

Reply to “(Im)Possible Frontiers: A Comment”

Thomas J. Brennan, Northwestern University

Andrew W. Lo, MIT Sloan

ABSTRACT

In Brennan and Lo (2010), a mean-variance efficient frontier is defined as “impossible” if every portfolio on that frontier has negative weights, which is incompatible with the Capital Asset Pricing Model (CAPM) requirement that the market portfolio is mean-variance efficient. We prove that as the number of assets n grows, the probability that a randomly chosen frontier is impossible tends to one at a geometric rate, implying that the set of parameters for which the CAPM holds is extremely rare. Levy and Roll (2014) argue that while such “possible” frontiers are rare, they are ubiquitous. In this reply, we show that this is not the case; possible frontiers are not only rare, but they occupy an isolated region of mean-variance parameter space that becomes increasingly remote as n increases. Ingersoll (2014) observes that parameter restrictions can rule out impossible frontiers, but in many cases these restrictions contradict empirical fact and must be questioned rather than blindly imposed.

Keywords: CAPM; Mean-Variance Analysis, Portfolio Optimization, Roll Critique, Shortselling, Long/Short.

JEL Codes: G11, G12

The views and opinions expressed in this article are those of the authors only, and do not necessarily represent the views and opinions of any institution or agency, any of their affiliates or employees, or any of the individuals acknowledged above. Research support from the MIT Laboratory for Financial Engineering and the Northwestern University School of Law Faculty Research Program is gratefully acknowledged.

We are immensely gratified—and not a little intimidated—by the time and interest devoted to our paper (Brennan and Lo (2010)) by Moshe Levy and Richard Roll (hereafter LR), and by Jonathan Ingersoll. LR’s comment focuses on the robustness of the main conclusion of our paper: for any fixed vector of expected returns $\boldsymbol{\mu}$,¹ a randomly selected covariance matrix $\boldsymbol{\Sigma}$ will almost always result in a mean-variance efficient frontier on which *every* portfolio contains negative weights. We called such frontiers “impossible” because they cannot be consistent with the Capital Asset Pricing Model (CAPM), in which the tangency portfolio is the market portfolio and has strictly positive weights on all assets.

While not contradicting our conclusions, LR claim that “parameter sets leading to possible frontiers resemble rational numbers: the probability of randomly sampling one from the real number line is zero, but for any point on the real number line there is always a rational number nearby.” What they are referring to in this analogy is a well-known property of rational numbers, which is that no matter where you look on the real line, there is always a rational number nearby. For example, π is not a rational number, but it can be shown that every interval $(\pi - \epsilon, \pi + \epsilon)$ around π contains a rational number, no matter how small ϵ is. In this sense, while the set of rational numbers is “small” (more formally, this set has measure zero), rational numbers are ubiquitous. Like mosquitoes on a warm summer evening, no matter where you are there are always a few nearby but they take up virtually no space and cannot easily be rounded up and eliminated. Sets of numbers that have this property are said to be “dense” for obvious reasons.

In Section 1, we show that this analogy does not hold. The set of so-called “possible” frontiers—those frontiers that contain at least one portfolio with strictly positive weights—is not only extremely rare, but they fall in very well-defined and localized regions of mean-variance parameter space, unlike the rational numbers on the real line. However, this is not to say that possible frontiers cannot be found, and we present a brief summary in Section 2 of the method we proposed in Brennan and Lo (2010, Section 5.3) for finding the “closest” frontier that is possible for any given set of parameters. We comment on LR’s empirical results in Section 3, and show

¹Throughout this article, we maintain the following notational conventions: (1) all vectors are column vectors unless otherwise indicated; (2) matrix transposes are indicated by prime superscripts, hence $\boldsymbol{\mu}'$ is the transpose of $\boldsymbol{\mu}$; and (3) vectors and matrices are always typeset in boldface, i.e., ω_i and σ_i are scalars and $\boldsymbol{\omega}$ and $\boldsymbol{\Sigma}$ are vectors or matrices.

that their construction of possible frontiers requires economically large shifts in expected-return parameters, which may be hard to justify from a professional portfolio manager’s perspective.

In contrast, Ingersoll presents a theoretical argument that impossible frontiers can be ruled out, in the same way that no-arbitrage conditions are used to rule out doubling strategies that generate free lunches asymptotically. In Section 4, we show that Ingersoll’s theoretical restrictions are tantamount to assuming away the problem, i.e., assuming that efficient-frontier portfolios have positive weights.

