

The Method of Moderation

Christopher D. Carroll^{a,b}, Karsten Chipeniuk^c, Alan Lujan^{a,b,*}, Kiichi Tokuoka^d, Weifeng Wu^e

^a*Johns Hopkins University, Krieger School of Arts and Sciences Baltimore, MD,*

^b*Econ-ARK,*

^c*Reserve Bank of New Zealand,*

^d*Japanese Ministry of Finance,*

^e*Fannie Mae, Washington, DC,*

Abstract

In a risky world, a pessimist assumes the worst will happen, while someone who ignores risk altogether is an optimist. Consumption is mathematically simple for both, since each behaves as if the world were riskless. A realist, who responds optimally to risk, faces a much harder problem but (under standard conditions) spends somewhere between the pessimist and the optimist. We use this fact to redefine the space in which the realist searches for optimal consumption rules. The resulting solution accurately represents the consumption rule over the entire range of feasible wealth with remarkably few computations.

Keywords: Dynamic Stochastic Optimization, Consumption-Saving Models, Numerical Methods

JEL classification: D14; C61; G11

1. Introduction

Solving a consumption-saving problem using numerical methods requires the modeler to choose how to represent a policy function. In the stochastic case, where analytical solutions are generally not available, a common approach is to use low-order polynomial splines that exactly match the function at a finite set of gridpoints, and then to assume that interpolated or extrapolated versions of that spline represent the function well at the continuous infinity of unmatched points. [Carroll \(2006\)](#) developed the endogenous gridpoints method (EGM), which has become a standard tool for computing consumption at gridpoints determined endogenously using the Euler equation.

Unfortunately, this endogenous gridpoints solution is not very well-behaved outside the original range of gridpoints, though other common solution methods are no better outside their own predefined ranges. [Figure 1](#) demonstrates

*Corresponding author

Email addresses: ccarroll@jhu.edu (Christopher D. Carroll), alujan@jhu.edu (Alan Lujan)

the point. The figure shows the approximated precautionary component of saving, the amount by which the realist consumes less than an optimist with the same expected income path. Theory proves that precautionary saving is always positive, yet the linearly extrapolated numerical approximation eventually predicts negative precautionary saving. Under uncertainty, however, the consumption-saving rule must be evaluated outside *any* prespecified grid, because large positive shocks push a sufficiently wealthy individual’s next-period assets beyond the grid boundaries.

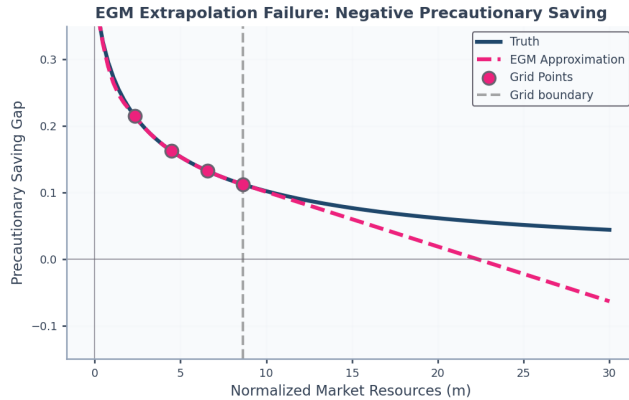


Figure 1: For Large Enough Resources m , Predicted Precautionary Saving is Negative (Oops!)

This error cannot be fixed by extending the upper gridpoint. While extrapolation techniques can prevent this from being fatal, the problem is often dealt with using inelegant methods whose implications for accuracy are difficult to gauge.

This paper argues that, in the standard consumption problem, a better approach is to rely upon the fact that without uncertainty, the optimal consumption function has a simple analytical solution. The key insight is that, under standard assumptions, the consumer who faces an uninsurable labor income risk will consume less than a consumer with the same path for expected income but who does not perceive any uncertainty. Following [Leland \(1968\)](#), [Sandmo \(1970\)](#), and [Kimball \(1990\)](#), the ‘realist’ consumer, who *does* perceive the risks, will engage in ‘precautionary saving,’ so the perfect foresight riskless solution provides an upper bound to the solution that will actually be optimal. A lower bound is provided by the behavior of a consumer who has the subjective belief that the future level of income will be the worst that it can possibly be. This consumer, too, behaves according to the convenient analytical perfect foresight solution, but his certainty is that of a pessimist perfectly confident in his pessimism.

We build on bounds for the consumption function and limiting MPCs established by [Stachurski and Toda \(2019\)](#), [Ma et al. \(2020\)](#), [Carroll \(2009\)](#), and [Ma and Toda \(2021\)](#) in buffer-stock theory. Using results from [Carroll and](#)

Shanker (2024), we show how to use these bounds to constrain the shape and characteristics of the solution to the ‘realist’ problem characterized by Carroll (1997). Imposition of these constraints clarifies and speeds the solution of the realist’s problem. For comparison, we use the endogenous gridpoints method (Carroll, 2006) as our benchmark, which computes consumption at gridpoints determined endogenously using the Euler equation.

After presenting the method in the baseline case and documenting its accuracy, we collect its refinements and extensions, a tighter theoretical bound, the value function, and an extension to rate-of-return risk, in the appendix.

2. The Realist’s Problem

Consider a consumer who correctly perceives all risks. The consumer has CRRA utility over consumption c with risk aversion parameter $\rho > 0$:

$$\mathbf{u}(c) = \begin{cases} \frac{c^{1-\rho}}{1-\rho} & \text{if } \rho \neq 1 \\ \log c & \text{if } \rho = 1. \end{cases} \quad (1)$$

This utility function satisfies prudence ($\mathbf{u}''' > 0$), which induces precautionary saving. The consumer maximizes expected lifetime utility, discounted by the time preference factor β :

$$\max_{c_t, a_t} \mathbb{E}_t \left[\sum_{n=0}^{T-t} \beta^n \mathbf{u}(c_{t+n}) \right] \quad (2)$$

subject to $c_t + a_t = m_t$, where m denotes market resources and a denotes assets. We focus on resources of the form $m_{t+1} = a_t R_{t+1} + y_{t+1}$, where R_{t+1} denotes the interest rate, and y_{t+1} labor income. Initially we take R_{t+1} to be deterministic, and relax this later.

