Traditional Optimization Is Not Optimal for Leverage-Averse Investors

Bruce I. Jacobs and Kenneth N. Levy

Bruce I. Jacobs
is a principal of Jacobs Levy Equity Management in Florham Park, NJ.
bruce.jacobs@jlem.com

Kenneth N. Levy
is a principal of Jacobs Levy Equity Management in Florham Park, NJ.
ken.levy@jlem.com

It has long been recognized that leveraging a portfolio increases risk. In order to mitigate the risk of leverage, investors using conventional mean–variance portfolio optimization often include a leverage constraint.[1]

In Jacobs and Levy [2012, 2013a], we discussed the unique risks of leverage and developed mean–variance–leverage portfolio optimization, which takes these unique risks into consideration.[2] The mean–variance–leverage optimization model incorporates a leverage-aversion term in the utility function, which allows investors to explicitly consider the economic trade-offs between expected return, volatility risk, and leverage risk. Investors can then determine the optimal amount of leverage, according to their particular level of leverage aversion.[3]

In this article, we contrast mean–variance–leverage portfolio optimization with the conventional approach: using a leverage constraint in mean–variance portfolio optimization. We consider the mean–variance investor who is averse to volatility risk. We develop mean–variance efficient frontiers using the conventional mean–variance utility function and optimizing with a series of leverage constraints. Looser constraints—that is, constraints at higher levels of leverage—provide greater mean–variance utility, until a utility peak is reached. The portfolio at this utility peak can be identified with mean–variance optimization without a leverage constraint. This portfolio has a very high level of leverage.

We then consider the mean–variance–leverage investor who is averse to both volatility risk and leverage risk. We can arrive at the portfolio that offers the highest utility for such an investor by either of two methods.

The first method determines the mean–variance–leverage utility that a leverage-averse investor would obtain from conventional leverage-constrained optimal mean–variance portfolios. By using a mean–variance–leverage utility function, we show how the leverage-averse investor could identify the leverage-constrained mean–variance portfolio that has the optimal level of leverage and offers the highest mean–variance–leverage utility. Note that, without knowing the investor's mean–variance–leverage utility function, we cannot determine the leverage-averse investor's optimal portfolio.

The second method demonstrates how a leverage-averse investor can use mean–variance–leverage optimization to directly determine the portfolio that has the optimal leverage level and offers the highest mean–variance–leverage utility. We show that both methods produce the same optimal portfolio.

We also show that, as an investor's leverage tolerance increases without bound, optimal mean–variance–leverage portfolios approach those determined by a conventional mean-variance utility function.

For an investor who is averse to leverage, conventional mean-variance optimization offers little guidance regarding the optimal leverage level, and is thus unable to identify the portfolio that offers the highest utility. More importantly, without knowing the leverage-averse investor's mean-variance-leverage utility function, using the conventional mean-variance utility function and optimizing with a leverage constraint is unlikely to lead to the portfolio that offers the highest utility.

MEAN-VARIANCE OPTIMIZATION WITH A LEVERAGE CONSTRAINT

Conventional mean-variance portfolio optimization identifies the portfolio that maximizes the following utility function:

$$U = \alpha_P - \frac{1}{2\tau_V}\sigma_P^2 \tag{1}$$

where α_p is the portfolio's expected active return (relative to benchmark), σ_p^2 is the variance of the portfolio's active return, and τ_V is the investor's risk tolerance with respect to the volatility of the portfolio's active return, which we will refer to as volatility tolerance. We use the terms *tolerance* and *aversion* with the understanding that each is the inverse of the other. We refer to the utility that derives from Equation (1) as MV(τ_V) utility, refer to investors who optimize using this utility function as MV(τ_V) investors, and refer to the portfolios that result from such optimization as MV(τ_V) portfolios.

We use active security returns and active security weights to calculate the portfolio's active return and variance. The active weight, x_i, of security i is equal to its holding weight, h_i, minus its benchmark weight, b_i:

$$x_i = h_i - b_i \tag{2}$$

The portfolio's expected active return is

$$\alpha_P = \sum_{i=1}^{N} \alpha_i x_i \tag{3}$$

where α_i is the expected active return of security i and N is the number of securities in the selection universe.

The variance of the portfolio's active return is

$$\sigma_P^2 = \sum_{i=1}^{N}\sum_{j=1}^{N} x_i \sigma_{ij} x_j \tag{4}$$

where σ_{ij} is the covariance between the active returns of securities i and j.

Using Equations (3) and (4), the utility function in Equation (1) is equivalent to the following:

$$U = \sum_{i=1}^{N} \alpha_i x_i - \frac{1}{2\tau_V}\sum_{i=1}^{N}\sum_{j=1}^{N} x_i \sigma_{ij} x_j \tag{5}$$

We define portfolio leverage as the sum of the absolute values of the portfolio holding weights minus 1:[4]

$$\Lambda = \sum_{i=1}^{N} |h_i| - 1 \tag{6}$$

For illustration, consider an enhanced active equity (EAE) portfolio structure, where E is the portfolio's enhancement and $E = \Lambda/2$. For example, a 130-30 EAE portfolio holds 130% of capital long and 30% short. The leverage, Λ, is 0.6, or 60%, and the enhancement, E, is 0.3, or 30%.

