively, the average and maximum of absolute residuals of the Euler equation across test points (in log10 units) on a stochastic simulation of 200,000 observations.

time and accuracy, of EGM against standard time iteration. The reported measures of accuracy are the average (L_1) and maximum (L_∞) absolute value of the log 10 residuals in the Euler equation (2), over all test points at which it theoretically holds (off-corners).

Columns 3 and 4 report results for time iteration using respectively the pre-decision and post-decision endogenous state. As anticipated, the post-decision state method is substantially faster (by a factor of 4). The trade-off is lower accuracy not only as reflected in a higher average of the Euler equation errors but also as a 5 per cent larger steady-state wealth-income ratio. Since, marginal utility is a convex function of consumption, the linear interpolation of marginal utility, as opposed to consumption, in the post-decision method results in an upward bias in future expected marginal utility and therefore in saving and the wealth-income ratio.

Column 1 and 2 report the same statistics for EGM using, respectively, the endogenous and exogenous grid for the simulation. In terms of execution time, EGM is about 3 times, 1.39 against 4.84 seconds, faster than post-decision time iteration and 14 times than pre-decision time iteration. When using the same (exogenous) grid for the simulation as time iteration, EGM has a very similar level of accuracy as pre-decision time iteration. When simulating using the endogenous grid, EGM displays a lower maximum Euler equation error because it never interpolates around the kink in the policy function.

Finally, the numbers in parenthesis for CPU time in the EGM columns refer to the case

in which the Euler equation is solved numerically, rather than analytically as in equation (4), for a . In this case, the execution times increases by about 75% to 2.43 seconds. EGM is still twice as fast as post-decision time iteration. This is due to the fact that, in the case of EGM, numerical root finding involves evaluation of the analytic function $u'(c)$ with known analytic derivative which is significantly less costly than repeated interpolation of the expected marginal utility in the case of time iteration.

The comparison is of more than of academic interest. As equation (4) makes clear, the ability to solve for a as a function of a' in closed form relies on both (a) marginal utility u' and (b) the total resource function $z(y, a) = Ra + y$ having an analytic inverse. Neither property is general. For example, in the neoclassical growth model total resources are given by the function $Af(k) + (1 - \delta)k$ which is not invertible with respect to capital k . However, Carroll (2006) shows that one can still eschew root finding (except at the very last iteration) by using z as an intermediate state variable to exploit the fact the Euler equation still admits an analytical solution for z .⁸ Numerical root finding cannot be avoided though when u' is not analytically invertible.⁹ Hallengreen, Jørgensen and Olesen (2024) show that significant gains in speed can be obtained by approximating the inverse marginal utility u'^{-1} by an interpolator which can be constructed before solving the model.

3 EGM with multiple continuous state variables

Section 2 has illustrated the gains and tradeoffs associated with EGM in simple problems with a single, continuous, endogenous state variables. The gains in computational speed compound with more complex problems, but some additional difficulties also arise. This section discusses these issues introducing a continuous durable choice into the consumer

⁸Barillas and Fernández-Villaverde (2007) use the neoclassical growth model with endogenous labour supply to generalise Carrol's method to models with multiple controls, though still only one endogenous state variable.

⁹This is the case, for example, in models with non-separable utility in consumption and housing services (e.g. Bajari, Chan, Krueger and Miller, 2013; Fella, 2014), or in the intra-household bargaining model in Mazzocco (2007) in which utility depends on both private and public consumption.

problem of the previous section. The resulting model has two continuous endogenous state variables, but our results generalise to a richer class of problems; namely, concave problems for which FOCs are both necessary and sufficient.

Let d denote the current stock of durables. The felicity function $u(c, d)$ is strictly increasing and concave and twice-differentiable. Durables depreciate at rate δ .

The consumer problem can be written in recursive form as

$$V(a, d, y) = \max_{a', d' \geq 0} u(y + Ra + (1 - \delta)d - d' - a', d) + \beta EV(a', d', y').$$

Assuming, for simplicity, an interior maximum the FOCs¹⁰ of the problem are

$$u_c(y + Ra + (1 - \delta)d - d' - a', d) = \beta EV_a(a', d', y') \quad (8)$$

and

$$u_d(y + Ra + (1 - \delta)d - d' - a', d) = \beta EV_d(a', d', y') \quad (9)$$

which can be combined into the intra-temporal condition

$$EV_a(a', d', y') = EV_d(a', d', y'). \quad (10)$$

There are two alternative ways to apply EGM to the above problem. The first one, which we just label EGM, is a straightforward extension of the univariate case. Let G_d denote a discrete grid for durables. EGM defines an exogenous grid $G_a \times G_d$ for the end-of-period endogenous state variables (a', d') and, for each income state y , solves for the associated values of beginning-of-period assets and durables (a^{end}, d^{end}) that satisfy the FOCs (8) and (9). If one denotes by I and J the sets of indices for grid points, respectively, in G_a and G_d the following pseudo-code applies.

¹⁰The envelope conditions are $V_a(a, d, y) = Ru_c$ and $V_d(a, d, y) = (1 - \delta)u_c + u_d$.

EGM

Given a guess of values $V_a^n(a, d, y)$ with $(a, d, y) \in G_{ady} = G_a \times G_d \times G_y$ iterate until convergence on:

1. For each point $(a', d', y) = (\alpha_i, d_j, y_k) \in G_{ady}$ find $a_{ijk}^{end} = a$ and $d_{ijk}^{end} = d$ satisfying the two FOCs (8) and (9).

These define the values of the policy functions $a'(s_{ijk}) = \alpha_i$ and $d'(s_{ijk}) = d_k$ in state $s_{ijk} = (a_{ijk}^{end}, d_{ijk}^{end}, y_k)$.

2. For each y_k interpolate $\{(s_{ijk}, a'(s_{ijk}))\}_{(i,j) \in I \times J}$ and $\{(s_{ijk}, d'(s_{ijk}))\}_{(i,j) \in I \times J}$ on $G_a \times G_d$ to obtain $a'(\cdot)$ and $d'(\cdot)$ on the, iteration-invariant, grid G_{ady} .
-

The second method, which we label hybrid EGM (HEGM), is due to Hintermaier and Koeniger (2010) and Ludwig and Schön (2018). It defines an exogenous grid for end-of-period assets and *beginning-of-period* durables (a', d) . For each income state y and $d \in G_d$: (i) it uses the intra-temporal condition (10) to solve for the optimal d' for each $a' \in G_a$; and (ii) it uses either (8) or (9) to find endogenously the values of beginning-of-period assets a^{end} given (a', d, d') . The corresponding pseudo-code follows.

