

Online Appendix for “The Sequential Endogenous Grid Method”^{*}

Alan Lujan^{a,b,*}

^a*Johns Hopkins University, Krieger School of Arts and Sciences Washington, DC,*
^b*Econ-ARK*

Keywords: endogenous grid method, dynamic programming, interpolation, computational economics

Part A specifies the health-investment interpolation testbed and its warp sweep; Part B reports the full-solve interpolation diagnostic’s protocol and readings. Section references point to the main paper.

Appendix A. The Health Investment Model as an Interpolation Testbed

The two-asset household block of Section 3 is the paper’s primary multidimensional application. For two reasons, this appendix retains a second two-dimensional model, the consumption-saving problem with health investment of Ludwig and Schön (2018) extended with the wage and depreciation risk of White (2015). First, the model is home ground for both alternatives: Ludwig and Schön (2018) introduced Delaunay interpolation on the version whose only risk is survival, and White (2015) introduced curvilinear interpolation on the stochastic version adapted here. Second, the health-investment stage is a concrete instance of the invertible-transition-marginal form of the paper’s Section 5.

Appendix A.1. A Health Investment Model

Each period $t \in \{0, \dots, T\}$, the household observes market resources m_t and health capital h_t and chooses consumption c_t and health investment n_t . Consumption yields CRRA utility; health investment produces health capital through a concave production function; and health capital both reduces mortality risk and raises labor income.

The recursive problem is

$$\begin{aligned} v_t(m_t, h_t) &= \max_{c_t \geq 0, n_t \geq 0} u(c_t) + \beta \mathcal{S}_t \mathbb{E}_t [v_{t+1}(m_{t+1}, h_{t+1})] \\ &\text{subject to} \\ H_t &= h_t + g(n_t) \\ a_t &= m_t - c_t - n_t \\ \mathcal{S}_t &= 1 - D_t / (1 + H_t) \\ y_{t+1} &= \omega_{t+1} H_t \\ h_{t+1} &= (1 - \delta_{t+1}) H_t \\ m_{t+1} &= R a_t + y_{t+1}, \end{aligned} \tag{A.1}$$

where the utility function and health production function are

^{*}This work was supported by the Alfred P. Sloan Foundation through the Econ-ARK project.

^{*}Corresponding author

Email address: alujan@jhu.edu (Alan Lujan)

URL: <https://advanced.jhu.edu/directory/alan-lujan/> (Alan Lujan)

$$u(c) = \frac{c^{1-\rho}}{1-\rho}, \quad g(n) = \frac{\gamma}{\varepsilon} n^\varepsilon. \quad (\text{A.2})$$

Here H_t is post-investment health capital, a_t end-of-period assets, $\mathcal{S}_t$ the survival probability (the maximum death probability D_t applies at $H_t = 0$), δ_{t+1} the stochastic health depreciation rate, ω_{t+1} the stochastic wage rate, and y_{t+1} labor income. Ludwig and Schön (2018) take this survival function from Hall and Jones (2007). With scale $\gamma > 0$ and elasticity $0 < \varepsilon < 1$, health investment has diminishing returns. The terminal condition $v_{T+1} \equiv 0$ makes the final-period household consume all remaining resources. Higher health investment raises $\mathcal{S}_t$ and thereby the weight on future value.

This model requires $\rho < 1$, so utility is non-negative for all $c > 0$. With a positive probability of zero wage, ending a nonterminal period with no liquid assets would risk zero consumption, so the liquidity constraint never binds before the terminal period. Strict concavity then makes the first-order conditions necessary and sufficient at interior solutions. The implementation also handles positive wage floors, like the benchmark calibration's unemployment-income floor, through an explicit liquidity-constrained segment at $a = 0$. As $\varepsilon \rightarrow 1$, g' approaches a constant, the marginal inversion becomes ill-conditioned, and at the linear limit the first-order condition no longer identifies health investment. This boundary case therefore falls outside the invertible-transition-marginal class.¹

Appendix A.2. Sequential EGM Solution

The problem splits into three stages: health investment, consumption, and a post-decision stage that handles expectations. I retain the stage superscripts of Section 2, so $v_t \equiv v_t^0$ is the first decision stage, and write w_t in place of v_t^2 for the post-decision value, following Powell (2011). Because value functions have multiple arguments, partials appear as superscripts ($w_t^a \equiv \partial w_t / \partial a$).

