Strategic 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 Cincinatti
email: saglammt@uc.edu

Preliminary Draft: August 15, 2013
This Draft: March 15, 2014

Abstract

We propose a simple approach to dynamic multiperiod portfolio choice with quadratic 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 a trading rule that is optimal if the problem is restricted to a broad class of strategies we define as 'linearity generating strategies.' When restricted to this parametric class the highly non-linear dynamic optimization problem reduces to a deterministic linear-quadratic optimization problem in the parameters of the trading strategies. We investigate realistic examples that show that the approach dominates several alternatives, especially in settings where the covariance matrix of returns is stochastic (*e.g.*, when there is a factor structure in returns or when returns have GARCH dynamics) or when transaction costs vary with the level of volatility.

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 the problem 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.

[1] Merton (1969), Merton (1971), Brennan, Schwartz, and Lagnado (1997), Kim and ohmberg (1991), Campbell (1999), Campbell, Chan, and Viceira (2003), Liu (2007), Detemple and Rindisbacher (2010) and many more. See Cochrane (2012) for a survey.

[2] The academic literature has documented numerous variables which forecast the cross-section of equity returns. Stambaugh, Yu, and Yuan (2011) provides a list of many of these variables, and also argue that the structure and magnitudes of this forecastability exhibits considerable time variation.

[3] Constantinides (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 multiasset case under CARA preferences and for *i.i.d.* returns.

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. First, we assume investors maximize the 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 costs of a given trade are 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 the standard set of return generating processes we consider 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 (2012)). 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." 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 on the security, its conditional variance, and variables which summarize the cost of trading this security.

Importantly each security's weight is also a function of these lagged exposures interacted with both it's own past returns and the past returns of a set of managed portfolios, which implies a very high dimensional optimization problem. One would anticipate that this high-dimensional problem would be unsolvable but, crucially, we show that for strategies in the LGS class, this weight optimization problem reduces to a deterministic linear-quadratic problem that we show can be solved very efficiently.

Another key question is whether the set of LGS's is sufficiently rich that the optimal solution LGS approximates the unconstrained optimum. This is an empirical question. We solve several realistic examples for which we find this is indeed the case. First, we compare the performance of our approach to that of several alternatives in two benchmark simulated economies: one we call a *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 value effects (Fama and French 1993). 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). On 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 (2012). Their linear-quadratic programming approaches 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 Litterman-Garleanu-Pedersen closed-form of solution perform almost equally well in the characteristics based economy. However, in the factor model economy, because the covariance matrix of returns is stochastic, the Litterman-Garleanu-Pedersen solution cannot be applied, since their approach relies on a constant covariance matrix. When we use their trading rule by plugging in either an unconditional estimate of the covariance matrix or the current estimate of the covariance matrix at every time step the resulting trading strategy 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. The LGS also outperforms a myopic mean-variance approach optimized for the presence of transaction-costs, which is often used by practitioners. This alternative approach consists in using the one-period mean-variance solution with transaction costs, but recognizing that this approach ignores the dynamic objective function, it adds a multiplier to the transaction costs incurred when trading. This t-cost multiplier is chosen so as to maximize the actual performance of the strategy across many simulations.

[4]One important difference is that to obtain a closed-form solution Litterman (2005) and Gârleanu and Pedersen (2012) specify their model for price changes and not returns and the objective function of the investor in terms of number of shares. 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.

We also perform an experiment with real return data. We analyze the performance of a trading strategy involving the 100 largest stocks traded on the NYSE over the time period from 1974 to 2012. We trade these stocks exclusively based on the short-term reversal factor, which is a well-known predictor of stock returns. Because the half-life of reversal is several days, portfolio turnover is high and strategy performance of a strategy based on this factor is highly dependent on transaction costs. Also, the literature suggests that strategy performance is dependent on volatility (Khandani and Lo (2007), Nagel (2012)). We therefore use a realistic return process that features GARCH in the common market factor as well as in the cross-sectional idiosyncratic variance. This captures salient empirical features of the reversal factor as documented in Collin-Dufresne and Daniel (2013). In our experiment the costs of trading shares of an individual firm depend on that firm's return volatility, consistent with the findings in the transaction cost literature. Thus, transaction costs are stochastic. We solve for the optimal trading strategy using our LGS and backtest our strategy in comparison with a myopic t-cost optimized strategy. We find that our approach outperforms this benchmark significantly.[5]

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 (2010) 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 closest related to two strands of literature: First, Brandt, Santa-Clara, and Valkanov (2009) model the portfolio weight on each asset directly as linear functions of a set of asset "characteristics" that are determined *ex-ante* to be useful for portfolio selection.[6] The vector of characteristic weights are optimized by by maximizing the average utility the investor would have

[5]Of course, this is not a guarantee that the performance would correspond to a real out of sample strategy performance, in particular, because the transaction cost function we chose, while based on estimates of the literature, does not correspond to actual t-cost paid for execution.

[6]See also Aït-Sahalia and Brandt (2001), Brandt and Santa-Clara (2006) and Moallemi and Saglam (2012).

obtained by implementing the policy over the 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.[7]

However, the BSV approach is limited in that the optimization is performed via numerical simulation, and therefore is limited to a relatively small number of predictive variables. Further, since the performance of the objective function is optimized in sample, restricting to a small number of parameters and predictors is desirable to avoid over-fitting.

Our contribution is that we identify a set of trading strategies for which the optimization can be performed in closed-form using deterministic linear quadratic control for very general return processes in a dynamic setting with transaction costs. We can thus achieve a greater flexibility in parameterizing the trading rule.

Second, Litterman (2005) and Gârleanu and Pedersen (2012) obtain a closed-form solution for the optimal portfolio choice in a model where price changes are linear in a set of predictor variables, the covariance matrix of price changes is constant, trading costs are a quadratic function of the number of shares traded, and investors have a linear-quadratic objective function expressed in terms of number of shares traded. Their approach relies heavily on linear-quadratic stochastic programming (*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, 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 Litterman or Garleanu-Pedersen. 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.

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

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 $S_{i,t}$ $i = 1, \ldots, N$ 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 of state variables X_t:

$$E_t[(R_{t+1}] = M(X_t, t) \tag{1}$$

$$E_t[(R_{t+1} - E_t[R_{t+1}])(R_{t+1} - \mathsf{E}_t[R_{t+1}])'] = \Sigma_{t\to t+1}(X_t, t) \tag{2}$$

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

It is important for our approach that the dynamics of X_t are known so that one can simulate the behavior of the conditional moments of stock returns. An example that nests many return generating processes used in the literature is:

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

for some functions $g(t, \cdot) : \mathbb{R} \to \mathbb{R}$, increasing in their second argument, and where we further introduce the following notation:

- $\beta_{i,t}$ is the $(k, 1)$ vector of exposures to the factors.
- F_{t+1} is the $(k, 1)$ vector of random (as of time t) factor realizations, with mean 0 that follows

a multivariate GARCH process with conditional covariance matrix $\Omega_{t,t+1}$.

- $\epsilon_{i,t+1}$ is the idiosyncratic risk of stock i.

  We assume that $\epsilon_{\cdot,t+1}$ 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.

In that case the vector of state variables $X_t = [\beta_{1,t}; \beta_{2,t}; \ldots \beta_{N,t}; \lambda_t; \Omega_{t,t+1}]$ has $N * k + 1 + k + k * (k + 1)/2$ 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:

1. The **"discrete exponential affine"** model for security returns in which log-returns are affine in factor realizations:[8]

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

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 these return generating processes. We emphasize that the approach does not rely on this factor structure assumption. All that is required is that there be some known relation between the conditional first and second moments of stock returns and the known state vector X_t so that conditional means and variances of stock returns can be simulated along with the state vector.

[8]The 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).

2.2 Cash and stock position dynamics

We assume discrete time dynamics. At the end of each period t the agent buys $u_{i,t}$ dollars of stock 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 stock $i = 1, \ldots N$ held at the end of each period t are thus given by:

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

In vector notation,

$$\begin{aligned} x_t &= x_{t-1} \circ R_t + u_t && (4) \\ w_t &= w_{t-1} R_{0,t} - \mathbf{1}^\top u_t - \frac{1}{2} u_t^\top \Lambda_t u_t && (5) \end{aligned}$$

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.[9]

The matrix Λ_t captures (possibly time-varying) quadratic transaction/price-impact costs, so that $\frac{1}{2} u_t^\top \Lambda_t u_t$ is the dollar cost paid when realizing a trade at time t of size u_t. Without loss of generality, we assume this matrix is symmetric. Gârleanu and Pedersen (2012) discuss some micro-economic foundations for such quadratic costs. It is also very convenient analytically.

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

$$\begin{aligned} x_{t+1} &= (x_t + u_t) \circ R_{t+1} && (6) \\ w_{t+1} &= (w_t - \mathbf{1}^\top u_t - \frac{1}{2} u_t^\top \Lambda_t u_t) R_{0,t+1} && (7) \end{aligned}$$

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 (2012) and thus makes the link between the two frameworks more transparent.

2.3 Objective function

We assume that the agent is endowed with a portfolio of dollar holdings in securities x_0 and an initial amount of cash w_0. The investor's objective function is to maximize a linear quadratic function of his terminal cash and stock positions $F(w_T, x_T) = w_T + a_1^\top x_T - \frac{1}{2} x_T^\top a_2\, x_T$, net of a risk-penalty which we take to be proportional to the per-period variance of the portfolio. We assume a_1 is a $(N, 1)$ vector and a_2 a (N, N) symmetric matrix.[10] So the objective function is simply:

$$\max_{u_1,\ldots,u_T} \mathsf{E}\left[F(w_T, x_T) - \sum_{t=0}^{T-1} \frac{\gamma}{2} x_t^\top \Sigma_{t\to t+1} x_t\right] \tag{8}$$

Recall that $\Sigma_{t\to t+1} = E_t[(R_{t+1} - E_t[R_{t+1}])(R_{t+1} - \mathsf{E}_t[R_{t+1}])']$ is the conditional one-period variance-covariance matrix of returns and γ can be interpreted as the coefficient of risk aversion.

The $F(\cdot,\cdot)$ function parameters can be chosen to capture different objectives, such as maximizing the terminal gross value of the position $(w_T + \mathbf{1}^\top x_T)$ or the terminal liquidation (i.e., net of transaction costs) value of the portfolio $(w_T + \mathbf{1}^\top x_T - \frac{1}{2} x_T^\top \Lambda_T x_T)$, or the terminal wealth penalized for the riskiness of the position $(w_T + \mathbf{1}^\top x_T - \frac{\gamma}{2} x_T^\top \Sigma_T x_T)$, or some intermediate situation.

By recursive substitution x_T and w_T can be rewritten as:

$$\begin{aligned} x_T &= x_0 \circ R_{0\to T} + \sum_{t=1}^{T} u_t \circ R_{t\to T} && (9) \\ w_T &= w_0 R_{0,0\to T} - \sum_{t=1}^{T} \left(u_t^\top \mathbf{1} R_{0,t\to T} + \frac{1}{2} u_t^\top \Lambda_t u_t R_{0,t\to T} \right) && (10) \end{aligned}$$

where we have defined the cumulative return between date t and T on security i as:

$$R_{i,t\to T} = \prod_{s=t+1}^{T} R_{i,s} \tag{11}$$

(with the convention that $R_{i,t\to t} = 1$) and the corresponding N-dimensional vector $R_{t\to T} = [R_{1,t\to T}; \ldots; R_{N,t\to T}]$.