However, despite these considerations, we conclude in Section 5 by arguing that our perspective may not be so different from LR’s and Ingersoll’s in that we all share an appreciation for the practical relevance of the CAPM as well as the significance of its limitations.

1 The Set of Possible Portfolios is not Dense

The starting point for LR’s analysis is an empirical exercise in which they solve for the set of parameters (μ^*, Σ^*) that yield a “possible frontier” (one with strictly positive weights) and which is closest to a given set of sample estimates $(\hat{\mu}, \hat{\Sigma})$ under a specific distance metric for mean-variance parameters of 100 randomly selected S&P 500 stocks from 1980 to 2005 (the same universe used by Brennan and Lo). These empirical results seem suggestive, but are not exactly a proof in the formal mathematical sense. This is not meant to be a “cheap shot”; the claim that there are always possible frontiers in every region of the mean-variance parameter space can be shown to be false. In particular, it is easy to show that there are many intervals around the parameters of impossible frontiers that do not contain any possible-frontier parameters. Consider, for example, the three-asset case of Brennan and Lo (2010, Section 3.2): our Proposition 2 shows that there exists a finite interval around the parameters in equation (6) for which all points in that subspace yield impossible frontiers. In fact, given that our construction of impossible frontiers involves continuous functions, it is easy to see that the subset of possible-frontier parameters cannot be dense for any number of assets $n \geq 2$.

Although it only takes one counter-example to prove that LR’s conjecture is false, there is a deeper sense in which the set of parameters yielding possible frontiers cannot be ubiquitous. Consider LR’s condition

(2), henceforth LR2, which specifies constraints on the parameters so that the resulting mean-variance efficient frontier is possible, i.e., has positive weights in all entries. With some algebraic manipulation, this condition can be rewritten as

$$\begin{bmatrix} \rho^{\text{sample}} \end{bmatrix} \begin{bmatrix} \omega_1 \\ \omega_2 \\ \vdots \\ \omega_n \end{bmatrix} = q \begin{bmatrix} (\mu_1 - r_z)/\sigma_1^2 \\ (\mu_2 - r_z)/\sigma_2^2 \\ \vdots \\ (\mu_n - r_z)/\sigma_n^2 \end{bmatrix} \quad (1)$$

where ρ^{sample} is the sample correlation matrix, r_z is the zero-beta expected rate of return, and $q > 0$ is a constant of proportionality. For a given set of parameters that yields an impossible frontier, LR's approach is to find the closest "possible" portfolio, i.e., one with strictly positive weights that satisfies (1).

In this simpler form, LR's constraint of strictly positive portfolio weights is, according to (1), essentially a restriction on the expected returns $\{\mu_i\}$ for a given fixed sample correlation matrix ρ^{sample} . Therefore, the region in the n -dimensional space of all possible portfolio weights $\{\omega_i\}$ that consists of vectors with strictly positive weights in all entries is just the positive orthant, which corresponds to a fraction of size $1/2^n$ of the entire space. According to (1), this region of possible-frontier weights is linearly mapped by multiplication by the correlation matrix to an open cone with a single connected component in μ -space.² Unlike the set of rational numbers, this cone of parameters actually has positive mass, so in this very technical sense, there are "more" parameters of this type than there are rational numbers.³ However, unlike the rational numbers, this cone is not ubiquitous and there

²A "cone" is a mathematical concept from linear algebra that refers to a subset of an n -dimensional coordinate system made up of a set of half-lines starting at the origin and extending out to infinity. A set is "open" if, for each point in the set, there is a neighborhood around that point that is completely contained in the set, and a set is "connected" if any two points in the set can be connected by a path that lies entirely within the set. The cone in μ -space is open and connected because it is the linear image of the positive orthant in ω -space, which is itself open and connected.