HEGM

Given a guess of values $V_a^n(a, d, y)$ with $(a, d, y) \in G_{ady} = G_a \times G_d \times G_y$ iterate until convergence on:

1. For each point $(a', d, y) = (\alpha_i, d_j, y_k) \in G_{ady}$, (i) find $\hat{d} = d'$ satisfying equation (10); (b) replace in either (8) or (9) to find $a_{ijk}^{end} = a$.

This defines the values of the policy functions $a'(s_{ijk}) = \alpha_i$ and $d'(s_{ijk}) = \hat{d}$ in state $s_{ijk} = (a_{ijk}^{end}, d_j, y_k)$.

2. For each (d_j, y_k) interpolate $\{(s_{ijk}, a'(s_{ijk}))\}_{i \in I}$ and $\{(s_{ijk}, d'(s_{ijk}))\}_{i \in I}$ on G_a to obtain $a'(\cdot)$ and $d'(\cdot)$ on the, iteration-invariant, grid G_{ady} .
-

At first sight, it may seem that EGM is faster. If u_c is analytically invertible for c , the FOCs (8) and (9) can be solved analytically for a^{end} and d^{end} .¹¹ Even if u_c does not have an analytic inverse, solving the two FOCs requires only evaluation of known functions. HEGM, instead, requires solving the intra-temporal condition (10) for d' which involves

¹¹Iskhakov (2015) characterises the class of concave problems that can be solved by EGM without root finding.

repeated interpolation as typically V and its derivatives are only known on a finite grid.

In fact, given that the infinite horizon setup considered requires interpolating on an iteration-invariant grid in order to establish convergence, the trade-off is more involved. The costs of interpolation in step 2 are significantly different between the two methods.

HEGM interpolates only in one dimension. The set of grid points s_{ijk} is off the, iteration-invariant, grid G_{ady} only along the asset dimension due to a_{ijk}^{end} being endogenous. So interpolation is straightforward, just as in the univariate problem in Section 2.

EGM, instead, interpolates with respect to both assets and durables using the irregular (non-rectilinear) grid of endogenous interpolating nodes $\{(a_{ijk}^{end}, d_{ijk}^{end})\}_{(i,j) \in I \times J}$. Unlike in a rectangular grid, one cannot locate the sector in which a point lies by univariate search along each grid dimension as the dimensions are not orthogonal. This is illustrated in Figure 1. A univariate search for the location of point X in the irregular, endogenous grid would incorrectly place it in the grey shaded sector, rather than the sector to its left. Locating the correct sector requires starting from some initial guess and moving from one sector to the other depending on which sector boundaries are violated. In general, such a procedure, typically referred to as “visibility walk”, may fail to converge and get stuck in a cycle.

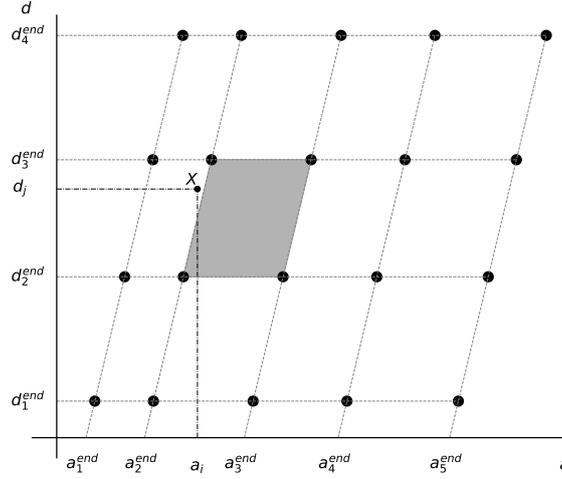
Ludwig and Schön (2018) propose to partition the endogenous grid into triangles using Delaunay triangulation which ensures convergence of the visibility walk. After locating the triangle containing the query point X the associated function value is computed using barycentric interpolation with interpolating nodes given by the vertices of the triangle.¹² Delaunay triangulation is costly¹³, though, and has to be carried out at each iteration. Ludwig and Schön (2018) report that EGM with Delaunay triangulation is actually slower than HEGM for dense (200x200) grids.

White (2015) proposes an alternative approach with substantially lower computational costs. Its key insight is that an irregular grid with sectors with 2^n vertices, such as the

¹²The barycentric weights ω for a point $X \equiv (a, d)$ in the triangle with coordinates (a_i^{end}, d_i^{end}) , $i = 1, 2, 3$ solve the system of equations $a = \omega_1 a_1^{end} + \omega_2 a_2^{end} + (1 - \omega_1 - \omega_2) a_3^{end}$ and $d = \omega_1 d_1^{end} + \omega_2 d_2^{end} + (1 - \omega_1 - \omega_2) d_3^{end}$.

¹³Its construction time increases at rate $O(n \log n)$ with the total number of grid points n .

Figure 1: Irregular grid



Notes: Interpolation on a non-rectilinear grid. Searching independently in each grid dimension does not locate the sector in which $X \in \mathbb{R}^2$ resides.

one in Figure 1, can be continuously mapped to the $[0, 1]^n$ hypercube. Multi-linear interpolation can then be applied on the transformed set of nodes. The computation costs increases only at approximately linear rate with the total number of nodes n . The additional advantage is that the method can be parallelized while Delaunay triangulation cannot due to its sequential construction. The method still involves a visibility walk to identify the sector in the endogenous grid to be mapped onto the unit hypercube. Its limitation is that, contrary to Delaunay triangulation, the visibility walk is guaranteed to converge only in a smaller, though still very large, class of problems for which policy functions are monotonic in the endogenous state variables.¹⁴

Finally, Druedahl and Jørgensen (2017) propose what they call a “local” triangulation that eschews the visibility walk altogether. Their method cleverly exploits the fact that there is a one-to-one mapping between points on the grid for the post-decision endogenous states and on the associated endogenous grid. In terms of the problem in this section, given $y = y_k$, each point $(a', d') = (\alpha_i, d_j) \in G_a \times G_d$ is associated with a unique pair $(a_{ijk}^{end}, d_{ijk}^{end})$. Therefore, a partition of the grid $G_a \times G_d$ into rectangular triangles induces a

¹⁴White (2015) also provides general condition for the applicability of EGM.