[10] The symmetry assumption on a_2 is without loss of generality.

Now note that

$$a_1^\top x_T = (a_1 \circ R_{0\to T})^\top x_0 + \sum_{t=1}^{T} (a_1 \circ R_{t\to T})^\top u_t \tag{12}$$

Substituting we obtain the following:

$$\begin{aligned}
F(w_T, x_T) &= F_0 + \sum_{t=1}^{T} \left\{ G_t^\top u_t - \frac{1}{2} u_t^\top P_t u_t \right\} - \frac{1}{2} x_T^\top a_2\, x_T && (13)\\
F_0 &= w_0 R_{0,0\to T} + (a_1 \circ R_{0\to T})^\top x_0 && (14)\\
G_t &= a_1 \circ R_{t\to T} + \mathbf{1} \circ R_{0,t\to T} && (15)\\
P_t &= \Lambda_t \circ R_{0,t\to T} && (16)
\end{aligned}$$

Substituting into the objective function given in equation (8) it can be rewritten as:

$$F_0 - \frac{\gamma}{2} x_0^\top Q_0 x_0 + \max_{u_1,\ldots,u_T} \sum_{t=1}^{T} \mathsf{E} \left[G_t^\top u_t - \frac{1}{2} u_t^\top P_t u_t - \frac{\gamma}{2} x_t^\top Q_t x_t \right] \tag{17}$$

subject to the non-linear dynamics given in equations (4) and (5) and where we have defined

$$Q_t = \begin{cases} \Sigma_{t\to t+1} & \text{for } t < T \\ \frac{1}{\gamma} a_2 & \text{for } t = T \end{cases} \tag{18}$$

2.3.1 A Useful Special Case

There is a particular choice of $a_1 = 1$ and $a_2 = 0$ that is useful to compare our objective function with previous literature. Indeed for that case, the objective function is simply

$$\max_{u_1,\ldots,u_T} \mathsf{E} \left[w_T + x_T^\top 1 - \sum_{t=0}^{T-1} \frac{\gamma}{2} x_t^\top \Sigma_{t\to t+1} x_t \right] \tag{19}$$

Note that by recursion we can write:[11]

$$x_T = x_0 + \sum_{t=0}^{T-1} x_t \circ r_{t+1} + \sum_{t=1}^{T} u_t \tag{20}$$

$$w_T = w_0 + \sum_{t=0}^{T-1} w_t r_{0,t+1} - \sum_{t=1}^{T} (u_t^\top \mathbf{1} + \frac{1}{2} u_t^\top \Lambda_t u_t) \tag{21}$$

where we have defined the *net return* $r_{t+1} = R_{t+1} - 1$ and corresponding expected net return $m_t = \mathsf{E}_t[R_{t+1} - 1] = M(X_t, t) - 1$ Inserting in the objective function, and defining the corresponding expected return and simplifying we find

$$\max_{u_1,\ldots,u_T} \mathsf{E}\left[\sum_{t=1}^{T} \left\{ (w_{t-1} m_{0,t-1} + x_{t-1}^\top m_{t-1}) - \frac{\gamma}{2} x_{t-1}^\top \Sigma_{t-1\to t} x_{t-1} - \frac{1}{2} u_t^\top \Lambda_t u_t \right\}\right] \tag{22}$$

We see that this objective function is very similar to that used in Litterman (2005) and Gârleanu and Pedersen (2012). It is the expected sum of local-mean-variance objectives, net of transaction costs paid. One notable difference is that the objective function here is expressed in terms of dollar holdings (x_t, w_t) and rates of returns on securities (r_t) as opposed to number of shares and price changes in their framework. One implication of working with number of shares for example, is that if expected returns (m_t), t-costs (Λ_t) and covariances (Σ_t) do not change, then it is optimal not to trade (since the optimal number of shares has not changed), even though the dollar position may have changed due to some random return realization. In other words, there is no standard rebalancing for diversification purposes in a framework that optimizes in terms of number of shares. Instead, our objective function will capture this rebalancing motive for trading (which is at the heart of the classic Merton (1969) dynamic portfolio optimization with constant investment opportunity set, for example).

We now turn to our proposed solution approach and, to that effect, introduce the set of 'linearity generating policies' that we consider.

[11] Indeed, $x_T = x_{T-1} \circ (R_T - 1) + x_{T-1} + u_T = x_{T-1} \circ (R_T - 1) + x_{T-2} \circ (R_{T-1} - 1) + x_{T-2} + u_{T-1} + u_T = \ldots$

2.4 Linearity generating policies

Even though the objective function is similar to that of a linear-quadratic problem which are known to be very tractable (*e.g.*, Litterman (2005), GP (2012)) our problem is not in that class because of the non-linearity introduced by the state equation, and because of the general return process, which may display stochastic volatility (and thus make the matrix Q_t stochastic). Thus the problem appears difficult to solve in full generality (even numerically). Instead, we introduce a specific set of 'linearity generating trading strategies' (LGS) for which the problem remains tractable.

The idea of restricting the set of strategies to make the problem tractable is not new, and has for example been used by Brandt, Santa-Clara, and Valkanov (2009). They propose some trading strategies that are linear in stock characteristics and to numerically optimize directly the empirical objective function on a sample of data over the parameters of the trading strategy. Because, their approach relies on a numerical in-sample optimization however, they have to specify fairly simple strategies so as to not over-fit the data. Instead, since with our approach the optimization is done in closed-form we can have a very rich class of path-dependent strategies, which is very useful to handle problems with transaction costs.[12]

The remarkable result we show below is that, for this class of LGS, the problem reduces to a deterministic linear-quadratic optimization problem in the parameters of the policy. While GP (2013) make some strong assumptions about the return generating process (no factor structure in the covariance matrix, no stochastic volatility) and the objective function (investors care about number of shares and not dollar exposures) to obtain a closed-form solution, we instead choose to specify more realistic return dynamics and work with the standard (non-linear) wealth dynamics, but to restrict the set of strategies that the investors can use in order to obtain a tractable solution.

Of course, it is an empirical question whether the set of LGS is sufficiently large to be useful. We present some empirical tests of our approach in the next section. First, we describe the strategy set we consider. Then we explain how the portfolio optimization can be done in closed-form, within that restricted set.

[12] One advantage of the Brandt, Santa-Clara, and Valkanov (2009) approach is that they dispense with specifying the return generating process altogether, instead relying on the empirical performance of there propose strategies. Instead, for our approach we need to specify the return generating process, and in particular, the way in which expected returns and variances depend on the characteristics used for the trading rule.

At this stage 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.

To define our set of LGS we first specify for each stock K variables $B_{i,t}$ which we call 'stock exposures.' These are typically non-linear transformations of the general state vector X_t (i.e., $B_{i,t} = h_i(X_t)$). For example, $B_{i,t}$ may include the individual stock return's conditional expected return divided by its conditional variance, which is a natural candidate (*e.g.*, Aït-Sahalia and Brandt (2001)). More generally, it would include stock specific factor exposures, conditional variances and other relevant information for portfolio formation. Then the LGS are further specified by K-dimensional vectors of parameters, $\pi_{i,s,t}$ and $\theta_{i,s,t}$, defined for all $i = 1, \ldots, N$ and for all $s \leq t$. These parameters will fully determine the dollar trades in asset i ($u_{i,t}$) and the corresponding positions ($x_{i,t}$) via the parametric relation:

$$u_{i,t} = \sum_{s=0}^{t} \pi_{i,s,t}^{\top} \mathcal{B}_{i,s\to t} \tag{23}$$

and

$$x_{i,t} = \sum_{s=0}^{t} \theta_{i,s,t}^{\top} \mathcal{B}_{i,s\to t} \tag{24}$$

where $\mathcal{B}_{i,s\to t}$ is defined as the K-dimensional vector of buy and hold returns between s and t on trading strategies that scale their positions at time s proportionally to the vector of exposures $B_{i,s}$:

$$\mathcal{B}_{i,s\to t} = B_{i,s} R_{_{i,s\to t}}. \tag{25}$$

So in effect, the dollar position at time t in asset i ($x_{i,t}$) can be thought of as a weighted average of simple buy and hold trading strategies that went long the stock at past dates ($s < t$) proportionally to past exposures and held the stock until date t. The parameter vector $\theta_{i,s,t}$ measures the weight in the current dollar position that is put on the time-s exposure trade. In other words, our LGS allow trades at time t to depend on current exposures $B_{i,t}$, but also on all past exposures weighted by their past holding period returns.

Intuitively, the dependence on current exposures is clearly important. In fact, in a no-transaction cost affine portfolio optimization problem where the optimal solution is well-known, the optimal solution will involve only current exposures (see, *e.g.*, Liu (2007)). Note that this is also the choice made by Brandt, Santa-Clara, and Valkanov (2009) for their 'parametric portfolio policies.' However, while Brandt, Santa-Clara, and Valkanov (2009) specify the loadings on exposure of individual stocks to be identical, we allow two stocks with identical exposures (and with perhaps different levels of idiosyncratic variance) to have different weights and trades.[13]

With transaction costs, allowing portfolio weights and trades to depend on past returns interacted with past exposures seems useful. The intuition for this comes from the path-dependence we observe in known closed-form solutions (Constantinides (1986), Davis and Norman (1990), Dumas and Luciano (1991), Liu and Loewenstein (2002) and others).

To proceed, we note that the assumed linear position and trading strategies in equations (23) and (24) have 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 $i = 1, \ldots, N$:

$$\theta_{i,s,t} = \theta_{i,s,t-1} + \pi_{i,s,t} \qquad \text{for } s < t \tag{26}$$

$$\theta_{i,t,t} = \pi_{i,t,t} \qquad \forall t \tag{27}$$

These restrictions are intuitive. They indicate that the position at time t loads on the exposure at date t only through the trade at time t. However, it may load on previous exposures through all previous trades. We can rewrite these policies in a concise matrix form. First, define the $NK(t+1)$-dimensional vectors π_t and θ_t as

$$\pi_t = [\pi_{1,0,t}; \ldots; \pi_{n,0,t}; \pi_{1,1,t}; \ldots; \pi_{n,1,t}; \ldots; \pi_{1,t,t}; \ldots; \pi_{n,t,t}] \tag{28}$$

$$\theta_t = [\theta_{1,0,t}; \ldots; \theta_{n,0,t}; \theta_{1,1,t}; \ldots; \theta_{n,1,t}; \ldots; \theta_{1,t,t}; \ldots; \theta_{n,t,t}] \tag{29}$$

Further, let's define the following (NK, N) matrices (defined for all $0 \leq s \leq t \leq T$) as the diagonal

[13](Note, for the Brandt, Santa-Clara, and Valkanov (2009) econometric approach it is useful to have fewer parameters. This is not an issue with our approach as our solution is closed-form.

concatenations of the N vectors $\mathcal{B}_{i,s\to t}$ $\forall i = 1, \ldots, N$:

$$\mathcal{B}_{s,t} = \left\{ \begin{array}{cccc} \mathcal{B}_{1,s\to t} & 0 & 0\ldots & 0 \\ 0 & \mathcal{B}_{2,s\to t} & 0\ldots & 0 \\ & \ldots & & \\ 0 & \ldots & 0 & \mathcal{B}_{n,s\to t} \end{array} \right\}$$