³More formally, the number of rationals is said to be countably infinite because there is a one-to-one correspondence between the rationals and the natural numbers (1, 2, etc.). However, the number of parameters in the cone (1) is said to be uncountably infinite because it can be shown that any attempt to create a one-to-one correspondence between the natural numbers and parameters in the convex cone will always miss an infinite number of elements.

are large regions of the μ parameter space where none of the parameters correspond to frontier portfolios with strictly positive weights. In fact, the complement of this cone consists of the boundary of the cone and an enormous open cone of impossible parameters. Specifically, the volume of n -dimensional μ -space containing impossible parameters is a fraction equal to $1 - C(\rho^{\text{sample}})/2^n$, where $C(\rho^{\text{sample}})$ is a constant that depends on the correlation matrix ρ^{sample} .⁴

In our paper, we show that for any arbitrary vector of means μ , the likelihood of finding a correlation matrix that satisfies (LR2) quickly becomes akin to finding a needle in a haystack. More formally, in addition to our analytical proof of impossibility in the limit as n goes to infinity, we compute lower bounds on the probability of encountering an impossible frontier when selecting correlation matrices at random (with a uniform distribution). For 100 assets, the likelihood of drawing an impossible frontier is 99.66%, a virtual certainty. These results explain why long-only portfolio managers always impose constraints when they employ mean-variance portfolio optimization techniques—without constraints, the optimal tangency portfolio will almost surely have negative weights, violating the long-only constraint.

In summary, our results show that the entire parameter space of all means and correlations is partitioned into two disjoint regions, one corresponding to impossible frontiers and the other corresponding to possible frontiers. The fact that these regions are so distinct implies that the set of possible-frontier parameters cannot be a dense set, which LR acknowledges, and is nowhere close to being a dense set, which LR incorrectly conjectures. Moreover, allowing the correlation parameters to vary only makes matters worse, contrary to LR’s proposed intuition. Adding degrees of freedom does not change the fact that the set of possible parameters is still its own distinct region, and it only serves to make the relative size of this region smaller, as we show in great detail in Brennan and Lo (2010).

⁴ The value of $C(\rho^{\text{sample}})$ is equal to the product of $\det(\rho^{\text{sample}})$, which lies between 0 and 1, and the ratio of the integral of $\exp(-\omega' \rho^{\text{sample}} \omega)$ to the integral of $\exp(-\omega' \omega)$, with both integrals taken over the positive orthant in ω -space.

2 Finding the Closest Frontier that is Possible

Nevertheless, determining the particular regions of the parameter space that do yield strictly positive frontier portfolios is still a worthwhile exercise. In Brennan and Lo (2010, Section 5.3), we provide a simple and general way to do this using the framework of Black and Litterman (1992), which is summarized in the Appendix of this article. In particular, starting with an arbitrary set of mean-variance parameters under the assumption of a linear one-factor (market) model, we derive the set of possible-frontier parameters that are “most consistent” with the initial set of parameter values, and show that this corresponds to the Black and Litterman (1992) weights. The two differences between this approach and LR’s are that: (1) ours is a theoretical result that holds in general, while LR’s is an empirical result that may or may not generalize; and (2) LR fix the correlations and vary the assets’ expected returns and variances while we fix the expected returns and vary the covariance matrix. However, our approach can easily be applied to the expected-return vector (as in the original Black and Litterman (1992) paper).

LR’s focus on means and variances rather than on covariances leads to some interesting insights into the impossible/possible frontiers distinction that are worth pointing out. First, the fact that LR choose to optimize over the means and variances is no accident. There are fewer degrees of freedom (only $2 \cdot n$ versus $n \cdot (n - 1) / 2$, or 200 versus 4,950 in the case of 100 stocks), and it is easier to develop economic intuition for means and variances versus correlations. However, in practice, correlations can change as abruptly as means and variances, and it is not clear that LR’s results will generalize to situations involving uncertainty regarding correlations. Put another way, by considering random correlation matrices, the parameter space is greatly enlarged, and much of that space contains impossible parameters.

3 Limitations of LR’s Impossible-to-possible Transformation

The limitations of LR’s approach to constructing possible frontiers can be observed in their empirical example summarized in their Tables 1 and 2. Although LR claims that “These results show that even a small adjustment to the sample parameters yields an efficient frontier with a rather substantial

segment of positive portfolios” (p. 6), in fact the adjustments are quite large. In Table 1 below, we reproduce LR’s parameters from their Table 1, but we have annualized the means (geometrically, by computing $(1 + r)^{12} - 1$) and the standard deviations (by multiplying by $\sqrt{12}$). Note that the adjustments to the means of stocks 6, 11, and 19 (highlighted in red) are 9.8% vs. 20.0%, 9.0% vs. 17.6%, and 1.3% vs. 19.6%. The first two adjustments require doubling the expected return while the third adjustment involves a 15-fold increase.

