Machine Learning about Optimal Portfolios

Theis Ingerslev Jensen, Bryan Kelly, Semyon Malamud, and Lasse Heje Pedersen*

Not for Distribution

This version: October 17, 2021

Abstract

Numerous studies document that returns can be predicted using various factors or machine learning, but does this translate to implementable portfolios? We find empirically that standard algorithms for overall return prediction focus on fleeting small-scale characteristics (e.g., 1-month reversal for small stocks) with poor net returns. We show instead how to build a prediction model focused on economically important characteristics, and document that this method produces superior net-of-transaction-cost out-of-sample returns. The method can be implemented in dual ways: either by combining predictions across multiple forecast horizons or by directly penalizing the objective function. The model gives rise to a new measure of "economic feature importance."

Keywords: asset pricing, factors, transaction costs, economic significance

JEL Codes: C5, C61, G00, G11, G12

*Jensen is at Copenhagen Business School. Kelly is at Yale School of Management, AQR Capital Management, and NBER; www.bryankellyacademic.org. Malamud is at Swiss Finance Institute, EPFL, and CEPR, and is a consultant to AQR. Pedersen is at AQR Capital Management, Copenhagen Business School, and CEPR; www.lhpedersen.com. AQR Capital Management is a global investment management firm, which may or may not apply similar investment techniques or methods of analysis as described herein. The views expressed here are those of the authors and not necessarily those of AQR.

1 Introduction

This paper studies how security information can be used for portfolio selection in a flexible and realistic setting with transaction costs. The goal is thus both to provide a powerful tool for portfolio choice and to shed new light on which security characteristics are economically important drivers of asset pricing.

Machine learning (ML) provides a flexible framework to combine several characteristics into a single measure of overall expected returns (e.g., Chen et al. (2020) and Gu et al. (2020)). At the same time, several ML give insights into each characteristic's relative "feature importance." However, standard ML predictors of short-run returns are often largely driven by short-lived characteristics that work well for small illiquid stocks (Avramov et al. (2021)). Indeed, such fleeting small-scale characteristics appear most prominently among standard measures of feature importance and, as a result, the implied portfolio strategies are associated with large transaction costs.

This paper seeks to generate predictions that are more economically useful. In particular, we develop a method designed to take into a account the turnover required for a strategy and the transaction costs of the securities that it trades. Further, we construct a measure of economically motivated feature importance that captures which characteristics provide most economic value. We find that the method produces significant out-of-sample performance improvement relative to a range of benchmarks and yields new insights on economically importance stock characteristics.

As a way of background, the goal of standard models of ML in finance is to find a function f of the characteristics x_{it} of each stock i that best explains its future returns, $r_{i,t+1} = f(x_{it}) + \varepsilon_{i,t+1}$. Clearly, this objective abstracts from transaction costs and turnover. As a result, investment strategies based on this approach produce negative returns net of transaction costs.

A more sophisticated approach uses these estimated expected returns, $E_t(r_{i,t+1}) \cong f(x_{it})$, as in input into a second-stage optimization that takes transaction costs into account. We show theoretically and empirically that this two-stage approach is also not optimal. To understand why, note that the first-stage ML procedure produces expected returns that

are dominated by short-term signals, and this method does not take into account which predictors are persistent and which have quick alpha decay. The second-stage optimization reduces turnover, especially for small stocks, which leads to a much better performance than portfolios that ignore transaction costs. However, this static approach can nevertheless be improved by recognizing the dynamic nature of the portfolio portfolio and considering expected returns of several return horizons or, equivalently, using information on alpha decay.

To improve upon this approach based on economically founded ideas, we develop a transaction-cost-aware ML approach that explicitly recognizes the dynamic alpha decay. There are two ways to implement our approach. First, we show how the objective function can be modified so that the ML algorithm directly estimates optimal portfolio weights, taking stocks' transaction costs and the strategy's turnover into account. As a second way of implementing the method, we show how the whole term structure of future expected returns can be used to compute the optimal portfolio. Indeed, with transaction costs, whatever you buy today, you will likely own for a while as you slowly exit this position, and, therefore, understanding the expected return both now and in the future is relevant.

Empirically, we find that the optimal ML predictor of near-term returns is indeed different from the optimal ML predictor of returns far into the future. In other words, if f_h is the function that best predicts returns h months into the future, $E_t(r_{i,t+h}) = f_h(x_{it})$, then this function is different across h. Given that the standard ML approach uses only f_1, we see that it misses the information contained in f_h at other horizons.

As a result, we find that our economically-motivated ML approach produces economically significant performance gains relative to standard ML approaches and even relative to two-stage methods that use standard ML approaches as an input into a second-stage transaction-cost optimization.

Lastly, we show how measures of feature importance change when we take transaction costs into account. Standard methods of machine learning about 1-month returns imply a large feature importance of signals such as 1-month return reversal. In contrast, our method, which seeks to optimally blend return predictability across multiple future horizons, especially for liquid stocks, implies a high feature importance of value and quality indicators.

Our paper is related to several large literatures. Indeed, portfolio selection has been at

the heart of financial economics since Markowitz (1952). Portfolio selection with transaction costs was studied by Constantinides (1986) and Davis and Norman (1990) in settings where returns are not predictable. Settings with predictable returns and transaction costs are studied in Gârleanu and Pedersen (2013) who derive an explicit solution when returns are driven by a factor model. Gârleanu and Pedersen (2016) consider more general dynamics in continuous time and Collin-Dufresne et al. (2020) extend the model to include different liquidity regimes. We present a model based on stationary dynamics (rather than dynamic programming as in the literature), allowing a general functional form for return predictability and for portfolio weights, thus creating a model well suited to ML implementation.

Our paper is also related to the large literature on return-predicting factors, summarized by Hou et al. (2020). Hou et al. (2020) challenge the replicability of this literature, but Chen and Zimmermann (2020) show that most factors can be replicated using pure replication, i.e., using the same data and methods as the original papers. Further, Jensen et al. (2021) find that most factors survive scientific replication, i.e., hold up with a consistent robust methodology, providing further evidence of internal validity. Jensen et al. (2021) also provide evidence of external validity, documenting out-of-sample evidence over time and across many countries, and using state-of-the-art Bayesian multiple testing adjustments. Our empirical results rely on the characteristics in the latter paper.

Several recent empirical papers point out that trading strategies based on these factors, as well as the literature on machine learning in asset pricing, as cited above, involves large transaction costs in practice. This literature includes Li et al. (2020), Chen and Velikov (2021), and Detzel et al. (2021). We complement this literature by developing a flexible method of portfolio optimization that lends itself to ML and corresponding measures of feature importance that take transaction costs into account.

2 Theory

We consider an economy with N securities traded in discrete time indexed by $t = ..., -2, -1, 0, 1, 2, ...$ The return of security n from time $t-1$ to t in excess of the risk-free rate r_t^f is denoted $r_{n,t}$ and the vector of all assets' excess returns is denoted r_t.

An investor observes a number of characteristics (or signals) for each security, denoted $x_{n,t} \in \mathbb{R}^K$, for example each asset's valuation ratio, momentum, size, and so on. The characteristics of all assets are collected in the matrix $x_t \in \mathbb{R}^{N\times K}$, and we assume that x_t and r_t are stationary and ergodic. The signals x_t fully characterize the investor's information about returns such that $E_t[r_{t+1}] = \mu(x_t)$ and $\text{Var}_t[r_{t+1}r'_{t+1}] = \Sigma(x_t)$, where μ and Σ are bounded Borel-measurable functions with Σ being positive definite. The investor has absolute risk aversion $\gamma(x_t)$, and we use throughout the notation $\Sigma_\gamma(x_t) = \gamma(x_t)\Sigma(x_t)$ for the product of risk-aversion and risk, for brevity.

The investor faces transaction costs due to her market impact. Specifically, if the investor chooses a position of $w_t \in \mathbb{R}^N$ at any time t, then the change in position $\Delta w_t = w_t - w_{t-1}$ leads to a market impact cost of $\Delta w'_t \Lambda(x_t)\Delta w_t$, where $\Lambda \in \mathbb{R}^{N\times N}$ is symmetric and positive semi-definite such that transaction costs are non-negative.[1]

We next consider how the investor chooses her portfolio without taking transaction costs into account, that is, behaving as if $\Lambda = 0$. Then we show how to account for transaction costs.

2.1 Standard Approach: Predicting Returns without T-Costs

The standard approach in the literature is to use different methods to predict returns, and then use these return predictions in portfolio sorts. In particular, suppose that we write the standard ML prediction problem as seeking to find a function f of stock characteristics $x_{n,t}$ that minimizes the mean-squared forecast errors for future 1-period (say, 1-month) excess returns $r_{n,t+1}$:

$$\min_{f:\mathbb{R}^K\to\mathbb{R}} \frac{1}{TN}\sum_{n,t}[r_{n,t+1} - f(x_{n,t})]^2 \tag{1}$$

This standard approach generates a function that approximates the conditional mean, $f(x_{i,t}) \cong E[r_{i,t+1}|x_{i,t}]$, when returns are stationary across time and assets, and the number of observations is large.

[1] The symmetry of Λ is without loss of generality since, if we start with non-symmetric $\tilde{\Lambda}$, we can define $\Lambda = \frac{1}{2}(\tilde{\Lambda} + \tilde{\Lambda}')$ and note that $\Delta w'_t \Lambda(x_t)\Delta w_t = \Delta w'_t \tilde{\Lambda}(x_t)\Delta w_t$

2.2 Machine Learning about Portfolios without T-Costs

How can we change the objective function so that the ML algorithm directly seeks to learn about the optimal portfolio? To accomplish this goal, we introduce the following objective function:

$$\max_{F:\mathbb{R}^{N\times K}\to\mathbb{R}^N} \frac{1}{T}\sum_t \left[r'_{t+1}F(x_t) - \frac{1}{2}F(x_t)'\Sigma_\gamma(x_t)F(x_t)\right] \tag{2}$$

where $\Sigma_\gamma(x_t) = \gamma(x_t)\Sigma(x_t)$ is the product of the risk aversion and the conditional variance-covariance matrix. Note that F is a vector, whereas f in (1) was a scalar. Here, F is seeking to approximate the conditional mean-variance optimal portfolio, not the conditional mean. In other words, the portfolio objective (2) generates the function $f(x_t) \cong Markowitz_t$ when the sample T is sufficiently large, where the $Markowitz$ portfolio (or tangency portfolio) is given by

$$Markowitz_t \ = \Sigma_\gamma(x_t)^{-1}E[r_{t+1}|x_t] = \ \Sigma_\gamma^{-1}\mu(x_t)\,. \tag{3}$$

We allow that signals for all stocks help predict the risk and return of any one stock, which introduces a dependence across different stocks.[2]

2.3 Machine Learning about Dynamic Portfolios with T-Costs

We next consider the case in which transaction costs matter, that is, $\Lambda \neq 0$ is a strictly positive semi-definite matrix. We are looking for the optimal portfolio F_t that maximizes

$$\max_F \frac{1}{T}\sum_t \left[r'_{t+1}F_t - \frac{1}{2}F'_t\Sigma_\gamma F_t - \frac{1}{2}\Delta F'_t\Lambda\Delta F_t\right] \tag{4}$$

Here, we consider the case with constant risk aversion and risk, Σ_γ, and transaction cost, Λ, but Section 2.6 analyzes stochastic volatility and liquidity risk.

With transaction costs, the optimal portfolio at time t not only depends on the cur-

[2] A closely related objective function is $\min_{f:\mathbb{R}^{N\times K}\to\mathbb{R}^N}\sum_t[1 - r'_{t+1}f(x_t)]^2$, which also shows the link to the standard ML objective (1).

rent characteristics, x_t, it also depends on the past characteristics, x_s, where $s < t$. This history dependence arises because the investor optimally chooses to trade gradually over time. Therefore, we are looking for an optimal portfolio of the form $w_t = F(\vec{x}_t)$, where $F : \mathbb{R}^{H\times K\times N} \to \mathbb{R}^N$ is a bounded function and $\vec{x}_t = x_{t-H+1}, ..., x_t$ is the history of security characteristics.

Taking the limit $T \to \infty$ is the above, we get the following objective:

$$\max_{F_t=F(\vec{x}_t)} E\left[r'_{t+1}F_t - \frac{1}{2}F'_t\Sigma_\gamma F_t - \frac{1}{2}\Delta F'_t\Lambda\Delta F_t\right] \tag{5}$$

using the Ergodic Theorem, where unconditional expectations are evaluated under the stationary ergodic distribution of x_t. Formula (5) implies that we do not need to deal with any value functions as the objective is stationary.

