

CONTEXTUAL AREAS

Technical Note—A Robust Perspective on Transaction Costs in Portfolio Optimization

 Alba V. Olivares-Nadal,^a Victor DeMiguel^b
^a Statistics and Operations Research, University Pablo de Olavide, 41013 Sevilla, Spain; ^b Management Science and Operations, London Business School, London NW1 4SA, United Kingdom

Contact: avolinad@upo.es,  <http://orcid.org/0000-0003-0004-2262> (AVO-N); avmiguel@london.edu,

 <http://orcid.org/0000-0002-7787-5185> (VD)

Received: August 11, 2016

Revised: March 23, 2017; July 11, 2017; September 4, 2017; September 20, 2017

Accepted: October 5, 2017

Published Online in Articles in Advance: May 22, 2018

Subject Classifications: finance: portfolio; programming: nonlinear, theory, duality

Area of Review: Financial Engineering

<https://doi.org/10.1287/opre.2017.1699>
Copyright: © 2018 INFORMS

Abstract. We prove that the portfolio problem with transaction costs is equivalent to three different problems designed to alleviate the impact of estimation error: a robust portfolio optimization problem, a regularized regression problem, and a Bayesian portfolio problem. Motivated by these results, we propose a data-driven approach to portfolio optimization that tackles transaction costs and estimation error simultaneously by treating the transaction costs as a regularization term to be calibrated. Our empirical results demonstrate that the data-driven portfolios perform favorably because they strike an optimal trade-off between rebalancing the portfolio to capture the information in recent historical return data and avoiding the large transaction costs and impact of estimation error associated with excessive trading.

Funding: The first author is supported in part by the Junta de Andalucía [Projects P11-FQM-7603 and FQM329] and the Ministerio de Economía y Competitividad, Spain [MTM2015-65915-R].

Supplemental Material: The online companion is available at <https://doi.org/10.1287/opre.2017.1699>.

Keywords: transaction costs • estimation error • robust optimization

1. Introduction

Markowitz (1952) showed that an investor who cares only about the mean and variance of portfolio returns should choose a portfolio on the efficient frontier. Mean-variance portfolios continue to be the workhorse of much of the investment management industry, but two crucial aspects in their successful implementation are estimation error and transaction costs. Estimation error is important because to implement the mean-variance portfolios in practice, one has to estimate the mean and the covariance matrix of asset returns, and it is well known that the resulting portfolios often perform poorly out of sample; see Michaud (1989), Chopra and Ziemba (1993), and DeMiguel et al. (2009a). A popular approach to alleviate the impact of estimation error is to use robust portfolio optimization; see, for instance, Goldfarb and Iyengar (2003), Garlappi et al. (2007), and Lu (2011a, b). This approach captures the uncertainty about the mean and covariance matrix of asset returns by assuming they may lie anywhere inside the so-called uncertainty sets. The robust portfolio is the one that maximizes the mean-variance utility with respect to the worst-case mean and covariance matrix of asset returns.

Transaction costs are important because they can easily erode the gains from a trading strategy. Transaction costs can be generally modelled with the p -norm of the portfolio trade vector. For small trades, which

do not impact the market price, the transaction cost is generally assumed to be proportional to the amount traded, and thus it can be approximated by the 1-norm of the portfolio trade vector. For larger trades, the literature has traditionally assumed that they have a linear impact on the market price, and thus they result in quadratic transaction costs that are captured by the 2-norm. Finally, several authors have recently argued that market impact costs grow as the square root of the amount traded (Almgren et al. 2005, Frazzini et al. 2015), and thus they are captured by the p -norm with $p = 1.5$.

We make a theoretical contribution and an empirical contribution. From a theoretical perspective, we show that the portfolio optimization problem with p -norm transaction costs can be equivalently reformulated as three different problems designed to alleviate the impact of estimation error: (i) a robust portfolio optimization problem, (ii) a regularized regression problem, and (iii) a Bayesian portfolio problem. These results demonstrate that incorporating a p -norm transaction cost term in the mean-variance portfolio problem may help to reduce the impact of estimation error. This observation motivates our empirical contribution: we propose a data-driven approach to portfolio selection that consists of using cross-validation to calibrate the transaction cost parameter from historical data and compute portfolios that perform well in a realistic scenario with both estimation error and transaction

costs. From a real-world perspective, combating estimation error by calibrating the transaction cost term has two advantages over the use of robust approaches. First, practitioners are used to incorporating transaction costs in their portfolio selection frameworks, and thus it may be easier for them to simply calibrate the transaction cost term instead of using a more sophisticated approach based on uncertainty sets. Second, a transaction cost term has a natural economic interpretation, and this facilitates the task of selecting a reasonable range of parameters to calibrate from.