The intuition for why such extreme adjustments are necessary is simple: holding correlations fixed, the adjustments needed to turn impossible-frontier means into possible-frontier means are to increase the lowest means enough to reduce the incentive to short those assets. While the adjustments in LR’s Tables 1 and 2 may seem “small” in terms of sampling variation, any portfolio manager will attest to the fact that a target expected return of 20.0% is quite different than 9.8%. A more striking comparison is given in the last two columns of our Table 1, which contains the Sharpe ratios of the sample estimates and the adjusted parameters. The Sharpe ratio of Stock 19 based on the estimated parameters is 0.03, but the Sharpe ratio of the adjusted parameters is 0.44—increasing an asset’s Sharpe ratio by a factor of 15 can hardly be viewed as a “small” adjustment. In fact, the LR adjustments from impossible-to-possible frontiers are only small in relation to the estimation errors of the parameters, but such errors are well known to be large for the expected returns of individual stocks (Merton (1980)). While the adjustments may not always be statistically significant because of the large standard errors surrounding expected-return estimates, they are certainly economically significant and difficult to justify from a professional portfolio manager’s perspective, not to mention clients’ and regulators’ perspectives.

In summary, LR’s impossible-to-possible-frontier adjustments involve changing the parameters so as to avoid short sales, which requires the reward-to-risk ratio of positively correlated assets to be roughly comparable. If they are too dissimilar, then it will be possible to construct a portfolio with higher risk-adjusted expected returns by shorting the less attractive assets and using the proceeds to leverage the positions of the more attractive assets. This is illustrated by comparing the cross-sectional standard deviations of the unadjusted and adjusted Sharpe ratios in Table 1, which is 0.071 for the sample and 0.043 for the possible-frontiers case; possible-frontier reward-to-risk ratios cannot be too diverse, otherwise short sales will be op-

Stock	$\hat{\mu}$	μ^*	$\hat{\sigma}$	σ^*	$\hat{\mu}/\hat{\sigma}$	μ^*/σ^*
1	15.0%	18.9%	29.9%	27.9%	0.50	0.68
2	19.6%	18.2%	24.0%	24.4%	0.81	0.75
3	15.4%	18.3%	28.1%	27.3%	0.55	0.67
4	17.3%	18.3%	33.6%	33.0%	0.51	0.55
5	24.5%	19.8%	37.7%	38.4%	0.65	0.52
6	9.8%	20.0%	51.5%	48.3%	0.19	0.41
7	19.7%	17.5%	22.5%	23.1%	0.88	0.76
8	17.7%	16.5%	18.0%	18.3%	0.98	0.90
9	15.4%	16.5%	17.2%	17.1%	0.89	0.96
10	14.7%	18.3%	23.7%	22.6%	0.62	0.81
11	9.0%	17.6%	26.3%	24.5%	0.34	0.72
12	18.6%	18.4%	24.7%	24.5%	0.75	0.75
13	18.4%	17.2%	17.4%	17.6%	1.06	0.98
14	20.0%	18.2%	28.4%	28.7%	0.70	0.63
15	20.7%	18.0%	22.2%	22.8%	0.93	0.79
16	17.0%	19.7%	33.8%	32.9%	0.50	0.60
17	19.0%	19.6%	32.2%	31.8%	0.59	0.62
18	13.5%	17.7%	27.9%	27.0%	0.48	0.66
19	1.3%	19.6%	48.4%	44.0%	0.03	0.44
20	13.0%	19.1%	38.1%	36.4%	0.34	0.53

Table 1: Comparison of annualized sample and possible-frontier parameters.

timal.⁵ We provide an analytical exposition of this intuition in the Appendix for the special case of n equally correlated assets, which shows that LR's possible-frontier restrictions amount to non-negativity restrictions on some transformation of the vector of mean-variance parameters. In this transformed coordinate space, non-negative coordinates lie in only one of 2^n possible orthants, hence this restricted portion of the parameter space is quite special and cannot be "close" to every possible set of sample parameters.

⁵We realize that there are only 20 out of 100 portfolio weights in Tables 1 and 2, but we did not have access to LR's supplementary information so these are all we had to work with.