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

$$\mathcal{B}_t = [\mathcal{B}_{0,t}; \mathcal{B}_{1,t}, \ldots, \mathcal{B}_{t,t}]$$

It is then straightforward to check that:

$$u_t = \mathcal{B}_t^\top \pi_t \tag{30}$$

$$x_t = \mathcal{B}_t^\top \theta_t \tag{31}$$

Further, in terms of these definitions the constraints on the parameter vector in (26) can be rewritten concisely as:

$$\theta_t = \theta_{t-1}^0 + \pi_t \tag{32}$$

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 'linear 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 equation (30) into our objective function we may rewrite the original problem given in equation (17) as follows:

$$F_0 - \frac{\gamma}{2} x_0^\top Q_0 x_0 + \max_{\pi_1,\ldots,\pi_T} \sum_{t=1}^{T} \mathcal{G}_t^\top \pi_t - \frac{1}{2}\pi_t^\top \mathcal{P}_t \pi_t - \frac{\gamma}{2}\theta_t^\top \mathcal{Q}_t \theta_t \tag{33}$$

$$s.t.\ \theta_t = \theta_{t-1}^0 + \pi_t \tag{34}$$

and where we define the vectors $\mathcal{G}_t$ and the matrices $\mathcal{P}_t$ and $\mathcal{Q}_t$ defined for all $t = 1, \ldots, T$ by

$$\mathcal{G}_t = \mathsf{E}[\mathcal{B}_t G_t] \tag{35}$$

$$\mathcal{P}_t = \mathsf{E}[\mathcal{B}_t P_t \mathcal{B}_t^\top] \tag{36}$$

$$\mathcal{Q}_t = \mathsf{E}[\mathcal{B}_t Q_t \mathcal{B}_t^\top] \tag{37}$$

Note that the time indices for the matrices $\mathcal{G}_t, \mathcal{P}_t, \mathcal{Q}_t$ also capture their size (index t denotes a square-matrix or vector of row-length $NK(t+1)$). The matrices $\mathcal{G}_t, \mathcal{P}_t, \mathcal{Q}_t$ can be solved for explicitly or by simulation depending on the assumptions made about the state vector X_t driving the return generating process R_t and the corresponding stock-specific exposure dynamics $B_{i,t}$. But once these expressions have been computed or simulated (and this only needs to be done once), then the explicit solution for the optimal strategy can be derived using standard deterministic linear-quadratic dynamic programming. We next derive the solution.

2.5 Closed form solution

Consider the deterministic linear-quadratic problem:

$$\max_{\pi_1,\ldots,\pi_T} \sum_{t=1}^{T} \mathcal{G}_t^\top \pi_t - \frac{1}{2}\pi_t^\top \mathcal{P}_t \pi_t - \frac{\gamma}{2}\theta_t^\top \mathcal{Q}_t \theta_t \tag{38}$$

$$s.t.\ \theta_t = \theta_{t-1}^0 + \pi_t \tag{39}$$

Define recursively the value function starting from $V(T) = 0$ for all $t < T$ by:

$$V(t-1) = \max_{\pi_t} \left\{ \mathcal{G}_t^\top \pi_t - \frac{1}{2}\pi_t^\top \mathcal{P}_t \pi_t - \frac{\gamma}{2}\theta_t^\top \mathcal{Q}_t \theta_t + V(t) \right\} \tag{40}$$

$$s.t.\ \theta_t = \theta_{t-1}^0 + \pi_t \tag{41}$$

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

$$V(t) = -\frac{\gamma}{2}\theta_t^\top M_t \theta_t + L_t^\top \theta_t + H_t \tag{42}$$

with M a symmetric matrix. Clearly $M(T) = 0$ and $L(T) = 0$ and $H(T) = 0$. To find the recursion plug the guess in the Bellman equation:

$$V(t) = \max_{\pi_t} \left\{ \mathcal{G}_t^\top \pi_t - \frac{1}{2}\pi_t^\top \mathcal{P}_t \pi_t - \frac{\gamma}{2}\theta_t^\top (\mathcal{Q}_t + M_t)\theta_t + L_t^\top \theta_t + H_t \right\} \tag{43}$$

$$s.t.\ \theta_t = \theta_{t-1}^0 + \pi_t \tag{44}$$

Now plugging in the constraint, we can simplify the Bellman equation using the following notation $\overline{x}$ is the vector (submatrix) obtained from x by deleting the last NK rows (rows and columns). In Matlab notation $\overline{x} = x[1 : end - NK, 1 : end - NK]$.

$$V(t) = \max_{\pi_t} \left\{ (\mathcal{G}_t + L_t)^\top \pi_t - \frac{1}{2}\pi_t^\top [\mathcal{P}_t + \gamma(\mathcal{Q}_t + M_t)]\pi_t - \frac{\gamma}{2}\theta_{t-1}^\top (\overline{\mathcal{Q}}_t + \overline{M}_t)\theta_{t-1} - \gamma\theta_{t-1}^{0\top}[\mathcal{Q}_t + M_t]\pi_t + \overline{L}_t^\top \theta_{t-1} + H_t \right\} \tag{45}$$

The first order condition gives:

$$\pi_t = [\mathcal{P}_t + \gamma(\mathcal{Q}_t + M_t)]^{-1} \left(\mathcal{G}_t + L_t - \gamma(\mathcal{Q}_t + M_t)^\top \theta_{t-1}^0 \right) \tag{46}$$

Plugging into the state equation we find

$$\theta_t = [\mathcal{P}_t + \gamma(\mathcal{Q}_t + M_t)]^{-1} \left(\mathcal{G}_t + L_t + \mathcal{P}_t^\top \theta_{t-1}^0 \right) \tag{47}$$

And plugging into the Bellman equation we find:

$$V(t) = \frac{1}{2}(\mathcal{G}_t + L_t - \gamma(\mathcal{Q}_t + M_t)^\top \theta^0_{t-1})^\top [\mathcal{P}_t + \gamma(\mathcal{Q}_t + M_t)]^{-1} \left(\mathcal{G}_t + L_t - \gamma(\mathcal{Q}_t + M_t)^\top \theta^0_{t-1}\right) \tag{48}$$

$$- \frac{\gamma}{2}\theta_{t-1}^\top(\overline{\mathcal{Q}}_t + \overline{M}_t)\theta_{t-1} + \overline{L}_t^\top \theta_{t-1} + H_t \tag{49}$$

Setting $\Psi_t = [\mathcal{P}_t + \gamma(\mathcal{Q}_t + M_t)]^{-1}$ and expanding we find:

$$V(t) = H_t + \frac{1}{2}(\mathcal{G}_t + L_t)^\top \Psi_t (\mathcal{G}_t + L_t) \tag{50}$$

$$- \gamma(\mathcal{G}_t + L_t)^\top \Psi_t(\mathcal{Q}_t + M_t)^\top \theta^0_{t-1} + \overline{L}_t^\top \theta_{t-1} \tag{51}$$

$$- \frac{\gamma}{2}\theta_{t-1}^\top[\overline{\mathcal{Q}}_t + \overline{M}_t - \gamma(\overline{\mathcal{Q}}_t + \overline{M}_t)^\top \overline{\Psi}_t(\overline{\mathcal{Q}}_t + \overline{M}_t)]\theta_{t-1} \tag{52}$$

Our guess is thus correct if the following recursion holds:

$$H_{t-1} = H_t + \frac{1}{2}(\mathcal{G}_t + L_t)^\top \Psi_t(\mathcal{G}_t + L_t) \tag{53}$$

$$L_{t-1} = \overline{L}_t - \gamma\overline{(\mathcal{Q}_t + M_t)\Psi_t(\mathcal{G}_t + L_t)} \tag{54}$$

$$M_{t-1} = \overline{\mathcal{Q}}_t + \overline{M}_t - \gamma(\overline{\mathcal{Q}}_t + \overline{M}_t)^\top \overline{\Psi}_t(\overline{\mathcal{Q}}_t + \overline{M}_t) \tag{55}$$

With initial condition

$$H_T = 0 \tag{56}$$

$$L_T = 0 \tag{57}$$

$$M_T = 0 \tag{58}$$

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 slightly restricted set of LGS strategies, where one looks back only a finite set of periods. This set of restricted 'finite lag' LGS is useful in practical applications when the time horizon is fairly long and for signals that have a relatively fast decay rate, where the dependence to past periods can be safely restricted. We show here that the same tractability obtains for finite lags.

2.6 LGS with finite number of lags

Suppose we want to use strategies that only look back at most ℓ lags. Let us first specify the trading rule to only trade based on at most ℓ lags, i.e. such that:

$$u_{i,t} = \sum_{s=t-\ell\vee 0}^{t} \pi_{i,s,t}^{\top} \mathcal{B}_{i,s\to t} \tag{59}$$

If we want the holdings to remain linear and of the form:

$$x_{i,t} = \sum_{s=0}^{t} \theta_{i,s,t}^{\top} \mathcal{B}_{i,s\to t} \tag{60}$$

Then we see that the linear constraints in equations (26) have to be modified so as to still satisfy the wealth dynamics in equations (4). Specifically, we require:

$$\theta_{i,s,t} = \pi_{i,s,t} \qquad \text{for } s = t \tag{61}$$

$$\theta_{i,s,t} = \theta_{i,s,t-1} + \pi_{i,s,t} \qquad \text{for } t - \ell \vee 0 \leq s < t \tag{62}$$

$$\theta_{i,s,t} = \theta_{i,s,t-1} \qquad \text{for } 0 < s < t - \ell \tag{63}$$

Since this is still a set of linear contstraints we can straightforwardly extend the previous method to derive the optimal LGS strategy with trades that only look back ℓ periods. Interestingly, inspecting these constraints we note that if we impose an additional linear constraint on the trading strategy such that $\theta_{i,t-\ell,t-1} + \pi_{i,t-\ell,t} = 0 \ \forall t > \ell$, then it follows that $\theta_{i,s,t} = 0 \ \forall 0 < s < t - \ell$. In other words, by imposing one additional linear constraint on the trading strategy one can find a set of LGS where both the trading strategy u_t and the dollar holdings x_t look-back at most ℓ periods. In other words, where

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

and

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

We summarize this second set of linear constraints as:

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

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.

It remains an empirical question whether this class of trading strategies is sufficiently large to be useful. We now present some numerical simulation experiments and an implementation with real data.

3 Simulation Experiment

In this section we present several experiments to illustrate the usefulness of our portfolio selection approach. We compare portfolio selection in a characteristics-based versus factors-based return generating environment.

As we show below the standard linear-quadratic portfolio approach proposed in Litterman (2005) and Gârleanu and Pedersen (2012) is well-suited to the characteristics-based environment, but in a factor-based environment, since it cannot adequately capture the systematic variation in the covariance matrix due to variations in the exposures it is less successful. Instead, our approach can handle this feature and thus performs better.