While our method can be adapted to a range of stochastic labor income processes, to fix ideas we suppose income evolves via the Friedman-Muth process (Friedman (1957) distinguished permanent from transitory income; Muth (1960) provided the stochastic framework). That is, $y_{t+1} = p_{t+1} \xi_{t+1}$ where p denotes permanent labor income and ξ_{t+1} a transitory component. Permanent income growth is given by $p_{t+1} = p_t \mathcal{G}_{t+1}$, where the gross growth factor $\mathcal{G}_{t+1} = G_{t+1} \psi_{t+1}$. Here G_{t+1} is the deterministic permanent income growth factor, while ψ_{t+1} are permanent shocks with mean unity and bounded support $[\underline{\psi}, \bar{\psi}]$ where $0 < \underline{\psi} \leq 1 \leq \bar{\psi} < \infty$. Transitory shocks ξ_{t+1} take value 0 with probability $\wp > 0$ (unemployment) or $\theta_{t+1}/(1-\wp)$ otherwise, with $\mathbb{E}_t[\theta_{t+1}] = 1$.

This problem can be rewritten (see Carroll (2020b) for a proof) in a more convenient form in which choice and state variables are normalized by the level of permanent income, e.g., replacing m_t with m_t/p_t . When that is done, the Bellman equation for the value function \mathbf{v} of the transformed problem is

$$\mathbf{v}_t(m_t) = \max_{c_t, a_t} \left(\mathbf{u}(c_t) + \beta \mathbb{E}_t[\mathcal{G}_{t+1}^{1-\rho} \mathbf{v}_{t+1}(m_{t+1})] \right) \quad (3)$$

subject to the normalized transition $m_{t+1} = (R/G_{t+1})a_t + \xi_{t+1}$, with Euler equation $\mathbf{u}'(c_t) = \beta R \mathbb{E}_t[\mathcal{G}_{t+1}^{-\rho} \mathbf{u}'(c_{t+1})]$.

Carroll and Shanker (2024) gives conditions for a finite solution of the problem with a Friedman-Muth process. Consider the case of time-invariant G_t , ψ_t , and R_t , and define the absolute patience factor $\Phi \equiv (\beta R)^{1/\rho}$. Then a finite solution requires: (i) finite-value-of-autarky condition (FVAC) $0 < \beta G^{1-\rho} \mathbb{E}[\psi^{1-\rho}] < 1$, (ii) absolute-impatience condition (AIC) $\Phi < 1$, (iii) return-impatience condition (RIC) $\Phi/R < 1$, (iv) growth-impatience condition (GIC) $\Phi/G < 1$, and (v) finite-human-wealth condition (FHWC) $G/R < 1$. These patience conditions ensure the consumption bounds and limiting MPCs used in our method.

For expositional simplicity, in the numerical results that follow, we assume $\rho \neq 1$ and set $G = 1$, $\psi = 1$ (no permanent income growth or shocks). The figures display the next-to-last period of a finite horizon problem, where the last-period analytical solution provides exact boundary conditions. The method extends naturally to the infinite-horizon case and to general parameter configurations.

3. The Method of Moderation

3.1. The Optimist, the Pessimist, and the Realist

As a preliminary to our solution, define \bar{h}^1 as end-of-period human wealth (the present discounted value of future labor income) for a perfect foresight version of the problem of a ‘risk optimist:’ a consumer who believes with perfect confidence that the shocks will always take their expected value of 1, $\xi_{t+n} = \mathbb{E}[\xi] = 1 \forall n > 0$. The solution to a perfect foresight problem of this kind takes the form

$$\bar{\mathbf{c}}(m) = (m + \bar{h})\underline{\kappa} \quad (4)$$

for a constant minimal marginal propensity to consume $\underline{\kappa}$.² We similarly define \underline{h}^3 as ‘minimal human wealth,’ the present discounted value of labor income if the shocks were to take on their worst value in every future period $\xi_{t+n} = \underline{\xi} \forall n > 0$. We refer to a consumer who expects to encounter this sequence of shocks as a ‘pessimist’. Their consumption decision rule is given by

$$\underline{\mathbf{c}}(m) = (m + \underline{h})\underline{\kappa}. \quad (5)$$

We will call a ‘realist’ the consumer who correctly perceives the true probabilities of the future risks and optimizes accordingly.

¹Setting $\xi_{t+n} = 1$ (the optimist’s assumption), human wealth in infinite-horizon is $\bar{h} = G/(R - G)$ if $R > G$. When $G = 1$, $\bar{h} = 1/(R - 1)$.

²The MPC of the perfect foresight consumer: infinite-horizon $\underline{\kappa} = 1 - \Phi/R$.

³Setting $\xi_{t+n} = \underline{\xi} \forall n > 0$, minimal human wealth is $\underline{h} = \underline{\xi}G/(R - G)$ if $R > G$. When $\underline{\xi} = 0$, $\underline{h} = 0$.

A first useful point is that, for the realist, a lower bound for the level of market resources is the natural borrowing constraint $\underline{m} = -\underline{h}$ derived by [Aiya-gari \(1994\)](#) and [Huggett \(1993\)](#), because if m equalled this value then there would be a positive finite chance (however small) of receiving $\xi_{t+n} = \underline{\xi}$ in every future period, which would require the consumer to set c to zero in order to guarantee that the intertemporal budget constraint holds. Since consumption of zero yields infinite marginal utility, [Zeldes \(1989\)](#) and [Deaton \(1991\)](#) show that the solution to the realist consumer’s problem is not well defined for values of $m \leq \underline{m}$, and the limiting value of the realist’s c is zero as $m \downarrow \underline{m}$ (established in [Carroll and Shanker \(2024\)](#)).

It is convenient to define ‘excess’ market resources as the amount by which actual resources exceed the lower bound, and ‘excess’ human wealth as the amount by which mean expected human wealth exceeds guaranteed minimum human wealth:

$$\begin{aligned}\Delta m &= m + \overbrace{\underline{h}}^{=-\underline{m}} \\ \Delta h &= \bar{h} - \underline{h}.\end{aligned}\tag{6}$$

We now rewrite the optimal consumption rules for the two perfect foresight problems in terms of excess resources and human wealth. The ‘pessimist’ perceives human wealth to be equal to its minimum feasible value \underline{h} with certainty, and so consumes a constant fraction of current excess resources

$$\underline{c}(m) = \Delta m \underline{\kappa}.\tag{7}$$

The ‘optimist,’ on the other hand, pretends that there is no uncertainty about future income, and therefore consumes the same fraction out of current excess resources *and* excess human wealth

$$\bar{c}(m) = (\Delta m + \Delta h) \underline{\kappa}.\tag{8}$$

3.2. The Moderation Ratio

The pessimist expects the worst possible income in every future period with certainty, and must finance all future consumption through saving; this generates the most aggressive saving behavior. A realist, by contrast, can reoptimize as uncertainty resolves each period, so need not prepare for the worst with certainty. At the same time, the adverse outcome remains possible, so even a wealthy realist maintains some precautionary saving: a sufficiently well-off individual is nearly (but never completely) self-insured, and thus mostly smooths consumption like the optimist. The realist therefore consumes strictly more than the pessimist but strictly less than the optimist at every wealth level, as shown in [Figure 2](#).