4 Ingersoll’s Critique

Subsequent to LR’s critique, Ingersoll (2014) has also provided his perspective on our results, for which we are equally grateful. As a consummate theorist, his analysis is naturally focused on theoretical restrictions that rule out impossible frontiers. For example, in his Section II, Ingersoll provides an example of a complete-markets economy in which arbitrage is asymptotically certain if no restrictions are imposed on the underlying parameters so as to yield positive state prices. Despite the unrealistic nature of this example, it highlights a major philosophical difference between Ingersoll’s—and, by extension, the typical theorist’s—view of the world and ours. In our view, theory is meant to be an approximation to a much more complex reality, hence the success or failure of theory should be measured by the magnitude of its approximation error rather than its aesthetic appeal. By asserting that equilibrium must hold under any circumstances—and irrespective of empirical evidence—theorists have summarily dismissed large regions of the space of parameters describing the economy, including regions that may very well be relevant for the economy that we currently inhabit.

This philosophical difference is most evident in Ingersoll’s more substantive and technical analysis in Section III, in which he begins with the premise that prices are given by:

$$\mathbf{p} = (\bar{\mathbf{x}} - a\Xi\mathbf{h})/b \quad (2)$$

for some arbitrary positive values of a and b , where $\boldsymbol{\mu} = \mathbf{P}^{-1}\bar{\mathbf{x}}$, $\boldsymbol{\Sigma} = \mathbf{P}^{-1}\Xi\mathbf{P}^{-1}$, \mathbf{p} is the vector of prices, and $\mathbf{P} \equiv \text{diag}(\mathbf{p})$ is the diagonal matrix of prices. The key result claimed by Ingersoll is that the values of the elements of \mathbf{h} may be specified arbitrarily, and that if they are all positive, then there is a portfolio with all positive weights on the mean-variance efficient frontier corresponding to the mean return vector $\boldsymbol{\mu}$ and the covariance matrix $\boldsymbol{\Sigma}$. In other words, the claim is that the frontier in question is in fact possible.

To elucidate Ingersoll’s analysis further, we note that the formula in (2) is algebraically equivalent to

$$\mathbf{P}\mathbf{h} = \boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu} - b\boldsymbol{\iota})/a. \quad (3)$$

Because prices are positive, (3) shows that the assumption that \mathbf{h} has all positive elements is equivalent to the assumption that the portfolio

$\omega_{I1} = \Sigma^{-1}(\mu - b\iota)/a$ has all positive elements. Ingersoll then asserts that a portfolio, ω_{I2} , which is proportional to ω_{I1} and hence also has all positive elements, lies on the relevant efficient frontier. We agree that such a portfolio ω_{I2} exists on the frontier, although it should be noted that ω_{I2} can lie on the efficient frontier only if $b \leq B/C$.⁶ More importantly, however, we note that restricting \mathbf{h} to have positive components is equivalent to restricting ω_{I2} to have positive components, a circular argument that essentially assumes the conclusion.

It is helpful to consider the literature referenced by Ingersoll to gain additional perspective. Ingersoll references Jensen and Long (1972) who offer a version of Ingersoll (2014, Equation (1)) without proof. Instead, Jensen and Long (1972) provide a lengthy literature review that cites Lintner (1965), also referenced in our paper, who derives the market portfolio in a world of mean-variance optimizing agents, and concludes that prices (and market weights) are the result of mean-variance optimization. This, of course, leads naturally to the result that the market portfolio is on the efficient frontier. If the market portfolio is mean-variance efficient, then, tautologically, this portfolio lies on the efficient frontier. The point of our paper was, however, to show that if we begin with expected returns and covariances as the starting point—as required by any empirical implementation of the CAPM—then the market portfolio generated by such parameters is almost never mean-variance efficient.

Finally, we do point out in Brennan and Lo (2010, Section 5) that when idiosyncratic risks are uncorrelated and a one-factor model applies, the resulting frontier becomes possible. This is a very specific example of a theoretical restriction that rules out impossible frontiers. In fact, as discussed above, we provide a much more general construction of possible frontiers via the Black and Litterman (1992) approach to computing a covariance matrix that is “closest” to the observed covariance matrix and