3.1 Characteristics versus Factor-based return generating model

We wish to compare the following two environments:

- The factor-based return generating process

$$R_{i,t+1} = \alpha_i + B_{i,t}^{\top}(F_{t+1} + \lambda) + \epsilon_{i,t+1} \tag{64}$$

- The characteristics based return generating process:

$$R_{i,t+1} = \alpha_i + B_{i,t}^{\top}\lambda + \omega_{i,t+1} \tag{65}$$

where in both cases we assume that there are three return generating factors corresponding to (1) short term (5-day) reversal, (2) medium term (1 year) momentum, (3) long-term (5 year) reversal (and potentially a common market factor).

Note the difference between the two frameworks. In the characteristics based framework, the conditional covariance of returns is constant $\Sigma_{t\to t+1} = \Sigma_\omega$ and is therefore not affected by the factor exposures. Instead, in the factor-based framework, the conditional covariance matrix of returns is time varying: $\Sigma_{t\to t+1} = B_t \Omega B_t^{\top} + \Sigma_\epsilon$ where $B_t = [B_{1,t}^{\top}; B_{2,t}^{\top}; \ldots; B_{n,t}^{\top}]$ is the (N, K) matrix of factor exposures.

We assume that the half-life of the 5-day factor is 3 days, that of the one-year factor is 150 days, that of the 5-year factor is 700 days. We define the exposure dynamics using the simple auto-regressive process:

$$B_{i,t+1}^k = (1 - \phi_k) B_{i,t}^k + \epsilon_{i,t+1}.$$

Note that the innovation in the factor exposure are driven entirely by idiosyncratic return shocks as expected given their interpretation as 'technical' return based factors. The AR1 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 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 $\phi_k = (\frac{1}{2})^{\hat{h}_k}$.

For the case, where we investigate the 'Characteristics based' model we set the constant covariance matrix Σ_ω so that it matches the unconditional covariance matrix of the factor based return generating process, i.e., we set

$$\Sigma_\omega = \mathsf{E}[B_t \Omega B_t^{\top} + \Sigma_\epsilon]$$

Table 1: Calibration results for λ and Ω.

	Fama-French Moments
$\boldsymbol{\lambda_1}$	-0.00726
$\boldsymbol{\lambda_2}$	0.00182
$\boldsymbol{\lambda_3}$	-0.00323
$\boldsymbol{\Omega_{11}}$	0.00103
$\boldsymbol{\Omega_{12}}$	0.00051
$\boldsymbol{\Omega_{13}}$	0.00154
$\boldsymbol{\Omega_{22}}$	0.00050
$\boldsymbol{\Omega_{23}}$	0.00081
$\boldsymbol{\Omega_{33}}$	0.00162

Note that

$$B_t \Omega B_t^\top = \sum_{l,m=1}^{K} \Omega_{l,m} B^l_{:,t} (B^m_{:,t})^\top$$

where $B^k_{:,t}$ is the factor values of each asset corresponding to the kTh factor at time t.

3.2 Calibration of main parameters

The number of assets in our experiment is 15. One can think of these as a collection of portfolios instead of individual stocks, *e.g.*, stock or commodity indices. Our trading horizon is 26 weeks with weekly rebalancing. Our objective is to maximize the net terminal wealth minus penalty terms for excessive risk. We thus set $a_1 = 1$ and $a_2 = 0$ in our objective function so it is similar to that used in previous papers (see section 2.3.1).

We calibrate the factor mean, λ, and covariance matrix, Ω, using the Fama-French 10 portfolios sorted on short-term reversal, momentum, and long term reversal. Using monthly returns, we compute the performance of the long-short portfolio for the highest and lowest decile in each factor data. Obtaining 3 long-short portfolios, we set λ to be their mean and Ω to be their covariance matrix. Table 1 illustrates the estimated values for λ and Ω.

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_ϵ, respectively. We assume that Σ_ϵ is a diagonal matrix *e.g.*, $\mathsf{diag}(\sigma_\epsilon)$. Each entry in σ_ϵ is set randomly at the beginning of the simulation according to a normal distribution with mean 0.20 and standard deviation 0.05.

The initial distribution for $B_{i,0}^k$ is given by the unconditional stationary distribution of $B_{i,t}^k$ which is given by a normal distribution with mean zero and variance $\frac{\sigma_{\epsilon,i}^2}{2\phi-\phi^2}$.

The transaction cost matrix, Λ is assumed to be a constant multiple of Σ_ω or Σ_ϵ with proportionality constant η in characteristics or factor-based return generating model respectively. We use a rough estimate of η according to widely used transaction cost estimates reported in the algorithmic trading community. We provide two regimes: low and high transaction cost environment. The slippage values for these two regimes are assumed to be around 4bps and 400bps respectively. Therefore, we expect that a trade with a notional value of $\$100,000$ results in \$40 and \$4000 of transaction costs in these regimes. In our model, $\eta\sigma_\epsilon^2 u^2$ measures the corresponding transaction cost of trading u dollars. Using $u = \$100,000$ and $\sigma_\epsilon = 0.20$, this yields an η is roughly around 5×10^{-6} and 5×10^{-4} for the low and high transaction cost regimes respectively.

Finally, we assume that the coefficient of risk aversion, γ equals 10^{-6}, which can be thought of as corresponding to a relative risk aversion of 1 for an agent with 1 million dollars under management.

3.3 Approximate policies

Due to the nonlinear dynamics in our wealth function, solving for the optimal policy even in the case of a concave objective function is intractable due to the curse of dimensionality. In this section, we will provide various policies that will help us compare the performance of the best linear policy to some alternative approaches.

Gârleanu & Pedersen Policy (GP): Using the methodology in Gârleanu and Pedersen (2012), we can construct an approximate trading policy that will work in our current set-up. A closed-form solution can be obtained if one works with linear dynamics in state and control variables:

$$\bar{r}_{t+1} = C_t B_t^{\mathsf{st}} + \bar{\epsilon}_{t+1}$$

$$B_{t+1}^{\mathsf{st}} = (I - \Phi)\, B_t^{\mathsf{st}} + \varepsilon_{t+1}$$

where $\bar{r}_{t+1} = S_{t+1} - S_t$ stores dollar price changes. Then, our objective function can be written as

$$\max \mathsf{E}\left[\sum_{t=1}^{T}\left(x_{t-1}^{\top}\bar{r}_t - \frac{\gamma}{2}x_t^{\top}\bar{\Sigma}_t x_t - \frac{1}{2}u_t^{\top}\bar{\Lambda}u_t\right)\right]$$
$$s.t. x_t = x_{t-1} + u_t$$

where $\bar{\Lambda}$ and $\bar{\Sigma}_t$ are deterministic and measured in dollars and given by

$$C_t = \mathsf{E}[\mathrm{diag}(S_t)]\left(\lambda^{\top} \otimes I_{N\times N}\right)$$
$$\bar{\Lambda} = \mathsf{E}[S_t S_t^{\top}]\Lambda$$
$$\bar{\Sigma}_t = \mathrm{Var}(\bar{r}_{t+1}).$$

The optimal solution to this problem is given by

$$x_t = \left(\bar{\Lambda} + \gamma\bar{\Sigma}_t + A_{xx}^{t}\right)^{-1}\left(\bar{\Lambda}x_{t-1} + \left(A_{xf}^{t}\left(I - \Phi\right)B_t^{\mathsf{st}}\right)\right)$$

with the following recursions:

$$A_{xx}^{t-1} = -\bar{\Lambda}\left(\bar{\Lambda} + \gamma\bar{\Sigma}_t + A_{xx}^{t}\right)^{-1}\bar{\Lambda} + \bar{\Lambda}$$
$$A_{xf}^{t-1} = \bar{\Lambda}\left(\bar{\Lambda} + \gamma\bar{\Sigma}_t + A_{xx}^{t}\right)^{-1}\left(A_{xf}^{t}\left(I - \Phi\right) + C_t\right)$$

Note that this policy uses expected price changes as an input. To obtain expected price changes we scaled our expected returns by the unconditional expected stock price. Of course, stock prices vary significantly across any given sample path. Therefore we also experimented in using conditional expected price changes as an input to the GP trading rule, where we multiply the expected returns used in the policy at every step by the ratio:[14]

$$D_t = \mathrm{diag}\left(\frac{S_{1,t}}{E\left[S_{1,t}\right]}, \ldots, \frac{S_{n,t}}{E\left[S_{n,t}\right]}\right).$$

[14] This somewhat ad hoc adjustment improves the performance of the GP policy which is designed for stationary price changes environment, whereas our simulation uses a 'log-normal' style return process.

Then our scaled policy uses

$$x_t = D_t^{-1}\left(\bar{\Lambda} + \gamma\bar{\Sigma}_t + A_{xx}^t\right)^{-1} D_t^{-1}\left(D_t\bar{\Lambda}D_t x_{t-1} + \left(D_t A_{xf}^t\left(I-\Phi\right)\right)B_t^{\mathsf{st}}\right).$$

Myopic Policy (MP): We can solve for the myopic policy using only one-period data. We solve the myopic problem given by

$$\max \mathsf{E}\left[\left(x_t^\top r_{t+1} - \frac{\gamma}{2}x_t^\top \Sigma_t x_t - \frac{1}{2}u_t^\top \Lambda u_t\right)\right].$$

Using the dynamics for r_{t+1}, the optimal myopic policy is given by

$$x_t = \left(\Lambda + \gamma\left(B_t\Omega B_t^\top + \Sigma_\epsilon\right)\right)^{-1}\left(B_t\lambda + \Lambda\left(x_{t-1}\circ R_t\right)\right)$$

Myopic Policy with Transaction Cost Multiplier (MP-TC): Since myopic policy only considers the current state of the return predicting factors, it incurs substantial transaction costs. This policy can be significantly improved by considering an another optimization problem where one adds a multiplier to the transaction cost matrix, which ultimately tries to control the amount of transaction costs incurred by the policy. This multiplier is optimized to maximize the unconditional performance (i.e., across all simulations) of the trading strategy. Thus, this policy uses

$$x_t = \left(\tau^*\Lambda + \gamma\left(B_t\Omega B_t^\top + \Sigma_\epsilon\right)\right)^{-1}\left(B_t\lambda + \tau^*\Lambda\left(x_{t-1}\circ R_t\right)\right)$$

where τ^* is given by

$$\underset{\tau}{\operatorname{argmax}}\, \mathsf{E}\left[\left(x_t^\top r_{t+1} - \frac{\gamma}{2}x_t^\top \Sigma_t x_t - \frac{1}{2}u_t^\top \Lambda u_t\right)\right]$$

with

$$x_t = \left(\tau\Lambda + \gamma\left(B_t\Omega B_t^\top + \Sigma_\epsilon\right)\right)^{-1}\left(B_t\lambda + \tau\Lambda\left(x_{t-1}\circ R_t\right)\right)$$

Best Linear Policy (BL): We define the relevant exposure variables for each stock to be the $B_{i,t}$. We then follow the methodology in Section 2 to find the optimal linear policy that satisfies our nonlinear state evolution:

$$u_t = \mathcal{B}_t^\top \pi_t^*$$
$$x_t = \mathcal{B}_t^\top \theta_t^*$$

where as before $\mathcal{B}_t$ is constructed from the $\mathcal{B}_{i,s\to t} = B_{i,s}R_{s\to t}$ managed strategy return. We find the optimal loadings on these past exposures at any time t by solving for the optimal π_t^* and θ_t^* in:

$$\max_{\pi_1,\ldots,\pi_T} \sum_{t=1}^{T} \mathcal{G}_t^\top \pi_t - \frac{1}{2}\pi_t^\top \mathcal{P}_t \pi_t - \frac{\gamma}{2}\theta_t^\top \mathcal{Q}_t \theta_t$$
$$s.t.\ \theta_t = \theta_{t-1}^0 + \pi_t$$

Restricted Best Linear Policy (RBL): Instead of using the whole history of factor exposures for our policy, we can restrict the best linear policy to use only a fixed number of periods. In this experiment, we will use only the last observed exposures in our position vector, x_t, and the last two period's exposures and the last period's return in our trade vector, u_t. Formally, we will let

$$x_{i,t} = \theta_{i,t}^\top \mathcal{B}_{i,t,t}, \tag{66}$$
$$u_{i,t} = \pi_{i,1,t}^\top \mathcal{B}_{i,t-1,t} + \pi_{i,2,t}^\top \mathcal{B}_{i,t,t}, \tag{67}$$

where we need

$$\pi_{i,2,t} = \theta_{i,t}, \tag{68}$$
$$\pi_{i,1,t} = -\theta_{i,t-1}, \tag{69}$$
$$\pi_{i,1,1} = 0, \tag{70}$$

in order to satisfy the nonlinear state dynamics in (4) and (5).

