Market Efficiency in the Age of Big Data[*]

Ian Martin[†]
London School of Economics
and CEPR

Stefan Nagel[‡]
University of Chicago,
NBER, CEPR, and CESIfo

December 2020

Modern investors face a high-dimensional prediction problem: thousands of observable variables are potentially relevant for forecasting. We reassess the conventional wisdom on market efficiency testing in light of this fact. In our model economy, N assets have cash flows that are a linear function of J firm characteristics, but with uncertain coefficients. Risk-neutral Bayesian investors impose shrinkage (ridge regression) or sparsity (Lasso) when they estimate the J coefficients of the model and use them to price assets. In equilibrium, when J is comparable in size to N, returns appear cross-sectionally predictable using firm characteristics to an econometrician who analyzes data from the economy ex post. A factor zoo emerges even without p-hacking and data-mining. Standard in-sample tests of market efficiency reject the no-predictability null with high probability, despite the fact that investors optimally use the information available to them in real time. In contrast, out-of-sample tests retain their economic meaning.

I. Introduction

Machine learning methods have proved useful in forecasting problems with huge numbers of predictor variables. High-dimensional prediction problems of this kind are faced not only by data scientists studying data as an outside observer, but also by economic decision-makers the marketplace. Many forward-looking economic decisions require predictions for which large numbers of variables could potentially be relevant, but the exact relationship between predictors and forecast target is unknown and must be learned from observed data. In this paper, we argue that to understand market outcomes in such settings, it is important to take into account the high-dimensional nature of decision-makers' prediction problem.

We demonstrate this in an asset-pricing setting. We show that properties of asset prices are strongly affected by the dimensionality of investors' prediction problem. Conventional notions of how to test market efficiency and how to interpret pricing anomalies break down in the high-dimensional case.

To price risky assets such as stocks, investors must forecast the future cash flows generated by these assets. In our model, cash-flow growth rates of a cross-section of N firms are a linear function of J firm characteristics that are fixed over time. Investors are Bayesian, homogeneous, risk-neutral, and price stocks based on the predictive distribution of cash flows. Realized asset returns in this setting are simply equal to investors' forecast errors. If investors knew the coefficients of the predictive model, they could form expectations of cash-flow growth, and hence price assets, in such a way that returns would not be predictable in the cross-section. This is the conventional rational expectations (RE) equilibrium that is the foundation for typical market efficiency tests. Similarly, if J is small relative to N, investors could estimate the parameters of their cash-flow forecasting model with great precision, leading to asset prices that are close to those in the RE equilibrium.

In reality, however, investors face a myriad of potential predictor variables that could be relevant in constructing such forecasts. In other words, J is not small relative to N.

1

As technology has improved, the set of available and potentially valuation-relevant predictor variables has expanded enormously over time. Textual analysis, satellite imagery, social media data, and many other new data sources yield a wealth of information. But in order to use these sources of information in a forecasting model, investors must estimate the relationship between these signals and future cash flows. This is a high-dimensional learning problem. The number of potential predictor variables could easily surpass the number of assets whose cash flow data is available to estimate this relationship.

Machine learning methods handle this issue by imposing some regularization on the estimation, for example by shrinking parameter estimates towards a fixed target or by searching for a sparse model representation that includes only a small subset of variables from a much larger set of potential predictors. With the goal of optimizing out-of-sample forecasting performance, regularization lets the learner trade off the costs of downweighting certain pieces of information against the benefit of reduced parameter estimation error. In a Bayesian interpretation, shrinkage reflects informative prior beliefs: when forecasters know, based on economic plausibility considerations, that forecasting model parameters cannot have arbitrarily large magnitudes, their posterior beliefs are shrunk towards zero.

Shrinkage ameliorates, but does not eliminate, the effects of parameter uncertainty on asset prices in the high-dimensional case. Relative to the RE equilibrium, asset prices are distorted by two components. First, noise in the past cash-flow growth observations that investors learn from will have, by chance, some correlation with the J predictor variables. This induces error in investors' parameter estimates, and hence an asset price distortion, that is correlated with the J predictor variables. Shrinkage downweights this estimation error component, but it also gives rise to a second distortion component because shrinkage implies underweighting the predictive information in the J predictors. Naturally, this second component, too, is correlated with the J predictor variables.

To stack the deck against return predictability, we endow investors with prior beliefs that are objectively correct in the sense that the coefficients of the cash-flow generating model

2

are drawn from this prior distribution. Investors also know that this model is linear. With this objective prior, the optimal amount of shrinkage exactly balances the two components in such a way that investors' forecast errors, and hence also asset returns, are unpredictable out-of-sample.

That returns are not predictable out-of-sample does not imply absence of in-sample predictability, however. An econometrician conducting an in-sample predictability test uses data that was not available to investors in real time when they priced assets. In an RE setting, this would not matter, because investors would already have perfect knowledge of model parameters. Approximately, the same would be true in a low-dimensional setting with small J and large N, where investors would be able to estimate forecasting model parameters with high precision. But in a high-dimensional setting, the econometrician's ability to see data realized ex post, after investors' pricing decisions are made, gives her a substantial advantage.

To show this, we consider an econometrician who collects asset price data from our model economy ex post and runs in-sample regressions to test whether the J firm characteristics cross-sectionally predict returns. When J is vanishing in size relative to N, there is almost no predictability: with $N \to \infty$ and J fixed, the predictability test would reject the null with test size close to the chosen significance level (e.g., 0.05). In contrast, in high-dimensional asymptotics, where $N, J \to \infty$ jointly, with their ratio J/N converging to a fixed number, the econometrician would reject the no-predictability null hypothesis, in the limit, with probability one. In simulations, we show that we also obtain rejection probabilities close to one for finite N when J is comparable in size to N. This overwhelming rejection of the no-predictability null occurs despite the fact that investors optimally use the information available to them in real time. Importantly, the high rejection rates in in-sample tests are not caused by distortions in the sampling properties of the econometrician's test statistics. Instead, the high rejection rates correctly reflect the fact that the properties of equilibrium asset prices change as $N, J \to \infty$ and that in-sample predictability truly exists.

The situation is different for out-of-sample tests. In our model economy, a portfolio

formed based on the econometrician's predictive regression estimates up to period t, with positive weights for stocks with positive predicted returns and negative weights for stocks with negative expected returns, has an average return of zero in the subsequent period $t+1$. In other words, returns are not predictable out-of-sample. This is true, too, in the high-dimensional asymptotic case. Intuitively, since Bayesian investors optimally use information available to them and price asset such that asset returns are not predictable under their predictive distribution, an econometrician restricted to constructing return forecasts using only data that was available to investors in real time is not able to predict returns out-of-sample either.

These results illustrate forcefully that the economic content of the (semi-strong) market efficiency notion that prices "fully reflect" all public information (Fama 1970) is not clear in this high-dimensional setting, even though we abstract from the joint hypothesis problem by assuming that investors are risk-neutral.[1] Does "fully reflect" mean that investors know the parameters of the cash-flow prediction model (i.e., the typical notion of RE in asset pricing and macroeconomics)? Or does "fully reflect" mean that investors employ Bayesian updating when they learn from data about the parameters of the cash-flow prediction model? The null hypothesis in a vast empirical literature in asset pricing—including return predictability regressions, event studies, and asset pricing model estimation based on orthogonality conditions—is the former version of the market efficiency hypothesis. Our results show that testing and rejecting it has little economic content in a high-dimensional setting. An apparent rejection of market efficiency might simply represent the unsurprising consequence of investors not having precise knowledge of the parameters of a data-generating process that involves thousands of predictor variables.

Empirical discoveries of new cross-sectional return predictors that are statistically significant according to conventional in-sample tests are therefore less interesting in a world in which investors have to take into account the valuation implications of a large number

---

1. The joint hypothesis problem (Fama 1970) refers to the problem that the econometrician studying asset prices does not know the model that determines risk premia required by risk-averse investors.

of forecasting variables. From the perspective of our model, it is not surprising that the technology-driven explosion in the number of available predictor variables has coincided with an explosion in the number of return predictors that are found significant in asset pricing studies (Cochrane 2011; Harvey, Liu, and Zhu 2016). Even without p-hacking, multiple testing, and data mining (Lo and MacKinlay 1990; Harvey, Liu, and Zhu 2016; Chordia, Goyal, and Saretto 2019), evidence of cross-sectional return predictability from in-sample regressions does not tell us much about the expected returns that investors perceived ex ante at the time they priced assets. Thus out-of-sample tests (such as those in McLean and Pontiff (2016)) gain additional importance in the age of Big Data.

Researchers are often skeptical of out-of-sample tests. In the case where a fixed underlying process is generating returns, as would be the case in many RE models, in- and out-of-sample methods test the same hypothesis—and in-sample tests are more powerful because they use the available data to the fullest extent. As a consequence, it is not clear why one would want to focus on out-of-sample tests (Inoue and Kilian 2005; Campbell and Thompson 2008; Cochrane 2008, Hansen and Timmermann 2015). In contrast, if investors face a high-dimensional learning problem, substantial in-sample predictability can coexist with absence of out-of-sample predictability. In-sample and out-of-sample tests examine fundamentally different hypotheses in this case. This provides a clear motivation for out-of-sample testing.

We illustrate the different perspectives provided by in- and out-of-sample tests with an empirical example. In the cross-section of U.S. stocks, we consider each stock's history of monthly simple and squared returns over the previous 120 months as a set of return predictors. Running a ridge regression over a full five decade sample, the in-sample coefficient estimates pick up the most prominent past return-based anomalies in the literature, including momentum (Jegadeesh and Titman 1993; Novy-Marx 2012), long-term reversals (DeBondt and Thaler 1985), and momentum seasonality (Heston and Sadka 2008). In other words, there is substantial in-sample predictability. In terms of out-of-sample predictability, the picture looks very different. Using rolling ridge regressions over 20-year windows to estimate

prediction model coefficients and then using those to predict returns in subsequent periods, we find that predictability is generally much weaker out of sample than in sample. Moreover, there is substantial decay over time. While some out-of-sample predictability exists in the early decades of the sample, it is basically nil in the most recent 15 years. This suggests that there may be little reason to seek risk-based or behavioral explanations of the cross-sectional predictability that shows up in the in-sample analysis.

That there was some out-of-sample predictability in the earlier parts of the sample may indicate the presence of ex-ante risk premia or mispricing at the time. It is also possible, however, that investors several decades ago were not able to process the information in each stock's price history as effectively as investors are able to do today. One can think of this as bounded rationality that induces excessive shrinkage or sparsity of investors' forecasting models, along the lines of Sims (2003), Gabaix (2014) and Molavi, Tahbaz-Salehi, and Vedolin (2020). We show in our simulations that sparsity or shrinkage beyond the level called for by objectively correct Bayesian priors leads to positive out-of-sample return predictability. With regard to empirical studies, this means that out-of-sample predictability in early parts of the sample may reflect the fact that some of the variables that researchers can retroactively construct today were not readily available to investors at the time, or could not feasibly be included in their forecasting models.

Overall, our results suggest that in-sample cross-sectional return predictability tests are ill-suited for uncovering return premia that require explanations based on priced risk exposures or behavioral biases. This is not to say that all of the documented patterns in the literature are explainable with learning and will not persist out of sample. But it is important to obtain other supporting evidence beyond in-sample predictability tests. If predictability associated with a predictor variable persists out of sample, if there is a compelling theoretical motivation, or if other types of data point to a risk or behavioral bias explanation (e.g., economic risk exposures or data on investor expectations), the case for a risk premium or a persistent behavioral bias is much stronger.

The insight that learning can induce in-sample return predictability in our setting relates to an earlier literature that studies time-series, rather than cross-sectional, return predictability (e.g., Timmermann 1993; Lewellen and Shanken 2002; Collin-Dufresne, Johannes, and Lochstoer 2016). This literature examines low-dimensional settings with few return predictors, but with samples that are sufficiently short in time so that learning effects matter. In contrast, in our model, we consider the case where a large cross-section would be sufficient to learn parameters precisely in a low-dimensional setting, but learning effects persist when the number of predictors is large.

Our approach has antecedents elsewhere in the literature. Aragones, Gilboa, Postlewaite, and Schmeidler (2005) and Al-Najjar (2009) treat decision makers as statisticians that have to learn from observed data in a non-Bayesian high-dimensional setting. Their focus is on conditions under which disagreement between agents can persist in the long-run. Klein (2019) and Calvano, Calzolari, Denicolò, and Pastorello (2018) focus on strategic interaction of machine learning pricing algorithms in product markets. Investors in our setting face a simpler learning problem within a Bayesian linear framework and without strategic interactions. Even in this simple setting, important pricing implications emerge.

All proofs are in the appendix.

II. Bayesian pricing in a high-dimensional setting

Consider an economy in discrete time, $t \in \{1, 2, ...\}$, with N assets. Each asset is associated with a vector of J firm characteristics observable to investors that we collect in the $N \times J$ matrix $\boldsymbol{X}$. The assets pay dividends, collected in the vector $\boldsymbol{y}_t$, whose growth $\Delta \boldsymbol{y}_t = \boldsymbol{y}_t - \boldsymbol{y}_{t-1}$ is partly predictable based on $\boldsymbol{X}$:

Assumption 1.

$$\Delta \boldsymbol{y}_t = \boldsymbol{X}\boldsymbol{g} + \boldsymbol{e}_t, \quad \boldsymbol{e}_t \sim N(\boldsymbol{0}, \boldsymbol{\Sigma}_e), \quad \text{rank}(\boldsymbol{X}) = J, \quad \frac{1}{NJ}\text{tr}\,\boldsymbol{X}'\boldsymbol{X} = 1. \tag{1}$$

The set of characteristics is potentially very large, but for simplicity we assume $J < N$. It would be relatively straightforward to extend the framework to allow for $J \geq N$, but the main points can be seen more clearly in the simpler $J < N$ setting. The set of characteristics in $\boldsymbol{X}$ exhausts the set of variables that investors can condition on. Due to technological change, this set could change as previously unavailable predictors become available. And so we will be concerned with the behavior of prices for various values of J.

The assumption that $\frac{1}{NJ} \operatorname{tr} \boldsymbol{X}'\boldsymbol{X} = 1$ is a normalization that defines a natural scale for the characteristics. For example, it holds if characteristics are scaled to have unit norm (i.e., if $\frac{1}{N} \sum_{n=1}^{N} x_{nj}^2 = 1$ for every characteristic j), as then

$$\frac{1}{NJ} \operatorname{tr} \boldsymbol{X}'\boldsymbol{X} = \frac{1}{J} \sum_{j=1}^{J} \frac{1}{N} \sum_{n=1}^{N} x_{nj}^2 = 1 \, .$$

We assume that the characteristics associated with a firm are constant over time for simplicity. In reality, firms' characteristics change. But as long as investors know the firm's characteristics at every point in time one can accommodate this in our setting by thinking of $\boldsymbol{y}_t$ as a vector of payoffs for hypothetical characteristics-constant firms. We would have to reshuffle firms each period so that each element of $\boldsymbol{y}_t$ is always associated with the same characteristics.

We further make the following assumption:

Assumption 2. *Investors are risk-neutral and the interest rate is zero.*

By abstracting from risk premia, we intentionally make it easy for an econometrician to test market efficiency in our setting. With risk-neutral investors, there is no joint hypothesis problem due to unknown risk pricing models. Yet, as we will show, interpretation of standard market efficiency tests is still tricky.

We focus on the pricing of one-period dividend strips, so that $\boldsymbol{p}_t$ represents the vector of prices, at time t, of claims to dividends paid at time $t+1$. We think of one period in this model as a long time span, say a decade, so that the errors in $\boldsymbol{e}_t$ are actually the averages of

the errors one would find if one sampled at higher frequencies over many shorter sub-periods. With this interpretation in mind, we can then think of the dividend strip payoff as a long-lived stock's cash flows compressed into a single cash flow at the typical duration of a stock (e.g., perhaps a decade).

The price vector is then equal to the vector of next-period expected dividends,

$$\boldsymbol{p}_t = \widetilde{\mathbb{E}}_t \boldsymbol{y}_{t+1} = \boldsymbol{y}_t + \widetilde{\mathbb{E}}_t \Delta \boldsymbol{y}_{t+1} = \boldsymbol{y}_t + \widetilde{\mathbb{E}}_t \left(\boldsymbol{X} \boldsymbol{g} + \boldsymbol{e}_{t+1} \right).$$

This formulation encompasses a range of possible assumptions about the process by which investors' expectations $\widetilde{\mathbb{E}}_t \left[\cdot \right]$ are formed.

In a rational expectations model, for example, investors know $\boldsymbol{g}$, so that $\widetilde{\mathbb{E}}_t \left(\boldsymbol{X} \boldsymbol{g} + \boldsymbol{e}_{t+1} \right) = \boldsymbol{X} \boldsymbol{g}$. The dividend strip price is therefore $\boldsymbol{p}_t = \boldsymbol{y}_t + \boldsymbol{X} \boldsymbol{g}$ and realized price changes $\boldsymbol{r}_{t+1} = \boldsymbol{y}_{t+1} - \boldsymbol{p}_t = \Delta \boldsymbol{y}_{t+1} - \boldsymbol{X} \boldsymbol{g} = \boldsymbol{e}_{t+1}$ are unpredictable with $\boldsymbol{X}$. This is the usual null hypothesis that underlies tests based on orthogonality conditions and Euler equations.