We compare the out-of-sample performance of the proposed data-driven portfolios on five empirical data sets with that of the mean-variance portfolios that ignore transaction costs as well as mean-variance portfolios that capture the nominal proportional transaction costs. We find that the proposed data-driven portfolios outperform the traditional portfolios in terms of their out-of-sample Sharpe ratio net of transaction costs. The data-driven portfolios perform well because they calibrate the transaction cost parameter to achieve intermediate levels of turnover that strike an optimal (data-driven) trade-off between two goals: (i) rebalancing the portfolio to capture the information in recent historical return data and (ii) avoiding the large transaction costs and impact of estimation error associated with excessive trading.

The remainder of this paper is organized as follows. Section 2 gives our theoretical results connecting transaction costs and robustness. Section 3 describes the proposed data-driven approach and evaluates its out-of-sample performance. The online companion contains two appendices: Appendix A presents the proofs of all results, and Appendix B presents supplementary tables and figures.

2. Transaction Costs and Robustness

We consider the following mean-variance problem with p -norm transaction costs:

$$\begin{aligned} \min_{\mathbf{w}} \quad & \left\{ \frac{1}{2} \gamma \mathbf{w}^T \Sigma \mathbf{w} - \boldsymbol{\mu}^T \mathbf{w} + \kappa \|\Lambda(\mathbf{w} - \mathbf{w}_0)\|_p^p \right\} \\ \text{s.t.} \quad & \mathbf{w}^T \mathbf{1}_N = 1, \end{aligned} \quad (1)$$

where $\gamma \in \mathbb{R}$ is the risk-aversion parameter; $\mathbf{w} \in \mathbb{R}^N$ is the portfolio weight vector; $\Sigma \in \mathbb{R}^{N \times N}$ is the estimated covariance matrix of asset returns; $\boldsymbol{\mu} \in \mathbb{R}^N$ is the estimated mean of asset returns; $\kappa \in \mathbb{R}$ is the transaction cost parameter; $\Lambda \in \mathbb{R}^{N \times N}$ is the transaction cost matrix, which we assume to be nonsingular; $\mathbf{w}_0 \in \mathbb{R}^N$ is the starting portfolio; $\|\mathbf{s}\|_p$ is the p -norm of vector \mathbf{s} ; $\|\mathbf{s}\|_p^p = \sum_{i=1}^N |s_i|^p$, $\mathbf{1}_N \in \mathbb{R}^N$ is the vector of ones; and the constraint $\mathbf{w}^T \mathbf{1}_N = 1$ requires that the portfolio weights sum up to 1.

The first two terms in the objective function capture the risk–return trade-off: the first term is the portfolio

return variance scaled by the risk-aversion parameter $(\gamma/2)\mathbf{w}^T \Sigma \mathbf{w}$, and the second term is the portfolio return mean $(\boldsymbol{\mu}^T \mathbf{w})$. More important for our purposes, the third term in the objective function is the p -norm transaction cost term, $\kappa \|\Lambda(\mathbf{w} - \mathbf{w}_0)\|_p^p$. Note that we allow for the portfolio trade vector $\mathbf{w} - \mathbf{w}_0$ to be transformed via a nonsingular transaction cost matrix Λ before computing the p -norm.

We now introduce some basic definitions and facts. For a given vector norm $\|\cdot\|$, its dual norm $\|\cdot\|^*$ is $\|\mathbf{x}\|^* = \max_{\|\mathbf{y}\| \leq 1} \mathbf{y}^T \mathbf{x}$. It is easy to show that the dual norm of the p -norm is the q -norm, where $1/p + 1/q = 1$; see Higham (2002, section 6.1). Let Λ be a nonsingular matrix; we define the (p, Λ) -norm of vector \mathbf{x} as $\|\mathbf{x}\|_{p, \Lambda} := \|\Lambda \mathbf{x}\|_p$. It is also easy to show that $\|\cdot\|_{p, \Lambda}$ is indeed a vector norm and $\|\cdot\|_{q, \Lambda^{-T}}$ is its dual norm.

The following proposition gives our main results.¹

Proposition 1. For every risk-aversion parameter $\gamma > 0$ and transaction cost parameter $\kappa \geq 0$, there exist $\delta, \kappa' \geq 0, \alpha > 0$, and μ_0 such that the mean-variance problem with p -norm transaction costs, problem (1), is equivalent to

(i) a robust portfolio problem:

$$\begin{aligned} \min_{\mathbf{w}} \quad & \left\{ \frac{\gamma}{2} \mathbf{w}^T \Sigma \mathbf{w} - \boldsymbol{\mu}^T \mathbf{w} + \max_{\hat{\boldsymbol{\mu}} \in U(\delta)} (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}})^T (\mathbf{w} - \mathbf{w}_0) \right\} \\ \text{s.t.} \quad & \mathbf{w}^T \mathbf{1}_N = 1, \end{aligned} \quad (2)$$

where the uncertainty set for mean asset returns is $U(\delta) = \{\hat{\boldsymbol{\mu}}: \|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}\|_{q, \Lambda^{-T}} \leq \delta\}$;