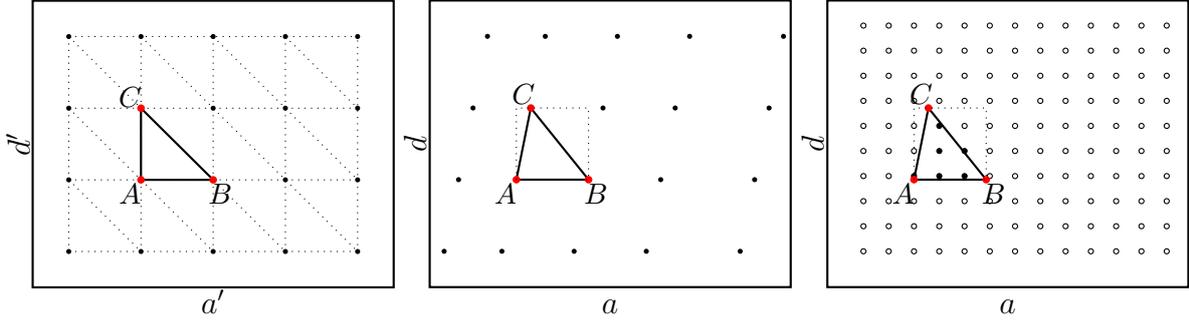
partition of the endogenous grid into corresponding, although not in general rectangular, triangles. This is illustrated in Figure 2. Panel 2a plots the grid for the post-decision endogenous variables (a', d') , its partition into rectangular triangles; e.g. A, B, C . Panel 2b plots the non-rectangular endogenous grid for the pre-decision variables (a, d) and the triangle associated with ABC in panel 2b and its bounding box, the (dotted) smallest rectangle that contains it. One can then use each such triangle and associated bounding box in the (a, d) space to interpolate at the points of an exogenous, iteration-invariant, rectilinear grid G_{ady} (panel 2c). For each point in G_{ady} which falls in the bounding box, one computes the barycentric interpolation weights with nodes given by the vertices of the triangle. If the barycentric weights are non-negative, the point is weakly interior to the triangle ABC (solid black circles in panel 2c) will be used for interpolation. If instead one of the barycentric weights is negative, the point is exterior to the triangle and there are two possibilities. In the first case, the point lies outside the boundary of the endogenous grid in panel 2a and the weights are used to extrapolate. In the alternative case, interpolation will not be carried out as the point is weakly interior to some other triangle.¹⁵

Although Druedahl and Jørgensen (2017) do not benchmark their EGM method for a concave problem of the kind considered in this section, it seems potentially fast compared to the alternatives considered here. The general message of this section, though, is that the extension of EGM to concave multidimensional problems is not trivial and the trade-off between coding and execution time not necessarily favourable for somebody not already heavily involved with the method.

¹⁵The method, as described in Druedahl and Jørgensen (2017), is actually more general, but more costly, that the one described here as it can accommodate, non-concave, problems in which the policy functions are not continuous and the induced triangulation in the endogenous grid is not necessarily a partition of it.

Figure 2: Druedahl and Jørgensen’s (2017) local triangulation

(a) Triangle in (a', d') space (b) Triangle in (a, d) space (c) Barycentric interpolation



Notes: The plots illustrates the “backward-induction” triangulation in Druedahl and Jørgensen (2017). A triangle in the exogenous grid for (a', d') (panel a) maps into a triangle in the endogenous irregular grid for (a, d) and associated bounding box (panel b). The bounding box identifies the candidate points in the exogenous, rectangular grid in (a, d) space for interpolation with nodes given by the vertices of the triangle. Interpolation is carried out only at points, lying inside the triangle (black dots), that have non-negative barycentric weights (panel c).

4 EGM with a discrete-continuous state

This section discusses the application of EGM to problems with both discrete and continuous endogenous state variables (DC problems in what follows). The number of discrete variables is immaterial since they can all be stacked into a single state vector. We are going to restrict to a single continuous state variable for tractability.

The model is basically the same as in the previous section with the only difference that now the durable choice is binary $d = \{0, 1\}$. Think of d as a housing size choice. Individuals decide when to upgrade their house. To simplify the notation, we assume the choice to upgrade is irreversible, there is no depreciation and utility from housing services starts accruing in the period the purchase takes place. The cost of the upgrade takes the form of a perpetual mortgage payment of κ units of consumption per period.

The problem of an individual who has upgraded is given by

$$V(a, 1, y) = \max_{a' \geq 0} u(y + Ra - \kappa - a', 1) + \beta EV(a', 1, y'). \quad (11)$$

Given that $d = 1$ is an absorbing state, the value function $V(a, 1, y)$ captures the opti-

mized utility contingent on having upgraded, whether in the current or past periods, and before choosing current savings.

In the case of an individual who starts the current period with $d = 0$, let $V(a, 0, y)$ denote her value contingent on the current choice of not upgrading. The corresponding Bellman equation is

$$V(a, 0, y) = \max_{a' \geq 0} u(y + Ra - a', 0) + \beta EW(a', 0, y') \quad (12)$$

where the continuation value is the expectation of

$$W(a, 0, y) = \max\{V(a, 0, y), V(a, 1, y)\}, \quad (13)$$

the value of choosing optimally between upgrading and not next period after observing the labor income realization.

While $V(a, 0, y)$ and $V(a, 1, y)$ are both continuous and increasing in the continuous state variable a under our maintained assumptions,¹⁶ only $V(a, 1, y)$ is also (strictly) concave and differentiable in a . The upper envelope $W(a, 0, y)$ and, by equation (12), $V(a, 0, y)$ are, in general, neither concave nor differentiable, as $W(a, 0, y)$ has a (“primary”) kink at the level of a for which the individual is indifferent between the two discrete choices. Therefore, the Euler equation

$$-u_c(Ra + y - a', 0) + \beta EW_a(a', 0, y) = 0 \quad (14)$$

is no longer sufficient for the optimal policy $a'(a, 0, y)$ which substantially complicates the solution. We illustrate these and other relevant issues associated with DC problems with the help of Figure 3.