⁶The portfolios on the efficient frontier corresponding to μ and Σ can be expressed in the form

$$\left\{ \omega: \omega = \frac{C}{D} \left(\mu_0 - \frac{B}{C} \right) \Sigma^{-1} \left(\mu - \left(\frac{B}{C} - \frac{D}{C^2 \cdot (\mu_0 - B/C)} \right) \iota \right), \text{ for } \mu_0 \geq \frac{B}{C} \right\}, \quad (4)$$

where $A = \mu' \Sigma^{-1} \mu$, $B = \mu' \Sigma^{-1} \iota$, $C = \iota' \Sigma^{-1} \iota$, and $D = AC - B^2$. This is a straightforward reformulation of the expression for the set of portfolios we set forth in our paper. It is clear that there exists an efficient portfolio proportional to ω_{I1} , provided that $b \leq B/C$. Note that Ingersoll uses gross returns rather than net returns, and so in his notation the correct way to express the upper bound for b would be $b - 1 \leq B/C$.

which yields a possible frontier. However, the question still remains as to whether or not such theoretical restrictions are plausible when confronted with empirical information to the contrary.

5 Conclusion

Despite the limitations of LR’s empirical analysis and Ingersoll’s theoretical perspective, we believe we share quite a bit of common ground with their conclusions, even if we have arrived at that shared perspective from a very different path. In particular, we certainly did not pronounce the CAPM “dead” in our paper. Like any model of economic equilibrium, the CAPM is a highly specialized parametrization that imposes many restrictions on market variables. The fact that these restrictions are a set of measure zero does not mean the CAPM is false, but rather that we need to have very strong prior beliefs before we take the CAPM literally. Accordingly, such priors should be made explicit to both portfolio managers and investors, and they should understand how sensitive their decisions are to those priors. In fact, the main thrust of Section 5.3 of Brennan and Lo (2010) was to illustrate just how we should make proper use of the CAPM, not to discard it completely.

In this respect, we believe our results are in the same spirit as the highly cited critique of the CAPM of Roll (1977), in which he observed that tests of the CAPM all amount to tests of the mean-variance efficiency of the market portfolio, and that this portfolio is practically unobservable because it must include all assets, e.g., real estate, art collections, and human capital, not just publicly traded stocks. Our findings may be viewed as extending the Roll critique in a rather different direction: we show that the set of parameters for which the market portfolio can be mean-variance efficient is extremely limited. Our results provide a compelling explanation for the problems that long-only portfolio managers face when applying portfolio optimization techniques, irrespective of the CAPM, and it is useful to know that these practical issues are not just due to estimation error, but are much more fundamental.

Finally, to add to LR’s defense of the CAPM, Sharpe’s beautiful model was the first to show how portfolio optimization could actually be implemented by everyone simultaneously in an equilibrium. The fact that this cannot hold for arbitrary mean-variance parameters should come as no

surprise, in much the same way that the probability of randomly played notes on a piano yielding Beethoven's *Moonlight Sonata* is a measure-zero event. The point of our analysis was not to discredit the CAPM. After all, all economic models are wrong; the more relevant question is how close an approximation any given model is to reality? Our main point was to show that the inconsistencies between the CAPM and empirical observation are more fundamental than just estimation error, so we need to be more explicit about imposing prior beliefs—as in Black and Litterman (1992), for example—when using this model.

Like Ingersoll, the legendary Fischer Black was also partial to theory and once said that if the data disagreed with the theory, he would discard the data. There is a certain wisdom to this striking conclusion given that financial data are noisy, dirty, and often fraught with fat tails, nonstationarities, and other theoretical infelicities. However, having grappled with financial data and institutional details that are typically ignored by theory for the sake of tractability, we have developed a somewhat greater respect for empirical phenomena, especially when they are at odds with theoretical implications.

We thank Professors Levy, Roll, and Ingersoll for their insightful comments on our work, and are grateful to *Critical Finance Review* for giving us this opportunity to engage in constructive and illuminating dialogue with such distinguished colleagues.

A Appendix

A.1 Finding the Nearest Possible Frontier

Suppose we begin with the assumption that asset returns follow a simple linear one-factor model:

$$\mathbf{R} = \iota R_f + \boldsymbol{\mu} \cdot (R_m - R_f) + \boldsymbol{\epsilon} \quad (\text{A1})$$

where R_m is the stochastic market return, $\boldsymbol{\mu}$ is an $(n \times 1)$ vector of constants, and $\boldsymbol{\epsilon}$ is an $(n \times 1)$ stochastic vector of idiosyncratic shocks. We assume that the expected value of $\boldsymbol{\epsilon}$ is zero.