However, we focus on the realistic case where investors don't know $\boldsymbol{g}$. They therefore face a learning problem in pricing assets. They can learn about $\boldsymbol{g}$ by observing the realizations of $\{\Delta \boldsymbol{y}_s\}_1^t$ and the characteristics $\boldsymbol{X}$. (We assume that investors know $\boldsymbol{\Sigma}_e$.) We then have $\widetilde{\mathbb{E}}_t \left(\boldsymbol{X} \boldsymbol{g} + \boldsymbol{e}_{t+1} \right) = \boldsymbol{X} \tilde{\boldsymbol{g}}_t$, where $\tilde{\boldsymbol{g}}_t$ represents investors' posterior mean of $\boldsymbol{g}$ at time t, after learning from historical data.

If J is close to (or perhaps even larger than) N, running an OLS regression to estimate $\boldsymbol{g}$ would not give investors useful forecasts. For example, with $J = N$, a cross-sectional regression of $\Delta \boldsymbol{y}_t$ on $\boldsymbol{X}$ exactly fits $\Delta \boldsymbol{y}_t$ in sample. Then $\widetilde{\mathbb{E}}_t \left(\boldsymbol{X} \boldsymbol{g} + \boldsymbol{e}_{t+1} \right) = \Delta \boldsymbol{y}_t$ so that $\boldsymbol{p}_t = \boldsymbol{y}_t + \Delta \boldsymbol{y}_t$ and $\boldsymbol{r}_{t+1} = \Delta \boldsymbol{y}_{t+1} - \Delta \boldsymbol{y}_t$. The forecast mean squared error (MSE) is then $\text{var} \left(\boldsymbol{e}_{t+1} - \boldsymbol{e}_t \right)$, i.e., twice the variance of the truly unpredictable $\boldsymbol{e}_{t+1}$.

For comparison, the naive "random walk" forecast that sets $\mathbb{E}_t \Delta \boldsymbol{y}_{t+1} = \boldsymbol{0}$ would result in $\boldsymbol{p}_t = \boldsymbol{y}_t$ and hence $\boldsymbol{r}_{t+1} = \Delta \boldsymbol{y}_{t+1}$. In this case, the forecast MSE is $\text{var} \left(\boldsymbol{X} \boldsymbol{g} + \boldsymbol{e}_{t+1} \right)$. If a relatively small component of cash-flow growth is predictable—that is, if $\text{var} \left(\boldsymbol{X} \boldsymbol{g} \right) \ll \text{var} \, \boldsymbol{e}_{t+1}$—then the random walk forecast MSE may be substantially lower than the OLS

forecast MSE.

II.A.   Priors and posteriors

The problem with least-squares regression forecasts is that the prior implicit in the least-squares estimator is economically unreasonable. The posterior mean equals the GLS estimator if investors' prior for $\boldsymbol{g}$ is diffuse. But a diffuse prior for $\boldsymbol{g}$ is not a plausible assumption. Economic reasoning should lead investors to realize that the amount of predictable variation in $\Delta \boldsymbol{y}_{t+1}$ must be limited. It does not make economic sense for investors to believe that arbitrarily large values for $\boldsymbol{g}$ are just as likely as values that give rise to moderate predictable variation in $\Delta \boldsymbol{y}_{t+1}$. While they might not have very precise prior knowledge of $\boldsymbol{g}$, it reasonable to assume that the distribution representing investors' prior beliefs about $\boldsymbol{g}$ is concentrated around moderate values of $\boldsymbol{g}$.

We therefore make the following specification of prior beliefs.

Assumption 3. *Before seeing data, investors hold prior beliefs*

$$\boldsymbol{g} \sim N(\boldsymbol{0}, \boldsymbol{\Sigma}_g).$$

That prior beliefs are centered around zero means that investors a priori don't know which characteristics predict cash-flow growth by how much. But they know that magnitudes of $\boldsymbol{g}$ elements cannot be too big. Economic restrictions on $\boldsymbol{\Sigma}_g$ that restrict the likely magnitudes of $\boldsymbol{g}$ elements will play an important role later on in our analysis.

Proposition 1. *After investors have observed dividend growth in a single period, $\Delta \boldsymbol{y}_1$, their posterior distribution of $\boldsymbol{g}$ is $\boldsymbol{g} | \Delta \boldsymbol{y}_1, \boldsymbol{X} \sim N(\tilde{\boldsymbol{g}}_1, \boldsymbol{D}_1)$, where*

$$\tilde{\boldsymbol{g}}_1 = \boldsymbol{D}_1 \boldsymbol{d}_1$$
$$\boldsymbol{D}_1^{-1} = \boldsymbol{\Sigma}_g^{-1} + \boldsymbol{X}' \boldsymbol{\Sigma}_e^{-1} \boldsymbol{X},$$
$$\boldsymbol{d}_1 = \boldsymbol{X}' \boldsymbol{\Sigma}_e^{-1} \Delta \boldsymbol{y}_1.$$

After observing data for t periods, the posterior mean is $\tilde{\boldsymbol{g}}_t = \boldsymbol{D}_t \boldsymbol{d}_t$ and

$$\boldsymbol{D}_t^{-1} = \boldsymbol{\Sigma}_g^{-1} + t\boldsymbol{X}'\boldsymbol{\Sigma}_e^{-1}\boldsymbol{X},$$

$$\boldsymbol{d}_t = t\boldsymbol{X}'\boldsymbol{\Sigma}_e^{-1}\overline{\Delta\boldsymbol{y}}_t,$$

where $\overline{\Delta\boldsymbol{y}}_t = \frac{1}{t}\sum_{s=1}^{t}\Delta\boldsymbol{y}_s$. Therefore

$$\tilde{\boldsymbol{g}}_t = \left[\frac{1}{t}\boldsymbol{\Sigma}_g^{-1} + \boldsymbol{X}'\boldsymbol{\Sigma}_e^{-1}\boldsymbol{X}\right]^{-1}\boldsymbol{X}'\boldsymbol{\Sigma}_e^{-1}\overline{\Delta\boldsymbol{y}}_t. \tag{2}$$

The posterior mean $\tilde{\boldsymbol{g}}_1$ takes the form of a Tikhonov-regularized (i.e., ridge) regression estimator, where the inverse of $\boldsymbol{D}_1^{-1}$ is "stabilized" by adding $\boldsymbol{\Sigma}_g^{-1}$ to $\boldsymbol{X}'\boldsymbol{\Sigma}_e^{-1}\boldsymbol{X}$. Thus our Bayesian framework connects to a large literature in machine learning in which Tikhonov regularization is used to deal with high-dimensional prediction problems (see, for example, Shalev-Shwartz and Ben-David (2014)).[2]

III. ASYMPTOTIC ANALYSIS

We now analyze the properties of asset prices in high-dimensional asymptotic analysis when $N, J \to \infty$, where $J/N \to \psi > 0$, where ψ is a fixed number.[3] This differs from the usual low-dimensional large N, fixed J (or large T, fixed N and fixed J) asymptotics that underlie

---

2. To interpret the posterior mean in terms of standard regression estimators, we can use the Woodbury identity to write

$$\tilde{\boldsymbol{g}}_1 = \boldsymbol{\Sigma}_g \left\{ \boldsymbol{\Sigma}_g + \left(\boldsymbol{X}'\boldsymbol{\Sigma}_e^{-1}\boldsymbol{X}\right)^{-1}\right\}^{-1}\tilde{\boldsymbol{g}}_{GLS,1}$$

$$\tilde{\boldsymbol{g}}_{GLS,1} = \left(\boldsymbol{X}'\boldsymbol{\Sigma}_e^{-1}\boldsymbol{X}\right)^{-1}\boldsymbol{X}'\boldsymbol{\Sigma}_e^{-1}\Delta\boldsymbol{y}_1.$$

This shows that the posterior mean is a weighted average of the prior mean (zero) and the GLS regression estimator $\tilde{\boldsymbol{g}}_{GLS,1}$.

3. See, e.g., Anatolyev (2012) and Dobriban and Wager (2018) for recent examples from the econometrics and statistics literature on high-dimensional regression that use this type of asymptotic analysis. This literature focuses on the asymptotic properties of estimators and statistical tests given an underlying data-generating model that stays fixed as N and J change. In contrast, in our case the nature of investors' learning problem changes as N and J change, and hence the properties of the data that the econometrician analyzes change as well.

most econometric methods in asset pricing.

We first simplify the setup by making

Assumption 4.

$$\boldsymbol{\Sigma}_e = \boldsymbol{I}$$
$$\boldsymbol{\Sigma}_g = \frac{\theta}{J}\boldsymbol{I}, \quad \theta > 0$$

Our results go through for a general (nonsingular) covariance matrix $\boldsymbol{\Sigma}_e$, i.e., with a factor structure in residuals, though at the cost of some extra notational complexity. By assuming $\boldsymbol{\Sigma}_e = \boldsymbol{I}$, we are making the learning problem easy for investors. With uncorrelated residuals, investors achieve a given posterior precision with a smaller J than if residuals were uncorrelated.

By assuming that $\boldsymbol{\Sigma}_g$ is proportional to the identity, we put all the predictor variables on an equal footing from the prior perspective; and it is essential that the variance of the elements of $\boldsymbol{g}$ should decline with J in order to consider sensible asymptotic limits. To see this, note that the covariance matrix of $\boldsymbol{Xg}$ is $\boldsymbol{X}\boldsymbol{\Sigma}_g\boldsymbol{X}'$, so the cross-sectional average prior variance of the predictable component of cash-flow growth satisfies

$$\frac{1}{N}\sum_{i=1}^{N}\left(\boldsymbol{X}\boldsymbol{\Sigma}_g\boldsymbol{X}'\right)_{ii} \stackrel{(A4)}{=} \frac{\theta}{JN}\sum_{i=1}^{N}\sum_{j=1}^{J}x_{ij}^2 \stackrel{(A1)}{=} \theta, \tag{3}$$

using Assumptions 1 and 4.

We view the matrix of characteristics, $\boldsymbol{X}$, as observable and form the eigendecomposition

$$\frac{1}{N}\boldsymbol{X}'\boldsymbol{X} = \boldsymbol{Q}\boldsymbol{\Lambda}\boldsymbol{Q}', \tag{4}$$

where $\boldsymbol{\Lambda}$ is diagonal with the eigenvalues λ_j of $\frac{1}{N}\boldsymbol{X}'\boldsymbol{X}$ on its diagonal and $\boldsymbol{Q}$ is an orthogonal matrix (i.e., $\boldsymbol{Q}'\boldsymbol{Q} = \boldsymbol{Q}\boldsymbol{Q}' = \boldsymbol{I}$) whose columns are the corresponding eigenvectors of $\frac{1}{N}\boldsymbol{X}'\boldsymbol{X}$. This is possible because $\frac{1}{N}\boldsymbol{X}'\boldsymbol{X}$ is symmetric. As it is also positive definite, the eigenvalues

satisfy $\lambda_j > 0$ for all j. Lastly, the normalization tr $\boldsymbol{X}'\boldsymbol{X} = NJ$ (Assumption 1) implies that the average eigenvalue equals one:

$$\frac{1}{J}\sum_{i=1}^{J}\lambda_i = \frac{1}{J}\operatorname{tr}\frac{1}{N}\boldsymbol{X}'\boldsymbol{X} = 1\,, \tag{5}$$

using the fact that the sum of the eigenvalues of a matrix equals its trace.

If $\boldsymbol{X}'\boldsymbol{X}$ has eigenvalues that are very close to zero, then the columns of $\boldsymbol{X}$ are roughly collinear. For, to find a linear combination $\boldsymbol{v} \in \mathbb{R}^J$ of columns of $\boldsymbol{X}$ with the property that $\boldsymbol{X}\boldsymbol{v}$ is small (where $\boldsymbol{v}$ is a unit vector, $\boldsymbol{v}'\boldsymbol{v} = 1$), we can choose $\boldsymbol{v}$ to be a unit eigenvector of $\boldsymbol{X}'\boldsymbol{X}$ with minimal eigenvalue—call it $\lambda_{\min}$—so that $(\boldsymbol{X}\boldsymbol{v})'(\boldsymbol{X}\boldsymbol{v}) = \lambda_{\min} \approx 0$. Thus if there are eigenvalues close to zero then some characteristics are approximately spanned by other characteristics. Our next assumption can therefore be thought of as ensuring that we really are in a Big Data environment.

Assumption 5. *The eigenvalues λ_j of $\frac{1}{N}\boldsymbol{X}'\boldsymbol{X}$ satisfy $\lambda_j > \varepsilon$ for all j, where $\varepsilon > 0$ is a uniform constant as $N \to \infty$.*

In fact all we need for our main results to go through is that λ_j does not tend to zero at a rate faster than $1/\sqrt{j}$. Assumption 5 is therefore stronger than we need, but it will be satisfied in our applications below.

Combining Assumption 4 with equation (2), we obtain

$$\tilde{\boldsymbol{g}}_t = \left[\frac{J}{\theta t}\boldsymbol{I} + \boldsymbol{X}'\boldsymbol{X}\right]^{-1}\boldsymbol{X}'\overline{\Delta\boldsymbol{y}}_t$$
$$= \boldsymbol{\Gamma}_t(\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\overline{\Delta\boldsymbol{y}}_t\,, \tag{6}$$

where

$$\boldsymbol{\Gamma}_t = \boldsymbol{Q}\left(\boldsymbol{I} + \frac{J}{N\theta t}\boldsymbol{\Lambda}^{-1}\right)^{-1}\boldsymbol{Q}' \tag{7}$$

is a symmetric matrix. The posterior mean $\tilde{\boldsymbol{g}}_t$ shrinks the naive OLS estimate (from a regression of $\overline{\Delta\boldsymbol{y}}_t$ onto the columns of $\boldsymbol{X}$) along the principal components. Shrinkage is a

13

consequence of the informative prior for $\boldsymbol{g}$. The prior's influence on the posterior is stronger if the observed data is less informative relative to the prior. To see explicitly what the degree of shrinkage depends on, note that $\left(\boldsymbol{I} + \frac{J}{N\theta t}\boldsymbol{\Lambda}^{-1}\right)^{-1}$ is diagonal with elements

$$\frac{\lambda_j}{\lambda_j + \frac{J}{N\theta t}}$$

along its diagonal. Thus shrinkage is strong if t or θ are small, or J/N is large, or along principal components with small eigenvalues.

We are now in a position to characterize the behavior of realized returns in equilibrium.

Proposition 2. *With assets priced based on $\tilde{\boldsymbol{g}}_t$, realized returns are*

$$\boldsymbol{r}_{t+1} = \boldsymbol{y}_{t+1} - \boldsymbol{p}_t = \boldsymbol{X}(\boldsymbol{I} - \boldsymbol{\Gamma}_t)\boldsymbol{g} - \boldsymbol{X}\boldsymbol{\Gamma}_t(\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\bar{\boldsymbol{e}}_t + \boldsymbol{e}_{t+1} \tag{8}$$

where $\bar{\boldsymbol{e}}_t = \frac{1}{t}\sum_{s=1}^{t}\boldsymbol{e}_s$.

Realized returns thus have three components. The first term on the right-hand side of (8) reflects the effect of "underreaction," due to shrinkage, to the fundamental information in $\boldsymbol{X}$. If investors had an uninformative prior ($\theta \to \infty$ and hence $\boldsymbol{\Gamma}_t \to \boldsymbol{I}$) as in many conventional low-dimensional Bayesian learning models, this term would not be present. But as we have argued, such an uninformative prior would imply that investors entertain an unreasonable amount of predictable variation in dividend growth. Under investors' informed prior beliefs, this part is still zero in expectation because the prior mean of $\boldsymbol{g}$ is zero, but for a given draw of $\boldsymbol{g}$ that generates the data that an econometrician would study, it is not zero.

The second term represents the effect of noise on investors' posterior mean. To the extent that the unpredictable shocks in $\bar{\boldsymbol{e}}_t$ in a given sample line up, by chance, with columns of $\boldsymbol{X}$, this induces estimation error that tilts investors cash-flow growth forecast away from $\boldsymbol{X}\boldsymbol{g}$. Shrinkage via $\boldsymbol{\Gamma}_t$ reduces this component, at the cost of generating the first term. Under Bayesian learning, $\boldsymbol{\Gamma}_t$ optimally trades off the pricing error arising from these two components.

14

The third term is the unpredictable shock $\boldsymbol{e}_{t+1}$. In the rational expectations case where $\boldsymbol{g}$ is known to investors, the realized return would simply be equal to $\boldsymbol{e}_{t+1}$ and the first two terms would not exist. In the Bayesian learning case, however, the first two terms are not zero, and, as a consequence, returns contain components correlated with the columns of $\boldsymbol{X}$. As we show now, these components may induce certain forms of return predictability.

