Efficient and Robust Numerical Maximization of Power Utility

Ludger Hentschel
Versor Investments

John B. Long, Jr.
Simon School of Business, University of Rochester

April 1, 2026

Abstract

We propose a safeguarded method for maximizing power utility that guarantees convergence to the economically valid optimum. In standard implementations, the implicit no-bankruptcy constraint is often omitted, allowing portfolio values to become non-positive in some sample periods. This relaxation fundamentally alters the problem: the objective becomes singular at zero wealth, leading to spurious local optima, unstable convergence, or undefined values depending on the specification of risk aversion. As a result, numerical routines may converge to economically meaningless portfolios that imply bankruptcy.

We address this problem by replacing the singularity at zero wealth with a smooth, curvature-matched extension of the utility function. The safeguarded objective coincides with the original power utility over the economically relevant domain while remaining globally concave and twice differentiable everywhere. The resulting optimization problem admits a unique solution, can be solved reliably from arbitrary starting points, and requires no explicit inequality constraints. When a feasible solution exists, the safeguarded method recovers the exact constrained optimum at substantially lower computational cost.

Contents

Acknowledgements
We gratefully acknowledge financial support from the John M. Olin Foundation.

1 Introduction

Maximizing expected power utility, or constant relative risk aversion (CRRA), conditional on a sample of asset returns, is a standard formulation of the portfolio problem. In practice, solving this problem requires numerical optimization.

Figure 1 illustrates the difficulty of this numerical optimization for a coefficient of relative risk aversion $\gamma = 3$. The top panel shows that the utility function has a large number of local optima; the middle panel shows that the gradient is zero at each of them; and the bottom panel shows that the investor goes bankrupt at every false optimum. Although only the global optimum marked x is economically sensible, standard numerical routines would confirm convergence at all of the local optima. The violation of the implicit no-bankruptcy constraint is not a technical curiosity: any solution that implies negative wealth is economically meaningless.

We propose a simple, fast, and reliable numerical method for maximizing

Figure 1: Power Utility with Risk Aversion of 3

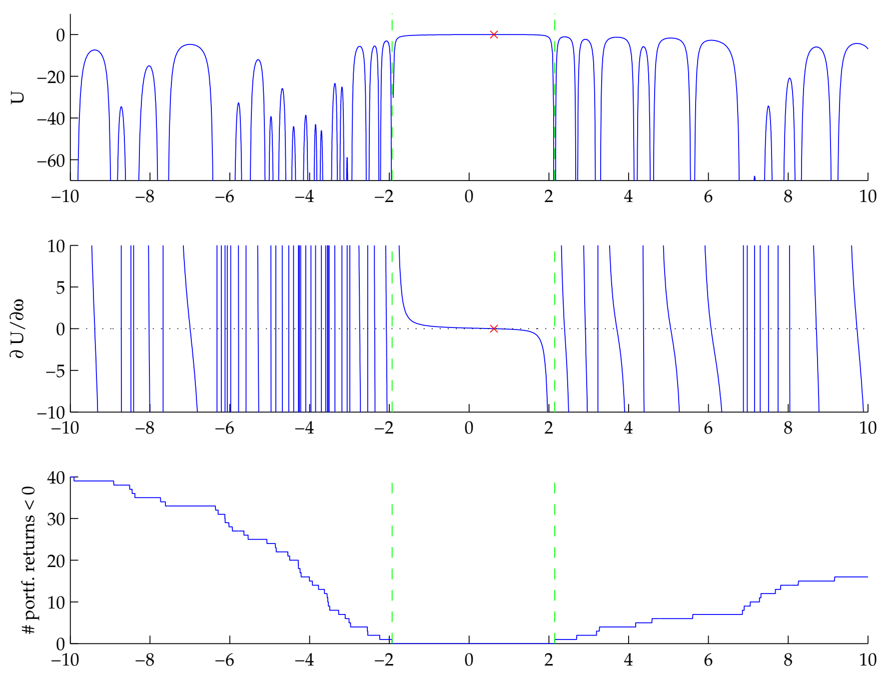

The figure plots power utility with relative risk aversion $\gamma = 3$ from investing a fraction ω in the CRSP value-weighted index and $(1 - \omega)$ in a risk-free asset with annual return of 3.9 percent. The top panel shows the expected utility, $E[U]$; the middle panel shows the derivative of utility with respect to ω, $\partial U/\partial\omega$; the bottom panel shows the number of periods with negative wealth. The true optimum is marked with an x. The other utility peaks are local optima that satisfy the mathematical optimality conditions. However, the associated portfolios violate the no-bankruptcy constraint in at least one period.

The CRSP value-weighted returns are annual returns including dividends from 1926 to 2003.

power utility over a sample of asset returns. The method guarantees that the resulting portfolio satisfies non-negativity of wealth by smoothly pasting a quadratic extension to the true utility function near zero wealth. This safeguarded objective behaves like standard power utility over the economically relevant region but remains continuous, monotonic, and concave even when wealth approaches zero or becomes negative.

Relative to other ways of keeping the numerical search within the feasible domain, such as adding explicit inequality constraints or performing an initial search for a feasible starting point, the safeguarded approach is faster and more stable. It avoids the large systems of constraints required by direct enforcement and eliminates the need to locate a feasible starting portfolio.