Myopic Policy without Transaction Costs (NTC): Without transaction costs, our trading problem is easy to solve, namely, the myopic policy will be optimal. Thus, using the myopic policy

Table 2: **Policy Performance: No common factor noise and low t-cost environment.** This table summarizes the performance of each policy in the case of no common factor noise and low transaction cost environment. For each policy, we report average terminal wealth, average objective value, variance of the terminal wealth, average terminal Sharpe ratio in the presence and absence of transaction costs and average weekly Sharpe ratio in the presence of transaction costs. (Dollar values are in thousands of dollars.)

	GP	**MP**	**MP-TC**	**RBL**	**BL**	**NTC**
Avg Wealth	263.9	573	574.6	547.5	568.5	594.3
Avg Objective	175.8	281.9	282.4	281.1	291.0	297.0
Variance	3.11e+04	1.37e+05	1.37e+05	1.23e+05	1.30e+05	1.44e+05
TC	5.65	9.967	11.71	14.66	13.45	0
Sharpe with TC	2.13	2.188	2.196	2.207	2.231	2.215
Sharpe w/o TC	2.15	2.194	2.204	2.22	2.244	2.215
Weekly Sharpe with TC	2.73	3.39	3.393	3.383	3.443	3.453

in the absence of transaction costs, i.e.,

$$x_t = \left(\gamma \left(B_t \Omega B_t^\top + \Sigma_\epsilon\right)\right)^{-1} (B_t \lambda)$$

and applying it to the objective function without the transaction cost terms will provide us an upper bound for the optimal objective value of the original dynamic program.

3.4 Simulation Results

We run the performance statistics of our approximate policies in the presence and lack of factor noise and low and high transaction costs. We observe that in all of these cases, the best linear policy performs very well compared to the other approximate policies and that when compared to the upper bound it achieves near-optimal performance.

Table 2 illustrates that when transaction costs are relatively small, myopic policies are also near-optimal but even in this case best linear policy dominates in terms of performance. The GP policy does not perform very well, we suspect mainly due to the simulated return dynamics, which do not exactly conform to the GP dynamics, and our adjustments to adapt the model to our (log-normal) environment appear insufficient. Table 3 underlines the amount of improvement introduced with the best linear policy. In this case, myopic policies perform significantly worse than the best linear

Table 3: **Policy Performance: No common factor noise and high t-cost environment.**

This table summarizes the performance of each policy in the case of no common factor noise and high transaction cost environment. For each policy, we report average terminal wealth, average objective value, variance of the terminal wealth, average terminal Sharpe ratio in the presence and absence of transaction costs and average weekly Sharpe ratio in the presence of transaction costs. (Dollar values are in thousands of dollars.)

	GP	**MP**	**MP-TC**	**RBL**	**BL**	**NTC**
Avg Wealth	114.2	180.9	52.19	74.76	232.2	594.3
Avg Objective	85.61	-98.41	25.31	59.58	138.1	297
Variance	8.96e+03	1.72e+05	2.28e+04	3.77e+03	3.14e+04	1.44e+05
TC	23.59	7.744	0.2971	44.43	43.97	0
Sharpe with TC	1.706	0.6168	0.4886	1.722	1.853	2.215
Sharpe w/o TC	1.753	0.5421	0.4896	2.222	1.94	2.215
Weekly Sharpe with TC	2.30	2.132	2.098	2.003	2.517	3.453

Table 4: **Policy Performance: Common factor noise and low t-cost environment.**

This table summarizes the performance of each policy in the case of common factor noise and low transaction cost environment. For each policy, we report average terminal wealth, average objective value, variance of the terminal wealth, average terminal Sharpe ratio in the presence and absence of transaction costs and average weekly Sharpe ratio in the presence of transaction costs. (Dollar values are in thousands of dollars.)

	GP	**MP**	**MP-TC**	**RBL**	**BL**	**NTC**
Avg Wealth	53.14	38	39.44	19.3	39.23	41.81
Avg Objective	-336.9	15.38	19.28	9.785	20.51	20.75
Variance	7.93e+04	2.25e+04	4.07e+03	2.04e+03	9.19e+03	4.21e+03
TC	7.57	1.186	0.98	0.2683	1.785	0
Sharpe with TC	0.27	0.3586	0.87	0.604	0.5786	0.911
Sharpe w/o TC	0.30	0.8848	0.9	0.6121	0.586	0.911
Weekly Sharpe with TC	0.40	0.9058	0.92	0.7347	0.8756	0.9436

policy.

Table 4 and Table 5 depict the impact of common factor shocks in the terminal wealth statistics. It is important to note that in this regime, Sharpe ratios are significantly lower. In both cases, the best linear policy achieves the best objective value statistics.

3.5 Discussion of the Trading Rules

In this section, we illustrate how the BL and GP strategies differ in terms of their trading strategy in the two different environments studied above (i.e., with factor risk and without).

Table 5: **Policy Performance: Common factor noise and high t-cost environment.** This table summarizes the performance of each policy in the case of common factor noise and low transaction cost environment. For each policy, we report average terminal wealth, average objective value, variance of the terminal wealth, average terminal Sharpe ratio in the presence and absence of transaction costs and average weekly Sharpe ratio in the presence of transaction costs. (Dollar values are in thousands of dollars.)

	GP	MP	MP-TC	RBL	BL	NTC
Avg Wealth	18.3	14.52	15.66	9.822	16.21	41.81
Avg Objective	-11.3	3.77	5.68	5.851	9.133	20.75
Variance	5.97e+03	4.94e+04	1.07e+04	8.12e+02	1.63e+03	4.21e+03
TC	5.27	1.916	2.91	1.881	2.059	0
Sharpe with TC	0.33	0.2919	0.21	0.4873	0.5674	0.911
Sharpe w/o TC	0.42	0.477	0.57	0.5738	0.6065	0.911
Weekly Sharpe with TC	0.43	0.5267	0.57	0.619	0.7386	0.9436

For simplicity we consider that there are only two return generating factors: momentum and value. Otherwise the return dynamics are as assumed in the previous section equation (64) for the economy with factor risk and equation (65) for the economy without factor risk.

We perform the following experiment. We start each predictor variable B_{mom}, B_{value} at its long run mean of zero at time 0 for each stock, and at time 0 we assume the idiosyncratic risk of one stock, i say, experiences a +1std deviation shock, i.e,. $\epsilon_i(0) = \sigma_i$. All other stocks are unaffected. Then we plot the time series of returns for stock i, and the trades, and position in stock i over time, assuming that at all future dates all realizations of future shocks $\epsilon_j(t) = 0 \ \forall t, j$ and $F_{mom}(t) = F_{val}(t) = 0$. This implies that for this path, realized returns are equal to expected returns.

Figure 1 plots the path of realized return on stock i for this experiment. A shock to idiosyncratic risk of stock i in period zero increases its exposure to the momentum factor which has a positive expected return, but it also increases its exposure to the value factor which has a negative expected return. The two exposures are plotted in the second panel and the net (expected which is equal to realized for this experiment) return that obtains in the first panel. The momentum exposure has a much shorter half-life, which translates into the value factor dominating in the longer term. As a result the net return becomes more and more negative as momentum fades away. The dollar trades and corresponding cumulative positions in stock i of four different strategies are plotted in panel

3 and 4. It is clear from the picture that for that path, GP and BL are almost exactly the same. One would expect that, since apart from modeling price changes versus dollar returns, the setting is very similar to that used in GP and for that setting their strategy is by definition optimal. This picture thus comforts us in confirming that the BL strategy can essentially closely match the first best, for the case where such a first best can be computed.

Furthermore, the picture also shows that a simple adjustment of the myopic policy, which is often used in practice (MP-TC which is described in the previous section) results in a cumulative exposure to stock i that is very close to the optimal one. This is due to the fact that we are plotting the trades in a low transaction cost environment, where as we showed in the previous section, the MP-TC strategy performs fairly well compared to BL and GP .

Instead, Figure 2 plots realized returns, trades and cumulative position in stock i for the economy with factor risk. The picture shows that there are substantial differences between the three strategies BL , GP , and MP-TC. Specifically, BL recognizes that the increase in expected returns also comes with a lot more systematic risk, and therefore trades least aggressively of all strategies. Indeed, if we recomputed the GP strategy allowing it to modify the covariance matrix and to use the current conditional covariance matrix instead of the unconditional average covariance matrix, then GP trades a lot more similarly to BL as shown in Figure 3. This explains the source of the improvement achieved by the BL strategy over the GP strategy observed in the previous section. The BL strategy recognizes that both conditional expected returns and the covariance matrix are changing over time, and adjusts its trading optimally to both. Instead the GP approach which requires a deterministic covariance matrix cannot react to stochastic changes in the covariance matrix.