III.A. In-sample predictability

We consider an econometrician who studies these realized returns with the usual tools of frequentist statistics. The econometrician looks for return predictability by regressing $\boldsymbol{r}_{t+1}$ on $\boldsymbol{X}$, using OLS.[4] The econometrician obtains a vector of cross-sectional regression coefficients

$$
\begin{aligned}
\boldsymbol{h}_{t+1} &= \left(\boldsymbol{X}'\boldsymbol{X}\right)^{-1}\boldsymbol{X}'\boldsymbol{r}_{t+1} \\
&= (\boldsymbol{I} - \boldsymbol{\Gamma}_t)\boldsymbol{g} - \boldsymbol{\Gamma}_t\left(\boldsymbol{X}'\boldsymbol{X}\right)^{-1}\boldsymbol{X}'\bar{\boldsymbol{e}}_t + \left(\boldsymbol{X}'\boldsymbol{X}\right)^{-1}\boldsymbol{X}'\boldsymbol{e}_{t+1}.
\end{aligned}
\tag{9}
$$

Following the logic of rational expectations econometrics, which assumes that investors price assets under knowledge of $\boldsymbol{g}$, the econometrician entertains $\boldsymbol{r}_{t+1} = \boldsymbol{e}_{t+1}$ as the no-predictability null hypothesis. Under this hypothesis, the first two terms in (9) would be zero. Given that the elements of $\boldsymbol{e}_{t+1}$ are distributed $N(0,1)$, it would follow, under this null, that

$$
\sqrt{N}\boldsymbol{h}_{t+1} \sim N\left(0, N(\boldsymbol{X}'\boldsymbol{X})^{-1}\right),
\tag{10}
$$

using the usual OLS asymptotic variance formulas.[5] In deriving these properties, the econometrician conditions on the observed predictors $\boldsymbol{X}$. From (10) it would follow—again, under

---

4. Given our assumption that $\boldsymbol{\Sigma}_e = \boldsymbol{I}$, OLS and GLS coincide. The econometrician could also use shrinkage methods like ridge regression, effectively imposing a prior that the coefficients in the return predictability regression cannot be too big. To the extent that the implied prior distribution of the coefficients is roughly in line with the true distribution of the coefficients, using such methods would strengthen the in-sample return predictability. We don't show formal results for this case, but we have explored the issue in simulations.

5. Recall that we assume $\boldsymbol{\Sigma}_e = \boldsymbol{I}$ (Assumption 4). More generally, if the econometrician has to estimate $\boldsymbol{\Sigma}_e$, this can be done based on the regression residual $\boldsymbol{\xi}_{t+1} = \boldsymbol{r}_{t+1} - \boldsymbol{X}\boldsymbol{h}_{t+1} = \left[\boldsymbol{I} - \boldsymbol{X}(\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\right]\boldsymbol{e}_{t+1}$, used to estimate the variance as $(\boldsymbol{X}'\boldsymbol{X})^{-1}\frac{1}{N-J}\boldsymbol{\xi}'_{t+1}\boldsymbol{\xi}_{t+1}$ which would estimate the variance consistently under conventional large-N, fixed-J asymptotics.

the econometrician's rational expectations null—that

$$h'_{t+1}(X'X)h_{t+1} \sim \chi^2_J. \tag{11}$$

As we want to characterize the properties of the econometrician's test under asymptotics where $N, J \to \infty$ and $J/N \to \psi > 0$, it is more convenient if we let the econometrician consider a scaled version of this test statistic:

$$T_{re} \equiv \frac{h'_{t+1}X'Xh_{t+1} - J}{\sqrt{2J}}. \tag{12}$$

Under the econometrician's rational expectations null, we would have, asymptotically,

$$T_{re} \xrightarrow{d} N(0,1) \quad \text{as } N, J \to \infty, \ J/N \to \psi > 0. \tag{13}$$

But the actual asymptotic distribution of T_{re} is influenced by the terms involving $\bar{e}_t$ and g in (9). These alter the asymptotic distribution and may lead the rejection probabilities of a test using $N(0,1)$ critical values based on (13), or χ^2 critical values based on (11), to differ from the nominal size of the test.

Our first result characterizes the properties of this test statistic under the true model, according to which returns follow equation (8). In this analysis, we assume that g is drawn from the prior distribution. This assumption is conservative in the sense that investors' prior beliefs about the distribution of g are objectively correct. If investors' priors deviated from this distribution, this would be another source of return predictability.

To assess the performance of the rational expectations econometrician's test statistic in our setting, it is helpful to write

$$\Sigma_{re} = \left(X'X\right)^{-1} \quad \text{and} \quad \Sigma_b = \mathbb{E}\left(h_{t+1}h'_{t+1}\right)$$

for the covariance matrices of the predictive coefficient estimates under the (incorrect) rational

16

expectations null hypothesis and the true model, respectively. When returns are generated under the true model (8), the rational expectations econometrician will use inappropriately small standard errors, in the sense that $\boldsymbol{\Sigma}_b - \boldsymbol{\Sigma}_{re}$, $\boldsymbol{\Sigma}_b\boldsymbol{\Sigma}_{re}^{-1} - \boldsymbol{I}$, and $\boldsymbol{\Sigma}_{re}^{-1} - \boldsymbol{\Sigma}_b^{-1}$ are all positive definite.[6]

Our first result shows that the first two cross-sectional moments of the eigenvalues of $\boldsymbol{\Sigma}_b\boldsymbol{\Sigma}_{re}^{-1}$ characterize the asymptotic behavior of T_{re}. These eigenvalues $(\zeta_{j,t})$ can be written explicitly in terms of the eigenvalues (λ_j) of $\frac{1}{N}\boldsymbol{X}'\boldsymbol{X}$ as

$$\zeta_{j,t} = 1 + \frac{\lambda_j}{t\lambda_j + \frac{\psi}{\theta}}. \tag{14}$$

As $\lambda_j > 0$ and $\psi/\theta > 0$, we have $1 < \zeta_{j,t} < 2$ for all $t \geq 1$. We write the limiting mean and variance of the eigenvalues as

$$\mu = \lim_{J\to\infty} \frac{1}{J}\sum_{j=1}^{J} \zeta_{j,t} \quad \text{and} \quad \sigma^2 = \lim_{J\to\infty} \frac{1}{J}\sum_{j=1}^{J} \zeta_{j,t}^2 - \mu^2.$$

By the "Big Data" Assumption 5, we have $\mu \in (1,2)$ and $\sqrt{\mu^2 + \sigma^2} \in (1,2)$ for all $t \geq 1$. Without it, we could have $\mu = 1$ and $\sigma = 0$ if $\lambda_j \to 0$.

Proposition 3. *If returns are generated according to (8) then in the large N, J limit*

$$\frac{\boldsymbol{h}_{t+1}'\boldsymbol{X}'\boldsymbol{X}\boldsymbol{h}_{t+1} - \sum_{j=1}^{J}\zeta_{j,t}}{\sqrt{2\sum_{j=1}^{J}\zeta_{j,t}^2}} \xrightarrow{d} N(0,1).$$

Thus the rational expectations case (13) is analogous to assuming $\zeta_{j,t} = 1$ for all j.

It follows that the test statistic T_{re} satisfies

$$\frac{T_{re}}{\sqrt{\mu^2 + \sigma^2}} - \frac{\mu - 1}{\sqrt{2(\mu^2 + \sigma^2)}}\sqrt{J} \xrightarrow{d} N(0,1)$$

---

6. For example, in the notation of the proof of Proposition 3, we can write $\boldsymbol{\Sigma}_b\boldsymbol{\Sigma}_{re}^{-1} - \boldsymbol{I} = N\mathbb{E}[\boldsymbol{h}_{t+1}\boldsymbol{h}_{t+1}']\frac{1}{N}\boldsymbol{X}'\boldsymbol{X} - \boldsymbol{I} = \boldsymbol{Q}\boldsymbol{\Omega}\boldsymbol{Q}'\boldsymbol{Q}\boldsymbol{\Lambda}\boldsymbol{Q}' - \boldsymbol{Q}\boldsymbol{Q}' = \boldsymbol{Q}(\boldsymbol{\Omega}\boldsymbol{\Lambda} - \boldsymbol{I})\boldsymbol{Q}'$. Thus $\boldsymbol{\Sigma}_b\boldsymbol{\Sigma}_{re}^{-1} - \boldsymbol{I}$ is symmetric and its eigenvalues are all positive (as the diagonal matrix $\boldsymbol{\Omega}\boldsymbol{\Lambda}$ has diagonal entries that are greater than one, as shown in the proof of Proposition 3). It is therefore positive definite.

where $1 < \mu < 2$ and $1 < \sqrt{\mu^2 + \sigma^2} < 2$.

We can therefore think of T_{re} as a multiple of a standard Normal random variable plus a term of order $\sqrt{J}$:

$$T_{re} \approx \sqrt{\mu^2 + \sigma^2}\, N(0,1) + \frac{\mu - 1}{\sqrt{2}} \sqrt{J}\,.$$

(For comparison, the rational expectations econometrician thinks T_{re} is asymptotically standard Normal, as in equation (13).) Thus the rejection probability tends rapidly to one as N and J tend to infinity.

Proposition 4. *In a test of return predictability based on the rational expectations null (13), we would have, for any critical value c_α and at any time t,*

$$\mathbb{P}\left(T_{re} > c_\alpha\right) \to 1 \ \ as \ N, J \to \infty, J/N \to \psi\,.$$

More precisely, for any fixed $t > 0$, the probability that the test fails to reject declines exponentially fast as N and J increase, at a rate that is determined by μ, σ, and ψ:

$$\lim_{N \to \infty} -\frac{1}{N} \log \mathbb{P}\left(T_{re} < c_\alpha\right) = \frac{(\mu - 1)^2 \psi}{4\left(\mu^2 + \sigma^2\right)}\,, \tag{15}$$

for any critical value c_α.

Thus in-sample predictability tests lose their economic meaning when J is not small relative to N, in the sense that the usual interpretation of in-sample return predictability evidence is not warranted.[7] The typical conclusion from rejections of the no-predictability null in studies of the cross-section of stock returns is that models of risk premia or mispricing due to imperfectly rational investors are needed to explain the evidence. Our model points

---

7. In the conventional low-dimensional asymptotics case, with J fixed as $N \to \infty$, the test rejects more often that it would under rational expectations, albeit to a small degree. In this case $\mathbf{\Gamma}_t \to \boldsymbol{I}$ so that the first term vanishes, but the second term in (9) is still relevant. This term shrinks with N at the same speed as the third term (which would be the only one under the rational expectations null), and leads to excessive rejection rates. Predictability vanishes in economic terms, however, as the return of the trading strategy that we will consider in Proposition 5 converges to zero if J is held fixed as $N \to \infty$.

to a third possibility: Investors are Bayesian, but in-sample return predictability arises even for large t because investors' forecasting problem is high-dimensional.

Another way to examine in-sample predictability is to consider a trading strategy with weights proportional to in-sample predicted returns, $\boldsymbol{Xh}_{t+1}$, according to the econometrician's regression coefficient estimates:

$$\boldsymbol{w}_{IS,t} = \frac{1}{N}\boldsymbol{Xh}_{t+1}.$$

We scale the predicted returns with $1/N$ so that if the distribution of predicted returns is kept fixed (normal, centered at zero, and fixed standard deviation), the expected sum of absolute portfolio weights remains constant as $N \to \infty$. This in turn means that, in expectation, the size of the combined long and short positions is constant as $N \to \infty$. To the extent that the expected portfolio return changes when $N, J \to \infty$, this would have to come from a change in the distribution of predicted returns as J grows rather than a mechanical effect due to greater N.

As we now show, however, the expected return converges to a constant asymptotically. More precisely, the objective expected return that an econometrician estimates by sampling repeatedly from this economy converges to a constant. In contrast, the Sharpe ratio explodes.

Proposition 5. *If returns are generated according to* (8), *then an in-sample trading strategy with weights* $\boldsymbol{w}_{IS,t} = \frac{1}{N}\boldsymbol{Xh}_{t+1}$ *and returns* $r_{IS,t+1} = \boldsymbol{r}'_{t+1}\boldsymbol{w}_{IS,t}$ *has expected returns that converge to a constant,*

$$\lim_{N,J\to\infty,J/N\to\psi} \mathbb{E}\, r_{IS,t+1} = \psi\mu > 0\,,$$

and a Sharpe ratio $SR_{IS} = \mathbb{E}\, r_{IS,t+1}/\operatorname{var}(r_{IS,t+1})^{1/2}$ *that grows at rate* $\sqrt{N}$,

$$\lim_{N,J\to\infty,J/N\to\psi} \frac{SR_{IS}}{\sqrt{N}} = \frac{\mu\sqrt{\psi}}{\sqrt{2\left(\mu^2 + \sigma^2\right)}}\,.$$

The expected return of the in-sample trading strategy also represents the amount of in-

19

sample predictable return variation in the cross-section because $\mathbb{E}[\boldsymbol{r}'_{t+1}\boldsymbol{w}_{IS,t}] = \frac{1}{N}\mathbb{E}[\boldsymbol{h}'_{t+1}\boldsymbol{X}'\boldsymbol{X}\boldsymbol{h}_{t+1}]$, i.e., the average portfolio return is equal to the average squared predicted returns. Our assumptions in (4) ensure—as is economically reasonable—that the amount of predictable variation in dividend growth stays bounded as N and J tend to infinity. As a result, the amount of predictable return variation approaches a finite limit. But the variance of the portfolio return shrinks with N and J, so that in the high-dimensional case in-sample moments seem to suggest incredible investment opportunities in mean-variance terms. But as we will see in Section III.C, these in-sample results do not represent returns of trading strategies that are implementable in real time.

III.B.   An example

With some additional assumptions on the matrix of firm characteristics $\boldsymbol{X}$, we can explicitly calculate the limit moments that appear in our propositions. We can then see, for example, how these limit moments depend on ψ, the limiting ratio of the number of predictors to number of observations.

Suppose that characteristics are determined at random, so that $\boldsymbol{X}$ has IID entries x_{ij} with mean zero, unit variance, and finite fourth moment. Nature generates this matrix once before investors start learning and it stays fixed thereafter. As before, investors know $\boldsymbol{X}$, and it stays fixed when we imagine an econometrician repeatedly sampling data by re-running the economy. Assumption 1 holds asymptotically, as $\frac{1}{NJ}\operatorname{tr}\boldsymbol{X}'\boldsymbol{X} = \frac{1}{NJ}\sum_{n,j}x_{n,j}^2 \to 1$ as $N, J \to \infty$ by the strong law of large numbers.

Moreover, we can use results from random matrix theory to characterize the distribution of the eigenvalues λ_i. In particular, the eigenvalue distribution converges to the Marchenko–Pastur distribution as $N, J \to \infty$ with $J/N = \psi$. For ψ close to one, this distribution features substantial probability mass on eigenvalues close to zero, indicating that many of the columns of $\boldsymbol{X}$ are close to being collinear. Nonetheless, the results of Yin, Bai, and Krishnaiah (1988) and Bai and Yin (1993) ensure that all the eigenvalues lie in a bounded interval that does not

20

contain the origin: $\lambda_j \in \left[\left(1 - \sqrt{\psi}\right)^2, \left(1 + \sqrt{\psi}\right)^2 \right]$ for all j. Thus Assumption 5 is satisfied.

Figure I shows histograms of the eigenvalue distributions in examples with $\boldsymbol{X}$ drawn randomly with $N(0,1)$ entries, setting $N = 1000$ and $J = 10$, 100, 500, and 900. Solid red lines in the figures illustrate the limiting Marchenko–Pastur distribution for the eigenvalues λ_j in each case; we also calculate the corresponding asymptotic distribution of the eigenvalues $\zeta_{j,t}$ by change of variable, using equation (33) in the appendix. When $\psi = J/N$ is close to one there is considerable mass near zero, implying that there are many approximately collinear relationships between the columns of $\boldsymbol{X}$. This is a realistic property that one would also find in actual empirical data if one assembled a huge matrix of firm characteristics.

Our next result characterizes the limiting cross-sectional mean, μ, and variance, σ^2, of the distribution of the eigenvalues $\zeta_{j,t}$.

Proposition 6. *The cross-sectional moments of $\zeta_{j,t}$ satisfy*

$$\mu = 1 + \frac{\psi + \theta t(\psi + 1) - \sqrt{[\psi + \theta t(\psi + 1)]^2 - 4\theta^2 t^2 \psi}}{2\theta t^2 \psi} \tag{16}$$

and

$$\sigma^2 = \frac{\theta^2 t^2 \psi - (\theta t + \psi)^2}{2\theta^2 t^4 \psi^2} + \frac{\theta t \psi \left(\theta^2 t^2 (\psi - 2) - \theta t \psi + \psi^2\right) + (\theta t + \psi)^3}{2\theta^2 t^4 \psi^2 \sqrt{[\psi + \theta t(\psi + 1)]^2 - 4\theta^2 t^2 \psi}}. \tag{17}$$

This result allows explicit calculation of the limit moments that appear in Propositions 3, 4, and 5. For example, Proposition 4 shows that the probability of rejecting the null of no predictability declines exponentially fast as N increases, and derives the rate of the exponential decay. Similarly, Proposition 5 shows that the expected return on a portfolio that uses in-sample information to construct portfolio weights is $\mathbb{E}\, r_{IS,t+1} = \psi\mu$. In both cases, Proposition 6 can be used to derive exact analytical expressions in terms of the model primitives θ, ψ, and t.

