THE UNIVERSITY OF CHICAGO

MACHINE LEARNING, QUANTITATIVE PORTFOLIO CHOICE, AND MISPRICING

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
CARTER KENT DAVIS

CHICAGO, ILLINOIS
JUNE 2022

To my incredible wife, Whitney Davis, whose selfless sacrifices keep our family afloat

TABLE OF CONTENTS

LIST OF FIGURES vii

LIST OF TABLES viii

ACKNOWLEDGMENTS ix

ABSTRACT xi

1 INTRODUCTION 1
- 1.1 Example with Pecunia 1
- 1.2 Antibiotic Resistant Bacteria Analogy 4
- 1.3 Approach and Main Findings 5
- 1.4 Additional Motivation 7
- 1.5 Description of the Model 8
- 1.6 Related Literature 13
- 1.7 Dissertation Outline 15

2 BASIC ECONOMICS OF THE MODEL 16
- 2.1 Simplified Model 17
- 2.2 Canonical Portfolio Choice Methods Generate Price Nonreactive Demand . . 21
  - 2.2.1 Canonical Portfolio Choice Decision 21
  - 2.2.2 Continuous Characteristic Weighting Instead of Discontinuous Jumps 24
  - 2.2.3 Summary of Demand Generated from Canonical Portfolio Choice Methods 25

3 HIGH DIMENSIONAL PROBLEM 26
- 3.1 Review of CARA Model with Full Knowledge of Parameters 26
- 3.2 Bayesian Learning 28
- 3.3 Learning from Portfolio Returns 31

4 FULL MODEL OVERVIEW 39
- 4.1 Two Kinds of Incumbents 39
- 4.2 Investor Problem 41
- 4.3 Equilibria 42
  - 4.3.1 Combined Equilibrium 43
  - 4.3.2 Incumbents-only Equilibrium 43
- 4.4 Example of Increasing Mispricing—Numerical Example with Investors 44

5 DIVIDEND PROCESS ESTIMATION 48
- 5.1 Data 50
- 5.2 Estimation of Size 50
- 5.3 Estimation of Return Process 51

5.3.1 Out-of-Sample Test . . . 51
5.3.2 Return Process Estimation . . . 52
5.4 Obtaining Incumbents-Only Prices from Market Equity . . . 53

6 DEMAND FUNCTION ESTIMATION . . . 54
6.1 Reason for using Estimated Demand Learners . . . 54
6.2 Reason for using Estimated Demand Incumbents . . . 54
6.3 Data . . . 55
6.4 Justification for using Linear Demand Function . . . 55
6.5 Estimation of Institutional Demand for Learners . . . 56
6.6 Estimation of Incumbent Asset Manager Demand . . . 60

7 LEARNERS . . . 62
7.1 Market Indexer . . . 62
7.2 Bayesian Learner . . . 63
7.3 True DGP Maximum Likelihood Learners . . . 63
7.4 True DGP MAP Learners . . . 64
7.5 True DGP Learners with Asymptotic Parameter Values . . . 64
7.6 Portfolio Optimizer . . . 65
7.7 Portfolio Optimizer with Shrinking Learner . . . 66
7.8 Asymptotic Portfolio Optimizer Learner . . . 67
7.9 DeMiguel et al. (2009) $1/N$ Learners . . . 68
7.10 Random Forest Learners . . . 68
7.11 Neural Network Learner . . . 69

8 MODEL SIMULATION RESULTS . . . 70
8.1 Mispricing can Increase When Sophisticated Learner Assets Managers Enter the Market . . . 72
8.2 Price Elasticites are Important Determinants in Mispricing Outcomes . . . 72
8.3 Shrinking Can Eliminate More Mispricing Than Knowing the True Covariance Matrix . . . 73
8.4 Other Interesting Results . . . 74

9 CONCLUSION . . . 75

BIBLIOGRAPHY . . . 77

A APPENDIX A: FIGURES . . . 79

B APPENDIX B: TABLES . . . 90

C APPENDIX C: FURTHER DETAILS AND PROOFS . . . 101
C.1 Lemmas and Proofs . . . 101
C.1.1 Lemma 1 . . . 101
C.1.2 Proof of Proposition 1.1 . . . 101

C.1.3 Proof of Proposition 1.2 . . . 106
C.1.4 Proof of Proposition 2.1 . . . 106
C.1.5 Proof of Proposition 2.2 . . . 107
C.1.6 Proof of Proposition 2.3 . . . 109
C.2 Additional Details . . . 110
C.2.1 Theoretical Model Details . . . 110
C.2.2 Learner Capital Percent . . . 119
C.2.3 Transformation of Characteristics . . . 120
C.2.4 General Neural Network . . . 121
C.2.5 Consistency of Main Econometric Model . . . 122
C.2.6 Matrix Identities for Econometric Method . . . 123
C.2.7 Calculating the Value of the Market . . . 124
C.2.8 Details of Example of Increasing Mispricing . . . 125
C.2.9 Overpricing and Information Ratio Components . . . 126
C.2.10 Brandt et al. (2009) and Kozak et al. (2020) Estimator . . . 127
C.2.11 Brandt et al. (2009) and Kozak et al. (2020) Estimator Consistency Given Data Generating Process . . . 130
C.2.12 $1/N$ Demand Rescaling . . . 132
C.2.13 Random Forest Estimator . . . 132

LIST OF FIGURES

A.1 Machine Learning and Investment Management Survey 79
A.2 Model Timeline 80
A.3 Vector Plot of Equilibrium Example 81
A.4 Mispricing and Learner Information Ratio Along Path From Incumbents-Only Investment to Equilibrium 82
A.5 Dollar Mispricing and Utility Along Path From Incumbents-Only Investment to Equilibrium 84
A.6 Mispricing Target 85
A.7 Mispricing Results Summary 86
A.8 Mispricing Results with Estimated Demand Learners 88

LIST OF TABLES

B.1 Definitions and Notation 90
B.2 Out-of-Sample Portfolio CAPM Regression 92
B.3 Description of Characteristics 93
B.4 Description of Exogenous Characteristics 94
B.5 Mispricing Summary Table 95
B.6 Mispricing with Bayesian Learners 96
B.7 Mispricing with CARA Incumbents and True DGP + Shrinking Learners 97
B.8 Mispricing with Estimated Demand Incumbents and True DGP + Shrinking Learners 98
B.9 Mispricing with Portolio Optimizer + Shrinking Learners 99

ACKNOWLEDGMENTS

I want to thank Lubos Pastor, Stefan Nagel, Ralph Koijen, and Niels Gormsen for their countless hours of guidance and mentoring. Lubos Pastor was my dissertation chair, and I was stunned by how much time he generously gave me. Being advised by these academic titans, including having Lubos as my chair, was one of the most incredible and singular privileges of my life. My debt of gratitude is deep and impossible to fully express.

I am deeply grateful for my wife, Whitney Davis. She selflessly supported me through my PhD program in countless ways. She is the glue in my life and our children's lives—with her we hold together and function and without her we fall apart. She is bright and full of life.

I also want to thank the PhD program office members: Malaina Brown, Cynthia Hillman, Kimberly Mayer, and Amity James. Their patience, diligence, and support for me and other Booth PhD students is both seemingly boundless and deeply appreciated. I cannot imagine a more competent and caring team.

I also want to acknowledge financial support from the Fama-Miller Center for Research in Finance and the John and Serena Liew Fama-Miller PhD Fellowship.

I also want to thank Michael Weber, Joachim Freyberger, and Andreas Neuhierl for providing me with excellent data. I want to thank Doug Diamond, Lars Hansen, John Heaton, Anthony Zhang, Bryan Kelly, Toby Moskowitz, Stefano Giglio, James Choi, Christopher Clayton, Paul Goldsmith-Pinkham, Gary Gorton, Song Ma, Elena Asparouhova, Jonathan Brogaard, Davidson Heath, Mark Jansen, Avner Kalay, Yihui Pan, J Li, Philip Bond, Stephan Siegel, Léa Stern, Christopher Hrdlicka, Yang Song, Adem Atmaz, Huseyin Gulen, Svetlana Bryzgalova, Roberto Gomez Cram, Marco Grotteria, Christopher Hennessy, Victor DeMiguel, Steven Kou, Andrew Lyasoff, Max Reppen, Scott Robertson, Lucy White, Hao Xing, Fernando Zapatero, Christian Heyerdahl-Larsen, Foti Grigoris, Preetesh Kantak, Fahiz Baba Yara, Craig Holden, Noah Stoffman, Daniel Carvalho, James H. Cardon, Brigham

Frandsen, Lars Lefgren, Olga Stoddard, Mark Showalter, Riley Wilson, Brian Boyer, Tyler Shumway, Karl Diether, Ryan Pratt, Craig Merrill, Darren Aiello, Taylor Nadauld and seminar participants at University of Utah, University of Washington, Yale, Purdue, London Business School, Boston University, Brigham Young University Economics Department, Brigham Young University Finance Department and Indiana University for their helpful comments and input. All remaining errors are my own.

ABSTRACT

What happens to mispricing when quantitative learners—asset managers who use quantitative methods to make portfolio choice decisions—enter the market? Mispricing can actually *increase* when these learners enter and trade against historical mispricing because estimation error and model error limit their ability to properly adapt to changing prices caused by their own asset demand. This causes some asset prices to be corrected relatively little, while other assets that are initially underpriced (overpriced) become overpriced (underpriced). In a model with an estimated dividend process and a Koijen and Yogo (2019) style demand system, learner entrants who invest with some canonical quantitative methods—such as Brandt et al. (2009), Kozak et al. (2020), and DeMiguel et al. (2009) methods—tend to increase mispricing. When mispricing does not increase, a substantial amount of mispricing remains even when the learners have access to a long time series of data.

CHAPTER 1

INTRODUCTION

[M]achine learning is not all about alpha. This is important because most discussions, and certainly most anecdotes, of machine learning applied to finance focus on the creation of alpha. Using new data and machine learning to build alpha (i.e., to find new, unique sources of return predictability) heads straight into the most competitive aspect of financial markets. As more investors enter the market with similar data and similar tools, the mispricing corrects and that alpha compresses.

Ronen Israel, Bryan Kelly and Tobias Moskowitz (2020)

The quote above from Israel et al. (2020) indicates that if enough traders have accurately estimated machine learning models in hand, mispricing corrects. Is machine learning a cure for mispricing?

Surprisingly, machine learning can in fact *increase* mispricing. Even when machine learning compresses alpha, substantial amounts of mispricing can remain. How can this be the case?

1.1 Example with Pecunia

Imagine, for the sake of concreteness, that an asset manager has built an advanced machine learning model to estimate the optimal mean-variance efficient weights for a large universe of assets—say 2,000 assets. This asset manager, call her Pecunia, takes her model in hand to the market. Pecunia is a mean-variance optimizer: she simply wants the highest Sharpe ratio. Other quantitative asset managers have built sophisticated models as well, and their demand is moving asset prices. Pecunia sees these price movements as she is taking long and short positions in the various assets. How should Pecunia react?

Pecunia needs to not only estimate the mean-variance efficient portfolio weights at the old prices, she also needs to know how the mean-variance efficient weights change as prices move. For instance, if the first asset's price increases, does this affect the mean-variance weights of the first asset? Does this price movement of the first asset affect the mean-variance weights of the second asset? In other words, the mean-variance efficient weights at the old prices are simply a 2,000 dimensional vector. However, understanding how to adapt investment decisions based on changing prices potentially involves a 2,000 $\times$ 2,000 matrix of slope terms that dictate how demand changes as prices change.

Thus Pecunia's model must answer the following two questions, referenced throughout the paper:

Q0: What are the mean-variance efficient weights at the original prices?

Q1: How do the mean-variance efficient weights change as prices move away from the original prices?

Suppose that Pecunia simply decides the mean-variance efficient weights do not move as prices move. This rather aggressive and determined trading strategy is akin to simply placing a market order for every asset. In other words, Pecunia is *insufficiently price reactive*. In fact, in this case where her demand does not change as prices change, her demand is *completely price nonreactive*. If enough asset managers have similar demand, there can easily be an overshooting effect. An overshooting effect occurs when assets that are originally underpriced (overpriced) are demanded (sold short) by traders who fail to sufficiently react as the prices become overpriced (underpriced). Overshooting can occur when demand is insufficiently reactive, and does not require completely price nonreactive demand. Notice that in this case, Pecunia's demand strongly reacts to historical mispricing, but reacts very little to *changing* prices away from historical prices. Thus, in order to avoid confusion about reacting strongly to historical prices and failing to react sufficiently to changing prices, this reactivity phraseology is adopted instead of the classic elasticity phraseology when discussing learner demand in this paper.

Suppose instead that Pecunia's model is relatively timid and conservative in terms of price movements. In other words, suppose Pecunia's model indicates that positive alpha evaporates with only minor price increases and negative alpha evaporates with only small drops in prices. Pecunia's demand in this case is *overly price reactive*. If the other traders trying to trade against alpha in the market behave similarly, very little alpha will be eliminated.

Perhaps Pecunia's model will produce demand that is overly price reactive for some assets and insufficiently price reactive for other assets. Pecunia's model will likely have error in the cross-price reaction terms as well, and with 2,000 assets there are almost 4 million of these cross-price terms. How will these errors affect alpha? Will overshooting still occur on the assets with demand that is insufficiently reactive?

As Pecunia and other asset managers like her enter the market, their Sharpe ratio is high relative to the market. Their portfolio returns are appealing. Investors allocate capital to them accordingly. As prices move away from the original prices due to this new investment, Pecunia and others make errors in responding to the changing prices. This can lead to overshooting effects with some assets, and overall mispricing can increase. These price movements also lower the Sharpe ratio of Pecunia's portfolio relative to the market.