The standard constraint set for an EAE portfolio is

$$\sum_{i=1}^{N} h_i = 1 \tag{7}$$

and

$$\sum_{i=1}^{N} h_i \beta_i = 1 \tag{8}$$

Equation (7) is the full-investment (net longs minus shorts) constraint, which requires that the sum of the signed holding weights equals 1. Equation (8) is the beta constraint (where β_i is the beta of security i relative to the benchmark), which requires that the portfolio's beta equals 1. In terms of active weights, these constraints are expressed as

$$\sum_{i=1}^{N} x_i = 0 \tag{9}$$

and

$$\sum_{i=1}^{N} x_i \beta_i = 0 \tag{10}$$

Using data for stocks in the S&P 100 index and constraining each security's active weight to be within 10 percentage points of its weight in the S&P 100 index benchmark, we plot, in Exhibit 1, six leverage-constrained efficient frontiers for leverage values ranging from 0% to 100% at 20% intervals.[5] These leverage levels correspond to enhancements ranging from 0% (an unleveraged, long-only portfolio) to 50% (a 150-50 EAE portfolio).

We implement the leverage constraints by setting Λ in Equation (6) equal to the constrained value, and including this as an additional constraint in the traditional mean-variance optimization. For example, the leverage constraint we use to achieve a 130-30 efficient frontier is $\Lambda = 0.6$. For expository purposes, we assume the strategy entails no financing costs.[6] To trace out each of these efficient frontiers, we employ a range of volatility tolerance (τ_V) values from near 0 to 2.[7]

Exhibit 1 illustrates the six efficient frontiers. For each frontier, as the investor's tolerance for volatility increases, the optimal portfolio moves out along the frontier, taking on higher levels of standard deviation of active return in order to earn higher levels of expected active return. The frontiers constrained to higher levels of leverage (and enhancement) provide higher expected active returns at each level of standard deviation of active return. It appears from this exhibit that the frontiers with greater leverage dominate those with less leverage. That is, a mean-variance investor would prefer the 150-50 EAE frontier to the 140-40 EAE frontier, and so on, with the 100-0 long-only, frontier being the least desirable frontier.[8]

We now locate the portfolio that is optimal for an investor with a volatility tolerance of 1—that is, the MV(1) portfolio—on each of the six efficient frontiers. These portfolios are shown in Exhibit 1, labeled "a"

EXHIBIT 1
Optimal MV(1) Portfolios for Various Leverage Constraints

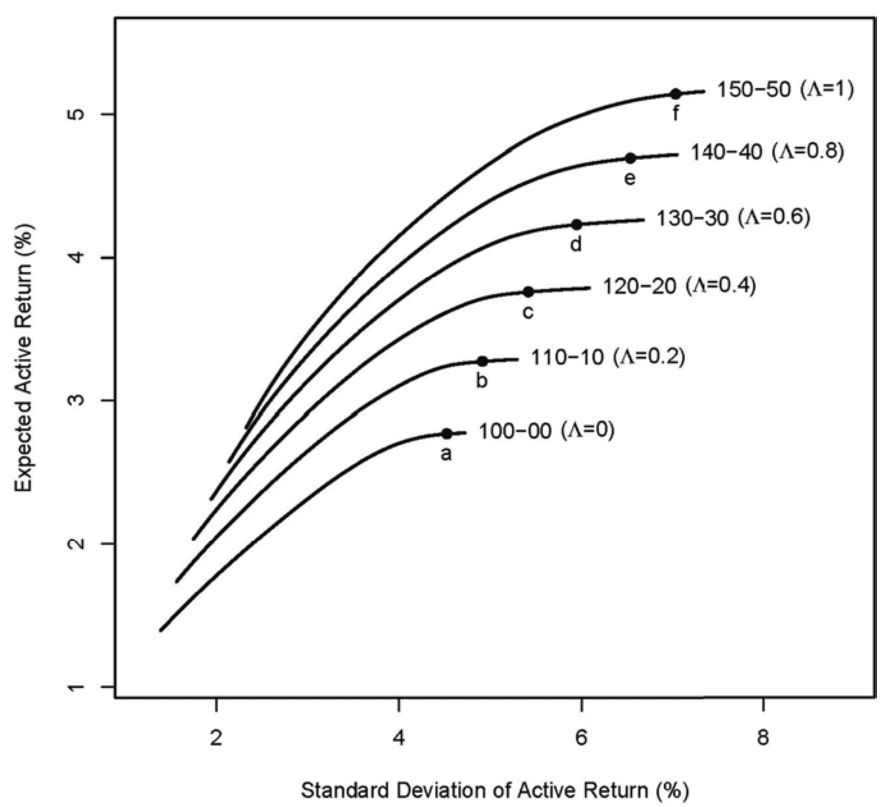

through "f." For instance, "c" on the 120-20 leverage-constrained efficient frontier is the portfolio on that frontier that offers the highest utility for a mean-variance investor with a volatility tolerance of 1.

Exhibit 2 extends the analysis of MV(1) portfolios for investors who are allowed higher levels of leverage. The solid line plots the MV(1) utility of optimal portfolios with security active weight constraints, as we increase the enhancement by steps of 1%, from 0% to beyond 400%. The exhibit shows portfolios "a" through "f." As securities reach the upper bounds of their security active weight constraints, the MV(1) utility peaks at portfolio "z." This portfolio is highly leveraged, with an enhancement of 392%, resulting in a 492-392 EAE portfolio with a leverage of 7.84 times net capital.[9] Portfolio "z" can also be obtained from a mean-variance optimization with security active weight constraints, but no leverage constraint.

