The Endogenous Grid Method for Epstein-Zin Preferences

Alan Lujan[1]
Johns Hopkins University
Econ-ARK

Wednesday 7th January, 2026

Abstract

The endogenous grid method (EGM) accelerates dynamic programming by inverting the Euler equation, but it appears incompatible with Epstein-Zin preferences where the value function enters the Euler equation. This paper shows that a power transformation resolves the difficulty. The resulting algorithm requires no root-finding, achieves speed gains of one to two orders of magnitude over value function iteration, and improves accuracy by more than one order of magnitude. Holding accuracy constant, the speedup is two to three orders of magnitude. VFI and time iteration face a speed-accuracy tradeoff; EGM sidesteps it entirely.

Keywords Endogenous grid method, Epstein-Zin preferences, Recursive utility, Dynamic programming, Consumption-savings

1 Introduction

The endogenous grid method (EGM) of Carroll [2006] inverts the Euler equation to obtain consumption as a function of end-of-period assets, eliminating root-finding and accelerating computation by orders of magnitude. EGM is now widely used in structural estimation of consumption-savings problems, including models of precautionary saving and liquidity constraints in the tradition of Deaton [1991]. Barillas and Fernández-Villaverde [2007] extend the method to multiple controls; Iskhakov et al. [2017] handle discrete-continuous choice.

Separately, Epstein-Zin preferences, developed by Epstein and Zin [1989], Epstein and Zin [1991], and Weil [1989], have become standard in macro-finance for their ability to separate risk aversion from intertemporal substitution. Building on the temporal lottery framework of Kreps and Porteus [1978], these preferences are theoretically well-grounded: Duffie and Epstein [1992] establish existence and uniqueness in continuous time; Ren and Stachurski [2020] provide modern discrete-time foundations. Yet EGM has not been applied to Epstein-Zin preferences. The apparent obstacle is that the value function enters the Euler equation. Existing approaches based on time iteration, such as Coeurdacier et al. [2020], treat consumption, value, and expected value as simultaneous control variables, requiring numerical root-finding on a system of non-linear equations at every grid point in every iteration.

[1]Correspondence to: alujan@jhu.edu

This paper shows that such computational expense is unnecessary. A power transformation converts the Epstein-Zin recursion into a form where EGM applies directly. The algorithm tracks two functions instead of one (consumption and transformed value) but requires no root-finding and maintains the standard EGM structure. The method extends to other recursive preferences used in dynamic models, including the risk-sensitive preferences of Hansen and Sargent [1995].

2 Model

Definition 2.1 (Consumption-Savings Problem with Epstein-Zin Preferences)**.** *An agent solves*

$$V(m, z) = \max_{c} \left[(1-\beta)c^{1-\rho} + \beta\left(\mathbb{E}_{z'|z}\left[V(m', z')^{1-\gamma}\right]\right)^{\frac{1-\rho}{1-\gamma}}\right]^{\frac{1}{1-\rho}} \tag{1}$$

subject to $a = m - c \geq 0$ and $m' = Ra + y(z')$, where m is cash-on-hand, c consumption, a end-of-period assets, R the gross interest rate, z the current income state, and $y(z')$ income as a function of the realized state z'.

The state z follows an AR(1) process with transition probabilities $\pi_{k\ell} = \Pr(z' = z_\ell \mid z = z_k)$, so the expectation $\mathbb{E}_{z'|z}$ is taken over next-period states conditional on the current state. The parameter ρ governs the elasticity of intertemporal substitution (EIS $= 1/\rho$), while γ governs risk aversion.[1] When $\rho = \gamma$, preferences collapse to standard CRRA expected utility.

We assume $\rho \neq 1$ and $\gamma \neq 1$ for the power transformation developed below; the limiting cases use logarithmic transformations but the method is otherwise analogous. For the infinite-horizon problem, standard conditions ensure the value function is well-defined and the consumption policy converges.[2]

2.1 Transformation

Definition 2.2 (Transformed Value Function)**.** *The transformed value function is $W(m, z) \equiv V(m, z)^{1-\rho}$.*

This transformation simplifies the certainty equivalent to a power mean.[3] Meyer-Gohde [2019] use the same transformation to analyze model uncertainty with recursive preferences.

Definition 2.3 (Certainty Equivalent)**.** *Let $\theta \equiv (1-\gamma)/(1-\rho)$. The certainty equivalent of next-period transformed value is*

$$\mu(a, z) \equiv \left(\mathbb{E}_{z'|z}\left[W(m', z')^{\theta}\right]\right)^{1/\theta} \tag{2}$$

where $m' = Ra + y(z')$.

The case $\theta = 1$ (equivalently $\gamma = \rho$) recovers time-separable CRRA expected utility; when $\gamma > \rho$, the agent prefers early resolution of uncertainty; when $\gamma < \rho$, late resolution.[4]

[1] We follow the original Epstein and Zin [1989] parameterization. An alternative convention, common in the asset pricing literature (e.g., Bansal and Yaron [2004]), defines the EIS directly as $\psi = 1/\rho$ and writes the recursion with ψ in place of $1/\rho$. The auxiliary parameter is then $\theta = (1-\gamma)/(1-1/\psi)$, which equals our $(1-\gamma)/(1-\rho)$.

[2] See Epstein and Zin [1989] for existence and uniqueness conditions. With $\theta = (1-\gamma)/(1-\rho)$, a sufficient condition is $\beta R^{\theta} < 1$. When $\theta < 0$ (as in our calibration with $\gamma > 1 > \rho$), this becomes $\beta < R^{|\theta|}$, which holds for typical parameterizations. The recursion also requires $V > 0$; this follows by induction from the terminal condition whenever consumption remains strictly positive.

[3] When $\rho > 1$, the transformation $W = V^{1-\rho}$ is decreasing, so maximizing V is equivalent to minimizing W. In the discussion below, the Euler equation characterizes the optimum in either case, and EGM inverts it directly. See Appendix: Robustness to $\rho > 1 fordetails.$

[4] See Backus et al. [2008] for a comprehensive treatment.