After mean-variance optimization, power-utility (CRRA) optimization is arguably the most common portfolio-construction framework in empirical finance.[1] Examples include Grauer and Hakansson (1987), Algoet and Cover (1988), Cover (1991), Brandt (1999), Barberis (2000), Aït-Sahalia and Brandt (2001), Hentschel and Long (2004), Driessen and Maenhout (2007), Malamud (2014), Faias and Santa-Clara (2017), and Gargano, Pettenuzzo, and Timmermann (2019). The existing literature rarely discusses the numerical pitfalls of this optimization problem or the need to enforce non-negative wealth explicitly. Especially for volatile assets, such as options, even moderate portfolio weights can lead to bankruptcy in some scenarios. When this occurs, unconstrained numerical optimizations can produce plausible-looking but economically nonsensical results.

A complementary motivation for power-utility maximization arises from Kelly optimality. When $\gamma = 1$, the CRRA objective reduces to log utility, and the resulting portfolio is the classic Kelly, or growth-optimal, portfolio that maximizes expected log wealth and, under weak regularity conditions, the almost-sure long-run growth rate of wealth.[2] In practice, many investors think in terms of "Kelly sizing" or fractional Kelly rules when allocating risk across trading strategies or asset classes. Power utility with $\gamma \neq 1$ provides a natural multi-asset generalization of this idea: it preserves the multiplicative, long-horizon structure of Kelly optimality while allowing the investor to increase risk aversion γ and thereby scale back from the extreme leverage that often characterizes the fully growth-optimal portfolio. From this perspective, fast and reliable numerical solutions to the power-utility problem are a prerequisite for implementing multi-asset Kelly-style position sizing in empirical applications.

---

[1] Campbell (2018) calls power utility a "workhorse model in the asset-pricing and macroeconomics literatures."

[2] See Kelly (1956) and Breiman (1961) for Kelly optimality and Long (1990) for the numeraire-portfolio characterization of the log-optimal (growth-optimal) portfolio.

Power-utility optimization over historical return samples, sometimes called full-scale optimization (see Cremers, Kritzman, and Page (2005); Hagströmer, Anderson, Binner, Elger, and Nilsson (2008); Hagströmer and Binner (2009)), has been explored before. We revisit this problem from a numerical perspective and show that, without safeguards, it can behave unpredictably even in simple settings.

The remainder of the paper proceeds as follows. Section 2 illustrates why discontinuities in the utility function and its gradient create serious difficulties for standard optimization routines. Section 3 presents the safeguarded utility function and demonstrates its efficiency. Section 4 compares the proposed approach to alternative solutions. Section 5 concludes.

2 The Optimization Problem and Its Pathologies

A common formulation of portfolio choice maximizes expected power utility, or constant relative risk aversion (CRRA), over a sample of observed asset returns. Let $\boldsymbol{\omega}$ denote portfolio weights on $N+1$ assets whose gross returns between t and $t+1$ are $\boldsymbol{R}_{t+1}$, with $R_{i,t+1} = P_{i,t+1}/P_{i,t}$. Terminal wealth is $W_{t+1} = W_t\, \boldsymbol{\omega}' \boldsymbol{R}_{t+1}$, and the investor solves

$$\max_{\boldsymbol{\omega}} E[U(W_{t+1})]\,, \tag{1}$$

for

$$U(W_{t+1}) = \begin{cases} \dfrac{W_{t+1}^{1-\gamma} - 1}{1-\gamma}, & \gamma > 0,\ \gamma \neq 1, \\ \ln(W_{t+1}), & \gamma = 1. \end{cases} \tag{2}$$

The objective is well defined only for non-negative wealth. Yet, unless this restriction is explicitly enforced, standard numerical routines can, and frequently do, converge to solutions that imply negative wealth in some sample periods. Because both the objective and its derivatives appear locally smooth and concave, the optimizer reports normal convergence with plausible utility values. Only a detailed check of the period-by-period wealth path reveals that some observations produce negative wealth, violating the implicit no-bankruptcy constraint.

2.1 Illustration: One risky asset

To illustrate, consider a risk-free asset with gross return $R_{f,t+1}$ and a single risky asset with return R_{t+1}. With portfolio weights summing to one,

$$W_{t+1} = R_{f,t+1} + \omega(R_{t+1} - R_{f,t+1}), \tag{3}$$

and the optimization problem becomes

$$\max_{\omega} E\left[\frac{\left(R_{f,t+1}+\omega(R_{t+1}-R_{f,t+1})\right)^{1-\gamma}-1}{1-\gamma}\right]. \tag{4}$$

The first-order condition,

$$E\left[(R_{t+1}-R_{f,t+1})\left(R_{f,t+1}+\omega(R_{t+1}-R_{f,t+1})\right)^{-\gamma}\right]=0, \tag{5}$$

has no closed-form solution but can be solved numerically.

Figure 1 plots the objective and its gradient for $\gamma = 3$ using CRSP value-weighted index returns and a 3.9% risk-free rate. The figure reveals dozens of false local maxima. Each corresponds to a portfolio that yields negative wealth in at least one period. The gradient is zero and the Hessian negative at every such point, so standard convergence tests confirm “success.” To a researcher inspecting only objective values or solver messages, these outcomes would appear perfectly acceptable; the pathology becomes visible only when wealth is computed explicitly.

2.2 Discontinuities in the utility and gradient

The pathology originates in the structure of (5). Whenever the gross portfolio return equals zero in some observation t,

$$\omega = -\frac{R_{f,t+1}}{R_{t+1}-R_{f,t+1}}, \tag{6}$$

the gradient is undefined and switches from $+\infty$ to $-\infty$. With T distinct returns, there are T such asymptotes. Between any two, the utility is locally concave. The implications depend on the value of γ:

Odd integer γ.