The proof is more difficult than might be imagined, but the necessary work is done in [Carroll and Shanker \(2024\)](#) so we will take the proposition as a fact:

$$\underline{c}(\underline{m} + \Delta m) < \hat{c}(\underline{m} + \Delta m) < \bar{c}(\underline{m} + \Delta m)\tag{9}$$

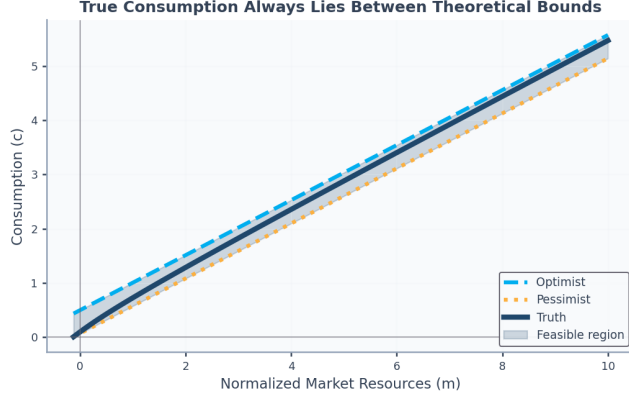


Figure 2: Moderation Illustrated: $\underline{c} < \hat{c} < \bar{c}$

Subtracting $\underline{c}(m + \Delta m)$ in each of these inequalities and using equations (7) and (8) gives

$$\begin{aligned}
 0 &< \hat{c}(m + \Delta m) - \underline{c}(m + \Delta m) < \Delta h \underline{\kappa} \\
 0 &< \underbrace{\left(\frac{\hat{c}(m + \Delta m) - \underline{c}(m + \Delta m)}{\Delta h \underline{\kappa}} \right)}_{\equiv \omega} < 1
 \end{aligned} \tag{10}$$

where the fraction in the middle of the last inequality is the moderation ratio measuring how close the realist's consumption is to the optimist's behavior (the numerator is the gap between the realist and pessimist) relative to the maximum possible gap between optimist and pessimist. When $\omega = 0$, the realist behaves like the pessimist (maximum precautionary saving); when $\omega = 1$, the realist behaves like the optimist (no precautionary saving). [Carroll and Kimball \(1996\)](#) and [Carroll and Shanker \(2024\)](#) establish that under bounded shocks, $\omega \in (0, 1)$ strictly for all $m > \underline{m}$. Defining $\mu = \log \Delta m$ (which can range from $-\infty$ to ∞), the object in the middle of the last inequality is

$$\omega(\mu) \equiv \left(\frac{\hat{c}(\underline{m} + e^\mu) - \underline{c}(\underline{m} + e^\mu)}{\Delta h \underline{\kappa}} \right), \tag{11}$$

and we now define

$$\begin{aligned}
 \chi(\mu) &= \log \left(\frac{\omega(\mu)}{1 - \omega(\mu)} \right) \\
 &= \log(\omega(\mu)) - \log(1 - \omega(\mu))
 \end{aligned} \tag{12}$$

which has the virtue that it is *asymptotically linear* in the limit as μ approaches $+\infty$.⁴ Since $\omega \in (0, 1)$, the ratio $\omega/(1 - \omega)$ is the odds ratio, and χ is

⁴Under the GIC, $\chi(\mu)$ is asymptotically linear with slope $\alpha \geq 0$ as $\mu \rightarrow +\infty$. We ex-

the log odds ratio, the same transformation that underpins logit regression in econometrics. The logit maps $\omega \in (0, 1)$ to $\chi \in (-\infty, \infty)$ with inverse sigmoid $\omega = 1/(1 + \exp(-\chi))$; log maps $(m - \underline{m}) \in (0, \infty)$ to $\mu \in (-\infty, \infty)$ with inverse $\Delta m = \exp(\mu)$.

Given χ , the consumption function can be recovered from

$$\hat{c} = \underline{c} + \overbrace{\frac{1}{1 + \exp(-\chi)}}^{=\omega} \Delta h \underline{\kappa}. \quad (13)$$

Thus, the method of moderation is to calculate χ at the points μ corresponding to the log of the Δm points defined above, and then using these to construct an interpolating approximation $\hat{\chi}$ from which we indirectly obtain our approximated consumption rule \hat{c} (an approximation to the true \hat{c}) by substituting $\hat{\chi}$ for χ in equation (13).

Because this method relies upon the fact that the problem is easy to solve if the decision maker has unreasonable views (either in the optimistic or the pessimistic direction), and because the correct solution is always between these immoderate extremes, we call our solution procedure the ‘method of moderation.’

Results are shown in Figure 3; a reader with very good eyesight might be able to detect the barest hint of a discrepancy between the Truth and the Approximation at the far right-hand edge of the figure, a stark contrast with the calamitous divergence evident in Figure 1.

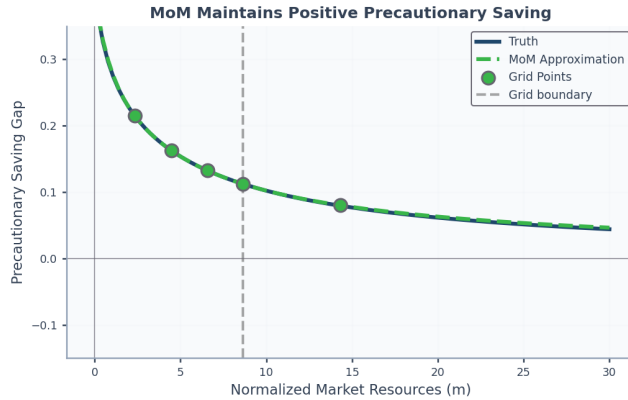


Figure 3: Extrapolated \hat{c} Constructed Using the Method of Moderation

trapolate χ linearly using the boundary slope, preserving $\omega \in (0, 1)$ and hence $\underline{c} < \hat{c} < \bar{c}$ throughout the extrapolation domain.

