European Journal of Operational Research 236 (2014) 160–176

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Stochastics and Statistics

Can long-run dynamic optimal strategies outperform fixed-mix portfolios? Evidence from multiple data sets

Daniele Bianchi [a,1], Massimo Guidolin [b,*]

[a] Department of Finance, Bocconi University, Via Roentgen 1, 20136 Milan, Italy
[b] Department of Finance and IGIER, Bocconi University, Via Roentgen 1, 20136 Milan, Italy

ARTICLE INFO

Article history:
Received 13 August 2012
Accepted 13 January 2014
Available online 29 January 2014

Keywords:
Finance
Investment analysis
Portfolio choice
Predictability

ABSTRACT

Using five alternative data sets and a range of specifications concerning the underlying linear predictability models, we study whether long-run dynamic optimizing portfolio strategies may actually outperform simpler benchmarks in out-of-sample tests. The dynamic portfolio problems are solved using a combination of dynamic programming and Monte Carlo methods. The benchmarks are represented by two typical fixed mix strategies: the celebrated equally-weighted portfolio and a myopic, Markowitz-style strategy that fails to account for any predictability in asset returns. Within a framework in which the investor maximizes expected HARA (constant relative risk aversion) utility in a frictionless market, our key finding is that there are enormous difference in optimal long-horizon (in-sample) weights between the mean–variance benchmark and the optimal dynamic weights. In out-of-sample comparisons, there is however no clear-cut, systematic, evidence that long-horizon dynamic strategies outperform naively diversified portfolios.

1. Introduction

One of the most important problems faced by investors involves the allocation of wealth among risky assets. The task of determining optimal portfolios is a complex problem that depends on the objective of the investor, her horizon, any constraints, as well as the dynamic properties (forecasts) of the set of investment opportunities. In this setting, the objective of the investor has typically (but not exclusively, see e.g., Liu, 1999) been the maximization of the expected value of a utility function depending on either terminal wealth or, more realistically, on the flow of consumption that a portfolio may finance over time.

A rich literature straddling the empirical finance and operations research fields (see Section 2) has investigated applications of stochastic programming approaches to dynamic portfolio problems. Such methods become of particular interest when asset returns are predictable, i.e., when the investment opportunity set is time-varying, as in Barberis (2000) or Campbell, Chan, and Viceira (2003). Using a number of recursive, out-of-sample (OOS) experiments applied to five different data sets, in our paper we document whether, how, and when long-term strategies based on dynamic optimal portfolios (DOPs) that exploit predictability may outperform simple—one may argue, naive—*fixed mix strategies* in which an investor may either ignore predictability, holding *constant weights* over her investment horizon, or even disregard the properties of asset returns altogether and assign *equal weights* (EW) to each of the assets in a menu. In our paper we capture a wide range of linear predictability patterns from past asset returns and a set of predictors that the finance literature has shown to forecast economic conditions (such as the dividend yield, the riskless term spread, and the default spread between Baa- and Aaa-rated bonds, see Fama & French, 1989).

In the empirical finance literature, it is has been long known that portfolios selected according to the mean–variance (MV) criterion are unlikely to perform as well as equally weighted portfolios in OOS tests (see e.g., DeMiguel, Garlappi, & Uppal, 2009; Duchin & Levy, 2009). In this paper, we recast the key question in this literature—i.e., "If we have at least some information on the expected returns, riskiness, and diversification properties of the assets, limited though it may be, why should we not expect optimization to improve on a naively diversified portfolio?" (see Kritzman, Page, & Turkington, 2010, p. 31)—with reference not to simple MV strategies, but to stochastic programming approaches to long-horizon DOPs that exploit predictability. Notice that from an ex-ante perspective, the odds that the answer to this question may be positive are not easily guessed. On the one hand, full DOPs are based on large-scale econometric models that are exposed to both the risks of model misspecification and to the problems caused by estimation uncertainty concerning a large number of parameters.

* Corresponding author. Tel.: +39 02 5836 3463.
E-mail addresses: daniele.bianchi@unibocconi.it (D. Bianchi), massimo.guidolin@unibocconi.it (M. Guidolin).
[1] Tel.: +39 02 5836 3354.

http://dx.doi.org/10.1016/j.ejor.2014.01.030
0377-2217/

This would make one think that DOPs may not stand a chance to out-perform the naive fixed mix EW strategy. On the other hand, if investment opportunities are truly time-varying and statistically predictable, then the entire literature that has claimed the failure of optimizing models over EW may simply derive from the fact that important features of the actual set of investment opportunities have been ignored, so that the Markowitz-style MV strategies would be biased for the true but unknown portfolio weights. Our paper sheds light on these important issues.

As in Moreno and Olmeda (2007), we follow a quite simple but extremely computationally intensive approach: we build a variety of linear models estimated using several alternative information sets and use them to obtain forecasts of a significant number of return series, along an extensive time span. Subsequently we employ test procedures to detect any differences among models, information sets as well as specifications. Within a framework in which the investor maximizes expected HARA (constant relative risk aversion) utility in a frictionless market, our key finding is that there are enormous differences in optimal long-horizon (in-sample) weights between the Markowitz MV benchmark and DOPs. While all fixed mix strategies imply zero hedging demands by construction, in the case of long-run DOPs, the dividend yield causes rather large hedging demands to be optimal in most applications. In OOS comparisons, there is however no clear-cut, *systematic* (i.e., common to all the data sets investigated) evidence that long-horizon (60-month) dynamic strategies that take linear predictability into consideration may outperform—according to a simple Sharpe ratio metric—naive strategies such as a myopic MV and EW portfolios.