For $\gamma \in \{1, 3, 5, \ldots\}$, the gradient changes sign across each asymptote:

$$g(\omega) \approx \frac{1}{T}\frac{(R_{t+1}-R_{f,t+1})^{1-\gamma}}{\theta^{\gamma}} > 0 \quad \text{just right of a discontinuity,} \tag{7}$$

$$g(\omega) \approx -\frac{1}{T}\frac{(R_{t+2}-R_{f,t+2})^{1-\gamma}}{\theta^{\gamma}} < 0 \quad \text{just left of the next.} \tag{8}$$

Hence, $g(\omega) = 0$ for some ω in each interval, producing the multiplicity of false optima observed in figure 1. All of these stationary points satisfy the first- and second-order conditions yet correspond to portfolios that bankrupt the investor in at least one period.

Because integer values of γ are easy to interpret and implement, odd integers (especially $\gamma = 1$) likely comprise a large share of applications in which the coefficient of relative risk aversion is fixed rather than estimated.

The logarithmic case, $\gamma = 1$, plays two important roles. In international finance, the log-optimal portfolio is the common portfolio held by all investors regardless of residence. (See Adler and Dumas (1983), for example.) In the context of stochastic discount factors, Long (1990) shows that the inverse gross return on the log-optimal portfolio is a legitimate stochastic discount factor in the sense of Ross (1978).

Even integer γ.

For $\gamma \in \{2, 4, 6, \ldots\}$, the utility function remains discontinuous, but the gradient retains the same sign across asymptotes. The optimizer does not settle at false local maxima; instead, it may oscillate or diverge as it steps across the discontinuities. Convergence is slow and unstable because the derivative approaches $\pm\infty$ near the asymptotes, and Newton-style updates may overshoot repeatedly before halting at an economically meaningless point.

Figure 2 illustrates the case $\gamma = 2$. The utility and its gradient remain discontinuous, but the gradient retains the same sign across asymptotes. Instead of multiple false optima, the optimizer oscillates or diverges as updates overshoot the discontinuities. Although these failures differ in appearance from those under odd γ, they reflect the same structural problem: the objective is undefined whenever portfolio wealth is non-positive.

Non-integer γ.

For non-integer γ, negative wealth makes $U(\cdot)$ complex-valued. Most optimizers fail outright once any observation produces $W_{t+1} \leq 0$. While such terminations prevent silent convergence at false optima, they greatly complicate finding the true optimum.

2.3 Scaling to higher dimensions

The same mechanism generates a combinatorial explosion of false optima in multi-asset portfolios. Figure 3 illustrates a three-asset example with a 3.9% risk-free rate and two CRSP size-decile portfolios (the 5th and the 10th, 1926-2003). Each line represents an asymptote from (6)Ña combination (ω_1, ω_2) for which portfolio wealth is zero in one sample period. The closed polygons delineate regions that contain a local maximum when γ is odd. But only the local maximum marked by x is economically sensible: all others

Figure 2: Power Utility with Risk Aversion of 2

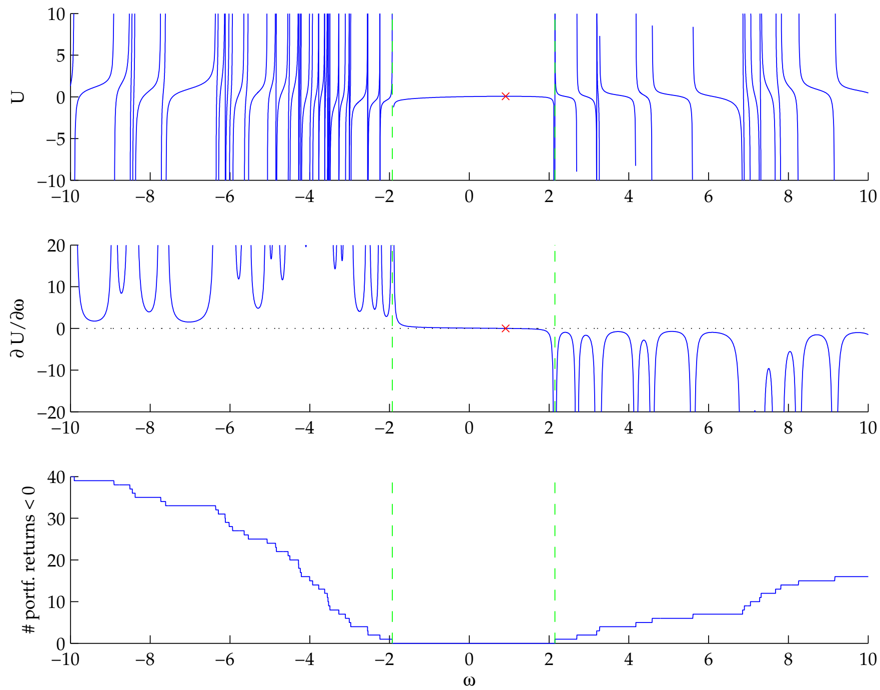

The figure shows the power utility from investing a fraction ω in the CRSP value-weighted index and a fraction $(1-\omega)$ in a riskless asset with annual return of 3.9 percent.
The top panel shows the expected utility, $E[U]$; the middle panel shows the derivative of utility with respect to ω, $\partial U/\partial\omega$; the bottom panel shows the number of negative gross returns for the portfolio.
The true optimum is marked with an X.
The CRSP value-weighted returns are annual returns including dividends from 1926 to 2003.

imply bankruptcy. As the number of assets or periods grows, the number of spurious regions rises combinatorially.[3]

These discontinuities imply that unconstrained numerical maximization of power utility is unreliable for *any* γ: odd integer values yield numerous false optima, even integer values converge slowly or fail, and non-integer values produce undefined utilities once any return is negative. Because the optimizer's diagnostics appear normal in all cases with convergence, the errors are easy to overlook and difficult to diagnose. The next section introduces a direct remedy: a smoothly safeguarded utility function that removes the discontinuities by construction while preserving the true CRRA objective over the economically relevant domain.