Figure II shows how the resulting expressions depend on ψ for $\theta = 1$ and $t = 1$. Panel (a) plots the rate function (15). For $\psi > 0.4$, the rate is higher than 0.015, indicating that

21

(A) Histogram of λ_j, $J/N = 0.05$.

(B) Histogram of $\zeta_{j,t}$, $J/N = 0.05$.

(C) Histogram of λ_j, $J/N = 0.5$.

(D) Histogram of $\zeta_{j,t}$, $J/N = 0.5$.

(E) Histogram of λ_j, $J/N = 0.9$.

(F) Histogram of $\zeta_{j,t}$, $J/N = 0.9$.

FIGURE I

Histograms of eigenvalue distributions in examples with $\theta = 1$, $t = 1$, $N = 1000$ and $J = 50, 500, 900$. The asymptotic distribution is shown as a solid line in each panel.

22

(A) The rate function given in equation (15)

(B) Expected return of in-sample trading strategy

Figure II

In-sample return predictability in the asymptotic limit with $\theta = 1$ and $t = 1$, for different values of ψ.

the probability of not rejecting the null is on the order of $\exp(-0.015N)$—which is a tiny number even for relatively small cross-sections of, say, $N \geq 300$. Panel (b) plots the expected in-sample portfolio return, $\mathbb{E}\, r_{IS,t+1} = \psi\mu$, from Proposition 5. As we noted earlier, $r_{IS,t+1}$ also represents the average squared in-sample predicted return. For comparison, the average squared in-sample unpredictable return is less than $\frac{1}{N} e'_{t+1} e_{t+1} \approx 1.0$. Therefore, the two panels show that for large ψ there is substantial in-sample predictability in, respectively, statistical and economic terms.

III.C.   (Absence of) out-of-sample return predictability

The situation looks very different with regards to out-of-sample predictability. We will now allow for the possibility that the econometrician only uses a subset of the variables in $\boldsymbol{X}$ that investors use to forecast fundamentals. Specifically, we assume that the econometrician's predictors are collected in the $N \times P$ matrix $\boldsymbol{C}$, where $P \leq J$ and $\text{rank}(\boldsymbol{C}) = P$.

Assumption 6. $\boldsymbol{X} = \begin{pmatrix} \boldsymbol{C} & \boldsymbol{M} \end{pmatrix}$ *where $\boldsymbol{M}$ is an $N \times (J - P)$ matrix and $\boldsymbol{C}'\boldsymbol{M} = \boldsymbol{0}$. Furthermore, the variables in $\boldsymbol{C}$ have been scaled such that $\text{tr}(\boldsymbol{C}'\boldsymbol{C}) = NP$. (As $\text{tr}\,\boldsymbol{X}'\boldsymbol{X} = \text{tr}\,\boldsymbol{C}'\boldsymbol{C} + \text{tr}\,\boldsymbol{M}'\boldsymbol{M}$, it follows that $\text{tr}(\boldsymbol{M}'\boldsymbol{M}) = N(J - P)$.)*

We view the econometrician's predictive variables $\boldsymbol{C}$ as fixed prior to observing the data. The assumption that $\boldsymbol{C}'\boldsymbol{M} = 0$ is without loss of generality, as we can think of the investors using the predictor variables observed by the econometrician, $\boldsymbol{C}$, together with extra unobserved variables $\boldsymbol{M}$ that are residualized with respect to $\boldsymbol{C}$.

The econometrician regresses $\boldsymbol{r}_{t+1}$ on $\boldsymbol{C}$, obtaining a coefficient vector $\boldsymbol{b}_{t+1}$ and predicted returns $\boldsymbol{C}\boldsymbol{b}_{t+1}$.

Proposition 7. *If returns are generated according to (8), then the predicted returns $\boldsymbol{C}\boldsymbol{b}_{s+1}$ from a regression estimated at time $s + 1$, are unrelated to returns at time $t + 1$ in the sense that*

$$\mathbb{E}\left[\boldsymbol{r}_{t+1}\left(\boldsymbol{C}\boldsymbol{b}_{s+1}\right)'\right] = 0 \tag{18}$$

whenever $t \neq s$.

The expression inside the expectation in equation (18) is an outer product, so Proposition 7 provides N^2 moment conditions. It implies in particular that any portfolio strategy whose weights are linear combinations of the elements of $\boldsymbol{C}\boldsymbol{b}_{s+1}$ will have expected return of zero whenever $t \neq s$. One useful special case of Proposition 7 that we will look at in more detail is a trading strategy with weights proportional to predicted returns

$$r_{OOS,t+1} = \boldsymbol{w}'_{OOS,s+1}\boldsymbol{r}_{t+1}, \qquad \boldsymbol{w}_{OOS,s+1} = \boldsymbol{C}\boldsymbol{b}_{s+1} \tag{19}$$

24

that is implementable at the end of period $s + 1$ based on the econometrician's regression coefficient estimates at that time. It follows directly from Proposition 8 that this strategy has expected return $\mathbb{E}\, r_{OOS,t+1} = 0$ whenever $t \neq s$.

In the forward prediction case $t > s$—using last period's estimated coefficients to form the portfolio—this is a natural result. Investors are Bayesian so the econometrician cannot "beat" investors in predicting returns as long as the econometrician is put on the same footing as investors in terms of the data that is available at the time of making the prediction.

That the result also applies backwards in time, with $t < s$, is more surprising. This case does not represent a tradable strategy, but it is interesting from an econometrician's perspective. The result suggests that the econometrician could conduct backwards out-of-sample tests. The fact that many cross-sectional asset-pricing anomalies do not hold up in backwards out-of-sample tests (Linnainmaa and Roberts 2018) could therefore be a consequence of investor learning, even without data-snooping on the part of researchers that published the original anomaly studies.

While the forward result is likely a general property of Bayesian learning (with objectively correct prior), the backwards result might be somewhat specific to the environment we have set up here (e.g., the assumption that cash-flow growth is IID over time). It is an interesting question for future research to what extent one can generalize the backwards result.

The absence of backwards OOS predictability is also interesting because it suggests that the common practice of using cross-validation for evaluating prediction models is justified in an environment like ours where investors face a learning problem. Many recent papers in the emerging machine learning literature in empirical asset pricing use cross-validation (e.g., Feng, Giglio, and Xiu 2020, Kozak, Nagel, and Santosh 2020, Bryzgalova, Pelger, and Zhu 2019). In cross-validation, the data is partitioned repeatedly into estimation and validation samples. For example, several blocks of calendar years are used for model estimation, and other blocks are held for OOS model validation. Then the blocks are switched and the procedure is repeated multiple times. The time ordering of blocks does not matter in this

procedure. Some validation blocks can therefore temporally precede some of the estimation blocks. This brings up the concern that cross-validation could pick up learning-induced in-sample predictability. Our backwards result in Proposition 7 suggests that this is not the case. Backwards prediction, forward prediction, and combinations of the two like in cross-validation, are equivalent. Further below, we explore whether cross-validation approaches could also be useful to assess the magnitude of OOS predictable cross-sectional variation in returns.

III.D.   *Testing for absence out-of-sample return predictability*

While the result in Proposition 7 that OOS trading strategies with weights based on return forecasts have zero expected returns is straightforward, empirical testing of this prediction is not. We now consider an econometrician who has observed a sample of returns from this economy and wants to test the $\mathbb{E}\, r_{OOS,t+1} = 0$ prediction of the trading strategy in (19). For this analysis, we focus on the case $s = t-1$. We now consider asymptotics where $N, J, P \to \infty$ with $J/N \to \psi$, as before, and $P/J \to \phi$, where $\phi > 0$ is a constant.

We obtain the following result for the asymptotic distribution of $r_{OOS,t+1}$:

Proposition 8. *If returns are generated according to* (8) *and* $s < t$, *then in the large* N, J, P *limit*

$$\frac{r_{OOS,t+1}}{\sqrt{\sum_{j=1}^{P} \zeta_{j,s}\zeta_{i,t}}} \xrightarrow{d} N(0,1)\,.$$

In principle, therefore, testing of the prediction $\mathbb{E}\, r_{OOS,t+1} = 0$ is straightforward, but in practice it is not. The econometrician could try to empirically approximate the asymptotic standard deviation in the denominator of the ratio in Proposition 8 from the time-series standard deviation of the r_{OOS} observations. As we suggested earlier, one could think of one period in this model as roughly a decade and one could imagine the econometrician observing intra-period r_{OOS} realizations that are sampled at higher frequency. Would a simple t-test, that compares the estimated mean of intra-period returns to the standard error calculated

from intra-period returns, provide a valid test of the $\mathbb{E}\,r_{OOS,t+1} = 0$ hypothesis?

Unfortunately, the answer is no. The reason is that the expectation in $\mathbb{E}\,r_{OOS,t+1} = 0$ is unconditional in the sense that it integrates over the distribution of $\boldsymbol{g}$. However, the econometrician only observes one sample for a given draw of $\boldsymbol{g}$. Nature draws $\boldsymbol{g}$ before any data is generated and the econometrician cannot re-run history for a different value of $\boldsymbol{g}$. Moreover, if the econometrician starts sampling the OOS strategy returns at the beginning of period $t+1$, the sample is also conditioned on the path of $\boldsymbol{e}_1, ..., \boldsymbol{e}_t$.

In this one sample available to the econometrician, the distribution of $r_{OOS,t+1}$ has mean $\mathbb{E}\left[r_{OOS,t+1}|\boldsymbol{g}, \boldsymbol{e}_1, ..., \boldsymbol{e}_t\right]$, which is generally not zero. In the same way, a sample variance of $r_{OOS,t+1}$ in this one sample would estimate $\mathrm{var}(r_{OOS,t+1}|\boldsymbol{g}, \boldsymbol{e}_1, ..., \boldsymbol{e}_t)$, not the unconditional variance under the square root in the denominator of the expression in Proposition 8. As a consequence, a t-statistic based on this sample mean and variance would not yield a correctly sized test of the Bayesian learning null hypothesis.

Heuristically, at least, there is a relatively simple solution. Since $\mathbb{E}\,r_{OOS,t+1} = 0$, we have $\mathrm{var}(r_{OOS,t+1}) = \mathbb{E}\,r_{OOS,t+1}^2$. Moreover, $\mathbb{E}\,r_{OOS,t+1}^2 = \mathbb{E}\,\mathbb{E}[r_{OOS,t+1}^2|\boldsymbol{g}, \boldsymbol{e}_1, ..., \boldsymbol{e}_t]$ and so $\mathrm{var}(r_{OOS,t+1}) = \mathbb{E}\,\mathbb{E}[r_{OOS,t+1}^2|\boldsymbol{g}, \boldsymbol{e}_1, ..., \boldsymbol{e}_t]$. Therefore, if the econometrician uses the mean of squared intra-period returns rather than the sample variance of returns as an estimate of $\mathrm{var}(r_{OOS,t+1})$, this estimate is on average, across samples with different $\boldsymbol{g}$ and $\boldsymbol{e}_1, ..., \boldsymbol{e}_t$, equal to the unconditional variance. The reason is that the mean of squared returns captures the deviations of $r_{OOS,t+1}$ from zero due to $\mathbb{E}\left[r_{OOS,t+1}|\boldsymbol{g}, \boldsymbol{e}_1, ..., \boldsymbol{e}_t\right] \neq 0$, while in a sample variance calculation these components are removed through demeaning.

III.E. *Out-of-sample moment conditions for risk premia estimation*

In our model so far, we have abstracted from risk premia or the possibility that investors' subjective beliefs may deviate from the prescriptions of Bayesian updating with objectively correct priors. Risk premia or belief distortions would induce additional predictable components in asset returns beyond those that we have (8). And an empiricist may want to

27

estimate to what extent characteristics $\boldsymbol{C}$ are associated with risk premia or belief distortions, while allowing, at the same time, for the possibility that investors may be learning about the predictive role of $\boldsymbol{C}$ for asset cash flows.

Suppose the characteristics are associated with a predictable component $\boldsymbol{X}\boldsymbol{\gamma}_x$ for some vector $\boldsymbol{\gamma}_x$ that represents risk premia (for brevity, we use the label "risk premia" from now on, but with the understanding that non-zero elements in $\boldsymbol{\gamma}_x$ could arise from belief distortions or frictions, too). In this case, adding this component to the returns in (8), we get

$$\boldsymbol{r}_{t+1} = \boldsymbol{X}\boldsymbol{\gamma}_x + \boldsymbol{X}(\boldsymbol{I} - \boldsymbol{\Gamma}_t)\boldsymbol{g} - \boldsymbol{X}\boldsymbol{\Gamma}_t(\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\bar{\boldsymbol{e}}_t + \boldsymbol{e}_{t+1}, \tag{20}$$

The econometrician wants to estimate to what extent the characteristics she observes, i.e., $\boldsymbol{C}$, are associated with risk premia. Projecting $\boldsymbol{X}$ on $\boldsymbol{C}$ we can decompose $\boldsymbol{X} = \boldsymbol{C}\boldsymbol{B} + \boldsymbol{E}$, where $\boldsymbol{E}'\boldsymbol{C} = \boldsymbol{0}$ and rewrite returns as

$$\boldsymbol{r}_{t+1} = \boldsymbol{C}\boldsymbol{\gamma} + \boldsymbol{E}\boldsymbol{\gamma}_x + \boldsymbol{X}(\boldsymbol{I} - \boldsymbol{\Gamma}_t)\boldsymbol{g} - \boldsymbol{X}\boldsymbol{\Gamma}_t(\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\bar{\boldsymbol{e}}_t + \boldsymbol{e}_{t+1} \tag{21}$$

where we defined $\boldsymbol{\gamma} = \boldsymbol{B}\boldsymbol{\gamma}_x$.

How can the econometrician estimate the risk premium component $\boldsymbol{C}\boldsymbol{\gamma}$? Under RE, with the third and fourth term in (21) absent, the econometrician could just regress $\boldsymbol{r}_{t+1}$ on $\boldsymbol{C}$ in an IS regression and estimate $\boldsymbol{C}\boldsymbol{\gamma}$ with the fitted value from this regression. But when investors are learning about $\boldsymbol{g}$, this does not work because the third and fourth term are not zero and they are correlated with $\boldsymbol{C}$. An OLS regression of $\boldsymbol{r}_{t+1}$ on $\boldsymbol{C}$ would be distorted by the presence of these terms. Along the same lines as in tests of the no-predictability null that we analyzed earlier, IS tests of hypotheses about $\boldsymbol{\gamma}$ would be distorted by the presence of the learning-induced components in returns.

Instead, the econometrician could use the moment conditions

$$\mathbb{E}\left[\left(\boldsymbol{r}_{t+1} - \boldsymbol{C}\boldsymbol{\gamma} - \boldsymbol{E}\boldsymbol{\gamma}_x\right)\left(\boldsymbol{C}\boldsymbol{b}_{s+1}\right)'\right] = 0 \quad \text{for } t \neq s. \tag{22}$$

That the expected values are zero follows from the fact that if we subtract $C\gamma + E\gamma_x$ from r_{t+1}, the remaining return components are the same as in (8) and so Proposition 7 applies. We can think of the return forecasts Cb_{s+1} here as instruments orthogonal to $r_{t+1} - C\gamma - E\gamma_x$. These moment conditions are not empirically implementable, though, because $E\gamma_x$ is unobservable to the econometrician.

We can solve this problem by focusing on a particular linear combination of moment conditions. For example, for the OOS return $r_{OOS,t+1} = w'_{OOS,s+1}r_{t+1}$, (22) implies

$$\mathbb{E}\left[r_{OOS,t+1} - \gamma'C'C\gamma\right] = 0 \quad \text{for } t \neq s. \tag{23}$$

where the $E\gamma_x$ component drops out because $C'E = 0$. We can solve for

$$\gamma'C'C\gamma = \mathbb{E}\left[r_{OOS,t+1}\right]. \tag{24}$$

Therefore, the econometrician could estimate the contribution of risk premia to return predictability, $\gamma'C'C\gamma$, by the sample analog of the expectation on the right-hand side.

Of course, statistical inference in this case runs into the same problem that we discussed for the OOS test in the previous subsection. Analogous to our earlier discussion, the econometrician could use squared intra-period realizations of $r_{OOS,t+1}$ to estimate $\text{var}(r_{OOS,t+1})$. In this case, with risk premia components in returns, the squared returns would also incorporate the squared risk premia. As a consequence, $\mathbb{E}\,r^2_{OOS,t+1} > \text{var}(r_{OOS,t+1})$. Standard errors based on these estimates of $\text{var}(r_{OOS,t+1})$ may therefore be somewhat conservative in the sense that they overstate the sampling variation on average.

