

Dynamic Asset Allocation with Predictable Returns and Transaction Costs*

Pierre Collin-Dufresne

École Polytechnique Fédérale de Lausanne

email: pierre.collin-dufresne@epfl.ch

Kent Daniel

Columbia University

email: kd2371@columbia.edu

Ciamac C. Moallemi

Columbia University

email: ciamac@gsb.columbia.edu

Mehmet Sağlam

University of Cincinnati

email: mehmet.saglam@uc.edu

Preliminary Draft: August 15, 2013

This Draft: June 15, 2015

Abstract

We propose a simple approach to dynamic multi-period portfolio choice with transaction costs that is tractable in settings with a large number of securities, realistic return dynamics with multiple risk factors, many predictor variables, and stochastic volatility. We obtain a closed-form solution for an optimal trading rule when the problem is restricted to a broad class of strategies we define as ‘linearity generating strategies’ (LGS). When restricted to this class, the non-linear dynamic optimization problem reduces to a deterministic linear-quadratic optimization problem in the parameters of the trading strategies. We show that the LGS approach dominates several alternatives in realistic settings, and in particular when the covariance structure and transaction costs are stochastic.

*For valuable comments and suggestions we thank Michael Brandt, Darrell Duffie, Bernard Dumas, Nicolae Gârleanu, and participants of seminars at Princeton, AQR, Barclays, Lombard-Odier, Kepos Capital, and the 2012 Netspar meetings, the FRIC conference at Copenhagen Business School, the Ninth Imperial College Hedge Fund Conference, the LSE Roundtable on Asset Management, and the 2015 AFA meetings.

1 Introduction

The seminal contribution of Markowitz (1952) has spawned a large academic literature on portfolio choice. The literature has extended Markowitz's one period mean-variance setting to dynamic multiperiod setting with a time-varying investment opportunity set and more general objective functions.¹ Yet there seems to be a wide disconnect between this academic literature and the practice of asset allocation, which still relies mostly on the original one-period mean-variance framework. Indeed, most MBA textbooks tend to ignore the insights of this literature, and even the more advanced approaches often used in practice, such as that of Grinold and Kahn (1999), propose modifications of the single period approach with *ad-hoc* adjustments designed to give solutions which are more palatable in a dynamic, multiperiod setting.

Yet the empirical evidence on time-varying expected returns suggests that the use of a dynamic approach should be highly beneficial to asset managers seeking to exploit these different sources of predictability.²

One reason for this disconnect is that the academic literature has largely ignored realistic frictions such as trading costs, which are paramount to the performance of investment strategies in practice. This is because introducing transaction costs and price impact in the standard dynamic portfolio choice problem tends to make it intractable. Indeed, most academic papers studying transaction costs focus on a very small number of assets (typically two) and limited predictability (typically none).³ Extending their approach to a large number of securities and several sources of predictability quickly runs into the curse of dimensionality.

In this paper we propose an approach to dynamic portfolio choice in the presence of transaction costs that can deal with a large number of securities and realistic return generating processes. For example, our approach can handle a large number of predictors, a general factor structure for returns, and stochastic volatility. The approach relies on three features.

¹Merton (1969), Merton (1971), Brennan, Schwartz, and Lagnado (1997), Kim and Omberg (1996), Campbell and Viceira (2002), Campbell, Chan, and Viceira (2003), Liu (2007), Detemple and Rindisbacher (2010) and many more. See Cochrane (2007) for a survey.

²The academic literature has documented numerous variables which forecast the cross-section of equity returns. Stambaugh, Yu, and Yuan (2012) provides a list of many of these variables, and also argue that the structure and magnitudes of this forecastability exhibits considerable time variation.

³Constantinides (1986), Davis and Norman (1990), Dumas and Luciano (1991), Shreve and Soner (1994) study the two-asset (one risky-one risk-free) case with iid returns. Cvitanić (2001) surveys this literature. Balduzzi and Lynch (1999) and Lynch and Balduzzi (2000) add some predictability in the risky asset. Lynch and Tan (2011) extend this to two risky assets at considerable computational cost. Liu studies the multi-asset case under CARA preferences and for *i.i.d.* returns.

First, we assume investors maximize their expected terminal wealth net of a risk-penalty that is linear in the variance of their portfolio return. Second, we assume that the total transaction cost for a given trade is quadratic in the dollar trade size. Third, we assume that the conditional mean vector and covariance matrix of returns are known functions of an observable state vector, and the dynamics of this state vector can be simulated. Thus, this framework nests most factor based models that have been proposed in the literature.

For a standard set of return generating processes, the portfolio optimization problem does not admit a simple solution because the wealth equation and return generating process introduce non-linearities in the state dynamics. Thus, the problem falls outside the linear-quadratic class which is known to be tractable (Litterman (2005), Gârleanu and Pedersen (2013)) even though we use the same objective function as they do. However, we identify a particular set of strategies, which we call “linearity generating strategies” (LGS), for which the problem admits a closed-form solution. An LGS is defined as a strategy for which the dollar position in each security is a weighted average of current and lagged stock “exposures” interacted with its own past returns (*i.e.*, it is effectively a linear combination of managed portfolios).

The exposures are selected *ex-ante* for each stock, and should include all stock specific state variables on which the optimal dollar position in each security depends: variables summarizing the conditional expected return and variance for each security, and variables summarizing the cost of trading this security. Note that the exposures can also include variables such as the vector of optimal security weights when transaction costs are zero, or the solution to a related optimization problem, such as that proposed by Litterman (2005) and Gârleanu and Pedersen (2013) or various rules of thumb (*e.g.*, Brown and Smith (2011)).

The optimal trade and position for each security will be a linear function of that security’s exposures, interacted with its past-returns, for a set of lags. This implies a very high dimensional optimization problem. While one would anticipate that this high-dimensional problem is difficult to solve, we show that for strategies in the LGS class this optimization problem reduces to a deterministic linear-quadratic problem that can be solved very efficiently.

Another key question is whether the set of LGS’s is sufficiently rich that the optimal LGS approximates the unconstrained optimum. This is an empirical question. However, assuming the specifications of the return generating process and transaction cost function are correct, the LGS can always be designed to perform as well as any alternative approach: the reason is that the solution of any other approach can be used as an input to the LGS approach. The magnitude of the improvement of the LGS will depend on the value of the additional exposures in getting closer to the unconstrained optimum.

We solve several realistic examples which allow us to study the magnitude of this improvement in different settings. First, we compare the performance of our approach to that of several alternatives in two benchmark simulated economies: one we call the *characteristics model* and the other the *factor model*. In both cases expected returns are driven by three characteristics which mimic the well-known reversal (Jegadeesh 1990), momentum (Jegadeesh and Titman 1993) and long-term-reversal/value (DeBondt and Thaler 1985, Fama and French 1993) effects. However, the economies differ in their covariance matrix of returns. The characteristics model assumes that the covariance matrix is constant (implying a failure of the APT in a large economy). In contrast, the factor model assumes that the three characteristics reflect loadings on common factors. Thus, they are reflected in the covariance matrix of returns. Since factor exposures are time-varying and drive both expected returns and covariances, in this model the covariance matrix is stochastic.

The characteristics model is similar to the return model used in the recent works of Litterman (2005) and Gârleanu and Pedersen (2013) (henceforth L-GP). Their linear-quadratic programming approach provides a useful benchmark since they solve for the exact closed-form solution for strategies with a similar objective function.⁴ Indeed, we find that the LGS and the L-GP closed-form of solution perform almost equally well in the characteristics based economy we simulate, as the covariance matrix is close to time-invariant.⁵

However, in the factor model economy, where the covariance matrix changes as the factor loadings of individual securities change, the L-GP solution is further from optimal, since their approach relies on a constant covariance matrix, and their trading rule significantly underperforms our approach based on LGS. This is because the latter explicitly takes into account the dual effect of higher factor exposures in both raising expected returns and covariances as well as their expected future dynamics.

We also investigate the performance of a trading strategy involving the 100 largest stocks traded on the NYSE over the time period from 1930 to 2014. We focus on the return predictability arising from short-term reversal, price-momentum and long-term reversal, which are a well-known predictors of stock returns.⁶ Since the half-life of these predictors are very

⁴One important difference is that to obtain a closed-form solution Litterman (2005) and Gârleanu and Pedersen (2013) specify their model for price changes and not returns and the objective function of the investor in terms of number of shares. They further assume the covariance matrix of price changes is constant. This allows them to retain a linear objective function avoiding the non-linearity in the wealth equation due to the compounding of returns over time.

⁵More precisely, the GP solution is optimal if the covariance matrix of changes in the dollar price per share is time invariant.

⁶See, respectively, Jegadeesh (1990) and Lehmann (1990), Jegadeesh and Titman (1993), and DeBondt and Thaler (1985).

different, ranging from a few days to several years, the optimal trading strategy is very sensitive to the presence of transaction costs. We document that our approach significantly outperforms an alternative often used in practice (e.g., Grinold and Kahn (1999)), which consists of the myopic mean-variance trading strategy where transaction costs are scaled by a multiplier, which is chosen to maximize the in-sample Sharpe ratio of the strategy. The t-cost multiplier is a reduced-form approach to account for the the half-life of expected returns (which depends on the half-life of the predictor variables). This reduced-form approach is dominated by our LGS because the t-cost multiplier can only capture the average half-life of stocks' expected returns. Instead, since the expected return is generated by several predictors with different half-lives, intuitively the optimal strategy attaches different t-cost multipliers at different times depending on the relative importance of each predictor in generating the expected return of stocks.

There is a growing literature on portfolio selection that incorporates return predictability with transaction costs. Balduzzi and Lynch (1999) and Lynch and Balduzzi (2000) illustrate the impact of return predictability and transaction costs on the utility costs and the optimal rebalancing rule by discretizing the state space of the dynamic program. Their approach runs into the curse of dimensionality and only applies to very few stocks and predictors. Brown and Smith (2011) discuss this issue and instead provide heuristic trading strategies and dual bounds for a general dynamic portfolio optimization problem with transaction costs and return predictability that can be applied to larger number of stocks.

Our approach is closely related to two strands of literature: First, Brandt, Santa-Clara, and Valkanov (2009, BSV) propose an approach in which the weight on each asset is a linear function of a set of asset “characteristics” that are specified *ex-ante* as likely to be useful for portfolio selection.⁷ The optimal vector of characteristic weights is found by maximizing the utility the investor would have obtained by implementing the policy over a historical sample period. The BSV approach explicitly avoids modeling the asset return distribution, and therefore avoids the problems associated with the multi-step procedure of first explicitly modeling the asset return distribution as a function of observable variables, and then performing portfolio optimization as a function of the moments of this estimated distribution.⁸ However, since the BSV approach relies on numerical optimization, the number of predictive variables is necessarily limited. Further, since the performance of the objective function is optimized in sample, to avoid over-fitting the number of parameters and predictors

⁷See also Ait-Sahalia and Brandt (2001), Brandt and Santa-Clara (2006) and Moallemi and Saglam (2012).

⁸See Black and Litterman (1991b), Chan, Karceski, and Lakonishok (1999), as well as references given in footnote 2 of BSV (p. 3412).

should be small.

In contrast, in the LGS approach the optimal solution is closed-form. We can thus achieve a greater flexibility in parameterizing the trading rule, something that is useful in settings where transaction costs play a role.⁹

As noted earlier, our approach is also closely related to the L-GP approach – as proposed by Litterman (2005) and Gârleanu and Pedersen (2013). L-GP obtain a closed-form solution for the optimal portfolio choice in a model where: (1) expected price change per share for each security is a linear, time-invariant function of a set of predictor variables; (2) the covariance matrix of price changes per share is time-invariant; and (3) trading costs are a quadratic function of the number of shares traded, and investors have a linear-quadratic objective function. Their approach relies heavily on linear-quadratic stochastic programming (see, *e.g.*, Ljungqvist and Sargent (2004)). Our approach considers a problem that is more general, in that our return generating process can allow for a general factor structure in the covariance matrix with stochastic volatility, the transaction costs can be stochastic, and our objective function is written in terms of dollar holdings. In general, such a problem does not belong to the linear-quadratic class and thus does not admit a simple closed-form along the lines of the L-GP solution. Our contribution is to find a special parametric class of portfolio policies, such that when the portfolio choice problem is considered in that class it reduces to a deterministic linear-quadratic program in the policy parameters.

2 Model

In this section we lay out the return generating process for the set of securities our agent can trade. Then we describe the portfolio dynamics in the presence of transaction costs. Finally, we present the agent’s objective function and our solution technique.

2.1 Security and factor dynamics

We consider a dynamic portfolio optimization problem where an agent can invest in N risky securities with price $S_{i,t}$ $i = 1, \dots, N$ and a risk-free cash money market with value $S_{0,t}$. We assume that security i pays a dividend $D_{i,t}$ at time t . The gross return to our securities is thus defined by $R_{i,t+1} = \frac{S_{i,t+1} + D_{i,t+1}}{S_{i,t}}$. We assume that the conditional mean return vector and covariance matrix of security returns are both known functions of an observable vector

⁹However, one concern in the LGS approach is that, if the return generating process is misspecified – such as by having inconsistent expected return and covariance processes (Black and Litterman 1991a), then the LGS approach will also “find” these solutions. A solution to this concern for the LGS approach is to first verify that the instantaneous MVE portfolio (*i.e.*, the zero t-cost solution) for the specified return generating process is reasonable. Once this is done, a large set of instruments can be used in constructing an optimal strategy in the LGS set.

of state variables X_t :

$$\mathbf{E}_t[R_{t+1}] = 1 + m(X_t, t) \quad (1)$$

$$\mathbf{E}_t[(R_{t+1} - \mathbf{E}_t[R_{t+1}])(R_{t+1} - \mathbf{E}_t[R_{t+1}])'] = \Sigma_{t \rightarrow t+1}(X_t, t) \quad (2)$$

The vector of observable state variable X_t may include both individual security characteristics (such as individual firms' book to market ratios, past returns or idiosyncratic volatilities) as well as common drivers of security returns (such as market volatility, and market or industry factors).

It is important for our approach that the dynamics of X_t be known – to implement the LGS approach requires that we be able to calculate the unconditional moments of security returns interacted with exposures.¹⁰ An example that nests many return generating processes used in the literature is:

$$R_{i,t+1} = g(t, \beta_{i,t}^\top (F_{t+1} + \lambda_t) + \epsilon_{i,t+1}) \quad i = 1, \dots, N \quad (3)$$

for some function $g(t, \cdot) : \mathbb{R} \rightarrow \mathbb{R}$, increasing in the second argument, and where:

- $\beta_{i,t}$ is the $(K, 1)$ vector of firm i 's factor exposures at time t .
- F_{t+1} is the $(K, 1)$ vector of random (as of time t) factor realizations over period $t + 1$. F_{t+1} is mean 0, and follows a multivariate GARCH process with conditional covariance matrix $\Omega_{t,t+1}$.
- $\epsilon_{i,t+1}$ is security i 's idiosyncratic return over period $t + 1$.
We assume that $\epsilon_{.,t+1}$ are mean zero, have a time-invariant covariance matrix Σ_ϵ , and are uncorrelated with the contemporaneous factor realizations.
- λ_t is the $(K, 1)$ vector of conditional expected factor returns at time t .

In this case the vector of state variables $X_t = [\beta_{1,t}; \beta_{2,t}; \dots \beta_{N,t}; \lambda_t; \Omega_{t,t+1}]$ has $NK + K + K \cdot (K + 1)/2$ elements. We further assume that $\beta_{i,t}$ and λ_t are observable and follow some known dynamics. In the empirical applications below, we assume that both λ_t and the $\beta_{i,t}$ follow Gaussian AR(1) processes.

Note that this setting captures two standard return generating processes from the literature:

¹⁰See Section 2.7. Note that these moments can either be calculated analytically, or via simulation.