The post-decision stage values the end-of-period states (a_t, H_t) before next period's shocks are realized:

$$w_t(a_t, H_t) = \beta \left(1 - \frac{D_t}{1 + H_t} \right) \mathbb{E}_t [v_{t+1}(m_{t+1}, h_{t+1})]. \quad (\text{A.3})$$

Building w_t once lets the earlier stages reuse it without repeated expectation calculations.²

After the health investment decision, the household has liquid wealth ℓ_t and post-investment health H_t . The consumption problem is

$$v_t^1(\ell_t, H_t) = \max_{c_t \geq 0} u(c_t) + w_t(a_t, H_t), \quad a_t = \ell_t - c_t. \quad (\text{A.4})$$

Post-investment health H_t passes through this stage because it enters the continuation value and the consumption decision leaves it unchanged. For each fixed H_t , this is a standard one-dimensional consumption-saving problem that EGM solves.³ On an exogenous grid $([a], [H])$ of post-decision states, the first-order condition $u'(c_t) = w_t^a$ inverts to $\mathbf{c}_t = (w_t^a)^{-1/\rho}$, which defines the endogenous liquid-wealth grid $\mathbf{l}_t = [a] + \mathbf{c}_t$. On that grid the envelope conditions give the investment stage the marginal values $(v_t^1)^\ell = u'(c_t)$ and $(v_t^1)^H = w_t^H$.

At the beginning of the period, the household has market resources m_t and health capital h_t . The health investment problem is

¹With $\rho > 1$, utility is negative and the product $\mathcal{S}_t \cdot \mathbb{E}[v_{t+1}]$ would reward mortality. The restriction mirrors the calibrations of Ludwig and Schön (2018) and White (2015) and is specific to this health formulation, not to EGMⁿ generally. Models with mortality risk and $\rho > 1$ can use a bequest motive or utility differences (see De Nardi (2004)).

²The marginal value functions $w_t^a(a_t, H_t) \equiv \partial w_t / \partial a$ and $w_t^H(a_t, H_t) \equiv \partial w_t / \partial H$ are computed once on an exogenous grid and interpolated as needed by earlier stages. The health marginal w_t^H includes both the survival-probability benefit of health (through $\partial \mathcal{S}_t / \partial H_t = D_t / (1 + H_t)^2$) and the continuation-value benefit through next period's health.

³Strict concavity of u and w_t ensures joint strict concavity of $v_t^1(\ell_t, H_t)$, so the recovered optimum is unique. Grid validity is separate: concavity alone does not deliver fold-freeness or index monotonicity, so the warp sweep below tests both directly on every solved mesh.

$$v_t^0(m_t, h_t) = \max_{n_t \geq 0} v_t^1(\ell_t, H_t), \quad \ell_t = m_t - n_t, \quad H_t = h_t + g(n_t). \quad (\text{A.5})$$

With no direct utility payoff at this stage, investment moves liquid wealth through the budget and post-investment health through the production function. Because g' has an explicit inverse, the EGM step can run before the transition equations recover the incoming states. The Inada condition $\lim_{n \rightarrow 0} g'(n) = \infty$ handles the constraint $n_t \geq 0$, keeping solutions interior whenever the marginal value of health is positive and resources permit.