III.F. Comparison with cross-validated penalized regression

Since the above estimation procedure amounts to looking for return components that are predictable OOS, one may wonder whether penalized regression with penalty choice based on cross-validation could be an alternative route towards a suitable estimator of quantities

like $\boldsymbol{\gamma}'\boldsymbol{C}'\boldsymbol{C}\boldsymbol{\gamma}$ that provide an assessment of the magnitude of risk premia. As we show now, there is, indeed, a close close connection between the estimator based on a sample analog of (24) and penalized regression with penalty choice based on cross-validation.

Penalized regression like ridge or lasso regression shrinks OLS coefficients $\boldsymbol{b}_t$ towards zero with some shrinkage matrix $\boldsymbol{\Phi}(\xi_t)$ that is a function of a penalty parameter ξ_t. Here we focus on ridge regression where the estimator of the regression coefficients is $(\boldsymbol{C}'\boldsymbol{C} + \xi_t\boldsymbol{I})^{-1}\boldsymbol{C}'\boldsymbol{r}_t$, which can be written as $\boldsymbol{\Phi}(\xi_t)\boldsymbol{b}_t$ with $\boldsymbol{\Phi}(\xi_t) = \left[\boldsymbol{I} + \xi_t(\boldsymbol{C}'\boldsymbol{C})^{-1}\right]^{-1}$. Cross-validation seeks to choose the penalty parameter that minimizes the OOS residual sum of squares, i.e.,

$$\xi_t^* = \underset{\xi_t}{\arg\min}\, \mathbb{E}\left[\{\boldsymbol{r}_{t+1} - \boldsymbol{C}\boldsymbol{\Phi}(\xi_{s+1})\boldsymbol{b}_{s+1}\}' \{\boldsymbol{r}_{t+1} - \boldsymbol{C}\boldsymbol{\Phi}(\xi_{s+1})\boldsymbol{b}_{s+1}\}\right], \tag{25}$$

where $s \neq t$ so that $\boldsymbol{r}_{t+1}$ is observed in a validation sample that does not overlap with the estimation sample used to estimate $\boldsymbol{b}_{s+1}$. We have $\partial\boldsymbol{\Phi}(\xi_t)/\partial\xi_t = -\boldsymbol{\Phi}(\xi_t)(\boldsymbol{C}'\boldsymbol{C})^{-1}\boldsymbol{\Phi}(\xi_t)$. Letting $\boldsymbol{\Phi}_t$ denote $\boldsymbol{\Phi}(\xi_t^*)$, the first-order condition of the problem (25) then yields

$$\mathbb{E}\left[\boldsymbol{b}'_{s+1}\boldsymbol{\Phi}_t(\boldsymbol{C}'\boldsymbol{C})^{-1}\boldsymbol{\Phi}_t\boldsymbol{C}'\boldsymbol{r}_{t+1}\right] = \mathbb{E}\left[\boldsymbol{b}'_{s+1}\boldsymbol{\Phi}_t\boldsymbol{C}'\boldsymbol{C}\boldsymbol{\Phi}_t(\boldsymbol{C}'\boldsymbol{C})^{-1}\boldsymbol{\Phi}_t\boldsymbol{b}_{s+1}\right]. \tag{26}$$

In the special case where $\boldsymbol{C}'\boldsymbol{C}$ is proportional to the identity matrix, $\frac{1}{N}\boldsymbol{C}'\boldsymbol{C} = \boldsymbol{I}$, the ridge regression shrinkage matrix takes the simple form $\boldsymbol{\Phi}(\xi_t) = \frac{1}{1+\xi_t N}\boldsymbol{I}$ and $\boldsymbol{b}_{s+1} = N\boldsymbol{C}'\boldsymbol{r}_{s+1}$. In this case, equation (26) simplifies to

$$\mathbb{E}\left[\boldsymbol{b}'_{s+1}\boldsymbol{C}'\boldsymbol{r}_{t+1}\right] = N\,\mathbb{E}\left[\frac{1}{1+\xi_t^* N}\boldsymbol{b}'_{s+1}\boldsymbol{b}_{s+1}\right]$$
$$= \mathbb{E}\left[\frac{1}{1+\xi_t^* N}\boldsymbol{b}'_{s+1}\boldsymbol{C}'\boldsymbol{r}_{s+1}\right]. \tag{27}$$

The term in brackets in the last line is the IS return on a portfolio where assets are weighted by the predicted return from the ridge regression based on the cross-validated penalty. We denote this portfolio return with $r_{CV,s+1}$. The term in brackets on the left-hand side is the

OOS portfolio return $r_{OOS,t+1}$, as defined earlier in (19). We therefore have

$$\mathbb{E}[r_{OOS,t+1}] = \mathbb{E}[r_{CV,s+1}]. \tag{28}$$

Therefore, in this special case, one can obtain the contribution of risk premia to cross-sectional variation of expected returns, $\boldsymbol{\gamma}'\boldsymbol{C}'\boldsymbol{C}\boldsymbol{\gamma}$ either from $\mathbb{E}[r_{OOS,t+1}]$, via the OOS moment condition (24), or, equivalently, from $\mathbb{E}[r_{CV,t+1}]$.

In the general case where $\boldsymbol{C}'\boldsymbol{C}$ is not proportional to the identity matrix, the relationship $\mathbb{E}[r_{OOS,t+1}]$ and $\mathbb{E}[r_{CV,t+1}]$ is more complicated, but the example in the special case above demonstrates that there is a close connection between cross-validated penalized regression and OOS moment conditions. In our empirical application below, we will compare both approaches.

IV. FINITE-SAMPLE ANALYSIS: SIMULATIONS

In this section, we report the results of finite-sample simulations. The high-dimensional asymptotics, with $N, J \to \infty$ and $J/N \to \psi$, are intended to provide an approximation for the properties of return predictability tests in the case where J is not small relative to N (just as conventional large N-fixed J asymptotics provide an approximation for the small J/N case). Finite-sample simulations provide some insight to what extent the asymptotic results provide a good approximation in a setting with realistic J and N. We set $N = 1000$ and let J vary from 1 to close to 1000. We draw the elements of $\boldsymbol{X}$ from a standard Normal distribution.

For this purpose of this numerical analysis, we also need to set the parameter θ that pins down the share of predictable variation in cash-flow growth through $\boldsymbol{\Sigma}_g = \frac{\theta}{J}\boldsymbol{I}$. What matters here is not the total level of cash-flow variance but rather the share that is predictable. For this reason, we normalize, as before, $\boldsymbol{\Sigma}_e = \boldsymbol{I}$. We then look for a value of θ that yields a plausible amount of predictable variation in cash-flow growth relative to this normalized

residual variance.

Based on our data-generating process for cash-flow growth in (1), annualized growth rates over a horizon of T periods are

$$\frac{1}{T} \sum_{t=1}^{T} \Delta \boldsymbol{y}_t = \boldsymbol{X}\boldsymbol{g} + \frac{1}{T} \sum_{t=1}^{T} \boldsymbol{e}_t.$$

We now evaluate the share that is predictable given knowledge of $\boldsymbol{g}$. This is an upper bound on the share that investors learning about $\boldsymbol{g}$ may be able to predict. The annualized variance of the predictable component (first term on the RHS) is constant w.r.t. T, while the variance of the residual component (second term on the RHS) shrinks at the rate $1/T$. As we indicated earlier, we think of one period in the model as representing roughly one decade. In this case, at a horizon of one decade, i.e. $T = 1$, the variance of the forecastable component is

$$\frac{1}{N} \mathbb{E}\left[\boldsymbol{g}'\boldsymbol{X}'\boldsymbol{X}\boldsymbol{g}\right] = \frac{1}{N} \operatorname{tr}\left(\boldsymbol{X}'\boldsymbol{X}\right) \frac{\theta}{J} \approx \theta$$

and so the ratio of forecastable to residual variance also equals θ.

We can do a back-of-the-envelope calculation to compare this to the growth rate evidence for various revenue and profit measures in Chan, Karceski, and Lakonishok (2003). When stocks are sorted based on IBES analysts forecasts (their Table IX), Chan et al. find an inter-decile spread of slightly around 10 percentage points (pp) for annualized growth rates over the next one, three, and five years. Extrapolating to 10 years, we would have 10 pp also at a 10-year horizon.[8] When they sort stocks instead based on their ex-post realized annualized growth rates over a one-year horizon, they find an inter-decile range of around 50 pp. Assuming normal distributions, these estimates imply a ratio of about 0.04 of forecastable to residual variance at a one-year horizon. An IID data-generating process for cash flows, as in our model, would imply that the residual variance shrinks at rate $1/T$ and hence the ratio

8. For this analysis, they don't report percentiles at sufficient detail to calculate the inter-decile spread, but they report means for quintile bins, and the spread between top and bottom quintile bin mean should correspond, approximately to the inter-decile range.

of forecastable to residual variance at a ten-year horizon is 0.4. This share of forecastable variance represents a lower bound as analysts' can predict only a less than full share of the total potentially predictable variance. For this reason, it seems reasonable to set the ratio of maximally predictable to residual variance somewhat higher than 0.4 in our simulations. Accordingly, we set $\theta = 1$.

We now simulate cash flows and, based on investors' Bayesian updating and pricing, the returns on the $N = 1000$ assets. We then consider an econometrician that samples these returns ex post and runs regressions of $\boldsymbol{r}_{T+1}$ on $\boldsymbol{X}$ after investors have learned about $\boldsymbol{g}$ for T periods. Figure IIIa presents the (in-sample) adjusted R^2 from these regressions. As J increases towards N, the adjusted R^2 also increases, and hence returns become more predictable in sample. If investors have learned for more periods, return predicability gets weaker.

Figure IIIb looks at the properties of a standard Wald test of the no-predictability null hypothesis, testing whether the coefficients on the J predictor variables are jointly equal to zero. The plot shows the rejection probabilities (actual size) from a χ^2-test based on the null distribution in (10). The dotted line shows the nominal size of 5% that the test would have, asymptotically, if investors priced assets under rational expectations with perfect knowledge of $\boldsymbol{g}$. The figure shows that the actual rejection probabilities can be far higher than 5%. The rejection probabilities go to one as J grows towards N. The increase is slower if investors have learned for more periods, but even with $T = 4$, the rejection probability exceeds 90% when $J > N/2$. Thus, the simulations confirm that the asymptotic result of rejection probabilities going to one is, indeed, a good approximation for the large J/N case with finite N and J.

V. Sparsity

So far we have assumed a setting in which shrinkage of coefficients towards zero is the optimal way for investors to deal with the large number of cash flow predictors. But investors do not impose sparsity—i.e., some coefficients of exactly zero—on the forecasting model. The

(A) Adjusted R^2

(B) Rejection probability of no-return-predictability null

FIGURE III

In-sample return predictability tests

Based on cross-sectional regression with $N = 1000$ assets and J predictor variables, predicting the last return in a sample of size $T + 1$, and where investors have learned about $\boldsymbol{g}$ from sample of size T. The test in panel (b) is a joint χ^2-test using all J predictors. It has an asymptotic 5% rejection probability under the rational expectations null hypothesis (where investors know $\boldsymbol{g}$). Actual size shows the rejection probability when this test is applied in setting where Bayesian investors with an objectively correct informative prior estimate $\boldsymbol{g}$.

absence of sparsity was a consequence of the normal prior distribution of $\boldsymbol{g}$. If, instead, (i) investors' prior is that the elements of $\boldsymbol{g}$ are drawn from a Laplace distribution; and (ii) investors price assets based on the mode rather than the mean of the posterior distribution (i.e., a maximum-a-posteriori, or MAP estimator), then asset prices reflect sparse cash flow forecasts in which some columns of $\boldsymbol{X}$ are multiplied with coefficients of zero. Their forecasts can then be represented as the fitted values from a Lasso regression (Tibshirani 1996).

That investors use the posterior mode in pricing is a deviation from the fully Bayesian framework. Our simulations will shed light on how much of a deviation this is in terms of how much additional return predictability results from it.

With a Laplace prior, elements g_j are IID with the distribution

$$f(g_j) = \frac{1}{2b} \exp\left(-\frac{|g_j|}{b}\right).$$

The variance is $2b^2$. To keep the variance the same as in the normal prior case, we set $2b^2 = \frac{\theta}{J}$. This Laplace distribution not only represents investors' prior, but we now also draw the elements of $\boldsymbol{g}$ in our simulations from this distribution so the prior is again objectively correct, as in the normal prior case we considered earlier.

Figure IV shows that the results in the Laplace prior case are extremely similar to those in Figure III for the normal prior case. In terms of how in-sample predictability strengthens with increasing J/N, it does not make much difference whether investors shrink prediction model coefficients with or without sparsity. For the sake of brevity, most of our results in the paper therefore focus on the normal prior case.

VI.   Excess shrinkage or sparsity

In our analysis up to this point, shrinkage or sparsity was purely due to prior knowledge reflected in investors' prior beliefs. Aside from such statistical optimality considerations, there could be other reasons for investors to shrink coefficients and impose sparsity on their

(A) Adjusted R^2

(B) Rejection probability of no-return-predictability null

FIGURE IV

Lasso: In-sample return predictability tests

Based on cross-sectional regression with $N = 1000$ assets and J predictor variables, predicting the last return in a sample of size $T + 1$, and where investors have learned about $\boldsymbol{g}$ from sample of size T. The test in panel (b) is a joint χ^2-test using all J predictors. It has an asymptotic 5% rejection probability under the rational expectations null hypothesis (where investors know $\boldsymbol{g}$). Actual size shows the rejection probability when this test is applied in setting where investors estimate $\boldsymbol{g}$ (which is drawn from a Laplace distribution) with Lasso.

forecasting models. For example, if variables are costly to observe, investors might prefer to discard a variable if it only offers a weak signal about cash flows. Relative to the frictionless Bayesian benchmark with objectively correct prior, such a model would be excessively sparse, but the reduction in forecast performance may be justified by the cost savings from model sparsity. Relatedly, Sims (2003) and Gabaix (2014) show that shrinkage or sparsity can be used to represent boundedly rational decision-making if attention to a variable generates an actual or psychological cost (that shrinkage or sparsity helps avoid).

Such variants of the model with excessive shrinkage can still be mapped into a Bayesian updating scheme, but the prior beliefs are concentrated more tightly around zero than in the case with objectively correct prior (where the prior distribution agrees with distribution that we draw g from in generating the data). For this reason, we label the benchmark case with objectively correct prior as DGP-consistent shrinkage or sparsity.

Figure V shows the consequences of excessive shrinkage or sparsity for out-of-sample return predictability. In all cases shown in the figure, the cash-flow data is generated, as before, with $\theta = 1$. However, investors' prior beliefs are now based on a value of θ that is different. In the excessive shrinkage and sparsity cases, we let investors form beliefs based on $\theta = 0.5$. This means that they have a prior distribution for the elements of g that is more tightly concentrated around zero than the actual distribution of g that generates the data. For comparison, we also consider a case where shrinkage is insufficient. In this case, investors assume $\theta = 2$. This can be interpreted as investors having a lack of confidence, and hence excessively wide dispersion, in their prior beliefs about g. In all cases, we show results for $T = 4$, which means that investors have learned for four periods up to the beginning of the period in which we measure the return on the out-of-sample portfolio.

The figure shows the out-of-sample return of a portfolio that weights assets by their predicted expected return as in (19). Panel (a) shows the results in the normal prior/ridge regression case. As expected from Proposition 7, the average OOS return in the DGP-consistent case is zero. In contrast, when shrinkage is excessive, investors end up downweighting too

(B) Lasso

FIGURE V

Out-of-sample portfolio returns when investors apply excess shrinkage or sparsity

Based on cross-sectional regression with $N = 1000$ assets and J predictor variables, predicting the last return in a sample of size $T + 1$, and where investors have learned about $\boldsymbol{g}$ from sample of size $T = 4$. Cash-flow data are always generated with $\theta = 1$. In the DGP-consistent prior case, investors' prior is based on $\theta = 1$. In the excess shrinkage (or sparsity) case, investors' prior is based on $\theta = 0.5$, in the insufficient shrinkage (or sparsity) case, investors' prior is based on $\theta = 2$. The DGP in all cases always features $\theta = 1$.

much pieces of information in $\boldsymbol{X}$ that predict cash flows. As a result, an econometrician sampling returns from this economy is able to forecast returns OOS. And the effect gets stronger with higher J. Forming a portfolio that weights assets based on their estimated expected returns from predictive regressions on data up to time T earns a predictably positive return in period $T + 1$.

Insufficient shrinkage also results in an OOS average return that differs zero, but the sign is negative. This means that assets that would be predicted to have positive expected returns, based on the econometrician's predictive regression estimates from data up to time T, actually end up having negative returns in $T + 1$ and vice versa. With insufficient shrinkage, the component $-\boldsymbol{X}\boldsymbol{\Gamma}_t \left(\boldsymbol{X}'\boldsymbol{X}\right)^{-1} \boldsymbol{X}'\bar{\boldsymbol{e}}_t$ in the expression for $\boldsymbol{r}_{t+1}$ in (8) plays a bigger role in than under DGP-consistent shrinkage. As a consequence, its negative covariance with the estimation error component $\boldsymbol{X} \left(\boldsymbol{X}'\boldsymbol{X}\right)^{-1} \boldsymbol{X}'\boldsymbol{e}_t$ of the fitted value $\boldsymbol{X}\boldsymbol{h}_t$ from the predictive regression in (9) dominates, which means the forecasts based on $\boldsymbol{X}\boldsymbol{h}_t$ tend to have the wrong sign out-of-sample.