1. The “**discrete exponential affine**” model for security returns in which log-returns are affine in factor realizations:¹¹

$$\log R_{i,t+1} = \alpha_i + \beta_{i,t}^\top (F_{t+1} + \lambda) + \epsilon_{i,t+1} - \frac{1}{2} (\sigma_i^2 + \beta_{i,t}^\top \Omega \beta_{i,t})$$

2. The “**linear affine factor model**” where returns (and therefore also excess returns) are affine in factor exposures:

$$R_{i,t+1} = \alpha_i + \beta_{i,t}^\top (F_{t+1} + \lambda_t) + \epsilon_{i,t+1}$$

As we show below, our portfolio optimization approach is equally tractable for both of these return generating processes. We emphasize that the approach does not rely on this factor structure assumption. We only require that there exists some known relation between the conditional first and second moments of security returns and the known state vector X_t so that conditional means and variances of security returns can be simulated along with the state vector.

2.2 Cash and security position dynamics

We assume discrete time dynamics. At the end of each period t the agent buys $u_{i,t}$ dollars of security i at price $S_{i,t}$. All trades in risky securities incur transaction costs which are quadratic in the dollar trade size. Trades in risky securities are financed using the cash money market position, which we assume incurs no trading costs. The cash position (w_t) and dollar holdings ($x_{i,t}$) in each security $i = 1, \dots, N$ held at the end of each period t are thus given by:

$$\begin{aligned} x_{i,t} &= x_{i,t-1} R_{i,t} + u_{i,t} & i = 1, \dots, N \\ w_t &= w_{t-1} R_{0,t} - \sum_{i=1}^N u_{i,t} - \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N u_{i,t} \Lambda_t(i, j) u_{j,t}, \end{aligned}$$

or, in vector notation,

$$x_t = x_{t-1} \circ R_t + u_t \tag{4}$$

$$w_t = w_{t-1} R_{0,t} - \mathbf{1}^\top u_t - \frac{1}{2} u_t^\top \Lambda_t u_t \tag{5}$$

¹¹The continuous time version of this model is due to Vasicek (1977), Cox, Ingersoll, and Ross (1985), and generalized in Duffie and Kan (1996). The discrete time version is due to Gourieroux, Monfort, and Renault (1993) and Le, Singleton, and Dai (2010).

where the operator \circ denotes element by element multiplication if the matrices are of same size or if the operation involves a scalar and a matrix, then that scalar multiplies every entry of the matrix.¹²

The matrix Λ_t captures (possibly time-varying and stochastic) quadratic transaction/price-impact costs, so that $\frac{1}{2}u_t^\top \Lambda_t u_t$ is the dollar cost paid given a vector of trades at time t of (dollar) size u_t . Without loss of generality, we assume this matrix is symmetric. Gârleanu and Pedersen (2013) discuss the micro-economic foundations for quadratic costs. This assumption is also very convenient analytically.

2.3 Objective function

We assume that the agent is endowed with an initial portfolio of dollar holdings in securities x_0 and cash of w_0 . We assume that the investor's objective function is to maximize his expected terminal wealth net of a risk penalty which, following L-GP, we take to be linear in the sum of per-period variances. For simplicity, we also assume that the risk-free rate is zero, i.e., $R_{0,t} = 1$.¹³ Thus the objective is:

$$\max_{u_1, \dots, u_{T+1}} \mathbb{E} \left[w_{T+1} + x_{T+1}^\top \mathbf{1} - \sum_{t=0}^T \frac{\gamma}{2} x_t^\top \Sigma_{t \rightarrow t+1} x_t \right] \quad (8)$$

Recall that $\Sigma_{t \rightarrow t+1} = \mathbb{E}_t [(R_{t+1} - \mathbb{E}_t[R_{t+1}])(R_{t+1} - \mathbb{E}_t[R_{t+1}])^\top]$ is the conditional one-period variance-covariance matrix of returns, and that γ can be interpreted as the coefficient of risk aversion.

¹²The timing convention could be changed so that the agent buy $u_{i,t}$ dollars of security i at price $S_{i,t}$ at the beginning of period t . In that case the dynamics would be:

$$x_{t+1} = (x_t + u_t) \circ R_{t+1} \quad (6)$$

$$w_{t+1} = (w_t - \mathbf{1}^\top u_t - \frac{1}{2} u_t^\top \Lambda_t u_t) R_{0,t+1} \quad (7)$$

All our results go through for this alternative timing convention. We make the choice in the text because, for one parameterization of our objective function identified below, it allows us to closely approximate the objective function of Litterman (2005) and Gârleanu and Pedersen (2013) and thus makes the link between the two frameworks more transparent.

In addition, note that we are assuming that all dividends are reinvested at zero cost.

¹³It is straightforward to extend our approach to a non-zero risk-free rate and to an objective function that is linear-quadratic in the position vector (i.e., $F(x_T, w_T) = w_T + a_1^\top x_T - \frac{1}{2} x_T^\top a_2 x_T$) rather than linear in total wealth. See Appendix A.

By recursion we can write:¹⁴

$$x_{T+1} = x_0 + \sum_{t=0}^T x_t \circ r_{t+1} + \sum_{t=1}^{T+1} u_t \quad (9)$$

$$w_{T+1} = w_0 - \sum_{t=1}^{T+1} (u_t^\top \mathbf{1} + \frac{1}{2} u_t^\top \Lambda_t u_t) \quad (10)$$

where we have defined the *net return* $r_{t+1} = R_{t+1} - 1$ with corresponding expected net return $m_t = \mathbb{E}_t[R_{t+1}] - 1$. Inserting in the objective function and simplifying:¹⁵ we find the optimization reduces to¹⁶

$$\max_{u_1, \dots, u_T} \mathbb{E} \left[\sum_{t=1}^T \left\{ x_t^\top m_t - \frac{\gamma}{2} x_t^\top \Sigma_{t \rightarrow t+1} x_t - \frac{1}{2} u_t^\top \Lambda_t u_t \right\} \right] \quad s.t. \quad \text{eq (4)} \quad (11)$$

We see that this objective function is very similar to that used in L-GP (see, *e.g.*, equation (4) of GP): we maximize the expected sum of local-mean-variance objectives, net of transaction costs paid. However, there are several notable and important differences. First, our objective function is in terms of dollar holdings (x_t, w_t) and dollar trades (u_t). In contrast, the L-GP objective function is in terms of number of shares held and traded (their x_t and Δx_t). For the price processes, our expected returns (m 's) and covariance matrix ($\Sigma_{t-1 \rightarrow t}$) are in terms of returns, while in the L-GP framework r_{t+1} and Σ necessarily denote the expected price change and the price-change variance, both on a per share basis. Finally, our approach can accommodate an arbitrary stochastic return covariance matrix which can be a function of the state variables, while the L-GP approach requires that the price-change covariance matrix be deterministic.

At first glance this may appear to be an innocuous change of units. However, to obtain an analytical solution, the L-GP framework requires a constant covariance matrix of price changes. This implies that the *return* variance will be inversely related to the security price squared: if a security's price falls from \$100/share to \$50/share, the return variance must quadruple. It also requires that the transaction cost function – as measured in the transaction costs per share traded – must be independent of the share price. This is generally inconsistent

¹⁴Indeed, $x_{T+1} = x_T \circ (R_{T+1} - 1) + x_T + u_{T+1} = x_T \circ (R_{T+1} - 1) + x_{T-1} \circ (R_T - 1) + x_{T-1} + u_T + u_{T+1} = \dots$

¹⁵Clearly, $u_{T+1} = 0$ is optimal. Thus the difference between the value functions of the two problems in equation 8 and equation 11 is constant.

¹⁶In the specification given here, γ is constant over time – risk-aversion does not change with the agent's wealth level. It is difficult to extend our framework to make γ a function of the agent's wealth at time t . However, it is also straightforward to extend γ to a time-varying and possibly stochastic parameter which is a function of time or the state variables.

with empirical evidence on security return dynamics. To better illustrate this, we first focus on the special case where expected return and variances are constant, which can be solved for in closed-form before turning to the more general case with predictability.

2.4 Comparing the optimal strategy when returns as opposed to price changes have constant expectation and variance

If m_t, Σ_t , and Λ_t are constant, then the optimal portfolio choice problem in equation (11) admits a closed-form solution. For simplicity, we focus on the one-asset case in an infinite horizon stationary model. In Appendix (A.1) we derive the solution to the following problem:

$$\max_{u_1, \dots,} \mathbb{E} \left[\sum_{t=1}^{\infty} \rho^t \left\{ x_t^\top m - \frac{\gamma}{2} x_t^\top \Sigma x_t - \frac{1}{2} u_t^\top \Lambda u_t \right\} \right] \quad s.t. \quad \text{eq (4)} \quad (12)$$

where $\rho < 1$ is a time discount factor.¹⁷ We show that the optimal dollar trade u_t is linear-affine in the current dollar position held in the security at the time of the trade, i.e.,

$$u_t = a_0 + (a_1 - 1)\bar{x}_t \quad (13)$$

where $\bar{x}_t = x_{t-1}R_t$ and the coefficients a_0, a_1 are given explicitly in equation (81) in Appendix A.1. Instead, if one assumes that the expectation and variance of price changes are constant, then the optimal policy would imply an optimal trade such that the number of shares traded h_t is linear affine in the number of shares held, n_t :¹⁸

$$h_t = b_0 + (b_1 - 1)n_t \quad (14)$$

where the coefficients b_0, b_1 are given in equation (85) in the appendix. Clearly, these two trading rules are inconsistent (since by definition $u_t = h_t S_t$ and $x_t = n_t S_t$ both equations (13) and (14) cannot both hold at the same time). As expected, the optimal trading strategy obtained for constant covariance of returns differs from that obtained for a constant covariance of price changes.

One important difference between the two solutions is that if the covariance of price changes is constant, then if at some point we hold the mean-variance optimal portfolio (*i.e.*, if $x_t = (\gamma\Sigma)^{-1}m$ or equivalently $n_t = (\gamma\Sigma_s)^{-1}\mu_s$ where $\Sigma_s = \Sigma * S_t^2$ and $\mu_s = m * S_t$ are defined as the variance and expectation of price changes respectively) then it is optimal to never trade hence-forth (see Appendix B.6). This implies that if we held the mean-variance optimal

¹⁷This stationary infinite horizon objective function is also used in Gârleanu and Pedersen (2013).

¹⁸Both are linked by the relation $n_t = n_{t-1} + h_t$.

portfolio, and the price of a security were to fall by a factor of two, the optimal solution would be not to trade. Intuitively, there is no trade to rebalance the portfolio because, given the assumed dynamics (constant expectation and variance of price changes), when the price halves, the security’s expected return and return volatility both double, meaning the optimal dollar holdings also halve, so there is no motive for rebalancing.

If instead we assume that the expectation and variance of returns (rather than price changes) is constant, then there is no position such that it is never optimal to trade at all future dates. This is because return shocks induce random changes in future dollar positions via equation (4), which in turn lead to deviations in dollar portfolio holdings from the first best, and thus to a rebalancing motive for trading even in the *i.i.d.* case. This rebalancing motive for trading is the one investigated in the traditional transaction cost literature (such as Constantinides (1986)). In addition, we point out in the appendix that in the *i.i.d.* case, there exists a position x_{no} given in equation (90) such that it is optimal not to trade for one period (i.e., if $x_t = x_{no}$ then $u_t = 0$). However, interestingly, if $\Lambda\rho \neq 0$ and $\Sigma + \mu^2 \neq \mu$, then this no-trade position is not equal to the mean-variance efficient portfolio (i.e., $x_{no} \neq \Sigma^{-1}m$). The intuition is that the current position does not reflect where it is expected to be in one period, since it will experience random return shocks. So even in the *i.i.d.* case, current optimal trades reflect a trade-off between where we are today and where we expect to be in the future given the return shocks we will experience.

While we can obtain a closed-form solution in the *i.i.d.* case, the general framework we lay out in the previous section allows for security price processes to have more general dynamics, with time-varying expected returns, variances and trading costs. In general, we are unable to obtain a closed-form solution. However, just as in the *i.i.d.* case the model will typically capture this rebalancing motive for trading (which is at the heart of the classic Merton (1969) dynamic portfolio optimization with constant investment opportunity set). The *i.i.d.* solution is also interesting as it motivates our choice of focusing on ‘linearity generating strategies.’ Indeed, combining the linearity of the trading rule in (13) and the dynamics of the state in (4) and iterating backwards we see that both the optimal trade and the optimal position are of the form

$$u_t = \sum_{s \leq t} \pi_{s,t} R_{s \rightarrow t} \tag{15}$$

$$x_t = \sum_{s \leq t} \theta_{s,t} R_{s \rightarrow t} \tag{16}$$

where we define the holding period returns $R_{s \rightarrow t} = R_{s \rightarrow s+1} R_{s+1 \rightarrow s+2} \dots R_{t-1 \rightarrow t}$. The optimal

loadings $\pi_{s,t}$ and $\theta_{s,t}$ are deterministic and obtained from the optimal solution. Specifically, we show in the appendix that

$$\theta_{s,t} = a_0 a_1^{t-s}.$$

From equation (4) $\pi_{s,t}$ is such that:

$$\begin{cases} \theta_{s,t} = \theta_{s,t-1} + \pi_{s,t} & \text{for } s < t \\ \theta_{t,t} = \pi_{t,t} & \text{for } s = t \end{cases}$$

Clearly, since $a_1 < 1$ (see the appendix) the optimal position loads on past returns at an exponentially decreasing rate given by $-\log a_1$. This decay rate is a function of the fundamental parameters of the model $(\mu, \Sigma, \Lambda, \rho)$. For example, the decay rate is lower the higher the transaction costs, which shows that the optimal position depends more on past holding period returns when transaction costs are higher.

In this simple example where the expected returns, return variances and the quadratic transaction cost parameter are all constant, the optimal loadings $\theta_{s,t}$ are deterministic. For the general case, where the investment opportunity set is time-varying, we will seek a solution within a set of LGS that has the same structure, but where the loadings on past holding period returns can increase or decrease depending on a set of instruments that can be stochastic. We now turn to the general case and introduce the set of ‘linearity generating strategies’ that we consider.

2.5 Definition of linearity generating strategies

Even though our objective function is similar to those in Litterman (2005) and Gârleanu and Pedersen (2013), the L-GP problems are linear-quadratic because of the restrictions that they place on the return generating process. Our problem is not in this class both because of the non-linearity introduced by the state equation, and because we allow for a far more general set of return generating processes with stochastic expected returns and covariances.

Thus our problem is difficult to solve in full generality, even numerically. Our approach is to instead solve a constrained problem, which is to find the optimal solution among a specific set of ‘linearity generating trading strategies’ (LGS) that is a natural generalization of the form we derived for the simple constant mean and variance problem in equations (15) and (16) above, and for which the problem remains tractable. As we discuss below, as long as we can specify a sufficiently rich set of LGSs, our solution will approach the globally optimal solution.

To define our set of LGS we first specify, for each security, a K -vector $B_{i,t}$ of ‘security

exposures.” The exposures are typically non-linear transformations of the general state vector X_t (*i.e.*, $B_{i,t} = h_i(t, X_t)$). For example, $B_{i,t}$ may include the individual security’s conditional expected return divided by its conditional variance (see, *e.g.*, Ait-Sahalia and Brandt (2001)), the optimal dollar position in the security in the absence of transaction costs given by the myopic solution, or a t-cost aware solution from another method. More generally, it would include security specific factor exposures, conditional variances and other relevant information for portfolio formation.