Example: Connection to the standard objective. Suppose that returns are independent across assets such that $\Sigma_\gamma = \text{diag}(\sigma_i^2)$ and that the market impact function is also diagonal $\Lambda = \text{diag}(\Lambda_i)$. Further, suppose that the expected return, of stock n only depends on its own signals, $x_{n,t}$, and that this dependence is the same for all stocks, $E_t(r_{n,t+1}) = \bar{\mu}(x_{n,t})$. Then the solution to (5) has the form $F = (f, ..., f)$, where $f : \mathbb{R}^{H\times N} \to \mathbb{R}$ is a scaler function of each stocks own signals, $\vec{x}_{n,t} = x_{n,t-H+1}, ..., x_{n,t}$ is the history of security characteristics.

$$\max_f \sum_{n=1}^N E\left[r_{n,t+1}f(\vec{x}_{n,t}) - \frac{1}{2}\sigma_n^2[f(\vec{x}_{n,t})]^2 - \frac{1}{2}\Lambda_n[\Delta f(\vec{x}_{n,t})]^2\right] \tag{6}$$

When all securities have equal risk, which we normalize to 1, then the objective is equivalent to (dropping constant terms)

$$\min_f \sum_{n=1}^N E\left[(r_{n,t+1} - f(\vec{x}_{n,t}))^2 + \Lambda_n[\Delta f(\vec{x}_{n,t})]^2\right] \tag{7}$$

We see that this objective is closely related to the standard objective (1), but with some important differences. This objective is clearly related to (1) because both are looking for a function f that predicts returns in the sense of producing small squared errors as seen from

the first term in (7). What is new is that we are penalizing functions that generate large turnover as seen from the second term in (7). Further, to create a turnover-efficient portfolio, we must allow that the function f depends on past and as well current characteristics. □

To solve for the optimal portfolio strategy, we need the following result, which provides a kind of discount factor, denoted m.

Lemma 1 *The matrix Riccati equation*

$$m^2 \; - \; (2I + \Lambda^{-1}\Sigma_\gamma)m \; + \; I \; = \; 0 \tag{8}$$

has a unique solution such that all eigenvalues are positive and are below one, given by

$$m \; = \; I + \frac{1}{2}\Lambda^{-1}\Sigma_\gamma \; - \; \Lambda^{-\frac{1}{2}}\left(\Lambda^{-\frac{1}{2}}\Sigma_\gamma\Lambda^{-\frac{1}{2}} + \frac{1}{4}(\Lambda^{-\frac{1}{2}}\Sigma_\gamma\Lambda^{-\frac{1}{2}})^2\right)^{\frac{1}{2}}\Lambda^{\frac{1}{2}} \tag{9}$$

Furthermore, when $\Lambda \to 0$, we have

$$m \; = \; \Sigma_\gamma^{-1}\Lambda \; + \; O(\|\Lambda\|^2) \tag{10}$$

We are ready to show explicitly how the optimal portfolio depends on current and past signals, discounted based on the known m coming from Lemma 1, and the solution naturally depends on the *Markowitz* portfolio defined in (3).

Theorem 1 (Optimal strategy) *The solution to the portfolio problem (5) is time separable in the sense there exists functions f_s such that $F(\vec{x}_t) \; = \; \sum_{s=0}^{H} f_s(x_{t-s})$. When $H \to \infty$, this solution converges to*

$$F(\vec{x}_t) \; = \; (I - m)\sum_{s=0}^{\infty} m^s A(x_{t-s}) \tag{11}$$

where $A(x_t)$ is the "aim portfolio" at time t, given by

$$A(x) \; = \; (I - m)\sum_{\tau=0}^{\infty} m^\tau E[Markowitz_{t+\tau}|x_t = x] \tag{12}$$

To understand this economics of this result, note that the optimal portfolio $F_t = F(\vec{x}_t)$ can be written as

$$F_t \ = m\, F_{t-1} + (I-m)A(x_t) = F_{t-1} + (I-m)[A(x_t) - F_t] \tag{13}$$

In other words, the optimal portfolio at time t is the weighted average of the optimal portfolio in the earlier time period $t-1$ and the current aim portfolio A, as seen from the first part of this equation. Said differently, the optimal portfolio is the portfolio inherited from the previous time period plus a trade in the direction of the aim portfolio, as seen from the last expression.

To understand the economics of the aim portfolio A, recall first that $\mu(x_t) = E_t(r_{t+1})$ is the short-term expected return, so that $E_t(\mu(x_{t+\tau})) = E_t(E_{t+\tau}(r_{t+\tau+1})) = E_t(r_{t+\tau+1})$. Using this identity, the aim portfolio can be written as

$$A(x_t) \ = \ (I-m)\sum_{\tau=0}^{\infty} m^\tau \Sigma_\gamma^{-1} E_t[r_{t+1+\tau}] \tag{14}$$

So, interestingly, the aim portfolio depends on expected returns across all future time horizons. Specifically, the aim portfolio is the weighted average of all future Markowitz portfolios (as the coefficients sum to I).

The coefficients in the "backward average" (11) are the same as the coefficients in the "forward average" (12), in both cases powers of m. This identity is no coincidence. Both results from the optimal trading behavior (13). To see how, note first that, when the investor has always traded gradually toward the aim, then the current portfolio ends up being a weighted average of past aims, yielding (11). Second, when the investor foresees behaving like this, she knows that, whatever she buys today, she will only gradually sell in the future, therefore she buys a portfolio that works well over the long term, yielding (12).

From a machine learning perspective, Theorem 1 is a powerful result. To see the power of this result, note that the theorem tranfers the ML problem from looking for a general function F all *current and past* signals to a problem of looking for a function A of only the *current* signals. The enormous reduction in dimension vastly simplifies the problem.

Further, Theorem 1 shows that there are dual ways to use machine learning about optimal portfolios (A and F) with transaction costs:

Method 1: ML based on the term-structure of expected returns. One approach to apply Theorem 1 empirically is to first use standard ML techniques as in (1), but do this for a range of forecasting horizons, thus producing proxies for $E_t[r_{t+1+\tau}]$ for all τ. Using these forecasts, the aim portfolio is given by (14) and the portfolio at each time can be computed recursively using (13).

Method 2: Penalized ML about short-term returns. Alternatively, we can use (5) directly as our ML objective, let F depend on A based on (11), and then search for the function A that maximizes this objective. This method uses the insights that (i) we can focus on the aim portfolio A, and (ii) the objective should penalize transaction costs.

These dual approaches seek to use machine learning to learn the optimal portfolio (A and F) are related to the standard ML approach (1), which is focused on learning about expected returns (μ), but the difference is that our methods are directly designed to be aware of transaction costs in a theoretically consistent manner.

2.4 Dynamic vs. Static Handling of T-Costs

To further see the difference between the standard approach and our method in Section 2.3, let us consider how transaction costs might be implemented in the standard approach. In particular, one can use the standard approach to first estimate expected returns $f(x_{i,t}) \cong E[r_{i,t+1}|x_{i,t}]$ via (1), and then account for a transaction costs in a second step. While this two-step procedure fails to account for the dynamic nature of the problem, it serves as an interesting benchmark for our fully dynamic method. To see how this works, consider the problem of choosing an optimal portfolio w based on the existing portfolio w_{t-1} and the vector of expected returns $\mu_t = (f(x_{1,t}), ..., f(x_{N,t}))'$:

$$\max_{w \in \mathbb{R}^N} \left\{ w'\mu - \frac{1}{2} w'\Sigma_\gamma w - \frac{1}{2\phi}(w - w_{t-1})'\Lambda(w - w_{t-1}) \right\}. \tag{15}$$

Here, the transaction costs are multiplied by a "transaction-cost amortisation parameter" ϕ to account for the static nature of the problem in an ad-hoc manner. A naive choice of this parameter is $\phi = 1$, which would mean that the objective (15) compares the returns earned over the next time period (a "flow" variable) with the transaction costs (a "stock" variable) paid today. This comparison is problematic if the portfolio is expected to be held for many time periods, so having the fudge factor ϕ is a simple way to address this problem. In particular, if the portfolio is expected to be held for ϕ time periods, say $\phi = 6$, then we can amortize the trading cost over these time periods, thus dividing the current transaction cost by 6.

The solution to the static objective (15) is:

$$\begin{aligned} w &= (\Sigma_\gamma + \Lambda_\phi)^{-1}(\mu + \Lambda_\phi w_{t-1}) = (\Sigma_\gamma + \Lambda_\phi)^{-1}(\Sigma_\gamma Markowitz_t + \Lambda_\phi w_{t-1}) \\ &= m^{static}\, w_{t-1} + (I - m^{static})\, Markowitz_t \end{aligned} \tag{16}$$

where $\Lambda_\phi = \frac{1}{\phi}\Lambda$ and $m^{static} = (\Sigma_\gamma + \Lambda_\phi)^{-1}\Lambda_\phi$. So we see that the solution actually has the same form as (13), but with two notable differences. First, and most importantly, the static approach treats the current Markowitz portfolio is the aim portfolio, thus failing the distinguish which signals are persistent and which have fast alpha decay. Second, the static approach uses another discount factor, m^{static}.

2.5 Economic Feature Importance

It is important to determine which characteristics are economically important. To address this issue, we first consider the value function V

$$V(x) = \lim_{H\to\infty} \max_{F_t = F(\vec{x}_t)} E\left[r'_{t+1}F_t - \frac{1}{2}F'_t\Sigma_\gamma F_t - \frac{1}{2}\Delta F'_t \Lambda \Delta F_t\right] \tag{17}$$

where the portfolio is allowed to depend on the full history of past signals ($H \to \infty$) and, for simplicity, we continue to focus on the case of constant Σ_γ and Λ. We see that V is the steady-state investor utility, and we write V as a function of the signals, x, that the investor has access to.

We next define the importance of any feature n as

$$\iota_n = V(x) - V(x^{-n}) \tag{18}$$

where x^{-n} is the set of signals x, except that we drop feature n at all times. In other words, the importance in feature n is the drop in utility when the investor no longer has access to this information. The following results provides an intuitive characterization of the drivers of feature importance.

Theorem 2 *In the limit when Λ is small, the investor's steady-state optimal utility is $V(x) = v(x) - c(x) + O(\|\Lambda\|^2)$, where*

$$v(x) = \frac{1}{2} E[Markowitz_t' \, \Sigma_\gamma \, Markowitz_t] \,, \tag{19}$$

is value function without transaction costs, and c measures the cost of time-variation in the Markowitz portfolio:

$$c(x) = \frac{1}{2} E[(Markowitz_{t+1} - Markowitz_t)' \, \Lambda \, (Markowitz_{t+1} - Markowitz_t)] \,. \tag{20}$$

The importance, ι_n, of feature n is

$$\iota_n = \underbrace{v(x) - v(x^{-n})}_{\text{efficiency loss}} - \underbrace{(c(x) - c(x^{-n}))}_{\text{cost reduction}} + O(\|\Lambda\|^2) \,. \tag{21}$$

We see that a feature is more important if it is an important contributor to the Markowitz portfolio (the first term in (21)) and if it is a persistent signal such that it reduces the turnover and, hence, the transaction costs (the second term in (21)).

2.6 Handling Stochastic Volatility and Liquidity Risk

We next look for the optimal portfolio $F(\vec{x}_t)$ that maximizes the objective:

$$\max_{F_t = F(\vec{x}_t)} E\left[r_{t+1}' F_t - \frac{1}{2} F_t' \Sigma_\gamma(x_t) F_t - \Delta F_t' \Lambda(x_t) \Delta F_t \right] \,, \tag{22}$$

where we allow randomness in the risk aversion and variance-covariance matrix, captured by $\Sigma_\gamma(x_t)$, and in transaction costs, $\Lambda(x_t)$. For the sake of brevity, we use the notation $\Lambda(x_t) = \Lambda_t$, $\Sigma_\gamma(x_t) = (\Sigma_\gamma)_t$, and so on.

We show that the solution in the general case with stochastic volatility and liquidity risk has a similar structure as in the simpler case discussed in Section 2.3. However, the discount factor m from (9) now becomes stochastic. In particular, the discount factor $M(x_t)$ must solve the following integral equation

$$M(x) \;=\; (\Lambda(x)^{-1}\Sigma_\gamma(x) + I + \bar{\Lambda}(x))^{-1} \left(E[\Lambda(x)^{-1}\Lambda(x_{t+1})M(x_{t+1})M(x)|x_t = x] + I\right) \quad (23)$$

where $\bar{\Lambda}(x) \;=\; E[\Lambda_t^{-1}\Lambda_{t+1}|x_t = x]$. The following lemma provides natural conditions that ensure such a solution and its uniqueness.

Lemma 2 *The integral equation* (23) *has a unique solution* M *with* $\|\Lambda(x)^{1/2}M(x)\Lambda(x)^{-1/2}\| \leq 1$ *when either of the following conditions hold:*

A. $\sup_x\{\|(\Lambda(x)^{-1/2}\Sigma_\gamma(x)\Lambda(x)^{-1/2} + I + \Lambda(x)^{1/2}\bar{\Lambda}(x)\Lambda(x)^{-1/2})^{-1}\|\,\lambda(x)\} \;<\; \frac{1}{2}$, *where* $\lambda(x) \;=\; E[\,\|\Lambda_t^{-1/2}\Lambda_{t+1}\Lambda_t^{-1/2}\|\,|\,x_t = x]$

B. all matrices in the set $\mathcal{M} \;=\; \{\Lambda(x)^{-1}\Sigma_\gamma(x),\ \Lambda(x)^{-1}\Lambda(y)\}_{x,y\in\mathcal{X}}$ *commute with each other.*

Here, the condition A in Lemma 2 is clearly satisfied in the case of constant risk and liquidity as in Section 2.3, but it is also satisfied in many other natural settings. Condition B is satisfied, for example, when $\Lambda(x) \;=\; \lambda(x)\Lambda$ and $\Sigma_\gamma(x) \;=\; \sigma(x)\,\Sigma_\gamma$ for general positive definite matrices Λ and Σ_γ and general scalar function $\lambda(x), \sigma(x) \in \mathbb{R}_{>0}$.