As panel (b) shows, the results in the Laplace prior/lasso regression case are similar. One difference is that the average OOS return in the DGP-consistent case is somewhat negative instead of zero. This is a consequence of the assumption that investors in this case use, by assumption, the mode of the posterior distribution of $\boldsymbol{g}$ (which induces sparsity) to price assets rather than the posterior mean (which would imply a forecasting model that is not sparse).

The take-away from all this is that OOS return predictability evidence can help shed light on whether investors apply excessive shrinkage or sparsity in their cash-flow forecasting models. For example, if actual or psychological costs of complex models induce additional sparsity, this should show up in the data as OOS-predictable returns. Similarly, if investors' prior distribution of $\boldsymbol{g}$ assumed a distribution of coefficients that was too tightly concentrated around zero compared with the true distribution that generated the data, this would show

up as OOS predictability.[9]

This analysis also provides a perspective on the likely effects of technological progress in data construction and data analysis on return predictability observed in empirical analyses. Many studies of the cross-section of stock returns use data that goes back to time periods in which data availability and analysis was much more constrained than it is today. A researcher today can construct many variables (say, through automated textual analysis of corporate filings) that were inaccessible to investors until not very long ago. In this sense, the forecasting models that investors used at the time when they priced stocks several decades ago may have been excessively sparse relative to the model that an empirical researcher could work with today. It is to be expected, therefore, that a researcher today can construct variables, or use combinations of large numbers of variables, that predict returns in the earlier years of stock return data sets, even in (pseudo-) OOS tests in which the researcher re-constructs investors' learning process, but without taking into account the additional model complexity constraints that investors faced in real time.

VII. Empirical Application: Predicting Stock Returns with Past Returns

To illustrate how our model provides an interpretation of discrepancies between IS and OOS stock return predictability, we look at an empirical application. The key prediction is that investor learning in a high-dimensional setting should lead to a substantial wedge between IS and OOS predictability. In the model, OOS predictability is zero, but the model abstracted from risk premia and (except for our analysis of excess shrinkage) from behavioral biases and bounded rationality that could induce OOS predictability. For this reason, in our empirical

---

9. Note that this would not mean that investors were irrational. Rational Bayesian reasoning does not require that prior beliefs are consistent with the true distribution of g, which would be unknown to investors. Existence of OOS return predictability evidence would be consistent not only with the bounded rationality explanation of excess shrinkage or sparsity, but also a tight-prior explanation. Changes over time in OOS predictability might allow to disentangle the two.

analysis we focus on the wedge between IS and OOS predictability, and its evolution over time, rather than on the absolute level of predictability.

For this exercise, we seek a large set of predictor variables that were, at least in principle, consistently available to investors over a long period of time. Many predictors based on accounting variables do not satisfy this criterion because they become available in Compustat data only in later decades. Furthermore, to stay close to our setting in the model where the econometrician studies a given set of predictor variables only once, without specification search and multiple testing, we do not want a set of predictors that may already be the product of past data mining efforts by earlier researchers. For example, the set of published predictors in the academic literature likely includes some that have been data-mined ex-post. To address both concerns, we use each stock's full price history to generate our predictor variables. More precisely we use the monthly returns and monthly squared returns in the past 120 months as predictors next month's stock return. Including squared returns allows the relationship between past returns and future returns to be nonlinear.

This price history was, at least in principle, available to investors even in the early parts of the sample. This eliminates the possibility that return predictability could show up ex post simply because a variable that we can construct today was not available at all to investors in real time when they priced stocks. Of course, the fact that the price history was available in principle does not necessarily mean that investors throughout the sample always had the ability to integrate all of these variables into their forecasting models. If they could not, we might find OOS return predictability in parts of the sample where technological constraints may have prevented them from doing so—consistent with our excess shrinkage results in the previous section.

Focusing on price-history-based predictors, without pre-selecting particular subsets of those based on earlier evidence of predictability, allows us to side-step, for the most part, the influence of earlier researchers' data mining. The only potential remaining problem is that our choice of considering price-history-based predictors as a class could be influenced by

existing evidence that subsets of these seem to have predictive power (e.g., momentum, long-run reversal). On the other hand, the class of price-history-based predictors would surely be a natural candidate, even in absence of any existing evidence, given that weak-form efficiency is the most basic market efficiency notion (Fama 1970).

A drawback of price-history-based predictors is that they don't perfectly map into our model. In our theoretical analysis, we worked with an exogenous cash-flow predictor matrix X. In contrast, past returns are an equilibrium outcome. One could, however, imagine an extension of the model in which cash-flow growth shocks could have persistent components at various lags. In this case, investors' set of potential cash-flow predictors would include the history of past cash-flow growth shocks and lagged returns would be correlated with these. In this sense, the distance from our model is not that big. In any case, the purpose of the empirical analysis is not to provide a formal test of the model but rather to illustrate in a simple setting with a large number of predictors the wedge between IS and OOS predictability.

We use all U.S. stocks in the CRSP database except small stocks that have market capitalization below the 20th NYSE percentile and price lower than one dollar at the end of month $t - 1$. To avoid picking up microstructure related issues, we skip the most recent month in our construction of the set of predictor variables. Thus, we use simple and squared returns in months $t - 2$ to $t - 120$, i.e., a total of 238 predictor variables, to predict returns in month t in a panel regression. We demean the dependent variable and all explanatory variables month by month to focus purely on cross-sectional variation. In addition, we cross-sectionally standardize all predictor variables to unit standard deviation each month. We weight the observations each month such that the panel regression gives equal weight to each month in the sample.

As a first step, to demonstrate that a regression with shrinkage delivers meaningful estimates with such a large number of predictor variables, we examine an in-sample panel ridge regression to predict monthly returns from the beginning of 1971 until end of June 2019. We show that the ridge regression automatically recovers many prominent predictability patterns

that have been documented in the existing literature for roughly this sample period or parts of it. We pick the penalty hyperparameter that determines the strength of shrinkage through leave-one-year-out cross-validation.[10]

Panel (a) in Figure VI presents the regression coefficients for each of the 119 simple return explanatory variables. It shows that a single ridge regression recovers several major anomalies related to past returns: the positive coefficients up to lags of 12 months capture momentum as in Jegadeesh and Titman (1993); the plot also shows that continuation of recent returns is concentrated in lags 7 to 12, as pointed out in Novy-Marx (2012); the mostly negative coefficients for lags beyond 12 months reflect long-term reversals as in DeBondt and Thaler (1985); the positive coefficients at lags equal to multiples of 12 reflect the momentum seasonality reported by Heston and Sadka (2008). Panel (b) reports the regression coefficients for the 119 lagged squared returns. At shorter lags, there is no clear pattern. But at longer lags beyond lag 50, the coefficients are predominantly positive, indicating a positive association of long-run lagged individual stock return volatility and future returns. The monthly in-sample R^2 in this ridge regression is 0.23%. With market-adjusted annualized volatility of roughly 50% for the dependent variable, this R^2 implies a cross-sectional annualized standard deviation of fitted in-sample expected returns of approximately 8%, which is quite big.

However, in terms of OOS predictability, the picture looks very different. To construct OOS portfolio returns, we run rolling OLS regressions in 20-year estimation windows. We use the OLS estimates $\boldsymbol{b}_t$ of the window ending in month t to forecast returns in month $t+1$. Given the 20-year estimation window and the up to 10-year lag of the predictors, the first month in which we have a prediction is January 1956, 30 years after the start of the CRSP database. We use the estimates $\boldsymbol{b}_t$ to calculate the OOS portfolio return in month $t+1$ with weights proportional to predicted returns as in (19) that goes long in stocks with positive

---

10. We compute the estimates using all but one year of the sample, we calculate the implied predicted returns in the year left out of the estimation, and we record the resulting R^2 in the left-out year. We then repeat with a different left-out year and again record the R^2 in the left-out years and repeat until each year of the sample has been left out once. At the end, we average the R^2 across all left-out years and we search for a penalty value that maximizes this cross-validated R^2.

(A) Coefficients for past returns

(B) Coefficients for past squared returns

FIGURE VI
Ridge Regression Coefficient Estimates

44

predicted returns and short in stocks with negative predicted returns.[11] To get returns with interpretable magnitudes, we divide the weights of the OOS portfolio $\boldsymbol{w}_{OOS,t} = \boldsymbol{C}\boldsymbol{b}_t$ in (19) by the sum of absolute values of the elements of $\boldsymbol{w}_{OOS,t}$. In this way, we obtain a portfolio that is \$0.50 long and \$0.50 short every month.

Figure VII plots the the OOS portfolio return averaged over 10-year moving windows. For comparison, the figure also shows the IS returns obtained by weighting returns in-sample with the fitted values from the cross-validated ridge regression and an OLS regression, with portfolio weights in both cases again scaled to be \$0.50 long and \$0.50 short every month. We average the time series of these monthly returns over 10-year moving windows.

As the figure shows, the in-sample portfolio return is very stable across time and reliably above zero. In contrast, the out-of-sample return is much lower and frequently close to zero or below. The shaded area of the figure shows two-standard error bands obtained from the mean of squared monthly portfolio returns within the 10-year moving windows, in line with our discussion following Proposition 8 on estimation of $\text{var}(r_{OOS,t+1})$. Towards the end of the sample, in the 10-year windows ending between 2014 and 2019, the OOS portfolio return has been less than two standard errors above zero.

Another interesting result in this figure is that r_{CV}, the IS return based on the cross-validated ridge regression, is close to r_{OOS}, as suggested by our earlier analysis and (28).[12] This confirms that an IS portfolio return based on penalized regression with cross-validated penalty can provide good estimates of the OOS portfolio return.

Overall, these results illustrate the substantial wedge between IS and OOS return predictability. As we showed in (24), the average OOS portfolio return can be interpreted as an estimate of $\boldsymbol{\gamma}'\boldsymbol{C}'\boldsymbol{C}\boldsymbol{\gamma}$, the contribution of risk premia (or return premia induced by behavioral

---

11. Since the predictors in our empirical application are close to having $\boldsymbol{C}'\boldsymbol{C}$ proportional to the identity matrix, ridge regression estimates are close to a scalar shrinkage factor times the OLS estimates. As a consequence, if we instead used ridge regression estimates to form the OOS portfolio , the shrinkage factor cancels when we scale by the sum of absolute values of the raw portfolio weights. Portfolio weights based on ridge regression would therefore be almost identical to the OLS-based portfolio weights.

12. To maintain the equality between $\mathbb{E}\, r_{OOS}$ and $\mathbb{E}\, r_{CV}$ as in (28), we scale the portfolio weights of the CV portfolio by the same scaling factor (sum of absolute values of the elements of $\boldsymbol{w}_{OOS,t}$) that we use for the portfolio weights of the OOS portfolio.
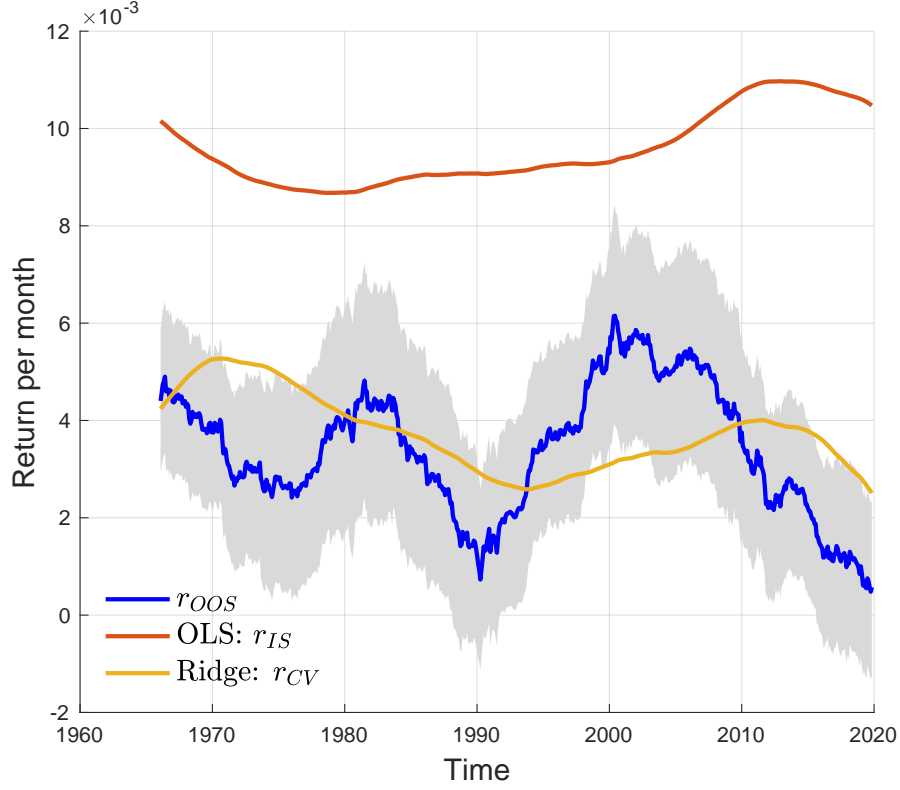
FIGURE VII

Rolling estimates of the risk premium/mispricing component of returns

The risk premium component $\boldsymbol{\gamma}'\boldsymbol{C}'\boldsymbol{C}\boldsymbol{\gamma}$ is estimated by the OOS portfolio return $r_{OOS,t+1} = \boldsymbol{r}'_{t+1}\boldsymbol{C}\boldsymbol{b}_t$. We obtain the OLS estimates $\boldsymbol{b}_t$ in backwards 20-year moving window up to month t and use them to construct weights applied to month $t + 1$ returns. The blue line in the figure shows $r_{OOS,t+1}$ averaged in 10-year moving windows. For comparison, the figure also shows the IS return based on cross-validated ridge regression estimates, r_{CV}, and OLS estimates, r_{IS}.

biases or frictions) to cross-sectional variation in expected returns. In a rational expectations model in which investors know the parameters of the cash flow process, the expected IS and OOS portfolio returns would both be equal to $\gamma'C'C\gamma$. But if investors learn about the cash flow process parameters over time, only the expected OOS portfolio return is equal to $\gamma'C'C\gamma$ while the IS portfolio return is distorted by learning-induced components that are not predictable OOS. The big wedge between IS and OOS shown in this figure suggest that the learning case is relevant. This underscores the message from our model that IS cross-sectional return predictability evidence is not a good motivation for seeking risk-based or behavioral economic explanations. Much of the IS predictability here does not carry over into OOS predictability and hence does not reflect risk premia demanded by investors ex-ante or persistent belief distortions.

There are a number of concerns one might have about this interpretation. First, wouldn't it be possible to pick, ex post, a much smaller number of lags of simple and/or squared returns that happen to do better, in-sample and in the (pseudo-) OOS tests from the rolling regressions? This may well be true. But what would be the ex-ante justification to pick those specific lags? Why not others? Pursuing this avenue would inevitably introduce data-snooping and multiple-testing concerns that cloud the interpretation of the evidence. As an example, consider that the literature on momentum has shifted from emphasizing momentum based on returns in months $t - 2$ to $t - 12$, as in Jegadeesh and Titman (1993), towards highlighting returns in months $t - 7$ to $t - 12$ as those that really matter for predicting returns (Novy-Marx 2012). It's not clear that one should seek deep economic reasons for in-sample predictability results that reflect such ex-post data-driven specification changes. Simply including all lags up to a certain point and and letting the shrinkage take care of preventing overfitting minimizes this data-snooping problem.

Second, OOS tests may have low power to detect return predictability. This point has been made by Inoue and Kilian (2005), Campbell and Thompson (2008), and Hansen and Timmermann (2015) in a time-series setting. However, their arguments are based on a ratio-

nal expectations framework in which investors know the parameters of the data-generating process and only the econometrician faces the problem of recovering its parameters from observed data. In their case, IS and OOS methods test the same economic hypothesis. But if investors are learning about parameters, especially in settings where the number of potentially relevant predictor variables is huge, the situation is fundamentally different. There is no fixed-parameter model that generates each period's returns. Instead, the properties of the data are evolving over time as investors learn. As a consequence, IS and OOS methods test different economic hypotheses. IS predictability tests basically lose their economic meaning they have in a rational expectations setting because they cannot discriminate between predictability induced by learning and predictability induced by risk premia or behavioral biases. Only OOS tests can do so. For this reason, even if OOS tests have low power, IS tests are simply not a viable alternative method because they test a different hypothesis without clear economic interpretation.

VIII. Conclusion

Our analysis provides a new perspective on markets in which decision-makers face high-dimensional prediction problems. Learning how to translate observed predictor variables into forecasts is hard when the number of predictors is comparable in size to the number of observations. To an econometrician studying these forecasts ex post, or the equilibrium prices that reflect these forecasts, the forecast errors look predictable, but they are not predictable to the decision-maker in real time. We developed this analysis in a cross-sectional asset pricing application, but the issue may be relevant more broadly in settings in which large numbers of variables can be relevant for forecasting.