Suppose that a mean return vector, $\boldsymbol{\mu}$, and a market-capitalization weight vector, $\boldsymbol{\omega}_m$, are given, and consider a covariance matrix $\boldsymbol{\Sigma}$ that is derived either empirically or from prior information, but which is not

necessarily compatible with $\boldsymbol{\mu}$ and $\boldsymbol{\omega}_m$ in the sense that $\boldsymbol{\omega}_m \neq \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$, as required by the CAPM. The matrix most compatible with the observed $\boldsymbol{\Sigma}$ but still conforming to the known values of $\boldsymbol{\mu}$ and $\boldsymbol{\omega}_m$ can be determined in the following manner. From Brennan and Lo (2010, Section 5.1), note that $\boldsymbol{\Sigma}$ may be expressed uniquely in the form

$$\boldsymbol{\Sigma} = \mathbf{A} \begin{pmatrix} w + \mathbf{x}'\mathbf{V}\mathbf{x} & \mathbf{x}'\mathbf{V} \\ \mathbf{V}\mathbf{x} & \mathbf{V} \end{pmatrix} \mathbf{A}', \tag{A2}$$

where \mathbf{A} is the matrix that takes the first coordinate vector, $\mathbf{e}_1 \equiv [1 \ 0 \ \dots \ 0]'$, to $\boldsymbol{\mu}$, and takes each other coordinate vector to itself, and where $(w, \mathbf{x}, \mathbf{V})$ can be thought of as “coordinates” for $\boldsymbol{\Sigma}$ satisfying the conditions that $w > 0$, $\mathbf{x} \in \mathcal{R}^{n-1}$, and \mathbf{V} is a covariance matrix of dimension $(n-1) \times (n-1)$. Also, following Brennan and Lo (2010, Section 5.3), define an alternative value for \mathbf{x} , written $\tilde{\mathbf{x}}$, by

$$\tilde{x}_i \equiv \frac{-\omega_{m,i+1}}{\mu_1 + (\mu_2 - \mu_1) \cdot \omega_{m,2} + \dots + (\mu_n - \mu_1) \cdot \omega_{m,n}}, \tag{A3}$$

and define an alternative value for $\boldsymbol{\Sigma}$, written $\tilde{\boldsymbol{\Sigma}}$, by

$$\tilde{\boldsymbol{\Sigma}} = \mathbf{A} \begin{pmatrix} w + \tilde{\mathbf{x}}'\mathbf{V}\tilde{\mathbf{x}} & \tilde{\mathbf{x}}'\mathbf{V} \\ \mathbf{V}\tilde{\mathbf{x}} & \mathbf{V} \end{pmatrix} \mathbf{A}'. \tag{A4}$$

This expression for $\tilde{\boldsymbol{\Sigma}}$ is the same as that for $\boldsymbol{\Sigma}$ above, except that \mathbf{x} is replaced by $\tilde{\mathbf{x}}$. This new covariance matrix, $\tilde{\boldsymbol{\Sigma}}$, is then compatible with $\boldsymbol{\omega}_m$ and $\boldsymbol{\mu}$ in that $\boldsymbol{\omega}_m$ is the tangency portfolio resulting from this mean and covariance. In addition, $\tilde{\boldsymbol{\Sigma}}$ is the covariance matrix most compatible with the specified values of $\boldsymbol{\mu}$ and $\boldsymbol{\omega}_m$ and the observed value of $\boldsymbol{\Sigma}$ in that it requires precisely the amount of alteration to $\boldsymbol{\Sigma}$ needed to make the three sets of parameters compatible.

Therefore, for those who have strong conviction that the CAPM must hold and that $\boldsymbol{\mu}$ and $\boldsymbol{\omega}_m$ are, in fact, the correct expected returns and market weights, and $\boldsymbol{\Sigma}$ is their best estimate of the covariance matrix, the covariance matrix they should adopt is $\tilde{\boldsymbol{\Sigma}}$ given in (A4). This is essentially the approach taken by Black and Litterman (1992).