Table 1: Maximum absolute approximation errors by interval. Orders of magnitude in parentheses.

Method	$[m_0, m_1]$	$[m_1, m_2]$	$[m_2, m_3]$	$[m_3, m_4]$	$[m_4, \bar{m}]$
EGM	8.6(-3)	1.8(-4)	2.5(-5)	7.3(-6)	1.1(-1)
MoM	2.9(-3)	4.3(-6)	6.6(-7)	1.3(-7)	2.4(-3)

4. Numerical Accuracy

Table 1 reports the method’s accuracy between each pair of gridpoints m_j, m_{j+1} and for extrapolation out to $\bar{m} = 30$, as the maximum absolute error on a dense subgrid of each interval. Except in the first interval, where it is about three times more accurate, the method of moderation is more than an order of magnitude more accurate than the basic endogenous gridpoints method, complementing related work on solution precision (Chipeniuk, 2020).

The construction extends further in the Appendix. A tighter upper bound holds near the natural borrowing constraint for low-wealth consumers ((D.3)). The same idea approximates the value function: with the inverse value transformation $\bar{\Lambda} = ((1 - \rho)\hat{v})^{1/(1-\rho)}$, we form a value moderation ratio $\hat{\Omega}$, interpolate its logit, and invert to recover $\hat{v} = \mathbf{u}(\hat{\Lambda})$ ((E.5)). It also accommodates rate-of-return risk: with i.i.d. returns and certain income the optimal rule is linear (Merton-Samuelson), so the optimist’s and pessimist’s bounds extend and bracket the realist (Stochastic Rate of Return).

5. Conclusion

The method is not universally applicable: it requires known upper and lower bounds on the ‘true’ solution. But many problems have obvious bounds, and there (as in our consumption example) it can substantially improve the accuracy and stability of solutions. The method is efficient because the transformed moderation ratio is better-behaved than consumption, so a given grid yields substantially more accurate approximations. The gains are largest in the extrapolation region, where standard methods fail most severely.

References

- Aiyagari, S.R., 1994. Uninsured Idiosyncratic Risk and Aggregate Saving. *Quarterly Journal of Economics* 109, 659–684. doi:[10.2307/2118417](https://doi.org/10.2307/2118417).
- Benhabib, J., Bisin, A., 2018. Skewed Wealth Distributions: Theory and Empirics. *Journal of Economic Literature* 56, 1261–1291. doi:[10.1257/jel.20161390](https://doi.org/10.1257/jel.20161390).
- Benveniste, L.M., Scheinkman, J.A., 1979. On the Differentiability of the Value Function in Dynamic Models of Economics. *Econometrica* 47, 727–732. doi:[10.2307/1910417](https://doi.org/10.2307/1910417).

- de Boor, C., 2001. A Practical Guide to Splines. volume 27 of *Applied Mathematical Sciences*. Revised ed., Springer. doi:[10.1007/978-1-4612-6333-3](https://doi.org/10.1007/978-1-4612-6333-3).
- Carroll, C.D., 1997. Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *Quarterly Journal of Economics* 112, 1–55. doi:[10.1162/003355397555109](https://doi.org/10.1162/003355397555109).
- Carroll, C.D., 2006. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters* 91, 312–320. doi:[10.1016/j.econlet.2005.09.013](https://doi.org/10.1016/j.econlet.2005.09.013).
- Carroll, C.D., 2009. Precautionary Saving and the Marginal Propensity to Consume Out of Permanent Income. *Journal of Monetary Economics* 56, 780–790. doi:[10.1016/j.jmoneco.2009.06.016](https://doi.org/10.1016/j.jmoneco.2009.06.016).
- Carroll, C.D., 2020a. CRRA utility with rate risk. techreport. Johns Hopkins University. Lecture notes.
- Carroll, C.D., 2020b. Solving microeconomic dynamic stochastic optimization problems. techreport. Johns Hopkins University. URL: <https://www.econ2.jhu.edu/people/ccarroll/SolvingMicroDSOPs.pdf>.
- Carroll, C.D., Kimball, M.S., 1996. On the Concavity of the Consumption Function. *Econometrica* 64, 981–992. doi:[10.2307/2171853](https://doi.org/10.2307/2171853).
- Carroll, C.D., Shanker, A., 2024. Theoretical Foundations of Buffer Stock Saving (Revise and Resubmit, Quantitative Economics). URL: <https://llorracc.github.io/BufferStockTheory/BufferStockTheory.pdf>. revise and Resubmit, Quantitative Economics.
- Carroll, C.D., Toche, P., 2009. A Tractable Model of Buffer Stock Saving. Working Paper 15265. National Bureau of Economic Research. doi:[10.3386/w15265](https://doi.org/10.3386/w15265).
- Chipeniuk, K.O., 2020. Optimal Grid Selection for the Numerical Solution of Dynamic Stochastic Optimization Problems. *Computational Economics* 56, 883–928. doi:[10.1007/s10614-019-09953-4](https://doi.org/10.1007/s10614-019-09953-4).
- Chipeniuk, K.O., Katz, N.H., Walker, T.B., 2021. Households, Auctioneers, and Aggregation. *European Economic Review* 141, 103997. doi:[10.1016/j.euroecorev.2021.103997](https://doi.org/10.1016/j.euroecorev.2021.103997).
- Deaton, A., 1991. Saving and Liquidity Constraints. *Econometrica* 59, 1221–1248. doi:[10.2307/2938366](https://doi.org/10.2307/2938366).
- Friedman, M., 1957. *Theory of the Consumption Function*. Princeton University Press, Princeton, NJ. doi:[10.1515/9780691188485](https://doi.org/10.1515/9780691188485).
- Fritsch, F.N., Butland, J., 1984. A Method for Constructing Local Monotone Piecewise Cubic Interpolants. *SIAM Journal on Scientific and Statistical Computing* 5, 300–304. doi:[10.1137/0905021](https://doi.org/10.1137/0905021).