In the cross-sectional asset pricing setting, in-sample tests of return predictability lose their economic meaning when investors are faced with a large number of potential predictors of asset cash flows. The usual economic interpretation that in-sample predictable returns represent priced risks or the effects of investors' behavioral biases does not apply in this

48

case. This is not a statistical problem with the sampling properties of the econometrician's predictability tests. Instead, it is a problem with the null hypothesis in these tests. As investors' learning problem becomes harder with increasing dimensionality of the set of potential predictors, the true properties of equilibrium prices change. Absence of risk premia and behavioral biases does not imply that the usual null hypothesis of a return-generating process with unpredictable increments applies. Investors' learning of the cash-flow forecasting model parameters leaves in-sample predictable components in returns that represent investors' real-time estimation error and the shrinkage they optimally apply to reduce this estimation error in a high-dimensional setting. This is true even though we kept investors' learning problem very simple: the potentially predictable component of future cash-flow growth is linear in predictors and investors know this linear functional form. If investors also had to entertain that the functional form could be nonlinear, this would further magnify the dimensionality of the prediction problem they face.

In contrast to in-sample tests, out-of-sample tests retain their economic meaning in the high-dimensional case. Our argument in favor of out-of-sample tests is different from those usually discussed in the econometrics literature. The usual case for out-of-sample tests motivates them as remedies against distortions of the sampling properties of in-sample tests or against data mining. As Inoue and Kilian (2005), Campbell and Thompson (2008), Cochrane (2008), and Hansen and Timmermann (2015) have pointed out, these arguments in favor of out-of-sample testing are questionable in settings where the null hypothesis is a population model with truly unpredictable returns. Our point is that when investors face a high-dimensional forecasting problem, this is not an economically interesting null hypothesis. Absence of risk premia and behavioral biases implies absence of out-of-sample predictability, but not absence of in-sample predictability. Since investors are arguably facing a high-dimensional prediction problem in the real world, researchers should give more emphasis to out-of-sample testing.

Our results offer a novel interpretation of the fact that in-sample return predictability

tests in the literature have produced hundreds of variables that appear to predict returns in the cross section. As the number of predictor variables that are available to researchers and investors has grown enormously, it is to be expected, even with Bayesian investors, that equilibrium price changes increasingly look predictable when viewed through the hindsight-tinted lenses of an econometrician running in-sample regressions. Our results show that many of these predictor variables show up as in-sample statistically significant cross-sectional return predictors. For this reason, it not clear that one should look for risk-based explanations or behavioral explanations for their in-sample predictive power.

A number of extensions of our work could be interesting. Our setting is a purely cross-sectional one with firm characteristics that are constant over time. But a similar learning problem also exists in the time dimension, e.g., at the aggregate stock market level. A huge number of macro variables could, jointly, be relevant for predicting aggregate stock market fundamentals. Furthermore, to keep the model simple and transparent, we have focused on learning about exogenous fundamentals with homogeneous investors. It would be interesting to extend this to a setting with heterogeneous investors, where investors observe a large number of exogenous and endogenous signals from which they can extract information not only about asset fundamentals, but also about the trading behavior of other investors.

References

Al-Najjar, Nabil I, 2009, "Decision Makers as Statisticians: Diversity, Ambiguity, and Learning," *Econometrica* 77, 1371–1401.

Anatolyev, Stanislav, 2012, "Inference in Regression Models with Many Regressors," *Journal of Econometrics* 170, 368–382.

Aragones, Enriqueta, Itzhak Gilboa, Andrew Postlewaite, and David Schmeidler, 2005, "Fact-Free Learning," *American Economic Review* 95, 1355–1368.

Bai, Z. D., and Y. Q. Yin, 1993, "Limit of the Smallest Eigenvalue of a Large Dimensional Sample Covariance Matrix," *Annals of Probability* 21, 1275–1294.

Bryzgalova, Svetlana, Markus Pelger, and Jason Zhu, 2019, "Forest Through the Trees: Building Cross-Sections of Stock Returns," Working paper, Stanford University.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello, 2018, "Artificial intelligence, Algorithmic Pricing and Collusion," Working paper, CEPR.

Campbell, John Y., and Samuel B. Thompson, 2008, "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?," *Review of Financial Studies* 21, 1509–1531.

Chan, Louis KC, Jason Karceski, and Josef Lakonishok, 2003, "The Level and Persistence of Growth Rates," *Journal of Finance* 58, 643–684.

Chordia, Tarun, Amit Goyal, and Alessio Saretto, 2019, "Anomalies and False Rejections," *Review of Financial Studies* 33, 2134–2179.

Cochrane, John H., 2008, "The Dog That Did Not Bark: A Defense of Return Predictability," *Review of Financial Studies* 21, 1533–1575.

Cochrane, John H., 2011, "Presidential Address: Discount Rates," *Journal of Finance* 66, 1047–1108.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A Lochstoer, 2016, "Parameter Learning in General Equilibrium: The Asset Pricing Implications," *American Economic Review* 106, 664–698.

DeBondt, Werner F.M., and Richard Thaler, 1985, "Does the Stock Market Overreact?," *Journal of Finance* 40, 793–805.

Dobriban, Edgar, and Stefan Wager, 2018, "High-Dimensional Asymptotics of Prediction: Ridge Regression and Classification," *Annals of Statistics* 46, 247–279.

Fama, Eugene, 1970, "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance* 25, 383–417.

Feng, Guanhao, Stefano Giglio, and Dacheng Xiu, 2020, "Taming the Factor Zoo: A Test of New Factors," *Journal of Finance* 75, 1327–1370.

Gabaix, Xavier, 2014, "A Sparsity-Based Model of Bounded Rationality," *Quarterly Journal of Economics* 129, 1661–1710.

Hansen, Peter Reinhard, and Allan Timmermann, 2015, "Equivalence between Out-of-Sample Forecast Comparisons and Wald Statistics," *Econometrica* 83, 2485–2505.

Harvey, Campbell R, Yan Liu, and Heqing Zhu, 2016, ". . . and the Cross-Section of Expected Returns," *Review of Financial Studies* 29, 5–68.

Heston, Steven L, and Ronnie Sadka, 2008, "Seasonality in the Cross-Section of Stock Returns," *Journal of Financial Economics* 87, 418–445.

Inoue, Atsushi, and Lutz Kilian, 2005, "In-Sample or Out-of-Sample Tests of Predictability: Which One Should We Use?," *Econometric Reviews* 23, 371–402.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, "Returns to Buying Winners and Selling Losers: Implications for Market Efficiency," *Journal of Finance* 48, 65–91.

Klein, Timo, 2019, "Autonomous Algorithmic Collusion: Q-Learning Under Sequential Pricing," Working paper 2018-15, University of Amsterdam.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2020, "Shrinking the Cross-Section," *Journal of Financial Economics* 135, 271–292.

Lewellen, Jonathan, and Jay Shanken, 2002, "Learning, Asset-Pricing Tests and Market Efficiency," *Journal of Finance* 57, 1113–1145.

Lindley, Dennis V, and Adrian FM Smith, 1972, "Bayes Estimates for the Linear Model," *Journal of the Royal Statistical Society: Series B (Methodological)* 34, 1–18.

Linnainmaa, Juhani T, and Michael R Roberts, 2018, "The History of the Cross-Section of Stock Returns," *Review of Financial Studies* 31, 2606–2649.

Lo, Andrew W., and A. Craig MacKinlay, 1990, "Data-Snooping Biases in Tests of Financial Asset Pricing Models," *Review of Financial Studies* 3, 431–467.

McLean, David R., and Jeffrey Pontiff, 2016, "Does Academic Research Destroy Stock Return Predictability?," *Journal of Finance* 71, 5–32.

Molavi, Pooya, Alireza Tahbaz-Salehi, and Andrea Vedolin, 2020, "Asset Pricing with Misspecified Models," Working paper, Boston University.

Novy-Marx, Robert, 2012, "Is Momentum Really Momentum?," *Journal of Financial Economics* 103, 429–453.

Shalev-Shwartz, Shai, and Shai Ben-David. *Understanding machine learning: From theory to algorithms* (Cambridge University Press, New York, NY, 2014).

Sims, Christopher A., 2003, "Implications of Rational Inattention," *Journal of Monetary Economics* 50, 665 – 690.

Tibshirani, Robert, 1996, "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society: Series B (Methodological)* 58, 267–288.

Timmermann, Allan G, 1993, "How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices," *Quarterly Journal of Economics*, 1135–1145.

Yin, Y. Q., Z. D. Bai, and P. R. Krishnaiah, 1988, "On the Limit of the Largest Eigenvalue of the Large Dimensional Sample Covariance Matrix," *Probability Theory and Related Fields* 78, 509–521.

A.   Proofs

Proof of Proposition 1. The result follows on making Assumptions 1 and 3 in the Bayesian linear model of Lindley and Smith (1972). □

Proof of Proposition 2. The realized return is

$$\boldsymbol{r}_{t+1} = \boldsymbol{y}_{t+1} - \boldsymbol{p}_t = \boldsymbol{y}_{t+1} - \left[\boldsymbol{y}_t + \widetilde{\mathbb{E}}_t\left(\boldsymbol{X}\boldsymbol{g} + \boldsymbol{e}_{t+1}\right)\right] = \Delta\boldsymbol{y}_{t+1} - \boldsymbol{X}\tilde{\boldsymbol{g}}_t,.$$

From Assumption 1 and equation (6), this becomes

$$
\begin{aligned}
\boldsymbol{r}_{t+1} &= \boldsymbol{X}\boldsymbol{g} + \boldsymbol{e}_{t+1} - \boldsymbol{X}\boldsymbol{\Gamma}_t(\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\overline{\Delta\boldsymbol{y}_t} \\
&= \boldsymbol{X}\boldsymbol{g} + \boldsymbol{e}_{t+1} - \boldsymbol{X}\boldsymbol{\Gamma}_t(\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\left[\boldsymbol{X}\boldsymbol{g} + \overline{\boldsymbol{e}}_t\right] \\
&= \boldsymbol{X}\boldsymbol{g} + \boldsymbol{e}_{t+1} - \boldsymbol{X}\boldsymbol{\Gamma}_t\boldsymbol{g} - \boldsymbol{X}\boldsymbol{\Gamma}_t(\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\overline{\boldsymbol{e}}_t \\
&= \boldsymbol{X}\left(\boldsymbol{I} - \boldsymbol{\Gamma}_t\right)\boldsymbol{g} - \boldsymbol{X}\boldsymbol{\Gamma}_t(\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\overline{\boldsymbol{e}}_t + \boldsymbol{e}_{t+1}
\end{aligned}
$$

as required. □

Proof of Proposition 3. We start by showing that $\boldsymbol{h}_{t+1}'\boldsymbol{X}'\boldsymbol{X}\boldsymbol{h}_{t+1}$ is asymptotically Normal after appropriate standardization. (Recall that we view $\boldsymbol{X}$ as fixed and given.) Note first that as $\boldsymbol{h}_{t+1}$ is a linear combination of normally distributed mean-zero random variables, it is itself normally distributed with $\mathbb{E}\,\boldsymbol{h}_{t+1} = 0$. From equation (9), we have

$$N\,\mathbb{E}[\boldsymbol{h}_{t+1}\boldsymbol{h}_{t+1}'] = \frac{N\theta}{J}(\boldsymbol{I} - \boldsymbol{\Gamma}_t)^2 + \frac{N}{t}\boldsymbol{\Gamma}_t\left(\boldsymbol{X}'\boldsymbol{X}\right)^{-1}\boldsymbol{\Gamma}_t + N\left(\boldsymbol{X}'\boldsymbol{X}\right)^{-1} \tag{29}$$

after imposing $\boldsymbol{\Sigma}_g = \frac{\theta}{J}\boldsymbol{I}$ and $\boldsymbol{\Sigma}_e = \boldsymbol{I}$ (Assumption 4). It follows from the eigendecomposition (4) $\frac{1}{N}\boldsymbol{X}'\boldsymbol{X} = \boldsymbol{Q}\boldsymbol{\Lambda}\boldsymbol{Q}'$ that the terms on the right-hand side can be expressed in the form $\boldsymbol{Q}\boldsymbol{\Xi}\boldsymbol{Q}'$ for appropriately chosen diagonal matrices $\boldsymbol{\Xi}$:

$$\boldsymbol{\Gamma}_t = \boldsymbol{Q}\left(\boldsymbol{I} + \frac{J}{\theta N t}\boldsymbol{\Lambda}^{-1}\right)^{-1}\boldsymbol{Q}' \tag{30}$$

$$\boldsymbol{I} - \boldsymbol{\Gamma}_t = \boldsymbol{Q}\left(\boldsymbol{I} + \frac{\theta N t}{J}\boldsymbol{\Lambda}\right)^{-1}\boldsymbol{Q}'. \tag{31}$$

We will repeatedly exploit the fact that $\boldsymbol{Q}$ is orthogonal, that is, $\boldsymbol{Q}\boldsymbol{Q}' = \boldsymbol{Q}'\boldsymbol{Q} = \boldsymbol{I}$.

Adding the three terms on the right-hand side of equation (29), we find

$$N\,\mathbb{E}[\boldsymbol{h}_{t+1}\boldsymbol{h}_{t+1}'] = \boldsymbol{Q}\boldsymbol{\Omega}\boldsymbol{Q}',$$

where $\boldsymbol{\Omega}$ is a diagonal matrix of eigenvalues with positive diagonal elements

$$\omega_{j,t} = \frac{1}{t\lambda_j + \frac{J}{\theta N}} + \frac{1}{\lambda_j}.$$

Now define $\boldsymbol{u}_{t+1} = \sqrt{N}\boldsymbol{\Omega}^{-1/2}\boldsymbol{Q}'\boldsymbol{h}_{t+1}$, so that $\sqrt{N}\boldsymbol{h}_{t+1} = \boldsymbol{Q}\boldsymbol{\Omega}^{1/2}\boldsymbol{u}_{t+1}$ and $\boldsymbol{u}_{t+1}$ is standard Normal: $\mathbb{E}\,\boldsymbol{u}_{t+1}\boldsymbol{u}'_{t+1} = N\boldsymbol{\Omega}^{-1/2}\boldsymbol{Q}'\frac{1}{N}\boldsymbol{Q}\boldsymbol{\Omega}\boldsymbol{Q}'\boldsymbol{Q}\boldsymbol{\Omega}^{-1/2} = \boldsymbol{I}$. We can then write

$$\boldsymbol{h}'_{t+1}\boldsymbol{X}'\boldsymbol{X}\boldsymbol{h}_{t+1} = \underbrace{\boldsymbol{u}'_{t+1}\boldsymbol{\Omega}^{1/2}\boldsymbol{Q}'}_{\sqrt{N}\boldsymbol{h}'_{t+1}}\underbrace{\boldsymbol{Q}\boldsymbol{\Lambda}\boldsymbol{Q}'}_{\frac{1}{N}\boldsymbol{X}'\boldsymbol{X}}\underbrace{\boldsymbol{Q}\boldsymbol{\Omega}^{1/2}\boldsymbol{u}_{t+1}}_{\sqrt{N}\boldsymbol{h}_{t+1}}$$
$$= \boldsymbol{u}'_{t+1}\boldsymbol{\Omega}^{1/2}\boldsymbol{\Lambda}\boldsymbol{\Omega}^{1/2}\boldsymbol{u}_{t+1}$$
$$= \boldsymbol{u}'_{t+1}\boldsymbol{\Omega}\boldsymbol{\Lambda}\boldsymbol{u}_{t+1}.$$

The last line exploits the fact that $\boldsymbol{\Omega}^{1/2}$ and $\boldsymbol{\Lambda}$ commute, as they are diagonal.

Thus

$$\boldsymbol{h}'_{t+1}\boldsymbol{X}'\boldsymbol{X}\boldsymbol{h}_{t+1} = \sum_{j=1}^{J}\zeta_{j,t}u_j^2, \tag{32}$$

where $\zeta_{j,t}$ are the diagonal entries of the diagonal matrix $\boldsymbol{\Omega}\boldsymbol{\Lambda}$ and u_j are independent $N(0,1)$ random variables (the entries of $\boldsymbol{u}_{t+1}$). Explicitly,

$$\zeta_{j,t} = \omega_{j,t}\lambda_j = \frac{\lambda_j}{t\lambda_j + \frac{J}{\theta N}} + 1. \tag{33}$$

(For comparison, under the rational expectations null we would have $\boldsymbol{\Omega} = \boldsymbol{\Lambda}^{-1}$ and hence $\zeta_{j,t} = 1$ for all j.) As we have $\lambda_j > 0$ by positive definiteness of $\boldsymbol{X}'\boldsymbol{X}$, it follows that for $t \geq 1$, $\zeta_{j,t} \in (1,2)$. Moreover, as $\lim_{J,N\to\infty}\frac{J}{N} = \psi > 0$ and (by Assumption 5) $\lambda_j > \varepsilon$, $\zeta_{j,t}$ is uniformly bounded away from 1 and 2. It follows that $\mu \in (1,2)$ and $\sqrt{\mu^2 + \sigma^2} \in (1,2)$.