Just for the sake of exposition, assume in what follows that labor income is deter-

¹⁶Monotonicity, strict concavity and differentiability of $V(a, 1, y)$ and continuity and monotonicity in a of $V(a, 0, y)$ and $W(a, 0, y)$ follow from standard results in dynamic programming; e.g. Theorems 9.7 and 9.8 in Stokey, Lucas and Prescott (1989)

ministic and the horizon is finite with N indexing the last period.¹⁷ Working backwards from N , panel (a) in Figure 3 plots the two d -contingent value functions $V^N(a, d, y)$ and their upper envelope $W^N(a, 0, y)$ as functions of a . The upper envelope $W^N(a, 0, y)$ has a kink at the level of wealth for which the discrete choice flips from 0 to 1. As the two d -contingent value functions are continuous and increasing, W^N is non-differentiable and non-concave, its slope increases discretely, at such a kink.

Working backward to period $N - 1$, consider the problem of solving for the optimal saving function contingent on $d = 0$. Panel (b) plots the (absolute values of the) two addenda of the Euler equation (14), as functions of end-of-period wealth a' , for an agent that has not upgraded in the current period. The thick, discontinuous curve is the discounted partial derivative of the continuation value W^N in the left panel. The thin, upward-sloping, curves are the marginal utility of consumption for two different values of beginning-of-period wealth a . For values of initial wealth for which the upward-sloping u_c curves intersect the derivative of the continuation value more than once, the right hand side of equation (12) has multiple local extrema and the Euler equation is not sufficient for a maximum. Yet, the Euler equation is still necessary, all maxima are zeros of the Euler equation, as proved in much greater generality by Clausen and Strub (2020). Intuitively, a point of discontinuity in the Euler equation such as α_6 is a local minimum, not a maximum, as the value of the Euler equation, the vertical difference between the thick line and an upward sloping curve, changes sign from negative to positive.¹⁸

Because DC problems are non-concave and non-differentiable, the typical approach in economics, short of convexifying them by resorting to lotteries or ad hoc uncertainty, is to discretize the continuous state variables and use value function iteration. The above discussion, though, implies that solution methods that rely on the Euler equation can be

¹⁷The finite horizon ensures that the d -contingent optimal saving in the last period N is zero and the value function is strictly concave and differentiable in wealth as it coincides with the utility of spending all cash at hand on non-durable consumption given the respective housing choice. The discussion applies unchanged to solving an infinite horizon problem starting from a strictly concave and differentiable guess V^N for the value function and iterating to convergence.

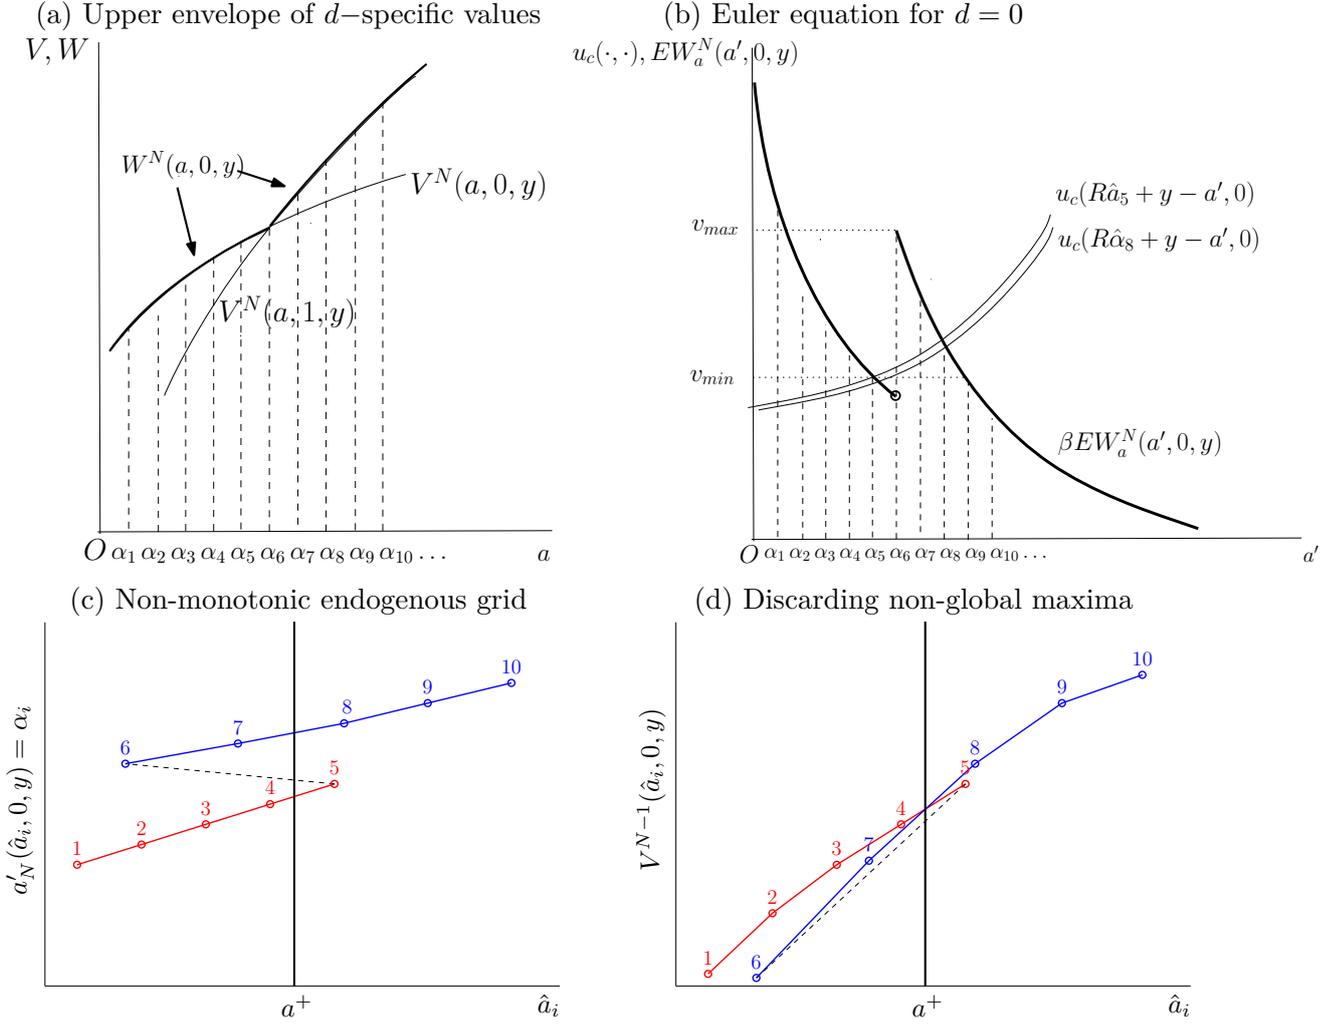
¹⁸To the best of my knowledge the first paper that made this, only *ex post* obvious, point was a 2005 version of Gallipoli and Nesheim (2013). Intuitively, it is never optimal to choose a level of saving which implies indifference between the two discrete choices next period.