Importantly, much of the literature, including estimation methods such as Brandt et al. (2009), Kozak et al. (2020), and DeMiguel et al. (2009), focus on Q1 above and not Q2. In other words, these methods focus on estimating mean-variance weights at original prices—prices contained in the data. These models are relatively elastic to perceived mispricing, but are insufficiently reactive to new prices. These methods do not focus on how the estimated mean-variance efficient weights adapt as a function of price movements away from the original prices.

In this paper, the asset managers using these quantitative models move prices with their own demand. These asset managers must then adapt to these price movements, and errors in how they adapt can lead to increases in mispricing.

1.2 Antibiotic Resistant Bacteria Analogy

Since the discovery by Alexander Fleming of penicillin in 1928, antibiotics have been an incredibly powerful and useful tool in maintaining human health against the onslaught of bacterial infection. However, heterogeneity across bacteria means that some bacteria survive treatments due to successful resistance mechanisms to antibiotics. While the vast majority of the bacteria may be killed, the surviving bacteria are able to multiple and become the dominant strain. According to the Centers for Disease Control and Prevention (2021), there are more than 2.8 million infections and 35,000 deaths due to antibiotic-resistant bacteria in the US each year. Humanity is in a continual arms race of adaptation against bacteria. We are caught in an endless cycle of new antibiotic treatment attacks and antibiotic resistance counterattacks.

This can be viewed as a machine learning analogy where the test data always changes when the learner starts to interact with the world. In this analogy, the learners represent the researchers who develop an antibiotic and the doctors who administer it. The current set of bacteria is the training dataset. Once the antibiotic is used in treatment, the set of bacteria differs from the pre-treatment set of bacteria due to adaptation. In machine learning terminology, the training dataset is not a good representation of the testing dataset because the learners directly impact the testing data. There is a feedback loop, an arms race, that naturally makes the learning problem difficult.

Estimating and trading on mean-variance optimal portfolio weights with historical returns is analogous to the constant battle between humans armed with antibiotics and antibiotic resistant bacteria. The learners actively learn from historical returns data and trade accordingly. By trading though, the learners fundamentally change future returns, requiring further adaptation. Similarly, by treating bacteria with antibiotics, doctors change the set of bacteria. It's a constant arms race of adaptation in both situations.

In the relatively tame environment with the original prices in the data, the quantitative portfolios appear very promising. They often have high Sharpe ratios compared to standard benchmarks. Similarly, antibiotic treatments can do very well during trial studies before the bacteria adapt. In both settings, deploying at scale erodes the efficacy of antibiotics and quantitative trading tools. While it is rational for any given person to use antibiotics, using them en masse seriously increases the risk of generating a public health crisis. Similarly, while the market impact of any given quantitative investor may be small, many quantitative investors acting in correlated ways can generate severe mispricing.

Thus machine learning tools often do outperform asset pricing benchmarks with prices seen in the data. However, if prices move due to this demand, will alpha shrink as Israel et al. (2020) argue? This model indicates that alpha does not necessarily shrink.

1.3 Approach and Main Findings

This paper uses a theoretical model in which incumbent asset managers generate both a time series of returns and mispricing initially. These incumbents can be thought of as traditional asset managers. Quantitative learners—which are asset managers investing according to quantitative models and referred to as just learners throughout the paper—then enter the market. Investors provide an endogenous mechanism to optimally allocate capital to the two types of asset managers.

The most important finding of this paper is that quantitative learner asset managers can *increase* mispricing when they enter the market. This occurs because these learners struggle to adapt as the data generating process of returns changes due to their own demand. As prices change, they struggle to react appropriately. In other words, these learners struggle to understand what the true alpha is as return dynamics shift due to their own demand. The following two results explain why the Israel et al. (2020) logic that learners will inevitably eliminate mispricing does not necessarily hold.

First, adapting to changes in the data generating process of returns due to learner demand is a fundamentally difficult problem. Many learners in the model have appealing portfolios when they are atomistic (they have not yet captured enough investment to change asset prices), but once their demand starts to shift prices, they struggle to adapt. Importantly, the model highlights that adapting is a difficult large dimensional problem.

Second, many canonical portfolio optimization methods, such as Brandt et al. (2009) and Kozak et al. (2020), struggle to adapt. These portfolio optimization methods are biased in favor of being insufficiently price reactive, which generates overshooting and can easily increase aggregate mispricing across assets. Also, even learners who learn from the true dividend process but overfit the model generate demand that struggles to adapt as prices change. They increase mispricing when they enter the market. These learners still create appealing portfolios when they are atomistic, but struggle to adapt as prices change.

Learners often still generate mispricing increases even after trading in the market for a relatively long period of time. Using model simulations, both learners and incumbent asset managers trade in the market together for over two decades, where learners are using monthly data, and many learners still increase mispricing relative to the incumbents-only equilibrium.

In the model, mispricing often increases when learner enter even with hyper-rational investors. These investors optimally allocate investment between the traditional incumbents and learner asset managers and know all asset managers' fund return distributional parameters. The paper provides a simple example of mispricing with three assets. Incumbents have relatively inelastic demand, and learners have demand that reacts to mispricing but is insufficiently reactive to changing prices. The learner aggressively invests in one of the assets in the example even as the price rises and it becomes severely overpriced. The other asset positions in the learners' portfolio provide compelling risk-return trade-offs for the investors of the learner fund, convincing them to put capital in the learner fund even as the price of the overpriced asset is pushed far from the efficient price.

Allowing incumbent demand price elasticity to deviate from traditional asset pricing model price elasticity is a key ingredient in generating mispricing increases. When incumbent asset managers react to changes in prices as constant absolute risk aversion (CARA) utility optimizers do, mispricing does not increase when learner asset managers enter. However, when incumbent demand function elasticities are estimated from institutional holdings, using Koijen and Yogo (2019) style methodology and data, mispricing does sometime increase when learners enter.

Learners who estimate all the parameters in the model but use strong priors to estimate the covariance matrix of returns (shrinking) can actually eliminate *more* mispricing than Bayesian learners who know the true covariance matrix. MacKinlay and Pástor (2000) show that shrinking the covariance matrix can actually outperform an investment strategy that uses the true covariance matrix. This paper gives a similar result, except in terms of mispricing instead of portfolio performance.

1.4 Additional Motivation

There is an increasing focus on machine learning and quantitative methods in the investment management industry. Morgan Stanley polled 400 of the largest investment managers in 2019 and found that 51% of managers say that machine learning is at least a component of their investment management strategy, up from 27% in 2016 (see Table A.1 for more details). As the use of machine learning increases in investment management, it is unclear how this affects price efficiency. As a result, this paper seeks to answer this question: how much mispricing do machine learners eliminate?

There are three main reasons this research question is important. First, there is an increasing use and focus on machine learning tools, as discussed above.

Second, this paper weakens the argument that mispricing does not exist or is very small. One long-standing debate in asset pricing is how much asset pricing model alpha is due to problems with the model or mispricing. Problems with the model include inappropriately taking into account risks, preferences, or differences between in-sample and out-of-sample tests. A common argument in favor of mispricing being either small or nonexistent is that sophisticated traders would have taken advantage of mispricing and by so doing prices would have been corrected. In this paper, I introduce sophisticated learners into a model with mispricing and show that a substantial amount of mispricing remains in the market. In other words, this model shows that mispricing may not be eliminated by sophisticated asset managers, even with optimal investment.

Third, mispricing can have important consequences for capital allocational efficiency and overall welfare. For example, van Binsbergen and Opp (2019) argue that persistent alpha can lead to capital allocational losses and negatively impact welfare. While this paper uses an exogenous dividend process, extending the model in this paper to allow endogenous cash flows that are a function of asset valuations could allow researchers to measure the welfare consequences of mispricing due to increasingly quantitative asset management.

1.5 Description of the Model

The model has an exogenous dividend process, with endogenous prices. For tractability, asset payoffs are just dividends in excess of the price of the asset. There are two kinds of asset managers: incumbents and learners. The periods of the model are divided into two eras: an incumbents-only era and a combined incumbents-and-learner era. In the incumbents-only era, investors can only invest with incumbents and the risk-free asset. Incumbent asset managers play a key role in the model by generating a time series of returns that learners can use to make decisions and mispricing that learners have a chance to eliminate.

I assume we live in an incumbents-only world during both eras, and thus I estimate the parameters of this model with data using the entire sample period. I estimate a return process at incumbents-only equilibrium prices and an exogenous measure of size. These estimated parameters combined with observed asset prices, an exogenous risk-free rate, and a risk aversion parameter are sufficient to obtain all parameters in the incumbents-only model. I use these parameter values in order to make inference throughout the paper. In order to maintain the realistic environment of changing asset characteristics through time in the model simulations, I use actual exogenous asset characteristics. Note that the combined era results are considered *counterfactual* because of the assumption that we only see the incumbents-only world. This assumption is not violated by asset managers in the data who actually use quantitative methods. This assumption just allows the model to use the observed data as a baseline to compare against counterfactual mixes of traditional incumbent and quantitative learner asset managers.

The key intuition of the model can be seen in a single equation:

$$x = \underbrace{a}_{\text{Q1}} + \underbrace{J}_{\text{Q2}}(p - \bar{p})$$

where x is the demand of an asset manager for the N assets, a is an N dimensional intercept term, J is an $N \times N$ dimensional slope (Jacobian) term, p is an N dimensional column vector of prices, and $\bar{p}$ is the N dimensional vector of original prices—prices observed in the data. The intercept term can be thought of as the estimated mean-variance efficient weights at the original prices [1] ($p = \bar{p}$). Thus the intercept represents Q1 above. In the asset pricing literature, hypothetical portfolios are created and examined, but it is assumed that the prices observed in the data are not impacted by these portfolios. For example, Fama and French (2015) looks at the returns of different portfolios sorted on various characteristics, with an

1. The units of demand are actually the share of the total value of each asset outstanding. So the intercept is actually the mean-variance portfolio weights divided by prices and rescaled.

implicit no-price-impact assumption. In fact, back-testing portfolios with a no-price-impact assumption is a component of a high fraction of asset pricing research. In other words, it's typically assumed that these hypothetical portfolios have no impact on prices and thus $p = \bar{p}$. In this case, $x = a$, and the J term does not affect the portfolio returns.

In this paper, learner demand affects prices. In this case, the slope matrix J now plays a role as p moves away from $\bar{p}$. This addresses Q2 discussed above. The matrix J represents how these learners respond to changes in prices—how they adapt to the changing environment. It represents how aggressive (insufficiently price reactive) or timid (overly price reactive) asset managers are with respect to price changes.

The key intuition of this paper is that while estimating mean variance efficient weights at original prices (a) is relatively possible (with some estimation error), adapting to changing prices can be quite difficult and lead to mispricing increases or at least substantial amounts of mispricing remaining. This is caused by overshooting effects as described above. Note that while a is N dimensional, J is $N \times N$ dimensional. Thus if there are 2,000 assets, a has only 2,000 terms while J has 4,000,000 terms.

In order to simulate the model with realistic parameters, I estimate the model return process uses a novel characteristic-based econometric model. The econometric model uses a Koijen and Yogo (2019) functional form for the mean and covariance matrix of returns and maximum likelihood to fit the parameters. I show using the Freyberger et al. (2020) data (also used in Kelly et al. (2019)) that this model produces large and statistically significant out-of-sample alphas with a monthly CAPM information ratio of 0.83. This provides some evidence that the model is able to capture some return dynamics relatively well, even out-of-sample. I estimate the return process at incumbent prices on a data set of exogenous characteristics using the entire sample period. I use a maximum a posteriori (MAP) estimator in order to shrink mispricing down to the target of 0.83 from the out-of-sample test.

Once the learners enter at the beginning of the combined incumbents-and-learners era, investors optimally allocate capital between the incumbent asset managers, learner asset managers, and the risk-free asset. Investors effectively act as a simple endogenous way for capital to be allocated between the two investment strategies. For simplicity, I assume investors know the true distributional parameters of the two types of asset managers' funds. Equilibrium is found by solving a fixed-point problem that iterates back and forth between optimal prices and optimal investor allocation between the two types of asset managers.

The learners in this model are a superset of typical agents in economic models (see Pástor and Veronesi (2009) for a review). A typical economic agent in the literature knows the true data generating process (DGP), and uses Bayesian learning to update their beliefs. While this paper considers Bayesian learning, it also considers learners who take a reasonable quantitative model that may or may not be consistent with the true DGP, estimate the parameters, and invest accordingly. In this sense, most learners in this paper are like empirically oriented economists in that they make assumptions, estimate a model, and use the parameter estimates to make inferences and decisions. In this paper, I estimate mispricing outcomes that result from various model assumptions and overall estimation error. I use learners that come from the academic literature and are shown to perform well out-of-sample, including Brandt et al. (2009), Kozak et al. (2020), and DeMiguel et al. (2009).

I interpret machine learning broadly, in line with basic machine learning courses and textbooks. For example, Mitchell (1997) and Hastie et al. (2001) describe linear regression, principle component analysis, penalty parameters in maximum a posteriori estimation, and other standard methods of statistical analysis as part of the machine learning tool set. Thus, using the term machine learners interchangeably with quantitative learner portfolio managers, or just learners throughout the paper, is in line with the field of machine learning.

