MSCI

Research Insight

Managing Odd Lot Trades with the Barra Optimizer

Roundlotting Tradeoffs in Portfolio Construction

Rong Xu and Scott Liu

Rong.Xu@msci.com

Scott.Liu@msci.com

September 2013

Introduction

Quantitative portfolio construction relies on a Markowitz (1952) framework or its variations. The popularity and appeal of the approach can be attributed to its theoretical elegance and computational tractability. One of its practical drawbacks, however, is that standard mean-variance optimization seldom yields round lots on its own. The presence of numerous odd lots in an optimized trade list[1] is not only a nuisance, but also could result in additional transaction costs when implemented. This casts doubt on the traditional solutions of the mean-variance framework.

Complementing the standard mean-variance optimization with round lot constraints may sound like an easy thing to do. However, solving such a problem to optimality is not a trivial task. In this paper, we reveal how effectively the Barra Optimizer handles the round lot optimization problem.

We would like to emphasize that optimal roundlotting[2] based on heuristics does not guarantee portfolio optimality. Nevertheless, a successfully returned "optimal" portfolio will definitely be feasible under all the user-imposed constraints.[3] By comparison, post-optimization roundlotting[4] is simple and straightforward, yet the resulting portfolio may violate one or more constraints. Barra Optimizer offers both optimal roundlotting and post-optimization roundlotting capabilities to assist managers in cost-effective portfolio construction.

Why is Round Lot Optimization Difficult?

In theory, standard mean-variance optimization belongs to the class of convex quadratic programming that possesses a number of desirable properties. One example is that any convex combination of feasible portfolios will remain feasible. Another one is that a local optimal solution must also be globally optimal. These properties make the search for an optimal solution potentially tractable and straightforward. Convex quadratic problems are well-behaved and efficient algorithms are available to solve them. A globally optimal portfolio, if it exists, will always be found and returned.

In practice, however, requiring the trades to be all in round lots complicates the optimization process significantly. Round lot constraints are discrete and non-convex in nature and, in general, they are mathematically complicated and computationally challenging. A globally optimal solution may not exist. And even if it does exist, it is usually difficult or extremely time-consuming to find.

If the problem size is small, integer programming algorithms such as the "branch-and-bound" or "branch-and-cut" can be utilized to solve the round lot optimization problem to optimality. As the problem size increases, however, these algorithms are usually unsuccessful in generating a feasible solution within a reasonable amount of time.

---

[1] Round lots in this paper refer to the trades instead of positions. When the initial portfolio is all cash or consists of all round lot positions, round lot trades will also lead to round lot positions in the final portfolio.

[2] Optimal Roundlotting refers to the process of solving the Round Lot Optimization problem—one that involves round lot constraints—to optimality.

[3] In this sense, optimal roundlotting can also be described as the "constraint-aware" roundlotting.

[4] Post-Optimization Roundlotting refers to the process of rounding the solutions of any portfolio optimization up or down to their nearest round lots.

How is Optimal Roundlotting Done?

In practice, the Barra Optimizer employs a handful of heuristics to tackle the round lot optimization problem. This approach essentially trades off optimality with speed. The following subsections address the main techniques used in these heuristics.

Linearization of the Objective Function

A basic idea in optimal roundlotting is to linearize the quadratic utility function over a subset of the feasible region bounded by round lot trade positions. Doing so transforms the original quadratic programming problem into a linear programming one.

We take advantage of the linear programming (LP) property that one of the optimal solutions, if it exists, must be a basic feasible solution that consists of basic and non-basic variables. In an optimal basic solution, the non-basic variables will appear at their boundaries. In the context of portfolio optimization, the number of basic variables is often defined by the number of linear constraints (e.g., constraints on factor exposures) or piecewise linear constraints (e.g., turnover or transaction cost limit constraints) in a given problem.