To state the main result of this section, we also need to define the cumulative discount factor $M_{t-s,t}$, corresponding to m^s in Section 2.3:

$$M_{t-s,t} \;=\; \prod_{\tau=1}^{s} M_{t-\tau+1}, \quad (24)$$

where $M_{t,t} = 1$ and $M_{t-1,t} \;=\; M_t$. We also need the following slightly different cumulative

discount rate

$$\tilde{M}_{t-s,t} \;=\; (I - M_{t-s})^{-2} \left(\prod_{\tau=1}^{s} M_{t-s+\tau-1} \Lambda_{t-s+\tau-1}^{-1} \Lambda_{t-s+\tau} \right) N_t, \tag{25}$$

where $\tilde{M}_{t,t} = 1$ and $N_t = I - M_t - M_t E_t[\Lambda_t^{-1} \Lambda_{t+1}(I - M_{t+1})]$. With the definitions in place, we are ready to state the main result of this section, showing how the optimal portfolio depends on current on future $Markowitz$ portfolios as defined in (3) as well as the stochastic discount rates defined in (24)–(25).

Theorem 3 *The portfolio problem (22) always has a solution $F(\vec{x}_t)$. Under the condition of Lemma 2.A or B, then, this solution is unique, and, in the limit as $H \to \infty$, F converges uniformly to*

$$F^*(x_{-\infty,t}) \;=\; \sum_{s=0}^{\infty} M_{t-s,t}(I - M_{t-s}) A_{t-s}\,, \tag{26}$$

with

$$A_t \;=\; (I - M_t) \sum_{\tau=0}^{\infty} E_t \left[\tilde{M}_{t,t+\tau}\, Markowitz_{t+\tau} \right]\,. \tag{27}$$

To understand this powerful result, note that the optimal portfolio can again be written as a convex combination of the previous portfolio and the current aim portfolio:

$$F^*(x_{-\infty,t}) \;= M_t F^*(x_{-\infty,t-1}) + (I - M_t) A_t\,. \tag{28}$$

In other words, the optimal portfolio at time t is the optimal portfolio in the earlier time period $t-1$, plus a trade toward the aim portfolio A.

The optimal portfolio admits a simpler analytical expressions when transaction costs are small as seen in the following result.

Theorem 4 *Suppose that $\Lambda(x) = \varepsilon \Lambda_*(x)$ where $\Lambda_*(x)$ is uniformly positive definite and*

$\varepsilon > 0$ is small. Then $M(x) = \varepsilon \Sigma_\gamma^{-1}(x)\Lambda(x) + O(\varepsilon^2)$ and the optimal portfolio is

$$F^*(x_{-\infty,t}) = (I - M_t)A_t + M_t A_{t-1} + O(\varepsilon^2), \tag{29}$$

where the aim portfolio A_t is given by

$$A_t = Markowitz_t + \Sigma_\gamma^{-1} E_t[\Lambda_{t+1}(Markowitz_{t+1} - Markowitz_t)] + O(\varepsilon^2) \tag{30}$$

3 Data and Empirical Methodology

3.1 Data

We use the dataset from Jensen et al. (2021), a publicly available dataset and replication code of stock returns and characteristics, where the underlying return data is sourced from CRSP and the accounting data is based on Compustat.[3] We restrict our focus to US common stocks (`shrcd`: 10, 11, and 12) traded on AMEX, NASDAQ or NYSE (`exchcd`: 1, 2, or 3) from 1963 to 2020.

To predict returns, we use the 153 characteristics (or features) studied in Jensen et al. (2021) and add 1-year trailing volatility for a total of 154 features.[4] We standardize the features each month, by mapping the cross-sectional rank into the [0,1] interval. We set missing values to 0.5, but only include stocks with non-missing market equity at the beginning of the month.

3.2 Empirical Methodology

Trading Cost Matrix

In the theoretical section, we assumed that trading cost are convex and equal to $\Delta w_t' \Lambda \Delta w_t$. In our empirical analysis, we further assume that the trading cost matrix is diagonal and

[3] The data, code, and documentation are available at https://github.com/bkelly-lab/ReplicationCrisis/tree/master/GlobalFactors.

[4] We add 1-year trailing volatility, because of its close connection to the covariance matrix. For example, if volatility is unrelated to expected returns, the optimal portfolio should have a higher allocation to low volatility assets all else equal.

calibrate it based on the estimates in Frazzini et al. (2018). Specifically, we assume that the average market impact is 0.1%, when trading 1% of the daily dollar volume. This assumption means that the ith diagonal entry in Λ_t is given by

$$\lambda_{i,t} = \frac{0.1}{V_{i,t}}, \tag{31}$$

where $V_{i,t}$ is the expected daily dollar volume of stock i at time t. For example, trading \$5 million over a day, in a stock with a daily volume of \$500 million has a total impact of $\frac{0.1}{500}5^2 = 0.005$, which corresponds to 0.01% of the traded amount. We follow Frazzini et al. (2018) and assume that the expected daily volume is equal to the average daily volume over the last 6 months. In our empirical implementation, we assume that the realized trading cost are the same as the expected trading cost. Furthermore, we assume that the investor can freely inject/withdraw money for the portfolio, to keep the invested wealth constant. This assumption means that we compute the traded dollar amount as $\Delta w_t = w_t - w_{t-1}$. With constant wealth, we have to consider that dollar volume has increased substantially over time. Therefore, we multiply dollar volume estimates at time t by the ratio of the average dollar volume among non-microcap stocks in December 2020 and the same average at time t. For example, the average dollar volume of non-microcap stock in January 1963 was \$0.14 millon compared to \$134 million in December 2020. The adjusted dollar volume of stocks in January 1963 is therefore roughly 1.000 times the actual dollar volume. The adjustment implies that the assumption of constant wealth is better interpreted as a constant ratio of wealth to market liquidity.

Covariance Matrix

This section describes how we create the covariance matrix estimator, Σ_t. We start by estimating the empirical covariance matrix, using 3 years of weekly return data. We only include stocks with non-missing data for all weeks in the estimation period.

Unfortunately, the empirical covariance matrix, is a poor estimator of the true covariance matrix when the number of assets is large, relative to the number of time periods (Ledoit and Wolf, 2003). To improve the estimator, we start by shrinking the empirical correlation

matrix, $\hat{\Omega}_t$, 50% towards a industry-block-correlation matrix, $\hat{\Omega}_t^B$,

$$\hat{\Omega}_t^S = \frac{1}{2}\hat{\Omega}_t + \frac{1}{2}\hat{\Omega}_t^B. \tag{32}$$

We create $\hat{\Omega}_t^B$ by computing the average correlation between pairs of stocks from each of the 49 Fama and French (1997) industries[5]. These averages then serve as the block estimate for any pair of stocks. For example, the block-correlation estimate for Facebook and Coca-Cola is the average correlation between pairs of stocks where one stock is in the "Computer Software" industry, and the other is in the "Candy and Soda" industry.

With the shrunk correlation matrix, we recompute our updated covariance matrix as

$$\hat{\Sigma}_t^S = D_t \hat{\Omega}_t^S D_t, \tag{33}$$

where D_t is a diagonal matrix with the trailing 252 day volatility multiplied by $\sqrt{21}$ in the diagonal.

Finally, we regularize the covariance matrix by adding a constant to the diagonal. Economically, adding a constant to the diagonal is the same as limiting the portfolio leverage. To see this point, suppose an investor optimizes mean-variance utility, without taking trading cost into account, subject to a constraint on the sum of squared portfolio weights, $w'w \leq l$. The solution to this investors maximization problem is

$$w = (\Sigma + \kappa I)^{-1} \mu \tag{34}$$

where κ is a Lagrange multiplier, I is the identity matrix, and μ is the vector of expected returns. In other words, limiting the squared portfolio weights is the same as adding a constant to the diagonal of the covariance matrix. Interestingly, Kozak et al. (2020) show that adding a constant to the diagonal of the covariance matrix, is also appropriate when the investor is uncertain about expected returns. Finally, the approach is also closely related to ridge regression and ensures that the matrix can be inverted even if Σ is not of full rank. We let the constant vary over time by setting it to $\kappa = c\bar{\sigma}_t^2$, where c is a tuning parameter

[5]We extract the mapping from SIC codes to Fama-French industries from Kenneth French's homepage.

and $\bar{\sigma}_t^2$ is the average diagonal of $\hat{\Sigma}_t^S$. Empirically, $c = 5$ appears to be optimal, in the sense of minimizing the difference between expected and realized portfolio variance.

Expected Returns

To estimate expected returns, we use the XGBoost model from Chen and Guestrin (2016). The model is based on an ensemble of decision trees and the gradient boosting methodology from Friedman (2001). Specifically, with a square loss objective as (1), each new decision tree, tries to predict the residuals of the previous trees in the ensemble. As such, the model directly tackles the bias-variance tradeoff by adaptively minimizing bias, while model variance is controlled by the ensemble feature. The XGBoost algorithm improves on standard gradient boosted machines, by adding more regularization parameters and by providing a more efficient implementation (Chen and Guestrin, 2016).

To ensure that we are using out-of-sample predictions, we split the data set into a training period from 1963 to 1993 and a test period from 1995 to 2020. We skip one year, to avoid information leakage between the training and test period[6]. We select hyper-parameters, by splitting the training data into three temporal folds with an equal number of months[7]. For each set of hyper-parameter, we fit the model to two of the folds and let it predict the remaining third. We repeat this procedure three times, such that each fold acts as a test set once. Finally, we select the hyper-parameters with the lowest mean squared error and re-fit the model with the full training data set.

In total, we fit seven models. For the purpose of estimating the aim portfolio, we estimate six models to predict $E[R_{t+\tau}|x_t]$ with $\tau \in \{1, \ldots, 6\}$. The model with $\tau = 1$ is special, and we sometimes refer to it as $\mu(x_t)$. The last model is designed to predict the average return over the next twelve months. We shrink model implied expected returns 50% towards zero. The shrinkage reflects our prior that expected returns are likely to be less dispersed out-of-sample compared to in-sample. For example, learning and competition among arbitrageurs imply that return predictability should decline over time.

[6]The choice of a one-year gap between the training and test period, is necessary because one of the models we fit, predicts returns over the next 12 months.

[7]Whenever a model uses information from overlapping periods, we make sure to leave an appropriate gap between the folds.

4 Empirical Results

This section implements our method as well as several static benchmarks for comparison. We consider an investor with a constant wealth of \$10 billion and a relative risk aversion of 10, which means that the investor has an absolute risk aversion of $\gamma = \frac{10}{10^{10}} = 10^{-9}$. As discussed in section 3.2, we adjust dollar volumes to be comparable to dollar volumes in December 2020, so the assumption of constant wealth is better interpreted as a constant ratio of wealth to market liquidity.

4.1 Predicting Short- vs. Long-Run Returns

We start by analyzing the difference between predicting short-term and long-term returns. Analyzing this difference between short- and long-run predictability is an important preamble before we dive deeply into the models. Indeed, if there is no difference between short-term and long-term expected return, then the aim portfolio in (14) becomes proportional to the standard Markowitz portfolio and the benefits of dynamic portfolio optimizing disappear.

Figure 1 shows the performance of a ML method trained to predict 1-month excess returns relative to that of a ML method trained to predict the 12-month average excess return across stocks. In both cases, the ML algorithm used is a gradient boosted decision tree implemented using the open-source algorithm called XGBoost. These prediction models are evaluated based on their correlations to the future realized excess returns in month $t+h$ with $h = 1, \ldots, 24$.

The left panel in Figure 1 shows the performance of these prediction models over the in-sample training period, 1963-1993, and the right panel shows the performance over the out-of-sample test period, 1995-2020. In the training period, the 1-month model performs the best for short-term returns while the 12-month model performs the best for long-term returns. Interestingly, this difference carries over to the test period, except that the long-run model tends to do as well or better than the short-term model.