We then define the set of LGS as strategies for which the dollar holdings and dollar trades of security i are linear functions of current and lagged exposures interacted with sets of K -dimensional vectors of parameters, $\pi_{i,s,t}$ and $\theta_{i,s,t}$, defined for all $i = 1, \dots, N$ and for all $s \leq t$. These parameters fully determine the dollar holdings ($x_{i,t}$) and the corresponding dollar trades ($u_{i,t}$) for each asset i via the parametric relations:

$$x_{i,t} = \sum_{s=0}^t \theta_{i,s,t}^\top \mathcal{B}_{i,s \rightarrow t} \quad \text{for } t = 0, \dots, T \quad (17)$$

$$u_{i,t} = \sum_{s=0}^t \pi_{i,s,t}^\top \mathcal{B}_{i,s \rightarrow t} \quad \text{for } t = 1, \dots, T \quad (18)$$

where $\mathcal{B}_{i,s \rightarrow t}$ is defined as the vector of time s exposures $B_{i,s}$, scaled by the gross-return on security i between s and t :

$$\mathcal{B}_{i,s \rightarrow t} = B_{i,s} R_{i,s \rightarrow t}. \quad (19)$$

In effect, the dollar trades and dollar positions in security i at time t in asset i ($x_{i,t}$) can be thought of as a weighted sum of simple buy and hold trading strategies that went long the security at past dates ($s < t$) proportionally to time s exposures and held the security until date t .¹⁹

However in the LGS framework, this time- s scaled exposure can be built up gradually after time s , and then sold gradually. Scaled exposure, because it is scaled by the firm’s cumulative gross return, is time invariant: if you bought one unit of scaled-exposure at time s and did not trade further, you would still hold one unit at all future times. The value of a unit of scaled time- s exposure at time t is given by $\mathcal{B}_{i,s \rightarrow t}$. The number of units of time- s exposure purchased at time $t \geq s$ is given by $\pi_{i,s,t}$, and the total number of units held at time t ($\theta_{i,s,t}$) is just the sum of the number of units purchased between s and t .

Perhaps the easiest way to illustrate this is to examine the equations for the dollar

¹⁹We see that LGS nest the closed-form solution obtained in equations (15) and (16) above for the special case of constant expected return and variances, which obtain with $B_{i,s} = 1$.

positions and trades of firm i at $t = 0, 1, 2$, as given below:

$$\begin{array}{r}
x_{i,0} = \theta_{i,0,0}^\top B_{i,0} \\
\hline
u_{i,1} = \pi_{i,0,1}^\top \mathcal{B}_{i,0 \rightarrow 1} \qquad + \pi_{i,1,1}^\top B_{i,1} \\
x_{i,1} = \underbrace{(\theta_{i,0,0} + \pi_{i,0,1})^\top}_{=\theta_{i,0,1}} \mathcal{B}_{i,0 \rightarrow 1} \qquad + \underbrace{\pi_{i,1,1}^\top}_{=\theta_{i,1,1}^\top} B_{i,1} \\
\hline
u_{i,2} = \pi_{i,0,2}^\top \mathcal{B}_{i,0 \rightarrow 2} \qquad + \pi_{i,1,2}^\top \mathcal{B}_{i,1 \rightarrow 2} \qquad + \pi_{i,2,2}^\top B_{i,2} \\
x_{i,2} = \underbrace{(\theta_{i,0,0} + \pi_{i,0,1} + \pi_{i,0,2})^\top}_{=\theta_{i,0,2}} \mathcal{B}_{i,0 \rightarrow 2} \qquad + \underbrace{(\pi_{i,1,1} + \pi_{i,1,2})^\top}_{=\theta_{i,1,2}} \mathcal{B}_{i,1 \rightarrow 2} \qquad + \underbrace{\pi_{i,2,2}^\top}_{=\theta_{i,2,2}^\top} B_{i,2}
\end{array}$$

The first equation gives the initial position as a function of the time 0 exposures. Since the initial position is generally not a choice variable, the vector $\theta_{i,0,0}$ must be constrained so that the first equation holds.²⁰

The second equation gives the first trade, $u_{i,1}$. Note that this trade is a function of both the lagged exposures for time 0, scaled by $R_{i,0 \rightarrow 1}$, and the current ($t = 1$) exposures. The dependence on the time zero exposure is important here, because the optimal trade at $t = 1$ and later are dependent on the initial position. Intuitively, if we are given a large initial position in a security, the strategy will start trading out of that position with the first trade at time 1 – how quickly it trades out will be determined by $\pi_{i,0,1}$.

The third equation gives the total dollar holdings of security i at $t = 1$. $x_{i,1}$ is equal to initial position, grossed up by the realized return on firm i from 0 to 1, plus $u_{i,1}$. However, note that this equation decomposes these holdings into the number of units of scaled time zero exposure $\theta_{i,0,1}$, and time 1 exposure $\theta_{i,1,1}$. Since the first time we purchase time 1 exposure is at time 1, $\theta_{i,1,1} = \pi_{i,1,1}$.

The fourth and fifth equations give, respectively, the time 2 trade and position. The trade is decomposed into the number of units of time 0, 1, and 2 scaled exposure we buy. The vector of costs of the exposures are given by the \mathcal{B} s. $\theta_{i,0,2}$ – the total number of units of time 0 scaled exposure held at time 2 – is the sum of the initial endowment ($\theta_{i,0,0}$) plus the number of units purchased at time 1 and at time 2. The number of units of time 1 exposure held at time 2 ($\theta_{i,1,2}$) is the sum of the number of units purchased at time 1 and 2.

In an environment with transaction costs, the position in the lagged return-scaled time s exposure will generally be accumulated gradually over time. That is, following a shock at time s to exposures that raises a security's expected return (holding its risk constant) the

²⁰In general, one of the elements of the vector $B_{i,0}$ will be a one, so a straightforward way to impose this constraint is to require that the corresponding elements of $\theta_{i,0,0}$ be equal to the initial dollar position $x_{i,0}$. Alternatively, one can have the initial position be an element of $B_{i,0}$ and restrict $\theta_{i,0,0}$ appropriately.

corresponding elements of $\pi_{i,s,t}$ will be positive for t slightly bigger than s , and then will turn negative as t increases, and then finally asymptote to zero. That is, it will be optimal to gradually trade into positions in securities, and then trade out of these positions as the expected return decays towards zero. We will illustrate this via simulation in Section 3.5.

As is apparent in the discussion above, $\theta_{i,s,t}$ and $\pi_{i,s,t}$ must be chosen so that holdings and trades are consistent. Specifically the trades and positions in equations (17) and (18), respectively, are required to satisfy the dynamics given in equations (4) and (5). It follows that the parameter vectors $\pi_{i,s,t}$ and $\theta_{i,s,t}$ have to satisfy the following restrictions, for all $i = 1, \dots, N$:

$$\begin{aligned} \theta_{i,s,t} &= \theta_{i,s,t-1} + \pi_{i,s,t} & \forall t \geq 1 \text{ and } 0 \leq s < t \\ \theta_{i,t,t} &= \pi_{i,t,t} & \forall t \geq 1 \end{aligned} \tag{20}$$

and initial conditions:

$$\begin{aligned} \theta_{i,0,0} B_{i,0} &= x_{i,0} \\ \pi_{i,0,0} &= 0 \end{aligned}$$

These restrictions are intuitive. The first specifies that the number of units of scaled time s exposure held at time t is equal to the number of units held at time $t - 1$ plus the number of units bought at time t . The second restriction specifies that the number of units of scaled time t exposure held at time t is the number bought at time t . Since $B_{i,t}$ is not in the information set until time t , you cannot buy time t exposure before time t . The last two conditions specify that the initial scaled-exposures must be chosen to match the initial holdings $x_{i,0}$, and that the time 0 trade is zero, consistent with the dynamics laid out in Section 2.2.

2.6 The LGS approach and alternative approaches

Intuitively, the dependence of LGS on current exposures is important. In a zero-transaction cost affine portfolio optimization problem where the optimal solution is well-known, the optimal holdings will involve only today's exposures (see, *e.g.*, Liu (2007)).²¹ With transaction costs, allowing today's weights and trades to also depend on lagged security exposures, scaled by each security's return up to today, is useful because these variables summarize the

²¹Note that this is also the choice made by Brandt, Santa-Clara, and Valkanov (2009) for their 'parametric portfolio policies.' However, while BSV specify the loadings on exposure of individual securities to be identical, we allow two securities with identical exposures (and with perhaps different levels of idiosyncratic variance) to have different weights and trades.

positions – held today – as a result of trades made in previous periods. When transactions costs are present, the optimal trades today will generally depend on positions held in past periods. This path-dependence is observed in known closed-form solutions in environments with transaction costs (as in, for example, Constantinides (1986), Davis and Norman (1990), Dumas and Luciano (1991), Liu and Loewenstein (2002) and also our simple closed-form example above).

As noted earlier, the idea of restricting the set of strategies to make the problem tractable is not new. For example, this is the insight underlying Brandt, Santa-Clara, and Valkanov (2009, BSV), who consider strategies which are restricted to be linear in security characteristics. BSV select the optimal characteristic weights by numerically optimizing their objective function over a set of historical returns. Because their approach relies on a numerical, in-sample optimization, they are necessarily restricted to low-dimensional strategies (*i.e.*, with a small number of characteristics/exposures). In contrast, with the LGS approach the optimization is done in closed-form. This allows the use of rich path dependent strategies, with lagged-scaled characteristics as exposures. In a high transaction-costs environment, incorporating lagged-characteristics allows the resulting strategy to slowly trade in and out of each asset in response to exposure shocks.

The only other approach in the literature that yields a closed form solution – the L-GP approach – makes some strong assumptions about the return generating process and the objective function to obtain a closed-form solution. Specifically, these approaches require that the covariance matrix of price changes per share and the per share transaction cost function be constant or, at most, deterministic. With these assumptions, the L-GP solution is the exact optimal solution. If the return generating process is close to the assumed process (constant price-change variance and constant transaction costs) then the L-GP approach will yield a good approximate solution. However, in many realistic settings the solution will be far from optimal, as we show below.

The advantage of the LGS approach is that we can determine the optimal solution given a wide range of security price dynamics. The drawback to our approach is that, for most return generating processes, the solution we derive is only optimal among the set of all solutions that are linear functions of the exposures we select.²² So the key to getting a good solution with the LGS methodology is selecting a set of exposures that come close to spanning the globally optimal solution. One advantage that our method has on this front is that virtually any variable in the information set can be used as an exposure. So, for example, the solution to the simple myopic or the more complex L-GP problem, or both

²²In selected settings, like that explored in Section 2.4, the LGS solution will be globally optimal.

can be chosen as exposures. In this case, our methodology will assign weights to additional exposures – including scaled-lagged exposures — if and only if they provide an improvement over and above what can be obtained with the myopic or L-GP solution. For example, in a setting where the L-GP solution was optimal, these additional exposures would add nothing, consequently they would get no weight and our solution would be identical to the L-GP solution.

The magnitude of the improvement of LGS over alternative solutions depends on how much improvement these additional exposures provide. In Section 3, we investigate this via simulations. First though, we explain how the portfolio optimization can be done in closed-form, within that restricted set.

2.7 The LGS optimization problem

We now solve for the set of exposure weights ($\theta_{i,s,t}$ and $\pi_{i,s,t}$ from equations (17) and (18)) that determine the optimal positions and trades. To proceed, we first rewrite the policies in equation (20) in a concise matrix form. It is convenient to introduce the following notation (inspired from *Matlab*): We write $[A; B]$ (respectively $[A B]$) to denote the vertical (respectively horizontal) concatenation of two matrices.

First, define the $NK(t+1)$ -dimensional vectors π_t and θ_t as

$$\pi_t = [\pi_{1,0,t}; \dots; \pi_{N,0,t}; \pi_{1,1,t}; \dots; \pi_{N,1,t}; \dots; \pi_{1,t,t}; \dots; \pi_{N,t,t}] \quad (21)$$

$$\theta_t = [\theta_{1,0,t}; \dots; \theta_{N,0,t}; \theta_{1,1,t}; \dots; \theta_{N,1,t}; \dots; \theta_{1,t,t}; \dots; \theta_{N,t,t}] \quad (22)$$

Also, we define the following (NK, N) matrices (defined for all $0 \leq s \leq t \leq T$) as the diagonal concatenations of the N vectors $\mathcal{B}_{i,s \rightarrow t} \forall i = 1, \dots, N$:

$$\mathcal{B}_{s,t} = \begin{pmatrix} \mathcal{B}_{1,s \rightarrow t} & 0 & \dots & 0 \\ 0 & \mathcal{B}_{2,s \rightarrow t} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \mathcal{B}_{N,s \rightarrow t} \end{pmatrix}$$

Finally, we define the $(NK(t+1), N)$ matrix \mathcal{B}_t by stacking the $t+1$ matrices $\mathcal{B}_{s,t} \forall s = 0, \dots, t$:

$$\mathcal{B}_t = [\mathcal{B}_{0,t}; \mathcal{B}_{1,t}; \dots; \mathcal{B}_{t,t}] \quad (23)$$

With these definitions, it is straightforward to verify that:

$$u_t = \mathcal{B}_t^\top \pi_t \quad (24)$$

$$x_t = \mathcal{B}_t^\top \theta_t \quad (25)$$

Further, in terms of these definitions the constraints on the parameter vectors in (20) can be rewritten concisely as:

$$\theta_t = \theta_{t-1}^0 + \pi_t \quad (26)$$

where we define $x^0 = [x; \mathbf{0}_{NK}]$ to be the vector x stacked on top of an NK -dimensional vector of zeros $\mathbf{0}_{NK}$.

The usefulness of restricting ourselves to this set of ‘linearity generating trading strategies’ is that optimizing over this set amounts to optimizing over the parameter vectors π_t and θ_t , and that, as we show next, that problem reduces to a **deterministic** linear-quadratic control problem, which can be solved in closed form.

Indeed, substituting the definition of our linear trading strategies from equations (24) and (25) into our objective function in equation (11) and then taking expectations gives:

$$\max_{\pi_1, \dots, \pi_T} \sum_{t=1}^T \theta_t^\top \bar{m}_t - \frac{1}{2} \pi_t^\top \bar{\Lambda}_t \pi_t - \frac{\gamma}{2} \theta_t^\top \bar{\Sigma}_t \theta_t \quad (27)$$

$$\text{subject to } \theta_t = \theta_{t-1}^0 + \pi_t \quad (28)$$

and where we define the vector \bar{m}_t and the square matrices $\bar{\Sigma}_t$ and $\bar{\Lambda}_t$ for $t = 0, \dots, T$ by

$$\bar{m}_t = \mathbf{E}_0[\mathcal{B}_t m_t] \quad (29)$$

$$\bar{\Lambda}_t = \mathbf{E}_0[\mathcal{B}_t \Lambda_t \mathcal{B}_t^\top] \quad (30)$$

$$\bar{\Sigma}_t = \mathbf{E}_0[\mathcal{B}_t \Sigma_{t \rightarrow t+1} \mathcal{B}_t^\top] \quad (31)$$

Note that the time indices also capture their size: \bar{m}_t is a vector of length $NK(t+1)$, and $\bar{\Sigma}_t$ and $\bar{\Lambda}_t$ are square matrices of the same dimensionality. Equation (27) is just the objective function (equation (11)) with the u_t ’s and x_t ’s rewritten as linear functions of the elements in \mathcal{B}_t , with coefficients π_t and θ_t , respectively. Since the policy parameters π_t and θ_t are set at time 0, they can be pulled outside of the expectation operator.

Intuitively equation (27) is a linear-quadratic function of the policy parameters π_t and θ_t , with \bar{m}_t , $\bar{\Lambda}_t$, $\bar{\Sigma}_t$, as the coefficients in this equation. These three components give, respectively, the effect on the objective function of (i) the expected portfolio returns resulting

from trades at time t , (ii) the expected transaction costs paid as a result of trades at time t , and (iii) the effect of the holdings at time t on the risk-penalty component of the objective function.

Since $\bar{m}_t, \bar{\Sigma}_t, \bar{\Lambda}_t$ are not functions of the policy parameters, they can be solved for explicitly. In some settings, this can be done analytically; however if this is not straightforward the moments can always be calculated using simulation. Note that the moments only need to be calculated once. Given these moments, the set of θ_t and π_t that maximize equation (27), these optimal θ_t and π_t will determine all future positions and trades as a function of the (as yet unknown) scaled exposures. Note that these moments do not depend on either the initial conditions, or on the assumptions made about the state vector X_t driving the return generating process R_t , or on the corresponding security-specific exposure dynamics $B_{i,t}$.

We next show how to solve equation (27) using standard methods. Again, this is possible since it is a linear-quadratic equation, albeit a high-dimensional one.