Raising the Bellman equation in Definition 2.1 to the power $1-\rho$ yields the transformed Bellman equation:

$$W(m, z) = \max_c \left[(1-\beta)c^{1-\rho} + \beta\mu(m-c, z)\right] \tag{3}$$

The transformation succeeds because it converts the CES aggregator in Definition 2.1 into an additive structure. The original Bellman equation involves V raised to fractional powers both inside and outside the expectation; the transformed equation (3) is additive, with the nonlinearity isolated in μ. This additive structure admits clean differentiation, yielding an Euler equation that depends on W and c but can be inverted for c in closed form.

The first-order condition equates the marginal utility of consumption to the marginal value of savings:

$$(1-\beta)(1-\rho)c^{-\rho} = \beta\frac{\partial\mu(a, z)}{\partial a} \tag{4}$$

The envelope theorem, combined with the first-order condition, gives

$$\frac{\partial W(m, z)}{\partial m} = (1-\beta)(1-\rho)c(m, z)^{-\rho} \tag{5}$$

Differentiating $\mu(a, z)$ with respect to a yields:

$$\frac{\partial\mu}{\partial a} = R\cdot\mu(a, z)^{1-\theta}\cdot\mathbb{E}_{z'|z}\left[W(m', z')^{\theta-1}\cdot\frac{\partial W(m', z')}{\partial m'}\right] \tag{6}$$

Substituting (6) into the FOC (4) and applying the envelope condition (5) to the next-period marginal value, the Euler equation becomes

$$c^{-\rho} = \beta R\cdot\mu(a, z)^{1-\theta}\cdot\Xi(a, z) \tag{7}$$

where

$$\Xi(a, z) \equiv \mathbb{E}_{z'|z}\left[W(m', z')^{\theta-1}\cdot c(m', z')^{-\rho}\right] \tag{8}$$

Given $W(\cdot, z')$ and $c(\cdot, z')$ for all z', both $\mu(a, z)$ and $\Xi(a, z)$ depend only on end-of-period assets a and the current state z. Inverting the Euler equation yields consumption as a function of (a, z):

Proposition 2.1 (Inverted Euler Equation)**.** *The Euler equation for Epstein-Zin preferences can be inverted to yield*

$$c(a, z) = \left(\beta R\cdot\mu(a, z)^{1-\theta}\cdot\Xi(a, z)\right)^{-1/\rho} \tag{9}$$

This closed-form expression is what enables EGM: given an exogenous grid over a, we compute $c(a, z)$ directly and recover $m = c + a$.

Remark 2.1 (Limiting Cases)**.** *When $\rho \to 1$, the transformation becomes $W = \log V$ and the certainty equivalent becomes $\mu(a, z) = \frac{1}{1-\gamma}\log\mathbb{E}_{z'|z}[\exp((1-\gamma)W(m', z'))]$, yielding the risk-sensitive preferences of Hansen and Sargent [1995]; Tallarini Jr [2000] demonstrates the equivalence in business cycle models. When $\gamma \to 1$, the certainty equivalent of V' becomes the geometric mean $\exp(\mathbb{E}_{z'|z}[\log V(m', z')])$; Giovannini and Weil [1989] show that this case yields myopic portfolio allocation.*[5]

3 Algorithm

One distinction from standard EGM emerges from the structure of Proposition 2.1: the Euler equation requires both the policy function $c(\cdot, z')$ and the value function $W(\cdot, z')$ evaluated at next-period states. In standard CRRA problems, the Euler equation depends only on $c(\cdot, z')$, so EGM can iterate on the policy function alone. Here, we must iterate until both functions converge simultaneously.

[5]The logarithmic limit requires renormalization: formally, $\lim_{\rho\to 1}(V^{1-\rho}-1)/(1-\rho) = \log V$ by L'Hôpital's rule. The Bellman equation and certainty equivalent formulae stated here are the limiting forms after this renormalization.

This makes the algorithm a hybrid of EGM and the Howard policy iteration technique introduced by Howard [1960]. EGM provides the root-finding-free policy update, while tracking $W(\cdot)$ alongside $c(\cdot)$ ensures the value function remains consistent with the policy. Each iteration updates both functions, and convergence requires both to stabilize.

Remark 3.1 (Storing V instead of W)**.** *While the algorithm is derived using $W = V^{1-\rho}$, numerical implementation benefits from storing and interpolating V directly, converting to W only when needed. When $\rho > 1$, the transformation $W = V^{1-\rho}$ maps small values of V to large values of W (since the exponent is negative), making W less stable for interpolation near the borrowing constraint. When $\rho < 1$, both are well-behaved. In either case, V is the natural object to store. The algorithm below stores (c, V), converts $V \to W$ for the Euler computation, and converts back after updating. The algorithm works for both $\rho < 1$ (EIS > 1) and $\rho > 1$ (EIS < 1): although θ changes sign at $\rho = 1$, the Euler equation and certainty equivalent formulas remain valid.*[6]

The algorithm proceeds as follows: Fix a grid of end-of-period asset values $\{a_j\}_{j=1}^J$ and an exogenous grid of cash-on-hand values $\{m_i\}_{i=1}^I$. Initialize with a policy such as $c^{(0)}(m, z) = m$ and a monotonically increasing, concave $V^{(0)}$ (e.g., $V^{(0)} = c^{(0)}$); the iteration will converge to the fixed point. For finite-horizon problems, the terminal condition is $c(m, z) = m$; for infinite-horizon, iterate until convergence.

Algorithm 1 EGM for Epstein-Zin

Require: Grid of end-of-period assets $\{a_j\}_{j=1}^J$, exogenous grid $\{m_i\}_{i=1}^I$, initial guess $(c^{(0)}, V^{(0)})$
Ensure: Converged policy $c(\cdot,\cdot)$ and value $V(\cdot,\cdot)$

```
Given (c^(n), V^(n)) from iteration n, for each income state z_k:
for each a_j do
    Compute m'_{jℓ} = R a_j + y(z'_ℓ) for all z'_ℓ
    Interpolate c^(n)(m'_{jℓ}, z'_ℓ) and V^(n)(m'_{jℓ}, z'_ℓ); compute W^(n) = (V^(n))^{1-ρ}
    Compute μ_j = ( Σ_ℓ π_{kℓ} W^(n)(m'_{jℓ}, z'_ℓ)^θ )^{1/θ}
    Compute Ξ_j = Σ_ℓ π_{kℓ} [ W^(n)(m'_{jℓ}, z'_ℓ)^{θ-1} · c^(n)(m'_{jℓ}, z'_ℓ)^{-ρ} ]
    Invert Euler: c_j = ( βR · μ_j^{1-θ} · Ξ_j )^{-1/ρ}
    Recover grid: m_j = c_j + a_j
end for
Append (0, 0) at the constraint
Interpolate c^(n+1)(·, z_k) from {(m_j, c_j)} to {m_i}
for each m_i do
    Compute a_i = m_i - c^(n+1)(m_i, z_k)
    Interpolate μ(a_i, z_k) from {μ_j}
    Update: W^(n+1)(m_i, z_k) = (1-β) c^(n+1)(m_i, z_k)^{1-ρ} + β μ(a_i, z_k)
    Convert: V^(n+1)(m_i, z_k) = (W^(n+1)(m_i, z_k))^{1/(1-ρ)}
end for
Iterate until ||c^(n) - c^(n-1)|| < ε
```