The dashed line in Exhibit 2 plots the MV(1) utility of optimal portfolios without security active weight constraints, as we increase the enhancement by the same 1% steps as before. Without a constraint on leverage, MV(1) utility peaks at an extremely leveraged portfolio, one that is literally off the chart. This is a 4,650-4,550 EAE portfolio with an enhancement of 4,550% and a leverage of 91 times net capital.

We have continued to assume no financing cost. Even so, the amount of leverage that the optimal mean-variance portfolio takes on is not unlimited. This is because the portfolio's volatility rises with leverage, and the volatility-aversion term in the mean-variance utility function eventually reduces utility by more than the expected return term increases utility.

Exhibit 3 gives the characteristics of the optimal portfolios identified in Exhibit 2. These portfolios have constraints on security active weights and leverage. Standard deviation of active return, expected active return, and utility all increase monotonically with the amount of leverage. Of portfolios "a" through "f," portfolio "f," the 150-50 portfolio, offers the mean-variance investor

EXHIBIT 2

MV(1) Utility of Optimal MV(1) Portfolios as a Function of Enhancement

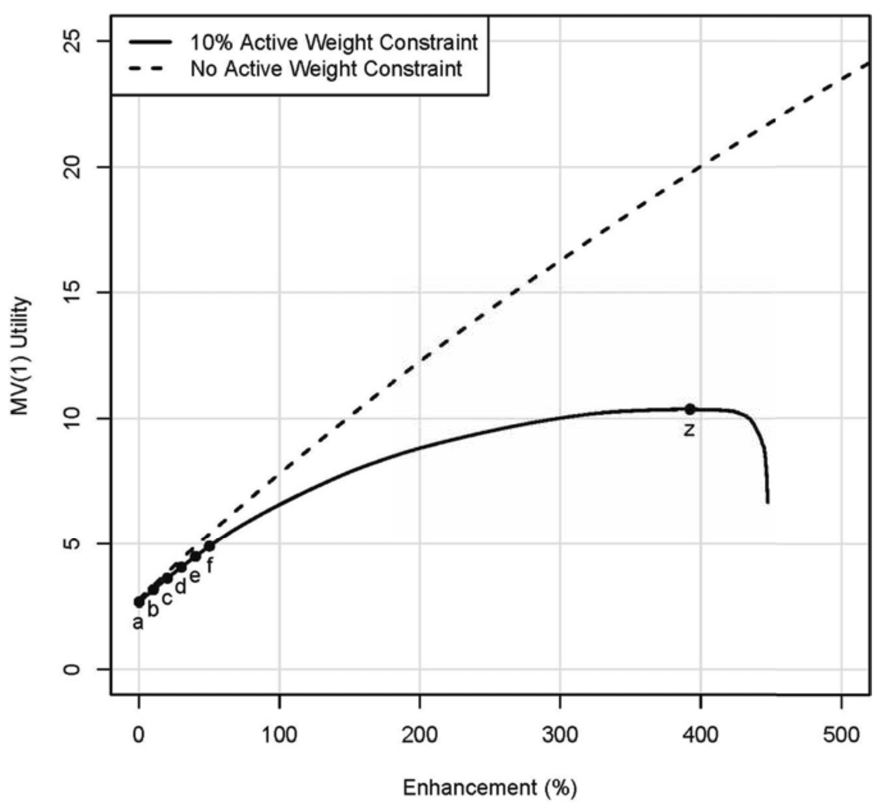

EXHIBIT 3
Characteristics of Optimal MV(1) Portfolios from the Perspective of an MV(1) Investor

Portfolio	EAE	Leverage	Standard Deviation of Active Return	Expected Active Return	Utility for an MV(1) Investor
a	100-0	0	4.52	2.77	2.67
b	110-10	0.2	4.91	3.27	3.15
c	120-20	0.4	5.42	3.76	3.61
d	130-30	0.6	5.94	4.23	4.06
e	140-40	0.8	6.53	4.70	4.48
f	150-50	1.0	7.03	5.14	4.90
z	492-392	7.84	15.43	11.55	10.36

the highest MV(1) utility. But portfolio "z," the 492-392 portfolio, offers the highest utility of all the MV(1) portfolios.[10]

These findings are consistent with those in Jacobs and Levy [2012], demonstrating that conventional mean-variance analysis implicitly assumes that investors have no aversion to (or, stated differently, have an infinite tolerance for) the unique risks of leverage. This lack of consideration by mean-variance analysis of investor aversion to these unique risks motivates the development of a mean-variance-leverage optimization model.

THE LEVERAGE-AVERSE INVESTOR'S UTILITY OF OPTIMAL MEAN-VARIANCE PORTFOLIOS

In Jacobs and Levy [2013a], we specified an augmented utility function that includes a leverage-aversion term:

$$U = \alpha_P - \frac{1}{2\tau_V}\sigma_P^2 - \frac{1}{2\tau_L}\sigma_T^2\Lambda^2 \tag{11}$$

where σ_T^2 is the variance of the leveraged portfolio's total return and τ_L is the investor's leverage tolerance. The leverage-aversion term assumes that the risks of leverage rise with the product of the variance of the leveraged portfolio's total return and the square of the portfolio's leverage.[11] We refer to the utility that derives from Equation (11) as MVL(τ_V,τ_L) utility, refer to investors who optimize their portfolios using this utility function as MVL(τ_V,τ_L) investors, and refer to the portfolios that result from such optimization as MVL(τ_V,τ_L) portfolios.