(ii) a regularized linear regression problem:

$$\begin{aligned} \min_{\mathbf{w}} \quad & \left\{ \|\mathbf{1}_T - R\mathbf{w}\|_2^2 + \kappa' \|\Lambda(\mathbf{w} - \mathbf{w}_0)\|_p^p \right\} \\ \text{s.t.} \quad & \mathbf{w}^T \boldsymbol{\mu} = \mu_0, \\ & \mathbf{w}^T \mathbf{1}_N = 1, \end{aligned} \quad (3)$$

where $R \in \mathbb{R}^{T \times N}$ is the matrix whose columns contain the historical returns for each of the N assets; or

(iii) a Bayesian portfolio problem, where the investor believes a priori that the variance of the mean-variance portfolio return has an independent distribution $\pi(\sigma^2)$, that asset returns are normally distributed, and that the mean-variance portfolio weights are jointly distributed as a multivariate exponential power distribution, with probability density function

$$\pi(\mathbf{w}) = \frac{p^N |\Lambda|}{2^N \alpha^N \Gamma(1/p)^N} e^{-\|\Lambda(\mathbf{w} - \mathbf{w}_0)\|_p^p / \alpha^p}, \quad (4)$$

where α is the scale parameter and $\Gamma(\cdot)$ is the gamma function.

A few comments are in order. Proposition 1(i) shows that the portfolio optimization problem with p -norm transaction costs is equivalent to a robust portfolio optimization problem where the mean of asset returns can take any value in an uncertainty set defined by the

q -norm, where $1/p + 1/q = 1$. For instance, a mean-variance portfolio problem with proportional transaction costs can be equivalently rewritten as a robust portfolio optimization problem where the mean can take any value in an uncertainty set defined by a box centered at the nominal mean return. This theoretical result provides theoretical justification for the well-known empirical observation that robust portfolio policies often result in low turnover; see Fabozzi et al. (2007). In essence, solving a robust portfolio optimization problem is equivalent to introducing a transaction cost on any trades. Therefore, in addition to alleviating the impact of estimation error, robustifying a portfolio optimization problem is also likely to reduce portfolio turnover.

Proposition 1(i) also provides compelling statistical motivation for the use of *quadratic* transaction cost terms to combat estimation error. To see this, note that the portfolio problem with quadratic transaction costs for the case with $\Lambda = \Sigma^{1/2}$, which Gârleanu and Pedersen (2013) argue is realistic, is equivalent to a robust portfolio problem with ellipsoidal uncertainty set for mean returns given by $\|\mu - \hat{\mu}\|_{2, \Sigma^{-1/2}} \leq \delta$. It is reassuring that this ellipsoidal uncertainty set actually coincides with the statistical confidence region for the sample estimator of mean returns under the assumption that returns are independent and identically distributed as a Normal distribution with covariance matrix Σ ; see Goldfarb and Iyengar (2003).

In addition, Proposition 1(i) shows that the p -norm transaction cost can be interpreted as the maximum *regret* the investor may experience (in terms of expected return) by trading from the starting portfolio \mathbf{w}_0 to portfolio \mathbf{w} , assuming the true mean belongs to the uncertainty set $U(\delta)$, which is defined in terms of the dual q -norm. To see this, note that Proposition 1(i) essentially shows that the transaction cost term can be rewritten as

$$\begin{aligned} \kappa \|\Lambda(\mathbf{w} - \mathbf{w}_0)\|_p^p &= \kappa \|\mathbf{w} - \mathbf{w}_0\|_{p, \Lambda}^p \\ &= \max_{\hat{\mu}: \|\mu - \hat{\mu}\|_{q, \Lambda^{-T}} \leq \delta} (\mu - \hat{\mu})^T (\mathbf{w} - \mathbf{w}_0), \end{aligned}$$

where μ is the estimated mean asset return vector, and $\hat{\mu}$ is the worst-case mean asset return vector for the given portfolio \mathbf{w} .

Proposition 1(ii) shows that the portfolio optimization problem with p -norm transaction costs is equivalent to a robust regression formulation of the mean-variance problem. It is well known that the mean-variance portfolio optimization problem can be equivalently reformulated as a linear regression problem; see, for instance, Britten-Jones (1999). We extend this result by showing that the transaction cost term $\kappa \|\mathbf{w} - \mathbf{w}_0\|_p^p$ in a mean-variance portfolio can be interpreted as a regularization term that reduces the impact of estimation error on the linear regression. In particular, for $p = 1$

this transaction cost term resembles a *lasso* regularization term, and for $p = 2$, a *ridge* regularization term; see James et al. (2013, chapter 6) for a discussion of regularization techniques in linear regression.