The algorithm requires no root-finding: each iteration involves only interpolation, expectation computation, and closed-form inversions. In contrast, time iteration methods must solve a non-linear system at each grid point; here, the Euler inversion yields consumption directly. The cost per iteration is comparable to standard EGM; the only addition is tracking the value function alongside consumption. Convergence is monitored

[6] See Appendix: Robustness to $\rho > 1fornumericalverification.$

via $\|c^{(n)} - c^{(n-1)}\|$; since V is uniquely determined by c through the Bellman equation, policy convergence implies value convergence.

To handle the borrowing constraint, we augment the endogenous grid with the point $(m, c) = (0, 0)$ to anchor interpolation. For typical calibrations, the resulting grid is monotonic and interpolation proceeds directly. If non-monotonicity occurs (i.e., $m_j > m_{j+1}$ for some j despite $a_j < a_{j+1}$), construct the consumption function via the upper envelope following Iskhakov et al. [2017]. In the constrained region where $a = 0$ binds, $c(m, z) = m$ and $W(m, z) = (1 - \beta)m^{1-\rho} + \beta\mu(0, z)$.

4 Benchmarks

We implement the algorithm for an infinite-horizon consumption-savings problem with stochastic income. Income follows an AR(1) process with persistence 0.95, consistent with estimates from labor market data (e.g., Storesletten et al. [2004]), discretized using the Tauchen [1986] method with 10 grid points. The asset grid uses 100 exponentially-spaced points with an upper bound of 20 times mean income.

Parameters. $\beta = 0.96$, $R = 1.02$, $\gamma = 10$, and EIS $= 1.5$ (so $\rho = 2/3$), following Bansal and Yaron [2004]. The meta-analysis of Havránek [2015] finds that micro estimates of the EIS often fall below unity, though macro and asset pricing calibrations typically use higher values.[7] Since $\gamma > \rho$, the agent prefers early resolution of uncertainty. The auxiliary parameter is $\theta = (1 - \gamma)/(1 - \rho) = -27$.

4.1 Speed

We compare three solution methods: EZ-EGM, time iteration (TI), and VFI. Rendahl [2015] proves that TI and VFI converge to the same solution under standard conditions, so the comparison is one of computational efficiency. For TI and VFI, we distinguish between *fast* mode (precompute μ on the asset grid, interpolate during search) and *accurate* mode (compute μ exactly at each candidate).[8] EGM has no such distinction: it evaluates μ on the endogenous grid, achieving accurate-mode precision without the computational cost.

Table 1: Speed comparison (fast modes)

Method	Time (ms)	Policy Iters	Euler Error
EZ-EGM	21	141	-4.8
TI (fast)	237	140	-3.6
VFI (fast)	1162	239	-3.3

Note: Euler error is mean $\log_{10}$ error on the ergodic distribution (more negative = more accurate). Fast modes precompute μ on the asset grid. Timings on CPU; results vary by hardware. All methods use JAX.

The speed advantage of EZ-EGM comes from eliminating numerical search. Time iteration (TI), introduced by Coleman [1991] and extended to Epstein-Zin preferences by Coeurdacier et al. [2020], requires bisection at every grid point; VFI requires golden-section search. EZ-EGM bypasses this entirely through analytic Euler inversion. At equal grid size (100 points), EZ-EGM is approximately 50 times faster than VFI and 10 times faster than TI. Moreover, EGM is more than one order of magnitude more accurate than both. Figure 1 shows that the resulting policy functions are smooth and well-behaved.

[7] High γ with low ρ implies aversion to contemporaneous consumption risk but tolerance for intertemporal variation. Whether this reflects preferences or serves as a modeling device for asset pricing remains debated; the method here is agnostic on calibration.

[8] Appendix: Alternative algorithms details both algorithms and the tradeoff between modes.

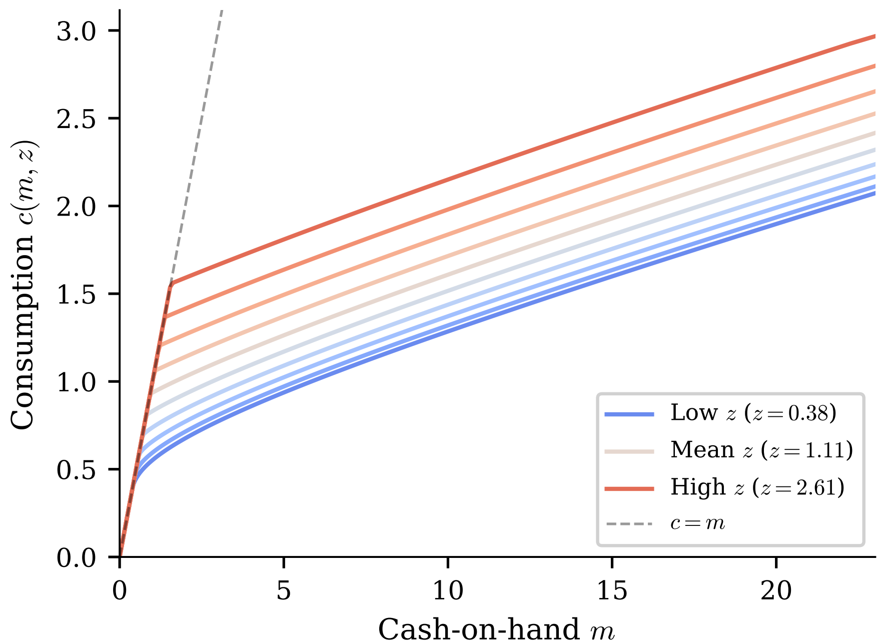

Figure 1: Consumption policy $c(m, z)$ for different income states. Higher income states (red) allow more consumption at each wealth level.