Defining q_{ij} as the covariance between the total returns of securities i and j, Equation (11) can be written as:

$$U = \sum_{i=1}^{N}\alpha_i x_i - \frac{1}{2\tau_V}\sum_{i=1}^{N}\sum_{j=1}^{N}x_i\sigma_{ij}x_j - \frac{1}{2\tau_L}\left(\sum_{i=1}^{N}\sum_{j=1}^{N}h_i q_{ij} h_j\right)\Lambda^2 \tag{12}$$

Using Equations (2) and (6), Equation (12) becomes:

$$U = \sum_{i=1}^{N}\alpha_i x_i - \frac{1}{2\tau_V}\sum_{i=1}^{N}\sum_{j=1}^{N}x_i\sigma_{ij}x_j - \frac{1}{2\tau_L}\left(\sum_{i=1}^{N}\sum_{j=1}^{N}\left(b_i + x_i\right)q_{ij}\left(b_j + x_j\right)\right)\left(\sum_{i=1}^{N}\left|b_i + x_i\right| - 1\right)^2 \tag{13}$$

We can use the mean-variance-leverage utility function specified in Equation (13) to calculate the utility that a leverage-averse investor would obtain from the MV(1) portfolios "a" through "f," shown in Exhibit 3. For illustration, we assume the leverage-averse investor has a volatility tolerance of 1, the same tolerance as the mean-variance investor, and a leverage tolerance of 1: that is, the investor is an MVL(1,1) investor. The portfolios' utilities are plotted as a function of their enhancement and labeled as "a" through "f" in Exhibit 4.

In order to trace the continuous curve shown in this exhibit, we determined more than 1,000 optimal leverage-constrained MV(1) portfolios by increasing the constrained amount of the leverage from 0% to more than 100%, in increments of 0.1% (corresponding to enhancements from 0% to more than 50% in increments of 0.05%). We then calculated the utility that each portfolio would offer an MVL(1,1) investor. The exhibit thus plots utilities for an MVL(1,1) investor over a range of leverage-constrained optimal MV(1) portfolios.

The resulting MVL(1,1) utility curve is shaped like an arch. This arch peaks at portfolio "g," which offers the MVL(1,1) investor the highest utility. It is a 129-29 EAE portfolio. This peak in investor utility occurs because, as the portfolio's enhancement increases beyond that of portfolio "g," the leverage and volatility aversion terms reduce utility by more than the expected return term increases utility.

Exhibit 5 displays these portfolios' characteristics. Although the standard deviation of active return and expected active return increase monotonically with leverage (note that they are the same values as in Exhibit 3), investor utility does not. For our leverage-averse investor, the leverage-constraint level corresponding to the 129-29 portfolio (portfolio "g") provides the highest utility. Other leverage constraints

Exhibit 4

MVL(1,1) Utility of Optimal MV(1) Portfolios as a Function of Enhancement

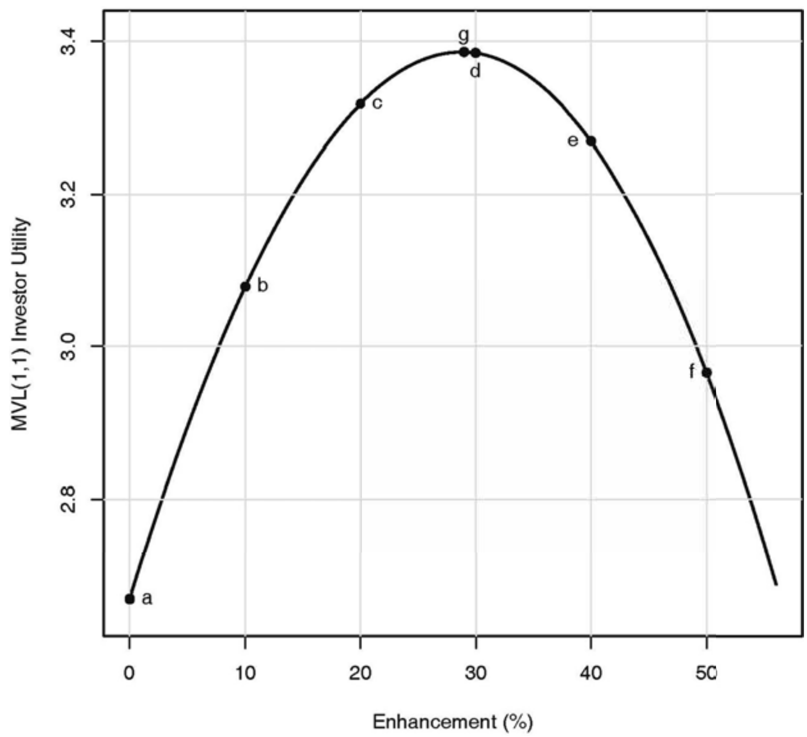

Exhibit 5

Characteristics of Optimal MV(1) Portfolios from the Perspective of an MVL(1,1) Investor