However, Sharpe ratios are, by construction, suitable risk-adjusted measures only for mean–variance investors and/or in the presence of normally distributed returns. Therefore the Certainty Equivalent Return (CER) represents, a more reliable OOS metric to rank competing strategies. Using a CER metric, we obtain a slightly more negative view of the actual chances of dynamic strategies that exploit linear predictability to outperform the benchmarks. Only in one of our five applications, there was evidence of DOP that exploit predictability within VAR frameworks producing useful portfolio signals, in terms of realized utilities. In three other applications it was instead the fixed mix strategy that outperformed all other models. However, only for one data set, the DOPs were systematically ranked below all benchmarks.

The structure of the paper is as follows. In Section 2 we review the literature, also to emphasize our contributions. In Section 3 we describe our research design and the structure of the statistical models that capture predictability, the dynamic portfolio problem, the solution methods, and the criteria used to test whether DOPs outperform the benchmarks. Section 4 describes the asset allocation strategies, with special emphasis on the conditions under which DOPs simplify to myopic, Markowitz-style ones. Section 5 describes our data sets. Section 6 compares ex-ante, recursive in-sample portfolio weights and hedging demands. Section 7 is the core of the paper as it documents realized OOS performances. Section 8 discusses our key results, a few additional robustness tests, and concludes.

2. Literature review

A rich literature straddling the empirical finance (see e.g., Campbell & Viceira, 2002, and references therein) and operations research fields (see, among many others, Barro & Canestrelli, 2005; Buckley, Saunders, & Seco, 2008; Çanakoğlu & Özekici, 2010; Consigli & Dempster, 1998; Dantzig & Infanger, 1993) has investigated the applications of stochastic programming approaches to dynamic portfolio problems, with recent forays into applications in which asset returns are forecastable and the investment opportunity set is time-varying. In particular, finance scholars have long established a number of unsettling and thought-provoking results concerning the comparison of the realized OOS performance of myopic optimizing portfolios—often of a simple MV, Markowitz-style type—vs. simple, fixed mix strategies, such as the equally weighted (EW, also called $1/N$) one. For instance, since Frankfurter, Phillips, and Seagle (1971) it has been known that portfolios selected according to a MV criterion are unlikely to be as effective as EW. DeMiguel et al. (2009) and Duchin and Levy (2009) have recently reported evidence on this point with reference to a large array of alternative data sets, and provided intuition on what the determinants of this troubling result may be.

Kritzman et al. (2010) have re-examined this evidence and argued that poor optimization results may arise from reliance on rolling 60- and 120-month historical windows fed in the estimation algorithms supporting asset allocation. Without assuming the existence of any predictability, but using 13 alternative data sets, they find that in a 1-month ahead portfolio framework, investors may benefit from optimizing strategies provided they use sufficiently long-time series to estimate the MV inputs. We extend these systematic OOS comparisons to examine whether DOPs may outperform MV and EW strategies, especially over long horizons. Because under predictability, a full DOP may imply a component of portfolio weights that represents a *hedging demand* that provides investors with (self-) insurance against adverse dynamics in investment opportunities, such a comparison boils down to asking whether hedging may actually improve ex-post, realized performance.

The problem of comparing the realized performance of alternative portfolio strategies requiring different degrees of computational effort is of course not new to operational researchers. For instance, Fleten, Høyland, and Wallace (2002) have examined the relative performance of dynamic stochastic vs. naive, fixed mix approaches in a portfolio (strategic asset and liability) framework, finding that the former approach dominates a fixed mix approach, but that the degree of outperformance is much smaller when the models are compared in OOS tests than when they are assessed in-sample. They attribute this modest differences in OOS performance to the fact that in a forecasting context, it is often the case that the random input data may be structurally different from the in-sample scenarios, i.e., optimizing methods may incorrectly but forcefully rely on the assumption that the underlying data generating process is correctly specified. Interestingly, their fixed mix approach encompasses both benchmark strategies used in our paper.[2] Fleten et al. (2002) correctly emphasize the need for the comparison between optimizing and fixed mix portfolio strategies to be "(...) done in a fair and realistic way. Realistic here means a situation that is close to the practical use of the models, where the models are rerun at regular intervals." Although highly realistic, their focus is on a multi-period, stochastic asset liability management model developed for a Norwegian life insurance company. We devote our attention, instead, to a large family of long-run DOPs that have been popular in the finance literature at least since the late 1990s (see, e.g., Campbell & Viceira, 2002) and that exploit rich patterns of linear predictability relating asset returns to well-known state variables.

Moreno and Olmeda (2007) have compared the realized OOS forecasting performance of a range of both linear (including vector autoregressive) and nonlinear (artificial neural networks) models

[2] In Fleten et al. (2002), fixed mix models consists of decision rules in which at every period, the portfolio is rebalanced to fixed proportions. In our paper we focus on EW strategies as these have been recently rediscovered by the empirical finance literature because of their strong realized OOS performance.