The first-order condition of the health investment problem rearranges to

$$g'(n_t) = \frac{(v_t^1)^\ell(\ell_t, H_t)}{(v_t^1)^H(\ell_t, H_t)}. \quad (\text{A.6})$$

Since $g'(n) = \gamma n^{\varepsilon-1}$ is strictly decreasing for $0 < \varepsilon < 1$, optimal health investment on an exogenous grid $([\ell], [H])$ of liquid wealth and post-investment health follows:

$$\mathbf{n}_t([\ell], [H]) = \left(\frac{(v_t^1)^\ell([\ell], [H])}{\gamma (v_t^1)^H([\ell], [H])} \right)^{1/(\varepsilon-1)}. \quad (\text{A.7})$$

From the transition equations, the endogenous pre-decision states then follow:

$$\begin{aligned} \mathbf{m}_t([\ell], [H]) &= [\ell] + \mathbf{n}_t([\ell], [H]), \\ \mathbf{h}_t([\ell], [H]) &= [H] - g(\mathbf{n}_t([\ell], [H])). \end{aligned} \quad (\text{A.8})$$

By the envelope theorem, the beginning-of-period marginal values are

$$\begin{aligned} (v_t^0)^m(m_t, h_t) &= (v_t^1)^\ell(\ell_t, H_t), \\ (v_t^0)^h(m_t, h_t) &= (v_t^1)^H(\ell_t, H_t). \end{aligned} \quad (\text{A.9})$$

Starting from a regular rectilinear grid of $([a], [H])$, the two EGM steps together produce the curvilinear endogenous grid of $(\mathbf{m}_t, \mathbf{h}_t)$ shown in Figure A.1.⁴

Index preservation across warping enables the interpolation strategy of Section 4. The mapping $(\ell, H) \mapsto (m, h)$ above is the recovery map whose sampled meshes the warp sweep below measures.

Appendix A.3. Interpolation Comparison

This appendix adds Delaunay triangulation to the comparison of Section 4, on a model where all three methods complete. Figure A.2 times the two structured methods across both time horizon (backward-induction cycles) and grid resolution.⁵ Apart from one small 20×15 grid, the grid sweep holds fifty health nodes and scales the asset axis from 44 to 176 points; the cycle sweep stops at fifteen cycles. Denser health grids or twenty cycles make the endogenous health coordinate non-monotone between the two lowest health nodes, which violates IPM, a condition ENGINE requires. On this model the comparison isolates the interpolation method cleanly, because the sequential and simultaneous solvers produce numerically identical policies (the maximum absolute difference in consumption is 3×10^{-13}). EGMⁿ with ENGINE solves 2.8 to 4.7 times faster than the simultaneous solver with the curvilinear method across both sweeps. A third configuration that swaps only the interpolant on the simultaneous solver reproduces most of that gap, so the interpolation method, not the decomposition, accounts for the bulk of the speedup.

⁴The figure uses $\rho = 0.5$, $\beta = 0.95$, $R = 1.03$, health production exponent $\varepsilon = 0.35$, health depreciation rate $\delta \sim U[0, 0.1]$, wage rate $\omega \sim \text{lognormal}$ with mean 0.1 and standard deviation 0.1, and a two-period model ($T = 1$, so the household lives periods $t \in \{0, 1\}$).

⁵The two codebases are compiled alike (Section 4's implementation footnote) and run on the same hardware and software as Section 3's benchmarks, both recorded in the replication archive's README.




Figure A.1: EGM grid transformation for the health investment model at $t = 0$ (the first period of economic life). Panel (a) shows the regular exogenous grid of post-decision assets (a_t) and post-investment health (H_t). Panel (b) shows the curvilinear endogenous grid of market resources (m_t) and beginning-of-period health (h_t), with points colored by optimal consumption and the plot windowed to $m_t \leq 50$ for legibility. The endogenous grid preserves the index ordering of the exogenous grid despite substantial warping near the constraint boundary.




Figure A.2: Solve time and speedup on the health investment model, by backward-induction cycles at a 44×50 grid (a, b) and by grid size at ten cycles (c, d). Annotations give the speedup of EGM* with ENGINE over the simultaneous solver with the curvilinear method of White (2015). Bands are 95 percent confidence intervals of the mean over ten timed runs.

The converged policies separate the structured methods from the triangulation just as the timings do. Each interpolation method runs inside its own backward solve (holding the simultaneous solver fixed), and interpolation differences propagate through the solution. On a uniform 50×40 grid of off-mesh query points spanning $m, h \in [0.5, 25]$, with the curvilinear solution as the reference, ENGINE differs by a mean of 5.8×10^{-5} in consumption (maximum 7.2×10^{-3}) and 1.8×10^{-3} in health investment (maximum 4.8×10^{-2}); Delaunay triangulation differs by means of 3.5×10^{-4} and 4.9×10^{-3} (maxima 1.5×10^{-1} and 3.5×10^{-1}).