employed to locate local maxima, but they have to be supplemented by a “global” step to eliminate solutions which are not a global optimum. Compared to the case with only continuous state variables, the computational costs of solving the Euler equation for a' using EXGM are substantially larger, as the Euler equation is both discontinuous and has multiple zeros as a function of a' . EGM is unaffected in this respect as, for given a' , the Euler equation has a unique zero and remains differentiable, and possibly analytical, in a . It is this insight that motivated the extensions of EGM to this type of problems by Fella (2014) and Iskhakov et al. (2017).

Consider its application to panel (b) of Figure 3 to compute the optimal saving and value functions in period $N-1$ contingent on $d = 0$. By backward induction, the derivative of the continuation value is known at the grid points $\alpha_i \in G_a$.¹⁹ For each α_i , the standard EGM step computes the endogenous value of initial assets \hat{a}_i for which α_i is a local optimum. Graphically, this is the value of \hat{a}_i associated with the upward sloping curve intersecting the thick curve at α_i . Panel (c) in Figure 3 plots the resulting linearly-interpolated correspondence with \hat{a}_i on the horizontal and α_i on the vertical axes. The numbering of the points corresponds to the numbering of grid points in panels (a) and (b). As noted above, the relationship between current assets and the corresponding local maximum in panel (c) is not a function as, for $1 < i < 9$, α_i may not be the only local maximum associated with \hat{a}_i . The correspondence is not even monotonic. In panel (b) the upward sloping u_c curve intersecting βEW_a at grid point α_6 is higher than that through α_5 —consumption is lower at any level of a' —which implies \hat{a}_6 is lower than \hat{a}_5 . It can be proved, though, that the optimal saving correspondence for the problem is monotonic (Proposition 1 in Fella (2014) and Theorem 2 in Iskhakov et al. (2017)). It follows that the optimal saving correspondence is discontinuous and jumps from the lower to the upper leg of the correspondence in panel (c) for some level of initial assets between \hat{a}_6 and \hat{a}_5 . Fella (2014) and Iskhakov et al. (2017) propose two alternative ways to discard suboptimal points generated by the EGM step. Panel (d) in Figure 3 illustrates the more

¹⁹Its value at α_6 is assumed to be v_{max} .

Figure 3: DC-EGM in period $N - 1$



Notes: The plots illustrate how endogenous points corresponding to local, but not global, optima for the saving function of individual with $d = 0$ in period $N - 1$ are eliminated in Iskhakov et al.'s (2017) DC-EGM algorithm. Panel (a) plots the continuation value in the Bellman maximand. Panel (b) plots its derivative (thick black curve) and that of the utility of consumption (thin line) for different values of initial wealth a , both as a function of end-of-period wealth a' . The endogenous grid point \hat{a}_i , computed by the standard EGM, ensures that the Euler equation is satisfied at $a' = \alpha_i$. Panel (c) plots the resulting correspondence from the endogenous grid points (on the horizontal axis) to the associated value of a' . Panel (d) plots the associated value correspondence, the right hand side of the Bellman equation, and shows how its upper envelope is used to discard non-global optima.

elegant and faster method in the latter paper. The panel plots the interpolated values of the right hand side of equation (12) at the pairs (\hat{a}_i, α_i) in panel (c). Since the value function of the original problem is increasing in initial wealth, its approximation on the grid in panel (d) is the upper envelope of the overlapping red and blue segments. The algorithm discards points that do not belong to the upper envelope in the following way.

it loops over the index i until it finds (a) the first grid point j such that $\hat{a}_{j+1} < \hat{a}_j$ and (b) the first grid point $k > j$ such that $\hat{a}_j < \hat{a}_k$. It then computes the upper envelope between the segments associated respectively with points $\{1, j\}$ and $\{j - 1, k\}$, discards the dominated points $i = \{j, \dots, k - 1\}$ and continues until the next instance. Intuitively, the upper the upper envelope can constructed from the piecewise linear interpolants of each of the two overlapping segment.²⁰ Faster implementations involve more technical details which are beyond the scope of this survey.²¹ The algorithm in Dobrescu and Shanker (2023) is the fastest I am aware of and quite intuitive despite involving some technical test conditions.

Finally, to improve the accuracy of the solution, and avoid interpolation around the discontinuity in the saving function, the refined grid can be supplemented by the approximated location of the discontinuity a^+ obtained as the intersection of the two lines going respectively through points $k - 1$ and k and $j - 1$ and j in panel (d).²² In the example in Figure 3 the resulting refined endogenous grid of interpolating nodes would be given $\{\hat{a}_1, \hat{a}_4, a^+ - \nu, a^+ + \nu, \hat{a}_8, \hat{a}_9, \hat{a}_{10}\}$ with ν an appropriately small number which is necessary as most interpolating algorithms do not accept grids with identical abscissas. The associated optimal saving values at $a^+ \pm \nu$ are obtained by interpolation. For example, in the case of $a^+ + \nu$ one could extrapolate using points 8 and 9. Having solved for the optimal saving function in period $N - 1$ contingent on $d = 0$ one obtains the associated value function $V^{N-1}(a, 0, y)$ from the Bellman equation (12) and continues iterating backward.²³

²⁰This is the approach in Bueren (2025) who exploits the envelope theorem to compute the slope of each segment, the derivative of the value function, from the marginal utility of (the known) consumption at the grid points.