Proposition 1(iii) shows that the portfolio optimization problem with p -norm transaction costs is equivalent to a Bayesian portfolio problem where the investor has a prior belief over the *portfolio weights*. This result generalizes the results by DeMiguel et al. (2009b), who provide a Bayesian interpretation for the 1-norm, 2-norm, and A -norm constrained portfolios. We extend their result to the portfolio problem with (p, Λ) -norm transaction cost, with $p \in [1, 2]$, by defining a new distribution, which we term the multivariate exponential power (MEP) distribution. The MEP prior distribution includes as particular cases the multivariate normal prior distribution for $p = 2$ and $\Lambda = \Sigma^{1/2}$, which corresponds to the quadratic transaction cost of Gârleanu and Pedersen (2013), and the Laplace prior distribution for $p = 1$ and $\Lambda = I$, where I is the identity matrix, which corresponds to proportional transaction costs.

Finally, it is easy to see that all three parts of Proposition 1 also hold for the case where there are additional constraints to the mean-variance problem with p -norm transaction costs by just adding these constraints to the robust portfolio problem, the regularized linear regression problem, and the Bayesian portfolio problem.

3. Data-Driven Portfolios

We now propose a data-driven approach to portfolio selection that consists of treating the transaction cost term as a regularization term and using cross-validation to calibrate the transaction cost parameter κ of problem (1) as if it were the penalty parameter in a regularization term.

We assume the investor faces proportional transaction costs of 50 basis points, an assumption that is consistent with the existing literature; see Balduzzi and Lynch (1999). We compare the out-of-sample performance of the data-driven portfolios with that of the portfolios that ignore transaction costs and the portfolios that consider the nominal proportional transaction costs of 50 basis points.²

We consider five empirical data sets with U.S. stock monthly return data, similar to those used in the literature; see DeMiguel et al. (2009b). Specifically, we consider four data sets downloaded from Ken French's website (http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html, accessed January 2014), covering the period from July 1963 to December 2013: the 10 industry-portfolio data set (10Ind), the 48 industry-portfolio data set (48Ind), the 6 portfolios of stock sorted by size and book-to-market (6FF), and the 25 portfolios of firms sorted by size and book-to-market (25FF). Finally, we consider a data set with returns on individual stocks downloaded

from the Center for Research in Security Prices (CRSP) database covering the period from April 1968 to April 2005. This data set is constructed as in DeMiguel et al. (2009b): in April of each year, we randomly select 25 assets among all assets in the CRSP data set for which there are return data for the previous 120 months as well as for the next 12 months. These randomly selected 25 assets become our asset universe for the next 12-month period.³

3.1. Description of the Portfolios

We compare the performance of four different types of portfolios defined in terms of how they capture transaction costs. First are the portfolios that ignore transaction costs, which are computed by solving problem (1) with transaction cost parameter $\kappa = 0$. Second are the portfolios that capture the nominal proportional transaction costs of 50 basis points, which are computed by solving problem (1) with $\kappa = 0.005$, $p = 1$, and $\Lambda = I$; that is, with a transaction cost term $0.005\|\Delta\mathbf{w}\|_1$, where $\Delta\mathbf{w} = \mathbf{w} - \mathbf{w}_0$. Finally are the portfolios with calibrated transaction costs, computed by solving problem (1) with a transaction cost (or penalty) parameter $\kappa = \kappa_{cv}$ calibrated with 10-fold cross-validation (we explain this in detail later); we consider two types of calibrated transaction costs. The first are proportional transaction costs ($p = 1$, $\Lambda = I$), which result in a transaction cost term $\kappa_{cv}\|\Delta\mathbf{w}\|_1$; the second are quadratic transaction costs ($p = 2$, $\Lambda = \Sigma^{1/2}$), which result in a transaction cost term $\kappa_{cv}\|\Sigma^{1/2}\Delta\mathbf{w}\|_2^2$. For each of these four different types of transaction costs, we compute four different portfolios: minimum-variance portfolio, shortsale-constrained minimum-variance portfolio, mean-variance portfolio, and shortsale-constrained mean-variance portfolio.

A few comments are in order. First, why consider data-driven portfolios with calibrated *quadratic* transaction costs when the nominal transaction costs are *proportional*? The answer is that the data-driven portfolios are designed to address not only transaction costs but also estimation error, and Section 2 argues that a quadratic transaction cost term is well suited to address estimation error.