Portfolio	EAE	Leverage	Standard Deviation of Active Return	Expected Active Return	Utility for an MVL(1,1) Investor
a	100-0	0	4.52	2.77	2.67
b	110-10	0.2	4.91	3.27	3.08
c	120-20	0.4	5.42	3.76	3.32
g	129-29	0.58	5.89	4.18	3.39
d	130-30	0.6	5.94	4.23	3.38
e	140-40	0.8	6.53	4.70	3.27
f	150-50	1.0	7.03	5.14	2.97

provide less utility, because they are either too tight (less than 129-29) or too loose (greater than 129-29). Either way, they are not optimal for the MVL(1,1) investor.[12]

In this analysis, by considering numerous optimal MV(1) portfolios—each constrained at a different leverage level—and applying an MVL(1,1) utility function to evaluate each portfolio, we were able to determine which leverage-constrained MV(1) portfolio offers an MVL(1,1) investor the highest utility. Leverage-constrained MV(1) optimization cannot locate this highest-utility portfolio if the leverage-averse investor's utility function is not known.

In the next section, we show that we can determine the optimal portfolio directly, by using the mean-variance-leverage optimization model.

MEAN-VARIANCE-LEVERAGE OPTIMIZATION VERSUS LEVERAGE-CONSTRAINED MEAN-VARIANCE OPTIMIZATION

An MVL(τ_V,τ_L) investor maximizes the utility function represented by Equation (13) to identify the optimal portfolio. We found the portfolios that maximize this utility function for a range of volatility and leverage-tolerance pairs (τ_V,τ_L).[13] As in Jacobs and Levy [2012], we chose 100×100 pairs of values for the volatility and leverage tolerances to cover an illustrative range [0.001, 2], for a total of 10,000 optimizations.[14,15]

The enhancements of the optimal portfolios obtained as a function of τ_V and τ_L are shown as the efficient surface in Exhibit 6.[16] At zero leverage tolerance, the optimal portfolios lie along the volatility-tolerance axis, having no leverage and hence no enhancement. At zero volatility tolerance, the portfolios lie along the leverage-tolerance axis, having no active return volatility and hence holding benchmark weights in each security.

To help identify other features of the efficient surface, Exhibit 7 plots a contour map of the efficient surface from Exhibit 6. Each contour line represents a slice of the efficient surface at a given level of enhancement, and shows the combinations of volatility tolerance and leverage tolerance for which a given level of enhancement is optimal (an iso-enhancement contour). Each contour line is labeled with its enhancement level, and its color corresponds to the same enhancement level on the efficient surface in Exhibit 6. The contour lines show that the optimal enhancement increases with leverage

EXHIBIT 6
Mean-Variance-Leverage Efficient Surface

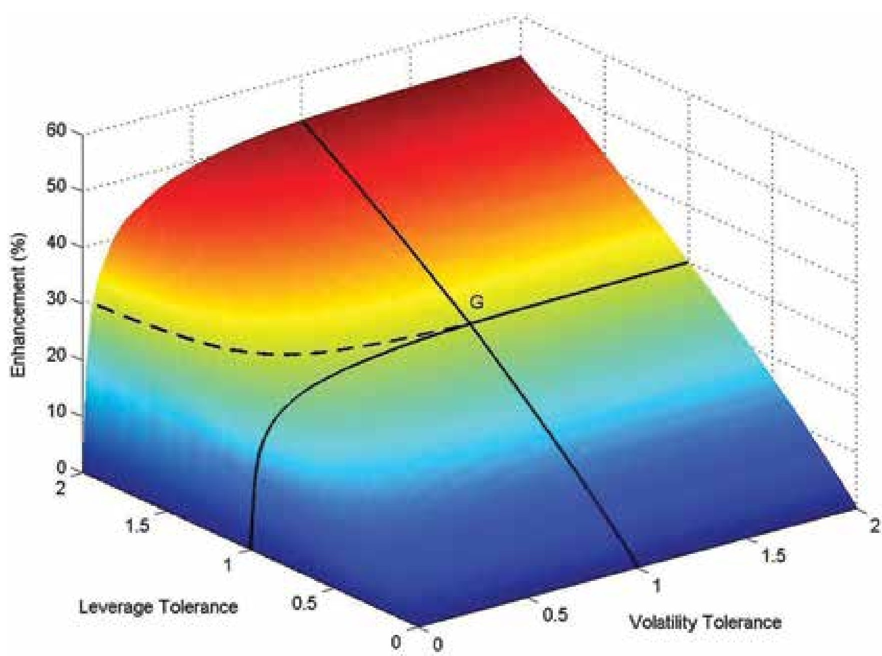

tolerance, but is approximately independent of volatility tolerance, if the latter is large enough.

The two solid black lines drawn on the efficient surface in Exhibits 6 and 7 correspond to optimal portfolios for investors having a volatility tolerance of 1 (and a range of values of leverage tolerance), and those for investors having a leverage tolerance of 1 (and a range of values of volatility tolerance). Consider an MVL(1,1) investor, whose optimal enhancement is the same as that of the iso-enhancement contour that passes through the intersection of the vertical and horizontal lines at point G. In this case, the optimal enhancement is 29%, resulting in a 129-29 EAE portfolio. This portfolio provides the MVL(1,1) investor the highest utility of all the portfolios on the efficient surface.

Portfolio "G," the optimal MVL(1,1) portfolio, has the same enhancement level as portfolio "g" in Exhibit 4. It also has the same standard deviation of active return and expected active return. In fact, portfolios "G" and "g" are identical: that is, they have the same holdings, and hence the same active weights.