²¹Iskhakov et al. (2017) too do not describe their upper envelope algorithm.

²²An alternative, that does not require keeping track of the dominated points $i = \{j, \dots, k - 1\}$, exploits the insight of Bueren (2025). One can obtain a^+ as the intersection of the two lines through $j - 1$ and k with respective slopes given by the marginal utility of consumption at points $j - 1$ and k .

²³The algorithm in Fella (2014) instead identifies the subset of the grid for a' where the Euler equation may have multiple local maxima; $A^{nc} = \{\alpha_2, \alpha_7\}$ in panel (b). For points $\alpha_i \in A^{nc}$ it uses the standard EGM step to compute the associated value of initial assets \hat{a}_i . It then verifies that α_i is a global maximum by maximizing the discretized right hand side of equation (12) given $a = \hat{a}_i$ with respect to $a' \in A^{NC}$. If, upon evaluating e.g. point α_5 in panel (b), it finds that the local maximum for the discretized problem is α_8, α_6 and α_7 can be discarded by monotonicity without the need for the EGM step.

One feature of DC problems is that kinks propagate backwards. This can be seen in our example. The primary kink in W^N at the level of wealth $a = \alpha_6$ for which the consumer is indifferent between the two discrete choices generates a “secondary” kink in the d -specific value function $V^{N-1}(a, 0, y)$ at $a = a^+$, the value of wealth for which saving in period N jumps up and the future discrete choice switches. It follows that the upper envelope $W^{N-1}(a, 0, y)$ will likely have two kinks: a primary one at the level of wealth for which $V^{N-1}(a, 0, y)$ and $V^{N-1}(a, 1, y)$ cross and a secondary one at a^+ . It should be clear that in the absence of uncertainty the number of kinks increases going backward. Note that this is a feature of the true model solution and constitutes a significant challenge for the numerical solution of DC problems. As is well known, the presence of uncertainty, for example in the form of stochastic income shocks (Gomes, Greenwood and Rebelo, 2001), introduces convexity and can smooth kinks. Fella and Gallipoli (2014) use EGM to solve a life-cycle model with an AR(1) labour income process and a discrete crime choice in each working year. In each period, the number of secondary (back-propagating) kinks never exceeds four despite a realistic working life of 40 years. Iskhakov et al. (2017) propose adding additive IID Extreme Value²⁴ taste shocks to improve the reliability of the EGM algorithm in DC problems where the inbuilt uncertainty provides insufficient smoothing. They show that the variance of such shocks can be chosen so that the perturbed model approximates the original model with an arbitrary degree of precision.²⁵

Fella (2014) and Iskhakov et al. (2017) show that, with the same number of grid points, DC-EGM is one to two orders of magnitudes faster and three orders of magnitude more accurate than value function iteration, the most commonly used algorithm to solve DC problems. Figure 4, adapted from Fella (2014), compares the performance of DC-EGM (top row) vs VFI²⁶ (bottom row) in solving a DC problem similar to the one discussed in

²⁴The advantage of Extreme Value shocks is that they imply that the expectation of the value function with respect to the shocks is available in closed form.

²⁵For example one can choose the variance of taste shocks so as to smooth secondary, but not primary, kinks.

²⁶To make the comparison as fair as possible, the version of VFI employed is more accurate than the discretized one typically used (e.g. Rust, 1987) in solving DC problems. It uses grid-search to bracket the maximum over the discrete grid and then switches to linear interpolation of the saving function on the bracketing interval.

this section, with the only difference that the durable choice can take seven values and is not irreversible and non-capital income is the sum of an AR(1) and white noise processes. From left to right, each column corresponds to a different grid size—200, 400 and 1000 points—for the continuous state variable. The figure reports execution time, average and maximum Euler error and plots the size distribution and the average of the Euler errors (red line) over a 50,000 observations simulated history.

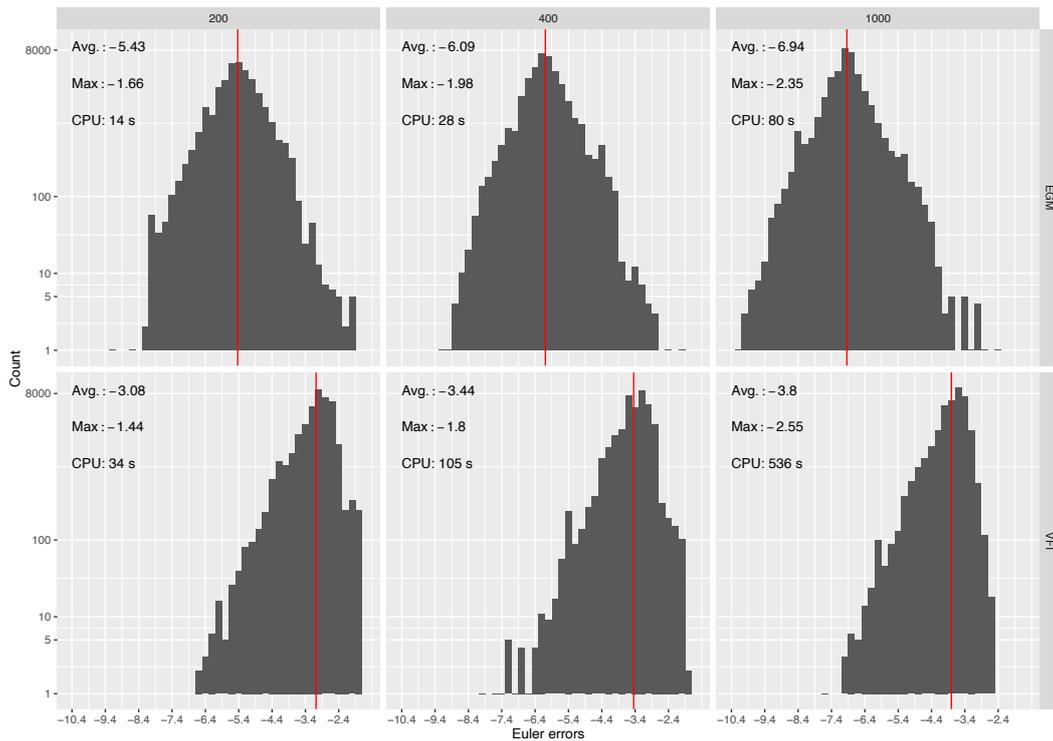
In terms of accuracy, not only EGM produces an average Euler error two to three orders of magnitude smaller than VFI, *for the same grid size*. The average approximation error of EGM with 200 grid points is nearly two order of magnitude smaller than VFI with 1000 grid points, against a difference in computational time of about 40 times. In fact, less than 0.1 per cent (less than 40 out of more than 40,000 observations off corners) of the Euler errors of EGM with 200 grid points lie to the right of the *average* Euler error of VFI with 1000 points.²⁷ The dramatic difference in accuracy between DC-EGM and VFI, compared to the much smaller difference in Table 1 in Section 2.1, is due to the fact that in Section 2.1 both solution methods, EGM and time iteration, solve for the zeros of the Euler equation. Instead VFI is less accurate since it maximises the right hand side of the Bellman operator without using the Euler equation. As we have discussed above, though, the computational costs of solving the Euler equations for a' using EXGM are too large for it to be a viable alternative in DC problems.²⁸