On the other hand, the number of non-basic variables equals the difference between the number of names in the investment universe and the number of basic variables. Typically, the number of names is much greater than the number of linear constraints, implying that the number of non-basic variables will be significantly more than the basic variables. By linearization and solving an LP problem where the asset bounds are set to round lot trade positions,[5] the number of odd lot trades in the optimal solution of the LP will be reduced to (at most) the number of linear or piecewise linear constraints in just one step. In short, linearization offers an effective way to reduce the number of odd lots rapidly.

There are various ways of linearizing the utility function, representing different approximations of the original problem. One is to linearize it at the optimal solution to the underlying standard portfolio optimization problem.[6] In the stage of solving such a linearized problem, we temporarily reduce the feasible region to a one-lot region that includes the optimal solution (see Figure 1) in order to minimize loss of accuracy. After the temporary reduction, we always restore the original feasible region.

[5] In general, the lower and upper bounds are set to different round lot trade positions.

[6] This refers to the problem when disregarding the round lot constraints.

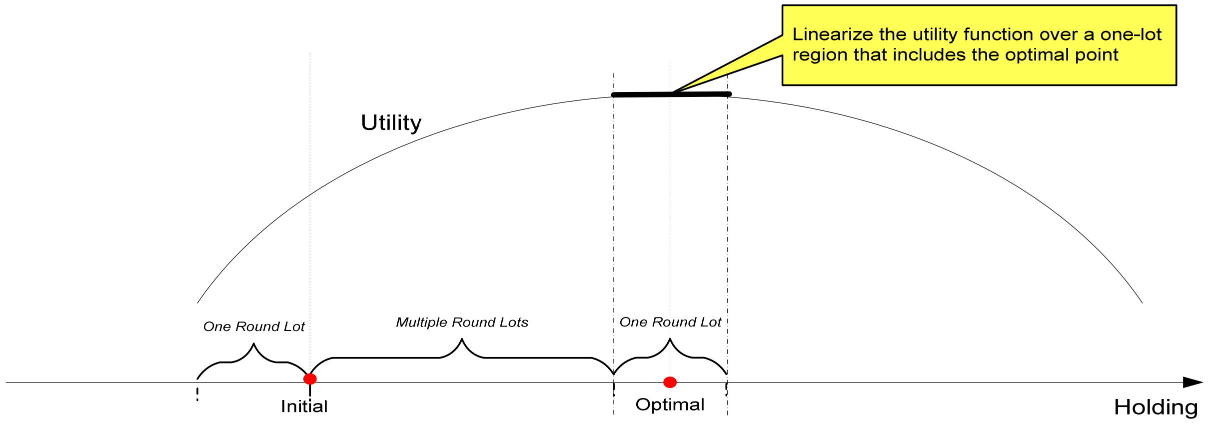

Figure 1. Linearization at the Optimal Point.

An alternative linearization strategy is to divide the original feasible region into a number of one-lot regions, and then linearize the utility function over each of these small regions (see Figure 2). This method essentially transforms the original quadratic function into a piecewise linear function. The actual feasible region can be set to include one or more of those one-lot regions.