4.2 Accuracy

We assess accuracy using the normalized Euler equation error, a standard metric whose magnitude bounds the policy function error. Santos [2000] establishes the theoretical foundation:

$$\varepsilon(m, z) = \log_{10} \left| 1 - \frac{\tilde{c}(m, z)}{c(m, z)} \right| \tag{10}$$

where $\tilde{c}(m, z)$ is consumption implied by the Euler equation given the computed policy. We report mean (L1) and maximum (L∞) errors following Maliar and Maliar [2014]. More negative values indicate higher accuracy; an error of -5 means the policy deviates from the Euler equation by approximately 10^{-5}, or 0.001%.

Errors should be evaluated at wealth levels agents actually visit, not arbitrary grid points. We simulate the ergodic distribution (10,000 agents, 500 periods) and compute errors at the 5th–95th percentiles of realized wealth. For this calibration, the ergodic distribution concentrates in the lower wealth range (median $m \approx 3$ versus grid maximum of 20), reflecting moderate impatience. Errors evaluated on the ergodic distribution are similar to those on a uniform grid, indicating the solution is accurate where agents spend time.[9]

Table 2: Accuracy comparison (accurate modes)

Method	Time (ms)	Policy Iters	Mean (L1)	Max (L∞)
EZ-EGM	21	141	−4.8	−3.2
TI (accurate)	2301	141	−4.8	−3.5
VFI (accurate)	8293	239	−3.4	−2.2

Note: Errors evaluated on ergodic distribution. Accurate modes compute μ exactly during search. TI-accurate matches EGM accuracy but is 110× slower. VFI-accurate is 400× slower yet less accurate.

[9] Points near the borrowing constraint are excluded, as the Euler equation holds as an inequality there; this is standard practice following Santos [2000].

Table 2 shows an important result: TI can match EGM's accuracy (mean error -4.8), but only by computing μ exactly during search, which makes it 110 times slower. VFI-accurate is 400 times slower than EGM yet still 1.4 orders of magnitude less accurate (-3.4 versus -4.8). The fundamental difference is that EGM and TI work with the Euler equation directly while VFI optimizes the Bellman equation; Euler-based methods achieve higher precision when μ is computed exactly. But TI solves the Euler equation numerically via bisection, whereas EGM inverts it analytically. This is why EGM sidesteps the speed-accuracy tradeoff entirely: the endogenous grid places evaluation points exactly where the Euler equation is satisfied, achieving accurate-mode precision at fast-mode speed.

Consumption-equivalent welfare costs (the permanent consumption an agent would sacrifice to avoid using the approximate policy) are below 0.1% for EGM, making numerical error economically negligible.[10]

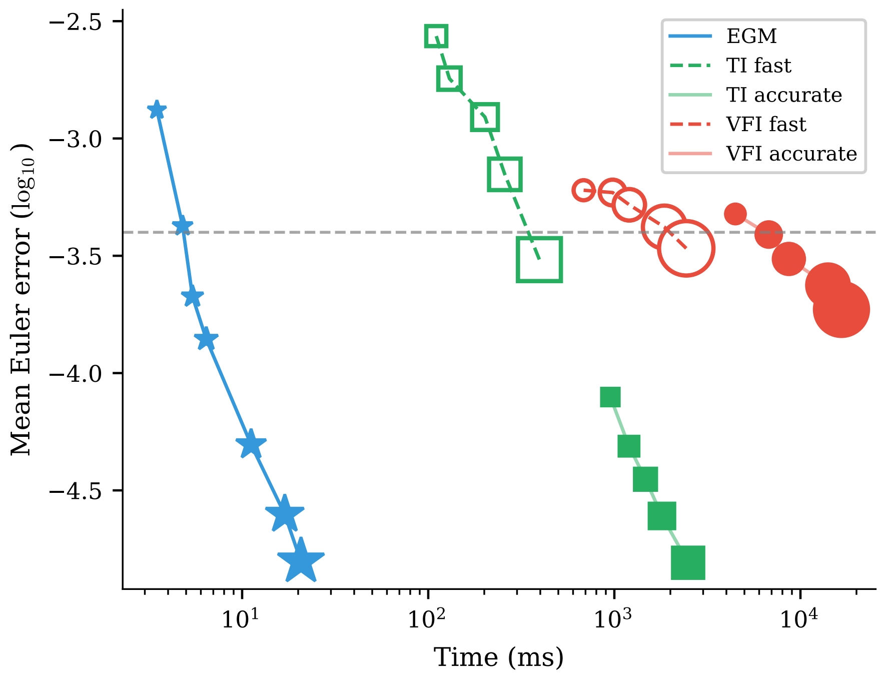

Figure 2: Speed-accuracy Pareto frontier across grid sizes. Marker size reflects grid resolution n. The dashed line marks a reference accuracy level; EGM achieves this with small grids while VFI and TI require substantially more computation.

Equal-accuracy comparison. A more informative comparison holds accuracy constant. Figure 2 shows that to match EGM's Euler errors with 20–25 grid points, VFI requires 50–300 points; the resulting speedup is 150–630 times.[11] This comparison isolates computational efficiency: both methods deliver the same solution quality, but EGM does so with far less computation. For applications requiring many model evaluations (estimation, uncertainty quantification, real-time decision support), such speedups translate directly into feasibility.

Speed-accuracy tradeoff. Time iteration offers a middle ground between EGM and VFI. Like EGM, TI works with the Euler equation; like VFI, it requires numerical optimization. As Figure 3 illustrates, both VFI and TI face a speed-accuracy tradeoff: computing μ exactly at each candidate during search is accurate but slow, as Table 2 documents; precomputing μ on the asset grid and interpolating is fast but introduces

[10]Computed by simulating 20,000 agents under the approximate and high-accuracy policies with identical shocks, then comparing ergodic average values.

[11]See Appendix: Equal-accuracy comparison for details.