4 Trading the Short-Term Reversal Effect

The short-term reversal anomaly was noted in Fama (1965), and was later explored more fully in Jegadeesh (1990) and Lehmann (1990). This negative serial correlation is generally interpreted as evidence consistent with incomplete liquidity provision, and much empirical evidence is consistent

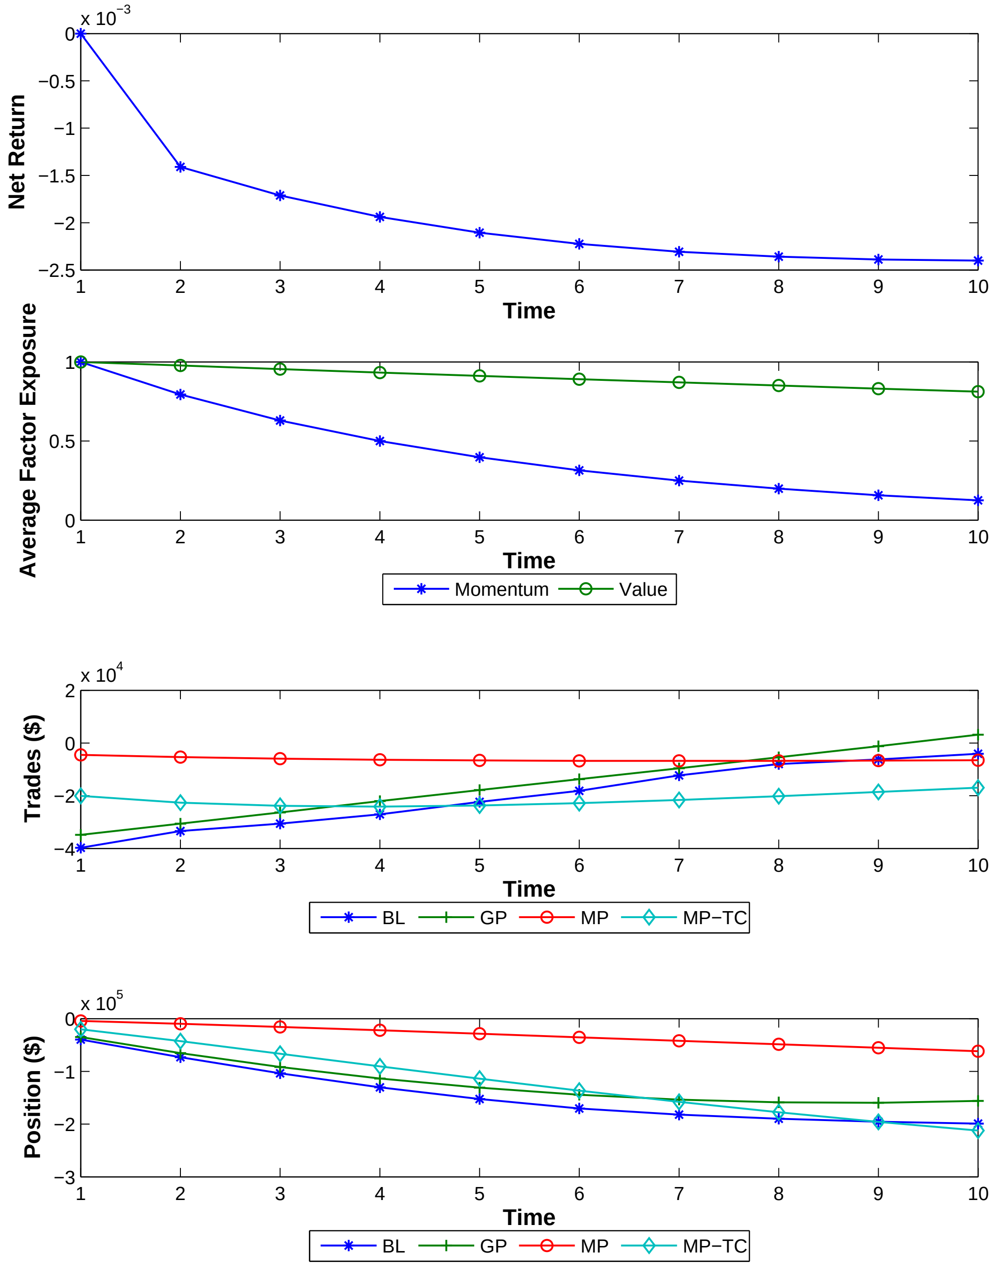

Figure 1: Path of expected returns (which is also equal to the actual realized returns) in an experiment in an economy without factor risk, where all future idiosyncratic shocks are set to zero except for time 0 and stock i which then experiences a one standard deviation idiosyncratic volatility shock. Lower panels represent the trades and cumulative position in stock i for various trading strategies described in the previous section.

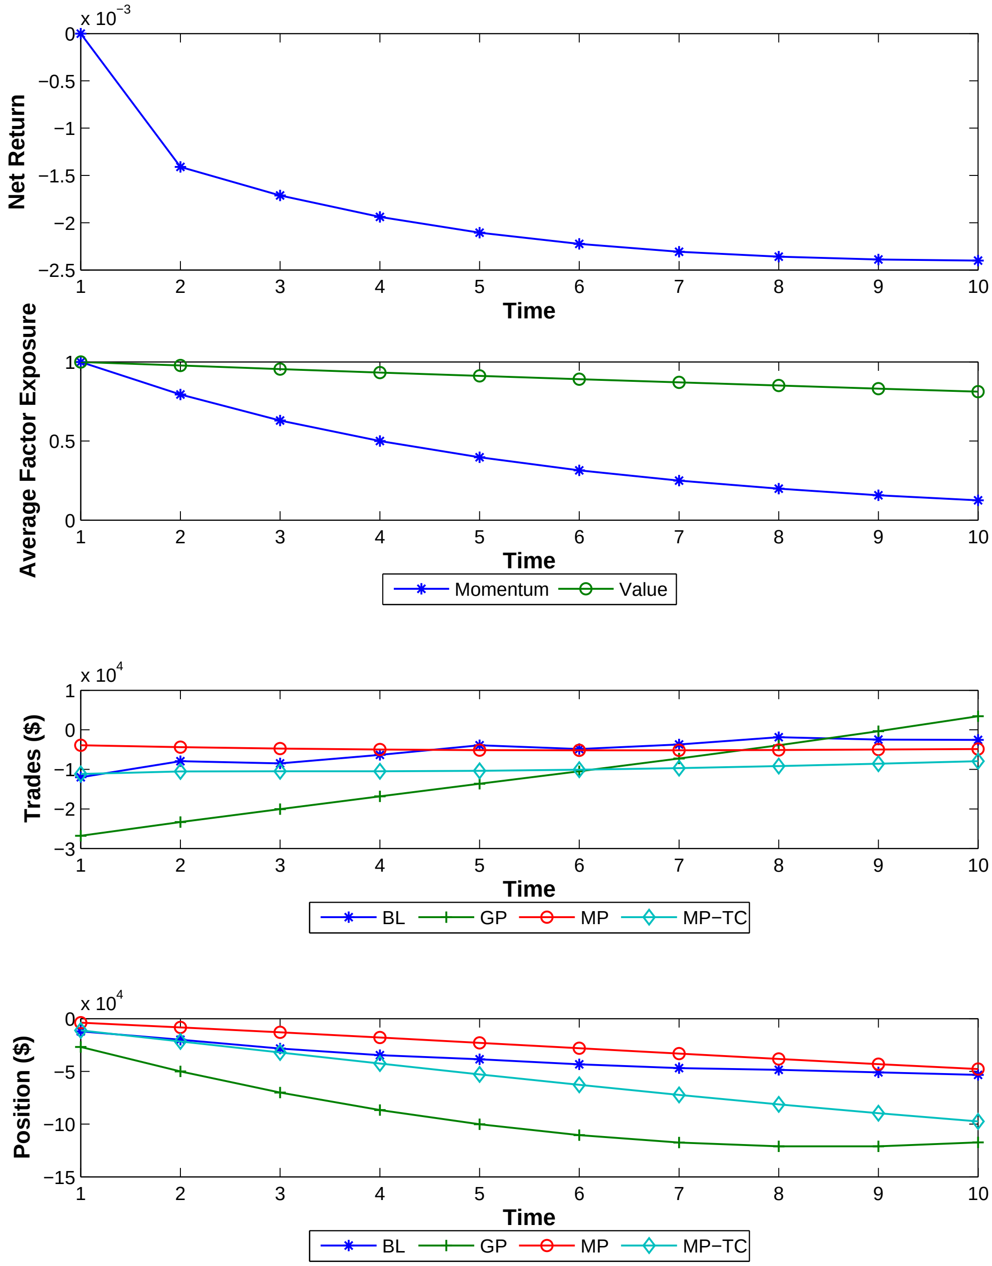

Figure 2: Path of expected returns (which is also equal to the actual realized returns) in an experiment in an economy with factor risk, where all future idiosyncratic as well as factor shocks are set to zero except for time 0 and stock i which then experiences a one standard deviation idiosyncratic volatility shock. Lower panels represent the trades and cumulative position in stock i for various trading strategies described in the previous section.

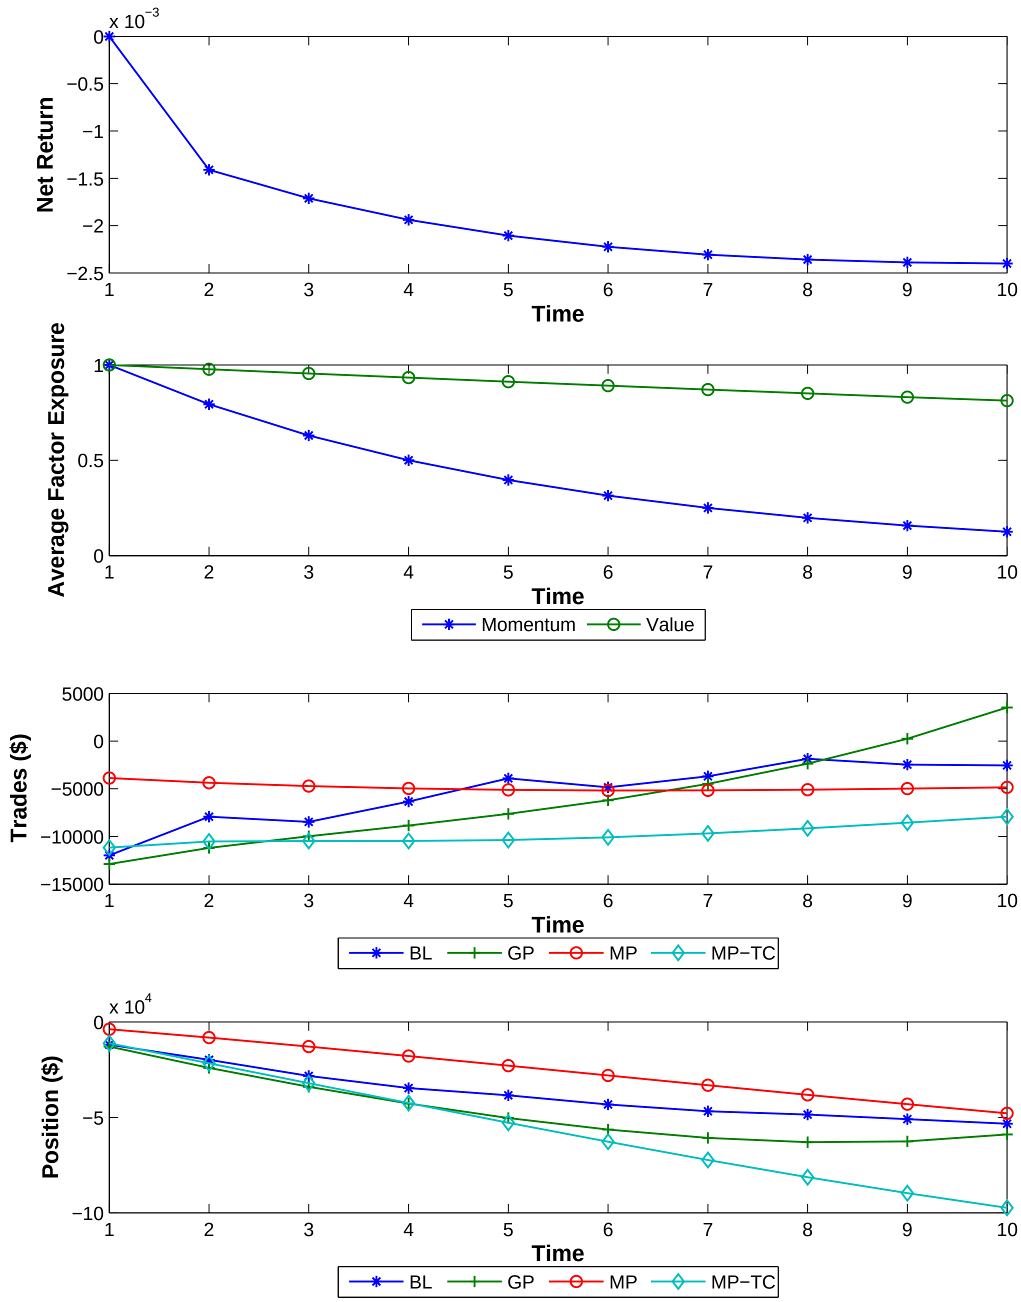

Figure 3: Path of expected returns (which is also equal to the actual realized returns) in an experiment in an economy with factor risk, where all future idiosyncratic as well as factor shocks are set to zero except for time 0 and stock i which then experiences a one standard deviation idiosyncratic volatility shock. Lower panels represent the trades and cumulative position in stock i for various trading strategies described in the previous section. In this case, we compute trades for the GP strategy with the updated conditional covariance matrix that obtains after the time 1 shock to factor returns instead of its time 0 unconditional average expected covariance matrix (as in Figure 2).

with this.[15] Collin-Dufresne and Daniel (2013, CD) develop a detailed dynamic model for the short-term reversal effect, and use this to examine the time variation in the reversal effect.