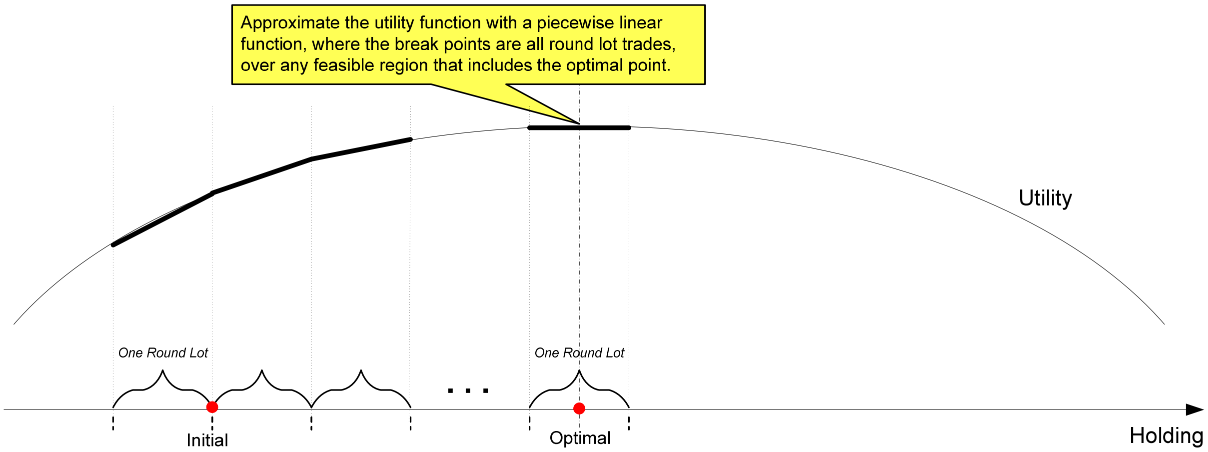

Figure 2. Piecewise Linearization over Multiple Round Lots.

Tuning the Remaining Odd Lots

Linearization alone is rarely sufficient to produce a complete round lot solution. Fine tuning techniques are often needed after linearization to eliminate the remaining odd lots. During our research, we have experimented with two fine-tuning techniques—"tuning by bound" and "tuning by penalty."

"Tuning by bound" is an iterative procedure. At the start of each iteration, we lock the assets that are already in round lot trade positions to their current positions. Then we select a small portion of the remaining odd lots and lock their positions to a nearby feasible round lot trade position. Next, we solve a standard or linearized portfolio optimization problem with these new asset bounds to obtain the "optimal" holdings for the rest of the odd lots. This process is repeated until a round lot solution is generated, or until the maximum number of tuning iterations is reached.

"Tuning by penalty" is a one-step procedure in which a quadratic penalty is imposed on any deviations away from the targeted round lot position. Often times, this target consists of the nearest round lot position to the current position for each asset. While the "tuning by bound" procedure enforces the round lot constraints by successive brute force, the "tuning by penalty" technique treats the round lot constraints as soft bounds. Our experience is that "tuning by penalty" alone is not very effective in getting a complete round lot solution, even with heavy penalties.

Ranking Criteria

In deciding which odd lot assets are locked first in the "tuning by bound" procedure, we have ranked assets by three criteria —"distance," "gradient," and "gradient*distance." These criteria provide a search direction for the optimal solution based on the current solution. The following are brief explanations of these criteria.

- "Distance" is simply the absolute difference between an asset's current position and its nearest round lot position.
- "Gradient" represents the slope of the tangent line to the utility function at the current position. When gradient is positive, the utility value will increase if we increase the asset holding level. Conversely, if the gradient is negative, then the utility will improve only if we decrease the asset holding. Hence, the sign of the gradient signals the direction of utility improvement, and the absolute value of the gradient estimates the potential of the improvement.
- "Gradient*distance" is the product of gradient and distance, which in theory should measure the potential change in utility more accurately by taking both angular and horizontal effects into consideration.

Heuristics Based on a Combination of Linearization and Tuning Techniques

We have designed nearly 40 heuristic procedures based on different combinations of the linearization and tuning techniques. Based on testing thousands of simulations and comparing approaches, we have chosen one of the piecewise linearization methods, combined with "tuning by bound" and prioritizing assets by "distance*gradient," as the main heuristic in Barra Optimizer's Optimal Roundlotting. If more than one feasible round lot solutions is generated by different heuristics when solving a single case, these will be used to compare with and update the "best available solution" in terms of the utility value. The feasible round lot solution with the highest utility value will be returned at the end.

Limitations of the Optimal Roundlotting

Due to the discrete nature of the roundlotting constraints, it may be impractical or even impossible to find a globally optimal portfolio. The heuristics discussed above do not guarantee to return a globally optimal portfolio, sometimes not even a locally optimal portfolio. However, the successfully returned solution will always satisfy all constrains including the roundlotting ones imposed on the problem. Given the difficult nature of this type of optimization problems, optimal roundlotting is valuable to portfolio managers, especially when it comes to conforming to constraints.

Post-Optimization Roundlotting is Convenient but Constraint-Blind

Post-optimization roundlotting is a convenient and computationally inexpensive alternative to optimal roundlotting. However, it pays no attention to the constraints, and is only appropriate when violating a few constraints is negligible or tolerable.