There are two variations of the model with different types of incumbent asset managers: 1) CARA utility managers and 2) estimated demand function managers using Koijen and Yogo (2019) style estimation methods. The CARA utility managers make mistakes about the expected return of assets, which generates CAPM alpha. However, while these CARA incumbents make mistakes about expected asset payoffs, their price elasticity of demand is right in the sense that it equals the optimal CARA utility demand price elasticity. In other words, while these asset managers make mistakes about asset payoffs, they respond appropriately to changes in prices. This assumption is relaxed with the second type of incumbent asset manager—the estimated demand function incumbents—which have identical mispricing in the incumbents-only equilibrium, but have demand price elasticities that are allowed to differ from optimal CARA utility preferences. The model indicates that mispricing tends to increase substantially when estimated demand incumbents are used as compared to CARA incumbents. Koijen and Yogo (2019) argue that estimated asset demand price elasticities differ from classic asset pricing model price elasticities, and thus allowing this deviation from the classic mean-variance model is an important step in judging the effect of learner demand on mispricing.

In order to use estimated demand functions in the model, both as learners and incumbents, institutional demand functions must be estimated. I follow Koijen and Yogo (2019) to estimate these demand functions, with some model specific variations. I estimate institutional demand functions for each institution in each period. In order to obtain aggregate market demand, I also estimate non-institutional asset demand as well. The sum of the institutional and non-institutional demand is by definition the entire market (incumbent) demand. While this incumbent demand results in identical mispricing at incumbents-only prices, when prices move away from incumbents-only prices, these estimated demand functions differ from the CARA incumbents' demand.

I consider 22 variations of the model—all combinations of two types of incumbents and eleven types of learners. The eleven learners I use in the paper are as follows: 1) market indexer, 2) Bayesian learner, 3) learners who learn from the true exogenous dividend DGP using maximum likelihood, 4) learners who learn from the true exogenous dividend DGP using maximum a posteriori (MAP), 5) learners who know the true DGP parameters referred to as asymptotic DGP learners, 6) portfolio optimizers that use a Brandt et al. (2009) estimator, 7) portfolio optimizers that use shrinking with a Kozak et al. (2020) estimator, 8) learners who plug in asymptotic Brandt et al. (2009) parameter values, 9) learners who invest simply according to a DeMiguel et al. (2009) $1/N$ rule, 10) learners who use a random forest method to form portfolios, and 11) learners who use a neural network.

1.6 Related Literature

There is a rich literature on rational traders destabilizing asset prices. For example, Hart and Kreps (1986) argue that rational speculators can easily make prices less stable. In their model, these speculators rationally increase price volatility due to potential small probabilities of high returns. They show that speculators' demand rarely "look[s] sensible ex post."

Similarly, Stein (1987) shows a model with information asymmetries where relatively uninformed arbitrageurs enter the market, and their actions increase the risk for other market participants and subsequently destabilize prices and decrease welfare. Hong and Stein (1999) also show how sophisticated investors who trade against the market's tendency to underreact to news can still destabilize prices.

Stein (2009) argues that sophisticated professional traders can destabilize asset prices through two different ways. First, arbitrageurs who are unsure of both fundamental values and how many peers are involved in a trade can end up destabilizing prices. Second, arbitrageurs can cause fire-sale externalities with leverage. Stein argued that "[a]rbitrageurs do not base their demand on an independent estimate of fundamental value. As a result, *their demand for an asset may be a nondecreasing function of the asset's price.* Strategies of this type are common in practice, and include many in which demand is independent of price." Thus Stein also emphasizes how traders' reaction to price changes can play a key role in mispricing outcomes. DeLong et al. (1990) also argue that rational speculative strategies that are an increasing function of an asset's price can increase price fluctuations. This paper contributes to this literature in five important ways.

First, this paper shows that in a world with many assets, adapting to changes in the data generating process of returns due to learner demand is difficult. The Israel et al. (2020) logic above argues that since many quantitative models do well at estimating historical alpha, they should push prices to be more efficient. This paper emphasizes that estimating alpha even as price dynamics shift is difficult relative to just estimating historical alpha. The key finding of this paper is that learners who can do well at estimating historical alpha, but struggle to adapt to new price dynamics, can make mispricing worse.

Second, this paper shows that canonical portfolio optimization methods—such as Kozak et al. (2020) and Brandt et al. (2009)—tend to struggle to adapt to changing return dynamics. These portfolio optimization methods learn from past returns, but are insufficiently price reactive as the data generating process of returns changes due to their own demand. In the model, even some learners who learn from the exogenous dividend process struggle to adapt properly to changing prices.

Third, Stein (2009) emphasizes arbitrageurs can make "unanchored" trades which exacerbate mispricing; however this paper shows that even relatively anchored trades—such as investing in value stocks—can exacerbate mispricing. This occurs simply because learners invest a fixed amount of capital into the value portfolio for example, and this amount does not vary even as the value return premium shifts. Thus while portfolio weights of assets within the value portfolio change, overall capital invested in the value portfolio does not shift. This results in insufficiently price reactive learner demand, resulting in long-short value fund underperformance.

Fourth, the learner, which plays a similar role to the arbitrageurs and speculators in the above models, simply estimates a model and uses the model to make portfolio choice decisions. In other words, the parameter estimates are plugged in as a demand function, and this replaces the unanchored trade, information asymmetries, and other economic mechanisms used in the papers listed above. Thus all mispricing outcomes are driven simply by estimation error and model error, rather than the economic mechanism discussed in these other papers.

Fifth, these learners can only invest capital allocated to them by rational mean-variance investors who know the true portfolio returns of the learners and incumbents. These ultra-rational well-informed investors tend to limit increases in mispricing relative to uninformed investors. Even with these investors, price efficiency can easily decrease.

This paper is also related to Martin and Nagel (2019), who show that in a high dimensional setting where many variables predict fundamentals, in-sample alpha will remain due to estimation error. While I similarly show that substantial amounts of alpha remain in-sample, I also show that alpha can actually increase due to estimation error and model error. I also highlight how the price elasticity of incumbent and learner asset managers plays an important role in mispricing outcomes.

1.7 Dissertation Outline

The paper proceeds as follows. Chapter 2 gives an extremely simple version of the model to describe the basic economic intuition. Chapter 3 shows the closed-form asymptotic results of the model, where mispricing remains in equilibrium even with an infinite learning horizon. Chapter 4 gives an overview of the full model. I outline how the dividend process is estimated in Chapter 5. Chapter 6 outlines the demand function estimation methodology. The various kinds of learners inserted in the model are described in Chapter 7. The model simulation results are described in Chapter 8. Finally, Chapter 9 concludes.

CHAPTER 2

BASIC ECONOMICS OF THE MODEL

The basic intuition that smart quantitative learners will eliminate all mispricing by trading against it pervades the academic finance literature. This powerful result is deeply intuitive, and has been used to argue that deviation from asset pricing models, such as the CAPM, must be driven by risk rather than pricing errors. If the deviations are pricing errors, so the logic goes, then these errors would be eliminated.

This paper argues that mathematically this is unlikely to be true in a world with a lot of data but limited time series from which to learn. Fundamentally, judging deviations from asset pricing models requires *a time series* of returns. This fundamentally makes it difficult to assess the *contemporaneous* deviation from a model, since mispricing is only viewed primarily from the lens of *historical* returns.

In this simple model, mispricing can persist *due purely to estimation error and misspecification error of the model* that investors use. The misspecification comes from learning from previous returns, rather than the true dividend process. While learning from the true dividend process make sense, in reality much of the variation in asset prices is actually driven by variation in discount rates (see Cochrane (2011)). Thus learning based on past returns is reasonable as long as the relationship between state variables and returns is fixed. If the investor is small and not affecting prices, and no other secular shifts are causing the relationship between returns and state variables to shift, then this is reasonable. However, if other investors are employing similar learning strategies, then the relationship between state variables and returns can shift due to crowded trading. If mispricing is estimated based on historical data, the contemporaneous affects of crowded trading will likely not be priced in appropriately.

In this model, pricing errors can persist in two ways:

1. The magnitude of historical alpha can be underestimated, causing insufficient arbi-

trage capital deployment, which means that alpha remains and keeps its original sign (positive or negative).

2. Historical alpha is estimated, arbitrage capital is deployed, but investors fail to understand how many other investors are deploying similar strategies. Since mispricing is primarily viewed through the lens of *historical returns*, the arbitrageurs fail to be sufficiently price responsive, and the sign of alpha flips. Thus mispricing remains, but the sign of alpha has flipped. This can be extreme enough to even cause the total magnitude of mispricing to increase (but with a different sign).

In this section, an extremely simplified version of model is presented to give the basic economic intuition of the paper.

2.1 Simplified Model

There are two types of investors: incumbents and learners. In this section, *there is no learning.* We simply plug in different values of perceived historical alpha and price reaction functions to characterize the equilibrium.

There is only $N = 1$ asset. In this section, the model is static: there is only a single period. There is an exogenously given mass of θ_L learners, and $1 - \theta_L$ incumbent investors. We endogenize θ later on in the paper. All of the learners and incumbents are price takers. We assume they have CARA utility demand with risk aversion coefficient $\gamma = 1$. The volatility of the single asset is set to one, and the risk-free rate is set to zero. The dollar or share return of the asset is simply $r = d - p$, where d is the payout of the asset, and p is the price paid for the asset.

Following Kozak et al. (2018), we add a mistake term to incumbent demand, to induce mispricing in an equilibrum with only incumbents ($\theta_L = 0$). This mistake is a simple reduce-form way to induce mispricing among incumbents. Kozak et al. (2018) interpret this mistake term, denoted as δ, as driven by sentiment. We take a more agnostic interpretation here. It could be driven by sentiment, preferences, beliefs, or a variety of other economic mechanisms. In summary, we insert this δ term into demand to cause mispricing, in order for the learners to potentially have a chance to enter and eliminate mispricing.

Incumbent demand is simply

$$x_I = \mathbb{E}[d] - p + \delta$$

Supply is normalized to unity, which means that equilibrium with only incumbents, referred to as the incumbents-only equilibrium, is given by

$$1 = \mathbb{E}[d] - p + \delta$$

Thus the incumbents-only price is $\bar{p} = \mathbb{E}[d] - 1 + \delta$. Mispricing is $\alpha = \mathbb{E}[d] - p - 1$, and thus incumbents-only mispricing is

$$\bar{\alpha} = -\delta$$

Thus incumbents-only mispricing is nonzero if and only if the incumbent mistake term, δ, is nonzero.

Assume learner demand is

$$x_L = \mathbb{E}[d] - \bar{p} + J_L(p - \bar{p})$$

If the learner appropriately adapts to changing prices, then $J_L = -1$. Assume, however, that the learner insufficiently reacts to changing prices but still has downward sloping demand, i.e. assume that $-1 \leq J_L \leq 0$. Notice that the learner demand intercept term is error-free. Thus the learner here perfectly knows the mean-variance efficient weights at the incumbents-only prices. In other words, the learner knows the answer to Q1 above without error.

This learner demand function can be rewritten as

$$x_L = \bar{\alpha} + 1 + J_L(p - \bar{p})$$

Written this way, the learner demand is reminiscent of the story of Pecunia above. Pecunia estimates the historical α, and adds some adjustment for changing prices, and invests accordingly. While there is no learning in this simple example, the intuition of investing based on historical alpha still shines through.

The combined equilibrium is described by

$$1 = (1 - \theta_L)(\mathbb{E}[d] - p + \delta) + \theta_L(\mathbb{E}[d] - \bar{p} + J_L(p - \bar{p}))$$

Thus the combined equilibrium price is

$$p = \mathbb{E}[d] + \delta - 1 - \frac{\theta_L \delta}{1 - (1 + J_L)\theta_L}$$

and the combined equilibrium mispricing is

$$\alpha = \frac{\theta_L \delta}{1 - (1 + J_L)\theta_L} - \delta$$

Thus if the incumbent mistake δ is nonzero, but learners capture the entire market ($\theta_L = 1$) and the learners perfectly adapt to changing prices ($J_L = -1$), then the mispricing is zero.

If δ is nonzero and

$$J_L > \frac{2 - 3\theta_L}{2\theta_L}$$

then $\bar{\alpha}$ has the opposite sign of α and $|\alpha|$ is *larger* than $|\bar{\alpha}|$. Thus the magnitude of mispricing actually *increases* when the learners enter.

For example, assume $\bar{\alpha}$ is positive in the incumbents-only equilibrium ($\delta < 0$), and thus the asset is underpriced. Then in this case, α is negative when the learners enter and the asset becomes overpriced.

In summary, in this example, an asset that is initially underpriced (overpriced) becomes overpriced (underpriced) to such an extent that overall mispricing actually increases. This is due to an overshooting effect, where the learner fails to appropriately react to changing prices.

Notice that in this case, the incumbents are CARA incumbents as discussed above. In other words, the incumbents are *not* too price inelastic relatively to utility-based demand functions. Thus if learners have downward sloping demand curves and mispricing increases, it must be the case that $\theta_L > 2/3$. In other words, for mispricing to increase, the learners must capture a substantial amount of the market. However, with more inelastic incumbent demand, learners can capture a smaller fraction of the market and still increase mispricing. In other words, inelastic incumbent demand makes it much more likely for learners to increase mispricing upon entry. To see this, assume that the incumbents-only price is the same as above, but assume that the incumbent demand is

$$x_I = \mathbb{E}[d] - \bar{p} + \delta + J_I(p - \bar{p})$$

This demand function is analogous to the estimated demand function for the incumbents used throughout the paper, in that the incumbent demand is allowed to be more inelastic than utility-based demand. Here, of course, J_I is not estimated, just allowed to differ from -1. Assume, like the learner, that $-1 \leq J_I \leq 0$.