Figure A.3 maps the consumption disagreement over the state space. ENGINE’s differences from the curvilinear solution are small and nearly uniform, between 10^{-8} and 10^{-4} over most of the domain. However, the worst case, 7.2×10^{-3} , occurs in the few visible spots at the lowest market resources where the mesh cells are most warped, with fainter disagreement at the lowest health capital. The triangulation’s differences are larger and spatially diffuse, following its simplex boundaries.




Figure A.3: Consumption disagreement with the curvilinear reference over the (m, h) query plane, on a shared logarithmic color scale. (Left) ENGINE. (Right) Delaunay triangulation. Circles mark each panel’s maximum.

Appendix A.4. Accuracy against Measured Warp

The comparison above fixes the calibration and varies the query point. How do the methods behave as the mesh itself becomes more warped? I vary the two model parameters that bend the endogenous mesh hardest. Across the sweep, the health production elasticity takes values in $\{0.20, 0.35, 0.50\}$, controlling the curvature of the investment inversion, and the maximum death probability D_t , held constant across ages in every sweep calibration, takes values in $\{0.12, 0.50, 0.82\}$, scaling the survival term in w_t^H . I solve each of the nine calibrations for five periods on a smaller grid under the protocol above (each interpolant inside its own backward solve, with the simultaneous solver held fixed). After every solve I test IPM slice by slice and measure the warp by the empirical constant $\hat{\alpha}$ of Section 4, the maximum bracket displacement between adjacent health slices.

Accuracy is scored against the model’s own optimality conditions rather than against any reference solution. At three thousand off-mesh query points I rebuild the post-decision marginals w^a and w^H by quadrature through the true continuation solution and evaluate both first-order conditions. The consumption Euler residual uses implied consumption capped at the $a \geq 0$ budget bound, so constrained queries are scored rather than discarded. The investment margin $g'(n)w^H = u'(c)$ holds at interior investment whether or not the liquidity constraint binds, since the Inada condition on g keeps investment interior. Queries where an interpolated policy violates feasibility, whether through the budget or through a negative control, are repaired minimally for the quadrature and counted, none of them dropped.

Every warp-sweep mesh the two structured methods produce satisfies IPM. Yet across those meshes $\hat{\alpha}$ ranges from 10 to 16, five to eight times ENGINE’s default search margin of two. A mismatch in grid spacing produces these values: twenty health nodes span $[0.01, 50]$ whereas the asset axis is finely log-spaced, so a query’s bracket can cross many asset cells between adjacent health slices. Without the containment fallback of Section 4, the default-margin search misses brackets on this sweep without raising an error, producing

budget-infeasible interpolated policies and inflated extreme residuals. Setting the margin to the measured constant removes every violation on the sweep’s declared finite query set at no measurable time cost. All sweep solves use that matched margin. The archived implementation also carries the fallback, which runs a full row search when a windowed bracket misses the query.

Every fold on the structured meshes is a sliver cell at the lowest post-decision health node. I start the health grid at $H_t = 0.01$. The HARK implementation of this model starts it at zero, where the investment inversion at zero assets divides one infinite marginal value by another. At $H_t = 0.01$, however, the investment the inversion assigns, even at zero assets, produces more health than the node holds, so every inverted state on that row has negative health and lies outside the state space. Whether a cell on that row folds depends on the calibration, so I count the folds directly. The final period, inverted on the analytic terminal solution with no interpolation, produces the same slivers. Evaluation points reach the cells only at $(\varepsilon, D_t) = (0.35, 0.82)$ and $(0.50, 0.82)$. There, up to 0.45 percent of continuation points fall in folded cells, and at $(0.50, 0.82)$ so do 0.37 percent of final queries for the curvilinear method and 0.40 percent for ENGINE. Although ENGINE’s fold-free requirement fails in those cells, the reported residuals include every query inside one.