DC-EGM can also be extended to models with risky borrowing and endogenous default. Contrary to incomplete market model with only risk-free debt, though, models with risky borrowing imply that the borrowing limit is a state-dependent, equilibrium object. Jang and Lee (2024) show how the DC-EGM algorithm can be adapted to deal

²⁷Comparing the error distributions suggests that the maximum Euler error is not very informative about the right tail of the error size distribution. This can be understood in light of the fact that in DC problems the true saving function is discontinuous. Independently from the solution algorithm used, as long as the simulation involves interpolating between two points bracketing a discontinuity in the true policy function (e.g. point 4 and 8 in panel (c) in Figure 3) the Euler equation evaluated at the approximate solution may be significantly violated.

²⁸Note that both DC-EGM and VFI use the distance between subsequent iterates of the value function as their convergence metric. The difference is just in the way the two approaches solve the inner loop for the optimal saving function.

Figure 4: Accuracy and speed of DC-EGM vs VFI. Source Fella (2014).



Notes: Avg. and Max are respectively the average and maximum Euler error (in log 10 units), off corners, on a stochastic simulation of 50,000 observations. CPU is CPU time in seconds. Both algorithms were written in Fortran 95 and executed on a single core of a Xeon X5570 processor.

with this complication. They exploit the theoretical result in Arellano (2008) that the zero-profit condition for credit intermediaries implies a mapping from the quantity of debt to the value (price times quantity) of debt that follows an inverted Laffer curve. In equilibrium, a borrower would never choose to borrow more than the minimum of the Laffer curve because there is an alternative contract that increases consumption today by the same amount while entailing a smaller liability next period. Therefore the borrowing limit in each income state coincides with minimum of the Laffer curve in that state. At each iteration, the algorithm computes the zero profit condition given the current guess for the default policy function and locates its minimum. Fella's (2014) DC-EGM is then used to solve for the saving function.

Although this section has confined attention to DC problems with a single continuous state variables, Druedahl and Jørgensen (2017) discuss how a combination of DC-EGM

with their triangulation procedure can accommodate non-concave and non-differentiable problems with multiple continuous state variables. They provide necessary and sufficient conditions on the model primitives that characterise the class of models to which their algorithm is applicable.

5 Conclusion

The endogenous gridpoint method has proven an important alternative to conventional numerical techniques in solving dynamic optimization problems. It leverages the insight that the Euler equation is a known, and often linear function, of the pre-decision continuous state. Therefore, fixing the value of post-decision states and solving the Euler equation for the value of pre-decision states significantly reduces interpolation errors and computational costs.

The wide variety of EGM implementations documented in the literature—and surveyed here—demonstrates the method’s flexibility and robustness across a large class of models. As this survey illustrates, extending the basic EGM framework to richer environments is both possible and advantageous, though not without challenges. As is often the case, the main trade-off is between computing and user programming time. Using EGM in the context of models with multiple state variables requires users to grapple with technical issues, such as interpolation on non-rectangular grids or eliminating dominated local but not global maxima, which typically require substantial investment in adapting available codes to the particular problem at hand. Exogenous grid methods, such as value function iteration, are much more off-the-shelf. The computational and accuracy advantages of EGM, though, are particularly large for problems with both continuous and discrete state variables. Unlike EGM, exogenous grid methods that rely on the Euler equation are computationally too costly for this class of problems, leaving VFI as the only, slower and less accurate, alternative. In these instances, resorting to EGM is often the efficient choice particularly when the model solution is part of a structural estimation

and, a fortiori, when the programming language used is a non-compiled one.

In sum, the endogenous gridpoint method stands as a key breakthrough for efficiently solving dynamic optimization problems. Its core advantage—efficiently solving the Euler equation minimizing the reliance on costly root-finding during each iteration—has helped researchers tackle increasingly realistic and high-dimensional models. I hope this survey serves not only as a reasonably-accessible introduction for newcomers but also as an invitation for experienced modelers to consider its use when desirable.

References

- Ameriks, John, Briggs, John, Caplin, Andrew, Shapiro, Matthew D. and Tonetti, Christopher (2020), ‘Long-term care utility and late-in-life saving’, *Journal of Political Economy* **128**(6), 2375–2451.
- Arellano, Cristina (2008), ‘Default risk and income fluctuations in emerging economies’, *American Economic Review* **98**(3), 690–712.
- Bajari, Patrick, Chan, Phoebe, Krueger, Dirk and Miller, Daniel (2013), ‘A dynamic model of housing demand: Estimation and policy implications’, *International Economic Review* **54**(2), 409–442.
- Barillas, Francisco and Fernández-Villaverde, Jesús (2007), ‘A generalization of the endogenous grid method’, *Journal of Economic Dynamics and Control* **31**, 2698–2712.
- Bayer, Christian, Luetticke, Ralph, Pham-Dao, Lien and Tjaden, Volker (2019), ‘Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk’, *Econometrica* **87**(1), 255–290.
- Bellman, Richard (1961), *Adaptive Control Processes: A Guided Tour*, Princeton University Press, Princeton, NJ.