- Fritsch, F.N., Carlson, R.E., 1980. Monotone Piecewise Cubic Interpolation. *SIAM Journal on Numerical Analysis* 17, 238–246. doi:[10.1137/0717021](https://doi.org/10.1137/0717021).
- Huggett, M., 1993. The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies. *Journal of Economic Dynamics and Control* 17, 953–969. doi:[10.1016/0165-1889\(93\)90024-m](https://doi.org/10.1016/0165-1889(93)90024-m).
- Hyman, J.M., 1983. Accurate Monotonicity-Preserving Cubic Interpolation. *SIAM Journal on Scientific and Statistical Computing* 4, 645–654. doi:[10.1137/0904045](https://doi.org/10.1137/0904045).
- Judd, K.L., Maliar, L., Maliar, S., 2017. Lower Bounds on Approximation Errors to Numerical Solutions of Dynamic Economic Models. *Econometrica* 85, 991–1020. doi:[10.3982/ecta12791](https://doi.org/10.3982/ecta12791).
- Kimball, M.S., 1990. Precautionary Saving in the Small and in the Large. *Econometrica* 58, 53–73. doi:[10.2307/2938334](https://doi.org/10.2307/2938334).
- Leland, H.E., 1968. Saving and Uncertainty: The Precautionary Demand for Saving. *Quarterly Journal of Economics* 82, 465–473. doi:[10.2307/1879518](https://doi.org/10.2307/1879518).
- Ma, Q., Stachurski, J., Toda, A.A., 2020. The Income Fluctuation Problem and the Evolution of Wealth. *Journal of Economic Theory* 187, 105003. doi:[10.1016/j.jet.2020.105003](https://doi.org/10.1016/j.jet.2020.105003).
- Ma, Q., Toda, A.A., 2021. A Theory of the Saving Rate of the Rich. *Journal of Economic Theory* 192, 105193. doi:[10.1016/j.jet.2021.105193](https://doi.org/10.1016/j.jet.2021.105193).
- Merton, R.C., 1969. Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case. *Review of Economics and Statistics* 51, 247–257. doi:[10.2307/1926560](https://doi.org/10.2307/1926560).
- Merton, R.C., 1971. Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory* 3, 373–413. doi:[10.1016/0022-0531\(71\)90038-X](https://doi.org/10.1016/0022-0531(71)90038-X).
- Milgrom, P., Segal, I., 2002. Envelope Theorems for Arbitrary Choice Sets. *Econometrica* 70, 583–601. doi:[10.1111/1468-0262.00296](https://doi.org/10.1111/1468-0262.00296).
- Muth, J.F., 1960. Optimal Properties of Exponentially Weighted Forecasts. *Journal of the American Statistical Association* 55, 299–306. doi:[10.1080/01621459.1960.10482064](https://doi.org/10.1080/01621459.1960.10482064).
- Samuelson, P.A., 1969. Lifetime Portfolio Selection by Dynamic Stochastic Programming. *Review of Economics and Statistics* 51, 239–246. doi:[10.2307/1926559](https://doi.org/10.2307/1926559).
- Sandmo, A., 1970. The Effect of Uncertainty on Saving Decisions. *Review of Economic Studies* 37, 353–360. doi:[10.2307/2296725](https://doi.org/10.2307/2296725).

Santos, M.S., 2000. Accuracy of Numerical Solutions Using the Euler Equation Residuals. *Econometrica* 68, 1377–1402. doi:[10.1111/1468-0262.00165](https://doi.org/10.1111/1468-0262.00165).

Stachurski, J., Toda, A.A., 2019. An Impossibility Theorem for Wealth in Heterogeneous-Agent Models with Limited Heterogeneity. *Journal of Economic Theory* 182, 1–24. doi:[10.1016/j.jet.2019.04.001](https://doi.org/10.1016/j.jet.2019.04.001). provides an explicit linear lower bound on consumption; Corrigendum 2020.

Zeldes, S.P., 1989. Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy* 97, 305–346. doi:[10.1086/261605](https://doi.org/10.1086/261605).

Appendix A. Patience Conditions Details

Each patience condition from the main text controls a distinct way the problem could misbehave. The FVAC $0 < \beta G^{1-\rho} \mathbb{E}[\psi^{1-\rho}] < 1$ guarantees that even autarky, saving nothing and consuming income as it arrives, delivers finite expected discounted utility, so the consumer has a reason to value resources at all. The AIC $\Phi < 1$ rules out indefinite deferral of consumption: under certainty the marginal utility of consuming now exceeds the discounted marginal utility of consuming later. Two further conditions bound wealth from opposite directions. The RIC $\Phi/R < 1$ holds asset growth below the patience-adjusted discount rate, so wealth cannot explode; the GIC $\Phi/G < 1$ holds consumption growth below permanent-income growth, which is what pins down a finite target wealth ratio. Finally, the FHWC $G/R < 1$ keeps the present value of future labor income finite. Where these conditions fail, behavior changes qualitatively: when all hold, the consumer runs a buffer stock around a target wealth ratio; when the AIC fails, consumption grows without bound; when the GIC fails but the RIC still holds, wealth grows without bound (Carroll, 1997, 2020b; Carroll and Shanker, 2024).

Appendix B. Human Wealth Formulas

The optimist’s human wealth (assuming $\xi_{t+n} = 1 \forall n > 0$) can be computed three ways: backward recursion $\bar{h}_T = 0$, $\bar{h}_t = (G/R)(1 + \bar{h}_{t+1})$; forward sum $\bar{h}_t = \sum_{n=1}^{T-t} (G/R)^n$; or infinite-horizon $\bar{h} = G/(R - G)$ when $R > G$. With $G = 1$, $\bar{h} = 1/(R - 1)$.

The pessimist’s human wealth (assuming $\xi_{t+n} = \underline{\xi} \forall n > 0$) follows similarly: backward recursion $\underline{h}_T = 0$, $\underline{h}_t = (G/R)(\underline{\xi} + \underline{h}_{t+1})$; forward sum $\underline{h}_t = \underline{\xi} \sum_{n=1}^{T-t} (G/R)^n$; or infinite-horizon $\underline{h} = \underline{\xi} G/(R - G)$. When $\underline{\xi} = 0$ (unemployment), $\underline{h} = 0$.

Appendix C. Marginal Propensity to Consume Formulas

The minimal MPC (perfect foresight consumer with horizon $T - t$) has three forms (Carroll, 2009): backward recursion $\underline{\kappa}_t = \underline{\kappa}_{t+1}/(\underline{\kappa}_{t+1} + \Phi/R)$ with $\underline{\kappa}_T = 1$;

forward sum $\underline{\kappa}_t = (\sum_{n=0}^{T-t} (\Phi/R)^n)^{-1}$; or infinite-horizon $\underline{\kappa} = 1 - \Phi/R = 1 - (\mathbf{R}\beta)^{1/\rho}/R$.