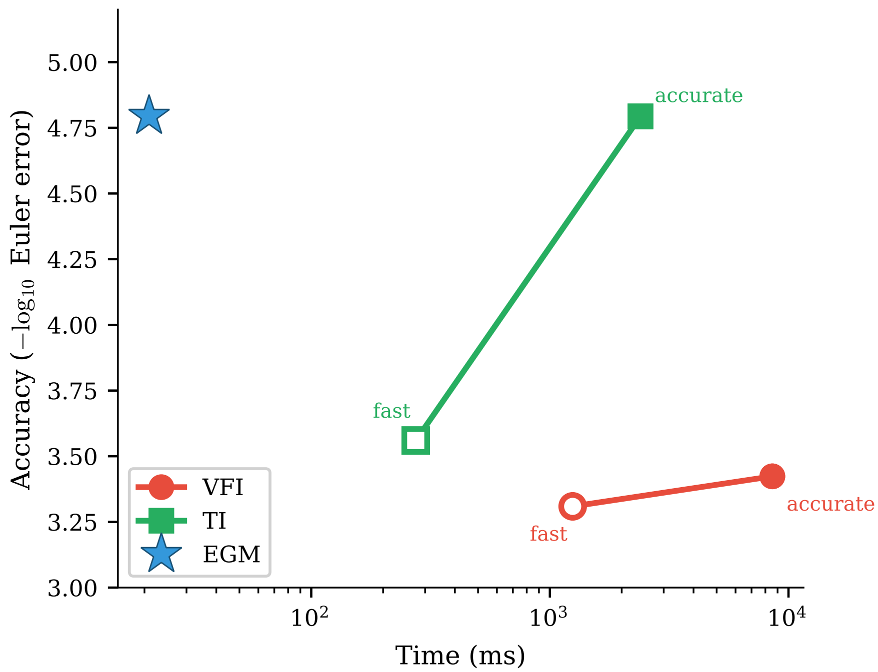

Figure 3: Speed-accuracy tradeoff. VFI and TI face a tradeoff between fast modes (open markers) and accurate modes (filled markers). EGM sidesteps this tradeoff: a single mode achieves both speed and accuracy.

approximation error, as Table 1 shows. The fast modes in Table 1 represent practical implementations. EGM sidesteps this tradeoff entirely: the endogenous grid places evaluation points exactly where the Euler equation is satisfied, achieving accurate-mode precision at fast-mode speed.

Howard acceleration. The baseline algorithm performs one policy update per value update. To accelerate convergence, fix the policy $c(\cdot,\cdot)$ and iterate on W alone:

$$W^{(k+1)}(m,z) = (1-\beta)c(m,z)^{1-\rho} + \beta\mu^{(k)}(m-c,z) \tag{11}$$

where $\mu^{(k)}$ uses $W^{(k)}$. After K value iterations, the policy is updated via EGM. The optimal K differs by method: VFI benefits from large K (30 or more), TI performs best with small K (3-5), while EGM is fastest at $K=1$ because its analytic policy step is already cheap.[12] With one additional value iteration ($K=2$), EGM becomes even more accurate while remaining faster than all VFI and TI configurations.

5 Conclusion

The endogenous grid method extends to Epstein-Zin preferences through the transformation $W = V^{1-\rho}$, which decouples the Euler equation from the value recursion.

The approach generalizes beyond Epstein-Zin. The method extends to the limiting case $\rho = 1$ (unit EIS) and to the risk-sensitive preferences of Hansen and Sargent [1995] used in robust control, both of which admit analogous transformations and closed-form Euler inversions.

The single-asset case presented here serves as a building block for richer models with multiple assets, portfolio choice, or labor supply decisions.

[12] Appendix: Howard acceleration parameter examines how K affects each method in both fast and accurate modes.

For practitioners solving consumption-savings models with recursive utility, the message is simple: EGM works. Where time iteration requires numerical root-finding at every grid point, EZ-EGM inverts the Euler equation analytically. The transformation is elementary, the implementation follows standard EGM patterns, and the speed gains are substantial.

This online appendix supplements the main text with algorithmic details and additional benchmarks. It presents pseudocode for value function iteration (VFI) and time iteration (TI), documents the speed-accuracy tradeoff these methods face, reports results with Howard policy improvement acceleration, provides equal-accuracy comparisons that hold solution quality constant while measuring computational cost, compares Euler error evaluation methods on the exogenous grid versus the ergodic distribution, and discusses the case $\rho > 1$ (EIS < 1) where the transformed problem becomes a minimization.

A Appendix: Alternative algorithms

Both value function iteration (VFI) and time iteration (TI) require numerical search at each grid point, unlike EGM which inverts the Euler equation analytically. Both methods face an inherent tradeoff: the certainty equivalent μ must be evaluated at candidate consumption values during search. Two approaches exist: compute μ exactly at each candidate (`accurate` mode), or precompute μ on the asset grid and interpolate during search (`fast` mode).

Value function iteration. VFI finds optimal consumption by maximizing the Bellman equation directly:

$$V(m,z) = \max_c \left[(1-\beta)c^{1-\rho} + \beta\mu(m-c,z)^{1-\rho}\right]^{1/(1-\rho)} \tag{12}$$

where $\mu(a,z) = \left(\mathbb{E}[V(Ra + y(z'))^{1-\gamma}|z]\right)^{1/(1-\gamma)}$ is the certainty equivalent. Golden-section search evaluates μ at many candidate consumption values.

Algorithm 2 VFI for Epstein-Zin

Require: Exogenous grid $\{m_i\}_{i=1}^I$, asset grid $\{a_j\}_{j=1}^J$, initial guess $V^{(0)}$, mode $\in \{\texttt{accurate}, \texttt{fast}\}$
Ensure: Converged value $V(\cdot,\cdot)$ and policy $c(\cdot,\cdot)$

```
Given V^(n) from iteration n
if mode = fast then
    Precompute μ_j = (Σ_ℓ π_kℓ V^(n)(R a_j + y(z'_ℓ))^(1−γ))^(1/(1−γ)) for all a_j
end if
for each income state z_k and each m_i do
    Define objective f(c) = [(1 − β)c^(1−ρ) + βμ(m_i − c, z_k)^(1−ρ)]^(1/(1−ρ))
    if mode = accurate then
        Compute μ by interpolating V^(n) at m' = R(m_i − c) + y(z'), then taking expectation
    else
        Interpolate μ from precomputed {μ_j} at a = m_i − c
    end if
    Maximize f(c) via golden-section search to get c^(n+1)(m_i, z_k) and V^(n+1)(m_i, z_k)
end for
Iterate until ‖V^(n) − V^(n−1)‖ < ε
```

Time iteration. Time iteration (TI), introduced by Coleman [1991], works directly with the Euler equation rather than the Bellman equation. Like EGM, TI exploits the first-order condition; unlike EGM, it cannot invert the Euler equation analytically and must use numerical root-finding. Rendahl [2015] proves that under standard conditions, TI converges to the same solution as VFI. Coeurdacier et al. [2020] extend TI to Epstein-Zin preferences, treating consumption, value, and expected value as simultaneous controls; our accurate mode follows this approach by interpolating decision rules at each candidate consumption during root-finding.