The reason that short- and long-run models are different is that different return predictors have different persistence as seen in Figure 2. A number of predictors are highly persistent as seen in the left panel. For example, a stock's beta, profitability, liquidity, size, and age

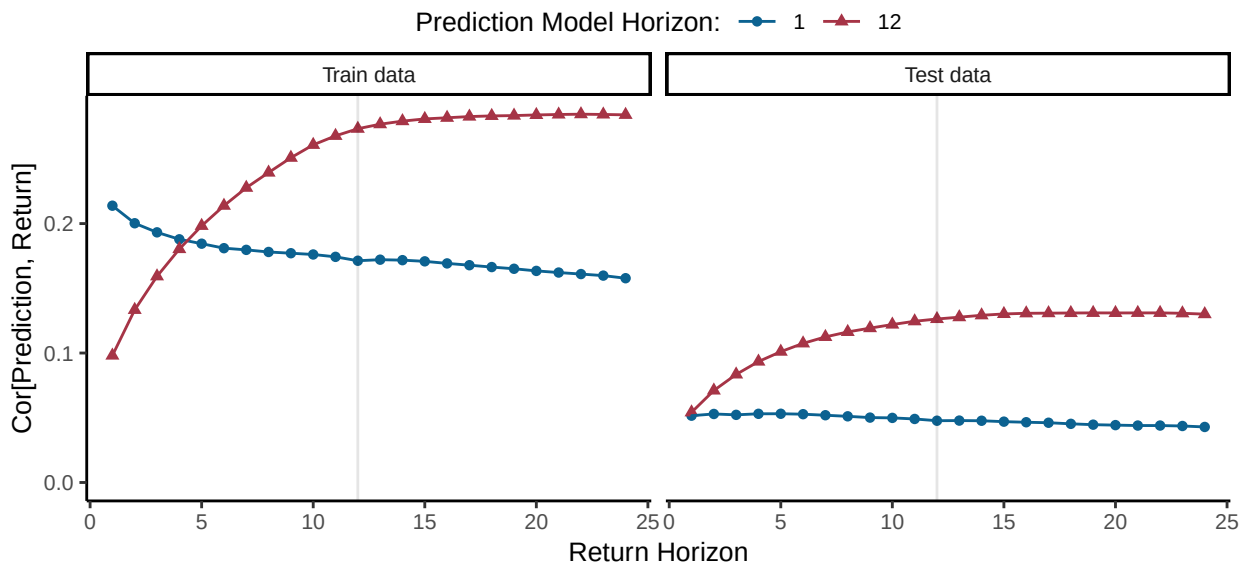


Figure 1: Predictive ability short vs. long-run model

Note: The figure shows the correlation between the ML model prediction in month t and the average realized excess return in month $t+h$ with $h = 1, \ldots, 24$. We compute these correlations for two ML models: The first (blue-circle) is trained to predict next month's excess return; The second (red-triangle) is trained to predict the average excess return over the next 12 months. The training period is from 1963 to 1993 and the test period is from 1995 to 2020, where we skip one year so that the test period is fully out of sample. The models are trained and the results are shown for the sample of all US stocks.

are among the highly persistent characteristics.

As seen in the right panel of Figure 2, other characteristics are quickly mean-reverting, such as momentum factors, and especially short-term reversal factors, and a minority of factors are even negatively autocorrelated. Naturally, such transient factors matter less for long-term returns than short-term returns, while the persistent factors may matter more similarly across horizons.

In summary, we find a difference in short-term and long-term return prediction models, and this difference sets the stage for our analysis of dynamic portfolio optimization in the next section.

4.2 Performance of Portfolio Strategies

We next consider the performance of different portfolio strategies. In particular, we analyze the following alternative specifications:

- **Markowitz-ML.** This specification uses the 1-month return predictions and the cur-

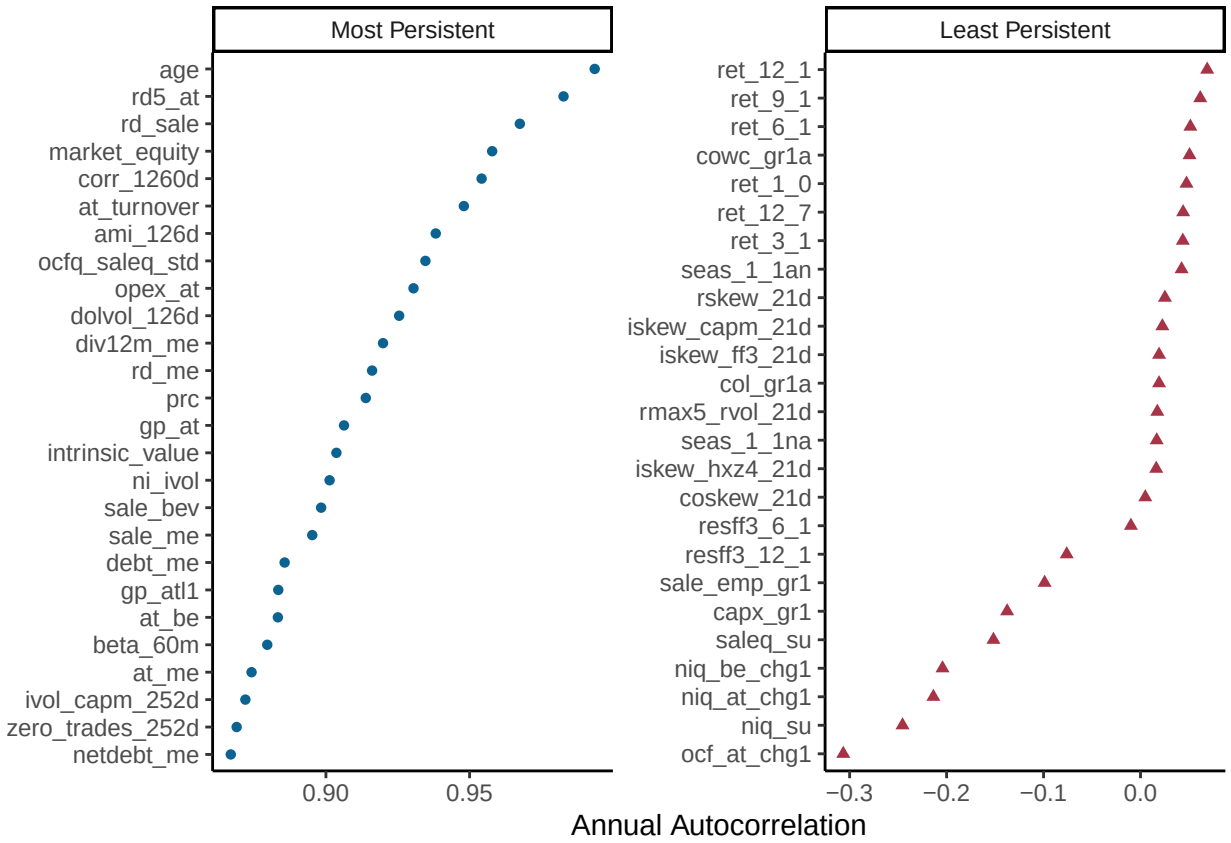


Figure 2: Feature persistence

Note: We report the annual autocorrelation of the most (left) and least (right) persistent feature. We use all observations with non-missing current and 12 month lagged feature values.

rent variance-covariance estimate to form a proxy for the Markowitz portfolio as described in Section 2.2.

- **Economic-ML.** This specification uses our economically motivated dynamic theory described in Section 2.3, method 1. Specifically, we use ML across a range of forecasting horizons, thus producing proxies for $E_t[r_{t+1+\tau}]$ for across τ. Using these forecasts and the variance-covariance matrix, the aim portfolio is computed via (14) and the portfolio at each time is computed recursively using (13).

- **Static-ML(ϕ).** This specification performs a static optimization in light of transaction costs as described in Section 2.4. The static optimization depends on the "fudge factor" ϕ that we denote the transaction cost amortization parameter. We perform this static

optimization for values $\phi \in \{1, 6, 12\}$.

Table 1 reports the performance of each of these benchmarks. In particular, the table reports the different portfolio construction specifications in the first column and a variety of performance statistics in the following columns. The performance statistics consist of the returns before transactions costs, the transaction costs, the net returns, the volatility of the portfolio returns, the value of objective function, and the leverage. These statistics are reported over the out-of-sample test period, and both in terms of expected ex-ante values and in terms of realized ex-post values. Specifically, the expected return is based on the estimated 1-month returns at each time period and the expected volatility is based on the ex ante variance-covariance matrix. In contrast, the realized return is the return that each portfolio actually delivers and the realized volatility is the full-sample realized portfolio volatility.

Focusing first on the expected ex ante values, we see that the Economic-ML approach performs the best in the sense that it delivers the largest utility as evaluated from the objective function that it is designed to maximize. This is what we hoped, but this result is far from mechanical since this analysis is performed out of sample. Figure 3 shows this ex ante outformance of Econonic-ML graphically. Table 1 shows that Economic-ML also delivers the highest ex ante Sharpe ratio net of transaction costs.

We next turn to the realized performance shown in the lower panel of Table 1 as well as in Figure 4. We see that Economic-ML also performs well in the sense of realized returns. However, the realized outperformance in terms of utility is smaller than expected. The expected outperformance 12% relative to the second-best model, Static-ML(1), whereas the realized objective for Economic-ML and Static-ML(1) is similar. The difference between expected and realized performance ultimately depends on the input parameters. The realized returns are slightly lower than the expected returns, which can be attributed to a biased expected return estimate. The volatility is slightly higher, which can be attributed to a biased covariance estimate. The realized Sharpe ratio after trading cost is 1 for Economic-ML and Static-ML(1), which is quite substantial considering that we only use out-of-sample inputs and include trading cost for a large investor with $10 billion invested. The realized Sharpe ratio after cost compares favorably to the realized Sharpe ratio before cost of 1.3 for

Table 1: Performance of Different Portfolio Strategies

	R	TC	R-TC	Vol.	Obj.	Lev.
Expected						
Economic-ML	0.13	0.013	0.12	0.10	0.065	2.1
Markowitz-ML	0.26	4.732	-4.47	0.16	-4.600	3.7
Static-ML(1)	0.11	0.003	0.11	0.10	0.058	1.8
Static-ML(6)	0.17	0.039	0.13	0.12	0.052	2.6
Static-ML(12)	0.19	0.094	0.10	0.13	0.009	2.8
Realized						
Economic-ML	0.12	0.013	0.11	0.11	0.049	2.1
Markowitz	0.20	4.732	-4.53	0.15	-4.638	3.7
Static-ML(1)	0.10	0.003	0.10	0.10	0.049	1.8
Static-ML(6)	0.14	0.039	0.10	0.12	0.032	2.6
Static-ML(12)	0.16	0.094	0.06	0.13	-0.017	2.8

Note: This table reports performance statistics of five different portfolio optimization methods. R is the average monthly return multiplied by 12. The expected R is based on the one-month model: $\mu(x_t)'w_t$. TC is the average monthly trading cost divided by \$10 billion multiplied by 12, R-TC is annualized return net of trading cost, Vol. is the annualized volatility. Expected Vol. is based on the estimated covariance matrix:$w_t'\Sigma_t w_t$. Obj. is the average monthly value of the objective from (2) divided by \$10 billion (investor wealth), multiplied by 12. It can be interpreted as the expected risk-adjusted return net of trading cost. Lev. is the average sum of absolute weights divided by \$10 billion. We consider five different optimization methods: "Economic-ML" is our dynamic approach based on Method 1 in Section 2.3, using a weighted-average of the expected tangency portfolios over the next 6 months. "Markowitz-ML" is the optimal portfolio when ignoring transaction costs based on the expected 1-month return. The "static-ML" portfolios use static optimization (16) based on the 1-month expected return and transaction costs, using a transaction cost amortization parameter of, respectively, $\phi \in \{1, 6, 12\}$. We use out-of-sample data from 1995 to 2020.

the Markowitz portfolio.

4.3 Economic Feature Importance

Section 2.5 shows that a feature is important if it is persistent and important for liquid stocks. We take this theory to the data, by comparing features importance of the 1-month model vs. the 12-month model.

Figure 5 shows a global feature importance measure from Lundberg et al. (2020). The measure is based on Shapley values, with an computational feasible implementation for tree-based models. Each bar can be interpreted as the marginal change in the model prediction, caused by the feature.[8] By this measure, features are important if it leads to a large dispersion

[8]The causal interpretation is only confined to the model. A feature causes the model to change its prediction. We do not mean to imply that the features necessarily caused a change in the expected return of a stock outside of the model.