In this section we apply the LGS methodology we develop in Section 2 to a trading strategy based on the CD model and a simple model for transaction costs. We then apply this model to estimate the effectiveness of an LGS-based strategy using historical equity returns and simulated transaction costs, and compare the performance of the LGS strategy with other approaches suggested in the literature.

We include this for several reasons: first, we demonstrate the steps necessary for implementing the LGS methodology in a setting like that would be faced by an investor. Second, this gives us an opportunity to compare the performance of the LGS methodology to other candidate methodology using actual rather than simulated data.

The short-term-reversal strategy is a good test-bed for dynamic-portfolio-choice methods. The strategy returns, gross of transaction costs, are large. But the returns associated with short-term-reversal die out quickly: CD show that the half life associated with this decay averages 2.4 days over this sample period.

4.1 Sample and Data

We implement our LGS-based trading strategy on the 100 largest firms by market capitalization in the CRSP universe, over the period from January 1974-March 2013. At close on the last trading day of each calendar year we select the 100 firms with the largest market capitalization. We then trade those firms over the next calendar year. What this means is that our cross-section is the very largest firms in the CRSP universe. We do this to ensure that there is high liquidity in each of the stocks in our sample on each trading day.

Because we trade only the largest firms in the CRSP universe, every one of the firms in our

[15] There is now a voluminous academic literature on this topic: Chan (2003) and Tetlock (2011) show that large stock price moves exhibit more reversal if they are not associated with news and Da, Liu, and Schaumburg (2013) show that within industry/residual based reversals have a far higher Sharpe ratio. In addition, Avramov, Chordia, and Goyal (2006) argue that the reversal effect is present only in small, illiquid stocks with high turnover, and Khandani and Lo (2007) document a dramatic decline over time in the efficacy of the strategy. Finally, the strength of the reversal strategy appears to depend on arbitrageurs ability to access capital: Nagel (2012) documents that the return of a \$1-long/\$1-short short-term-reversal portfolio is far higher when market volatility is higher. He attributes this time variation to limited investor capital in times of high volatility.

sample has a fairly liquid market. Almost every one of the firms in our sample trades every day. There are a relatively small number of delistings, and for those firms that are delisted, a reliable delisting return is available.[16]

Note the firms in our sample often change at close on the last trading day of each calendar year. If a firm leaves our sample (because it is no longer one of the 100 largest firms) we close out our position at the closing price on that date. Similarly, if a firm enters the sample of the 100 largest firms, we allow our simulation to trade into that firm starting at close on the last trading day of the year.

4.2 Model Specification

Our baseline specification for the excess market return ($\tilde{r}_{m,t+1}$) and the individual firm excess returns ($\tilde{r}_{i,t+1}$) is:

$$\tilde{r}_{m,t+1} = \mu + h_{m,t}\tilde{v}_{t+1}, \qquad \tilde{v}_{t+1} \sim \mathcal{N}(0,1) \tag{71}$$

$$\tilde{r}_{i,t+1} = \beta_{i,m}\tilde{r}_{m,t+1} + \underbrace{B_{i,t}\lambda_t + \sigma_i h_{\epsilon,t}\tilde{\epsilon}_{i,t+1}}_{\equiv \tilde{u}_{i,t+1}} \qquad \tilde{\epsilon}_{i,t+1} \sim \mathcal{N}(0,1) \tag{72}$$

where $E_t[\tilde{v}_{t+1}\tilde{\epsilon}_{i,t+1}] = 0 \;\forall\; i, E_t[\tilde{\epsilon}_{i,t+1}\tilde{\epsilon}_{j,t+1}] = 0 \;\forall\; i \neq j$, and where $\frac{1}{n}\sum_{i=1}^{n}\sigma_i^2 = 1$.

The market return $\tilde{r}_{m,t+1}$ has a constant mean return and time varying volatility. We specify that the market volatility follows a Glosten, Jagannathan, and Runkle (1993) GARCH process:

$$h_{m,t}^2 = \omega + \beta h_{m,t-1}^2 + (\alpha + \gamma \mathrm{I}[(\tilde{r}_{m,t} - \mu) < 0])\, u_{m,t}^2 \tag{73}$$

Individual firm excess return $\tilde{r}_{i,t+1}$ have a single common factor (the market). Firm i has a loading of $\beta_{i,m}$ on the market. $u_{i,t+1}$ denotes firm i's residual return at time $t+1$. Unlike in standard models, it is convenient for us to specify that the residual return $u_{i,t+1}$ is not mean zero. This is consistent with the literature showing that residual returns exhibit reversal.

[16]In tests of our trading strategy, if a firm is delisted, up until the delisting takes place, we trade as if we are unaware that the delisting will take place. On the delist date, we assume that any holdings, long or short, of that firm's shares earn the return on that date plus the CRSP delisting return.

4.2.1 Mean-return specification

Equation (72) shows that firm i's expected residual return is the product of a firm specific "exposure" $B_{i,t}$ and a common "premium" λ_t. A firm's exposure ($B_{i,t}$) is governed by an auto-regressive process:

$$B_{i,t} = \beta_r B_{i,t-1} + (1-\beta_r)\tilde{u}_{i,t-1} = (1-\beta_r)\sum_{l=1}^{\infty} \beta_r^l u_{i,t-l} \tag{74}$$

Each firm's time t exposure is an exponentially weighted sum of its lagged daily residuals.[17] Collin-Dufresne and Daniel (2013, CD) show that, following a positive residual shock to its price, the price of a firm generally declines exponentially over the next few weeks. CD's estimation of equation (74) gives a $\hat{\beta}_r = 0.720$, corresponding to a half-life of 2.4 days.

Again, according to the specification in equation (72), the *ex-ante* expected excess return $E_t[\tilde{u}_{i,t+1}]$ is the product of the exposure $B_{i,t}$ and the *ex-ante* expected premium (per unit of exposure) λ_t. λ_{t+1} is updated each period based on the *ex-post* premium in period $t+1$, which we denote with $\lambda_{t+1}^{\mathrm{xp}}$, which is the OLS estimator of $\lambda_{t+1}^{\mathrm{xp}}$ in the cross-sectional regression equation:

$$\tilde{u}_{i,t+1} = \lambda_{t+1}^{\mathrm{xp}} B_{i,t} + \sigma_i h_{\epsilon,t} \tilde{\epsilon}_{i,t+1},$$

where $\tilde{u}_{i,t+1}$ is the dependent variable and $B_{i,t}$ the independent variable. This estimator is given by:

$$\hat{\lambda}_t^{\mathrm{xp}} = \left(\sum_i B_{i,t-1}^2\right)^{-1} \sum_i B_{i,t-1}\tilde{u}_{i,t}.$$

CD specify that the *ex-ante* premium per unit exposure at time $t+1$, λ_{t+1}, is a weighted average of the time t premium and the *ex-post* realized premium at time $t+1$: evolves according to the following process:

$$\lambda_{t+1} = \alpha\lambda_t + (1-\alpha)\,\lambda_{t+1}^{\mathrm{xp}} \tag{75}$$

CD estimate the AR(1) coefficient in equation (75) to be $\hat{\alpha} = 0.99748$, corresponding to a half life of 274 days, suggesting that the premium varies very slowly over time.

[17] With this specification, the expected residual return at time $t+1$ is a function of residuals at times $t-1$ and before. This is, at least in part, to avoid potential bid-asked bounce effects. Note that every firm in our sample trades (almost) every day. See Collin-Dufresne and Daniel (2013) for details.

4.2.2 Volatility Specification

We specify in equation (72) that the volatility of firm i's residual return is $\sigma_i h_{\epsilon,t}$ – the product of a time-invariant but firm-specific term σ_i, and the (common) level of cross-sectional volatility, $h_{\epsilon,t}$.

We'll first discuss the common level of idiosyncratic volatility, $h_{\epsilon,t}$. Recall that the CD specification incorporates the restriction that $\frac{1}{n}\sum_{i=1}^{n}\sigma_i^2 = 1$. Given this restriction,

$$E_t\left[\sum_{i=1}^{n}\left(\tilde{u}_{i,t+1} - B_{i,t}\lambda_t\right)^2\right] = E_t\left[\sum_{i=1}^{n}\left(\sigma_i h_{\epsilon,t}\tilde{\epsilon}_{i,t+1}\right)^2\right] = h_{\epsilon,t}^2$$

Thus $h_{\epsilon,t}$ is simply the cross-sectional volatility of individual firm residual returns.

Our specification is consistent with Kelly, Lustig, and Van Nieuwerburgh (2012), who argue that time variation in individual firm idiosyncratic volatilities are largely captured by a single factor structure. Here, all changes are a result of changes in the common level $h_{\epsilon,t}$.

We specify the dynamics of $h_{\epsilon,t}$ as:

$$h_{\epsilon,t}^2 = \kappa_\epsilon + \alpha_\epsilon h_{\epsilon,t-1}^2 + \mu_\epsilon(\hat{h}_{\epsilon,t}^*)^2 \tag{76}$$

where $\hat{h}_{\epsilon,t}^*$ is the realized cross-sectional volatility at time t:

$$(\hat{h}_{\epsilon,t}^*)^2 \equiv \sum_{i=1}^{n}\left(\tilde{u}_{i,t+1} - B_{i,t}\lambda_t\right)^2.$$

CD estimate $\hat{\alpha}_\epsilon = 0.75$, corresponding to a half-life of 2.4 days.

For an individual firm, $E_t\left[\left(\tilde{u}_{i,t+1} - B_{i,t}\lambda_t\right)^2\right] = \sigma_i^2 h_{\epsilon,t}^2$. That is, σ_i^2 is the ratio of a firm's residual variance to the average residual variance of the n firms in our sample. Following CD, in our empirical implementation, we estimate σ_i at the beginning of each year, and then hold it fixed from the first trading day of the year through the last, consistent with the specification in equation (72). Empirically, CD find considerable cross-sectional variation in σ_i.