Given current guesses $(c^{(n)}, V^{(n)})$, TI finds consumption at each (m, z) by solving the Euler equation via bisection, then updates the value function. The certainty equivalent uses V directly: $\mu(a, z) = (\mathbb{E}[V(m', z')^{1-\gamma}])^{1/(1-\gamma)}$, exactly as in VFI.

Algorithm 3 Time Iteration for Epstein-Zin

Require: Exogenous grid $\{m_i\}_{i=1}^{I}$, asset grid $\{a_j\}_{j=1}^{J}$, initial guess $(c^{(0)}, V^{(0)})$, mode $\in \{\texttt{accurate}, \texttt{fast}\}$
Ensure: Converged policy $c(\cdot,\cdot)$ and value $V(\cdot,\cdot)$

Given $(c^{(n)}, V^{(n)})$ from iteration n
Precompute $c^{(n)}$ and $V^{(n)}$ at $m' = Ra_j + y(z'_\ell)$ for all (a_j, z'_ℓ)
if mode = $\texttt{fast}$ **then**
    Precompute $\mu_j = (\sum_\ell \pi_{k\ell} V^{(n)}(m'_{j\ell})^{1-\gamma})^{1/(1-\gamma)}$ for all a_j
end if
for each income state z_k and each m_i **do**
    Define Euler residual $r(c)$ using μ and marginal utilities
    if mode = $\texttt{accurate}$ **then**
        Compute μ by interpolating $V^{(n)}$ at $m' = R(m_i - c) + y(z')$, then taking expectation
    else
        Interpolate μ from precomputed $\{\mu_j\}$ at $a = m_i - c$
    end if
    Solve $r(c) = 0$ via bisection to get $c^{(n+1)}(m_i, z_k)$
    Update $V^{(n+1)}(m_i, z_k)$ using the Bellman equation
end for
Iterate until $\|c^{(n)} - c^{(n-1)}\| < \varepsilon$

B Appendix: Speed-accuracy tradeoff

Table 3: Accurate-mode comparison: EZ-EGM vs numerical search ($K = 1$)

Method	Time (ms)	Euler Mean	Euler Max
EZ-EGM	25	−4.8	−3.4
TI (accurate)	2371	−4.8	−3.4
VFI (accurate)	8385	−3.5	−2.3

Note: Grid size 100, $K = 1$ (no Howard acceleration). Accurate modes compute μ exactly during search (no precomputation). Errors evaluated on uniform grid ($\bar{\varepsilon}_G$); see the main text for ergodic errors ($\bar{\varepsilon}_\pi$).

Table 3 reveals two main findings. First, TI-accurate achieves the same accuracy as EZ-EGM (mean Euler error -4.8), confirming that when μ is computed exactly during search, time iteration matches EGM's precision. Second, this accuracy comes at substantial cost: TI-accurate is nearly 100 times slower than EZ-EGM, because bisection requires many μ evaluations per grid point, each involving interpolation and expectation. VFI-accurate is slower still (over 300 times) yet less accurate than both EGM and TI, reflecting the underlying disadvantage of optimizing over the Bellman equation rather than working with the Euler equation.

EGM's advantage is that it achieves accurate-mode precision at fast-mode speed by avoiding numerical search entirely. The main text compares EGM to the fast modes of VFI and TI, which represent practical implementations; this appendix shows that even the accurate modes cannot match EGM's speed.

C Appendix: Howard acceleration parameter

The main text reports baseline results without Howard acceleration ($K = 1$). This appendix examines how the choice of K affects performance in both fast and accurate modes.

Note on iteration counts: For EGM, each "policy iteration" consists of one Euler equation inversion (the EGM step) plus up to K value function updates. For TI, each policy iteration consists of one bisection solve at every grid point plus up to K value updates. For VFI, each "policy iteration" consists of one numerical optimization over the consumption grid plus up to K value function updates. All three methods use early termination: if the value function converges before K updates, the remaining updates are skipped. The EGM policy step is analytic and thus much cheaper than VFI's optimization or TI's bisection.

EGM. EGM has no fast/accurate distinction because it evaluates μ exactly on the endogenous grid.

Table 4: EZ-EGM: Effect of Howard acceleration parameter K

K	Time (ms)	Policy Iters	Euler Mean	Euler Max
1	25	141	−4.8	−3.4
2	49	99	−4.9	−3.4
3	62	86	−4.9	−3.4
4	70	78	−4.9	−3.4
5	78	70	−4.9	−3.4

For EZ-EGM, $K = 1$ is fastest. Additional value iterations reduce policy iterations (from 141 to 70) but increase total time because the EGM policy step is already cheap and accurate. The slight accuracy improvement (-4.8 to -4.9) does not justify the 3x time increase.

VFI. VFI benefits substantially from Howard acceleration in both modes. The optimal $K \approx$ 30-40 reduces time by 7x in fast mode and 13x in accurate mode.

Table 5: VFI fast mode: Effect of Howard acceleration parameter K

K	Time (ms)	Policy Iters	Euler Mean	Euler Max
1	1190	239	−3.3	−2.4
10	247	31	−3.3	−2.4
20	183	16	−3.3	−2.4
30	181	11	−3.3	−2.4
40	177	9	−3.3	−2.4
50	193	8	−3.3	−2.4

VFI-accurate benefits more from Howard than VFI-fast because each policy optimization is more expensive in accurate mode (computing μ exactly versus interpolating). At $K = 40$, VFI-accurate (626ms) is faster than VFI-fast at $K = 1$ (1190ms), though still less accurate (-3.5 versus -3.3). This reflects the cost structure: golden-section search is the bottleneck, and reducing policy iterations helps more when each policy step is expensive.