2.8 Closed form solution

We begin with the linear-quadratic problem defined by equations (27) and (28). Define recursively the value function starting from $V(T) = 0$ for all $1 < t < T$ by:

$$V(t-1) = \max_{\pi_t} \left\{ \theta_t^\top \bar{m}_t - \frac{\gamma}{2} \theta_t^\top \bar{\Sigma}_t \theta_t - \frac{1}{2} \pi_t^\top \bar{\Lambda}_t \pi_t + V(t) \right\}$$

subject to $\theta_t = \theta_{t-1}^0 + \pi_t$

Then it is clear that $V(0)$ gives the solution to the problem we are seeking. To solve the problem explicitly, we guess that the value function is of the form:

$$V(t) = -\frac{\gamma}{2} \theta_t^\top M_t \theta_t + L_t^\top \theta_t + H_t \tag{32}$$

with M_t a symmetric matrix. Since $V(T) = 0$, it follows that $M_T = 0$, $L_T = 0$ and $H_T = 0$. To find the recursion plug the guess in the Bellman equation:

$$V(t-1) = \max_{\pi_t} \left\{ \theta_t^\top \bar{m}_t - \frac{1}{2} \pi_t^\top \bar{\Lambda}_t \pi_t - \frac{\gamma}{2} \theta_t^\top (\bar{\Sigma}_t + M_t) \theta_t + L_t^\top \theta_t + H_t \right\} \tag{33}$$

$$\text{subject to } \theta_t = \theta_{t-1}^0 + \pi_t \tag{34}$$

Substituting the constraint (equation (34)) into the value function (equation (33)), we obtain:

$$V(t-1) = \max_{\theta_t} \left\{ \theta_t^\top \bar{m}_t - \frac{1}{2} (\theta_t - \theta_{t-1}^0)^\top \bar{\Lambda}_t (\theta_t - \theta_{t-1}^0) - \frac{\gamma}{2} \theta_t^\top (\bar{\Sigma}_t + M_t) \theta_t + L_t^\top \theta_t + H_t \right\} \quad (35)$$

The first order condition gives the optimal position vector:

$$\theta_t = [\bar{\Lambda}_t + \gamma(\bar{\Sigma}_t + M_t)]^{-1} (\bar{m}_t + L_t + \bar{\Lambda}_t \theta_{t-1}^0),$$

and plugging into the state equation (28), gives the optimal trade vector:

$$\pi_t = [\bar{\Lambda}_t + \gamma(\bar{\Sigma}_t + M_t)]^{-1} (\bar{m}_t + L_t - \gamma(\bar{\Sigma}_t + M_t) \theta_{t-1}^0).$$

Substituting these optimal policies into the Bellman equation in (33) gives another expression for the value function, given the conjectured specification in equation (32):

$$V(t-1) = \frac{1}{2} (\bar{m}_t + L_t + \bar{\Lambda}_t \theta_{t-1}^0)^\top [\bar{\Lambda}_t + \gamma(\bar{\Sigma}_t + M_t)]^{-1} (\bar{m}_t + L_t + \bar{\Lambda}_t \theta_{t-1}^0) + H_t - \frac{1}{2} (\theta_{t-1}^0)^\top \bar{\Lambda}_t \theta_{t-1}^0 \quad (36)$$

Comparing this equation and equation (32) shows that this specification will be correct if H_t , L_t , and M_t are chosen to satisfy the recursions:

$$\begin{aligned} H_{t-1} &= H_t + \frac{1}{2} (\bar{m}_t + L_t)^\top [\bar{\Lambda}_t + \gamma(\bar{\Sigma}_t + M_t)]^{-1} (\bar{m}_t + L_t) \\ L_{t-1}^\top &= \frac{(\bar{m}_t + L_t)^\top [\bar{\Lambda}_t + \gamma(\bar{\Sigma}_t + M_t)]^{-1} \bar{\Lambda}_t}{\gamma} \\ \gamma M_{t-1} &= \frac{\bar{\Lambda}_t - \bar{\Lambda}_t [\bar{\Lambda}_t + \gamma(\bar{\Sigma}_t + M_t)]^{-1} \bar{\Lambda}_t}{\gamma} \end{aligned}$$

with initial conditions $H_T = 0$, $L_T = 0$ and $M_T = 0$ and where $\underline{\mathcal{M}}$ denotes the vector (or matrix) obtained from \mathcal{M} by deleting the last NK rows (or rows and columns).

We have thus derived the optimal value function and the optimal trading strategy in the LGS class.

Before discussing some specific examples it is useful to introduce a set of LGS strategies which uses the exposures lagged at most ℓ periods. This set of “restricted lag” LGS is useful in applications when the time horizon is fairly long, and for signals that have a relatively fast decay rate, so that the dependence on lagged exposures can be restricted without a significant cost. We next show that the same tractability obtains for the restricted lag setting.

2.9 LGS with finite number of lags

In the baseline LGS, trades and positions are a linear function of return-scaled-exposures (*i.e.*, $\mathcal{B}_{i,s,t}$ for $0 \leq s < t$). In most settings, we would expect the coefficients in both the position and the trade equations ($\theta_{i,s,t}$ and $\pi_{i,s,t}$) to converge to zero for $s \ll t$. Indeed, this is the case for the closed-form solution examined in Section 2.4. Further, we shall show via impulse response functions in Section 3.5 that this is also the behavior we observe for more general return generating processes.²³ Thus, to reduce complexity it can be advantageous to use strategies for which the trades are dependent on scaled exposures lagged at most ℓ periods.

We first specify that the trading rule will only trade based on at most ℓ lags, *i.e.* such that:

$$u_{i,t} = \sum_{s=t-\ell \vee 0}^t \pi_{i,s,t}^\top \mathcal{B}_{i,s \rightarrow t} \quad (37)$$

where $t - \ell \vee 0$ denotes the maximum of $t - \ell$ and 0. If we want the holdings to remain linear and of the form:

$$x_{i,t} = \sum_{s=0}^t \theta_{i,s,t}^\top \mathcal{B}_{i,s \rightarrow t} \quad (38)$$

Then we see that the linear constraints in equations (20) have to be modified so as to still satisfy the wealth dynamics in equations (4) and (5). Specifically, we require that:

$$\begin{aligned} \theta_{i,t,t} &= \pi_{i,t,t} & \forall t &\geq 1 \\ \theta_{i,s,t} &= \theta_{i,s,t-1} + \pi_{i,s,t} & \text{for } t - \ell \vee 0 &\leq s < t \\ \theta_{i,s,t} &= \theta_{i,s,t-1} & \text{for } 0 < s &< t - \ell \end{aligned} \quad (39)$$

Since this is still a set of linear constraints we can straightforwardly extend the previous method to derive the optimal LGS strategy with trades that only look back ℓ periods.

However, it is also generally the case that the weights on scaled-exposures will approach zero when they are sufficiently old. Inspecting these constraints, we see that if we impose the additional constraint that ($\pi_{i,t-\ell,t} = -\theta_{i,t-\ell,t-1}$) $\forall t > \ell$ (*i.e.*, that we completely trade out of any remaining time- $(t-\ell)$ scaled-exposure at time t), then it follows that $\theta_{i,s,t} = 0 \forall 0 < s \leq t - \ell$. In other words, by imposing one additional linear constraint on the trading strategy one can find a set of LGS where the trading strategy u_t looks back at most ℓ periods and

²³See, in particular, Figures 2 and 4 and the related discussion.

the dollar position x_t looks back at most $\ell - 1$ periods. Formally, we have

$$u_{i,t} = \sum_{s=t-\ell \vee 0}^t \pi_{i,s,t}^\top \mathcal{B}_{i,s \rightarrow t}$$

and

$$x_{i,t} = \sum_{s=t-\ell+1 \vee 0}^t \theta_{i,s,t}^\top \mathcal{B}_{i,s \rightarrow t}$$

We summarize this second set of linear constraints as:

$$\begin{aligned} \theta_{i,t,t} &= \pi_{i,t,t} && \forall t \geq 1 \\ \theta_{i,s,t} &= \theta_{i,s,t-1} + \pi_{i,s,t} && \forall \text{ and } t - \ell \vee 1 \leq s < t \\ \pi_{i,s,t} &= -\theta_{i,s,t-1} && \text{for } 0 < s = t - \ell \\ \theta_{i,s,t} &= 0 && \text{for } 0 < s \leq t - \ell \end{aligned}$$

Because these constraints are linear, we can follow the approach above and derive the optimal trading strategy coefficients by solving a deterministic dynamic programming problem.

3 Simulation Experiment

How much our proposed method improves on the approaches proposed in the literature is an empirical question, and will clearly depend on the economic environment studied. In this section we present several experiments that allow us to examine where the LGS approach will provide the largest improvements. Specifically, we compare methods in a setting where the return generating process is “characteristics-based” and in a second setting where the return generating process is “factor-based.” As we show below the standard linear-quadratic portfolio approach proposed in Litterman (2005) and Gârleanu and Pedersen (2013) is fairly well-suited to the characteristics-based environment we examine, as the price-change covariance matrix is approximately stationary over short horizons. However, in the factor-based environment, where the covariance structure changes with the exposures to short-lived factors change, the LGS approach significantly outperforms the L-GP approach.

3.1 Characteristics versus factor-based return generating model

We wish to examine the following two environments:

- The factor-based return generating process with excess return and exposure dynamics

$$\begin{aligned} r_{i,t+1} &= \beta_{i,t}^\top (F_{t+1} + \lambda) + \epsilon_{i,t+1}, \\ \beta_{i,t+1}^k &= (1 - \phi_k) \beta_{i,t}^k + \epsilon_{i,t+1}. \end{aligned} \quad (40)$$

- The characteristics based return generating process with excess return and exposure dynamics

$$\begin{aligned} r_{i,t+1} &= \beta_{i,t}^\top \lambda + \epsilon_{i,t+1} \\ \beta_{i,t+1}^k &= (1 - \phi_k) \beta_{i,t}^k + \nu \epsilon_{i,t+1}. \end{aligned} \quad (41)$$

where in both cases we assume that $\beta_{i,t}$ is a $(3, 1)$ vector with elements corresponding to firm i 's exposure to (1) short term reversal (Jegadeesh 1990, Lehmann 1990), (2) medium term momentum (Jegadeesh and Titman 1993), and (3) long-term reversal (DeBondt and Thaler 1985), which we henceforth label *str*, *mom* and *ltr*. We set the half-life of the *str* factor to be 5 days, that of the *mom* factor to be 150 days, and that of the *ltr* factor to be 700 days. These half lives are designed to roughly match the documented horizons at which short-term reversal, momentum, and long-term reversal are typically found.

In both frameworks, expected returns are the product of the *ex-ante* observable factor exposures and the factor premia, $\beta_{i,t}^\top \lambda$. However, in the characteristics based framework, we assume that the conditional covariance matrix of security returns is constant, *i.e.* $\Sigma_{t \rightarrow t+1} = \mathbf{E}_t[\epsilon_{t+1} \epsilon_{t+1}^\top] = \Sigma$. In contrast, in the factor-based framework, the residual covariance matrix is constant, $\mathbf{E}_t[\epsilon_{t+1} \epsilon_{t+1}^\top] = \Sigma$, but the conditional covariance matrix of returns is a function of the time varying vector of factor loadings β_t :

$$\Sigma_{t \rightarrow t+1} = \beta_t \Omega \beta_t^\top + \Sigma \quad (42)$$

where $\beta_t = [\beta_{1,t}^\top; \beta_{2,t}^\top; \dots; \beta_{n,t}^\top]$ is the (N, K) matrix of factor exposures, and $\Omega = \mathbf{E}_t[F_{t+1} F_{t+1}^\top]$ is the (assumed time-invariant) (K, K) factor covariance matrix. Finally, ν is a free parameter used to match the Sharpe ratios generated in both environments for a myopic investor trading costlessly.

Note that the innovations in the factor exposure are driven entirely by idiosyncratic return shocks consistent with their interpretation as ‘technical’ return based factors. The AR(1) representation has the convenient representation as a weighted average of past shocks where the weights depend on the ϕ_k . This makes the interpretation as short, medium and

Table 1: **Parameters for Simulation Experiment**

The table presents the parameters estimated using the procedure described in Appendix C, and used in the simulation exercise. The three factors are designed to capture the short-term reversal, momentum, and long-term reversal effects. \hat{h}_k is the factor half-life, ϕ_k is the factor decay rate, $\hat{\lambda}_k$ the factor premium, and $\hat{\sigma}_{f,k}$ the factor volatility (all in daily terms). The final three columns give the estimated factor correlations. The factor covariance matrix Ω is equal to $\text{diag}(\sigma_f)\rho\text{diag}(\sigma_f)$.

k	Factor	\hat{h}_k	ϕ_k	$\hat{\lambda}_k$	$\hat{\sigma}_{f,k}$	$\hat{\rho}$ (correlations)		
						1	2	3
1	str	3	0.206299	-0.093482	0.406887	1	-0.366	0.167
2	mom	150	0.004610	0.001484	0.006999	-0.366	1	-0.576
3	ltr	700	0.000990	-0.000400	0.001764	0.167	-0.576	1

long-term return based factors transparent.

The value of ϕ_k is tied to its half-life (expressed in number of days) \hat{h}_k by the simple relation $\phi_k = 1 - (\frac{1}{2})^{1/\hat{h}_k}$.

3.2 Calibration of main parameters

The number of assets in our experiment is 15. Our trading horizon is 26 weeks with weekly rebalancing. Our objective is to maximize the net terminal wealth minus penalty terms for excessive risk (see Section 2.3).

We calibrate the factor mean, λ , and covariance matrix, Ω , using the Fama-French decile portfolios sorted on short-term reversal, momentum, and long term reversal. The calibration is described in Appendix C. The parameters obtained from this calibration and used in the simulation are given in Table 1.

For our simulations, we assume that both F and ϵ vectors are serially independent and normally distributed with zero mean and covariance matrix Ω and Σ , respectively. We calibrate Σ using historical daily return data on 100 largest firms measured by market capitalization from 1974 to 2012. We randomly choose 15 stocks, estimate the daily variance-covariance matrix from their returns, and calibrate Σ by converting it to its weekly counterpart. We set initial exposures to zero, i.e., $\beta_{i,0}^k = 0 \forall i, k$. Finally, ν is computed to be 0.2498 so that the Sharpe ratios generated in both models in the absence of transaction costs are equal.

The transaction cost matrix Λ is a constant multiple (η) of the conditional covariance matrix in both factor-based and characteristics environments. Since in the characteristics environment the covariance matrix of returns is constant, the cost of trading a specific dollar amount of any security is time invariant. However, from equation (42), in the factor-based environment the return covariance matrix is stochastic (because β_t is stochastic), which

results in stochastic variation in transaction costs. In our simulations, we examine three transaction costs regimes: low, medium, and high, with values of η equal to 1×10^{-7} , 2×10^{-7} and 4×10^{-7} respectively.²⁴ Finally, we set the coefficient of risk aversion to $\gamma = 10^{-8}$, which can be interpreted as a relative risk aversion of 1 for an agent with \$100 million ($= \10^8) is assets.

3.3 Approximate policies

As discussed previously, solving for the globally optimal policy in our general model is intractable due to the curse of dimensionality. Thus to assess the performance of the LGS, we compare it to alternative policies suggested in the literature or used in practice. In this section, we lay out how we implement these policies and discuss the implementation of the optimal LGS, which we label the Best Linear or BL strategy.

3.3.1 Myopic Policy (MP):

The myopic policy maximizes the single period expected return net of transaction costs and with a penalty for the (single period) portfolio variance:

$$\max_{x_t} \mathbf{E} \left[\left(x_t^\top r_{t+1} - \frac{\gamma}{2} x_t^\top \Sigma_{t \rightarrow t+1} x_t - \frac{\eta}{2} u_t^\top \Sigma_{t \rightarrow t+1} u_t \right) \right] \quad s.t. \quad eq.(4). \quad (43)$$

Substituting equation (4), which gives the dynamics for x_t , into this expression and taking the first order condition yields the closed form solution:

$$x_t^{\text{MP}} = ((\eta + \gamma) \Sigma_{t \rightarrow t+1})^{-1} (\beta_t \lambda + \eta \Sigma_{t \rightarrow t+1} (x_{t-1} \circ R_t)) \quad (44)$$

3.3.2 Myopic Policy with Transaction Cost Multiplier (MP-TC):

An issue with the myopic optimization of equation (43) is that R_{t+1} and Σ have units of time^{-1} (*i.e.*, return or return variance per unit time), but transaction costs are unitless. Thus, the myopic policy may give nonsensical solutions, particularly if the period length does not line up with the units in which expected returns and variances are measured. For this reason, it is common among practitioners to modify the myopic policy by scaling the transaction-cost term in (43) by an amortization factor τ (with units of time^{-1}).²⁵ In our

²⁴Moallemi, Saglam, and Sotiropoulos (2014) find that the average slippage for algorithmic trading firms is ≈ 5 bps/trade. If we assume that this is the slippage for a \$2 million trade on a security with a weekly volatility of 0.05, then since the dollar cost of a trade in our model is $\frac{1}{2} \eta \sigma^2 u^2$, this implies $\eta = 2 \times 10^{-7}$. We use this value of η as the multiplier for the “middle” regime. The η s for the high and low regime are a factor of 2 higher and lower, respectively.