Repeating the above analysis yields the equilibrium price

$$p = \mathbb{E}[d] + \delta - 1 + \frac{\theta_L \delta}{(1 - \theta_L)J_I + \theta_L J_L}$$

and the equilibrium mispricing is

$$\alpha = -\frac{\theta_L \delta}{(1-\theta_L)J_I + \theta_L J_L} - \delta$$

Thus in this case, if

$$J_L > -\frac{2(1-\theta_L)J_I + \theta_L}{2\theta_L}$$

then mispricing increases when the learners enter. Thus, if incumbent demand is relatively inelastic, then the learners can capture a smaller share of the market and still increase mispricing even with downward sloping demand. In the extreme case where the incumbent is perfectly price inelastic (i.e. $J_I = 0$) and $-1/2 < J_L < 0$, then mispricing increases for any $\theta_L > 0$. In other words, in this case, the learners can capture only a tiny share of the market and still increase mispricing overall.

In summary, if incumbent demand is relatively inelastic, a larger range of parameters can lead to increases in mispricing when the learners enter.

2.2 Canonical Portfolio Choice Methods Generate Price Nonreactive Demand

In this subsection, I consider how the canonical portfolio choice generates a simple demand function for an asset, and how price reactive this demand function is for the asset.

2.2.1 Canonical Portfolio Choice Decision

In the previous example, J_L represents how the learners react to changes in prices. I showed that if J_L is sufficiently larger than -1 and the learners capture enough investment then an asset price can overshoot the efficient price to such a degree that mispricing can worsen. Why would J_L be larger than -1? In other words, why would learner demand react so little to changing prices? This subsection shows how this can occur using standard portfolio choice tools. As the above example shows, this insufficiently price reactive learner demand is the key to the overshooting economic mechanism of this paper.

Consider, for the sake of concreteness, a learner who optimizes investment in the market portfolio and a long-short value portfolio. I consider a single asset, which is contained in both the market fund and the value fund. In other words, I consider a market portfolio (composed of many assets), a long-short value portfolio (composed of many assets), and one asset in particular that is potentially in both the market and value portfolios.

As discussed above, the number of shares outstanding has no economic significance, and thus I assume that there is one share outstanding for the asset considered here[1]. Just as above, let p denote the price of this asset, which is the same as market equity if this asset is a stock because there is only one share outstanding.

Importantly, in this example, each learner invests a fixed fraction w of her money into the value fund and $1 - w$ into the market. The variable w is *chosen by considering optimal portfolio weights in the data* and critically *does not vary as prices change.* Brandt et al. (2009) and Kozak et al. (2020) are portfolio optimization methods where the *investment weights for each portfolio do not vary as prices change.* This is the key reason that portfolio optimization techniques lead to insufficiently price reactive demand.

The weight of the asset in the market portfolio is p/A, where A is the value of the market—i.e. A is the sum of the prices of all assets. Assume that the learner has A_L dollars to invest. Assume that the asset's weight in the value fund is $u(p)p/A$, where $u(p)$ is a function that equals either 1, 0, or -1 depending on whether the asset is in the long portion of the portfolio sort (i.e. $u(p) = 1$), the asset has no weight in the portfolio (i.e. $u(p) = 0$), or the asset is in the short portion of the portfolio sort (i.e. $u(p) = -1$). Notice that this portfolio is a difference of value-weighted portfolios. In other words, this is a standard

1. Investors can invest with fractional shares.

valued-weighted long-short portfolio.

In this case, this learner invests the following amount of dollars into the asset:

$$(1-w)A_L\frac{p}{A} + wA_L\frac{u(p)p}{A}$$

where the first term represents the dollar amount invested in this asset *because the asset is in the market portfolio* and the second term represents the dollar amount investment in this asset *because the asset is in the value portfolio.*

Demand in this paper is terms of the fraction of the share of the asset, and thus it follows that the learner demand for this asset is

$$\theta_L x_L = \frac{1}{p}\left((1-w)A_L\frac{p}{A} + wA_L\frac{u(p)p}{A}\right) = (1-w)\frac{A_L}{A} + wA_L\frac{u(p)}{A}$$

This can be further simplified by assuming that the approximation $\theta_L = A_L/A$ holds, which means learner demand is

$$x_L = (1-w) + wu(p)$$

Notice that except for a few portfolio sort break points where the derivative is not defined, $\frac{du}{dp} = 0$. Thus except for these break points, $\frac{dx_L}{dp} = 0$. In the last section $J_L = \frac{dx_L}{dp}$. Thus, except for a few break points, this demand function is *completely price nonreactive.*

It is important to note that if this strategy instead was a combination of investing in the market and a profitability portfolio for example, then $u(p)$ would not vary with prices, because market equity does not affect profitability portfolio sorts. In other words, in this example, this demand function is *completely price nonreactive* everywhere. Thus, this key result, that the classic quantitative portfolio choice strategy results in completely price non-reactive demand almost everywhere holds for any combination of the market portfolio and long-short value-weight portfolios.

This result—that canonical portfolio choice methods result in completely price nonreactive demand almost everywhere—is striking. To re-emphasize, if w varies as p changes, then this does not hold. However, Brandt et al. (2009) and Kozak et al. (2020), as well as classic portfolio choice methods, assume fixed weights in portfolios. Thus even if the portfolios themselves are valued-weighted, investing with fixed weights in each portfolio results in completely price nonreactive demand almost everywhere.

2.2.2 Continuous Characteristic Weighting Instead of Discontinuous Jumps

The above u function is characterized by a few discontinuous jumps. Suppose instead, that this u function is linearized. Assume that p is known to fall between some small price p_L—a low price that makes the asset an extreme value stock—and some large price p_U—a high price that makes the asset an extreme growth stock. Assume $p_U - p_L$ is large, which means that the plausible range for p is large and the difference between a large growth price and a small value price is large. Define u to be

$$u(p) = 2\frac{p_U - p}{p_U - p_L} - 1$$

This linearized version of u represents a portfolio that is *characteristic-weighted*, like Brandt et al. (2009) and Kozak et al. (2020). In other words, the asset in the portfolio is weighted by the degree to which the asset is a value stock or a growth stock.

In this case,

$$\frac{dx_L}{dp} = -\frac{2w}{p_U - p_L}$$

Since the growth price p_U is much larger than the value price p_L, this quantity is small. Thus, even in this linearized case, learner demand is still relatively price nonreactive.

Thus the key intuition of this subsection is *not* the fact that portfolio sorts generate price nonreactive demand. The key intuition of this subsection is that investing a fixed fraction w

into the value portfolio *based on historical data and not based on the current prices of value and growth stocks leads to price nonreactive demand*. It's the fixed w that this example is meant to emphasize.

2.2.3 Summary of Demand Generated from Canonical Portfolio Choice Methods

In conclusion, canonical quantitative portfolio choice methods, which include Brandt et al. (2009) and Kozak et al. (2020), tend to produce relatively price nonreactive demand because these canonical methods choose constant weights on various portfolios.

CHAPTER 3

HIGH DIMENSIONAL PROBLEM

In the very basic model above, learning about the mispricing of a single asset appears simple and easy. Readers may understandably think that persistent mispricing cannot actually be occurring in financial markets since learning about just a few parameters cannot take long. In this chapter, we imagine a world with many assets, or alternatively many asset characteristics. We show that in a world with many assets or many asset characteristics, mispricing can remain in the model *even with an infinite time horizon.* These asymptotic results are key to understanding the existence of persistent mispricing highlighted by this paper.

We first review the basics of CARA utility demand and the resulting equilibrium when everything is known. Then we discuss CARA utility demand with a Bayesian investors who learn from the exogenous dividend process. In this model, if the number of assets is a similar order of magnitude to the number of time periods to learn from, mispricing remains in the model. Finally, we explore a learner who learns from returns instead of the true dividend process. We show that mispricing remains in this model as well.

In this chapter, we eject the incumbents from the model and retain the learners only. Thus we consider the learner-only equilibrium, with $\theta_L = 1$. This lends tractability to the results, giving us closed-form solutions. The amount of learners in the market, θ_L, is endogenized in the next section.

3.1 Review of CARA Model with Full Knowledge of Parameters

There are N assets, and d_t is the N dimensional column vector of dividends in period t. The N dimensional column vector p_t contains the prices of the assets in period t. I define $P_t = \text{diag}(p_t)$, which in words means that P_t is the $N \times N$ dimensional matrix with zeros on the off-diagonal and p_t on the diagonal. Share excess returns, dollar excess returns, or excess dollar payoffs,[1] are defined as $r_{t+1} = d_{t+1} - (1 + r_f)p_t$ where r_f is an exogenous risk free rate. The asset payoffs are driven by the dividends, and there is no price of the asset in the following period. This allows the covariance matrix of returns to be completely exogenous and not a function of endogenous prices in the next period. This makes the model much more tractable. The conditional covariance matrix of returns is denoted as $\Lambda_t = \mathrm{Var}_t[d_{t+1}] = \mathrm{Var}_t[r_{t+1}]$. Dividends follow a multivariate normal distribution. Let ι be an N dimensional column vector of ones. We normalize supply to be unity for each risky asset, so that quantities are normalized to be a fraction of ownership of the entire asset.

Table B.1 gives the notation and definitions for this chapter and throughout the paper. This table shows the definitions of variables both in terms of dollar returns and typical returns. Many papers with methods used in this paper, such as Brandt et al. (2009), Kozak et al. (2020), and Koijen and Yogo (2019) use asset excess returns and portfolio weights, but canonical asset pricing models, which Kozak et al. (2018) is a variation, typically use dollar excess returns and asset shares.

There is a unit mass of identical investors. Each is a price taker. Demand is given as

$$x_t = \frac{1}{\gamma}\Lambda^{-1}(\mathbb{E}_t[d_{t+1}] - (1 + r_f)p_t)$$

Since supply is normalized to unity, equilibrium is given by solving the price p_t that solves:

$$\iota = \frac{1}{\gamma}\Lambda^{-1}(\mathbb{E}_t[d_{t+1}] - (1 + r_f)p_t)$$

Thus the equilibrium price is

$$p_t = \frac{1}{1 + r_f}\left(\mathbb{E}_t[d_{t+1}] - \gamma\Lambda\iota\right) \tag{3.1}$$

1. The terms share excess returns, dollar excess returns, and excess dollar payoffs are all used in the literature. This is standard to work with returns in this form in a model rather than percentage returns.

As shown in Table B.1, the CAPM α_t is given as

$$\alpha_t = \underbrace{\mathbb{E}_t[r_{t+1}]}_{\text{excess return}} - \underbrace{\frac{\Lambda\iota}{\iota'\Lambda\iota}}_{\beta} \underbrace{\iota'\mathbb{E}_t[r_{t+1}]}_{\text{market excess return}} = \left(I - \frac{\Lambda\iota\iota'}{\iota'\Lambda\iota}\right)(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t) \tag{3.2}$$

Plugging equation (3.1) into equation (3.2), we find that $\alpha_t = 0$ in this equilibrium. In fact, we find that the *only price vector that eliminates mispricing* α_t is a price in the form of

$$p_t = \frac{1}{1+r_f}\left(\mathbb{E}_t[d_{t+1}] - c\Lambda\iota\right) \tag{3.3}$$

for any constant c.

This is a canonical result, which we only review to outline the notation and to contrast to the two models below.

3.2 Bayesian Learning

Assume that $\mu \equiv \mathbb{E}_t[d_{t+1}]$ is constant. Thus we can write $d_t \sim N(\mu, \Lambda)$. Now our investors in this model are Bayesian learners who knows Λ but have the prior:

$$\mu \sim N(\mu_0, \Lambda_0)$$

After observing T periods of dividends, the learners have the following posterior predictive:

$$d_{T+1} \sim N\left((\Lambda_0^{-1} + T\Lambda^{-1})^{-1}\left(\Lambda_0^{-1}\mu_0 + T\Lambda^{-1}\bar{d}_T\right),\ (\Lambda_0^{-1} + T\Lambda^{-1})^{-1} + \Lambda\right)$$

where

$$\bar{d}_T = \frac{1}{T}\sum_{t=1}^{T} d_t$$

Thus their belief, referred to as the posterior predictive, about dollar returns are:

$$r_{T+1} \sim N\left((\Lambda_0^{-1} + T\Lambda^{-1})^{-1} \left(\Lambda_0^{-1}\mu_0 + T\Lambda^{-1}\bar{d}_T \right) - (1+r_f)p_T,\ (\Lambda_0^{-1} + T\Lambda^{-1})^{-1} + \Lambda \right)$$

Thus the learner's time T demand is

$$x = \frac{1}{\gamma}\left((\Lambda_0^{-1} + T\Lambda^{-1})^{-1} + \Lambda \right)^{-1} \left((\Lambda_0^{-1} + T\Lambda^{-1})^{-1} \left(\Lambda_0^{-1}\mu_0 + T\Lambda^{-1}\bar{d}_T \right) - (1+r_f)p_T \right)$$

Since supply is unity, the equilibrium is the price that satisfies:

$$\iota = \frac{1}{\gamma}\left((\Lambda_0^{-1} + T\Lambda^{-1})^{-1} + \Lambda \right)^{-1} \left((\Lambda_0^{-1} + T\Lambda^{-1})^{-1} \left(\Lambda_0^{-1}\mu_0 + T\Lambda^{-1}\bar{d}_T \right) - (1+r_f)p_T \right)$$