In the Barra Optimizer, post-optimization roundlotting is offered as a callable function. Users decide if and when they want to apply the function. It is a one-step mechanical procedure that can be applied to the solution of any single portfolio optimization — i.e., standard mean-variance optimization, or maximizing the Sharpe or Information Ratio, paring optimization, risk constrained optimization, long-short optimization, after-tax optimization, 5/10/40 rules, or even optimal roundlotting.[7]

While optimal roundlotting may fail to find a feasible all-round-lot solution satisfying asset bounds and all the other constraints, post-optimization roundlotting will never fail to produce a solution because it ignores all the constraints. It simply rounds the optimal trades up or down to their nearest round lots. To preserve the portfolio value, cash position will be adjusted accordingly, but no consideration is placed on minimizing constraint violations or maximizing the objective function. In particular, cash or other asset bounds may be violated.

In many cases, the constraint violations due to post-optimization roundlotting may be small, inconsequential, or negligible. Portfolio managers may be satisfied just by using the post-optimization roundlotting to obtain a round lot trade list.

Illustrative Test Cases

We illustrate the effects of optimal versus post-optimization roundlotting using two simple cases. Each case consists of three real assets plus cash, with the initial portfolio equally weighted and the upper bound on any real asset being 60 percent of the portfolio value. Both cases also share the common data as given in Table 1. For simplicity, we assume no transaction costs and no constraints other than the holding constraint and the long-only constraint.[8] Risk aversion is 0.75.

Table 1. Basic Case Data.

Asset ID	Alpha (%)	Lot Size (shares)	Price ($/share)	Asset-by-Asset Covariance Matrix (%)				Initial Weights (%)	Initial Shares	Initial Value ($)
				Cash	1	2	3			
cash	0.00	1	1	0.00	0.00	0.00	0.00	25	1,000,000	1 MM
1	13.06	100	100	0.00	17.71	1.39	3.00	25	10,000	1 MM
2	13.83	100	100	0.00	1.39	3.77	0.74	25	10,000	1 MM
3	15.30	100	100	0.00	3.00	0.74	5.59	25	10,000	1 MM

[7] When optimal roundlotting fails to produce an all-round-lot solution or only produces a partially round lot solution.

[8] The holding constraint requires that the sum of all asset weights in a portfolio be 1. The long-only constraint restricts all asset bounds to be greater or equal to 0.

We consider two cases:

- Case 1: Invest all cash (i.e., the upper and lower bounds on the cash asset are both zero).
- Case 2: Allow at most 20 percent cash (i.e., the upper bound on the cash asset is 20 percent of the portfolio value, while the lower bound is zero).

We compare the results from optimal roundlotting, post-optimization roundlotting, and standard optimization (i.e., without roundlotting).

As shown in Table 2 and Table 3, the solutions of both the optimal and post-optimization roundlotting are always in round lot trades, while the solutions of the standard optimization are in odd lots. For Case 1, the solution of the post-optimization roundlotting violates the cash lower bound. For Case 2, the solutions of both the optimal and post-optimization roundlotting satisfy all constraints, including the lower and upper bounds on cash; however, the solution of the optimal roundlotting outperforms the solution of the post-optimization roundlotting in terms of utility.

Table 2. Case 1—Invest All Cash.

Asset ID	Standard Optimization			Optimal Roundlotting			Post-Optimization Roundlotting		
	Weights (%)	Shares	Round Lots Traded	Weights (%)	Shares	Round Lots Traded	Weights (%)	Shares	Round Lots Traded
cash	0.00	-	-	0.00	-	-	-0.25	(10,000.00)	(1,010,000.00)
1	3.69	1,475.46	(85.25)	3.75	1,500.00	(85.00)	3.75	1,500.00	(85.00)
2	59.90	23,961.18	139.61	59.75	23,900.00	139.00	60.00	24,000.00	140.00
3	36.41	14,563.36	45.63	36.50	14,600.00	46.00	36.50	14,600.00	46.00
Portfolio Return (%)	14.340			14.341			14.376		
Portfolio Risk (%)	16.075			16.076			16.116		
Utility	-1.795			-1.795			-1.804		

Table 3. Case 2—Allow 20 Percent Cash.