Second, note that the penalties κ_{cv} corresponding to the data-driven portfolios with proportional versus quadratic costs are not easy to compare. Fortunately, the data-driven portfolios with proportional transaction costs can be equivalently calibrated in terms of trading volume or turnover, $\|\mathbf{w} - \mathbf{w}_0\|_1 \leq \tau$. Moreover, Kourtis (2015) shows that the optimal portfolio for a mean-variance investor with quadratic transaction costs is a convex combination of the starting portfolio and the mean-variance portfolio in the absence of transaction costs. Consequently, for these portfolios, one can also calibrate the trading volume τ instead of the transaction cost parameter κ . Summarizing, to facilitate the comparison between the two

types of data-driven portfolios (with proportional and quadratic costs), we calibrate these portfolios by selecting their trading volume.⁴

Third, we calibrate the data-driven portfolios using the bootstrap methodology of 10-fold cross-validation; see Efron and Gong (1983). Specifically, we divide the estimation window of M returns into 10 intervals of $M/10$ returns each. For j from 1 to 10, we remove the j th interval from the estimation window and use the remaining sample returns to compute the data-driven portfolio for each value of the trading volume τ from 0%, 0.5%, 1%, 2.5%, 5%, and 10%. We then evaluate the return of the resulting portfolios (net of transaction costs of 50 basis points) on the j th interval. After completing this process for each of the 10 intervals, we have the M “out-of-sample” portfolio returns for each value of τ . Finally, we compute the variance of these out-of-sample returns and select the value of τ that corresponds to the portfolio with smallest variance.⁵

3.2. Out-of-Sample Performance

We use a rolling-horizon methodology similar to that used in DeMiguel et al. (2009a) and (2009b) to compare the performance of the different portfolios. We use an estimation window of $M = 120$ monthly returns. To test the statistical significance of the differences between the out-of-sample Sharpe ratios of the different portfolios and those of the minimum-variance portfolio that ignores transaction costs, we use the bootstrap methodology employed by DeMiguel et al. (2009b), which is based on the work by Ledoit and Wolf (2008).

Table 1 reports the out-of-sample Sharpe ratio for each of the 16 portfolio policies considered. Panel A reports the Sharpe ratios for the portfolios that ignore transaction costs. This panel shows that minimum-variance portfolios generally outperform mean-variance portfolios. This is explained by the well-known difficulties associated with estimating mean returns from historical data. Imposing shortsale constraints on the minimum-variance portfolio helps for only two of the data sets with the largest number of assets (48Ind and CRSP). This makes sense, as estimating the covariance matrix of asset returns is harder for data sets with many assets, and under these circumstances, the shortsale constraints will help to alleviate the impact of estimation error. Imposing shortsale constraints on the mean-variance portfolio helps for every data set because the unconstrained mean-variance portfolio is very sensitive to estimation error.⁶

Panel B shows that capturing nominal proportional transaction costs of 50 basis points generally helps to improve the performance of the mean-variance portfolios, but it helps to improve only the minimum-variance portfolio for the CRSP data set in the shortsale-unconstrained case and the 25FF in the shortsale-constrained case. The reason for this is that the nominal

Table 1. Sharpe Ratios

Strategy	10Ind	48Ind	6FF	25FF	CRSP
Panel A: Portfolios that ignore transaction costs, $\kappa = 0$					
Minimum-variance, shortsale unconstrained	0.3007	0.1167	0.3480	0.3124	0.3781
Minimum-variance, shortsale constrained	0.2953	0.2452*	0.2493*	0.2390**	0.3974
Mean-variance, shortsale unconstrained	0.0686*	-0.0890*	0.2142*	-0.0076*	-0.0091*
Mean-variance, shortsale constrained	0.2128**	0.1782	0.2502*	0.2382***	0.2194**
Panel B: Portfolios with nominal transaction costs, $0.005\ \Delta\mathbf{w}\ _1$					
Minimum-variance, shortsale unconstrained	0.2959	0.0955	0.3026**	0.3063	0.3987
Minimum-variance, shortsale constrained	0.2420***	0.2601*	0.2374*	0.2318**	0.3925
Mean-variance, shortsale unconstrained	0.2074**	-0.0523*	0.2631**	-0.0467*	0.1005*
Mean-variance, shortsale constrained	0.2214**	0.2018**	0.2505*	0.2588	0.2681
Panel C: Data-driven portfolios with calibrated proportional transaction costs, $\kappa_{cv}\ \Delta\mathbf{w}\ _1$					
Minimum-variance, shortsale unconstrained	0.3281***	0.1505	0.3284	0.3745**	0.3977
Minimum-variance, shortsale constrained	0.3006	0.2925*	0.2479*	0.2563***	0.3929
Mean-variance, shortsale unconstrained	0.2443	0.0039**	0.2436**	-0.0442*	0.1161*
Mean-variance, shortsale constrained	0.2693	0.2248**	0.2477*	0.2442***	0.2613***
Panel D: Data-driven portfolios with calibrated quadratic transaction costs, $\kappa_{cv}\ \Sigma^{1/2}\Delta\mathbf{w}\ _2^2$					
Minimum-variance, shortsale unconstrained	0.3234***	0.2349*	0.3481	0.3761*	0.3995
Minimum-variance, shortsale constrained	0.2983	0.2762*	0.2446*	0.2460***	0.3930
Mean-variance, shortsale unconstrained	0.2565	0.0105***	0.2424**	0.0464*	0.1193*
Mean-variance, shortsale constrained	0.2748	0.2561*	0.2497*	0.2514***	0.2579***

Notes. This table reports the monthly out-of-sample Sharpe ratio and the corresponding p -value that the Sharpe ratio for each of the portfolios is different from that for the minimum-variance portfolio. The highest Sharpe ratio for each data set is highlighted in boldface.