Portfolio "g" was determined by considering numerous leverage-constrained MV(1) portfolios and selecting the one that has the highest utility for an MVL(1,1) investor. In contrast, we determined portfolio "G" directly from a leverage-unconstrained MVL(1,1) optimization. We will now show the equivalence of portfolios "g" and "G" by considering the efficient surface and contour map.

The solid black line representing optimal portfolios on the efficient surface or contour map at a volatility tolerance of 1 can be extended for levels of leverage tolerance beyond 2. Consider an MVL(1,∞) investor—that is, an investor with infinite leverage tolerance, or no leverage aversion. This investor is identical to an MV(1) investor with no leverage constraint. Now consider subjecting this investor to a leverage constraint, such that the enhancement is equal to 29%. With this constraint, portfolio "G" is the optimal portfolio for an MV(1) investor, as it is for a leverage-unconstrained MVL(1,1) investor.[17]

Alternatively, consider the 29% iso-enhancement contour in Exhibit 7 (or the dashed line in Exhibit 6).[18] This contour represents all portfolios on the efficient surface that have an enhancement of 29%. When the enhancement is constrained to equal 29%, the optimal

EXHIBIT 7
Contour Map of the Efficient Surface

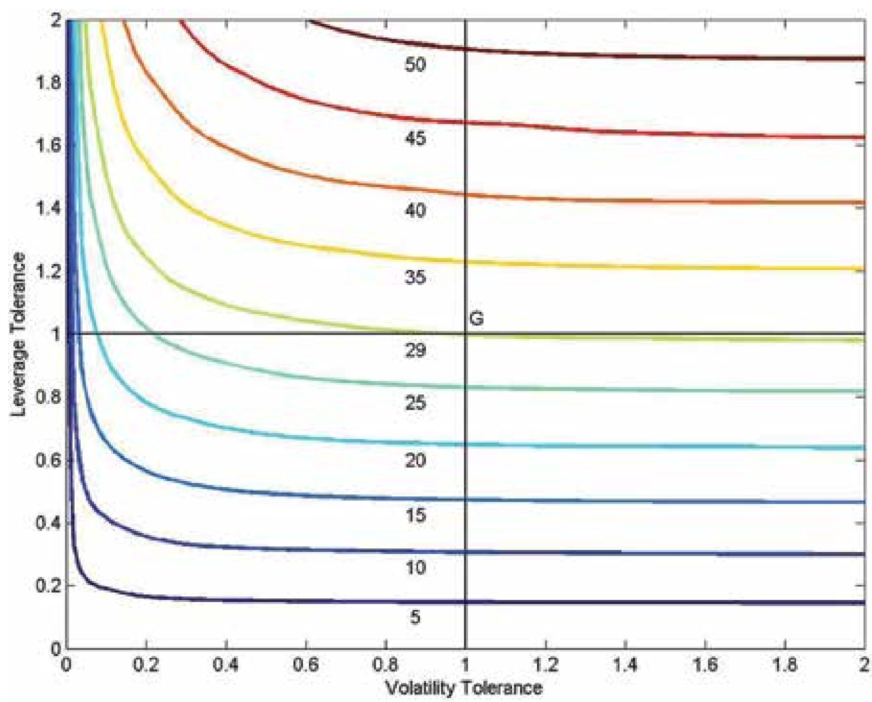

EXHIBIT 8

Characteristics of Optimal MVL(1,τ_L) Portfolios

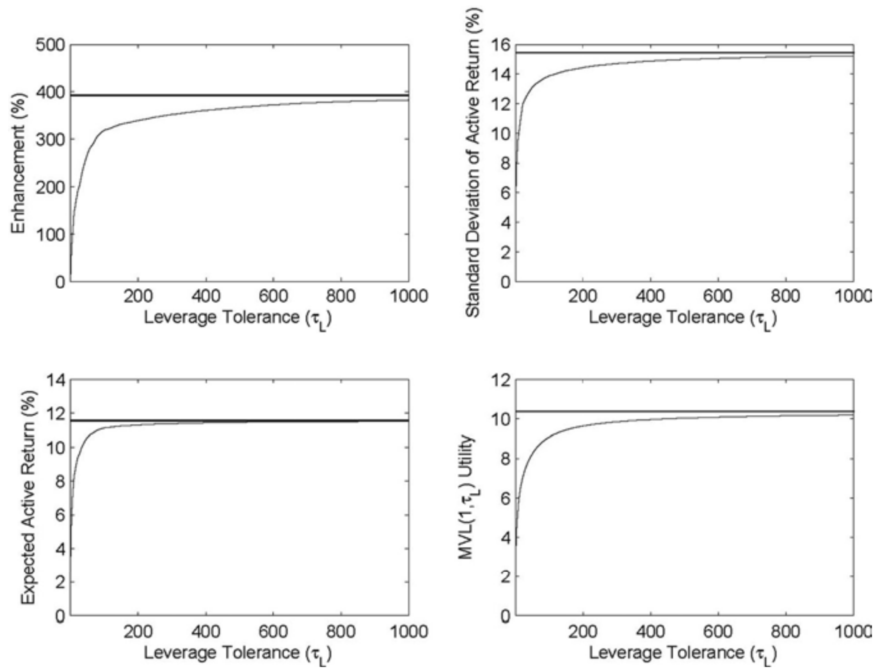

portfolio must be somewhere on the 29% contour. Optimal portfolios for investors with a volatility tolerance of 1 (whatever their leverage tolerance) lie on the solid black vertical line representing a volatility tolerance of 1. Thus, portfolio "G" (the point at which the 29% contour intersects the solid vertical line representing a volatility tolerance of 1) is optimal for an MV(1) investor who constrains the enhancement at 29%. Portfolios that are on the 29% contour, but not on the solid vertical line (representing a volatility tolerance of 1) would have lower utility than portfolio "G," because the implied volatility tolerance of those portfolios would either be less than or greater than 1, departing from the investor's volatility tolerance.