Table 6: **Maximum Likelihood Estimates of Market Process Parameters**
This table presents the estimates, standard errors and t-statistics from the joint estimation GJR-GARCH process (equations (71) and (73)). All parameter estimates are obtained from an iterative ML procedure run on daily market returns, where the market is defined as the equal-weighted average of the returns of the 100 largest firms, measured at the beginning of each year. The sample period is January 2, 1974 through March 28, 2013.

Joint Estimation of Equations (71) and (73)

Variable	ML. Est.	std. err.	t-stat
$\mu^\dagger$	2.4800	0.8405	2.9506
$\omega^\dagger$	1.5403	0.3657	4.2115
α	0.0156	0.0043	3.6336
γ	0.0992	0.0170	5.8293
β	0.9215	0.0108	85.3438

†The coefficients and std. errors for μ are $\times 10^4$, and for ω are $\times 10^6$.

4.3 Model Estimation

In this section we present the results of the estimation of the reversal model described in Section 4.2 for our sample of the 100 largest CRSP firms, and over our sample period of 1974:01-2013:03. The details of the model estimation are reported in Collin-Dufresne and Daniel (2013).

We begin with the estimation of equations (71) and (73), which were estimated jointly using a numerical maximum likelihood procedure. The results of this estimation are reported in Table 6

We also estimate the parameters of the individual firm process (equations (74), (75) and (76)) using an iterative MLE procedure, and report the results of this estimation in Table 7.

In the next section, we develop a LGS based on: (1) the process specification laid out in Section 4.2, (2) the estimated parameters, (3) a simple model of transaction costs.

4.4 An LGS-based methodology applied to short-term reversal

In this section we document the construction of the objective function for the short term reversal trading strategy based on the methodology developed in Section 2, and the estimated model of short-term-reversal described in Sections 4.2 and 4.3.

4.4.1 Objective function

We use an objective function similar to that used in recent literature as discussed in section 2.3.1.

Table 7: **Maximum Likelihood Estimates of Individual Firm Model Parameters**
This table presents the estimates, standard errors and t-statistics for the model parameters in equations (74), (75) and (76). All estimates are obtained from an iterative ML procedure run on daily returns from the 100 largest market capitalization at the start of each year. The sample period is January 2, 1974 through March 28, 2013. Note that $\hat{\lambda}_0$ is the starting value of λ_t that maximizes the the likelihood function.

param.	ML-est.	std. err	t-stat
$\hat{\beta}_r$	0.719534		
λ_0	-0.073254	0.126814	-0.57
α	0.997574	0.003025	329.82
$\kappa_\epsilon^\dagger$	5.871475	0.311785	18.85
α_ϵ	0.749041	0.017648	42.44
μ_ϵ	0.259744	0.018973	13.69

†The coefficient and std. error for κ_ϵ are $\times 10^{10}$.

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

where $\Sigma_{t\to t+1} = h_{m,t}^2 \beta\beta^\top + h_{e,t}^2 \operatorname{diag}(\sigma_1, \ldots, \sigma_N)$ from the dynamics of the security returns. With this objective function we assume that transaction cost of trading a particular asset varies linearly with the variance of only this asset's return but does not depend on the variance of the remaining stock returns (i.e., no cross-impact). We multiply $\Sigma_{t\to t+1}$ in the transaction cost term by I_N in order to ensure this, i.e., diagonal transaction cost matrix. Note that this is chosen for brevity as our methodology is general enough to incorporate any quadratic function, deterministic or stochastic.

We calibrate η to be 5×10^{-4} as in the simulation experiment. Note that as η approaches to zero, the complexity of the problem decreases drastically and myopic policy becomes optimal. We illustrate the empirical results of the no-transaction-cost scenario in Section 4.5. Finally, we assume that the coefficient of risk aversion, γ, equals 5×10^{-8} which we can think of as corresponding to a relative risk aversion of 1 for an agent with 500 millions under management.

4.4.2 Policies

We compare the gains from trading according to myopic policy and restricted best linear policy.

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,

$$x_t^{\mathsf{MP}} = \left(I_N \Sigma_{t\to t+1} + \gamma \Sigma_{t\to t+1}\right)^{-1} \left(B_t \lambda_{0,t} + \Lambda \left(x_{t-1}^{\mathsf{MP}} \circ R_t\right)\right),$$

where $\Sigma_{t\to t+1} = h_{m,t}^2 \beta\beta^\top + h_{e,t}^2 \operatorname{diag}(\sigma_1, \ldots, \sigma_N)$. Note that in the absence of transaction costs, x_t^{MP} simplifies to $\frac{1}{\gamma}\Sigma_{t\to t+1}^{-1} B_t \lambda_{0,t}$.

We will compare the myopic policy with restricted best linear policy that uses factor information of two most recent trading days. Formally, we solve for the optimization

$$\begin{aligned}
&\max_{\theta,\pi} \sum_{t=1}^{T} \mathsf{E}\left[x_t^\top r_{t+1} - \frac{\gamma}{2} x_t^\top \Sigma_{t\to t+1} x_t - \frac{\eta}{2} u_t^\top I_N \Sigma_{t\to t+1} u_t\right], \\
\text{subject to } & x_{i,t} = \theta_{i,t} \frac{B_{i,t}\lambda_{0,t}}{h_{\epsilon,t}}, \\
& u_{i,t} = \pi_{i,1,t} \frac{B_{i,t-1}\lambda_{0,t-1} R_{i,t}}{h_{\epsilon,t-1}} + \pi_{i,2,t} \frac{B_{i,t}\lambda_{0,t}}{h_{\epsilon,t}}, \\
& \pi_{i,1,t} = \theta_{i,t} \text{ for } t > 1, \\
& \pi_{i,2,t} = -\theta_{i,t-1} \text{ for } t > 1, \\
& \pi_{i,1,1} = 0.
\end{aligned}$$

Note that since idiosyncratic volatilities are stochastic and linear on $h_{\epsilon,t}$, we scale our trading policy with $h_{\epsilon,t}$. We do not need to put the exact volatility but just the stochastic part as the constants are picked up by the optimal parameters, θ and π.

4.4.3 Results

We run the performance statistics of our myopic and restricted best linear policy for the short-term reversal experiment. Table 8 illustrates that the restricted best linear policy outperforms the myopic policy in terms of average objective value and Sharpe ratios for the one-year and daily trading profits.

Table 8: **Short-Term Reversal Experiment: Policy Performance.**

For the myopic policy (MP) and the restricted best-linear policy (RBL), we report average terminal wealth, average objective value, Sharpe ratio for terminal wealth at the end of each trading year and Sharpe ratio for daily trading gains. (Dollar values are in thousands of dollars.)

	MP	**RBL**
Avg Wealth	6223.01	7126.09
Avg Objective	4339.73	5210.56
TC	50.94	429.33
Sharpe Ratio of Terminal Wealth	0.80	0.92
Sharpe Ratio of Daily Profit	0.85	1.01

Table 9: **Short-Term Reversal: Performance when Transaction Costs are Zero.**

This table provides summary statistics for the myopic policy (MP) and the restricted best-linear policy (RBL) for the short-term reversal experiment then transaction costs are zero. For each policy, we report average terminal wealth, average objective value, Sharpe ratio for terminal wealth at the end of each trading year and Sharpe ratio for daily trading gains. Dollar values are in thousands of dollars.

	MP	**RBL**
Avg Wealth	646481.09	412498.28
Avg Objective	332172.64	220061.32
TC	15114.00	10125.14
Sharpe Ratio of Terminal Wealth	1.39	1.25
Sharpe Ratio of Daily Profit	3.55	3.17

4.5 Short-term Reversal Experiment - No Transaction Costs

To check the importance of transaction costs for trading a signal that has such a short half life, we compare the performance statistics of a myopic strategy to the restricted best linear policy for the short-term reversal experiment in the absence of transaction costs. In this scenario, we know that myopic policy is optimal so we are certain that restricted best linear policy will be sub-optimal. Table 9 illustrates that restricted best linear policy indeed underperforms the myopic policy terms of average objective value and Sharpe ratios for the one-year and daily trading profits. Comparing the results in this table to the previous one reveals how important transaction costs are when trading on such a signal. Daily Sharpe ratios drop by a factor of three.

5 Conclusion and Future Directions

The LGS framework we propose here accommodates complex return predictability models studied in the literature in multiperiod models 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, 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 LGS performs similarly to the linear-quadratic solutions of Litterman (2005) Gârleanu and Pedersen (2012) when the return generating process is restricted to be homoscedastic and transaction costs are deterministic. However, the LGS framework performs better when returns display stochastic volatility, e.g., when they are driven by a factor model. We also investigate the performance of the LGS framework when trading a high turnover strategy based on return reversal, for which transaction costs are a first order concern. The benefits to using our dynamic framework are large 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.

Avramov, Doron, Tarun Chordia, and Amit Goyal, 2006, Liquidity and autocorrelations in individual stock returns, *The Journal of Finance* 61, 2365–2394.

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, 1991, 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, D. B., and J. E. Smith, 2010, Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds, Working paper.

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.

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.

Chan, Wesley S., 2003, Stock price reactions to news and no-news: Drift and reversal after headlines, *Journal of Financial Economics* 70, 223–260.

Collin-Dufresne, Pierre, and Kent Daniel, 2013, Liquidity and return reversals, Columbia GSB working paper.

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.

Cvitanić, 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).

Da, Zhi, Qianqiu Liu, and Ernst Schaumburg, 2013, A closer look at short-term return reversal, *Management Science,* forthcoming.

Davis, M. H. A., and A. R. Norman, 1990, Portfolio selection with transaction costs, *Mathematics of Operations Research* 15, 676–713.

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., 1965, The behavior of stock-market prices, *The Journal of Business* 38, 34–105.

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.

Gârleanu, Nicolae, and Lasse H. Pedersen, 2012, Dynamic trading with predictable returns and transaction costs, Haas School of Business working paper.

Glosten, Lawrence R., Ravi Jagannathan, and David Runkle, 1993, On the relation between the expected value and the volatility of the nominal excess return on stocks, *Journal of Finance* 48, 1779–1801.

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.

Kelly, Bryan, Hanno N. Lustig, and Stijn G. Van Nieuwerburgh, 2012, The volatility factor structure, Chicago Booth Research paper 12-54.

Khandani, Amir E., and Andrew W. Lo, 2007, What Happened To The Quants in August 2007?, *Journal of Investment Management* 5, 5–54.

Le, Ahn, Kenneth J. Singleton, and Qiang Dai, 2010, Discrete-time affine$^{\mathbb{Q}}$ 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, 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.

Nagel, Stefan, 2012, Evaporating liquidity, *Review of Financial Studies* 25, 2005–2039.

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, 2011, The short of it: Investor sentiment and anomalies, *Journal of Financial Economics* 104, 288–302 Wharton School Working Paper.

Tetlock, Paul C., 2011, All the news that's fit to reprint: Do investors react to stale information?, *Review of Financial Studies* 24, 1481–1512.

Vasicek, Oldrich A., 1977, An equilibrium characterization of the term structure, *Journal of Financial Economics* 5, 177–188.