[3] If we have k risky assets for T periods, the asymptotes imply an arrangement of T affine hyperplanes in $\mathbb{R}^k$. Barring degenerate cases such as coincident planes or multiple intersections at the same point, Zaslavsky (1975) and Mahajan (2010) show that this partitions the space into $N = \sum_{i=0}^{k} \binom{T}{i}$ regions. Among these, there are $N_c = \binom{T-1}{k}$ fully enclosed regions that each can contain a local maximum. Imposing the weights-sum constraint reduces the effective dimensionality by one.

Figure 3: Power Utility with Risk Aversion of 3 – Level Curves

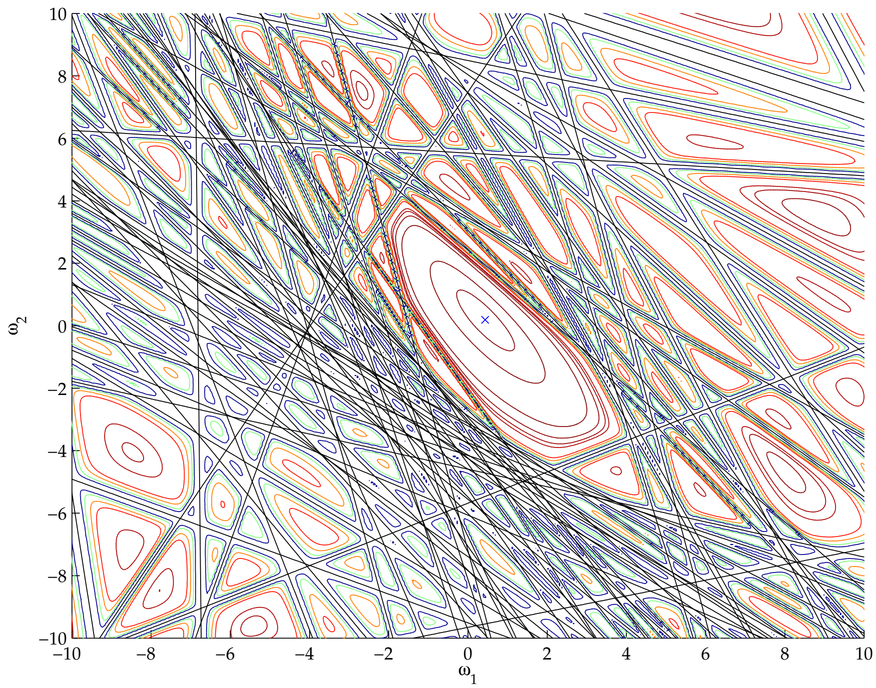

The figure shows level curves for power utility with $\gamma = 3$ for an example with 2 risky assets. The investor is given the opportunity to invest in a riskless asset with annual return of 3.9 percent and two of the CRSP decile portfolios, the 5th and 10th. (The 10th decile contains the largest firms.) Portfolio weight ω_1 is for the 5th decile, ω_2 is for the 10th decile. The weight on the riskless asset is implied as $1 - \omega_1 - \omega_2$. The CRSP portfolio returns are annual returns including dividends from 1926 to 2003.
The true optimum for an investor with risk aversion of 3 is marked with an x. The *ex post* optimal portfolio with constant weights for this investor places 38% into the riskless asset, 34% into the 5th decile portfolio and 19% into the 10th decile portfolio. The remainder is invested in the risk-free asset.
The remaining cells contain local optima that satisfy the mathematical optimality conditions. However, the associated portfolios violate the no-bankruptcy constraint in at least one period.

3 The Safeguarded Utility Function

The difficulties described above arise because the CRRA utility function is undefined for non-positive wealth and its gradient is discontinuous at zero. Standard numerical optimizers do not enforce non-negativity of portfolio returns, so the objective surface becomes fragmented and non-concave. We now introduce a simple modification that restores global concavity and ensures a unique, economically sensible optimum. The safeguard eliminates the possibility of hidden numerical errors: once applied, any optimizer will converge to the true solution or clearly fail to improve the objective.