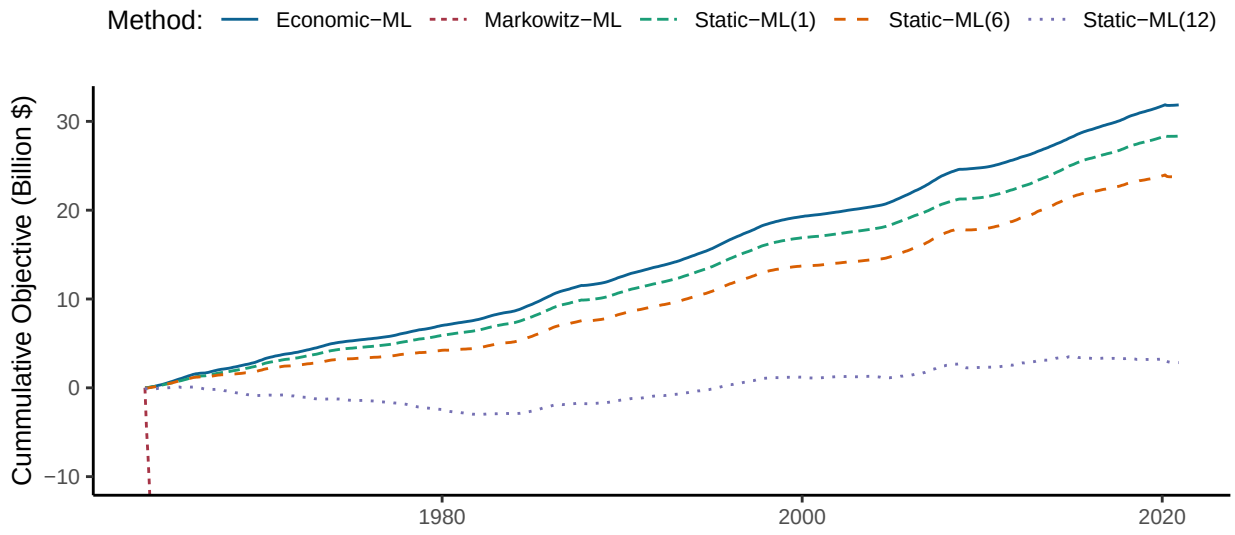


Figure 3: Portfolio Optimization Comparison

Note: This figure shows the cumulative objective (4), which can be interpreted as the cumulative expected risk-adjusted return after transaction costs. We consider five different optimization methods: "Economic-ML" is our dynamic approach based on Method 1 in Section 2.3, using a weighted-average of the expected tangency portfolios over the next 6 months. "Markowitz-ML" is the optimal portfolio when ignoring transaction costs based on the expected 1-month return. The "static-ML" portfolios use static optimization (16) based on the 1-month expected return and transaction costs, using a transaction cost amortization parameter of, respectively, $\phi \in \{1, 6, 12\}$.

in expected returns. Focusing on the 1-month model standard to the literature. The most important feature is short-term reversal. The second most important feature is the price of a stock, which is highly correlated to its size. In summary, the figure agrees with the finding in Gu et al. (2020) the price trend, liquidity, and volatility features are important for predicting returns. Value features also seem to be important for 1-month expected returns. The important variables for the 12-month model are very different from the 1-month model. Most notably, short-term reversal is completely unimportant for longer-run returns. Similarly, momentum features lose their importance. Based on figure 2, this finding is unsurprising since neither short-run reversal nor medium-term momentum is very persistent. In contrast, liquidity, volatility, and value features remain important for the 12-month model. Profitability variables also become more important as the horizon increases, consistent with the fact that a profitable firm today is most likely also profitable in a year.

The other important aspect for investors is whether the feature works for large stocks. By definition, this aspect makes liquidity variables less important, as they require trading illiquid

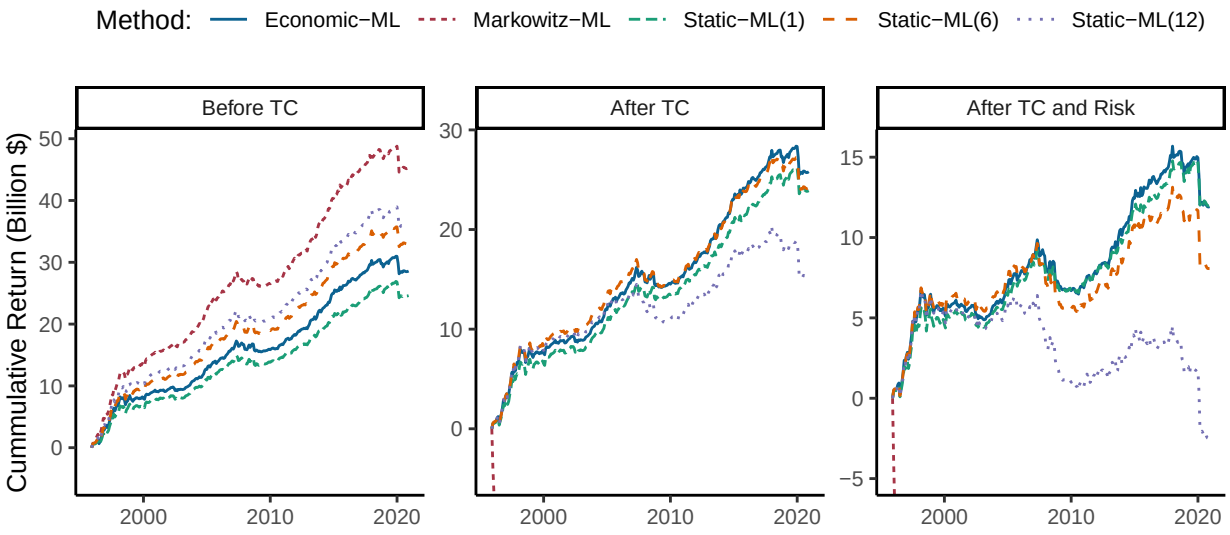


Figure 4: Cumulative portfolio return

Note: We compute the cumulative return before and after trading cost over the test period from 1995-2020. We consider five different optimization methods: "Economic-ML" is our dynamic approach based on Method 1 in Section 2.3, using a weighted-average of the expected tangency portfolios over the next 6 months. "Markowitz-ML" is the optimal portfolio when ignoring transaction costs based on the expected 1-month return. The "static-ML" portfolios use static optimization (16) based on the 1-month expected return and transaction costs, using a transaction cost amortization parameter of, respectively, $\phi \in \{1, 6, 12\}$.

stocks. More generally, feature importance might interact with stock liquidity. Features that are important for illiquid stocks are not necessarily important for liquid stocks. Figure 6 shows the marginal feature importance for the five most important features in each model, where each observation is colored based on the underlying stocks six months trailing dollar volume. Perhaps most notably, the relation between value and expected returns seem to be more pronounced for illiquid stocks. High volatility predicts negative returns equally for illiquid and liquid stocks, but high volatility is much more likely among illiquid stocks.

5 Conclusion

We develop a method for optimal portfolio selection in light of transaction costs when returns are predictable via an *arbitrary* function of security characteristics. The intuitive solution to the portfolio problem shows how the optimal strategy can be computed via machine learning in a tractable way. Empirically, we find significant out-of-sample gains from this method.

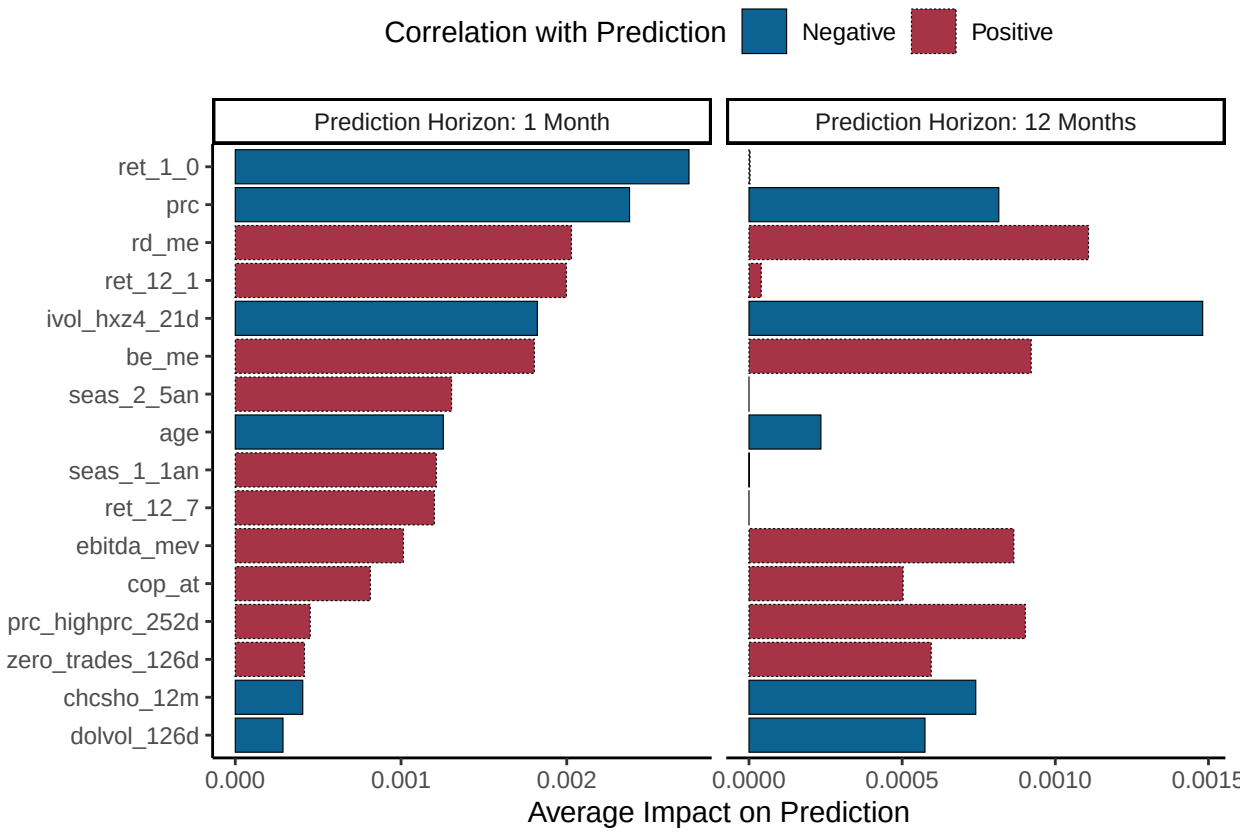


Figure 5: Global feature importance short vs. long-run model

Note: We report the global feature importance for two ML models: In the left panel, the model is trained to predict next month's excess return; In the right panel, the model is trained to predict the average excess return over the next 12 months. The feature importance measure is based on the SHAP method from Lundberg et al. (2020). Specifically, we report the average absolute SHAP values for each feature, across all observations. The bars can be interpreted as the average marginal contribution of a feature, to the monthly expected return of a stock. A blue-solid (red-dotted) bar, indicates that the Spearman correlation between the feature and SHAP value is negative (positive). The models use 154 features, but we only report features which are among the 10 most important in either model.

Further, the method implies a novel view of which securities are important. Indeed, while standard methods that ignore transaction costs are focused on transient features that work well on paper for small stocks, our method naturally selects persistent features of economic importance.

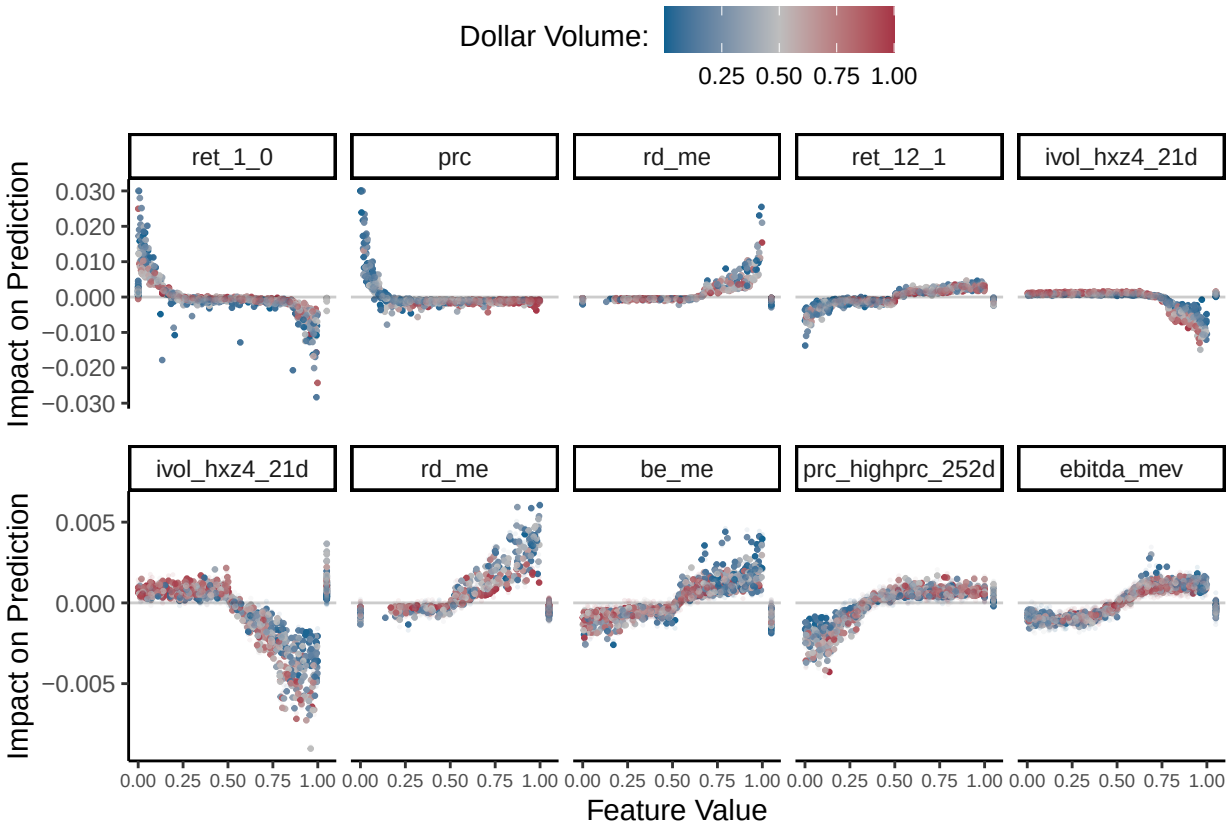


Figure 6: Marginal feature importance

Note: We report marginal feature importance for two ML models: In the upper panel, the model is trained to predict next month's excess return; In the lower panel, the model is trained to predict the average excess return over the next 12 months. The feature importance measure is based on the SHAP method from Lundberg et al. (2020). We sample 5000 observations, and plot individual SHAP values with the corresponding feature value. The features are mapped to the [0, 1] interval. We also report the SHAP value of missing features, by assigning these a value of 1.05. As such, the values to the right of each sub-plot is the SHAP value of when the feature is missing. To capture interaction effects with liquidity, the points are colored according to the 6 months trailing dollar volume of the stock.