Table 6: VFI accurate mode: Effect of Howard acceleration parameter K

K	Time (ms)	Policy Iters	Euler Mean	Euler Max
1	8385	239	−3.5	−2.3
10	1312	31	−3.5	−2.3
20	802	16	−3.5	−2.3
30	641	11	−3.5	−2.3
40	626	10	−3.5	−2.3
50	650	10	−3.5	−2.3

TI. TI shows a distinctive pattern: moderate K helps, but larger values hurt performance. The optimal K differs between modes.

Table 7: TI fast mode: Effect of Howard acceleration parameter K

K	Time (ms)	Policy Iters	Euler Mean	Euler Max
1	252	140	−3.2	−2.7
2	229	100	−3.2	−2.7
3	213	88	−3.2	−2.7
4	213	81	−3.2	−2.7
5	481	174	−3.2	−2.7

Table 8: TI accurate mode: Effect of Howard acceleration parameter K

K	Time (ms)	Policy Iters	Euler Mean	Euler Max
1	2371	141	−4.8	−3.4
2	1716	98	−4.9	−3.4
3	1540	86	−4.9	−3.4
4	1427	78	−4.9	−3.4
5	1290	70	−4.9	−3.4

TI-fast is optimal at $K = 3$-4, with time falling from 252ms to 213ms. Beyond $K = 4$, performance degrades sharply: $K = 5$ nearly doubles time to 481ms. This instability occurs because TI's value function updates can overshoot when the policy is held fixed, destabilizing convergence.

TI-accurate is optimal at $K = 5$, with time falling from 2371ms to 1290ms, a 1.8x speedup. The higher optimal K reflects the greater cost of each bisection step in accurate mode.

The contrast with VFI is instructive: VFI benefits from large K (up to 30-50) because golden-section search produces noisier policy updates, requiring more value iterations to stabilize. TI's bisection produces more accurate policies, reducing the benefit of additional value iterations and causing instability at high K.

D Appendix: Equal-accuracy comparison

The main text compares EZ-EGM and VFI at fixed grid size (100 points), where EGM is both faster and more accurate. A more informative comparison holds accuracy constant: what grid size does VFI require to match EGM's Euler errors, and how do solve times compare?

We find grid sizes where mean Euler errors are approximately equal:

Table 9: Speed comparison at equal accuracy

VFI n	EGM n	Mean Error	VFI (ms)	EGM (ms)	Speedup
50	20	-3.2	672	5	145×
100	20	-3.3	1214	5	261×
150	20	-3.4	1793	5	386×
200	20	-3.5	2281	5	490×
300	25	-3.6	3433	5	627×

At equivalent accuracy levels, EZ-EGM is 150–630 times faster than VFI. The speedup grows with target accuracy: to achieve mean Euler error near -3.6, VFI requires 300 grid points while EGM needs only 25. Euler errors weighted by the ergodic distribution (5th–95th percentiles of simulated wealth) are similar to uniformly-weighted errors; the ergodic distribution concentrates wealth in the lower range (median around 3 versus grid maximum of 20), but both weighting schemes yield mean errors near -4.9.

These results use $\beta R = 0.98 < 1$, so agents converge to finite target wealth. For more impatient agents ($\beta R \approx 0.92$), target wealth falls but accuracy remains similar. For very patient agents ($\beta R \geq 1$), target wealth diverges and the grid must be extended accordingly; the accuracy comparison in that regime warrants separate investigation.

E Appendix: Euler error evaluation methods

Two approaches exist for evaluating Euler equation errors: uniform grid evaluation and ergodic distribution weighting; we report both for completeness.

Standard approach (Santos 2000). Santos [2000] establishes the theoretical foundation for using Euler equation residuals to bound policy function errors. The standard implementation evaluates errors on a uniform grid spanning the state space, typically excluding boundary regions where the constraint binds. For our baseline (100 grid points), we evaluate at the 10th–90th percentiles of the grid:

$$\bar{\varepsilon}_G = \frac{1}{|G|} \sum_{(m,z)\in G} \varepsilon(m, z) \tag{13}$$

where $G \subset \mathcal{M} \times \mathcal{Z}$ is the set of interior grid points.

Ergodic distribution approach. An alternative evaluates errors at wealth levels agents actually visit in the long run. We simulate the ergodic distribution (10,000 agents, 500 periods, 200 burn-in) and compute errors at the 5th–95th percentiles of realized wealth:

$$\bar{\varepsilon}_\pi = \frac{1}{|S|} \sum_{(m,z)\in S} \varepsilon(m, z) \tag{14}$$

where $S \sim \pi$ is a sample from the stationary distribution $\pi(m, z)$. This approach weights accuracy by economic relevance: errors at rarely-visited wealth levels matter less than errors where agents spend time.

For this calibration ($\beta R = 0.98 < 1$), the stationary distribution π concentrates in the lower wealth range: the median is approximately $m = 3$ versus the grid maximum of 20. Despite this difference in evaluation regions, EGM yields similar error statistics under both $\bar{\varepsilon}_G$ and $\bar{\varepsilon}_\pi$. TI shows improved accuracy under $\bar{\varepsilon}_\pi$ (-3.6 versus -3.2), suggesting its errors concentrate at high wealth levels that agents rarely visit.

EGM achieves the best accuracy because it satisfies the Euler equation by construction at the endogenous grid points; interpolation introduces error, but linear interpolation on a fine grid (100 points) keeps this

Table 10: Euler errors by evaluation method ($\log_{10}$)

	$\bar{\varepsilon}_G$ (grid)		$\bar{\varepsilon}_\pi$ (ergodic)	
Method	Mean	Max	Mean	Max
EZ-EGM	-4.8	-3.4	-4.8	-3.2
TI	-3.2	-2.7	-3.6	-2.7
VFI	-3.3	-2.4	-3.3	-2.3

Note: $\bar{\varepsilon}_G$ evaluates at 10th-90th percentiles of grid ($m \in [0.4, 16.5]$). $\bar{\varepsilon}_\pi$ evaluates at 5th-95th percentiles of simulated wealth ($m \in [0.7, 9.6]$); median wealth ≈ 3. Errors exclude constrained points where the Euler equation holds as inequality. All methods use fast mode (precomputed μ).

small. TI and VFI both require numerical search, which compounds interpolation error. VFI shows the worst max error under $\bar{\varepsilon}_\pi$ (-2.3), reflecting that its optimization-based approach is less precise at low wealth levels where the policy function has higher curvature and agents spend the most time.