We will apply Lyapunov's version of the central limit theorem to $\sum_{j=1}^{J}\zeta_{j,t}u_j^2$, which here requires that for some $\delta > 0$

$$\lim_{N,J\to\infty}\frac{1}{s_J^{2+\delta}}\sum_{j=1}^{J}\zeta_{j,t}^{2+\delta}\,\mathbb{E}\left[\left|u_j^2 - 1\right|^{2+\delta}\right] = 0 \quad \text{where} \quad s_J^2 = 2\sum_{j=1}^{J}\zeta_{j,t}^2.$$

It is enough to show that this holds for $\delta = 2$. But as $\mathbb{E}\left[(u_j^2 - 1)^4\right] = 60$ and $\zeta_{j,t} \in (1,2)$,

$$\frac{1}{s_J^4}\sum_{j=1}^{J}\zeta_{j,t}^4\,\mathbb{E}\left[\left|u_j^2 - 1\right|^4\right] = \frac{60\sum_{j=1}^{J}\zeta_{j,t}^4}{\left(2\sum_{j=1}^{J}\zeta_{j,t}^2\right)^2} \leq \frac{960J}{4J^2} \to 0 \quad \text{as } J \to \infty,$$

as required. Therefore the central limit theorem applies for $\boldsymbol{h}'_{t+1}\boldsymbol{X}'\boldsymbol{X}\boldsymbol{h}_{t+1} = \sum_{j=1}^{J}\zeta_{j,t}u_j^2$ after appropriate standardization by mean and variance, which (as the u_j are IID standard

54

Normal) are $\sum_{j=1}^{J} \zeta_{j,t}$ and $2 \sum_{j=1}^{J} \zeta_{j,t}^2$, respectively. Thus we have

$$T_b \equiv \frac{h'_{t+1} X' X h_{t+1} - \sum_{j=1}^{J} \zeta_{j,t}}{\sqrt{2 \sum_{j=1}^{J} \zeta_{j,t}^2}} \xrightarrow{d} N(0,1).$$

The remaining results follow immediately. $\qquad\square$

Proof of Proposition 4. The first statement follows from the second. To prove the second, note that Proposition 3 implies that

$$\mathbb{P}\left(T_{re} < c_\alpha\right) = \mathbb{P}\left(\frac{T_{re}}{\sqrt{\mu^2 + \sigma^2}} - \frac{\mu - 1}{\sqrt{2(\mu^2 + \sigma^2)}}\sqrt{J} < \frac{c_\alpha}{\sqrt{\mu^2 + \sigma^2}} - \frac{\mu - 1}{\sqrt{2(\mu^2 + \sigma^2)}}\sqrt{J}\right)$$

$$\to \Phi\left(\frac{c_\alpha}{\sqrt{\mu^2 + \sigma^2}} - \frac{\mu - 1}{\sqrt{2(\mu^2 + \sigma^2)}}\sqrt{J}\right),$$

where $\Phi(\cdot)$ denotes the standard Normal cumulative distribution function. The result follows from the well-known inequalities $\frac{e^{-x^2/2}}{|x + \frac{1}{x}|\sqrt{2\pi}} < \Phi(x) < \frac{e^{-x^2/2}}{|x|\sqrt{2\pi}}$, which hold for $x < 0$. $\qquad\square$

Proof of Proposition 5. By decomposing returns into in-sample predicted component and residual, we can write the portfolio return as

$$r_{IS,t+1} = \frac{1}{N} h'_{t+1} X' X h_{t+1} + \frac{1}{N} h'_{t+1} X' \left(I - X'(X'X)^{-1}X'\right) e_{t+1}$$

$$= \frac{1}{N} h'_{t+1} X' X h_{t+1}$$

$$= \frac{1}{N} \sum_{j=1}^{J} \zeta_{j,t} u_j^2,$$

where the last line comes from equation (32) and the u_j are IID standard Normal. Thus

$$\mathbb{E}\, r_{IS,t+1} = \frac{\sum_{j=1}^{J} \zeta_{j,t}}{N} \quad \text{and} \quad \operatorname{var} r_{IS,t+1} = \frac{2 \sum_{j=1}^{J} \zeta_{j,t}^2}{N^2},$$

so that $\mathbb{E}\, r_{IS,t+1} \to \psi\mu$ and $N \operatorname{var} r_{IS,t+1} \to 2\psi\left(\mu^2 + \sigma^2\right)$ as $N, J \to \infty$ with $J/N \to \psi$. This gives the result. $\qquad\square$

Proof of Proposition 6. When $t > 0$, the cross-sectional moments of $\zeta_{j,t}$ can be computed using equation (14) and the fact that the eigenvalues λ_j follow (in the asymptotic limit) the Marchenko–Pastur distribution, whose probability density function $f_\lambda(x)$ takes the form

$$f_\lambda(x) = \frac{1}{2\pi} \frac{\sqrt{\left[\left(1 + \sqrt{\psi}\right)^2 - x\right]\left[x - \left(1 - \sqrt{\psi}\right)^2\right]}}{\psi x}$$

if $\left(1 - \sqrt{\psi}\right)^2 \le x \le \left(1 + \sqrt{\psi}\right)^2$, and $f_\lambda(x) = 0$ elsewhere. The relevant integrals can be calculated explicitly, giving equations (16) and (17); and we can calculate the probability density function of $\zeta_{j,t}$ in the asymptotic limit by change of variable using the relationship between $\zeta_{j,t}$ and λ_j given in equation (33). $\qquad\square$

Proof of Proposition 7. As $\boldsymbol{b}_{s+1} = (\boldsymbol{C}'\boldsymbol{C})^{-1}\boldsymbol{C}'\boldsymbol{r}_{s+1}$, we have

$$\mathbb{E}\left[\boldsymbol{r}_{t+1}\left(\boldsymbol{C}\boldsymbol{b}_{s+1}\right)'\right] = \mathbb{E}\left[\boldsymbol{r}_{t+1}\boldsymbol{r}'_{s+1}\right]\boldsymbol{C}\left(\boldsymbol{C}'\boldsymbol{C}\right)^{-1}\boldsymbol{C}'.$$

We will show that $\mathbb{E}\left[\boldsymbol{r}_{t+1}\boldsymbol{r}'_{s+1}\right] = 0$ when $s \ne t$—in other words, that all non-contemporaneous autocorrelations and cross-correlations are zero. Henceforth we assume that $s < t$ without loss of generality. From equation (8),

$$\mathbb{E}\left[\boldsymbol{r}_{t+1}\boldsymbol{r}'_{s+1}\right] = \frac{\theta}{J}\boldsymbol{X}(\boldsymbol{I} - \boldsymbol{\Gamma}_t)(\boldsymbol{I} - \boldsymbol{\Gamma}_s)\boldsymbol{X}' + \frac{1}{t}\boldsymbol{X}\boldsymbol{\Gamma}_t(\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{\Gamma}_s\boldsymbol{X}' - \frac{1}{t}\boldsymbol{X}\boldsymbol{\Gamma}_t(\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'.$$

We have used the facts that (i) $\boldsymbol{\Sigma}_g = (\theta/J)\boldsymbol{I}$, (ii) $\mathbb{E}\,\bar{\boldsymbol{e}}_t\boldsymbol{e}'_{s+1} = (1/t)\boldsymbol{I}$, (iii) $\mathbb{E}\,\boldsymbol{e}_{t+1}\bar{\boldsymbol{e}}'_s = 0$, and (iv), $\mathbb{E}\,\bar{\boldsymbol{e}}_t\bar{\boldsymbol{e}}'_s = (1/t)\boldsymbol{I}$. This expression can be rearranged as

$$\mathbb{E}\left[\boldsymbol{r}_{t+1}\boldsymbol{r}'_{s+1}\right] = \frac{1}{t}\boldsymbol{X}\left[\frac{\theta t}{J}(\boldsymbol{I} - \boldsymbol{\Gamma}_t)\boldsymbol{X}'\boldsymbol{X} - \boldsymbol{\Gamma}_t\right](\boldsymbol{X}'\boldsymbol{X})^{-1}(\boldsymbol{I} - \boldsymbol{\Gamma}_s)\boldsymbol{X}'.$$

As $\frac{\theta t}{J}(\boldsymbol{I} - \boldsymbol{\Gamma}_t)\boldsymbol{X}'\boldsymbol{X} = \boldsymbol{\Gamma}_t$ (from equations (30) and (31) and the fact that $\boldsymbol{X}'\boldsymbol{X} = N\boldsymbol{Q}'\boldsymbol{\Lambda}\boldsymbol{Q}$), the term in square brackets vanishes, and the result follows. $\qquad\square$

Proof of Proposition 8. As $\boldsymbol{r}'_{t+1}\boldsymbol{C}\boldsymbol{b}_{s+1} = \boldsymbol{b}'_{t+1}\boldsymbol{C}'\boldsymbol{C}\boldsymbol{b}_{s+1}$, we want to understand the behavior of $\boldsymbol{b}_{t+1}\boldsymbol{C}'\boldsymbol{C}\boldsymbol{b}_{s+1}$. We assume $s < t$. Assumption 6 implies that

$$\frac{1}{N}\boldsymbol{X}'\boldsymbol{X} = \begin{pmatrix} \frac{1}{N}\boldsymbol{C}'\boldsymbol{C} & 0 \\ 0 & \frac{1}{N}\boldsymbol{M}'\boldsymbol{M} \end{pmatrix}$$

is block-diagonal (where zeros indicate conformable matrices of zeros). Eigendecomposing $\frac{1}{N}\boldsymbol{C}'\boldsymbol{C} = \boldsymbol{Q}_C\boldsymbol{\Lambda}_C\boldsymbol{Q}'_C$ and $\frac{1}{N}\boldsymbol{M}'\boldsymbol{M} = \boldsymbol{Q}_M\boldsymbol{\Lambda}_M\boldsymbol{Q}'_M$ where $\boldsymbol{Q}'_i\boldsymbol{Q}_i = \boldsymbol{I}$ for $i = C, M$, we have

$$\frac{1}{N}\boldsymbol{X}'\boldsymbol{X} = \underbrace{\begin{pmatrix} \boldsymbol{Q}_C & 0 \\ 0 & \boldsymbol{Q}_M \end{pmatrix}}_{\boldsymbol{Q}}\underbrace{\begin{pmatrix} \boldsymbol{\Lambda}_C & 0 \\ 0 & \boldsymbol{\Lambda}_M \end{pmatrix}}_{\boldsymbol{\Lambda}}\underbrace{\begin{pmatrix} \boldsymbol{Q}'_C & 0 \\ 0 & \boldsymbol{Q}'_M \end{pmatrix}}_{\boldsymbol{Q}'}.$$

It follows that $\boldsymbol{\Gamma}_t = \begin{pmatrix} \boldsymbol{\Gamma}_{C,t} & 0 \\ 0 & \boldsymbol{\Gamma}_{M,t} \end{pmatrix}$ is block-diagonal, with $\boldsymbol{\Gamma}_{i,t} = \boldsymbol{Q}_i\left(\boldsymbol{I} + \frac{J}{N\theta t}\boldsymbol{\Lambda}_i^{-1}\right)^{-1}\boldsymbol{Q}'_i$ and $\boldsymbol{I} - \boldsymbol{\Gamma}_{i,t} = \boldsymbol{Q}_i\left(\boldsymbol{I} + \frac{N\theta t}{J}\boldsymbol{\Lambda}_i\right)^{-1}\boldsymbol{Q}'_i$ for $i = C, M$.

As $\mathbb{E}\,\boldsymbol{r}_{t+1}\boldsymbol{r}'_{t+1} = \frac{\theta}{J}\boldsymbol{X}\left(\boldsymbol{I} - \boldsymbol{\Gamma}_t\right)\boldsymbol{X}' + \boldsymbol{I}$, we have

$$\begin{aligned}
\mathbb{E}\,\boldsymbol{b}_{t+1}\boldsymbol{b}'_{t+1} &= \left(\boldsymbol{C}'\boldsymbol{C}\right)^{-1}\boldsymbol{C}'\,\mathbb{E}\,\boldsymbol{r}_{t+1}\boldsymbol{r}'_{t+1}\boldsymbol{C}\left(\boldsymbol{C}'\boldsymbol{C}\right)^{-1} \\
&= \frac{\theta}{J}\left(\boldsymbol{C}'\boldsymbol{C}\right)^{-1}\boldsymbol{C}'\boldsymbol{X}\left(\boldsymbol{I} - \boldsymbol{\Gamma}_t\right)\boldsymbol{X}'\boldsymbol{C}\left(\boldsymbol{C}'\boldsymbol{C}\right)^{-1} + \left(\boldsymbol{C}'\boldsymbol{C}\right)^{-1}.
\end{aligned}$$

56

As $C'M = 0$, we have $(C'C)^{-1} C'X (I - \Gamma_t) X'C (C'C)^{-1} = I - \Gamma_{C,t}$. Hence $\mathbb{E}\, b_{t+1}b'_{t+1} = \frac{\theta}{J}(I - \Gamma_{C,t}) + (C'C)^{-1}$. Summarizing,

$$
N \,\mathbb{E}\, b_{t+1}b'_{t+1} = Q_C \underbrace{\left[\left(\frac{J}{N\theta}I + t\Lambda_C\right)^{-1} + \Lambda_C^{-1}\right]}_{\Omega_{C,t}} Q'_C\,.
$$

Set $v_{t+1} = \sqrt{N}\Omega_{C,t}^{-1/2}Q'_C b_{t+1}$. Then $\sqrt{N}b_{t+1} = Q_C\Omega_{C,t}^{1/2}v_{t+1}$ and $\mathbb{E}\, v_{t+1}v'_{t+1} = I$. We also have $\mathbb{E}\, v_{s+1}v'_{t+1} = 0$ whenever $s \neq t$ (because, as shown in the proof of Proposition 7, $\mathbb{E}\, r_{s+1}r'_{t+1} = 0$ and hence $\mathbb{E}\, b_{s+1}b'_{t+1} = 0$). Thus v_{t+1} and v_{s+1} are independent standard Normal random vectors. We have

$$
b'_{t+1}C'Cb_{s+1} = v'_{t+1}\Omega_{C,t}^{1/2}Q'_C Q_C\Lambda_C Q'_C Q_C\Omega_{C,s}^{1/2}v_{s+1} = v'_{t+1}\Omega_{C,t}^{1/2}\Lambda_C\Omega_{C,s}^{1/2}v_{s+1}\,.
$$

As $\Omega_{C,t}^{1/2}\Lambda_C\Omega_{C,s}^{1/2}$ is a $P \times P$ diagonal matrix with ith diagonal entry $\sqrt{\zeta_{i,t}\zeta_{i,s}}$, we can write

$$
b'_{t+1}C'Cb_{s+1} = \sum_{i=1}^{P}\sqrt{\zeta_{i,t}\zeta_{i,s}}w_i\,,
$$

where w_i denotes the product of the ith entries of v_{t+1} and v_{s+1}. The w_i are independent of each other, and each is the product of two independent standard Normal random variables. Therefore each w_i has zero mean and unit variance.

We wish to apply Lyapunov's version of the central limit theorem to $\sum_{i=1}^{P}\sqrt{\zeta_{i,t}\zeta_{i,s}}w_i$, which here requires that for some $\delta > 0$

$$
\lim_{N,J,P\to\infty}\frac{1}{s_P^{2+\delta}}\sum_{p=1}^{P}\mathbb{E}\left[\left|\sqrt{\zeta_{i,t}\zeta_{i,s}}w_i\right|^{2+\delta}\right] = 0 \quad\text{where}\quad s_P^2 = \sum_{p=1}^{P}\zeta_{i,t}\zeta_{i,s}.
$$

It is enough to show that this holds when $\delta = 2$. In this case, as the fourth moment of a standard Normal random variable equals 3, we have $\mathbb{E}\left[w_i^4\right] = 9$, and so indeed

$$
\frac{1}{s_P^4}\sum_{p=1}^{P}\mathbb{E}\left[\left|\sqrt{\zeta_{i,t}\zeta_{i,s}}w_i\right|^4\right] = \frac{9\sum_{p=1}^{P}\zeta_{i,t}^2\zeta_{i,s}^2}{\left(\sum_{p=1}^{P}\zeta_{i,s}\zeta_{i,t}\right)^2} \leq \frac{144P}{P^2} \to 0 \quad\text{as } P\to\infty.
$$

(The inequality follows because $\zeta_{i,t} \in (1,2)$ for all i and $t \geq 1$.) Hence the central limit theorem applies for $\sum_{i=1}^{P}\sqrt{\zeta_{i,t}\zeta_{i,s}}w_i$ after appropriate standardization by mean and variance, which are 0 and $\sum_{i=1}^{P}\zeta_{i,t}\zeta_{i,s}$, respectively. Thus

$$
\frac{b'_{t+1}C'Cb_{s+1}}{\sqrt{\sum_{i=1}^{P}\zeta_{i,t}\zeta_{i,s}}} \xrightarrow{d} N(0,1)\,. \qquad\qquad \square
$$