References

Avramov, D., S. Cheng, and L. Metzker (2021). Machine learning versus economic restrictions: Evidence from stock return predictability. *Management Science, forthcoming*.

Chen, A. Y. and M. Velikov (2021). Zeroing in on the expected returns of anomalies. *Working paper, Federal Reserve Board*.

Chen, A. Y. and T. Zimmermann (2020). Open source cross-sectional asset pricing. *Working paper, Board of Governors of the Federal Reserve System*.

Chen, L., M. Pelger, and J. Zhu (2020). Deep learning in asset pricing. *Available at SSRN 3350138*.

Chen, T. and C. Guestrin (2016). Xgboost: A scalable tree boosting system. *CoRR abs/1603.02754*.

Collin-Dufresne, P., K. Daniel, and M. Sağlam (2020). Liquidity regimes and optimal dynamic asset allocation. *Journal of Financial Economics 136*(2), 379–406.

Constantinides, G. M. (1986). Capital market equilibrium with transaction costs. *Journal of Political Economy 94*, 842–862.

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

Detzel, A. L., R. Novy-Marx, and M. Velikov (2021). Model selection with transaction costs. *Available at SSRN 3805379*.

Fama, E. F. and K. R. French (1997). Industry costs of equity. *Journal of financial economics 43*(2), 153–193.

Frazzini, A., R. Israel, T. J. Moskowitz, and T. J. Moskowitz (2018). Trading costs. *Available at SSRN 3229719*.

Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics 29*(5), 1189 – 1232.

Gârleanu, N. and L. H. Pedersen (2013). Dynamic trading with predictable returns and transaction costs. *The Journal of Finance 68*(6), 2309–2340.

Gârleanu, N. and L. H. Pedersen (2016). Dynamic portfolio choice with frictions. *Journal of Economic Theory 165*, 487516.

Gu, S., B. Kelly, and D. Xiu (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies 33*(5), 2223–2273.

Hou, K., C. Xue, and L. Zhang (2020). Replicating anomalies. *The Review of Financial Studies 33*(5), 2019–2133.

Jensen, T. I., B. T. Kelly, and L. H. Pedersen (2021). Is there a replication crisis in finance? *Journal of Finance, forthcoming*.

Kozak, S., S. Nagel, and S. Santosh (2020). Shrinking the cross-section. *Journal of Financial Economics 135*(2), 271–292.

Ledoit, O. and M. Wolf (2003). Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of empirical finance 10*(5), 603–621.

Li, S. A., V. DeMiguel, and A. Martin-Utrera (2020). Which factors with price-impact costs? *Victor and Martin-Utrera, Alberto, Which Factors with Price-Impact Costs*.

Lundberg, S. M., G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence 2*(1), 2522–5839.

Markowitz, H. M. (1952). Portfolio selection. *Journal of Finance 7*, 77–91.

A Proofs

Proof of Lemma 1. Multiplying by $\Lambda^{\frac{1}{2}}$ from the left and $\Lambda^{-\frac{1}{2}}$ from the right, and defining $\hat{\Sigma_\gamma} = \Lambda^{-\frac{1}{2}}\Sigma_\gamma\Lambda^{-\frac{1}{2}}$ and $\hat{M} = \Lambda^{\frac{1}{2}}M\Lambda^{-\frac{1}{2}}$, we get

$$\begin{aligned} 0 &= \Lambda^{\frac{1}{2}}M^2\Lambda^{-\frac{1}{2}} - \Lambda^{\frac{1}{2}}(2I + \Lambda^{-1}\Sigma_\gamma)M\Lambda^{-\frac{1}{2}} + \Lambda^{\frac{1}{2}}I\Lambda^{-\frac{1}{2}} \\ &= \hat{M}^2 - (2I + \hat{\Sigma_\gamma})\hat{M} + I = 0\,, \end{aligned} \tag{35}$$

where $\hat{\Sigma_\gamma}$ is symmetric and positive semi-definite. Thus, the unique solution with eigenvalues below one is given by

$$\hat{M} = \frac{1}{2}((2I + \hat{\Sigma_\gamma}) - (4\hat{\Sigma_\gamma} + \hat{\Sigma_\gamma}^2)^{\frac{1}{2}}) \tag{36}$$

and hence

$$M = \Lambda^{-\frac{1}{2}}\hat{M}\Lambda^{\frac{1}{2}}\,. \tag{37}$$

Furthermore, if λ_i are the eigenvalues of $\hat{\Sigma_\gamma}$ then the eigenvalues of $\hat{M}$ are

$$\xi_i = \frac{1}{2}((2 + \lambda_i) - (4\lambda_i + \lambda_i^2)^{\frac{1}{2}})$$

and, by direct calculation, $\xi_i > 0$ and $\xi_i < 1$. Furthermore, M and $\hat{M}$ are similar and therefore have identical eigenvalues. This ensures the transversality condition. □

Proof of Theorem 1. Standard convexity arguments imply that it suffices to derive and verify the first order conditions for our candidate solution.

Let $\Phi = \Phi(x_{t-k,t})$ be a candidate optimal policy and $\sigma^2 = E[tr(\Sigma_\gamma(x_t))]$. Without loss of generality, we may assume that x_t has a transition density $p(x_t, x_{t+1})$ with respect to the Lebesgue measure. We will also use $p_*(x)$ to denote the stationary distribution. We will also use $x_{t-} = x_{t-k,t-1}$ to denote the k-history of the process. Then, by direct calculation, we have

$$\begin{aligned} \mathcal{O}(\Phi) - E^*[tr(\Sigma_\gamma(x_t))] &= E^*\Bigg[- \mu(x_t)'\Phi(x_t, x_{t-}) + \frac{1}{2}\|\Phi(x_t, x_{t-})'\Sigma_\gamma\Phi(x_t, x_{t-})\| \\ &+ \frac{1}{2}\int (\Phi(x_{t+1}, x_{t+1-} - \Phi(x_t, x_{t-})'\Lambda(\Phi(x_{t+1}, x_{t+1-} - \Phi(x_t, x_{t-}))p(x_t, x_{t+1})dx_{t+1}\Bigg] \\ &= E^*\Bigg[- \mu(x_t)'\Phi(x_t, x_{t-}) + \frac{1}{2}\Phi(x_t, x_{t-})'(\Sigma_\gamma + 2\Lambda)\Phi(x_t, x_{t-})\Bigg] \\ &- E^*[\Phi(x_t, x_{t-})'\Lambda\Phi(x_{t+1}, x_{t+1-}]\,, \end{aligned} \tag{38}$$

where we have used that, by the ergodicity of the p_* measure, we have

$$E^*[\Phi(x_t, x_{t-})'(\Sigma_\gamma + 2\Lambda)\Phi(x_t, x_{t-})] = E^*[\Phi(x_{t+1}, x_{t+1-})'(\Sigma_\gamma + 2\Lambda)\Phi(x_{t+1}, x_{t+1-})]$$

In order to compute the first order conditions, we need to calculate the Frechet derivative of (49) with respect to Φ. To this end, we consider a small perturbation $\Phi \to \Phi + \varepsilon Y$ and

calculate the first order term in ε, so that

$$\mathcal{O}(\Phi_t + \varepsilon Y_t) \;=\; \mathcal{O}(\Phi_t) \;+\; \varepsilon\, E^*[\mathcal{D}(\Phi_t)'\, Y_t] \;+\; O(\varepsilon^2) \tag{39}$$

and $\mathcal{D}(\Phi)$ is the Frechet derivative. To this end, we compute

$$\begin{aligned}
&E^*[\mu(x_t)'(\Phi(x_t, x_{t-}) + \varepsilon Y_t)] \;=\; E^*[\mu(x_t)'\Phi(x_t, x_{t-})] + \varepsilon E^*[\mu(x_t)'Y_t] \\
&E^*[(\Phi(x_t, x_{t-}) + \varepsilon Y_t)'(\Sigma_\gamma + 2\Lambda)(\Phi(x_t, x_{t-}) + \varepsilon Y_t)] \;=\; E^*[\Phi(x_t, x_{t-})'(\Sigma_\gamma + 2\Lambda)\Phi(x_t, x_{t-})] \\
&+\; 2\varepsilon\, E^*[\Phi(x_t, x_{t-})'(\Sigma_\gamma + 2\Lambda)Y_t] \;+\; O(\varepsilon^2) \\
&E^*[(\Phi(x_t, x_{t-}) + Y_t)'\Lambda(\Phi(x_{t+1}, x_{t+1-}) + Y_{t+1})] \\
&=\; E^*[\Phi(x_t, x_{t-})'\Lambda\Phi(x_{t+1}, x_{t+1-})] \;+\; \varepsilon\Bigg(E^*[Y_t'\Lambda E_t[\Phi(x_{t+1}, x_{t+1-})]] \;+\; E[\Phi(x_t, x_{t-})\Lambda Y_{t+1}]\Bigg)
\end{aligned} \tag{40}$$

Furthermore, by stationarity,

$$E[\Phi(x_t, x_{t-})Y_{t+1}] \;=\; E[\Phi(x_{t-1}, x_{t-1-})Y_t]$$

We now observe that Y_t is a function of $x_t, x_{t-1}, \cdots, x_{t-k}$ and, hence, only information about x_{t-k-1} is "lost." Thus, by stationarity

$$E[\Phi(x_{t-1}, x_{t-1-})Y_t] \;=\; E[Y_t\, E[\Phi(x_{t-1}, x_{t-1-})|(x_t, \cdots, x_{t-k})]]$$

and we conclude that the Frechet derivative is given by

$$\begin{aligned}
\mathcal{D}(\Phi) \;&=\; -\mu(x_t) \;+\; (\Sigma_\gamma + 2\Lambda)\Phi(x_t, x_{t-}) \\
&-\; \Lambda E_t[\Phi(x_{t+1}, x_{t+1-})] \;-\; \Lambda E[\Phi(x_{t-1}, x_{t-1-})|(x_t, \cdots, x_{t-k})]\,,
\end{aligned} \tag{41}$$

implying that a bounded Φ is optimal if it satisfies the *integral equation*

$$\begin{aligned}
\Phi(x_t, x_{t-}) \;&=\; (\Sigma_\gamma + 2\Lambda)^{-1}\Bigg(\mu(x_t) \;+\; \Lambda E_t[\Phi(x_{t+1}, x_{t+1-})] \\
&+\; \Lambda E[\Phi(x_{t-1}, x_{t-1-})|(x_t, \cdots, x_{t-k})]\Bigg)\,.
\end{aligned} \tag{42}$$

Substituting the Ansatz

$$\Phi \;=\; \sum_{\theta=0}^{k} F_\theta(x_{t-\theta}) \tag{43}$$

into this equation, we get the following system for F_θ : This gives a system of equations

$$\begin{aligned}
F_0(x_t) \;&=\; (\Sigma_\gamma + 2\Lambda)^{-1}\Big(\mu(x_t) \;+\; \Lambda(E[F_0(x_{t+1})|x_t] \;+\; F_1(x_t))\Big) \\
F_\theta(x_{t-\theta}) \;&=\; (\Sigma_\gamma + 2\Lambda)^{-1}\Lambda(F_{\theta+1}(x_{t-\theta}) + F_{\theta-1}(x_{t-\theta})),\; \theta \in [1, k-1] \\
F_k(x_{t-k}) \;&=\; (\Sigma_\gamma + 2\Lambda)^{-1}\Lambda(F_{k-1}(x_{t-k}) + E[F_k(x_{t-k-1})|x_{t-k}])
\end{aligned} \tag{44}$$