The triangulation’s own meshes violate IPM at five of the nine calibrations, the three with the lowest production elasticity (0.20) and the other two with the highest mortality (0.82). At the remaining four the triangulation’s $\hat{\alpha}$ is 12 or 13, against 10 to 12 for the structured meshes. In each of the five, between 82 and 120 of the mesh’s 820 nodes are undefined. The continuation queries of the highest health node fall above the convex hull of the next period’s mesh, where the boundary triangles stretch to outlying mesh nodes, and linear extrapolation from those triangles returns values outside the admissible range, such as negative consumption. The first-order conditions are then undefined for that health column, whose nodes drop out of the next triangulation. As the hull shrinks, the undefined nodes spread one period back at a time. By contrast, the two structured methods extrapolate along the mesh’s own index lines; every node of their meshes is defined at all nine calibrations.




Figure A.4: Consumption Euler residuals at off-mesh queries (median, 95th percentile, and maximum) against the measured bracket-displacement constant $\hat{\alpha}$ of each solve’s own endogenous mesh, across nine calibrations varying the production elasticity and the maximum death probability. Markers are horizontally jittered for legibility, and the three panels use different vertical scales; the triangulation’s five IPM-violating meshes have no defined $\hat{\alpha}$ and are absent from the panels.

Figure A.4 reports the consumption residuals against this measured warp. Median accuracy, between 3×10^{-4} and 3.4×10^{-3} across the sweep, moves by an order of magnitude with the calibration but by less than a factor of two across methods at any single calibration. The interpolant therefore matters far less for the median than the calibration’s warp does at this grid resolution.

The triangulation’s maximum residual sits between 0.25 and 0.66 at every calibration, including the five IPM-violating meshes the figure cannot place, repeating the simplex-boundary artifacts in Figure A.3. ENGINE has both a smaller 95th-percentile and a smaller maximum consumption residual than the curvilinear method in exactly three calibrations: $(\varepsilon, D_t) = (0.35, 0.82), (0.50, 0.50), (0.50, 0.82)$. At the last of these, its two consumption statistics, 0.019 and 0.20, are below the curvilinear method’s 0.033 and 0.41, and its maximum investment residual falls from 0.60 to 0.15. The constant $\hat{\alpha}$ alone does not rank the tails: among the three calibrations where ENGINE’s mesh has $\hat{\alpha} = 16$, ENGINE’s maximum consumption residual is higher at $(0.20, 0.82)$, lower at $(0.35, 0.82)$, and roughly half as large only at $(0.50, 0.82)$. Across the full sweep, ENGINE’s 95th percentile is smaller at seven of nine calibrations.

Every solve in the sweep completes. I compare timings only on the larger grids of Figure A.2.

Appendix B. The Full-Solve Interpolation Diagnostic

This appendix records the protocol behind Section 4’s full-solve-time comparison of ENGINE with the sector-walking curvilinear method of [White \(2015\)](#). Online Appendix A adds the Delaunay triangulation of [Ludwig and Schön \(2018\)](#).

I solve the Section 3 household block with each interpolator as the Stage-0 realign step, holding every other component fixed. Productivity follows the same three-state Rouwenhorst chain in every solve, with persistence 0.966 and unconditional standard deviation 0.92 for log productivity. Each iteration starts from the same marginal-value initialization and stops when the larger normalized sup-norm change in the beginning-of-period marginal values v_a^0 and v_b^0 (the arrays **Va** and **Vb** of the replication code) falls below 10^{-9} , subject to a cap of 8,000 backward steps; the normalization for either marginal is one plus the sup norm of its previous iterate. At coarse (24×20), medium (48×40), and fine (96×80) grids, both methods use linear interpolation and sixteen Numba threads. Solve time is the best of three repetitions after a separate warm-up, so compilation is excluded. The residual compares consumption with the minimum of the Euler-implied and binding-constraint levels at off-grid queries, so no constraint-bound query is excluded. The queries are the axis midpoints below 85 percent of each grid maximum, crossed over both states and every income state.