* p -value < 0.1; ** p -value < 0.05; *** p -value < 0.01.

transaction cost term helps to combat estimation error to a certain extent, and this is helpful for mean-variance portfolios, which are very sensitive to estimation error. Minimum-variance portfolios, on the other hand, are more resilient to estimation error, and a nominal transaction cost parameter is too conservative to strike the right balance between estimation error and transaction costs. This seems to indicate that using a data-driven approach to calibrate the transaction cost (or penalty) parameter may help to improve the performance.

We conclude that the shortsale-unconstrained minimum-variance portfolio that ignores transaction costs is the best of the portfolios in panels A and B, and thus we consider it the *benchmark portfolio* for the data-driven portfolios in panels C and D.

Panel C shows that the data-driven approach based on *proportional* transaction costs generally helps to improve the performance of the traditional portfolios. Specifically, the data-driven shortsale-unconstrained minimum-variance portfolio outperforms the benchmark portfolio for every data set except 6FF, with an improvement in the Sharpe ratio that ranges from 5% to 29% for the different data sets, and is statistically significant for two of the five data sets.

Panel D shows that the data-driven approach based on quadratic transaction costs also helps to improve the performance of the traditional portfolios. Specifically, the data-driven shortsale-unconstrained minimum-variance portfolio with calibrated quadratic transaction costs outperforms the benchmark portfolio for

every data set except 6FF, with an improvement in Sharpe ratio that ranges from 5% to 101%, and is statistically significant for three out of five data sets.⁷

Finally, comparing panels C and D, we observe that the data-driven minimum-variance portfolios based on quadratic transaction costs outperform the data-driven minimum-variance portfolios based on proportional transaction costs on every data set except 10Ind. This result suggests that a quadratic cost term may be better suited to address estimation error than a proportional transaction cost.

Table 2 reports the turnovers for the different portfolio policies. The table shows that taking transaction costs into account helps to reduce turnover. For instance, for the 25FF data set, the monthly turnover of the shortsale-unconstrained minimum-variance portfolio is 80.84%. Including a calibrated quadratic transaction cost term reduces this turnover to 6.44%; including a calibrated proportional transaction cost term, to 1.54%; and including a nominal proportional transaction cost term, to 0.08%.

Comparing the turnover of the portfolios with nominal proportional transaction costs (panel B) with that of the data-driven portfolios with calibrated proportional transaction costs (panel C), we observe that the nominal transaction cost term induces an *all-or-nothing* trading pattern, whereas the data-driven portfolios are associated with intermediate levels of turnovers. For instance, the unconstrained mean-variance portfolio with nominal costs has huge turnovers of 528%

Table 2. Turnover

Strategy	10Ind	48Ind	6FF	25FF	CRSP
Panel A: Portfolios that ignore transaction costs, $\kappa = 0$					
Minimum-variance, shortsale unconstrained	0.1566	0.7881	0.2310	0.8084	0.1801
Minimum-variance, shortsale constrained	0.0424	0.1157	0.0492	0.0915	0.0773
Mean-variance, shortsale unconstrained	1.0660	64.7887	1.8394	11.7604	1.5782
Mean-variance, shortsale constrained	0.1507	0.2434	0.1142	0.2447	0.1770
Panel B: Portfolios with nominal transaction costs, $0.005\ \Delta\mathbf{w}\ _1$					
Minimum-variance, shortsale unconstrained	0.0000	0.0401	0.0000	0.0008	0.0002
Minimum-variance, shortsale constrained	0.0003	0.0027	0.0000	0.0007	0.0007
Mean-variance, shortsale unconstrained	0.0118	16.3364	0.0319	5.2762	0.0967
Mean-variance, shortsale constrained	0.0045	0.0112	0.0000	0.0001	0.0032
Panel C: Data-driven portfolios with calibrated proportional transaction costs, $\kappa_{cv}\ \Delta\mathbf{w}\ _1$					
Minimum-variance, shortsale unconstrained	0.0117	0.0213	0.0155	0.0154	0.0015
Minimum-variance, shortsale constrained	0.0091	0.0090	0.0118	0.0082	0.0017
Mean-variance, shortsale unconstrained	0.0117	0.0716	0.0241	0.0217	0.0271
Mean-variance, shortsale constrained	0.0051	0.0098	0.0086	0.0124	0.0120
Panel D: Data-driven portfolios with calibrated quadratic transaction costs, $\kappa_{cv}\ \Sigma^{1/2}\Delta\mathbf{w}\ _2^2$					
Minimum-variance, shortsale unconstrained	0.0203	0.0591	0.0314	0.0644	0.0022
Minimum-variance, shortsale constrained	0.0103	0.0144	0.0127	0.0113	0.0006
Mean-variance, shortsale unconstrained	0.0176	0.0905	0.0467	0.0666	0.0288
Mean-variance, shortsale constrained	0.0076	0.0063	0.0110	0.0138	0.0186