The maximal MPC (Carroll and Toche, 2009) satisfies backward recursion $\bar{\kappa}_t = \bar{\kappa}_{t+1}/(\bar{\kappa}_{t+1} + \varphi^{1/\rho}\Phi/R)$ with $\bar{\kappa}_T = 1$; forward sum $\bar{\kappa}_t = (\sum_{n=0}^{T-t} (\varphi^{1/\rho}\Phi/R)^n)^{-1}$; or infinite-horizon $\bar{\kappa} = 1 - \varphi^{1/\rho}(\Phi/R)$.

Appendix D. A Tighter Upper Bound

The method in the main text does not guarantee that the approximation respects $\hat{c}(m) < \bar{\kappa}\Delta m$, where $\bar{\kappa}$ is the MPC at the natural borrowing constraint; near the constraint the optimist's bound is loose because it is calibrated to the low MPC that prevails at high wealth. Carroll and Shanker (2024) derives the maximal MPC $\bar{\kappa} = 1 - \varphi^{1/\rho}(\Phi/R)$, where φ is the unemployment probability of Carroll and Toche (2009), extending the limiting-MPC formulas of Ma and Toda (2021). Strict concavity implies $\hat{c}(m) < \bar{\kappa}\Delta m$ for low wealth, while the optimist's bound $\hat{c}(m) < (\Delta m + \Delta h)\underline{\kappa}$ is tighter for high wealth.

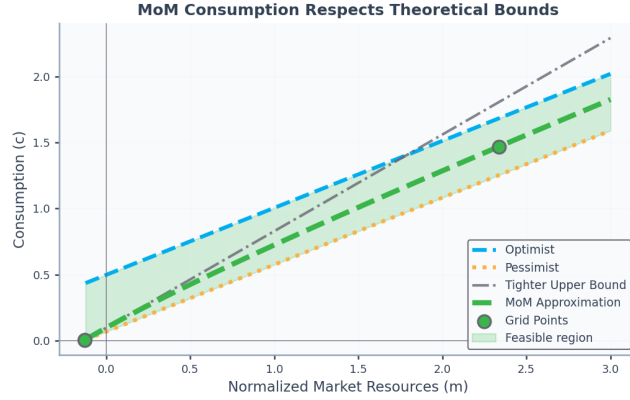


Figure D.4: A Tighter Upper Bound

As Figure D.4 shows, the two upper bounds intersect at the cusp point m^* where

$$\begin{aligned}
 (\Delta m^* + \Delta h)\underline{\kappa} &= \bar{\kappa}\Delta m^* \\
 \Delta m^* &= \frac{\underline{\kappa}\Delta h}{\bar{\kappa} - \underline{\kappa}} \\
 m^* &= -\underline{h} + \frac{\underline{\kappa}(\bar{h} - \underline{h})}{\bar{\kappa} - \underline{\kappa}},
 \end{aligned} \tag{D.1}$$

where $\Delta m^* \equiv m^* - \underline{m} > 0$ since $\bar{\kappa} > \underline{\kappa}$. For $m \in (\underline{m}, m^*]$, the tighter upper bound yields

$$\begin{aligned}
\Delta m \underline{\kappa} &< \hat{\mathbf{c}}(\underline{m} + \Delta m) &< \bar{\kappa} \Delta m \\
0 &< \hat{\mathbf{c}}(\underline{m} + \Delta m) - \Delta m \underline{\kappa} &< \Delta m (\bar{\kappa} - \underline{\kappa}) \\
0 &< \left(\frac{\hat{\mathbf{c}}(\underline{m} + \Delta m) - \Delta m \underline{\kappa}}{\Delta m (\bar{\kappa} - \underline{\kappa})} \right) &< 1.
\end{aligned} \tag{D.2}$$

This motivates the low-resource moderation ratio, defined for $m \in (\underline{m}, m^*]$ as

$$\dot{\omega}(\boldsymbol{\mu}) = \frac{\hat{\mathbf{c}}(\underline{m} + e^\mu) e^{-\mu} - \underline{\kappa}}{\bar{\kappa} - \underline{\kappa}}. \tag{D.3}$$

Since $e^{-\mu} = 1/\Delta m$, the right-hand side equals $(\hat{\mathbf{c}}/\Delta m - \underline{\kappa})/(\bar{\kappa} - \underline{\kappa})$, which lies in $(0, 1)$ for $m \in (\underline{m}, m^*]$: the lower bound is the minimal MPC $\underline{\kappa}$ and the upper bound is the maximal MPC $\bar{\kappa}$, with strict inequality at the upper end following from $\hat{\mathbf{c}} < \bar{\kappa} \Delta m$. Applying the logit transformation and interpolating as before yields consumption satisfying both upper bounds throughout. For computational robustness we combine the pieces into a three-part approximation: the tighter bound below the cusp, the optimist's bound above, and a Hermite segment (below) bridging the cusp, where the two bounds meet at equal levels but different slopes. Because the Hermite segment is matched to the level and the slope of the adjacent piece at each of its endpoints, the combined consumption function is continuous and differentiable and respects all theoretical constraints.

Appendix E. Value Function Derivation

Under perfect foresight, consumption grows at constant rate equal to the absolute patience factor Φ : $\mathbf{c}_{t+n} = \mathbf{c}_t \Phi^n$. The present discounted value of consumption, discounting the stream at the return R , satisfies $\text{PDV}_t^T(\mathbf{c}) = \sum_{n=0}^{T-t} R^{-n} \mathbf{c}_t \Phi^n = \mathbf{c}_t \sum_{n=0}^{T-t} (\Phi/R)^n$. Dividing by consumption yields the PDV-to-consumption ratio $\mathcal{C}_t^T = \text{PDV}_t^T(\mathbf{c})/\mathbf{c}_t = \sum_{n=0}^{T-t} (\Phi/R)^n = \underline{\kappa}_t^{-1}$, which is unchanged for normalized variables. Defining $\mathcal{C} \equiv \lim_{T \rightarrow \infty} \mathcal{C}_t^T$, this yields the key identity $\mathcal{C} = \underline{\kappa}^{-1}$, connecting the infinite-horizon PDV-to-consumption ratio to the minimal MPC.