F Appendix: Robustness to $\rho > 1$

The baseline uses $\rho = 2/3$ (EIS $= 1.5$), but the algorithm applies equally to $\rho > 1$ (EIS < 1). We verify this by solving the same model with $\rho \in \{0.5, 0.9, 1.1, 1.5, 2, 3\}$, holding $\gamma = 10$ fixed. Across all values, EZ-EGM achieves mean Euler errors near -5 and max errors near -3.5 (in $\log_{10}$ units). The case $\rho > 1$ is empirically relevant: the meta-analysis of Havránek [2015] finds that micro estimates of the EIS often fall below unity.

When $\rho > 1$, the transformation $W = V^{1-\rho}$ is decreasing, so maximizing V is equivalent to minimizing W. No algorithmic modification is required: the first-order condition $\partial W/\partial c = 0$ is sufficient and characterizes the optimum regardless of whether it is a maximum or minimum, and EGM inverts the Euler equation analytically. VFI is unaffected because it works with V directly; the CES aggregator remains increasing in c for all $\rho \neq 1$. Time iteration, like EGM, works with the Euler equation.

Acknowledgments

I am grateful to Christopher D. Carroll for years of mentorship, for detailed feedback on this paper, and for sustained support throughout my career. I also thank Matthew N. White for extensive discussions and collaboration that shaped this work. This work was supported by the Alfred P. Sloan Foundation [G-2025-79177].

References

Christopher D Carroll. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320, 6 2006. doi:10.1016/j.econlet.2005.09.013.

Angus Deaton. Saving and liquidity constraints. *Econometrica*, 59(5):1221–1248, 9 1991. doi:10.2307/2938366.

Francisco Barillas and Jesús Fernández-Villaverde. A Generalization of the Endogenous Grid Method. *Journal of Economic Dynamics and Control*, 31(8):2698–2712, 8 2007. doi:10.1016/j.jedc.2006.08.005.

Fedor Iskhakov, Thomas H Jørgensen, John Rust, and Bertel Schjerning. The Endogenous Grid Method for Discrete-Continuous Dynamic Choice Models with (or without) Taste Shocks. *Quantitative Economics*, 8 (2):317–365, 7 2017. doi:10.3982/QE643.

Larry G Epstein and Stanley E Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969, 7 1989. doi:10.2307/1913778.

Larry G Epstein and Stanley E Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *Journal of Political Economy*, 99(2):263–286, 4 1991. doi:10.1086/261750.

Philippe Weil. The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics*, 24(3):401–421, 11 1989. doi:10.1016/0304-3932(89)90028-7.

David M Kreps and Evan L Porteus. Temporal Resolution of Uncertainty and Dynamic Choice Theory. *Econometrica*, 46(1):185–200, 1 1978. doi:10.2307/1913656.

Darrell Duffie and Larry G Epstein. Stochastic Differential Utility. *Econometrica*, 60(2):353–394, 3 1992. doi:10.2307/2951600.

Guanlong Ren and John Stachurski. Dynamic programming with recursive preferences: Optimality and applications. Preprint 1812.05748, arXiv, 6 2020. URL `https://arxiv.org/abs/1812.05748`.

Nicolas Coeurdacier, Hélène Rey, and Pablo Winant. Financial Integration and Growth in a Risky World. *Journal of Monetary Economics*, 112:1–21, 2020. doi:10.1016/j.jmoneco.2019.01.022.

Lars Peter Hansen and Thomas J Sargent. Discounted linear exponential quadratic Gaussian control. *IEEE Transactions on Automatic Control*, 40(5):968–971, 5 1995. doi:10.1109/9.384242.

Ravi Bansal and Amir Yaron. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *The Journal of Finance*, 59(4):1481–1509, 8 2004. doi:10.1111/j.1540-6261.2004.00670.x.

Alexander Meyer-Gohde. Generalized entropy and model uncertainty. *Journal of Economic Theory*, 183: 312–343, 9 2019. doi:10.1016/j.jet.2019.06.004.

David K Backus, Bryan R Routledge, and Stanley E Zin. *Recursive preferences*. Palgrave Macmillan, London, 2 edition, 2008. doi:10.1057/978-1-349-95121-5_1963-1.

Thomas D Tallarini Jr. Risk-sensitive real business cycles. *Journal of Monetary Economics*, 45(3):507–532, 6 2000. doi:10.1016/S0304-3932(00)00012-X.

Alberto Giovannini and Philippe Weil. Risk Aversion and Intertemporal Substitution in the Capital Asset Pricing Model. Working Paper 2824, National Bureau of Economic Research, 1 1989.

Ronald A Howard. *Dynamic Programming and Markov Processes*. MIT Press, Cambridge, MA, 1960.

Kjetil Storesletten, Christopher I Telmer, and Amir Yaron. Cyclical Dynamics in Idiosyncratic Labor Market Risk. *Journal of Political Economy*, 112(3):695–717, 6 2004. doi:10.1086/383105.

George Tauchen. Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177–181, 1986. doi:10.1016/0165-1765(86)90168-0.

Tomáš Havránek. Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting. *Journal of the European Economic Association*, 13(6):1180–1204, 12 2015. doi:10.1111/jeea.12133.

Pontus Rendahl. Inequality Constraints and Euler Equation-based Solution Methods. *The Economic Journal*, 125(585):1110–1135, 6 2015. doi:10.1111/ecoj.12115.

Wilbur John Coleman. Equilibrium in a Production Economy with an Income Tax. *Econometrica*, 59(4): 1091–1104, 7 1991. doi:10.2307/2938175.

Manuel S Santos. Accuracy of Numerical Solutions Using the Euler Equation Residuals. *Econometrica*, 68 (6):1377–1402, 11 2000. doi:10.1111/1468-0262.00165.

Lilia Maliar and Serguei Maliar. *Numerical Methods for Large-Scale Dynamic Economic Models*, volume 3, pages 325–477. Elsevier, 2014. doi:10.1016/B978-0-444-52980-0.00007-4.