Figure 4: Safeguarded Power Utility

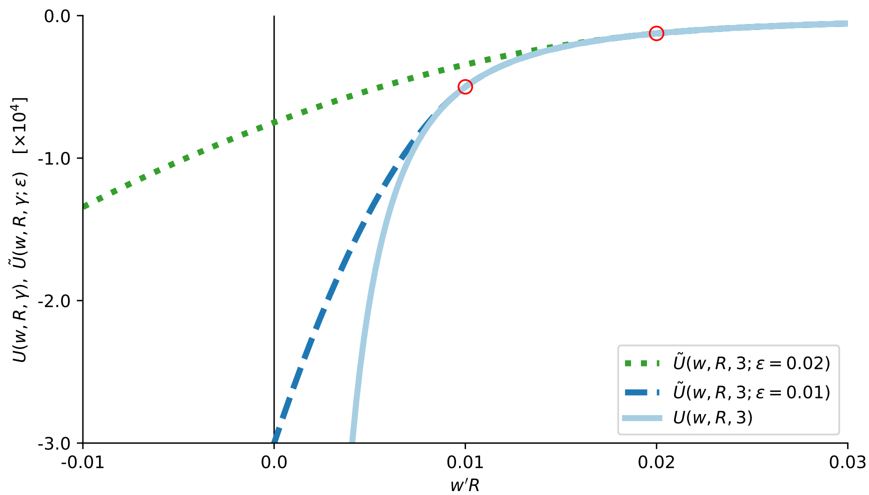

The figure plots the quadratic extensions to the power utility function for two values of the contact point ϵ: the dotted line is $\widetilde{U}(\boldsymbol{\omega}, \boldsymbol{R}, 3; \epsilon = 0.02)$ and the dashed line is $\widetilde{U}(\boldsymbol{\omega}, \boldsymbol{R}, 3; \epsilon = 0.01)$. The solid line shows the true power utility function for comparison. All three functions use $\gamma = 3$. The smooth pasting points, where the quadratic extensions join the power utility function, are marked by circles.
For portfolio gross returns above the pasting points the true power utility function and its safeguarded version coincide. The true power utility function is undefined for non-positive gross returns, where the investor goes bankrupt. The safeguarded version is defined for all gross returns to simplify numerical optimization.

3.1 Idea and construction

We redefine the objective function so that it coincides with true power utility for all portfolio returns above a small positive threshold ϵ, and smoothly attaches a quadratic extension below ϵ. The extension matches the value, slope, and curvature of the original function at the pasting point, making the overall utility twice differentiable and globally concave. For a portfolio with gross return $G = \boldsymbol{\omega}'\boldsymbol{R}$, define

$$\widetilde{U}(G; \gamma, \epsilon) = \begin{cases} \dfrac{G^{1-\gamma} - 1}{1-\gamma}, & G \geq \epsilon, \\ \dfrac{\epsilon^{1-\gamma} - 1}{1-\gamma} + \left(z - \frac{\gamma}{2} z^2\right) \epsilon^{1-\gamma}, & G < \epsilon, \end{cases} \qquad z \equiv \frac{G}{\epsilon} - 1, \tag{9}$$

where $\epsilon > 0$ is a small safeguard level (e.g., $\epsilon = 0.0001$ corresponds to a 99.99% loss of wealth). The two pieces meet with identical value, first derivative, and second derivative at $G = \epsilon$, guaranteeing smooth pasting. In the limit as ϵ goes to 0, $\widetilde{U}(G; \gamma, \epsilon)$ converges to $U(G; \gamma)$ for all $G > 0$.

Figure 4 illustrates the construction. The safeguard preserves the true utility over the feasible domain and introduces a continuous penalty below it. This change has two crucial effects: (1) the function becomes globally concave with a monotonic gradient in $\boldsymbol{\omega}$; and (2) it is defined for all real portfolio returns, avoiding complex or undefined values.

3.2 Derivatives and concavity

Because the extension is quadratic in G, analytical derivatives are straightforward. The gradient and Hessian of the safeguarded utility with respect to portfolio weights are

$$\widetilde{\boldsymbol{g}}(\boldsymbol{\omega}) = \begin{cases} G^{-\gamma}\boldsymbol{R}, & G \geq \epsilon, \\ \left[1 - \gamma\left(\dfrac{G}{\epsilon} - 1\right)\right]\epsilon^{-\gamma}\boldsymbol{R}, & G < \epsilon, \end{cases} \tag{10}$$

$$\widetilde{\boldsymbol{H}}(\boldsymbol{\omega}) = \begin{cases} -\gamma\, G^{-(1+\gamma)}\, \boldsymbol{R}\boldsymbol{R}', & G \geq \epsilon, \\ -\gamma\, \epsilon^{-(1+\gamma)}\, \boldsymbol{R}\boldsymbol{R}', & G < \epsilon. \end{cases} \tag{11}$$

Both are continuous at $G = \epsilon$, and $\widetilde{\boldsymbol{H}}$ is negative semi-definite everywhere, establishing global concavity.[4]

Because $\widetilde{U}$ equals U over the economically relevant region, the optimum of the safeguarded problem coincides exactly with the true constrained optimum when all $G_t = \boldsymbol{\omega}'\boldsymbol{R}_t$ are positive. If any $G_t < \epsilon$, the optimizer receives a smooth but steep penalty that pushes the solution back toward the feasible region. Hence, the algorithm can start from any initial portfolio and converge reliably without explicit constraints.

3.3 Economic interpretation

The safeguard acts as a continuous implementation of the limited-liability condition. Rather than imposing a hard inequality constraint, it embeds the constraint directly in the objective by attaching a rapidly decreasing penalty for negative wealth. This modification is economically consistentÑan investor who faces ruin loses almost all utilityÑand numerically stable because the objective and gradient remain continuous. Unlike the standard CRRA function, which can hide violations of limited liability behind apparently well-behaved gradients, the safeguarded version makes such violations impossible by construction. Unlike generic barrier or penalty methods (e.g., Forsgren et al., 2002; Boyd and Vanderberghe, 2004), no additional tuning parameter is required: the extension preserves exact concavity and yields the same optimum once the portfolio remains solvent.