²⁵See, *e.g.*, Grinold and Kahn (1999)

implementation, we choose this multiplier so as to maximize the unconditional performance (i.e., across all simulations) of the trading strategy. This modified myopic problem has the solution:

$$x_t^{\text{MP-TC}} = ((\tau^* \eta + \gamma) \Sigma_{t \rightarrow t+1})^{-1} (\beta_t \lambda + \tau^* \eta \Sigma_{t \rightarrow t+1} (x_{t-1}^{\text{MP-TC}} \circ R_t))$$

where τ^* is given by

$$\begin{aligned} \tau^* = \operatorname{argmax}_{\tau} \mathbb{E} \left[\left(x_t^\top r_{t+1} - \frac{\gamma}{2} x_t^\top \Sigma_{t \rightarrow t+1} x_t - \frac{\tau \eta}{2} u_t^\top \Sigma_{t \rightarrow t+1} u_t \right) \right], \\ \text{subject to } x_t = ((\tau \eta + \gamma) \Sigma_{t \rightarrow t+1})^{-1} (\beta_t \lambda + \tau \eta \Sigma_{t \rightarrow t+1} (x_{t-1} \circ R_t)). \end{aligned}$$

3.3.3 Unconditional Gârleanu & Pedersen Policy (GP-U):

As noted earlier, the L-GP solution is not optimal in either of our environments, since it requires a constant covariance matrix of price changes. However, we can implement their model following the methodology used in their empirical application (Section VI in GP). This approach relies on estimating from historical (or in our case simulated) data an unconditional covariance matrix of price changes and assuming it remains constant throughout the entire trading process. In most real world settings, and here in our simulation experiment in Section 3.4, this assumption is violated and we do not expect the GP-U method to perform well – something we verify in Section 3.4. Nonetheless we present these results to illustrate the importance of re-calibrating to the best estimate of the conditional covariance matrix of price-changes at each rebalancing date when implementing the L-GP method; we will present a “re-optimized” version of the L-GP solution – GP-R– in the next section.

Specifically, to obtain our ‘unconditional’ (GP-U) policy, we simulate data from our characteristics and factor-based framework. Then assuming an initial stock price of \$1 for each security and using percentage returns from the simulated data, we obtain the price change vector $\Delta S_{t+1} = S_{t+1} - S_t$. We then estimate the predictive ability of the each characteristic, ℓ^k from the following regression:

$$\Delta S_{i,t+1} = \ell^1 \beta_{i,t}^1 + \ell^2 \beta_{i,t}^2 + \ell^3 \beta_{i,t}^3 + \varepsilon_{i,t+1}. \quad (45)$$

We further estimate the constant covariance matrix of price changes, $\bar{\Sigma}^{pc}$, taking the unconditional covariance of price changes, that is to say, $\bar{\Sigma}^{pc} = \text{Var}(\Delta S_t)$. Since Gârleanu and Pedersen (2013) also uses an AR(1) representation for exposure dynamics, we use the same decay rate parameters (ϕ) as in our specification. For the constant transaction cost

matrix ($\bar{\Lambda}^{pc}$) used in the computation of the GP policy, we use $\eta\bar{\Sigma}^{pc}$, a constant multiple of the covariance matrix.

Using these estimated parameters, we obtain the trading policy that gives the optimal number of shares, h_t , to hold to maximize the following objective:

$$\max_{h_1, \dots, h_T} \mathbb{E} \left[\sum_{t=1}^T \left(h_t^\top \Delta S_{t+1} - \frac{\gamma}{2} h_t^\top \bar{\Sigma}^{pc} h_t - \frac{1}{2} n_t^\top \bar{\Lambda}^{pc} n_t \right) \right] \quad (46)$$

$$\text{subject to } h_t = h_{t-1} + n_t \quad (47)$$

The optimal solution to this problem is derived in Gârleanu and Pedersen (2013), and is given by

$$h_t = (\bar{\Lambda}^{pc} + \gamma\bar{\Sigma}^{pc} + A_{xx}^t)^{-1} (\bar{\Lambda}^{pc} h_{t-1} + (C + A_{xf}^t (I - \Phi)) \beta_t^{\text{st}})$$

where $\beta_t^{\text{st}} = [\beta_{:,t}^1; \dots; \beta_{:,t}^3]$ is the stacked vector of factor exposures, $C = \ell^\top \otimes I_{N \times N}$ and $\Phi = \text{diag}(\phi \otimes I_{N \times 1})$ and A_{xx}^{t-1} and A_{xf}^{t-1} satisfy the following recursions,

$$\begin{aligned} A_{xx}^{t-1} &= -\bar{\Lambda}^{pc} (\bar{\Lambda}^{pc} + \gamma\bar{\Sigma}^{pc} + A_{xx}^t)^{-1} \bar{\Lambda}^{pc} + \bar{\Lambda}^{pc}, \\ A_{xf}^{t-1} &= \bar{\Lambda}^{pc} (\bar{\Lambda}^{pc} + \gamma\bar{\Sigma}^{pc} + A_{xx}^t)^{-1} (A_{xf}^t (I - \Phi) + C), \end{aligned}$$

with $A_{xx}^T = \mathbf{0}$ and $A_{xf}^T = \mathbf{0}$.

3.3.4 Re-optimized Gârleanu & Pedersen Policy (GP-R):

As noted above, the GP-U policy described above is not likely to yield a reasonable solution, as the GP-U solution requires a constant covariance matrix of price-changes. Thus, we can improve on the performance of the GP-U method by re-optimizing each period. Specifically, at each time t we calculate a new (and accurate) price-change covariance matrix based on the time t conditional (return) covariance matrix and the level of prices.

Note that this solution is numerically intensive, as it requires re-calculating the Riccati recursions at each time t using this updated estimate of the covariance matrix of price changes. Note also that even this re-optimized solution ignores the future risk dynamics: the GP-R method assumes that the price-change covariance matrix will remain constant from time t forward.²⁶

We set the conditional covariance matrix of price changes to be $\bar{\Sigma}_t^{pc} = \text{diag}(S_t) \Sigma \text{diag}(S_t)$

²⁶This approach is very similar to the ‘anticipated utility’ concept of Kreps (1998) and Cogley and Sargent (2008).

in the characteristics based model and $\bar{\Sigma}_t^{pc} = \text{diag}(S_t) (\beta_t \Omega \beta_t^\top + \Sigma) \text{diag}(S_t)$ in the factor-based model. The transaction cost matrix, $\bar{\Lambda}_t^{pc}$, is also time-varying and set to $\eta \bar{\Sigma}_t^{pc}$. With this parameterization, the GP-R policy has the following form:

$$h_t = (\bar{\Lambda}_t^{pc} + \gamma \bar{\Sigma}_t^{pc} + A_{xx}^{t,t})^{-1} (\bar{\Lambda}_t^{pc} h_{t-1} + (C_t + A_{xf}^{t,t} (I - \Phi)) \beta_t^{\text{st}})$$

where $\beta_t^{\text{st}} = [\beta_{:,t}^1; \dots; \beta_{:,t}^3]$ is the stacked vector of factor exposures, $C_t = \lambda^\top \otimes \text{diag}(S_t)$ and $\Phi = \text{diag}(\phi \otimes I_{N \times 1})$ and $A_{xx}^{t,t}$ and $A_{xf}^{t,t}$ is the solution of the following recursions ($\forall t < n \leq T$),

$$\begin{aligned} A_{xx}^{t,n-1} &= -\bar{\Lambda}_t^{pc} (\bar{\Lambda}_t^{pc} + \gamma \bar{\Sigma}_t^{pc} + A_{xx}^{t,n})^{-1} \bar{\Lambda}_t^{pc} + \bar{\Lambda}_t^{pc}, \\ A_{xf}^{t,n-1} &= \bar{\Lambda}_t^{pc} (\bar{\Lambda}_t^{pc} + \gamma \bar{\Sigma}_t^{pc} + A_{xx}^{t,n})^{-1} (A_{xf}^{t,n} (I - \Phi) + C_t), \end{aligned}$$

with $A_{xx}^{t,T} = \mathbf{0}$ and $A_{xf}^{t,T} = \mathbf{0}$. Here, we have double time superscripts in $A_{xx}^{t,n}$ and $A_{xf}^{t,n}$ to underscore that we are re-solving the Riccati recursion (in n) at every time step (t).

3.3.5 Best Linear Policy (BL):

We define the relevant stock exposure variables for each security to be the stock specific myopic portfolio holdings and a constant term, i.e., $B_{i,t} = [x_{i,t}^{\text{MP}}; 1]$. We then follow the methodology developed in Section 2 to determine the optimal LGS satisfying our nonlinear state evolution:

$$\begin{aligned} u_t^{\text{BL}} &= \mathcal{B}_t^\top \pi_t^* \\ x_t^{\text{BL}} &= \mathcal{B}_t^\top \theta_t^* \end{aligned}$$

where as before \mathcal{B}_t is constructed from the return-scaled exposures $\mathcal{B}_{i,s \rightarrow t} = B_{i,s} R_{s \rightarrow t}$, where π_t^* and θ_t^* solve:

$$\begin{aligned} \max_{\pi_1, \dots, \pi_T} \quad & \sum_{t=1}^T \theta_t^\top \bar{m}_t - \frac{1}{2} \pi_t^\top \bar{\Lambda}_t \pi_t - \frac{\gamma}{2} \theta_t^\top \bar{\Sigma}_t \theta_t \\ \text{subject to} \quad & \theta_t = \theta_{t-1}^0 + \pi_t \end{aligned}$$

3.4 Simulation Results

We now discuss the performance of the approximate policies and the best linear (LGS) policies in the simulation for both the factor- and the characteristics-models, for low, medium and high transaction costs. We also provide performance statistics for a zero transaction-cost setting as a benchmark case.

Table 2: **Policy performance: characteristics environment.**

This table summarizes the performance of each policy in the characteristics environment (no common factors) for four different levels of transaction costs. For each policy, we report the average across the 10,000 runs of: the objective function, terminal wealth and transaction costs paid (in $\$ \times 10^5$); the information ratio using the myopic policy as a benchmark, and the annualized Sharpe Ratio. The final column reports the difference between the BL and the GP-R strategy results and the associated standard errors.

	MP	MP-TC	GP-U	GP-R	BL	BL-GP-R
Zero Transaction Costs						
Avg Objective	2739.75	2739.75	1704.41	2739.75	2739.65	-0.10
Std Err	22.59	22.59	8.98	22.59	22.59	0.34
Avg Wealth	5487.37	5487.37	2159.97	5487.37	5488.35	0.98
Std Err	22.01	22.01	8.87	22.01	22.01	0.34
TC	0.00	0.00	0.00	0.00	0.00	0.00
IR	NA	NA	-3.47	0.05	0.04	NA
SR	3.53	3.53	3.44	3.53	3.53	0.04
Low Transaction Costs ($\eta = 1 \times 10^{-7}$)						
Avg Objective	254.34	255.08	207.59	327.39	329.55	2.16
Std Err	5.37	5.11	1.46	3.46	3.47	0.56
Avg Wealth	427.69	412.19	221.89	404.08	407.30	3.21
Std Err	5.06	4.82	1.44	3.38	3.39	0.53
TC	226.71	194.49	40.65	252.25	249.86	-2.39
IR	NA	-0.79	-0.74	-0.12	-0.10	NA
SR	1.20	1.21	2.19	1.69	1.70	0.09
Medium Transaction Costs ($\eta = 2 \times 10^{-7}$)						
Avg Objective	147.83	147.92	119.73	187.74	190.04	2.30
Std Err	3.39	3.25	0.95	2.20	2.24	0.55
Avg Wealth	219.03	213.83	125.82	218.98	222.82	3.84
Std Err	3.24	3.12	0.94	2.16	2.19	0.53
TC	123.79	111.54	25.55	157.38	155.94	-1.44
IR	NA	-0.57	-0.52	0.00	0.03	NA
SR	0.96	0.97	1.90	1.43	1.44	0.10
High Transaction Costs ($\eta = 4 \times 10^{-7}$)						
Avg Objective	85.28	85.33	66.26	103.22	105.46	2.24
Std Err	1.96	1.99	0.60	1.34	1.42	0.52
Avg Wealth	110.00	110.70	68.67	114.92	118.77	3.85
Std Err	1.91	1.93	0.59	1.32	1.39	0.51
TC	64.93	66.96	15.07	92.02	91.34	-0.69
IR	NA	0.40	-0.40	0.06	0.10	NA
SR	0.82	0.81	1.65	1.23	1.20	0.11

3.4.1 Characteristics Model Simulation Results

The upper panel of Table 2 shows the results when the simulated returns are generated according to the characteristics model in equation (41), (*i.e.*, when there are no common factors, and thus all return variance is idiosyncratic), and when transaction costs are zero. Because there are no transaction costs, the myopic policy is optimal. Indeed, all policies that nest the unconstrained myopic policy (*i.e.*, MP-TC, GP-R, and BL, that is all except for GP-U) achieve the same objective function and corresponding (high) Sharpe ratio of 3.53. GP-U is not able to approach the first best strategy even in the no-transaction cost case, because it ignores the dynamics in the covariance matrix entirely. While its Sharpe ratio is also very good in the absence of transaction costs, the difference in performance relative to the other strategies that nest the conditional myopic strategy is significant (this is still more apparent when comparing the difference in average objective functions).

Note that BL nests the myopic strategy because we use as one of the stock exposures the myopic strategy holdings. This illustrates the necessity of choosing a large enough set of exposures for the LGS to be able to approach the first best.

The second panel of Table 2 shows that even when transaction costs are relatively low, the dynamic strategies outperform the myopic strategies. The objective functions and the Sharpe ratios of GP-R and BL are significantly higher than the corresponding values achieved with either the MP or MP-TC strategies. However, GP-U continues to underperform even the myopic model in terms of average objective function (as well as in terms of average wealth), because the model ignores any dynamics in the covariance matrix of price changes. We note that the Sharpe ratio of GP-U is actually high, which shows that the Sharpe ratio is a misleading performance measure in the presence of transaction costs. GP-U achieves that higher (net-of-t-costs) Sharpe ratio because it trades very little (transaction costs are less than one sixth that of the BL strategy for example) thus generating little average wealth (net of t-costs) and little volatility. The correct comparison for the different trading strategies is the average objective function (the expression of which is the same for all trading strategies considered).

The third and fourth panels show that, as t-costs increase, the performance of all strategies fall, as would be expected. However, BL and the GP-R continue to outperform the myopic strategies. BL outperforms GP-R only very slightly, but the differences and standard-errors reported in the final columns show that the differences in the objective function are statistically significant, as would be expected. The small performance difference between GP-R and BL is likely due to the fact that the log-normal return dynamics we simulate in the charac-

teristics environment (which assumes a constant covariance matrix returns) do not conform to the assumed normal return dynamics assumed in the GP-R solution (which assumes a constant covariance matrix of price changes).

3.4.2 Factor Model Simulation Results

We now turn to the simulations which are run for the factor model environment in which cross-sectional variation in expected returns is linked to common factor loadings (equation (40)). The upper panel of Table 3 presents the set of strategy performance measures when trading costs are zero. As before, with zero transaction costs the myopic strategy is optimal, and thus all strategies (except GP-U) achieve the same average objective function, since they all nest the myopic strategy.