Which means that the equilibrium price is

$$p_T = \frac{1}{1+r_f}\left[(\Lambda_0^{-1} + T\Lambda^{-1})^{-1} \left(\Lambda_0^{-1}\mu_0 + T\Lambda^{-1}\bar{d}_T \right) - \gamma\left((\Lambda_0^{-1} + T\Lambda^{-1})^{-1} + \Lambda \right)\iota \right] \quad (3.4)$$

Assume for simplicity that $\Lambda_0 = \kappa^{-1}\Lambda$ for some scalar $\kappa > 0$ and that $\mu_0 = \mu$. That is, we assume that investors' prior about the mean of beliefs is proportional to the actual covariance of returns Λ, and that their belief about the mean of returns is correct. Plugging in these prior beliefs into equation (3.4) for the price, and then plugging this price into equation (3.2) yields

$$\alpha_T = \left(I - \frac{\Lambda\iota\iota'}{\iota'\Lambda\iota} \right) \left(\frac{T}{\kappa+T}\mu - \frac{T}{\kappa+T}\bar{d}_T + \gamma\frac{1+\kappa+T}{\kappa+T}\Lambda\iota \right) \quad (3.5)$$

Now, α_T is an N dimensional column vector, but $\alpha_T'\Lambda^{-1}\alpha_T$ is the standard measure of alpha across all assets. This is the key component in the Gibbons et al. (1989) test statistic. It is also the Hansen and Jagannathan (1991) distance in this setting. Finally, we have the classic equation:

$$\underbrace{\mathbb{E}_t[r_{t+1}]'\Lambda^{-1}\mathbb{E}_t[r_{t+1}]}_{\text{MVE Sharpe Ratio Squared}} = \underbrace{\alpha_T'\Lambda^{-1}\alpha_T}_{\text{Information Ratio Squared}} + \underbrace{\frac{(\iota'\mathbb{E}_t[r_{t+1}])^2}{\iota'\Lambda\iota}}_{\text{Market Sharpe Ratio Squared}}$$

Thus the total amount of mispricing, $\alpha_T'\Lambda^{-1}\alpha_T$, is bounded above by the highest achievable Sharpe ratio squared, $\mathbb{E}_t[r_{t+1}]'\Lambda^{-1}\mathbb{E}_t[r_{t+1}]$.

The proposition below shows that expected mispricing in this model is a function of both the number of time periods the Bayesian learning has witnessed, T, as well as the number of assets N.

Proposition 1.1

Under the assumptions of the model above, we have

$$E[\alpha_T'\Lambda^{-1}\alpha_T] = \left(\frac{T}{\kappa + T}\right)^2 \left(\frac{N-1}{T}\right) \tag{3.6}$$

The proof is in the appendix below.

Interestingly, the mispricing *grows* with the number of assets. Thus, in the earlier model with a single asset, the expected mispricing dissipates to zero quickly. However, with many assets it takes longer.

In fact, if the number of assets is a similar order of magnitude to the number of time periods, the expected mispricing *fails to dissipate*. It remains in the model asymptotically. This is the heart of the proposition below.

Proposition 1.2

$$\underset{\substack{N/T=c \\ N,T\to\infty}}{\text{plim}} \; \alpha_T'\Lambda^{-1}\alpha_T = c$$

The proof is in the appendix below.

In this context, there is an important alternative interpretation for N. Instead of thinking of N as the number of assets, we can think of N as the number of characteristic-sorted portfolios. This means that if the number of characteristics is large relative to the number of time periods, mispricing should remain. This is similar to the findings from Martin and Nagel (2019).

It is important to remember that in this model, the investors are entirely rational. Mispricing remains simply because the problem is difficult to estimate, not because of any risk or behavioral failure.

In the model below, we explore the idea of characteristic-weighted portfolios, used extensively in the literature, such as in Brandt et al. (2009) and Kozak et al. (2020).

3.3 Learning from Portfolio Returns

In this model, it is still the case that $d_t \sim N(\mu, \Lambda)$, excess returns are $r_{t+1} = d_{t+1} - (1+r_f)p_t$, and supply for each asset is normalized to one.

At the heart of many asset pricing paper is the set of asset characteristics. For example, the Fama and French (2015) five factor model uses four characteristics (market equity, book to market ratio, profitability, and investment) to form five portfolios (market, size, and value, profitability, and investment portfolios).

Let Z_t be a $N \times K$ matrix of portfolio weights based on K asset characteristics. All asset characteristics in Z_t are known to investors at time t. For example, in the Fama and French (2015) five factor model, $K = 5$. In other words, in the Fama and French (2015) five factor model, each of the five columns of Z_t represent N portfolio weights of each of the five factors.

Portfolio excess returns are written as the K dimensional column vector:

$$F_{t+1} = Z_t' r_{t+1}$$

Importantly, in practice, some of the columns *depend on the price* of the assets, while other columns of Z_t *do not depend on the price* of the assets. For example, the Fama and French (2015) value and size portfolios depend on the price of the assets. If the price of the asset changes, the asset may move to another value or size bin, changing its portfolio weights. However, other portfolios in the Fama and French (2015) five factor model do not change based on the price of the asset, such as profitability and investment. These portfolio weights are formed based only on accounting variables, and not a function of the current price.

Note that here we are still using dollar or share returns, not percentage returns. Again, Table B.1 outlines how to convert between asset shares and portfolio weights, and between share returns and percentage returns. One difference, for example, is that market portfolio weights are proportional to market equity, but market portfolio shares compose a portfolio that just holds an equal *share* of each asset.

For this model, consider only asset characteristics that are either not a function of price, or a *linear* function of price. Now the Fama and French (2015) portfolios that depend on price are nonlinear functions of prices. We can think of linear approximations of these portfolio weights, as discussed in the previous chapter. We denote the portfolio weight of asset i at time t in portfolio k to be $a_{i,k} + m_{i,k}p_{i,t}$. It could of course be the case where $m_{i,k} = 0$ (like in the profitability and investment portfolios). For simplicity, we assume that the intercept terms, $a_{i,k}$, and the intercept terms, $m_{i,k}$, are constant through time.

We can write Z_t as

$$Z_t = \begin{bmatrix} a_{1,1} + m_{1,1}p_{1,t} & a_{1,2} + m_{1,2}p_{1,t} & \cdots & a_{1,K} + m_{1,K}p_{1,t} \\ a_{2,1} + m_{2,1}p_{2,t} & a_{2,2} + m_{2,2}p_{2,t} & \cdots & a_{2,K} + m_{2,K}p_{2,t} \\ \vdots & \vdots & \ddots & \vdots \\ a_{N,1} + m_{N,1}p_{N,t} & a_{N,2} + m_{N,2}p_{N,t} & \cdots & a_{N,K} + m_{N,K}p_{N,t} \end{bmatrix}$$

We can also write "intercept" and "slope" matrices:

$$Z_I = \begin{bmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,K} \\ a_{2,1} & a_{2,2} & \dots & a_{2,K} \\ \vdots & \vdots & \ddots & \vdots \\ a_{N,1} & a_{N,2} & \dots & a_{N,K} \end{bmatrix}, \quad Z_S = \begin{bmatrix} m_{1,1} & m_{1,2} & \dots & m_{1,K} \\ m_{2,1} & m_{2,2} & \dots & m_{2,K} \\ \vdots & \vdots & \ddots & \vdots \\ m_{N,1} & m_{N,2} & \dots & m_{N,K} \end{bmatrix}$$

This allows us to write

$$Z_t = Z_I + P_t Z_S$$

where $P_t = \text{diag}(p_t)$.

This learner learns from *returns* and not *dividends*. In reality, with equities, learning from just dividends misses much of the variation due to discount rate variation. Thus much of the asset pricing literature uses historical data to learn from returns, such as Fama and French (1993), Fama and French (2015), Kelly et al. (2019), Kozak et al. (2020), and DeMiguel et al. (2009).

The key assumption that these investors make is that, as in Brandt et al. (2009) and Kozak et al. (2020), mean-variance efficient portfolio weights are linear in asset characteristics:

$$\Lambda^{-1}\mathbb{E}[r_{t+1}|Z_t] = Z_t b \tag{3.7}$$

where b is a constant $K \times 1$ vector of portfolio weights of characteristic-weighted portfolios. The learner does not know b and uses historical returns to choose b. Thus the learner updates their view of b through time. We will replace b with b_t, which is estimated from historical returns *but not a function of* p_t. The learner knows Λ.

It is important to note that the learner's model is potentially misspecified. This is a key component of this paper. Learning about historical returns often requires the relationship between returns and some state variables to be fixed. However, if investors are making similar assumptions, and many of them together are moving prices, then the relationship between the state variables and returns can shift. Thus we allow investors in this model to pursue classic portfolio choice investment decisions implied by well-known factor models, such as Fama and French (1993) and Fama and French (2015).

We can multiply both sides of equation (3.7) by $Z_t'\Lambda$, which yields

$$Z_t'\mathbb{E}[r_{t+1}|Z_t] = Z_t'\Lambda Z_t b$$

The learner replaces $Z_t'\mathbb{E}[r_{t+1}|Z_t]$ and $Z_t'\Lambda Z_t$ with historical means. Priors can be included, but we want to consider mispricing when the number of time periods that the learner has experienced, T, is relatively large and the prior has a negligible impact. Thus we exclude the prior for convenience. When replaced by historical averages, and when b is replaced with b_T (to indicate the learners' view of b at time T) we have

$$\frac{1}{T}\sum_{t=0}^{T-1} Z_t' r_{t+1} = \frac{1}{T}\sum_{t=0}^{T-1} Z_t'\Lambda Z_t b_T$$

The learner solves for b_T to get:

$$b_T = \left(\frac{1}{T}\sum_{t=0}^{T-1} Z_t'\Lambda Z_t\right)^{-1} \frac{1}{T}\sum_{t=0}^{T-1} Z_t' r_{t+1}$$

In the section above, all mispricing was due to estimation error. The Bayesian learner couldn't eliminate all mispricing when the number of assets was large just because of the high-dimensional nature of the problem. Here, the learner will not eliminate all estimation error because of estimation error and misspecification error. This is driven the the distribution of b_T. If the distribution of b_T has a large variance, then it will be difficult for mispricing to be eliminated, just as the variance of $\bar{d}_T$ above caused the mispricing to remain in equilibrium. Here we make a strong assumption, but conservative in that it reduces in the variance of b_T.

Our assumption is that the prices historically have reach a long-run steady state such that

$$p = p_{T-1} = p_{T-2} = ...$$

Thus if we denote $P = \text{diag}(p)$ and $\bar{Z} = Z_I + PZ_S$, then we can rewrite b_T as

$$b_T = \left(\bar{Z}'\Lambda\bar{Z}\right)^{-1}\bar{Z}'\frac{1}{T}\sum_{t=0}^{T-1} r_{t+1}$$

We assume they have CARA utility demand, and using equation (3.7), we can write demand as

$$x = \frac{1}{\gamma}Z_T b_T = \frac{1}{\gamma}(Z_I + P_T Z_S)b_T$$

Thus equilibrium is given by the price that solves:

$$\iota = \frac{1}{\gamma}(Z_I + P_T Z_S)b_T$$

Recall above that prices in the form of equation (3.1) have zero alpha, and that only prices in the form of equation (3.3) have zero alpha. It is convenient for us to use a different wedge rather than α_T when computing mispricing here, since the distribution of standard mispricing $\alpha_T'\Lambda^{-1}\alpha_T$ is relatively intractible mathematically. We define our pricing wedge in the proposition below.

Proposition 2.1 Assume that $Z_S b_T$ do not have any elements that equal exactly zero. Prices equal the zero-alpha prices of equation (3.1) if and only if $l_t'\Lambda l_t = 0$, where

$$l_t = -(1 + r_f)\gamma\iota + \tilde{Z}b_t$$

and

$$\tilde{Z} = (1 + r_f)Z_I + (U - \gamma\lambda)Z_S, \quad U = \text{diag}(\mu), \quad \lambda = \text{diag}(\Lambda\iota)$$

The proof is in the appendix below.

We can use this wedge to determine if mispricing remains in the model in expectation, which is the content of the proposition below:

Proposition 2.2 If the assumptions above hold, then

$$\mathbb{E}[l_t'\Lambda l_t] = \frac{K}{T}\bar{\phi} + N^2(1+r_f)^2\gamma^2\bar{\chi} + K\bar{\psi} + 2(1+r_f)\gamma K\bar{\omega}$$

where

- ϕ_i is the element in the i^{th} row and i^{th} column of $\tilde{Z}'\Lambda\tilde{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}$.
- $\chi_{i,j}$ is the element in the i^{th} row and j^{th} column of Λ.
- ψ_i is the element in the i^{th} row and i^{th} column of $\bar{Z}'((1+r_f)p-\mu)((1+r_f)p-\mu)'\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\tilde{Z}'\Lambda\tilde{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}$.
- ω_i is the element in the i^{th} row and i^{th} column of $\tilde{Z}'\Lambda\iota((1+r_f)p-\mu)'\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}$.

and

$$\bar{\phi} = \frac{1}{K}\sum_{i=1}^{K}\phi_i,\ \ \bar{\chi} = \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N}\chi_{i,j},\ \ \bar{\psi} = \frac{1}{K}\sum_{i=1}^{K}\psi_i,\ \ \bar{\omega} = \frac{1}{K}\sum_{i=1}^{K}\omega_i$$

The proof is in the appendix below.

Thus the mispricing that remains in expectation, at least according to our measure of mispricing, is a function of T, K, and N. The proposition below simplifies the math considerably with an additional assumption.