Solving the second (difference) equation gives

$$F_\theta \;=\; F_k\,\frac{m^\theta-\hat{m}^\theta}{m^k-\hat{m}^k} \;+\; F_0\,\frac{m^k\hat{m}^\theta-\hat{m}^k m^\theta}{m^k-\hat{m}^k} \tag{45}$$

so that

$$\begin{aligned} F_1 \;&=\; F_k\,\frac{m-\hat{m}}{m^k-\hat{m}^k} \;+\; F_0\,\frac{m^k\hat{m}-\hat{m}^k m}{m^k-\hat{m}^k} \;=\; F_kA_k+F_0B_k \\ F_{k-1} \;&=\; F_k\,\frac{m^{k-1}-\hat{m}^{k-1}}{m^k-\hat{m}^k} \;+\; F_0\,\frac{m^k\hat{m}^{k-1}-\hat{m}^k m^{k-1}}{m^k-\hat{m}^k} \end{aligned} \tag{46}$$

where $m,\ \hat{m}$ are solutions to the Riccati equation

$$m^2-(2I+\Lambda^{-1}\Sigma_\gamma)m+I \;=\; 0\ ,\ m>1>\hat{m}>0$$

In the limit as $k\to\infty$, we get

$$F \;=\; \sum_{\theta=0}^{\infty} m^\theta F_0(x_{t-\theta})$$

while

$$F_1 \;=\; F_0 m$$

and hence

$$F_0(x_t) \;=\; (\Sigma_\gamma+2\Lambda)^{-1}(\mu(x_t) \;+\; \Lambda(E[F_0(x_{t+1})|x_t] \;+\; mF_0(x_t)))$$

so that, using the identity

$$(\Sigma_\gamma+2\Lambda)^{-1} \;=\; (\Lambda^{-1}\Sigma_\gamma+2I)^{-1}\Lambda^{-1}$$

and

$$F_0(x_t) \;=\; (I-(\Lambda^{-1}\Sigma_\gamma+2I)^{-1}m)^{-1}(\Lambda^{-1}\Sigma_\gamma+2I)^{-1}\Lambda^{-1}(\mu(x_t) \;+\; \Lambda E[F_0(x_{t+1})|x_t])$$

and the claim follows because

$$(I-(\Lambda^{-1}\Sigma_\gamma+2I)^{-1}m)^{-1}(\Lambda^{-1}\Sigma_\gamma+2I)^{-1} \;=\; m \tag{47}$$

so that

$$\begin{aligned} F_0(x_t) \;&=\; m\Lambda^{-1}(\mu(x_t) \;+\; \Lambda E[F_0(x_{t+1})|x_t]) \;=\; m\Lambda^{-1}\mu(x_t) \;+\; mE[F_0(x_{t+1})|x_t] \\ &=\; \Lambda^{-1}\Sigma_\gamma m(\Sigma_\gamma^{-1}\mu(x_t)) \;+\; mE[F_0(x_{t+1})|x_t] \end{aligned} \tag{48}$$

where we have used the fact that m and $\Lambda^{-1}\Sigma_\gamma$ commute. Furthermore,

$$m^2-2m+I \;=\; \Lambda^{-1}\Sigma_\gamma m$$

so that

$$(I-m)^2 \;=\; \Lambda^{-1}\Sigma_\gamma m$$

and, hence,

$$I-m \;=\; (I-m)^{-1}\Lambda^{-1}\Sigma_\gamma m$$

implying the required identity. □

Proof of Lemma 2. By direct calculation, the map

$$M(x) \to (\Sigma_\gamma(x)+\Lambda(x)+\bar{\Lambda}(x))^{-1}\,(E[\Lambda(x_{t+1})M(x_{t+1})|x_t=x]M(x)+\Lambda(x))$$

is a contraction on the set of M with $\|M\| \leq 1$. □

Proof of Theorem 3. Standard convexity arguments imply that it suffices to derive and verify the first order conditions for our candidate solution.

Let $\Phi = \Phi(x_{t-k,t})$ be a candidate optimal policy and $\sigma^2 = E[tr(\Sigma_\gamma(x_t))]$. Without loss of generality, we may assume that x_t has a transition density $p(x_t, x_{t+1})$ with respect to the Lebesgue measure. We will also use $p_*(x)$ to denote the stationary distribution. We will also use $x_{t-} = x_{t-k,t-1}$ to denote the k-history of the process. Then, by direct calculation, we have

$$\begin{aligned}
\mathcal{O}(\Phi) \;-\; E^*[tr(\Sigma_\gamma(x_t))] \;&=\; E^*\Big[-2\,\mu(x_t)'\Phi(x_t,x_{t-}) \;+\; \Phi(x_t,x_{t-})'\Sigma_\gamma(x_t)\Phi(x_t,x_{t-}) \\
&+\; \int (\Phi(x_{t+1},x_{t+1-}-\Phi(x_t,x_{t-})'\Lambda(x_{t+1})(\Phi(x_{t+1},x_{t+1-}-\Phi(x_t,x_{t-}))p(x_t,x_{t+1})dx_{t+1}\Big] \\
&=\; E^*\Big[-2\,\mu(x_t)'\Phi(x_t,x_{t-}) \;+\; \Phi(x_t,x_{t-})'(\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar{\Lambda}(x_t))\Phi(x_t,x_{t-})\Big] \\
&- 2E^*[\Phi(x_t,x_{t-})'\Lambda(x_{t+1})\Phi(x_{t+1},x_{t+1-}]\,,
\end{aligned} \tag{49}$$

where we have used that, by ergodicity of the p_* measure, we have

$$E^*[\Phi(x_t,x_{t-})'\Lambda(x_{t+1})\Phi(x_t,x_{t-})] \;=\; E^*[\Phi(x_t,x_{t-})'\bar{\Lambda}(x_t)\Phi(x_t,x_{t-})]$$

and

$$E^*[\Phi(x_{t+1},x_{t+1-})'\Lambda(x_{t+1})\Phi(x_{t+1},x_{t+1-})] \;=\; E^*[\Phi(x_t,x_{t-})'\Lambda(x_t)\Phi(x_t,x_{t-})]$$

In order to compute the first order conditions, we need to calculate the Frechet derivative of (49) with respect to Φ. To this end, we consider a small perturbation $\Phi \to \Phi + \varepsilon Y$ and calculate the first order term in ε, so that

$$\mathcal{O}(\Phi_t+\varepsilon Y_t) \;=\; \mathcal{O}(\Phi_t) \;+\; \varepsilon\, E^*[\mathcal{D}(\Phi_t)'\,Y_t] \;+\; O(\varepsilon^2) \tag{50}$$

and $\mathcal{D}(\Phi)$ is the Frechet derivative. To this end, we compute

$$\begin{array}{l} E^*[\mu(x_t)'(\Phi(x_t,x_{t-})+\varepsilon Y_t)] \;=\; E^*[\mu(x_t)'\Phi(x_t,x_{t-})]+\varepsilon E^*[\mu(x_t)'Y_t] \\ E^*[(\Phi(x_t,x_{t-})+\varepsilon Y_t)'(\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar{\Lambda}(x_t))(\Phi(x_t,x_{t-})+\varepsilon Y_t)] \\ =\; E^*[\Phi(x_t,x_{t-})'(\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar{\Lambda}(x_t))\Phi(x_t,x_{t-})] \\ +\; 2\varepsilon\, E^*[\Phi(x_t,x_{t-})'(\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar{\Lambda}(x_t))Y_t] \;+\; O(\varepsilon^2) \\ E^*[(\Phi(x_t,x_{t-})+Y_t)'\Lambda(x_{t+1})(\Phi(x_{t+1},x_{t+1-})+Y_{t+1})] \\ =\; E^*[\Phi(x_t,x_{t-})'\Lambda(x_{t+1})\Phi(x_{t+1},x_{t+1-})] \\ +\; \varepsilon\left(E^*[Y_t'E_t[\Lambda(x_{t+1})\Phi(x_{t+1},x_{t+1-})]] \;+\; E[\Phi(x_t,x_{t-})\Lambda(x_{t+1})Y_{t+1}]\right) \end{array} \tag{51}$$

Furthermore, by stationarity,

$$E[\Phi(x_t,x_{t-})Y_{t+1}] \;=\; E[\Phi(x_{t-1},x_{t-1-})Y_t]$$

We now observe that Y_t is a function of $x_t, x_{t-1}, \cdots, x_{t-k}$ and, hence, only information about x_{t-k-1} is "lost." Thus, by stationarity

$$E[\Phi(x_{t-1},x_{t-1-})\Lambda(x_t)Y_t] \;=\; E[Y_t\,\Lambda(x_t)E[\Phi(x_{t-1},x_{t-1-})|(x_t,\cdots,x_{t-k})]]$$

and we conclude that the Frechet derivative is given by

$$\begin{array}{l} \mathcal{D}(\Phi) \;=\; -2\mu(x_t) \;+\; 2(\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar{\Lambda}(x_t))\Phi(x_t,x_{t-}) \\ -\; 2E_t[\Lambda(x_{t+1})\Phi(x_{t+1},x_{t+1-})] \;-\; 2\Lambda(x_t)E[\Phi(x_{t-1},x_{t-1-})|(x_t,\cdots,x_{t-k})]\,, \end{array} \tag{52}$$

implying that a bounded Φ is optimal if it satisfies the *integral equation*

$$\begin{array}{l} \Phi(x_t,x_{t-}) \;=\; (\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar{\Lambda}(x_t))^{-1}\Big(\mu(x_t) \;+\; E_t[\Lambda(x_{t+1})\Phi(x_{t+1},x_{t+1-})] \\ +\; \Lambda(x_t)E[\Phi(x_{t-1},x_{t-1-})|(x_t,\cdots,x_{t-k})]\Big)\,. \end{array} \tag{53}$$

Substituting the Ansatz

$$\Phi \;=\; \sum_{\theta=0}^{k}\prod_{\tau=1}^{\theta} M(X_{t-\tau+1})Q(x_{t-\theta}) \tag{54}$$

into this equation, we get the following system for F_θ : This gives a system of equations

$$\begin{aligned}
&Q(x_t) \;=\; (\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar\Lambda(x_t))^{-1}\Big(\mu(x_t) \;+\; (E[\Lambda(x_{t+1})Q(x_{t+1})|x_t] \\
&+\; E[\Lambda(x_{t+1})M(x_{t+1})|x_t]\,Q(x_t))\Big) \\
&M_{t-\theta,t}Q(x_{t-\theta}) \\
&=\; (\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar\Lambda(x_t))^{-1}(E_t[\Lambda(x_{t+1})M_{t-\theta,t+1}]Q(x_{t-\theta})+\Lambda(x_t)\,M_{t-\theta,t-1}Q(x_{t-\theta})) \\
&\theta\in[1,k-1] \\
&M_{t-k,t}Q(x_{t-k}) \\
&=\; (\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar\Lambda(x_t))^{-1}(E_t[\Lambda(x_{t+1})M_{t-k,t-1}]Q(x_{t-k})+\Lambda(x_t)\,M_{t-k-1,t-1}E[Q(x_{t-k-1})|x_{t-k}])
\end{aligned} \tag{55}$$

Multiplying the second equation by $(M_{t-\theta,t-1}Q(x_{t-\theta}))^{-1} = (M(X_{t-1})M(X_{t-2})\cdots M(X_{t-\theta+1})Q(x_{t-\theta}))^{-1}$ *from the right,* we get

$$M(x) \;=\; (\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar\Lambda(x_t))^{-1}(E_t[\Lambda(x_{t+1})M(X_{t+1})M(X_t)]+\Lambda(x))\;,\;\theta\in[1,k-1] \tag{56}$$

giving the required integral equation for M.

Then, the final equation is

$$\begin{aligned}
&Q(x_t) \;=\; (\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar\Lambda(x_t))^{-1}\Big(\mu(x_t) \;+\; (E[\Lambda(x_{t+1})Q(x_{t+1})|x_t] \\
&+\; E[\Lambda(x_{t+1})M(x_{t+1})|x_t]\,Q(x_t))\Big)
\end{aligned} \tag{57}$$

But, from the equation for M, we get

$$(I-(\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar\Lambda(x_t))^{-1}E[\Lambda(x_{t+1})M(x_{t+1})|x_t]) \;=\; (\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar\Lambda(x_t))^{-1}\Lambda(x)M(x)^{-1} \tag{58}$$

and can therefore rewrite (57) as

$$\begin{aligned}
&(\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar\Lambda(x_t))^{-1}\Lambda(x_t)M(x_t)^{-1}Q(x_t) \\
&=\; (\Sigma_\gamma(x_t)+\Lambda(x_t)+\bar\Lambda(x_t))^{-1}\Big(\mu(x_t) \;+\; E[\Lambda(x_{t+1})Q(x_{t+1})|x_t])\Big)
\end{aligned} \tag{59}$$

so that

$$Q(x) \;=\; M(x)\Lambda(x)^{-1}\Big(\mu(x_t) \;+\; E[\Lambda(x_{t+1})Q(x_{t+1})|x_t=x])\Big) \tag{60}$$

By direct calculation,

$$M(x_t)\Lambda(x_t)^{-1}\Sigma_\gamma(x_t) \;=\; N(x_t) \tag{61}$$

and therefore

$$Q(x) \;=\; N(x)\Sigma_\gamma(x)^{-1}\mu(x) \;+\; E[M(x)\Lambda(x)^{-1}\Lambda(x_{t+1})\,Q(x_{t+1})|x_t=x])\,. \tag{62}$$

□

Proof of Theorem 4. We have

$$N(x_t) \;=\; I - M(x_t) - M(x_t)E_t[\Lambda_t^{-1}\Lambda_{t+1}] \;+\; O(\varepsilon^2)$$

and

$$\begin{aligned} A(x_t) \;&=\; (I - M(x_t))(I - M(x_t))^{-2}\Bigg(N(x_t)Markowitz_t \\ &+\; M_t\Lambda_t^{-1}E_t[\Lambda_{t+1}Markowitz_{t+1}]\Bigg) \;+\; O(\varepsilon^2) \\ &=\; (I + M_t)\Bigg((I - M_t)Markowitz_t \\ &+\; M_t\Lambda_t^{-1}E_t[\Lambda_{t+1}(Markowitz_{t+1} - Markowitz_t)] \;+\; O(\varepsilon^2) \\ &=\; Markowitz_t \;+\; M_t\Lambda_t^{-1}E_t[\Lambda_{t+1}(Markowitz_{t+1} - Markowitz_t)] \;+\; O(\varepsilon^2) \end{aligned} \tag{63}$$