- Bueren, Jesús (2025), A simple algorithm for identifying non-optimal points in non-concave problems. mimeo.
- Carroll, Christopher D. (2006), ‘The method of endogenous gridpoints for solving dynamic stochastic optimization problems’, *Economics Letters* **91**(3), 312–320.
- Clausen, Andrew and Strub, Carlo (2020), ‘Reverse calculus and nested optimization’, *Journal of Economic Theory* **187**, 105019.
- De Nardi, Mariacristina, Fella, Giulio and Paz-Pardo, Gonzalo (2024), ‘Wage risk and government and spousal insurance’, *The Review of Economic Studies* p. rdae042.
- Dobrescu, Loretta I. and Shanker, Akshay (2023), A fast upper envelope scan method for discrete-continuous dynamic programming. W.P. 2023/044 ARC Centre of Excellence in Population Ageing Research (CEPAR).
- Druedahl, Jeppe and Jørgensen, Thomas Høgholm (2017), ‘A general endogenous grid method for multi-dimensional models with non-convexities and constraints’, *Journal of Economic Dynamics and Control* **74**, 87–107.
- Druedahl, Jeppe and Jørgensen, Casper Nordal (2018), ‘Precautionary borrowing and the credit card debt puzzle’, *Journal of Economic Behavior & Organization* **152**, 1–14.
- Ejrnæs, Mette and Jørgensen, Thomas H. (2020), ‘Family planning in a life-cycle model with income risk’, *Journal of Applied Econometrics* **35**(5), 567–586.
- Fella, Giulio (2014), ‘A generalized endogenous grid method for non-smooth and non-concave problems’, *Review of Economic Dynamics* **17**, 329–344.
- Fella, Giulio and Gallipoli, Giovanni (2014), ‘Education and crime over the life cycle’, *The Review of Economic Studies* **81**(4), 1484–1517.
- Gallipoli, Giovanni and Nesheim, Lars (2013), Non-convexities in dynamic programming problems. mimeo.

- Gomes, Joao, Greenwood, Jeremy and Rebelo, Sergio (2001), ‘Equilibrium unemployment’, *Journal of Monetary Economics* **48**(1), 109–152.
- Guerrieri, Veronica and Lorenzoni, Guido (2017), ‘Credit crises, precautionary savings, and the liquidity trap*’, *The Quarterly Journal of Economics* **132**(3), 1427–1467.
- Hallengreen, Adam, Jørgensen, Thomas H. and Olesen, Annasofie M. (2024), The endogenous grid method without analytical inverse marginal utility. Mimeo.
- Hintermaier, Thomas and Koeniger, Winfried (2010), ‘The method of endogenous gridpoints with occasionally binding constraints among endogenous variables’, *Journal of Economic Dynamics and Control* **34**(10), 2074 – 2088.
- Iskhakov, Fedor (2015), ‘Multidimensional endogenous gridpoint method: Solving triangular dynamic stochastic optimization problems without root-finding operations’, *Economics Letters* **135**, 72–76.
- Iskhakov, Fedor and Keane, Michael (2016), ‘An analysis of the Australian social security system using a life-cycle model of labor supply with asset accumulation and human capital’, *Quantitative Economics* **7**(3), 699–750.
- Iskhakov, Fedor, Jørgensen, Thomas H, Rust, John and Schjerning, Bertel (2017), ‘The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks’, *Quantitative Economics* **18**(2), 317–365.
- Jang, Youngsoo and Lee, Soyoung (2024), ‘A generalized endogenous grid method for default risk models’, *Journal of Money, Credit and Banking* **n/a**(n/a).
- Judd, Kenneth L. (1998), *Numerical Methods in Economics*, MIT Press, Cambridge, MA.
- Judd, Kenneth L., Maliar, Lilia, Maliar, Serguei and Tsener, Inna (2017), ‘How to solve dynamic stochastic models computing expectations just once’, *Quantitative Economics* **8**(3), 851–893.

- Ludwig, Alexander and Schön, Matthias (2018), ‘Endogenous grids in higher dimensions: Delaunay interpolation and hybrid methods’, *Computational Economics* **51**, 463–492.
- Mazzocco, Maurizio (2007), ‘Household intertemporal behaviour: A collective characterization and a test of commitment’, *The Review of Economic Studies* **74**(3), 857–895.
- Powell, Warren B (2007), *Approximate Dynamic Programming: Solving the curses of dimensionality*, Vol. 703, John Wiley & Sons.
- Rust, John (1987), ‘Optimal replacement of gmc bus engines: An empirical model of harold zurcher’, *Econometrica* **55**(5), 999–1033.
- Stokey, Nancy L., Lucas, Robert E. and Prescott, Edward C. (1989), *Recursive Methods in Economic Dynamics*, Harvard University Press, Cambridge, MA.
- Van Roy, Benjamin, Bertsekas, Dimitri P., Lee, Yoon and Tsitsiklis, John N. (1997), A neuro-dynamic programming approach to retailer inventory management, in ‘Proceedings of the 36th IEEE Conference on Decision and Control’, San Diego, California, pp. 4052–4057.
- White, Matthew N (2015), ‘The method of endogenous gridpoints in theory and practice’, *Journal of Economic Dynamics and Control* **60**, 26–41.
- Wright, Brian D. and Williams, Jeffrey C. (1982), ‘The roles of public and private storage in managing oil import disruptions’, *The Bell Journal of Economics* **13**(2), 341–353.
- Wright, Brian D. and Williams, Jeffrey C. (1984), ‘The Welfare Effects of the Introduction of Storage’, *The Quarterly Journal of Economics* **99**(1), 169–192.
- Yao, Jing, Fagereng, Andreas and Natvik, Gisle (2021), ‘Housing, debt and the marginal propensity to consume’, *Journal of Political Economy* **129**(12), 3369–3416.