As we have discussed, as leverage tolerance approaches infinity, the optimal portfolios approach those determined by a conventional mean-variance utility function. Exhibit 8 shows the characteristics of optimal MVL(1, τ_L) portfolios as investor leverage tolerance, τ_L, increases in steps of 0.2, from near 0 to 1000. As before, the security active weights in these portfolios are constrained to be within 10 percentage points of the security weights in the benchmark index. The displayed characteristics are enhancement, standard deviation of active return, expected active return, and MVL(1, τ_L) utility. The horizontal asymptotes represent the levels associated with the optimal MV(1) portfolio "z," shown in Exhibit 3.

All the characteristics initially rise rapidly and continue to increase as they converge asymptotically to those of portfolio "z," as leverage tolerance approaches infinity. Except in the case of extreme leverage tolerance, the characteristics of the optimal MVL(1, τ_L) portfolios are quite different from those of the optimal MV(1) portfolio, which are represented by the asymptotes. Exhibit 8 shows that only by assuming an unreasonably large value for leverage tolerance would the solution to the MVL(1, τ_L) problem be close to that of the MV(1) portfolio.

CONCLUSION

Leverage entails a unique set of risks. In order to mitigate these risks, a leverage-averse investor can impose a leverage constraint in conventional mean-variance portfolio optimization. But mean-variance optimization provides the investor with little guidance as to where to set the leverage constraint.

In the absence of a leverage constraint and security active weight constraints, and given a level of volatility tolerance, an investor's mean–variance utility increases with leverage, up to an extremely high leverage level. Even in the presence of security active weight constraints, investor utility increases as leverage increases, up to a high leverage level. In either case, a mean–variance approach cannot identify the portfolio that offers the highest utility for a leverage-averse investor, because it does not consider the unique risks of leverage.

The optimal portfolio that offers the highest utility for a leverage-averse investor can only be identified if the investor's mean–variance–leverage utility function is known. The optimal portfolio and its leverage level can be determined by considering numerous conventional, leverage-constrained optimal mean–variance portfolios and evaluating each one by using the investor's mean–variance–leverage utility function to determine which portfolio offers the greatest utility. A more direct approach is to use mean–variance–leverage optimization to determine the optimal portfolio for a leverage-averse investor. Mean–variance–leverage optimization balances the portfolio's expected return against the portfolio's volatility risk *and* its leverage risk.

We have demonstrated that these two methods produce the same optimal portfolio. However, without knowing the investor's mean–variance–leverage utility function, conventional mean–variance optimization with a leverage constraint leads to the optimal portfolio for a leverage-averse investor only by chance.

ENDNOTES

We thank Judy Kimball, David Landis, and David Starer for helpful comments.

[1]Markowitz [1959] showed how to use individual security and portfolio constraints in optimization. With mean–variance optimization, constraints on leverage may be used to ensure compliance with regulations (Reg T, for instance) or client guidelines (such as for a 130-30 long-short portfolio). Such constraints can also be used with mean–variance–leverage optimization.

[2]The unique risks of leverage include the risks and costs of margin calls, which can force borrowers to liquidate securities at adverse prices due to illiquidity, losses exceeding the capital invested, and the possibility of bankruptcy.

[3]In response to the mean–variance–leverage optimization model proposed in Jacobs and Levy [2013a], Markowitz [2013] suggested a broader application of the general mean–variance portfolio selection model that differs from conventional usage. In particular, he suggested the development of a stochastic margin call model, which is yet to be developed. For a response to this suggestion, see Jacobs and Levy [2013b].

[4]Leverage is measured in excess of 1, that is, in excess of 100% of net capital.

[5]The data and estimation procedures are the same as those in Jacobs and Levy [2012]. Note that the specific numerical results in Exhibit 1 and throughout this article are dependent on the data and estimation procedures used, but the conclusions hold more generally.

[6]In practice there would be financing costs (such as stock loan fees); hard-to-borrow stocks may entail higher fees. For more on EAE portfolios, see Jacobs and Levy [2007].

[7]A value of $\tau_V = 0$ corresponds to an investor who is completely intolerant of active volatility risk, and a value of $\tau_V \approx 1$ causes quadratic utility of return to be equivalent to log-utility of wealth, a utility function often used in the finance literature (Levy and Markowitz [1979]). We used a range from 0.02 to 2 to generate Exhibit 1.

[8]The long-only efficient frontier converges to the origin: an index fund. The other frontiers cannot converge to a zero standard deviation of active return, because they are constrained to have an active enhancement, unless untrim positions are allowed. For a definition of untrim positions, see Jacobs et al. [2005, 2006].