Notes. This table reports the monthly turnover of the different portfolios. Turnover is the average percentage of wealth traded in each period and is equal to the sum of the absolute value of the rebalancing trades across the N available assets and over the $T - M - 1$ trading dates, normalized by the total number of trading dates.

for the 25FF data set and 1,634% for the 48Ind data set, whereas the counterpart data-driven portfolios with calibrated proportional costs have reasonable turnovers ranging between 1.17% and 7.16%. On the other hand, the unconstrained minimum-variance portfolio with nominal transaction costs is effectively a *buy-and-hold* portfolio (with almost zero turnover) for every data set except 48Ind, whereas the counterpart data-driven portfolios with calibrated proportional costs have reasonable monthly turnovers ranging between 0.15% and 2.13% for the different data sets.

The mathematical intuition behind why the nominal transaction cost term induces an all-or-nothing trading pattern is that the proportional transaction cost term is a piecewise linear term, which when combined in the objective function with the linear-quadratic mean-variance objective, results in policies that advise either large trading or no trading. This all-or-nothing trading pattern is indeed optimal in the absence of estimation error. Constantinides (1986), Davis and Norman (1990), and Muthuraman and Kumar (2006), among others, show that the optimal portfolio policy in the presence of proportional transaction costs is characterized by a no-trade region: if the portfolio is inside this region, then it is optimal not to trade, and if it is outside, then it is optimal to trade to the boundary of this region.

The Sharpe ratio results in Table 1, however, show that this all-or-nothing trading pattern leads to poor performance when, in addition to transaction costs, the investor is also facing estimation error. *All trading* leads to poor performance because the resulting portfolio

policies are too sensitive to recent historical data, which leads to large transaction costs and sensitivity to estimation error. *No trading* results in poor performance because buy-and-hold policies essentially ignore the information available in recent historical data—they stick to the portfolio weights obtained from the earliest estimation window. The data-driven portfolios, on the other hand, allow reasonable amounts of turnover that strike an optimal trade-off between incorporating the information in recent historical return data and avoiding the large transaction costs and impact of estimation error associated with large turnovers. Therefore, although one would expect that the data-driven portfolios would always result in smaller turnover compared to the portfolios that capture nominal transaction costs, our results show that from a data-driven perspective, it is optimal to calibrate the transaction cost parameter to achieve intermediate levels of turnover.

Acknowledgments

The authors thank Alberto Martin-Utrera, Francisco J. Nogales, Raman Uppal, and Gah-Yi Ban for their comments.

Endnotes

¹The result in part (i) of Proposition 1 is related to results by Gotoh and Takeda (2011) and Bertsimas et al. (2004). The proof of part (iii) uses results from Schwartz (1954).

²Note that all portfolio policies are *evaluated* in terms of their out-of-sample returns net of the nominal proportional transaction costs of 50 basis points, even though the data-driven portfolios are *computed* using a calibrated proportional or quadratic transaction cost term.

³We have also considered the cases with $N = 50$ and $N = 100$ stocks, and we find that the relative performance of the different portfolios is robust to changing the number of CRSP stocks.

⁴Kourtis' (2015) observation applies only to the unconstrained mean-variance portfolio. For computational convenience, however, we also approximate the shortsale-constrained mean-variance and the shortsale-constrained and unconstrained minimum-variance data-driven portfolios with quadratic transaction costs by taking a convex combination of the starting portfolio and the target portfolio in the absence of transaction costs.

⁵We have also tried using the Sharpe ratio of returns net of transaction costs as the calibration criterion. In addition, we have used 10-fold cross-validation and generalized cross-validation as defined by Fu (1998, section 5) to calibrate simultaneously the transaction cost parameter (κ) and the type of transaction cost (proportional or quadratic). However, we find that the results are qualitatively similar, and thus we do not report them to conserve space.

⁶This is illustrated by Figure 1 in Appendix B of the online companion, which shows the out-of-sample monthly returns of the different types of shortsale-unconstrained mean-variance portfolios and shortsale-constrained mean-variance portfolios. While the shortsale-unconstrained mean-variance portfolios are very sensitive to estimation error—and consequently, their returns may be quite extreme in some months—the shortsale-constrained mean-variance portfolios are more resilient to estimation error and thus result in more stable out-of-sample returns.