□

Lemma 3 *Suppose that* $\Lambda(x) = \Lambda_0 + \varepsilon\Lambda_*(x) + O(\varepsilon^2)$ *and* $\Sigma_\gamma(x) = \Sigma_\gamma^0 + \varepsilon\Sigma_\gamma^*(x) + O(\varepsilon^2)$. *Let* M_0 *be the unique solution to*

$$M_0^2 - (\Lambda_0^{-1}\Sigma_\gamma + 2I)M_0 \;+\; I \;=\; 0$$

with eigenvalues below one. Then,

$$M(x) \;=\; E\left[\sum_{\tau=0}^{\infty} M_0^\tau \Gamma(x_\tau) M_0^\tau | x_0 = x\right] \;+\; O(\varepsilon^2)\,, \tag{64}$$

with

$$\Gamma(x) \;=\; M_0 - M_0^3 - M_0\Lambda_0^{-1}M_0 \;+\; M_0\Lambda_0^{-1}\Big((\Sigma_\gamma(x) + \Lambda(x) + \bar{\Lambda}(x))^{-1}(\Lambda(x) + \bar{\Lambda}(x)M_0^2)\Big) \tag{65}$$

Proof. Suppose that $\Lambda(x) = \Lambda_0 + \varepsilon\Lambda_*(x)$ and $\Sigma_\gamma(x) = \Sigma_\gamma^0 + \varepsilon\Sigma_\gamma^*(x)$. Then,

$$M(x) \;=\; M_0 + \;\varepsilon M_*(x) \;+\; O(\varepsilon^2)$$

where M_0 solves

$$\Lambda_0^{-1}(\Sigma_\gamma + 2\Lambda_0)M_0 \;=\; M_0^2 + I$$

and then substituting $M_0 + \;\varepsilon M_*(x)$ into the equation and equating the ε-order terms, we

get

$$
\begin{aligned}
M_0 + \ \varepsilon M_*(x) \ &= \ \Big((\Sigma_\gamma^0 + 2\Lambda_0)^{-1} - \varepsilon(\Sigma_\gamma^0 + 2\Lambda_0)^{-1}(\Sigma_\gamma^*(x) + \Lambda_*(x) + \bar{\Lambda}_*(x))(\Sigma_\gamma^0 + 2\Lambda_0)^{-1}\Big) \\
&\times (E[(\Lambda_0 + \varepsilon\Lambda_*(x_{t+1}))(M_0 + \varepsilon M_*(x_{t+1}))(M_0 + \varepsilon M_*(x))|x_t = x] + \Lambda(x)) \ + \ O(\varepsilon^2) \\
&= \ \Big((\Sigma_\gamma^0 + 2\Lambda_0)^{-1} - \varepsilon(\Sigma_\gamma^0 + 2\Lambda_0)^{-1}(\Sigma_\gamma^*(x) + \Lambda_*(x) + \bar{\Lambda}_*(x))(\Sigma_\gamma^0 + 2\Lambda_0)^{-1}\Big) \\
&\times \big(E[\Lambda_0 M_0^2 + \varepsilon(\Lambda_0(M_0 M_*(x) + M_*(x_{t+1})M_0) + \varepsilon\Lambda(x_{t+1})M_0^2|x_t = x] + \Lambda_0 + \varepsilon\Lambda_*(x)\big) \ + \ O(\varepsilon^2) \\
&= \ (\Sigma_\gamma^0 + 2\Lambda_0)^{-1}\Lambda_0 M_0^2 + \Lambda_0 \\
&+ \ \varepsilon\Bigg((\Sigma_\gamma^0 + 2\Lambda_0)^{-1}\Big(\Lambda_0 M_0 M_*(x) + \Lambda_0 E[M_*(x_{t+1})|x_t = x]M_0 + E[\Lambda(x_{t+1})|x_t = x]M_0^2 + \Lambda_*(x)\Big) \\
&- (\Sigma_\gamma^0 + 2\Lambda_0)^{-1}(\Sigma_\gamma^*(x) + \Lambda_*(x) + \bar{\Lambda}_*(x))(\Sigma_\gamma^0 + 2\Lambda_0)^{-1}(\Lambda_0 M_0^2 + \Lambda_0)\Bigg) \ + \ O(\varepsilon^2)
\end{aligned} \tag{66}
$$

As a result,

$$
\begin{aligned}
M_*(x) \ &= \ (\Sigma_\gamma^0 + 2\Lambda_0)^{-1}\Big(\Lambda_0 M_0 M_*(x) + \Lambda_0 E[M_*(x_{t+1})|x_t = x]M_0 + \bar{\Lambda}_*(x)M_0^2 + \Lambda_*(x)\Big) \\
&- (\Sigma_\gamma^0 + 2\Lambda_0)^{-1}(\Sigma_\gamma^*(x) + \Lambda_*(x) + \bar{\Lambda}_*(x))(\Sigma_\gamma^0 + 2\Lambda_0)^{-1}(\Lambda_0 M_0^2 + \Lambda_0)
\end{aligned} \tag{67}
$$

so that

$$
(I - (\Sigma_\gamma^0 + 2\Lambda_0)^{-1}\Lambda_0 M_0)M_*(x) \ = \ (\Sigma_\gamma^0 + 2\Lambda_0)^{-1}(\Xi(x) \ + \ \Lambda_0 E[M_*(x_{t+1})|x_t = x]M_0) \tag{68}
$$

where we have defined

$$
\begin{aligned}
\Xi(x) \ &= \ (\Sigma_\gamma^0 + 2\Lambda_0)^{-1}(\Lambda_*(x) + \bar{\Lambda}_*(x)M_0^2) \\
&- (\Sigma_\gamma^0 + 2\Lambda_0)^{-1}(\Sigma_\gamma^*(x) + \Lambda_*(x) + \bar{\Lambda}_*(x))(\Sigma_\gamma^0 + 2\Lambda_0)^{-1}(\Lambda_0 M_0^2 + \Lambda_0) \\
&= \ \varepsilon^{-1}((\Sigma_\gamma(x) + \Lambda(x) + \bar{\Lambda}(x))^{-1}(\Lambda(x) + \bar{\Lambda}(x)M_0^2) - (\Sigma_\gamma^0 + 2\Lambda_0)^{-1}(\Lambda_0 M_0^2 + \Lambda_0)) \ + \ O(\varepsilon) \\
&= \ \varepsilon^{-1}((\Sigma_\gamma(x) + \Lambda(x) + \bar{\Lambda}(x))^{-1}(\Lambda(x) + \bar{\Lambda}(x)M_0^2) - M_0)
\end{aligned} \tag{69}
$$

Furthermore,

$$
(I - (\Sigma_\gamma^0 + 2\Lambda_0)^{-1}\Lambda_0 M_0)^{-1}(\Sigma_\gamma^0 + 2\Lambda_0)^{-1} \ = \ M_0\Lambda_0^{-1} \tag{70}
$$

so that we get the equation

$$
M_*(x) \ = \ M_0\Lambda_0^{-1}(\Xi(x) + \Lambda_0 E[M_*(x_{t+1})|x_t = x]M_0) \ = \ M_0\Lambda_0^{-1}\Xi(x) + M_0 E[M_*(x_{t+1})|x_t = x]M_0 \tag{71}
$$

and hence

$$
\begin{aligned}
M_0 + \varepsilon M_*(x) \;&=\; M_0 \;+\; M_0\Lambda_0^{-1}\Big((\Sigma_\gamma(x) + \Lambda(x) + \bar{\Lambda}(x))^{-1}(\Lambda(x) + \bar{\Lambda}(x)M_0^2) - M_0\Big) \\
&+\; M_0 E[M_0 + \varepsilon M_*(x_{t+1}) - M_0 | x_t = x]M_0 \\
&=\; M_0 - M_0^3 - M_0\Lambda_0^{-1}M_0 \;+\; M_0\Lambda_0^{-1}\Big((\Sigma_\gamma(x) + \Lambda(x) + \bar{\Lambda}(x))^{-1}(\Lambda(x) + \bar{\Lambda}(x)M_0^2)\Big) \\
&+\; M_0 E[M_0 + \varepsilon M_*(x_{t+1})|x_t = x]M_0
\end{aligned}
\tag{72}
$$

□

Proof of Proposition 2. Let $\nu_t = Markowitz_t$. With small transaction costs, using the approximation $m \approx \Sigma_\gamma^{-1}\Lambda$ and

$$
F_t \;=\; (I - m)A_t \;+\; mA_{t-1} \;+\; O(\|\Lambda\|^2)
$$

and

$$
\begin{aligned}
A_t \;&=\; \nu_t + \Sigma_\gamma^{-1}\Lambda E_t[\nu_{t+1} - \nu_t] \;+\; O(\|\Lambda\|^2) \\
&=\; \nu_t + \Sigma_\gamma^{-1}\Lambda E_t[\nu_{t+1} - \nu_t] \;+\; O(\|\Lambda\|^2)
\end{aligned}
\tag{73}
$$

so that

$$
\begin{aligned}
F_t \;&=\; (I - m)A_t \;+\; mA_{t-1} \;+\; O(\|\Lambda\|^2) \;=\; (I - m)(\nu_t + \Sigma_\gamma^{-1}\Lambda E_t[\nu_{t+1} - \nu_t]) \\
&+\; m\nu_{t-1} \;+\; O(\|\Lambda\|^2) \\
&=\; \nu_t + \Sigma_\gamma^{-1}\Lambda E_t[\nu_{t+1} - \nu_t] + m(\nu_{t-1} - \nu_t) + \; O(\|\Lambda\|^2) \\
&=\; \nu_t + \Sigma_\gamma^{-1}\Lambda(E_t[\nu_{t+1} - \nu_t] + (\nu_{t-1} - \nu_t)) + \; O(\|\Lambda\|^2)
\end{aligned}
\tag{74}
$$

Hence

$$
\begin{aligned}
&E[\mu(x_t)'F_t - 0.5F_t'\Sigma_\gamma F_t - 0.5(F_t - F_{t-1})'\Lambda(F_t - F_{t-1})] \\
&=\; E[\mu(x_t)'F_t - 0.5F_t'(\Sigma_\gamma + 2\Lambda)F_t + F_{t-1}'\Lambda F_t] \\
&=\; E[\mu_t'(\nu_t + \Sigma_\gamma^{-1}\Lambda(E_t[\nu_{t+1} - \nu_t] \;+\; (\nu_{t-1} - \nu_t))) \\
&- 0.5(\nu_t + \Sigma_\gamma^{-1}\Lambda(E_t[\nu_{t+1} - \nu_t] \;+\; (\nu_{t-1} - \nu_t)))(\Sigma_\gamma + 2\Lambda)(\nu_t + \Sigma_\gamma^{-1}\Lambda(E_t[\nu_{t+1} - \nu_t] \;+\; (\nu_{t-1} - \nu_t))) \\
&+ (\nu_{t-1} + \Sigma_\gamma^{-1}\Lambda(E_{t-1}[\nu_t - \nu_{t-1}] \;+\; (\nu_{t-2} - \nu_{t-1})))'\Lambda(\nu_t + \Sigma_\gamma^{-1}\Lambda(E_t[\nu_{t+1} - \nu_t] \;+\; (\nu_{t-1} - \nu_t)))] \\
&=\; 0.5E[\mu_t'\Sigma_\gamma^{-1}\mu_t] \;- E[\nu_t'\Lambda\nu_t] \\
&+\; E[\nu_t'\Lambda(E_t[\nu_{t+1} - \nu_t] \;+\; (\nu_{t-1} - \nu_t))] \\
&- E[\nu_t'\Lambda(E_t[\nu_{t+1} - \nu_t] \;+\; (\nu_{t-1} - \nu_t))] \\
&+ E[\nu_{t-1}\Lambda\nu_t] \;+\; O(\|\Lambda\|^2) \\
&=\; 0.5E[\mu_t'\Sigma_\gamma^{-1}\mu_t] \;- 0.5(2E[\nu_t'\Lambda\nu_t] + 2E[\nu_{t-1}\Lambda\nu_t]) \;+ O(\|\Lambda\|^2) \\
&=\; 0.5E[\mu_t'\Sigma_\gamma^{-1}\mu_t] \;- 0.5(E[\nu_t'\Lambda\nu_t] + E[\nu_{t-1}'\Lambda\nu_{t-1}] + 2E[\nu_{t-1}\Lambda\nu_t]) \;+ O(\|\Lambda\|^2) \\
&=\; 0.5E[\mu_t'\Sigma_\gamma^{-1}\mu_t] \;- 0.5E[(\nu_t - \nu_{t-1})'\Lambda(\nu_t - \nu_{t-1})] + O(\|\Lambda\|^2)\,.
\end{aligned}
$$

□