The optimist's value function satisfies

$$\begin{aligned}
\bar{\mathbf{v}}_{T-1}(m_{T-1}) &\equiv \mathbf{u}(c_{T-1}) + \beta \mathbf{u}(c_T) \\
&= \mathbf{u}(c_{T-1}) (1 + \beta \Phi^{1-\rho}) \\
&= \mathbf{u}(c_{T-1}) (1 + \Phi/R) \\
&= \mathbf{u}(c_{T-1}) \mathcal{C}_{T-1}^T
\end{aligned} \tag{E.1}$$

The infinite horizon expression becomes

$$\begin{aligned}
\bar{\mathbf{v}}(m) &= \mathbf{u}(\bar{\mathbf{c}}(m))\mathcal{C} \\
&= \mathbf{u}(\bar{\mathbf{c}}(m))\underline{\kappa}^{-1} \\
&= \mathbf{u}((\Delta m + \Delta h)\underline{\kappa})\underline{\kappa}^{-1} \\
&= \mathbf{u}(\Delta m + \Delta h)\underline{\kappa}^{-\rho}.
\end{aligned} \tag{E.2}$$

This can be transformed as

$$\begin{aligned}
\bar{\boldsymbol{\Lambda}} &\equiv ((1 - \rho)\bar{\mathbf{v}})^{1/(1-\rho)} \\
&= c\mathcal{C}^{1/(1-\rho)} \\
&= (\Delta m + \Delta h)\underline{\kappa}^{-\rho/(1-\rho)}.
\end{aligned} \tag{E.3}$$

The transformation $\boldsymbol{\Lambda} \equiv ((1 - \rho)\mathbf{v})^{1/(1-\rho)}$ is the inverse utility \mathbf{u}^{-1} ; for log utility ($\rho = 1$) it becomes $\boldsymbol{\Lambda} = \exp(\mathbf{v})$, the $\rho \rightarrow 1$ limit, so the construction carries over to log utility unchanged.

The pessimist's inverse value follows by the same steps, with $\underline{\mathbf{c}} = \Delta m \underline{\kappa}$:

$$\underline{\boldsymbol{\Lambda}} = \underline{\mathbf{c}}\mathcal{C}^{1/(1-\rho)} = \Delta m \underline{\kappa}^{-\rho/(1-\rho)}, \tag{E.4}$$

so that $\bar{\boldsymbol{\Lambda}} - \underline{\boldsymbol{\Lambda}} = \Delta h \underline{\kappa} \mathcal{C}^{1/(1-\rho)}$, the denominator that normalizes the value moderation ratio below.

For the realist's problem, we define $\hat{\boldsymbol{\Lambda}} = ((1 - \rho)\hat{\mathbf{v}}(m))^{1/(1-\rho)}$. At each m the values are ordered $\underline{\mathbf{v}} < \hat{\mathbf{v}} < \bar{\mathbf{v}}$: the realist's true income process stochastically dominates the pessimist's worst-case income and is dominated, for a risk-averse agent, by the optimist's certain expected income. Because the inverse-value transform is monotonic, the same ordering holds for $\underline{\boldsymbol{\Lambda}} < \hat{\boldsymbol{\Lambda}} < \bar{\boldsymbol{\Lambda}}$, and we define

$$\hat{\boldsymbol{\Omega}}(\boldsymbol{\mu}) = \left(\frac{\hat{\boldsymbol{\Lambda}}(\underline{m} + e^\mu) - \underline{\boldsymbol{\Lambda}}(\underline{m} + e^\mu)}{\Delta h \underline{\kappa} \mathcal{C}^{1/(1-\rho)}} \right) \tag{E.5}$$

and the logit-transformed counterpart:

$$\begin{aligned}
\hat{\boldsymbol{X}}(\boldsymbol{\mu}) &= \log \left(\frac{\hat{\boldsymbol{\Omega}}(\boldsymbol{\mu})}{1 - \hat{\boldsymbol{\Omega}}(\boldsymbol{\mu})} \right) \\
&= \log(\hat{\boldsymbol{\Omega}}(\boldsymbol{\mu})) - \log(1 - \hat{\boldsymbol{\Omega}}(\boldsymbol{\mu}))
\end{aligned} \tag{E.6}$$

Inverting these approximations yields

$$\hat{\boldsymbol{\Lambda}} = \underline{\boldsymbol{\Lambda}} + \overbrace{\left(\frac{1}{1 + \exp(-\hat{\boldsymbol{X}})} \right)}^{=\hat{\boldsymbol{\Omega}}} \Delta h \underline{\kappa} \mathcal{C}^{1/(1-\rho)} \tag{E.7}$$

from which the value function approximation is $\hat{\mathbf{v}} = \mathbf{u}(\hat{\boldsymbol{\Lambda}})$.

Appendix F. Hermite Interpolation

The numerical accuracy of the method of moderation depends critically on the quality of function approximation between gridpoints (Santos, 2000). Our bracketing approach complements work that bounds numerical errors in dynamic economic models (Judd et al., 2017). Although linear interpolation that matches the level of \hat{c} at the gridpoints is simple, Hermite interpolation (Fritsch and Carlson, 1980) offers a considerable advantage.

The moderation ratio derivative measures how quickly the realist approaches the optimist as resources increase. Differentiating (11) with respect to μ we obtain

$$\frac{\partial \omega}{\partial \mu} = \frac{\Delta m (\partial \hat{c} / \partial m - \underline{\kappa})}{\underline{\kappa} \Delta h}. \quad (\text{F.1})$$

Rearranging this yields a moderation form for the marginal propensity to consume:

$$\frac{\partial \hat{c}}{\partial m} = (1 - \eta) \underline{\kappa} + \eta \bar{\kappa} \quad (\text{F.2})$$

where

$$\eta = \frac{\underline{\kappa}}{\bar{\kappa} - \underline{\kappa}} \cdot \frac{\Delta h}{\Delta m} \cdot \partial \omega / \partial \mu. \quad (\text{F.3})$$

Carroll and Shanker (2024) guarantees $\underline{\kappa} \leq \partial \hat{c} / \partial m \leq \bar{\kappa}$ at gridpoints where the Euler equation holds, so $\eta \in [0, 1]$ and the expression above is indeed a convex combination of $\underline{\kappa}$ and $\bar{\kappa}$. At very high wealth, $\eta \rightarrow 0$ and the MPC approaches $\underline{\kappa}$; near the borrowing constraint, $\eta \rightarrow 1$ and the MPC approaches $\bar{\kappa}$.