⁷From the out-of-sample mean and variance for the different portfolios, reported in Tables 3 and 4 of Appendix B in the online companion, we observe that considering transaction costs generally does not help to reduce the variance of portfolio returns, but it helps to increase the mean. Thus the gains from using the proposed data-driven approaches are obtained from improvements in out-of-sample means, rather than variances. Moreover, from Tables 5 and 6 of Appendix B in the online companion, we find that the relative performance of the different portfolios in terms of *tail risk* as measured by value at risk and kurtosis is similar to that in terms of standard risk as measured by variance.

References

- Almgren R, Thum C, Hauptmann E, Li H (2005) Direct estimation of equity market impact. *Risk* 18(7):58–62.
- Balduzzi P, Lynch AW (1999) Transaction costs and predictability: Some utility cost calculations. *J. Financial Econom.* 52(1):47–78.
- Bertsimas D, Pachamanova D, Sim M (2004) Robust linear optimization under general norms. *Oper. Res. Lett.* 32(6):510–516.
- Britten-Jones M (1999) The sampling error in estimates of mean-variance efficient portfolio weights. *J. Finance* 54(2):655–671.
- Chopra VK, Ziemba WT (1993) The effect of errors in means, variances, and covariances on optimal portfolio choice. *J. Portfolio Management* 19(2):6–11.
- Constantinides GM (1986) Capital market equilibrium with transaction costs. *J. Political Econom.* 94(4):842–862.
- Davis MHA, Norman AR (1990) Portfolio selection with transaction costs. *Math. Oper. Res.* 15(4):676–713.

- DeMiguel V, Garlappi L, Uppal R (2009a) Optimal versus naive diversification: How inefficient is the $1/N$ portfolio strategy? *Rev. Financial Stud.* 22(5):1915–1953.
- DeMiguel V, Garlappi L, Nogales FJ, Uppal R (2009b) A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Sci.* 55(5):798–812.
- Efron B, Gong G (1983) A leisurely look at the bootstrap, the jack-knife, and cross-validation. *Amer. Statistician* 37(1):36–48.
- Fabozzi FJ, Kolm PN, Pachamanova DA, Focardi SM (2007) Robust portfolio optimization. *J. Portfolio Management* 33(3):40–48.
- Frazzini A, Israel R, Moskowitz TJ (2015) Trading costs of asset pricing anomalies. Fama-Miller working paper, University of Chicago Booth School of Business, Chicago.
- Fu WJ (1998) Penalized regressions: The bridge versus the lasso. *J. Comput. Graphical Statist.* 7(3):397–416.
- Garlappi L, Uppal R, Wang T (2007) Portfolio selection with parameter and model uncertainty: A multi-prior approach. *Rev. Financial Stud.* 20(1):41–81.
- Gârleanu N, Pedersen L-H (2013) Dynamic trading with predictable returns and transaction costs. *J. Finance* 68(6):2309–2340.
- Goldfarb D, Iyengar G (2003) Robust portfolio selection problems. *Math. Oper. Res.* 28(1):1–38.
- Gotoh J-Y, Takeda A (2011) On the role of norm constraints in portfolio selection. *Comput. Management Sci.* 8(4):323–353.
- Higham NJ (2002) *Accuracy and Stability of Numerical Algorithms* (SIAM, Philadelphia).
- James G, Witten D, Hastie T, Tibshirani R (2013) *An Introduction to Statistical Learning* (Springer, New York).
- Kourtis A (2015) A stability approach to mean-variance optimization. *Financial Rev.* 50(3):301–330.
- Ledoit O, Wolf M (2008) Robust performance hypothesis testing with the Sharpe ratio. *J. Empirical Finance* 15(5):850–859.
- Lu Z (2011a) A computational study on robust portfolio selection based on a joint ellipsoidal uncertainty set. *Math. Programming* 126(1):193–201.
- Lu Z (2011b) Robust portfolio selection based on a joint ellipsoidal uncertainty set. *Optim. Methods Software* 26(1):89–104.
- Markowitz H (1952) Portfolio selection. *J. Finance* 7(1):77–91.
- Michaud RO (1989) The Markowitz optimization enigma: Is “optimized” optimal? *Financial Anal. J.* 45(1):31–42.
- Muthuraman K, Kumar S (2006) Multidimensional portfolio optimization with proportional transaction costs. *Math. Finance* 16(2): 301–335.
- Schwartz J (1954) The formula for change in variables in a multiple integral. *Amer. Math. Monthly* 61(2):81–85.

Alba V. Olivares-Nadal is lecturer of statistics and operations research at University Pablo de Olavide. Her research interests include robust optimization for portfolio and inventory problems and the application of mathematical optimization techniques to enhance sparsity in multivariate time-series and linear regression models.

Victor DeMiguel is professor of management science and operations at London Business School. His research focuses on portfolio optimization in the presence of parameter uncertainty, transaction costs, and taxes.