Proposition 2.3 If the historical price p equals the zero alpha price of equation (3.1), and we define

- $LL' = (\bar{Z}'\Lambda\bar{Z})^{-1}$ to be the cholesky decomposition of $(\bar{Z}'\Lambda\bar{Z})^{-1}$
- G to be the vector $G = L'\bar{Z}'\Lambda\iota$

- $\bar{g} = \frac{1}{K}\sum_{i=1}^{K} G_i^2$

then

$$\mathbb{E}[l_t' \Lambda l_t] = \underbrace{\frac{K}{T}(1+r_f)^2}_{\text{estimation error}} + \underbrace{(1+r_f)^2\gamma^2\left(N^2\bar{\chi} - K\bar{g}\right)}_{\text{model error}}$$

The proof is in the appendix.

This equation is more transparent. Note that $(1+r_f)^2$ is of course positive, $\bar{\chi}$ is positive since Λ is positive-definite, and $\bar{g}$ is of course positive. Thus mispricing only equals zero if $K\bar{g}$ is large enough. This is unlikely if $N^2 >> K$.

The first component, labeled estimation error above, is similar to what drives mispricing in the model above. This term is larger if K is similar to T. The second component of mispricing is the model error component. This is due to the fact that the set of portfolios may not be able to completely span the mean-variance efficient frontier. If $K = N$, then this component becomes zero. Thus both model error and estimation error can remain even with quantitative investors who have a large sample size.

The below corollaries show how mispricing can remain in the model even asymptotically.

Corollary 2.4 Assume that the historical price p equals the zero alpha price of equation (3.1). Then

$$\lim_{T\to\infty} \mathbb{E}[l_t' \Lambda l_t] = (1+r_f)^2\gamma^2(N^2\bar{\chi} - K\bar{g})$$

Thus given appropriate parameter values $\bar{\chi}$ and $\bar{g}$ and large enough N compared to K, mispricing remains asymptotically even as the learning horizon expands *even if K and N are not large*. A similar corollary can be shown even if the historical prices are different than the zero-alpha price from equation (3.1), but of course the equation is slightly less transparent.

Corollary 2.5 Assume

- $\bar{\phi} = \mathcal{O}(1)$ as $N, T \to \infty$

- $\bar{\psi} = \mathcal{O}(1)$ as $N, T \to \infty$

- $\bar{\omega} = \mathcal{O}(1)$ as $N, T \to \infty$

- $\lim\limits_{\substack{N/T=c \\ N,T\to\infty}} \bar{\chi} = \chi > 0$

then

$$\lim_{\substack{N/T=c \\ N,T\to\infty}} \frac{1}{N^2}\mathbb{E}[l_t' \Lambda l_t] = (1+r_f)^2\gamma^2\chi > 0$$

Thus even if N and T get large, our measure of mispricing divided by N^2 still remains positive in the limit.

In summary, when investing in a set of portfolios using traditional portfolio management techniques replete throughout the academic literature, mispricing can still easily remain in equilibrium indefinitely.

CHAPTER 4

FULL MODEL OVERVIEW

In this chapter, I outline the full model with endogenous amounts of learners and incumbents. We switch to thinking of the incumbents and learners to be *asset managers* and the investors allocate capital between the two types of asset managers.

The incumbents-only era, also referred to the incumbents-only model, is similar to the model in Kozak et al. (2018). Let $\bar{p}_t$ denote the price that solves the incumbents-only equilibrium. There is a unit mass each of incumbent asset managers and investors. Investors are all identical. While incumbent heterogeneity is allowed in the model under some assumptions, as shown in Appendix C.2.1, only the aggregate incumbent demand functions play a role in the model. Incumbents have aggregate linear demand for assets, as discussed in Appendix C.2.1, denoted as

$$x_{I,t} = \frac{1}{\gamma}\left(a_{I,t} + J_{I,t}(p_t - \bar{p}_t)\right)$$

where $a_{I,t}$ is an $N \times 1$ intercept term, and $J_{I,t}$ is an $N \times N$ slope matrix. While a linear demand function may seem restrictive, I show below that linear demand functions arise naturally from a CARA utility framework.

Similar to the notation for incumbent demand, aggregate learner demand is also linear:

$$x_{L,t} = \frac{1}{\gamma}\left(a_{L,t} + J_{L,t}(p_t - \bar{p}_t)\right)$$

4.1 Two Kinds of Incumbents

There are two kinds of incumbents used in the model. Note that both incumbents are never in the model at the same time. The two types of incumbents are given below:

1. CARA utility incumbent demand has an intercept term that generates mispricing (written as “mistake” below) and a slope matrix that is consistent with classic CARA utility

demand:

$$x_{I,t} = \frac{1}{\gamma}\left(\underbrace{a_{I,t}}_{\text{CARA utility demand + mistake}} + \underbrace{J_{I,t}}_{\text{CARA utility slope}}(p_t - \bar{p}_t)\right)$$

2. Estimated demand (ED) has the intercept term identical to the CARA utility demand intercept. Thus mispricing is the same across incumbent types when only incumbents are in the model. However, the slope term is estimated using a Koijen and Yogo (2019) style method:

$$x_{I,t} = \frac{1}{\gamma}\left(\underbrace{a_{I,t}}_{\text{same intercept}} + \underbrace{J_{I,t}}_{\text{estimated slope}}(p_t - \bar{p}_t)\right)$$

Each CARA incumbent, knowing investors have CARA utility, generates CARA utility asset demand in order to deliver risk return trade-offs for their clients demand. This is described in more detail in Appendix C.2.1. In aggregate, CARA incumbents believe expected dollar excess returns on assets are $\mathbb{E}_t[r_{t+1}] + \delta_t$ instead of just $\mathbb{E}_t[r_{t+1}]$. The vector δ_t is a wedge that drives the incumbents to make mistakes and generate mispricing. I assume that the average values of δ_t are zero, that is $\frac{1}{N}\iota'\delta_t = 0$. Thus although incumbent asset managers make mistakes, the average mistake is zero.

Incumbent heterogeneity is allowed given the assumptions in Appendix C.2.1. The aggregate CARA incumbent demand, as derived in Appendix C.2.1 is

$$x_{I,t} = \frac{1}{\gamma}\Lambda_t^{-1}\left(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t + \delta_t\right)$$

Thus for CARA incumbents, the intercept and slope terms are

$$a_{I,t} = \Lambda_t^{-1}\left(\mathbb{E}_t[d_{t+1}] - (1+r_f)\bar{p}_t + \delta_t\right), \quad J_{I,t} = -(1+r_f)\Lambda_t^{-1}$$

4.2 Investor Problem

Investors provide an endogenous mechanism for capital allocation between the two types of asset managers. Investors can only invest with incumbent asset managers, learner asset managers, and the risk-free asset.

Investors maximize CARA utility, with the same risk aversion coefficient γ, by choosing how much to invest with each kind of asset manager. While incumbents generate α, and do not know about the parameters, investors know the 2×1 column vector of asset manager returns and the 2×2 covariance matrix of asset manager returns. Note that these returns are a function of price.

Thus this model assumes asset managers either make mistakes or are unaware of model parameters. At the same time the model assumes individual investors know the true return distributions of the funds, are hyper-rational, and are simply constrained to only invest with the asset managers and the risk-free asset. While this may appear to be backwards, this model is focused on investigating mispricing outcomes, and the assumption of hyper-rational investors is quite conservative in terms of mispricing outcomes. A model with both investors and asset managers that make mistakes or learn model parameters can easily generate much higher mispricing in equilibrium. Thus the model is both conservative in terms of mispricing outcomes and biases the model against mispricing remaining in the model.

Define $q_{I,t}(p_t) = \gamma x_{I,t}(p_t)$ and $q_{L,t}(p_t) = \gamma x_{L,t}(p_t)$. Thus these are not functions of γ. Define X_t and Q_t to be the $N \times 2$ dimensional matrices

$$X_t = [x_{I,t}(p_t) \;\; x_{L,t}(p_t)] \text{ and } Q_t = [q_{I,t}(p_t) \;\; q_{L,t}(p_t)]$$

Define

$$U_t(\theta_t, p_t) = \theta_t' X_t' (\mathbb{E}_t[d_{t+1}] - (1 + r_{f,t}) p_t) - \frac{\gamma}{2} \theta_t' X_t' \Lambda_t X_t \theta_t$$

Investors choose how much to invest with each type of asset manager, $\theta_t = (\theta_{I,t}, \theta_{L,t})'$, in order to maximize $U_t(\theta_t, p_t)$ such that

$$0 \leq \theta_{I,t}, \theta_{L,t} \leq 1$$

The remaining portion of investor wealth not used to finance positions with the two different asset managers is invested in the risk-free asset.

The values of θ_t represent the fractions of each type of learner that each investor chooses to invest with. For example, $\theta_t = (\theta_{I,t}, \theta_{L,t})' = (0.3, 0.6)'$ indicates that each investor chooses to invest with 30% of incumbent asset managers and with 60% of learner asset managers.

This solution to optimizing this objective function is the same as optimizing

$$\theta_t' Q_t' (\mathbb{E}_t[d_{t+1}] - (1 + r_{f,t}) p_t^*) - \frac{1}{2} \theta_t' Q_t' \Lambda_t Q_t \theta_t$$

which is not a function of γ. Thus the assumption that learners and incumbents have the same risk aversion coefficient γ is one of convenience, and it is ultimately the investor risk aversion that matters. Note that the constraint that the elements of θ_t be between 0 and 1 is due to the assumption that there is a unit mass of each type of asset manager. More details in are contained in Appendix C.2.1. I prove in Appendix C.2.1 that an equilibrium exists with some regularity conditions and if θ_t is restricted to $E = [\epsilon_1, 1] \times [\epsilon_2, 1]$ where either $\epsilon_1 > 0$ or $\epsilon_2 > 0$. The positive ϵ terms can be very small, and thus restrict the range of model outcomes in very minimal ways.

4.3 Equilibria

There are two eras in the model (each encapsulating multiple periods): 1) the incumbents-only era and the combined era (with both incumbents and learners). Thus there are two types of equilibria defined and discussed below. First, the more general combined equilibrium, with learners and incumbents is discussed. Then, the incumbents-only equilibrium is explained.

4.3.1 Combined Equilibrium

Supply of each asset is normalized to unity. Equilibrium is defined as a price p_t^* and investment levels θ_t^* such that θ_t^* solves the investor problem and prices p_t^* solve the following equation:

$$\iota = \frac{\theta_{I,t}}{\gamma}\left(a_{I,t} + J_{I,t}(p_t - \bar{p}_t)\right) + \frac{\theta_{L,t}}{\gamma}\left(a_{L,t} + J_{L,t}(p_t - \bar{p}_t)\right) \tag{4.1}$$

Notice that since supply is normalized to one, prices are the same as market equity if assets are equities.

Solving for an equilibrium in this model consists of finding a solution to a fixed point problem. For any given θ_t in the bounds, let $g_t(\theta_t) = p_t$ be the equilibrium price for that θ_t. For any given price p_t, define $h_t(p) = \theta_t$ to be the equilibrium θ_t demanded by investors for that price p_t. Define $f_t(\theta_t)$ to be $f(\theta_t) = h_t(g_t(\theta_t))$. Then equilibrium is found by solving for the fixed point θ_t^* such that $f(\theta_t^*) = \theta_t^*$. The equilibrium price is thus $p_t^* = g_t(\theta_t^*)$. Although prices can be easily solved for any θ_t and any θ_t can be solved for any given prices p_t, no closed form solution is available in general. Thus model simulations with estimated parameters are employed. Additional details about this equilibrium are contained in Appendix C.2.1.

4.3.2 Incumbents-only Equilibrium

This equilibrium is identical to the one above, except that $\theta_{L,t}$ is constrained to be zero.

Appendix C.2.1 proves that in the incumbents-only equilibrium, $\theta_{I,t} = 1$. That is, investors fully invest with the entire unit mass of incumbent asset managers[1].

1. Appendix C.2.1 actually proves that each investor may choose not to invest with some incumbent asset managers, but the set of incumbents asset managers the investor chooses not to invest with has measure zero.

4.4 Example of Increasing Mispricing—Numerical Example with Investors

This example of increasing mispricing is similar to the example above. Here, however, investors are included. Even with mean-variance optimizing investors, mispricing can still increase, as shown in this section. Once again, the t subscripts are dropped.

Since the investor knows all parameters and have two degrees of freedom (θ_I and θ_L), a reasonable example of mispricing increasing when learners are introduced needs three or more assets. If there are two assets and two types of asset managers, investors can effectively invest in any combination of assets (ignoring the constraints on θ) by investing through the funds themselves. Thus the investors tend to decrease mispricing in a two asset settings. In this example, I use three assets.

In this example, incumbents-only alpha is $\bar{\alpha} = (0.03, -0.18, 0.15)'$ and equilibrium alpha is $\alpha^* = (-0.148, -0.039, 0.187)'$. All details about this example are contained in Appendix C.2.8.

Critically, there is no learning in this example. The learners have a perfect intercept term—which means that the learners know the mean-variance efficient weights at the original prices. In other words, the learners have an error-free answer to Q1 above. If one of these learners back-tests the trading strategy—assuming no price impact—on similar data in the past, the portfolio Sharpe ratio will be the highest achievable Sharpe ratio.

However, the slope terms in the learner demand function in this example are filled with error. This slope term is chosen in order to give intuition about the model mechanisms. While this slope term is somewhat contrived, this is done to communicate model intuition. This learner struggles to adapt appropriately to changing prices caused by the entire group of learners.