Both solves use the same regular-grid Euler inversion. On the endogenous mesh, Stage 1 selects b' , inverts to $c = (w^{b'})^{-1/\rho}$, and obtains $\ell = b' + c$. On the regular pre-decision grid the solver again inverts the Euler equation on the interpolated continuation marginal, rather than interpolating the kinked consumption policy. With a' fixed, the Stage 1 policy supplies the evaluation point. One application of the inversion matches NEGM and the curvilinear interpolant off-grid, whereas a second, diagnostic application lowers on-grid Euler residuals to 10^{-6} to 10^{-7} without changing off-grid accuracy.

The grid axes are the model’s own log-spaced asset grids, which concentrate points near the borrowing constraints. The diagnostic therefore exercises the mesh conditions on strongly non-uniform spacing by construction rather than on uniform grids, which are probably the most favorable to them. Each method runs with its native boundary behavior, so neither receives an out-of-hull correction. To measure how often that behavior is used, I classify every regular Stage-0 target at the converged fixed point as inside or outside the polygon bounded by the verified fold-free bilinear mesh. Table B.1 reports the pooled count over all three income states at each grid. The replication CSV also reports each income state separately. Boundary points count as interpolation, so the reported outside share measures boundary-slope extrapolation rather than contact with the mesh boundary.

Table B.1: Stage-0 extrapolation frequency at the converged fixed point: regular (b, a) targets outside the endogenous mesh, pooled over income states. Boundary points count as interpolation.

Grid	Outside / total	Outside (percent)
(24, 20)	1 / 1440	0.069
(48, 40)	1 / 5760	0.017
(96, 80)	1 / 23040	0.004

Only one target lies outside the mesh at every resolution. Because the number of targets grows with the grid, the outside share falls from 0.069 to 0.004 percent. This is a fixed-point result only; the iteration path was not recorded.

ENGINE’s median $\log_{10}$ robust Euler error falls from -1.79 to -2.47 to -3.14, about fivefold per grid doubling. The 95th percentile remains four to six times the median, and the fine-grid worst-query error is below 1 percent. Both structured interpolants reach these levels on this metric, so linear interpolation at these grid densities probably sets them on this model, in which each of the two stacked first-order conditions uses marginal values interpolated at off-grid midpoints. Inside the 85-percent interior query set, both passes bracket rather than extrapolate. In the complementary high-wealth band, the worst residual improves monotonically from -1.08 to -1.69 to -2.32 in $\log_{10}$ units. The band median, however, is not monotone, moving from -2.24 to -1.74 before falling to -2.75. At every resolution the band maximum is no larger than

the interior maximum to within a fraction of a percent, with the two methods within one percent of each other.

With a 192×160 solution as the benchmark, median consumption-policy error falls from 9.3 to 3.7 to 0.9 percent and the worst query from 18 to 3.3 percent. The median robust Euler residual is about six to twelve times smaller than the median policy error because the policy is locally optimal given its grid’s continuation value. Extending both state-grid maxima by half barely moves the typical medium-grid policy (a median of 0.05 percent), but the 4.3 percent maximum rules out treating grid placement as irrelevant at every query. Table B.2 reports both checks.

Table B.2: Consumption-policy refinement and upper-bound sensitivity, as relative differences in percent over a common set of coarse-grid interior midpoints and all income states. Refinement rows compare each grid with the 192×160 benchmark; the bounds row compares the medium grid with both state-grid maxima scaled by 1.5.

Exercise	Grid	Median	p95	Maximum
Refinement	(24, 20)	9.25	17.84	18.47
Refinement	(48, 40)	3.69	8.35	8.75
Refinement	(96, 80)	0.89	2.77	3.31
Bounds $\times 1.5$	(48, 40)	0.05	0.39	4.27

Section 4’s pass-order reversal comparison uses this same protocol, with each pass order solved to convergence; the replication archive records its results at every grid.

The component table (Table B.3) assigns about one percent of serial time to expectations, a little over one third to the two stage inversions, five percent to ENGINE, and the remainder to the regular-grid Euler inversion and envelope updates. Components are timed on one thread at the Section 3 grids with the full-solve diagnostic’s code, whose ENGINE configuration reproduces the paper’s EGMⁿ solver exactly. Summed, they track Section 3’s timings within two percent at every grid.