[9]For enhancement levels beyond 392%, the expected active returns fall sharply, because the additional leverage needs to be met with additional security positions while still satisfying the active security weight constraint. This requires taking positions in securities that are detrimental to expected active returns.

[10]An MV(1) investor would be indifferent to a choice between each of the portfolios shown having a particular volatility risk and expected return, and a hypothetical portfolio having zero volatility risk and offering a certain return that is equal to the utility level shown.

[11]As we noted in Jacobs and Levy [2012], we assume that investors have the same aversion to leveraged long positions that they have to short positions. This assumption may not be the case in practice, because short positions have potentially unlimited liability and are susceptible to short squeezes. One could model the aversion to long and short positions asymmetrically, but this would complicate the algebra, so for simplicity we use a common leverage tolerance.

[12]An MVL(1,1) investor would be indifferent to a choice between each of the portfolios shown with a particular volatility risk, leverage risk, and expected return, and a hypothetical portfolio with zero volatility risk and zero leverage risk that offers a return that is equal to the utility level shown.

[13]Portfolios are subject to the standard EAE constraint set and the constraint that each security's active weight must be within 10 percentage points of its benchmark weight.

[14]In practice, the utility function in Equation (13) is difficult to optimize, because the leverage-risk term requires powers up to and including the fourth order in the x_i terms. To solve for optimal portfolios with this utility function, we use fixed-point iteration, as discussed in Jacobs and Levy [2013a].

[15]Tolerances for volatility and leverage can be greater than 2. As leverage tolerance approaches infinity, optimal portfolios will approach those determined by a conventional mean-variance utility function. This is because the augmented utility function (Equation (11)) reduces to the mean-variance utility function (Equation (1)) as the investor's leverage tolerance increases without limit.

[16]To estimate their tolerances for volatility and leverage, investors could select different portfolios from the efficient surface, and for each portfolio run a Monte Carlo simulation that generates a probability distribution of ending wealth. Investors could then infer their volatility and leverage tolerances based on their preferred ending wealth distribution. Alternatively, investors could use asynchronous simulation, which can account for the occurrences of margin calls, including security liquidations at adverse prices (see Jacobs et al. [2004, 2010]).

[17]Note that, in the absence of leverage constraints, an MVL(1,0) investor (with zero leverage tolerance) will hold a long-only portfolio at the intersection of the line for a volatility tolerance of 1 and the line for a leverage tolerance of 0. This MVL(1,0) investor is identical to an MV(1) investor with a leverage constraint of zero (long-only portfolio). At the other extreme of the line for a volatility tolerance of 1, consider an MVL(1,∞) investor. This investor is identical to an MV(1) investor with no leverage constraint. We have shown that for such an investor (subject to security active weight constraints of 10 percentage points), a leverage of 7.84 times net capital provides the highest utility. The optimal portfolio for an MVL(1,∞) investor (or an MV(1) investor with no leverage constraint) is located in the far distance on the MVL(1,τ_L) line. Between these two extremes are MVL(1,τ_L) investors with leverage tolerances, τ_L, between zero and infinity, or equivalently, MV(1) investors with leverage constraints between zero and 7.84. Thus, given an enhancement constraint that equals 29%, portfolio "G" is optimal for an MV(1) investor or for an MVL(1,τ_L) investor with any level of leverage tolerance τ_L.

[18]To the right of portfolio "G" in Exhibit 6, the dashed line is slightly below the solid line, but is visually indistinguishable from it.

REFERENCES

Jacobs, B., and K. Levy. "20 Myths about Enhanced Active 120-20 Strategies." *Financial Analysts Journal*, Vol. 63, No. 4 (2007), pp. 19-26.

——. "Leverage Aversion and Portfolio Optimality." *Financial Analysts Journal*, Vol. 68, No. 5 (2012), pp. 89-94.

——. "Leverage Aversion, Efficient Frontiers, and the Efficient Region." *The Journal of Portfolio Management*, Vol. 39, No. 3 (2013a), pp. 54-64.

——. "A Comparison of the Mean-Variance-Leverage Optimization Model and the Markowitz General Mean-Variance Portfolio Selection Model." *The Journal of Portfolio Management*, Vol. 40, No. 1 (2013b), pp. 1-5.

Jacobs, B., K. Levy, and H. Markowitz. "Financial Market Simulation." *The Journal of Portfolio Management*, Vol. 30, No. 5 (2004), pp. 142-151.

——. "Portfolio Optimization with Factors, Scenarios, and Realistic Short Positions." *Operations Research*, Vol. 53, No. 4 (2005), pp. 586-599.

——. "Trimability and Fast Optimization of Long-Short Portfolios." *Financial Analysts Journal*, Vol. 62, No. 2 (2006), pp. 36-46.

——."Simulating Security Markets in Dynamic and Equilibrium Modes." *Financial Analysts Journal*, Vol. 66, No. 5 (2010), pp. 42-53.

Levy, H., and H. Markowitz. "Approximating Expected Utility by a Function of Mean and Variance." *The American Economic Review*, Vol. 69, No. 3 (1979), pp. 308-317.

Markowitz, H., *Portfolio Selection: Efficient Diversification of Investments*, Yale University Press, 1959; 2nd ed. Cambridge, MA: Blackwell Publishing, 1991.

——. "How to Represent Mark-to-Market Possibilities with the General Portfolio Selection Model." *The Journal of Portfolio Management*, Vol. 39, No. 4 (2013), pp. 1-3.