Mispricing increases due to an overshooting effect. Asset one is initially underpriced (negative alpha), and asset demand for asset one is insufficiently reactive to price changes for both the incumbent and learner. In other words, as the price changes, demand changes very little. The learners have appealing risk-return trade-offs, so investors start to increase learner investment and decrease incumbent investment. As this occurs, the price of asset one increases because the learner has a relatively high demand for this asset. However, because demand is insufficiently price reactive for this asset, even as the price adjusts to the mean-variance efficient price and then exceeds it, learner demand is still high. Thus the price overshoots the mean-variance efficient price. Investors provide learners with the capital to create these price distortions because the other two asset positions in the learner portfolio deliver returns that are worth the cost of the position in the now-overpriced asset one. Thus in summary, insufficiently price reactive demand for asset one and appealing returns on the rest of the learner portfolio positions lead investors to allocate money into the learner funds, which in turn causes the price of asset one to overshoot the mean-variance efficient price and increase overall mispricing.

Figure A.3 gives a vector plot of $f(\theta) - \theta$ for each θ. In other words, for each θ, the arrow points in the direction of desired investor capital allocation at the given θ. If, for example, an arrow points up and left, the investor at that θ wants to invest less with the incumbent fund and more with learner asset managers. The size of the arrow is proportional to the magnitude of the $f(\theta) - \theta$. The blue dot signifies the equilibrium θ, or the point where $f(\theta) - \theta = 0$. The incumbents-only equilibrium is in the bottom right corner.

Figure A.4 gives more details about this example. Panel A shows the same vector plot, but shows a linear path with an arrow from the incumbents-only equilibrium in the bottom right corner to the equilibrium. The other four panels in this figure plot various quantities along this linear path. Panel B plots mispricing as a percentage of total incumbents-only mispricing along this path. Mispricing initially falls and then rises as asset one becomes severely overpriced. Panel C shows the information ratio of the learner portfolio, which is positive even as mispricing is increasing and terminates at zero, indicating the equilibrium has been reached. Panel D shows the overpricing of the three assets. Overpricing is defined in Appendix C.2.9, but the sum of the absolute value of the overpricing of the three assets equals the overall mispricing. This panel shows that while asset one is initially underpriced (below zero), the overpricing increases and drastically overshoots the mean-variance efficient price. Panel E shows the asset-specific components of the learner information ratio, which sum up to the overall learner information ratio. This is defined and explained further in Appendix C.2.9. Note that the information ratio is decreasing in part because of the overpricing of asset one, but the other two asset components offset this, which causes the investors to increase capital allocation to the learners.

Panel A of Figure A.5 shows that the sum of deviations from mean-variance efficient prices also increases when learners enter in this example. Panel B shows utility along the same path. Utility initially *decreases* along the path initially. Utility along the path can be written as $u(x) = U(\theta(x), p(\theta(x)))$, where $\theta(x) = (1 - x + x\theta_I^*, x\theta_L^*)$ and $p(\theta)$ is the equilibrium price that solves equation (4.1) above (i.e. $p(\theta) = g(\theta)$). Thus the slope of the curve in Panel B is

$$\frac{du}{dx} = \frac{\partial U}{\partial \theta}\frac{\partial \theta}{\partial x} + \frac{\partial U}{\partial p}\underbrace{\frac{\partial p}{\partial \theta}}_{= 0 \text{ for investors}}\frac{\partial \theta}{\partial x}$$

Note that because individual investors have no impact on prices, investors perceive that $\partial p/\partial\theta = 0$. Thus while investors are choosing θ to maximize utility, utility can actually decrease along the path if

$$\frac{\partial U}{\partial p}\frac{\partial p}{\partial \theta}\frac{\partial \theta}{\partial x} < -\frac{\partial U}{\partial \theta}\frac{\partial \theta}{\partial x}$$

Panel B also shows that in this example, utility in both the incumbents-only and combined equilibrium is equal. Recall that

$$U(\theta, p) = \theta' X'(\mathbb{E}[d_{t+1}] - (1 + r_{f,t})p) - \frac{\gamma}{2}\theta' X' \Lambda X \theta$$

However, market clearing dictates that $X\theta = \iota$. Thus equilibrium utility is

$$\iota'(\mathbb{E}[d] - (1 + r_f)p) - \frac{\gamma}{2}\iota'\Lambda\iota$$

Note that $\iota'(\mathbb{E}[d] - (1 + r_f)p)$ is the aggregate market dollar excess return. Thus if aggregate market returns across both types of equilibria are identical—which tends to happen except in some corner cases—then investor utility will be identical across equilibria.

The fact that investors' utility tends to be identical across both types of equilibria is directly caused by the fact that dividends are exogenous in this model. This model is designed to examine mispricing outcomes, and the exogenous dividends assumption allows the model to be very tractable. However, the exogenous asset dividends assumption necessarily trivializes interesting welfare analysis of model investors. Relaxing the exogenous dividends assumption would potentially allow this paper's mispricing results to be bridged to the mispricing and welfare literature (e.g. van Binsbergen and Opp (2019)). Dropping this assumption however comes at the cost of reduced tractability. Despite these difficulties, allowing dividends to vary as a function of asset prices could potentially lead to a promising area of future research.

However, welfare analysis of atomistic investors with perfect knowledge of model parameters and the ability to invest in individual assets is possible. Panel C of Figure A.5 shows the utility of these atomistic unconstrained investors. While the model investors can be thought of as constrained arbitrageurs, these unconstrained investors are unconstrained arbitrageurs. In other words, this example indicates that constrained arbitrage can both increase mispricing and increase the utility of unconstrained arbitrageurs.

In summary mispricing can increase when learners are introduced and demand is insufficiently price reactive, which leads to an overshooting effect of some assets. Other asset positions in the learner portfolio deliver high returns in expectation, convincing investors to optimally invest even though doing so ultimately increases mispricing overall.

CHAPTER 5

DIVIDEND PROCESS ESTIMATION

This chapter explains how the dividend process is estimated. The dividend process parameters naturally pin down the CARA utility demand parameters, which is also discussed in this chapter.

In this chapter and beyond, variables with bars over them represent variables in the incumbents-only equilibrium. For example, $\bar{p}_t$ is the incumbents-only equilibrium price of assets, and $\bar{r}_t$ is the excess dollar returns with the incumbents-only equilibrium price of assets. Variables with tildes represent variables estimated from data or observed directly in data. For example, $\tilde{p}_t$ represents actual asset prices, or market equity, observed in the data.

Assume dollar returns at incumbent prices, $\bar{r}_t$, follow

$$\bar{r}_t \sim N(\bar{\mu}_t, \Lambda_t)$$

It turns out that

$$\bar{\alpha}_t = -\delta_t$$

Thus

$$\delta_t = \gamma \Lambda_t \iota - \bar{\mu}_t$$

So the wedge δ_t is pinned down by $\bar{\mu}_t$ and Λ_t. I use this fact in the estimation section below. By estimating a return process, the wedge and therefore the mispricing is pinned down in the model. Thus like Kozak et al. (2018), the wedge generates CAPM alphas, which allows learners the opportunity to potentially reduce mispricing in the model.

The five parameters that are sufficient to pin down everything else in the model are $\bar{\mu}_t$, Λ_t, $\bar{p}_t$, γ, and r_f. I set $\gamma = 1$ and $r_f = 0.2\%$ per month. A risk free return of 0.2% is approximately the average t-bill rate over the sample period. In this chapter, I explain how the parameters $\bar{\mu}_t$, Λ_t, $\bar{p}_t$ are estimated or retrieved from data.

There are two main sources of variation in dollar returns at incumbents-only prices, denoted as $\bar{r}_{i,t}$: 1) variation in actual returns and 2) variation in asset size. Thus I assume:

$$\bar{r}_{i,t} = s_{i,t}\tilde{r}_{i,t}$$

where $\tilde{r}_t$ is normally distributed with mean μ_t and covariance matrix Σ_t, denoted as $\tilde{r}_t \sim N(\mu_t, \Sigma_t)$. The parameters μ_t and Σ_t are estimated from actual returns data as explained below. The variable $s_{i,t}$, and its associated vector s_t, is an exogenous measure of size, which accounts for the the variation in asset size. Let $S_t = \text{diag}(s_t)$. Then

$$\bar{r}_t \sim N(S_t\mu_t, S_t\Sigma_t S_t), \text{ thus } \bar{\mu}_t = S_t\mu_t \text{ and } \Lambda_t = S_t\Sigma_t S_t$$

There are three key ingredients to the model:

1. estimates of size s_t
2. parameters μ_t and Σ_t estimated from excess returns
3. incumbents-only prices $\bar{p}_t$ derived from observed market equity values

Plugging in these new variables, I rewrite asset demand as

$$x_{I,t} = \frac{1}{\gamma}S_t^{-1}\Sigma_t^{-1}(\mu_t - (1+r_f)S_t^{-1}(p_t - \bar{p}_t) + S_t^{-1}\delta_t)$$

Thus the new equilibrium condition is

$$\iota = \frac{1}{\gamma}S_t^{-1}\Sigma_t^{-1}(\mu_t - (1+r_f)S_t^{-1}(p_t - \bar{p}_t) + S_t^{-1}\delta_t)$$

where again the left-hand side is the supply of assets and the right-hand side is the demand of incumbents. Note that because S_t is diagonal, $S_t^{-1}p_t$ is just a price to size ratio, with similar intuition as a market equity to book value ratio. If $S_t^{-1} p_t$ is low (high), the price is low (high) relative to its size.

I first describe the data used for estimation, and then I describe the estimation steps used to obtain these three ingredients.

5.1 Data

Let Z_t be a $N \times K$ matrix of K characteristics for the N assets known at time t. I assume Z_t has a column of ones for an intercept term. I use the monthly equity returns and stock characteristics data from Freyberger et al. (2020). I use the same characteristic transformations as Kozak et al. (2020) and Kelly et al. (2019) to populate this matrix, as described in more detail in Appendix C.2.3. For a description of the variables, see Table B.3. For more details see Freyberger et al. (2020). The sample period I use is from 1970 to June 2014. The dataset includes characteristics before 1970, but I do not use this sample period due to data sparsity. Note that this dataset has been used in Kelly et al. (2019) as well.

5.2 Estimation of Size

It is important that asset size does not vary as prices move, otherwise the model becomes intractable. In order to obtain an exogenous measure of size, I follow the procedure below.

Every period t I run the following cross sectional regression, where $\tilde{p}_{i,t}$ denotes observed market equity:

$$\log(\tilde{p}_{i,t}) = a_t + b_{1,t} \log(\text{book equity}_{i,t-1}) + b_{2,t} \log(\text{sales}_{i,t-1}) + e_{it}$$

Then observed exogenous size $\tilde{s}_{i,t}$ is defined as

$$\tilde{s}_{it} = \exp(a_{t-1} + b_{1,t-1} \log(\text{book equity}_{i,t-1}) + b_{2,t-1} \log(\text{sales}_{i,t-1}))$$

The actual size vector s_t is set such that $s_t = \kappa_t \tilde{s}_t$ where the scalar κ_t is described below.

5.3 Estimation of Return Process

I assume a Koijen and Yogo (2019) style mean and covariance matrix:

$$\mu_t = Z_t \pi$$

$$\Sigma_t = \Gamma_t \Gamma_t' + \zeta I, \quad \Gamma_t = Z_t \phi$$

where ζ is a scalar, π and ϕ are $K \times 1$ vectors. This a very flexible characteristics-based model. This can accommodate non-linear functions of characteristics by plugging in transformation of Z_t. In Appendix C.2.4, I explain how this can be expanded to a neural network with an arbitrary number of hidden layers which expands and increases the flexibility of the model further. Even with this simple functional form, the econometric model flexibility is high, as discussed below.

I use maximum likelihood with a multivariate normal distribution to estimate these parameters. In Appendix C.2.6 I outline how some matrix identities can be used to make maximum likelihood calculations computationally easy and fast.

I prove in Appendix C.2.5, with some typical regularity conditions, that this estimator is consistent.

Before estimating these parameters, I use a simple out-of-sample test common in the literature to validate this econometric model. Then I explain the procedure of estimating the model parameters.

5.3.1 Out-of-Sample Test

There is a large literature on using various econometric models to estimate mean-variance efficient portfolio weights. This paper uses some of these specifically as learners, including Brandt et al. (2009), Kozak et al. (2020), and DeMiguel et al. (2009). Some others include Kelly et al. (2019), Gu et al. (2020), and Pástor (2000). In all of these papers, a common test used to validate the econometric model is to create out-of-sample portfolios, and then regress the portfolio on a common asset pricing factor model benchmark to see how the model performs. I do this same test with this model.

I use a standard expanding window estimation procedure to produce a series of estimates and portfolio weights. At the end of every calendar year, I estimate the model using only data from the current and all prior years, and then create estimated mean-variance portfolio weights for the next year by using Z_t and the estimated parameter values. I do this procedure for years 1989 through 2013 to create an excess return series from 1990 to June 2014. The period of 1970 to the end of 1989 is used as a training period. I follow Kozak et al. (2020) and scale the return series to have the same volatility as the market during the sample period.

I regress these excess returns on CRSP market excess returns to estimate a CAPM alpha and beta as usual. Table B.2 shows the results. The estimated CAPM alpha is a monthly 3.65%. This is not an annualized alpha. The t-statistic is 14.12, indicating this alpha is highly statistically significant. The information ratio of this regression is 0.83. I use this information ratio as a target for mispricing, as explained below.

5.3.2 Return Process Estimation

In order to estimate ζ, π, and ϕ I use the entire sample period from 1970 to June 2014. I cannot use this Freyberger et al. (2020) dataset without modifications because the characteristics include price ratio variables. I need a dataset of characteristics that are not a function of price. Thus I eliminate and replace some characteristics, as explained in Table B.4. I use this new dataset of 55 exogenous characteristics to populate Z_t for the rest of the paper.