Table B.3: Runtime in seconds by component of the EGMⁿ backward solve, summed over a converged one-thread solve (shares in parentheses). Stage inversions are Stage 1’s consumption EGM and the Stage 0 inversion; mesh-to-grid interpolation is the ENGINE regrid; regular-grid inversion is corner selection, Euler inversion, and envelope updates.

Grid	Expectation	Stage inversions	Mesh-to-grid	Regular-grid inversion
(40, 60)	0.01 (1%)	0.28 (36%)	0.04 (5%)	0.45 (58%)
(80, 120)	0.03 (1%)	1.15 (35%)	0.15 (5%)	1.92 (59%)
(160, 240)	0.10 (1%)	4.67 (35%)	0.63 (5%)	8.09 (60%)

The Auclert et al. (2021) block also has a native hybrid NEGM/G2EGM solver, at the version pinned in the replication archive. At matched grids and tolerance (Table B.4), EGMⁿ uses fewer iterations (395 to 424 versus 464 to 478) and has off-grid robust Euler medians within fifty percent of the native solver’s, smaller at the smallest grid and larger at the two finer grids. The native solver’s distributed array code runs 3.0 to 4.6 times faster than one-thread EGMⁿ, while sixteen-thread EGMⁿ is about five percent faster on the coarse grid and faster by factors of 1.4 and 1.7 on the two finer grids. Its EGMⁿ times, from a separate run of the configuration in Table 1, lie within three percent of that table’s. The native benchmark does not isolate the method. For that comparison, Section 3 uses NEGM with matched threading and shared components.

References

- Auclert, A., Bardóczy, B., Rognlie, M., Straub, L., 2021. Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. *Econometrica* 89, 2375–2408. doi:[10.3982/ECTA17434](https://doi.org/10.3982/ECTA17434).
- De Nardi, M., 2004. Wealth Inequality and Intergenerational Links. *The Review of Economic Studies* 71, 743–768. doi:[10.1111/j.1467-937X.2004.00302.x](https://doi.org/10.1111/j.1467-937X.2004.00302.x).

Table B.4: Time in seconds and converged-policy accuracy: the two-asset block’s native solver versus EGM^n , at matched grids, tolerance, and linear interpolation. The native solver is serial; parentheses give its time over EGM^n ’s at sixteen threads. Iterations are backward steps to the shared tolerance; accuracy is the off-grid robust Euler median ($\log_{10}$), EGM^n first.

Grid	Native, serial	EGM^n		Iterations $\text{EGM}^n/\text{native}$	Accuracy $\text{EGM}^n/\text{native}$
		1 thread	16 threads		
(40, 60)	0.25	0.76	0.24 (1.1x)	395/478	−2.86/−2.78
(80, 120)	0.78	3.22	0.55 (1.4x)	410/467	−3.44/−3.58
(160, 240)	2.94	13.41	1.74 (1.7x)	424/464	−4.21/−4.36

Hall, R.E., Jones, C.I., 2007. The Value of Life and the Rise in Health Spending. *The Quarterly Journal of Economics* 122, 39–72. doi:[10.1162/qjec.122.1.39](https://doi.org/10.1162/qjec.122.1.39).

Ludwig, A., Schön, M., 2018. Endogenous Grids in Higher Dimensions: Delaunay Interpolation and Hybrid Methods. *Computational Economics* 51, 463–492. doi:[10.1007/s10614-016-9611-2](https://doi.org/10.1007/s10614-016-9611-2).

Powell, W.B., 2011. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley Series in Probability and Statistics. 2nd ed., Wiley, Hoboken, NJ. doi:[10.1002/9781118029176](https://doi.org/10.1002/9781118029176).

White, M.N., 2015. The method of endogenous gridpoints in theory and practice. *Journal of Economic Dynamics and Control* 60, 26–41. doi:[10.1016/j.jedc.2015.08.001](https://doi.org/10.1016/j.jedc.2015.08.001).