The lower three panels present performance measures for the low, medium, and high t-cost environments. In each environment, we find that BL achieves the highest objective among the strategies. BL now more significantly outperforms GP-R. Interestingly, while the difference in the average objective functions (in the last column) declines with increasing t-costs, the percentage difference increases, and the BL objective function is 14% higher than the average GP-R in the high t-cost environment.²⁷ Also, GP-R is now much closer in performance to MP-TC, and in fact in the low transaction cost underperforms MP-TC. Recall that the t-cost multiplier for the MP-TC strategy is chosen so as to maximize the objective function across all simulations.

The underperformance of GP-R in the factor-based environment is likely a result of the fact that GP-R strategy does not take into account information on the expected future dynamics of that covariance matrix, and the corresponding expected future transaction-cost dynamics.²⁸ In contrast, BL takes into account expected future covariance and transaction-cost dynamics. To confirm this intuition, in the next section we examine impulse response functions for the various strategies.

3.5 Policy responses to return shocks

In this section, we construct impulse response functions for the BL, GP-R, MP and MP-TC policies described in Section 3.3. We do this for both the factor (equation (40)) and

²⁷Note that the Sharpe and information ratios for the strategies do not always line up with the average objective functions. The reason is that these ratios are not the objective function that is optimized and hence can be a misleading performance criterion. For example, in the medium transaction cost case MP achieves the second-lowest objective (i.e., very low average wealth net of t-cost and risk) but has the highest Sharpe ratio.

²⁸Recall that, GP-R is re-optimized, so it does use the correct conditional covariance matrix at each step. However, it cannot take into account information about future changes in the covariance structure.

Table 3: **Policy performance: common factor environment**

This table summarizes the performance of each policy in the factor model environment for four different levels of transaction costs. For each policy, we report the average across the 10,000 runs of: the objective function, terminal wealth and transaction costs paid (in $\$ \times 10^5$); the information ratio using the myopic policy as a benchmark, and the annualized Sharpe Ratio. The final column reports the difference between the BL and the GP-R strategy results and the associated standard errors.

	MP	MP-TC	GP-U	GP-R	BL	BL-GP-R
Zero Transaction Costs						
Avg Objective	3001.31	3001.31	-98689.21	3001.31	3000.60	-0.71
Std Err	24.39	24.39	391.02	24.39	24.40	0.32
Avg Wealth	6007.09	6007.09	37659.88	6007.09	6006.76	-0.34
Std Err	24.37	24.37	163.16	24.37	24.38	0.34
TC	0.00	0.00	0.00	0.00	0.00	0.00
IR	NA	NA	3.17	0.59	-0.01	NA
SR	3.49	3.49	3.26	3.49	3.49	-0.01
Low Transaction Costs ($\eta = 1 \times 10^{-7}$)						
Avg Objective	429.61	452.37	-8036.32	447.94	490.57	42.63
Std Err	3.18	4.02	37.02	3.21	4.19	1.78
Avg Wealth	489.25	546.45	-5756.90	509.32	593.44	84.12
Std Err	3.16	4.00	31.14	3.18	4.14	1.72
TC	234.41	385.42	10297.69	236.30	385.78	149.48
IR	NA	0.94	-3.01	0.17	0.72	NA
SR	2.19	1.93	-2.61	2.26	2.03	0.69
Medium Transaction Costs ($\eta = 2 \times 10^{-7}$)						
Avg Objective	229.21	242.62	-5002.27	251.86	279.06	27.20
Std Err	1.74	2.23	22.25	2.09	2.62	1.56
Avg Wealth	247.82	272.87	-4200.74	280.71	323.65	42.94
Std Err	1.74	2.23	20.05	2.07	2.59	1.53
TC	132.27	223.98	6812.61	158.05	232.64	74.59
IR	NA	0.71	-3.28	0.31	0.59	NA
SR	2.02	1.73	-2.96	1.91	1.77	0.40
High Transaction Costs ($\eta = 4 \times 10^{-7}$)						
Avg Objective	118.48	125.57	-2990.79	136.65	155.83	19.17
Std Err	0.92	1.19	12.92	1.44	1.79	1.42
Avg Wealth	123.88	134.51	-2727.93	152.10	179.26	27.16
Std Err	0.92	1.19	12.15	1.44	1.77	1.41
TC	70.44	120.45	4169.41	98.60	130.81	32.21
IR	NA	0.55	-3.44	0.31	0.50	NA
SR	1.90	1.60	-3.18	1.49	1.44	0.27

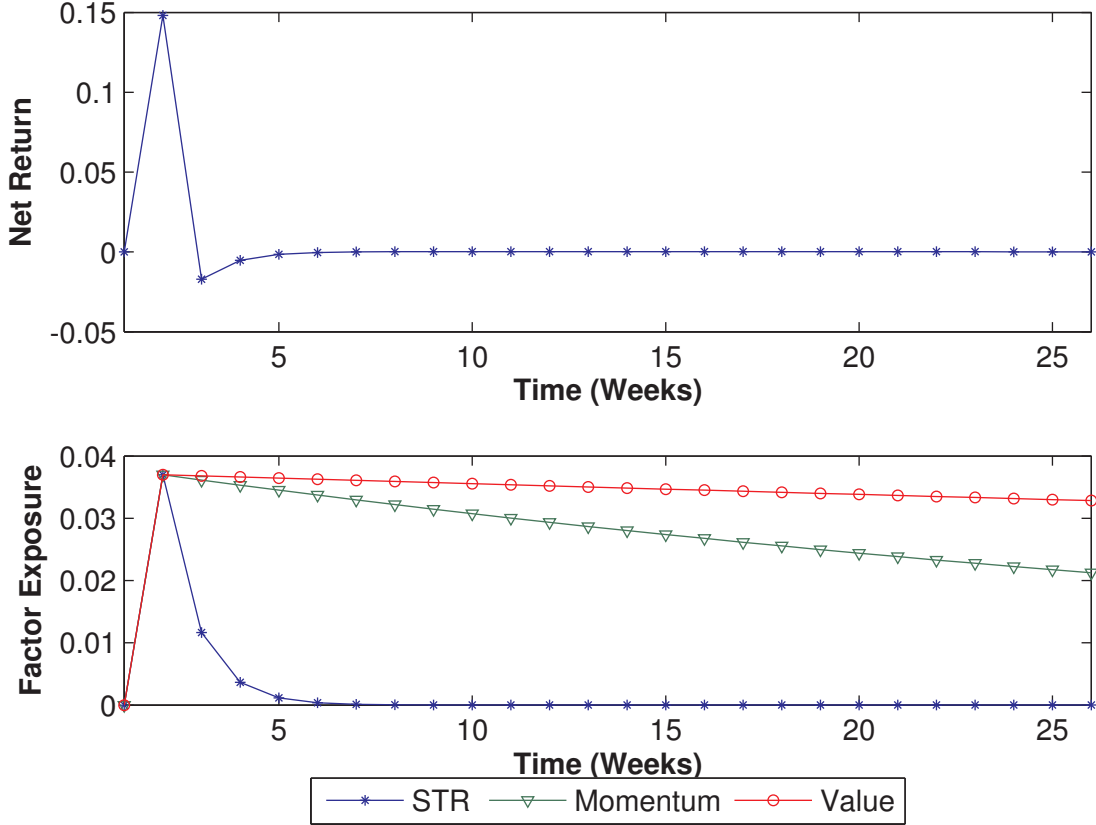


Figure 1: **Security returns and factor exposures – characteristics environment**
The top panel plots the realized returns for a security in response to a 2σ return shock in week 2. The bottom panel plots the factor exposures for reversal, momentum and value for security following this time 2 return shock. This plot is for the characteristics environment.

characteristic (equation (41)) environments described in Section 3.4. Our analysis provides some insights into the differences in performance uncovered in our analysis in Section 3.4.

The basic environment is the same as in the preceding section. We begin by setting the time 1 positions and exposures for each security equal to their long run mean of zero: $x_{j,1} = \beta_{j,rev,1} = \beta_{j,mom,1} = \beta_{j,value,1} = 0 \forall j$.

We further constrain the residuals for all securities over week 1 to be zero. In week 2, we “shock” the idiosyncratic return of security i with a positive 2-standard-deviation shock, *i.e.*, $\epsilon_{i,2} = 2\sigma_i$, but set the idiosyncratic shocks for all other assets to zero ($\epsilon_{j,2} = 0 \forall i \neq j$). From week 3 to week 26, all future shocks are set to zero so that the path of realized returns is equal to the path of expected returns.

3.5.1 Characteristics model results

The upper panel of Figure 1 plots the realized returns of security i for this experiment. The positive return at time 2 is the shock itself. As a result of this shock’s effect on the return generating process, the expected return at time 3 is negative, but then decays quickly toward

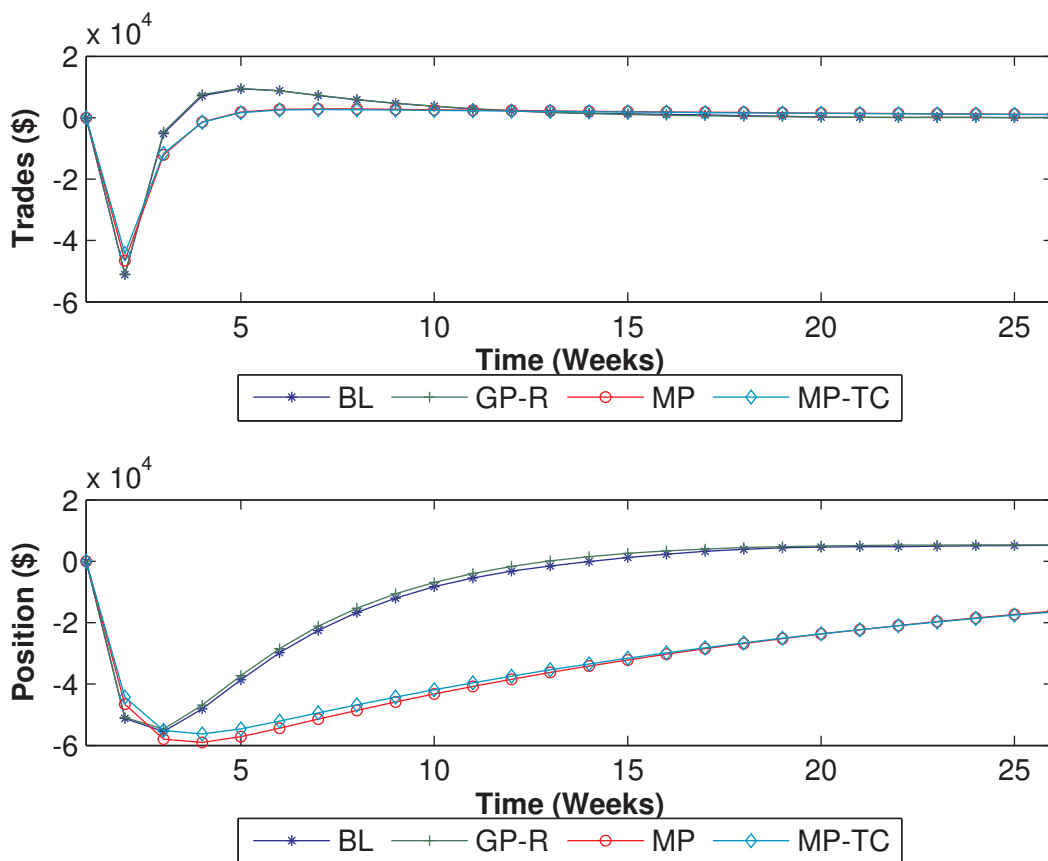


Figure 2: **Trades and positions – characteristics environment**

The upper panel plots the dollar size of trades, and the lower panel the dollar size of positions in security i for various trading policies, following a two standard deviation idiosyncratic volatility shock in week 2. The characteristic-based return generating process is used (equation (41)). The transaction costs parameter corresponds to “medium” (*i.e.*, $\eta = 2 \times 10^{-7}$)

zero, and eventually becomes very slightly positive – something that is difficult to see in this plot.

This pattern of expected returns is a result of the interplay between reversal, momentum and value. The lower panel of this Figure illustrates how this comes about. This plot shows the security i exposures to the three factors. At the end of week 2, all exposures are equal to approximately one-fourth of the idiosyncratic shock (recall that $\nu = 0.2498$ per equation (41)), but then decay at very different rates. In the determination of the expected return, the reversal effect dominates from week 3 to week 11 resulting in a negative expected return for security i . After week 11, the positive (but much smaller) premium for momentum generates a positive expected return, but because the premium for momentum is about two orders of magnitude smaller than that for reversal – as seen in Table 1 – the momentum effect is difficult to see in the plot. Of course, because momentum is much longer-lived than reversal, the cumulative effects are more comparable.

Figure 2 plots the dollar trades and corresponding positions in security i for the four policies in the characteristics-based setting. The transaction cost parameter η is set to 2×10^{-7} , corresponding to the “medium” cost regime. Consistent with the strategy results discussed in Section 3.4, the trades and positions of the two forward-looking strategies (GP-R and BL) are nearly identical in this characteristics environment, as are the trades and positions of the two “myopic” strategies (MP and MP-TC).

A comparison of trades/positions of myopic and forward-looking strategies is instructive in understanding the performance differential evident in Table 2. While the myopic strategies trades into the position at about the same rate as the BL strategies at time 2, the forward-looking strategies trade out of the position much more quickly. This is because the myopic policy trade is based only on the expected returns and covariance at any point in time, and not on how quickly the expected return and covariances are expected to change. In contrast, both GP-R and BL optimally incorporate the expected return dynamics of the security in how they trade at every step. Even the MP-TC which has one additional free parameter that helps account for the ‘expected return horizon’ cannot approach the optimal trading strategy when there are several factors with different decay rate driving expected returns.

3.5.2 Factor model results

We now examine the strategy trades when the return generating process for security returns is a factor model (equation (40)) rather than a characteristics model. Recall that in this environment risk dynamics are far more complex, in that a security’s covariance with risk factors changes as its factor loadings change. What we will see is that, since the BL method anticipates the changes in risk (and transactions costs) while the GP-R method does not, the BL outperforms the GP-R method by far larger amounts. In general, both of these method outperform the myopic policies, which anticipate neither future changes in expected returns nor future risk changes.

The upper panel of Figure 3 plots the path of realized returns of security i and the lower panel of Figure 3 shows the path of the factor exposures. The main difference relative to Figure 1 is that, at the end of week 2, all three factor exposures are equal to the value of the idiosyncratic shock (per equation (40)) which leads to four-times the magnitude of the expected return when compared to characteristics-based model (but there is also more risk since the return variance increases with factor exposures). The sign and pattern of realized and expected returns are the same as in the previous case.