Figure 5 demonstrates the effect in the one-asset case from figure 1. The safeguarded function is now smooth and globally concave, with a single

[4] Formally, $\partial^2\widetilde{U}/\partial G^2 = -\gamma\, G^{-(1+\gamma)}$ for $G \geq \epsilon$ and $-\gamma\, \epsilon^{-(1+\gamma)}$ for $G < \epsilon$, both ≤ 0.

Figure 5: Safeguarded Power Utility with Risk Aversion of 3

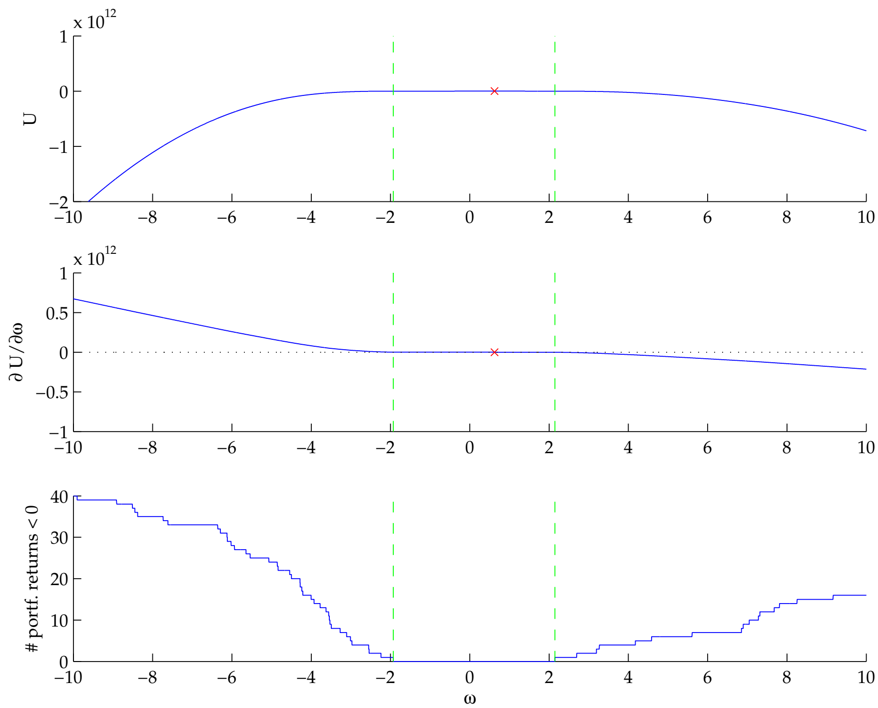

The figure shows the *safeguarded* power utility with $\gamma = 3$ for the same data as figure 1.
The utility is a function of investing a fraction ω in the CRSP value-weighted index and a fraction $(1 - \omega)$ in a riskless asset with annual return of 3.9 percent. We paste a quadratic extension to the power utility function at points where the gross portfolio return is less than 0.01.
The top panel shows the safeguarded utility, $\widetilde{U}$; the middle panel shows the derivative of the augmented utility with respect to ω, $\partial U/\partial\omega$; the bottom panel shows the number of negative gross returns for the portfolio. The true optimum is marked with an x.
The CRSP value-weighted index returns are annual returns including dividends from 1926 to 2003.

stationary point corresponding to the economically valid optimum. The gradient no longer exhibits discontinuities, and numerical maximization converges in a few iterations from any starting weight. This is true for any γ.

Figure 6 shows the corresponding bivariate case from figure 3. The safeguarded function eliminates the multitude of false optima: only one smooth peak remains. Together, these figures demonstrate that the safeguard provides a direct, exact, and computationally efficient solution to the numerical instability of CRRA optimization.

4 Alternative Remedies and Their Limitations

Several standard techniques have been proposed to stabilize numerical optimization when the objective function has defects similar to power utility. These approaches can mitigate failures in specific settings but do not provide