Asset ID	Standard Optimization			Optimal Roundlotting			Post-Optimization Roundlotting		
	Weights (%)	Shares	Round Lots Traded	Weights (%)	Shares	Round Lots Traded	Weights (%)	Shares	Round Lots Traded
cash	20.00	800000.00	(200,000.00)	20.00	800,000.00	(200,000.00)	19.75	790000.00	(210,000.00)
1	2.94	1174.90	(88.25)	2.75	1,100.00	(89.00)	3.00	1200.00	(88.00)
2	47.90	19161.28	91.61	48.00	19,200.00	92.00	48.00	19200.00	92.00
3	29.16	11663.83	16.64	29.25	11,700.00	17.00	29.25	11700.00	17.00
Portfolio Return (%)	11.473			11.475			11.508		
Portfolio Risk (%)	12.860			12.861			12.901		
Utility	-1.126			-1.126			-1.133		

Conclusion

Both optimal and post-optimization roundlotting are useful tools in the Barra Optimizer for portfolio construction. The key difference is that optimal roundlotting is constraint-aware, while post-optimization roundlotting is constraint-blind. Tradeoffs should be considered when evaluating which one to use.

Optimal roundlotting has its limitations, but when successful, it satisfies all user-constraints and provides better or equivalent utility compared to post-optimization roundlotting. Post-optimization roundlotting is simple, fast, widely applicable, and may be used when violations to constraints can be tolerated.

Acknowledgements

The authors would like to thank Jay Dermody, Stacy Cuffe, and Martin Soboszek for their valuable comments and feedback on the initial drafts of this paper.

References

Markowitz, Harry M. (1952) "Portfolio Selection." *Journal of Finance,* Vol. 7, No. 1, pp. 77–91.

Client Service Information is Available 24 Hours a Day

clientservice@msci.com

Americas

Americas	1.888.588.4567 (toll free)
Atlanta	+ 1.404.551.3212
Boston	+ 1.617.532.0920
Chicago	+ 1.312.675.0545
Montreal	+ 1.514.847.7506
Monterrey	+ 52.81.1253.4020
New York	+ 1.212.804.3901
San Francisco	+ 1.415.836.8800
Sao Paulo	+ 55.11.3706.1360
Stamford	+1.203.325.5630
Toronto	+ 1.416.628.1007

Europe, Middle East & Africa

Cape Town	+ 27.21.673.0100
Frankfurt	+ 49.69.133.859.00
Geneva	+ 41.22.817.9777
London	+ 44.20.7618.2222
Milan	+ 39.02.5849.0415
Paris	0800.91.59.17 (toll free)

Asia Pacific

China North	10800.852.1032 (toll free)
China South	10800.152.1032 (toll free)
Hong Kong	+ 852.2844.9333
Seoul	00798.8521.3392 (toll free)
Singapore	800.852.3749 (toll free)
Sydney	+ 61.2.9033.9333
Tokyo	+ 81.3.5226.8222

About MSCI

MSCI Inc. is a leading provider of investment decision support tools to investors globally, including asset managers, banks, hedge funds and pension funds. MSCI products and services include indices, portfolio risk and performance analytics, and governance tools.

The company's flagship product offerings are: the MSCI indices with close to USD 7 trillion estimated to be benchmarked to them on a worldwide basis[1]; Barra multi-asset class factor models, portfolio risk and performance analytics; RiskMetrics multi-asset class market and credit risk analytics; IPD real estate information, indices and analytics; MSCI ESG (environmental, social and governance) Research screening, analysis and ratings; ISS governance research and outsourced proxy voting and reporting services; and FEA valuation models and risk management software for the energy and commodities markets. MSCI is headquartered in New York, with research and commercial offices around the world.

[1] *As of September 30, 2012, as published by eVestment, Lipper and Bloomberg on January 31, 2013* *Apr 2013*