Figure 4 plots security i ’s trades and positions for the four policies in the factor-based environment. Comparing this to Figure 2, we see that there are now substantial differences

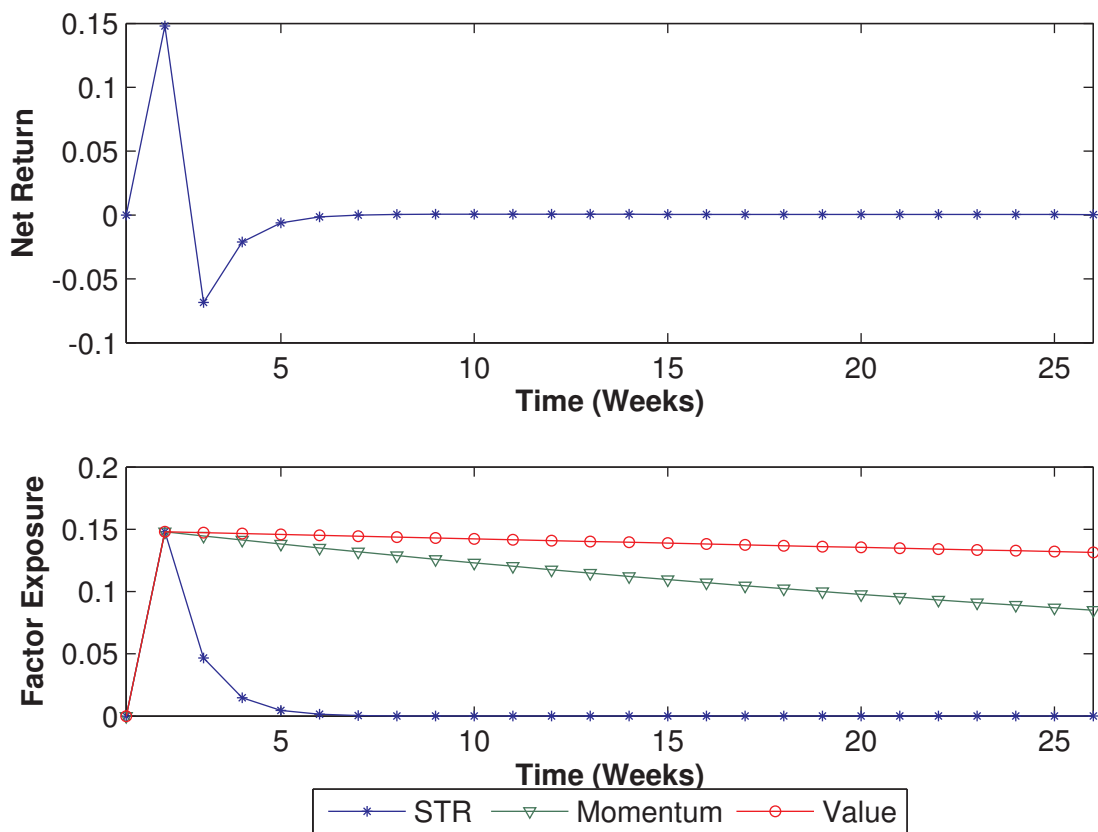


Figure 3: **Security returns and factor exposures – factor model environment**

The top panel plots the realized returns for a security in response to a 2σ return shock in week 2. The bottom panel plots the factor exposures for reversal, momentum and value for security following this time 2 return shock. This plot is for the factor environment.

between the BL and GP-R trades immediately following the shock. BL trades more aggressively and builds larger short position in the first few weeks (due to short-term reversal) and over time builds up a larger positive position in security i (due to momentum). This more aggressive trading allows BL to eventually outperform GP-R.

In the last section, we saw that when returns were generated by a characteristics model, both BL and GP-R outperformed the myopic strategies. The reason was that the BL and GP-R trades both anticipated future changes in expected returns, while the myopic strategies did not. In the factor-model setting, the covariance matrix and expected returns are both affected by factor shocks. The BL method takes into account the future dynamics associated with this changing covariance matrix. In contrast the GP-R method cannot, as it implicitly assumes that the price-change covariance matrix will not change going forward – an assumption that is clearly violated in the factor environment. This is why the week 2 trade in response to the shock is smaller for GP-R than for BL: the GP-R methodology implicitly assumes that the high risk for security i at time 2 will continue indefinitely. In contrast, the BL trade incorporates the fact that, as the factor loading decays over time, risk will decrease and

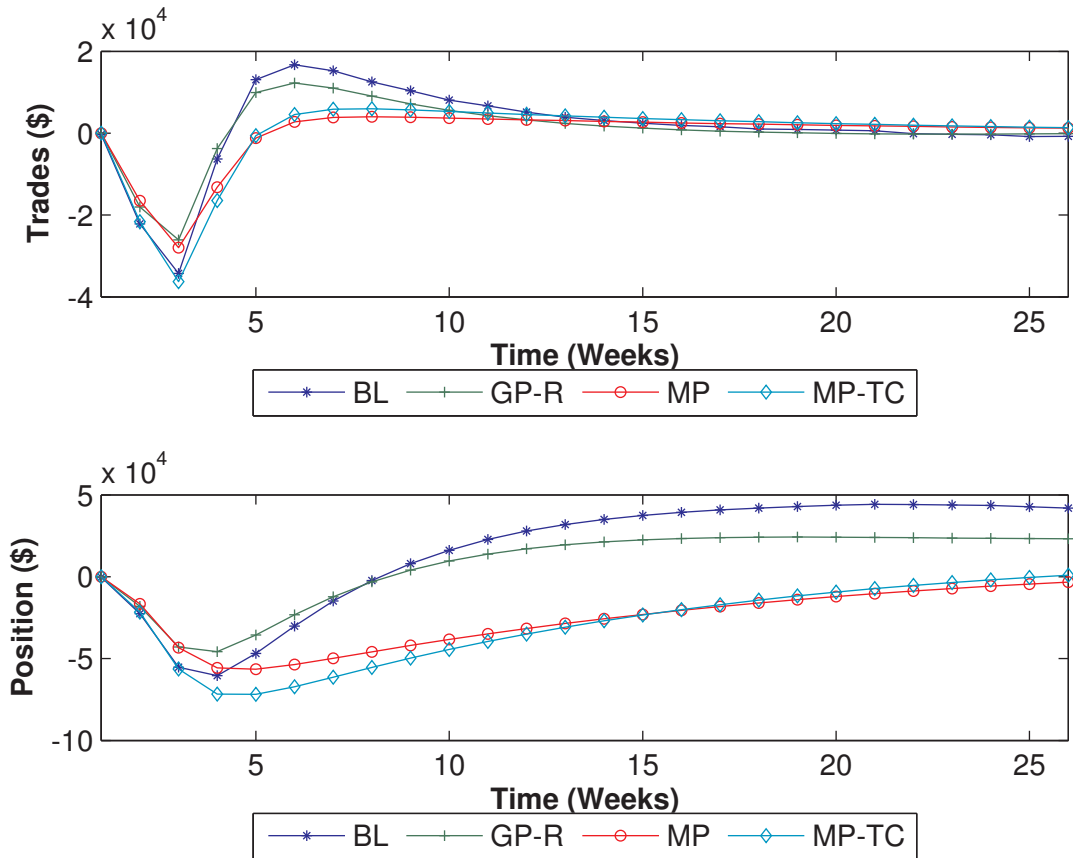


Figure 4: **Trades and positions – factor model environment**

The upper panel plots the dollar size of trades, and the lower panel the dollar size of positions in security i for various trading policies, following a two standard deviation idiosyncratic volatility shock in week 2. The factor-model-based return generating process is used (equation (40)), and the transaction costs parameter corresponds to “medium” (*i.e.*, $\eta = 2 \times 10^{-7}$)

therefore trades more aggressively.

The analysis of this section shows that, when the covariance matrix or transaction costs are highly dynamic, it is important to use a rule that calculates optimal trades and positions taking into account the forecastable future changes in risk or transaction costs.

4 Applying the LGS methodology to US equities

Section 3 examined the performance of the LGS method relative to existing method with simulated data as a way of characterizing where large performance gains were likely.

In this section we investigate the performance of the LGS when trading the 100 largest US equities over the 1930:01-2014:03 time period. In contrast to the simulation approach presented above here explore the performance of the LGS using real world data, meaning that we need to first, estimate the data generating process from empirical data.

In this analysis we trade a zero-investment portfolio where the investment universe is the

100 largest US common stocks. We develop a trading rule to exploit return predictability that arises from the short-term-reversal effect, price-momentum, and long-term reversal.²⁹ It also relies on an estimate of the covariance structure of the returns of these 100 securities.

We divide the 84.25 year sample period into five-year sample-periods and one final 4.25 year sample-period, for a total of 17 samples. We assume that our agent begins each period with the objective of maximizing wealth (net of transaction costs) at the end of that 5-year period, minus a penalty for variance, as specified earlier.

At the beginning of each 5-year (60 month) period, our portfolio has a value of \$0, and has an asset weight vector which is all zeros. So, for example, our first period starts on the last trading day of 1929, at which point the agent trades into the optimal portfolio based on the trading rule. At the end of each month, the agent observes the performance of each of the 100 securities over that month and, based on the revised portfolio holding and updated return forecasts and transaction cost estimates, trades into the new portfolio based on the trading rule. This pattern continues until the end of the last month in each 60-month period, at which point we evaluate the performance of the portfolio over that 60-month period, which allows us to compare the performance of the trading rules over the seventeen 5-year samples.

Our setting is not entirely realistic in that we assume our agent's information set contains the realized 'in-sample' covariance matrix, the coefficients from a projection of monthly residual returns onto lagged monthly returns, and the firm's market betas at the start of each period. We endow the agent with this information, as it allows us to abstract away from the question of how best to forecast future returns and covariances, and concentrate on the relevant question for this paper, which is how one would construct an optimal portfolio given these forecasts.

4.1 Data and trading setup

We proceed as follows. From CRSP, we extract monthly returns for all firms listed on the NYSE, AMEX or NASDAQ. We exclude ADRs, etc., by requiring a share code of 10 or 11.

We perform our analysis on this set of firms one five-year period at a time, starting with 1930:01-1934:12, and ending with the 2010:01-2014:03 period. For each 5-year period, we select the firms which have no missing returns from 30 months prior to the start of the period, up through the end of the period. Of these firms, we select the 100 largest, measured by equity market capitalization at the start of the five-year period.³⁰

²⁹See, respectively, Jegadeesh (1990) and Lehmann (1990), Jegadeesh and Titman (1993), and DeBondt and Thaler (1985).

³⁰Again, we note that the restriction that we have valid returns over each month of the coming 5-year period means that this is not an implementable strategy.

In each 5 year period, for each of the 100 firms, we calculate market betas and residual returns. Market betas come from a regression of monthly excess returns of each firm on the returns of the CRSP value-weighted market excess return, *i.e.*,

$$\tilde{R}_{i,t} = \alpha_i + \beta_i \tilde{R}_{m,t} + \tilde{r}_{i,t}$$

where $\tilde{R}_{i,t}$ and $\tilde{R}_{m,t}$ are, respectively, firm i 's and the market's return net of the one month T-Bill rate, and the regression residual $r_{i,t}$ is firm i 's residual return.³¹

In each 5-year period, we then run Fama and MacBeth (1973) regressions for the 100 firms. That is, for each month t in the 5-year period, we perform an OLS cross-sectional regression with the set of 100 month- t residual returns as the dependent variables, and the corresponding residual returns from month $t-\tau$, for $\tau = 1, \dots, 30$ as the independent variables:

$$r_{i,t} = \lambda_{0,t} + \sum_{\tau=1}^{30} \lambda_{\tau,t} r_{i,t-\tau} + \epsilon_{i,t} \quad (48)$$

We then average the estimated coefficients $\hat{\lambda}_{\tau,t}$ from the 60 monthly cross-sectional regressions in the five-year period to obtain our estimates of λ_{τ} for this period:

$$\lambda_{\tau} \equiv \frac{1}{\#\mathcal{T}} \sum_{t \in \mathcal{T}} \hat{\lambda}_{\tau,t},$$

where \mathcal{T} is the set of months in this 5-year sample period. We further define the 30×1 vector $\boldsymbol{\lambda}$ as the stacked λ_{τ} s: $\boldsymbol{\lambda} = [\lambda_1, \lambda_2, \dots, \lambda_{30}]^{\top}$.

4.1.1 Return generating process specification

For comparison with our simulation analysis of the previous section, we can rewrite the return generating process (RGP) as follows:

$$r_{i,t+1} = \boldsymbol{\beta}_{i,t}^{\top} \boldsymbol{\lambda} + \epsilon_{i,t+1} \quad (49)$$

$$\boldsymbol{\beta}_{i,t+1} = \mathbf{A} \boldsymbol{\beta}_{i,t} + \mathbf{B} \epsilon_{i,t+1} \quad (50)$$

³¹The series of one-month t-bill rates comes from Ken French's data library at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

Here $\beta_{i,t}$ is (30×1) . \mathbf{A} is (30×30) and \mathbf{B} is (30×1) , and are given by:

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 & 0 \\ 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \end{bmatrix} \quad \mathbf{B} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

In this specification $\beta_{i,t}$ is the vector of lagged unexpected returns for firm i , and the matrix \mathbf{A} acts as a shift operator.

The other element of the return generating process that we need to specify is the residual covariance matrix $\Sigma_{t \rightarrow t+1} = \mathbf{E}_t[\epsilon_{t+1}\epsilon_{t+1}^\top]$. In our model of the return generating process we assume that this covariance matrix is time-invariant over each 5-year period, and is equal to the realized covariance matrix, but we shrink each of the off diagonal elements of the covariance matrix by a factor of $(1/3)$ to ensure that the matrix is non-singular.

4.2 Model performance with zero transaction costs

To assess how well this simple specification captures the return generating process, we analyze the performance of a mean-variance-efficient portfolio based on this RGP specification. Using the expected return estimates from equation (49) above and the covariance matrix, calculated as described in Section 4.1.1 above, we generate a portfolio with weights:

$$\mathbf{w}_t^{MVE} = \Sigma^{-1} (\beta_t^\top \boldsymbol{\lambda}).$$

where column i of β_t is $\beta_{i,t}$, as discussed in Section 4.1.1 above, and $(\beta_t^\top \boldsymbol{\lambda})$ is therefore the (100×1) vector of expected residual returns over period $t \rightarrow t + 1$. The returns to this portfolio in this period are then just:

$$r_{t+1}^{MVE} = (\mathbf{w}_t^{MVE})^\top \tilde{r}_{t+1}.$$

Figure 5 plots the monthly strategy returns to this strategy over the full period from 1930:01-2014:03, where the returns are normalized have 19% annualized volatility. The figure shows that the strategy returns are, on average, well above zero. The annualized Sharpe ratios in the five-year periods range from a minimum of 1.135 to a maximum of 4.395. The full period annualized Sharpe ratio is 2.525.

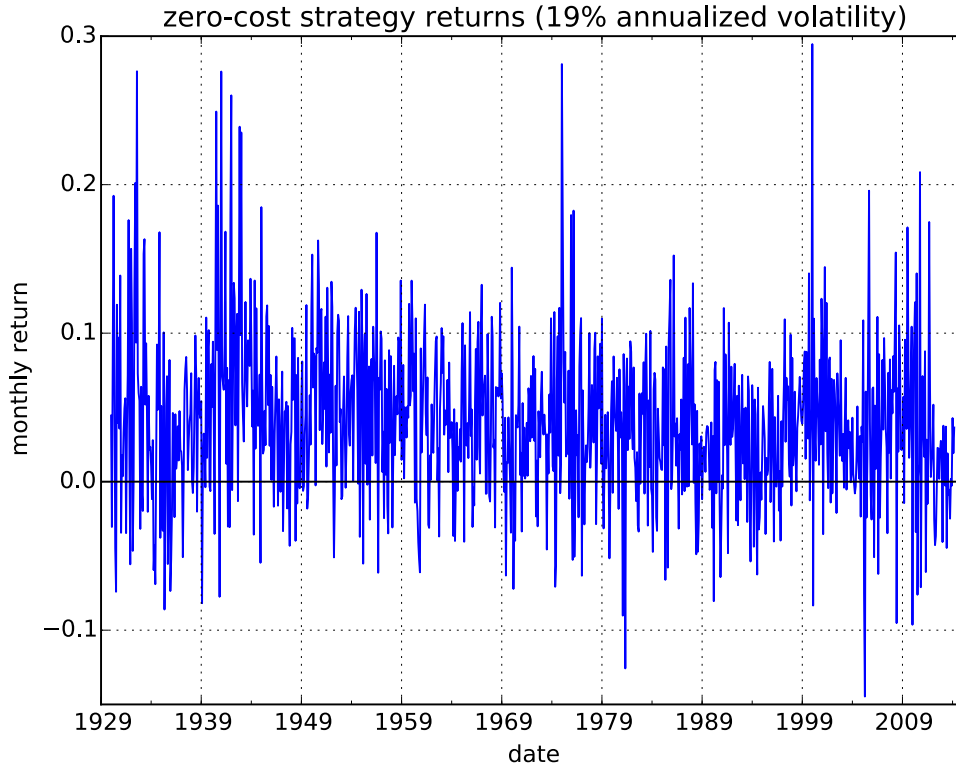


Figure 5: **Zero t-cost strategy monthly returns**

This figure plots the monthly returns to the zero t-cost strategy described in Section 4. Returns are scaled to the an annualized *ex-ante* volatility of 19%.

4.3 An LGS-based methodology applied to real-world data

In this section we document the construction of an LGS-based methodology developed in Section 2 using the real-world data. We also utilize the fixed-lag policy implementation described in detail in Section 2.9.