Figure 6: Safeguarded Power Utility with Risk Aversion of 3 – Level Curves

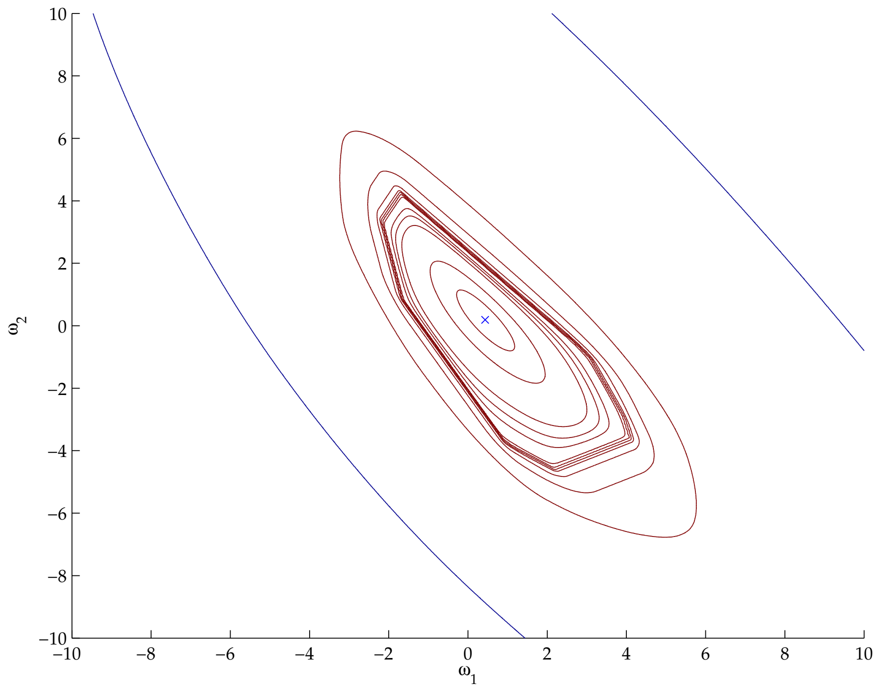

The figures shows level curves of the *safeguarded* power utility function with $\gamma = 3$ for 2 risky assets using the same data as figure 3.
We splice a quadratic extension to the power utility function at points where $\boldsymbol{\omega}'\boldsymbol{R} = 0.001$. The power utility function has risk aversion $\gamma = 3$. The figures shows level curves at $\{0.05, 0, -0.25, -0.5, -1, -10, -100, -10^6, -10^7, -10^8, -10^9, -10^{10}, -10^{12}\}$.
The investor has the opportunity to invest in a riskless asset with annual return of 3.9 percent and two of the CRSP decile portfolios, the 5th and 10th. (The 10th decile contains the larges firms.) portfolio weight ω_1 is for the 5th decile, ω_2 is for the 10th decile. The weight on the riskless asset is implied as $1 - \omega_1 - \omega_2$. The CRSP portfolio returns are annual returns including dividends from 1926 to 2003.
The true optimum for an investor with risk aversion of 3 is marked by the x. The optimal portfolio for this investor places 38% into the riskless asset, 34% into the 5th decile portfolio and 19% into the 10th decile portfolio. The remainder is invested in the risk-free asset.

a general or efficient solution. We briefly summarize three common strategies and explain why each is likely to fall short for the CRRA problem.

Feasible starting points.

Without the safeguarded utility function, we must begin the search from an economically feasible portfolio. Starting from there, we can hope that the search remains in the economically feasible region. The all-cash portfolio or the minimum-variance portfolio are natural choices. Young (1998) advocates the "minimax" portfolio that just avoids bankruptcy. All three ensure positive wealth initially but offer no guarantee that intermediate steps remain feasible. Because the gradient of the CRRA objective is discontinuous, an optimization started from a feasible point can easily leave the region of positive wealth after a few iterations and converge to a meaningless solution.

Explicit inequality constraints.

Another remedy is to impose non-negativity of portfolio wealth directly as a set of linear inequalities, one for each observation:

$$-\boldsymbol{R}_t'\boldsymbol{\omega} \leq 0, \qquad t = 1, \ldots, T. \tag{12}$$

Modern nonlinear optimizers, such as those described in Boyd and Vandenberghe (2004), can handle thousands of such constraints, but the computational cost rises quickly with the sample length T. Constraint-reduction methods by Boot (1962), Telgen (1983), Caron, McDonald, and Ponic (1989), or Paulraj and Sumathi (2010) can remove redundancies, but they require solving many auxiliary linear programs and are unlikely to yield net speed gains. In large samples, the constrained formulation becomes impractically slow.

Penalty and barrier functions.

Penalty and barrier methods are standard tools in nonlinear optimization (Forsgren, Gill, and Wright, 2002; Luenberger, 1984; Boyd and Vanderberghe, 2004). They modify the objective by adding a term that diverges when a constraint is violated. While these methods can maintain feasibility, they require problem-specific tuning parameters that balance accuracy against stability. In the CRRA setting, inappropriate tuning distorts curvature near the constraint boundary and can shift the optimum even when all portfolio returns are positive.

Summary.

Each of these techniques addresses part of the difficulty but none eliminates it. Feasible starts do not prevent the optimizer from stepping into bankruptcy; explicit constraints guarantee correctness only at high computational cost; and generic penalty functions require tuning that changes the problem's economic meaning. The next section demonstrates that a simple modification of the utility function itself, a smoothly safeguarded objective, efficiently resolves these issues by while preserving the true CRRA preferences over the feasible domain.

5 Discussion and Conclusion

Power-utility maximization is a foundational tool in empirical finance and asset-pricing research, yet its numerical implementation can fail silently when the non-negativity of wealth is not enforced. The discontinuities in the CRRA objective generate false local optima, undefined gradients, and spurious convergence, features that are not obvious from standard first- and second-order checks. These failures are not purely technical: they produce economically impossible portfolios that violate limited liability and invalidate welfare comparisons based on the resulting utility values. Because these failures often leave no visible trace in optimization outputs, empirical implementations may report apparently valid but economically meaningless optima.

We propose a simple safeguard that restores the theoretical and numerical integrity of the problem. By smoothly pasting a quadratic extension to the utility function near zero wealth, the safeguarded objective becomes globally concave and continuously differentiable while remaining identical to the true CRRA utility over the economically relevant region. This modification requires no constraints, no search for feasible starting points, and no tuning parameters. It enforces the limited-liability principle implicitly, through the shape of the objective itself.

Beyond computational convenience, the safeguard has conceptual value. It embeds the no-bankruptcy condition, an essential economic assumption, directly in the utility function. The resulting objective preserves global concavity and yields a unique optimum, making CRRA optimization reliable even in large-scale, high-volatility, or near-arbitrage settings. Because the adjustment leaves the utility unchanged for feasible wealth levels, all empirical interpretations of power utility remain valid.