For Hermite interpolation, compute $\partial \omega / \partial \mu$ at gridpoints, then derive $\partial \chi / \partial \mu = \partial \omega / \partial \mu / [\omega(1 - \omega)]$ for slope data. By matching both the level and the derivative of \hat{c} at the gridpoints, where the derivative is obtained from the envelope condition (Benveniste and Scheinkman, 1979; Milgrom and Segal, 2002) together with the EGM Euler equation, the interpolated consumption rule satisfies the Euler equation exactly at each solved gridpoint. These techniques extend naturally to the value function approximation.

For monotone cubic Hermite schemes (Fritsch and Carlson, 1980; Fritsch and Butland, 1984; de Boor, 2001), theoretical slopes may be adjusted to enforce monotonicity (Hyman, 1983). The Fritsch-Carlson algorithm modifies slopes at local extrema, while Fritsch-Butland uses harmonic mean weighting. Both preserve the shape-preserving property essential for consumption functions that must be strictly increasing.

Appendix G. Stochastic Rate of Return

For i.i.d. returns with $\log \mathbf{R} \sim \mathcal{N}(r + \pi - \sigma_{\mathbf{r}}^2/2, \sigma_{\mathbf{r}}^2)$,⁵ Samuelson (1969); Merton (1969, 1971) showed that for a consumer without labor income (or with perfectly forecastable labor income) the consumption function is linear, with an MPC = $1 - (\beta \mathbb{E}[\mathbf{R}^{1-\rho}])^{1/\rho}$, which is positive under the stochastic return impatience condition $\beta \mathbb{E}[\mathbf{R}^{1-\rho}] < 1$ (the i.i.d.-return analogue of the RIC $\Phi/\mathbf{R} < 1$). See Carroll (2020a); Benhabib and Bisin (2018); Chipeniuk et al. (2021) for extensions. The pessimist and the optimist face certain income but the same stochastic return, so the Merton-Samuelson result applies to both and their consumption functions remain linear. The realist faces both labor income and return risk, and the moderation ratio captures their combined precautionary response. In this case the previous analysis applies once we substitute this MPC for the one that characterizes the perfect-foresight problem without rate-of-return risk. As Figure G.5 shows, consumption remains bounded between the pessimist and the optimist, each of which (for $\rho > 1$) consumes slightly less in the face of return uncertainty; for $\rho < 1$ the effect reverses.

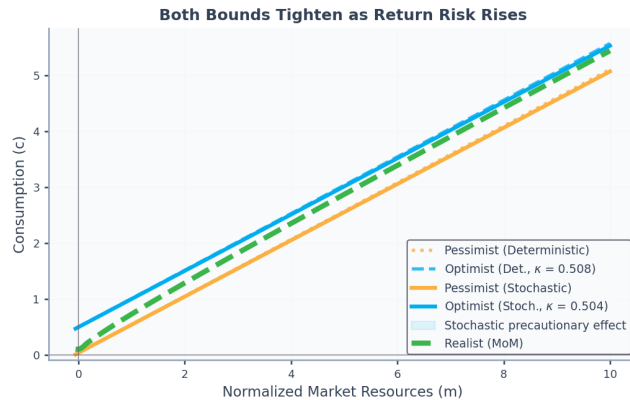


Figure G.5: Effect of Return Uncertainty on Consumption Bounds

The fact that a linear consumption function with an MPC = $1 - (\beta \mathbb{E}[\mathbf{R}^{1-\rho}])^{1/\rho}$ satisfies the Euler equation with i.i.d. returns and no labor income can be derived by the method of undetermined coefficients. In particular, assume that $\bar{c}(m) = m\kappa$, with a time-independent MPC κ to be determined. Substituting this into the Euler equation, we have

⁵Here r is the log risk-free rate and π is the equity premium (the expected excess log return). This parametrization ensures $\mathbb{E}[\mathbf{R}] = \exp(r + \pi)$, so that increasing $\sigma_{\mathbf{r}}^2$ constitutes a mean-preserving spread of the level of the return.

$$\begin{aligned}
1 &= \beta \mathbb{E}_t \left[\mathbf{R}_{t+1} \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} \right] \\
&= \beta \mathbb{E}_t \left[\mathbf{R}_{t+1} \left(\frac{m_{t+1}}{m_t} \right)^{-\rho} \right]
\end{aligned} \tag{G.1}$$

where the second equality uses the assumed form of the consumption function. Since there is no labor income, $m_{t+1} = \mathbf{R}_{t+1}(m_t - c_t)$. Substituting this into the above we obtain

$$1 = \beta \mathbb{E}_t \left[\mathbf{R}_{t+1} (\mathbf{R}_{t+1}(1 - \underline{\kappa}))^{-\rho} \right] \tag{G.2}$$

Solving for $\underline{\kappa}$ and recalling that returns are i.i.d. gives $\underline{\kappa} = 1 - (\beta \mathbb{E}[\mathbf{R}^{1-\rho}])^{1/\rho}$.

In the particular case of lognormal returns, the MPC can be written in closed form. The moment generating function (MGF) of the normal variable $X = \log \mathbf{R}$ provides the key formula. For $X \sim \mathcal{N}(\mu, \sigma^2)$, the MGF is $\mathbb{E}[e^{sX}] = \exp(\mu s + \sigma^2 s^2/2)$. Setting $s = 1 - \rho$ and $\mu = r + \pi - \sigma_{\mathbf{r}}^2/2$ yields⁶

$$\mathbb{E}[\mathbf{R}^{1-\rho}] = \exp \left((1 - \rho) \left(r + \pi - \frac{\sigma_{\mathbf{r}}^2}{2} \right) + \frac{(1 - \rho)^2 \sigma_{\mathbf{r}}^2}{2} \right). \tag{G.3}$$

Simplifying the variance terms: $(1 - \rho)^2 \sigma_{\mathbf{r}}^2/2 - (1 - \rho) \sigma_{\mathbf{r}}^2/2 = (1 - \rho)[(1 - \rho) - 1] \sigma_{\mathbf{r}}^2/2 = -\rho(1 - \rho) \sigma_{\mathbf{r}}^2/2$, giving the final form

$$\mathbb{E}[\mathbf{R}^{1-\rho}] = \exp \left((1 - \rho) \left(r + \pi - \rho \sigma_{\mathbf{r}}^2/2 \right) \right). \tag{G.4}$$

⁶Here we can interpret π as the risk premium, that is, the additional average return from holding a risky asset compared to the risk-free rate r . Adjusting the average log return by the asset volatility ensures that increasing $\sigma_{\mathbf{r}}^2$ constitutes a mean-preserving spread of the level of return.