As discussed in Section 2.3, we assume that the investor's objective function is to maximize his expected terminal wealth net of t-costs and net of a quadratic risk penalty. Since the conditional covariance is not time-varying for our model, the objective function can be cast as follows:

$$\max \sum_{t=1}^T \mathbb{E} \left[x_t^\top r_{t+1} - \frac{\gamma}{2} x_t^\top \Sigma x_t - \frac{\eta}{2} u_t^\top \Sigma u_t \right],$$

where $\Sigma = \text{Var}(r_{t+1})$ from the dynamics of the security returns (using shrunk estimates). For simplicity, we assume that the transaction cost matrix is a constant multiple of the covariance matrix of the returns as in the simulation experiment.

We calibrate η using the same methodology described in the simulation experiment. We present results for three (low, medium and high) transaction cost regimes. We assume that

the coefficient of risk aversion γ equals 10^{-9} . We also set the transaction cost multiplier (η) as in our simulation experiment so that the average slippage values in the three transaction cost regimes correspond to 2.5 bps, 5 bps and 10 bps respectively. Using monthly volatility of $\sigma_\epsilon = 0.1$, this yields an η roughly around 2.5×10^{-9} , 5×10^{-9} and 10×10^{-9} for the low, medium and high transaction cost regimes respectively.³²

We compare the gains from trading according to a myopic policy with transaction cost multiplier (MP-TC) and LGS-based fixed-lag Best Linear (BL) policy using the methodology developed in Section 2.9. We evaluate the performance of the policies in each of the 17 five-year trading horizons from 1930 to 2014.

We use a similar approach undertaken in Section 3.3 to compute both trading policies. Let x_t^{MP} be the vector of dollar positions that the myopic policy chooses in each asset. Then,

$$x_t^{\text{MP}} = ((\eta + \gamma)\Sigma)^{-1} (\beta_t \lambda + \eta \Sigma (x_{t-1}^{\text{MP}} \circ R_t)).$$

We will then choose an optimal multiplier τ^* so as to maximize the unconditional performance (i.e., across all simulations) of the trading strategy. Formally, this modified myopic strategy has a solution:

$$x_t^{\text{MP-TC}} = ((\tau^* \eta + \gamma) \Sigma)^{-1} (\beta_t \lambda + \tau^* \eta \Sigma (x_{t-1}^{\text{MP-TC}} \circ R_t))$$

where τ^* is given by

$$\begin{aligned} \tau^* = \operatorname{argmax}_{\tau} \mathbb{E} \left[\left(x_t^\top r_{t+1} - \frac{\gamma}{2} x_t^\top \Sigma x_t - \frac{\tau \eta}{2} u_t^\top \Sigma u_t \right) \right], \\ \text{subject to } x_t = ((\tau \eta + \gamma) \Sigma)^{-1} (\beta_t \lambda + \tau \eta \Sigma (x_{t-1}^{\text{MP-TC}} \circ R_t)). \end{aligned}$$

We will compare MP-TC with a fixed-lag best linear policy that uses at most two lags in security exposures. Therefore, our position and trade vectors will take the following form:

$$x_{i,t} = \theta_{i,t,t}^\top B_{i,t} \text{ and } u_{i,t} = \pi_{i,t-1,t}^\top B_{i,t-1 \rightarrow t} + \pi_{i,t,t}^\top B_{i,t}$$

We define the relevant stock exposure variables for each security to be the stock specific myopic portfolio holdings, i.e., $B_{i,t} = [x_{i,t}^{\text{MP}}]$. We then follow the methodology developed in Section 2.9 to determine the optimal parameters of our LGS strategy.

³²In our model, $\frac{1}{2} \eta \sigma_\epsilon^2 u^2$ measures the transaction cost of trading u dollars. Therefore, our choice of parameters implies that a trade with a notional value of \$20 million results in \$5,000, \$10,000 and \$20,000 of transaction costs in the three regimes.

Table 4: **Real World Experiment: Policy Performance.**

This table reports the average objective value, terminal wealth, transaction costs paid, and the standard errors (all in millions of dollars), and the average annualized Sharpe ratio. For all measures, averages are taken across the 17 five-year periods in the sample (1930:01-2014:03) in our real world experiment. The two policies are the myopic policy with the optimal t-cost multiplier (MP-TC), and the LGS fixed-lag policy (BL). The final column reports the difference between the BL and MP-TC metrics, and the standard errors of these differences.

	MP-TC	BL	BL-MP-TC
Low Transaction Costs			
Avg Objective	1317	5444	4126
Std Err	2598	2024	1046
Avg Wealth	15046	13999	-1047
Std Err	3219	2627	790
TC	13688	10383	-3305
Avg SR	0.69	0.80	0.11
Medium Transaction Costs			
Avg Objective	382	3011	2630
Std Err	1962	1525	700
Avg Wealth	8050	7453	-597
Std Err	2412	1876	666
TC	8617	6243	-2374
Avg SR	0.50	0.59	0.09
High Transaction Costs			
Avg Objective	258	1623	1365
Std Err	1443	1106	456
Avg Wealth	4280	3810	-469
Std Err	1760	1308	534
TC	4829	3449	-1380
Avg SR	0.36	0.44	0.08

4.3.1 Results

Table 4 shows the performance statistics of the myopic and LGS-based policies across the 17 five year samples. The results show that our LGS-based policy significantly outperforms statistically and economically the adjusted myopic policy in terms of average objective value and Sharpe ratios in all three transaction cost regimes. We would expect the outperformance to increase if we were to allow for more lags in the position and trade vectors of the LGS policy.

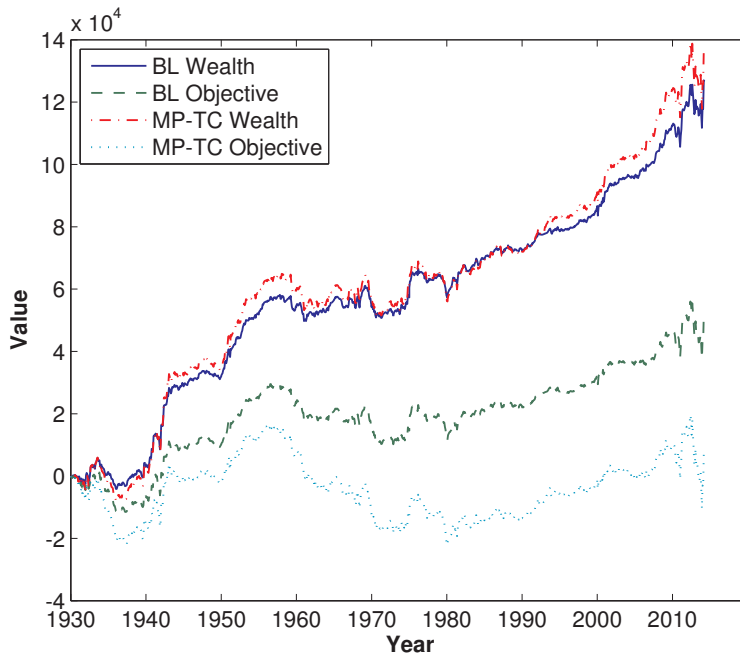


Figure 6: Real world experiment: cumulative wealth and objective function

This figure plots the time series of wealth levels and objective function levels from the real world experiment described in Section 4. We do this from a cumulative gains perspective by aggregating over time the statistics from the five-year investment horizons. We assume a medium transaction cost environment. Dollar values are in millions.

Even though the average terminal wealth values are similar between two policies, MP-TC seems to take substantially higher risk. This gets reflected in a higher variance of the terminal wealth and in much higher transaction costs paid. It appears that the adjusted myopic policy trades too aggressively compared to LGS policy. Furthermore, the outperformance of the LGS policy seems robust as the average statistics are not driven by any single five-year period performance. Actually, in 16 out of 17 five-year investment periods, the LGS-based policy achieves a better objective value than the myopic strategy.

Figure 6 illustrates the wealth and objective dynamics of an investor using the LGS based fixed-lag policy, BL, and the adjusted myopic policy, MP-TC, in each of the five-year investment horizons³³. We assume that the investors are in the medium transaction cost environment. Consistent with the earlier statistics, while having similar wealth evolution, we observe that the outperformance of the BL policy as measured by the cumulative objective value is robust over time.

³³We emphasize that these returns are not to be taken literally, since the transaction costs charged do not correspond to actual transaction costs that would have been paid, and since we used an in-sample estimate of the covariance matrix of returns as discussed previously. The graph is useful to compare the performance of two strategies having access to the same covariance matrix forecast and same transaction cost structure.

5 Conclusion

The LGS framework we propose accommodates complex models of return predictability in a multiperiod setting with transaction costs. Our return predicting factors do not need to follow any pre-specified model but instead can have arbitrary dynamics. We allow for factor dependent covariance structure in returns driven by common factor shocks or stochastic/GARCH volatility, as well as time varying transaction costs.

The main insight is that for the class of LGS the optimal policy can be computed in closed-form by solving a deterministic linear quadratic problem, which is computationally very efficient.

Numerical experiments show that the performance of the linear-quadratic solutions of Litterman (2005) and Gârleanu and Pedersen (2013) come close to the LGS solution when the covariance matrix of price changes is approximately constant (where L-GP provide the optimal solution). However, when returns display stochastic volatility the superiority of the LGS approach is stronger. We also investigate the performance of the LGS framework when trading a strategy based on short-term reversal, momentum and long-term reversal. These three predictor variables have very different half-lives and thus transaction costs are a first order concern. The benefits to using a dynamic framework appear significant compared to a widely used approach that relies on a myopic objective function with a transaction cost multiplier that is chosen to maximize the in-sample performance.

References

- Aït-Sahalia, Yacine, and Michael W Brandt, 2001, Variable selection for portfolio choice, *The Journal of Finance* 56, 1297–1351.
- Balduzzi, Pierluigi, and Anthony W Lynch, 1999, Transaction costs and predictability: Some utility cost calculations, *Journal of Financial Economics* 52, 47–78.
- Black, Fischer, and Robert Litterman, 1991a, Global asset allocation with equities, bonds, and currencies, Goldman, Sachs & Co. Fixed Income Research Report.
- Black, Fischer, and Robert Litterman, 1991b, Global portfolio optimization, *Financial Analysts' Journal* 48, 28–43.
- Brandt, Michael W, and Pedro Santa-Clara, 2006, Dynamic portfolio selection by augmenting the asset space, *The Journal of Finance* 61, 2187–2217.
- Brandt, Michael W., Pedro Santa-Clara, and Rossen Valkanov, 2009, Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns, *Review of Financial Studies* 22, 3411–3447.
- Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1997, Strategic asset allocation, *Journal of Economic Dynamics and Control* 21, 1377–1403.
- Brown, David B, and James E Smith, 2011, Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds, *Management Science* 57, 1752–1770.
- Campbell, John Y, Yeung Lewis Chan, and Luis M Viceira, 2003, A multivariate model of strategic asset allocation, *Journal of Financial Economics* 67, 41–80.
- Campbell, John Y, and Luis M Viceira, 2002, *Strategic asset allocation: portfolio choice for long-term investors*. (Oxford University Press Oxford).
- Chan, Louis KC, Jason Karceski, and Josef Lakonishok, 1999, On portfolio optimization: Forecasting covariances and choosing the risk model, *Review of Financial Studies* 12, 937–974.
- Cochrane, John H, 2007, Portfolio theory, University of Chicago Working Paper.
- Cogley, Timothy, and Thomas J Sargent, 2008, Anticipated utility and rational expectations as approximations of bayesian decision making*, *International Economic Review* 49, 185–221.
- Constantinides, George M., 1986, Capital market equilibrium with transaction costs, *The Journal of Political Economy* pp. 842–862.
- Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross, 1985, An intertemporal general equilibrium model of asset prices, *Econometrica* 53, 363–384.

- Cvitanic, Jaksa, 2001, Theory of portfolio optimization in markets with frictions, in Elyès Jouini, Jaksa Cvitanic, and Marek Musiela, eds.: *Option Pricing, Interest Rates and Risk Management* (Cambridge Univ Press, Cambridge, UK).
- Davis, M. H. A., and A. R. Norman, 1990, Portfolio selection with transaction costs, *Mathematics of Operations Research* 15, 676–713.
- DeBondt, Werner F. M., and Richard H. Thaler, 1985, Does the stock market overreact?, *Journal of Finance* 40, 793–808.
- Detemple, Jérôme, and Marcel Rindisbacher, 2010, Dynamic asset allocation: Portfolio decomposition formula and applications, *Review of Financial Studies* 23, 25–100.
- Duffie, Darrell, and Rui Kan, 1996, A yield-factor model of interest rates, *Mathematical finance* 6, 379–406.
- Dumas, Bernard, and Elisa Luciano, 1991, An exact solution to a dynamic portfolio choice problem under transactions costs, *Journal of Finance* 46, 577–595.
- Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Fama, Eugene F., and James MacBeth, 1973, Risk, return and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.
- Gârleanu, Nicolae, and Lasse H. Pedersen, 2013, Dynamic trading with predictable returns and transaction costs, *The Journal of Finance* 68, 2309–2340.
- Gourieroux, C., A. Monfort, and E. Renault, 1993, Indirect inference, *Journal of applied econometrics* 8, S85–S118.
- Grinold, Richard C., and Ronald N. Kahn, 1999, *Active portfolio management: a quantitative approach for providing superior returns and controlling risk*. (McGraw-Hill Companies New York, NY).
- Jegadeesh, Narasimhan, 1990, Evidence of predictable behavior of security returns, *Journal of Finance* 45, 881–898.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.
- Kim, Tong Suk, and Edward Omberg, 1996, Dynamic nonmyopic portfolio behavior, *Review of Financial Studies* 9, 141–161.
- Kreps, David M, 1998, Anticipated utility and dynamic choice, *Econometric Society Monographs* 29, 242–274.
- Le, Ahn, Kenneth J. Singleton, and Qiang Dai, 2010, Discrete-time affine^Q term structure models with generalized market prices of risk, *Review of Financial Studies* 23, 2184–2227.

- Lehmann, Bruce N., 1990, Fads, martingales, and market efficiency, *Quarterly Journal of Economics* 105, 1–28.
- Litterman, Robert, 2005, Multi-period portfolio optimization, working paper.
- Liu, Hong, and Mark Loewenstein, 2002, Optimal portfolio selection with transaction costs and finite horizons, *Review of Financial Studies* 15, 805–835.
- Liu, Jun, 2007, Portfolio selection in stochastic environments, *Review of Financial Studies* 20, 1–39.
- Ljungqvist, Lars, and Thomas J. Sargent, 2004, *Recursive Macroeconomic Theory*. (The MIT press Cambridge, MA).
- Lynch, Anthony W, and Pierluigi Balduzzi, 2000, Predictability and transaction costs: The impact on rebalancing rules and behavior, *The Journal of Finance* 55, 2285–2309.
- Lynch, Anthony W, and Sinan Tan, 2011, Labor income dynamics at business-cycle frequencies: Implications for portfolio choice, *Journal of Financial Economics* 101, 333–359.
- Markowitz, Harry M., 1952, Portfolio selection, *Journal of Finance* 7, 77–91.
- Merton, Robert C, 1969, Lifetime portfolio selection under uncertainty: The continuous-time case, *The review of Economics and Statistics* 51, 247–257.
- Merton, Robert C., 1971, Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory* 3, 373–413.
- Moallemi, Ciamac, and Mehmet Saglam, 2012, Dynamic portfolio choice with linear rebalancing rules, Available at SSRN 2011605.
- Moallemi, Ciamac C, Mehmet Saglam, and Michael Sotiropoulos, 2014, Short-term predictability and price impact, Available at SSRN 2463952.
- Shreve, Steven E, and H. Mete Soner, 1994, Optimal investment and consumption with transaction costs, *The Annals of Applied Probability* 4, 609–692.
- Stambaugh, Robert F., Jianfeng Yu, and Yu Yuan, 2012, The short of it: Investor sentiment and anomalies, *Journal of Financial Economics* 104, 288–302.
- Vasicek, Oldrich A., 1977, An equilibrium characterization of the term structure, *Journal of Financial Economics* 5, 177–188.