In practical terms, the safeguarded utility allows researchers to implement full-scale CRRA optimization for hundreds or thousands of assets without numerical fragility. It provides a direct, economically consistent alternative to ad-hoc inequality constraints or barrier penalties. Future work could extend the safeguard to dynamic or recursive utility models, or to other bounded-domain preferences used in behavioral and macro-finance applications.

Overall, the safeguarded formulation offers a small but essential correction to a widely used methodology. It ensures that power-utility maximization, a workhorse of portfolio theory, is not only theoretically sound but also numerically well behaved.

6 References

Adler, Michael, and Bernard Dumas, 1983, International portfolio choice and corporation finance: A synthesis, *Journal of Finance* 38, 925–984.

Aït-Sahalia, Yacine, and Michael W. Brandt, 2001, Variable selection for portfolio choice, *Journal of Finance* 56, 1297–1351.

Algoet, Paul H., and Thomas M. Cover, 1988, Asymptotic optimality and asymptotic equipartitioning properties of log-optimum investment, *Annals of Probability* 16, 876–898.

Andersen, Martin S., Joachim Dahl, and Lieven Vanderberghe, 2020, CvxOpt: A Python package for convex optimization, Working paper, https://cvxopt.org.

Barberis, Nicholas, 2000, Investing for the long run when returns are predictable, *Journal of Finance* 55, 225–264.

Boot, Johan C. G., 1962, On trivial and binding constraints in programming problems, *Management Science* 8, 419–441.

Boyd, Stephen, and Lieven Vanderberghe, 2004, *Convex Optimization* (Cambridge University Press, Cambridge, England).

Brandt, Michael W., 1999, Estimating portfolio and consumption choice: A conditional Euler equations approach, *Journal of Finance* 54, 1609–1645.

Breiman, Leo, 1961, Optimal gambling systems for favorable games, *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability* 1, 65–78.

Campbell, John Y., 2018, *Financial Decisions and Markets: A Course in Asset Pricing* (Princeton University Press, Princeton, NJ).

Caron, Richard J., James F. McDonald, and C.M. Ponic, 1989, A degenerate extreme point strategy for the classification of linear constraints as redundant or necessary, *Journal of Optimization Theory and Applications* 62, 225–237.

Cover, Thomas M., 1991, Universal portfolios, *Mathematical Finance* 1, 1–29.

Cremers, Jan-Hein, Mark Kritzman, and Sébastien Page, 2005, Optimal hedge fund allocations, *Journal of Portfolio Management* 31, 70–81.

Driessen, Joost, and Pascal Maenhout, 2007, An empirical portfolio perspective on option pricing anomalies, *Review of Finance* 11, 561–603.

Faias, José A., and Pedro Santa-Clara, 2017, Optimal option portfolio strategies: Deepening the puzzle of index option mispricing, *Journal of Financial and Quantitative Analysis* 52, 277–303.

Forsgren, Anders, Philip E. Gill, and Margaret H. Wright, 2002, Interior methods for nonlinear optimization, *SIAM Review* 44, 525–597.

Gargano, Antonio, Davide Pettenuzzo, and Allan Timmermann, 2019, Bond return predictability: Economic value and links to the macroeconomy, *Management Science* 65, 508–540.

Grauer, Robert R., and Nils H. Hakansson, 1987, Gains from international diversification: 1968–85 returns on portfolios of stocks and bonds, *Journal of Finance* 42, 721–739.

Hagströmer, Björn, Richard G. Anderson, Jane M. Binner, Thomas Elger, and Birger Nilsson, 2008, Mean-variance versus full-scale optimization: Broad evidence for the UK, *The Manchester School* 76, 134–156.

Hagströmer, Björn, and Jane M. Binner, 2009, Stock portfolio selection with full-scale optimization and differential evolution, *Applied Financial Economics* 19, 1559–1571.

Hentschel, Ludger, and John B. Long, Jr., 2004, The size and source of gains from international diversification, Working paper, Simon School, University of Rochester, Rochester, Ny.

Kelly, John L., Jr., 1956, A new interpretation of information rate, *Bell System Technical Journal* 35, 917–926.

Long, John B., Jr., 1990, The numeraire portfolio, *Journal of Financial Economics* 26, 29–69.

Luenberger, David G., 1984, *Linear and Nonlinear Programming*, second edition (Addison-Wesley Publishing Company, Reading, Ma).

Mahajan, Sanjoy, 2010, *Street-Fighting Mathematics* (MIT Press, Cambridge, MA).

Malamud, Semyon, 2014, Portfolio selection with options, Working paper, Swiss Finance Institute, Lausanne, Switzerland.

Paulraj, Sujava, and Poobal Sumathi, 2010, A comparative study of redundant constraints identification methods in linear programming problems, *Mathematical Problems in Engineering* 1–16.

Ross, Stephen A., 1978, A simple approach to the valuation of risky streams, *Journal of Business* 3, 453–476.

Telgen, Jan, 1983, Identifying redundant constraints and implicit equalities in systems of linear constraints, *Management Science* 29, 1209–1222.

Young, Martin R., 1998, A minimax portfolio selection rule, *Management Science* 44, 673–683.

Zaslavsky, Thomas, 1975, *Facing up to Arrangements: Face-Count Formulas for Partitions of Space by Hyperplanes*, volume 1 of *Memoirs of the American Mathematical Society* (American Mathematical Society, Providence, RI).