A.2 Equi-Correlated Example of Impossible-to-Possible Frontiers

The empirical observation in Table 1 can be made analytically in the very simple but instructive case of n equally correlated stocks with pairwise cor-

relation ρ and identical variances σ^2 , which yields the following covariance matrix:

$$\Sigma = \sigma^2 \cdot (1 - \rho)\mathbf{I} + \sigma^2 \cdot \rho \mathbf{u}\mathbf{u}', \quad -\frac{1}{n-1} < \rho < 1 \quad (\text{A5})$$

where \mathbf{I} is the $(n \times n)$ identity matrix, \mathbf{u} is an $(n \times 1)$ -vector of 1's, and the inequality constraint on ρ ensures that Σ is positive definite. The mean-variance optimal portfolio is proportional to $\Sigma^{-1}\boldsymbol{\mu}$, and in this simple case, we can evaluate this expression explicitly as:

$$\boldsymbol{\omega}^{\text{optimal}} \propto \Sigma^{-1}\boldsymbol{\mu} = \frac{1}{\sigma^2 \cdot (1 - \rho)} \left[\mathbf{I} - \frac{\rho}{1 + (n-1) \cdot \rho} \mathbf{u}\mathbf{u}' \right] \boldsymbol{\mu} \quad (\text{A6})$$

$$\omega_i^{\text{optimal}} \propto \mu_i - \frac{n \cdot \rho}{(n-1) \cdot \rho + 1} \bar{\mu}, \quad \bar{\mu} \equiv \boldsymbol{\mu}'\mathbf{u}/n. \quad (\text{A7})$$

This expression shows clearly the intuition that the stocks with negative mean-variance optimal weights are those with lower expected returns. For large n , (A7) implies that all stocks with expected returns below the cross-sectional average expected return $\bar{\mu}$ will be shorted, and those with expected returns above this average will be held long. As a function of the correlation ρ , (A7) confirms the intuition that higher correlation implies more short positions, other things equal. For example, when $\rho = 0$ all weights are positive (assuming all expected returns are positive) and proportional to the stocks' expected returns. As ρ increases, the threshold that μ_i needs to exceed to ensure a positive mean-variance optimal portfolio weight also increases, making it more difficult to maintain all non-negative weights. For $n = 100$ and $\rho = 0.5$, all stocks with μ_i less than $0.99 \times \bar{\mu}$ will have negative weights. If the cross-sectional distribution of expected returns $\{\mu_i\}$ were symmetric, we would expect approximately half of the stocks to be below the cross-sectional mean $\bar{\mu}$ and half to be above, implying a large number of short positions.

In fact, the possible-frontier restriction on the set of expected returns $\{\mu_i\}$ can be expressed compactly as:

$$\mathbf{A}\boldsymbol{\mu} \geq 0, \quad \mathbf{A} \equiv \mathbf{I} - \frac{\rho}{1 + (n-1) \cdot \rho} \mathbf{u}\mathbf{u}' \quad (\text{A8})$$

which is simply the positive orthant of the space of expected returns under the new coordinate system defined by \mathbf{A} . This interpretation shows that LR's claim that possible-frontier parameters are "like the rationals" is

incorrect. In general, the positive orthant of any coordinate system will, by construction, only occupy $1/2^n$ of the total space (for example, in two dimensions, the positive orthant is only one of four quadrants; in three dimensions, it is only one of eight), and before transformation by \mathbf{A} , the fraction of μ -space corresponding to possible frontiers is $C(\Sigma)/2^n$.⁷ As n gets larger, the region of possible-frontier parameters does get smaller and more “localized” geometrically fast, as shown more formally and generally in Brennan and Lo (2010).

This example can be generalized considerably to arbitrary correlation matrices and heterogeneous variances (along the lines of the Brennan and Lo (2010) Section 4 coordinate transformation), but the thrust is the same: LR’s possible-frontier restrictions amount to non-negativity restrictions on some transformation of the vector of mean-variance parameters, hence that restricted portion of the parameter space cannot be “close” to every possible set of sample parameters.

References

- Black, F. and R. Litterman. 1992. “Global Portfolio Optimization”. *Financial Analysts Journal*. 47: 28–43.
- Brennan, T. and A. Lo. 2010. “Impossible Frontiers”. *Management Science*. 56: 905–923.
- Ingersoll Jr., J. E. 2014. “Always Possible Frontiers”. *Critical Finance Review*.
- Levy, M. and R. Roll. 2014. “(Im)Possible Frontiers: A Comment”. *Critical Finance Review*.
- Merton, R. 1980. “On Estimating the Expected Return on the Market: An Exploratory Investigation”. *Journal of Financial Economics*. 8: 323–361.
- Roll, R. 1977. “A Critique of the Asset Pricing Theory’s Tests Part I: On Past and Potential Testability of the Theory”. *Journal of Financial Economics*. 4: 129–176.

⁷See footnote 4 for the precise definition of $C(\Sigma)$.