If I use a simple maximum likelihood estimation as used in the previous section, I end up with average mispricing that is much higher than the out-of-sample mispricing of 0.83 estimated above. This large mispricing is due to the flexibility of the model, which leads to overfitting.

In order to alleviate this, I use a MAP estimator by subtracting a prior (penalty) term from the maximum likelihood objective function: $\lambda(\pi'_{-1}\pi_{-1} + \phi'_{-1}\phi_{-1})$ where π_{-1} and ϕ_{-1} are the slope parameters of π and ϕ respectively[1]. This corresponds to a normal prior distribution for each

1. The notation π_{-1} and ϕ_{-1} is used because the first element of π and ϕ are intercept parameters. The

slope parameter of π and ϕ with mean zero and precision λ. I set λ in order to target an average mispricing of 0.83. Figure A.6 shows how mispricing varies as the penalty parameter λ varies. From this plot, it's apparent that $\lambda = 13,125$ corresponds to mispricing of 0.83. I use this value of λ to estimate the return process parameters and use these estimates for the model.

5.4 Obtaining Incumbents-Only Prices from Market Equity

Thus far, all variables have been set relative to incumbents-only prices. This intentionally fails to pin down incumbents-only prices. I plug in incumbents-only prices $\bar{p}_t$ to be proportional to observed market equity, denoted as $\tilde{p}_t$. In the full model below, incumbents-only prices are known by asset managers and fixed.

As shown in Appendix C.2.7, the sum of prices can be written as

$$\iota' p_t = \frac{\mathbb{E}_t[r_{M,t+1}]}{\gamma \mathrm{Var}_t[r_{M,t+1}]}$$

Thus either the sum of prices can be normalized or risk aversion γ can be normalized. In the model, since Z_t varies over time because I plug in observed Z_t in the simulation, $\mathbb{E}_t[r_{M,t+1}]/\mathrm{Var}_t[r_{M,t+1}]$ varies slightly over time. Thus in order to maintain a constant risk aversion, I set $\gamma = 1$ and normalize the sum of prices. This provides an appropriate scalar, κ_t, used to rescale prices

$$\kappa_t = \frac{\mathbb{E}_t[r_{M,t+1}]}{(\iota' \tilde{p}_t)\gamma \mathrm{Var}_t[r_{M,t+1}]}$$

Thus I set

$$\bar{p}_t = \kappa_t \tilde{p}_t, \quad s_t = \kappa_t \tilde{s}_t$$

where the exogenous size variable requires rescaling so that price to size ratios are maintained for model consistency.

rest of the vectors, π_{-1} and ϕ_{-1}, are slope parameters

CHAPTER 6

DEMAND FUNCTION ESTIMATION

This section lays out the reasons for using estimated demand functions for both learners and incumbents, as well as the data and estimation methods used.

6.1 Reason for using Estimated Demand Learners

There are a number of frictions that asset managers may face by trying to implement quantative portfolio choice methods discussed in this paper, such as trading costs, obtaining the necessary expertise, and correlated arbitrage capital (see Cho (2020)). Instead of explicitly modeling all of these frictions, I estimate the demand function of institutional investors from holdings data in each period, and plug in these demand functions into the model as learners.

6.2 Reason for using Estimated Demand Incumbents

As discussed above, Koijen and Yogo (2019) argue that observed asset demand is inelastic relative to utility-based standard asset pricing demand functions. Thus it is important to use estimated demand in order to understand how mispricing is affected by learners if incumbent demand differs substantially from utility based demand.

It is important to note that at incumbents-only prices, CARA incumbents and estimated demand incumbents produce identical mispricing because they have the same demand. As prices move away from the incumbents-only equilibrium, CARA incumbents respond according to the dictates of classic mean-variance utility, while estimated demand learners have estimated price elasticty estimates. Thus while demand is identical at incumbents-only prices, demand differs when the prices move away from the incumbents-only prices.

6.3 Data

We observe long-only holdings of institutional investors in the SEC 13F filings data. See Koijen and Yogo (2019) for more details about this data. I use only the sample period from 1990 to 2014 to keep the results comparable to the rest of the simulation results. I merge these holdings to the exogenous characteristics from Freyberger et al. (2020) described in Table B.4.

Let $x_{i,t}^j$ be the demand of institution j of asset i in period t. To keep the same units of demand in the model, demand is just the total number of shares held divided by total shares outstanding. Note that $x_{i,t}^j$ is not observed, only $(x_{i,t}^j)^+ = \max\{x_{i,t}^j, 0\}$ is observed because short positions are not in the 13F holdings data. Thus the data is censored, and I use a censored data estimation approach.

6.4 Justification for using Linear Demand Function

The demand function for a machine learner who knows all the parameters is

$$x_{L,t} = \frac{1}{\gamma} S_t^{-1} \Sigma_t^{-1} (\mu_t - (1 + r_f) S_t^{-1} (p_t - \bar{p}_t))$$

$$= \frac{1}{\gamma} S_t^{-1} \left[\frac{1}{\zeta} (\mu_t - (1 + r_f) S_t^{-1} (p_t - \bar{p}_t)) + c_t \Gamma_t \right]$$

where

$$c_t = -\frac{\Gamma_t' (\mu_t - (1 + r_f) S_t^{-1} (p_t - \bar{p}_t))}{\zeta(\zeta + \Gamma_t' \Gamma_t)}$$

Note that c_t is a scalar. If c_t is assumed to be a constant, like Koijen and Yogo (2019), then demand can be written to be linear in augmented characteristics $\hat{Z}_t$. Augmented characteristics are the set of all exogenous characteristics in addition to the price-to-size ratio ($S_t^{-1} p_t$) and the incumbents-only-price-to-size ratio ($S_t^{-1} \bar{p}_t$) as follows:

$$\hat{Z}_t = \begin{bmatrix} S_t^{-1} p_t & S_t^{-1} \bar{p}_t & Z_t \end{bmatrix}$$

If the parameters Π, Φ, and b_t are defined as

$$\Pi = \begin{bmatrix} -(1+r_f) \\ 1+r_f \\ \pi \end{bmatrix}, \quad \Phi = \begin{bmatrix} 0 \\ 0 \\ \phi \end{bmatrix}, \quad b_t = \frac{1}{\zeta}\Pi + c_t\Phi$$

Then machine learner demand, with perfect knowledge of the parameters, is

$$x_{L,t} = \frac{1}{\gamma} S_t^{-1} \hat{Z}_t b_t \tag{6.1}$$

6.5 Estimation of Institutional Demand for Learners

Following Koijen and Yogo (2019), I calculate the investment universe of every institution. Let Υ_t^j denote the investment universe of institution j at time t, which is the subset of asset indices $\{1, ..., N\}$ that are in the investment universe of institution j according to Koijen and Yogo (2019). Let

$$\upsilon_{i,t}^j = \begin{cases} 1 & i \in \Upsilon_t^j \\ 0 & \text{otherwise} \end{cases}$$

I take the above linear demand function equation and modify it in four ways, as described below, to get the key estimation equation:

$$s_{i,t} x_{i,t}^j = \beta_{1,t}^j \frac{p_{i,t}}{s_{i,t}} + Z_{i,t}\beta_{2,t}^j + e_{i,t}^j \text{ for all } i \in \upsilon_t^j$$

where $\beta_{1,t}^j$ is a price coefficient, $\beta_{2,t}^j$ is a column vector of coefficients, $Z_{i,t}$ represent the i^{th} row of Z_t, and $\epsilon_{i,t}^j$ represents manager specific preferences *not* based on the observed characteristics. For notational simplicity, assume the estimates below use only the assets with the investment universe of each institution. Writing the above equation in vector form, we can write

$$S_t x_t^j = \beta_{1,t}^j S_t^{-1} p_t + Z_t \beta_{2,t}^j + e_t^j \tag{6.2}$$

Equation (6.1) is modified in the following four ways to yield equation (6.2):

1. both sides are multiplied by size S_t

2. $\hat{Z}_t$ is replaced by Z_t and $S_t^{-1}p_t$. Note that demand is estimated from real data, which is assumed to be the incumbents-only world. Thus observed price-to-size ratios are the same as incumbents-only-price-to-size ratios (i.e. $\tilde{S}_t^{-1}\tilde{p}_t = S_t^{-1}p_t = S_t^{-1}\bar{p}_t$). Thus the change from $\hat{Z}_t$ to Z_t and $S_t^{-1}p_t$ is simply to eliminate perfect multicollinearity.

3. I replaced b^j/γ^j with $\beta_t^j = (\beta_{1,t}^j, (\beta_{2,t}^j)')'$, except for the element of b^j associated with the extra column of $\hat{Z}_t$.

4. I add manager specific preferences *not* based on the observed characteristics, denoted as ϵ_t^j.

The most straightforward identifying assumption to estimate the parameters β_t^j would be

$$\mathbb{E}[e_t^j | S_t^{-1}p_t, Z_t] = 0$$

Koijen and Yogo (2019) argue that correlated demand shocks violate this assumption. They use the instrument, which I denote as $m_{i,t}^j$, which is the market equity of asset i if manager j did not exist and all other managers held an equally weighted portfolios of assets in their investment universe. This is written as

$$m_{i,t}^j = \kappa_t \sum_{l \neq i} A_t^l \frac{\upsilon_{i,t}^l}{\sum_m \upsilon_{m,t}^l}$$

where A_t^j represents the assets under management of institution j.

This yields the first stage regression

$$S_t^{-1}p_t = a_{1,t}^j S_t^{-1} m_t^j + Z_t a_{2,t}^j + \upsilon_t^j$$

I denote $(\widehat{S_t^{-1}p_t})^j$ as the predicted dependent variable from this first stage regression. Note that an intercept is not needed in this equation because Z_t contains a column of ones. The coefficients are the scalar $a^j_{1,t}$ and K dimensional vector $a^j_{2,t}$. The identifying assumption is

$$\mathbb{E}[e^j_t|(\widehat{S^{-1}_t p_t})^j, Z_t] = 0$$

which yields the second stage regression equation:

$$S_t x^j_t = \beta^j_{1,t}(\widehat{S^{-1}_t p_t})^j + Z_t \beta^j_{2,t} + e^j_t$$

where $\beta^j_{1,t}$ is the first element of β^j_t and $\beta^j_{2,t}$ is the vector containing the remaining elements.

The dependent variable in the demand equation is censored. In order to identify the parameters in a censored data setting, I assume that the error term e^j_t is conditionally symmetric, which means that

$$\text{median}(e^j_t|(\widehat{S^{-1}_t p_t})^j, Z_t) = 0$$

This allows me to use the Censored Least Absolute Deviations (CLAD) model from Powell (1984) in order to estimate the regression parameters from the second stage. The first and second stage regressions are estimated for each institution in each period. Following Koijen and Yogo (2019), I restrict the coefficient on the price-to-size ratio to be non-positive, in order to get a negative price elasticity. Also, following Koijen and Yogo (2019), if institutions have too few strictly positive holdings, they are grouped together by size and institution type. Koijen and Yogo (2019) use 1,000 strictly positive holdings as a minimum cutoff, while I use 500 since the Powell (1984) model estimator does not require as large of a sample size as the Koijen and Yogo (2019) estimator.

Let

$$\hat{\epsilon}^j_{i,t} = \begin{cases} x^j_{i,t} - \frac{1}{s_{i,t}}\left(\hat{\beta}^j_{1,t}\frac{p_{i,t}}{s_{i,t}} + Z_{i,t}\hat{\beta}^j_{2,t}\right) & \text{if } x^j_{i,t} > 0 \\ 0 & \text{otherwise} \end{cases}$$

Then the natural demand function for institution j at time t is

$$\text{diag}(v^j_t) S^{-1}_t \left(\hat{\beta}^j_{1,t} S^{-1}_t p_t + Z_t \hat{\beta}^j_{2,t} + \hat{\epsilon}^j_t\right)$$

where υ_t^j is a vector with its i^{th} element equal to $\upsilon_{i,t}^j$. The diag(υ_t^j) just ensures that the firm only demands assets in its investment universe. Plugging this demand function into the model as a learner potentially leads to issues if the institution is small, because the $\theta_{L,t} \leq 1$ could potentially be quite binding. In order to avoid this by creating an estimated demand function that could potentially become a large part of the market, let

$$\varrho_t^j = \frac{\sum_l A_t^l}{A_t^j}$$

Then the natural demand function is

$$x_{L,t}^j = \varrho_t^j \mathrm{diag}(\upsilon_t^j) S_t^{-1} \left(\hat{\beta}_{1,t}^j S_t^{-1} p_t + Z_t \hat{\beta}_{2,t}^j + \hat{\epsilon}_t^j \right)$$

I consider three alternatives to estimated demand function of learners to judge the affect of manager preferences not associated with characteristics and the investment universe restriction. The following three restrictions are listed below:

1. Plain, exactly as described above:

$$x_{L,t}^j = \varrho_t^j \mathrm{diag}(\upsilon_t^j) S_t^{-1} \left(\hat{\beta}_{1,t}^j S_t^{-1} p_t + Z_t \hat{\beta}_{2,t}^j + \hat{\epsilon}_t^j \right)$$

2. A version with no residual term:

$$x_{L,t}^j = \varrho_t^j \mathrm{diag}(\upsilon_t^j) S_t^{-1} \left(\hat{\beta}_{1,t}^j S_t^{-1} p_t + Z_t \hat{\beta}_{2,t}^j \right)$$

3. A version with no residual and no investment universe restriction:

$$x_{L,t}^j = \varrho_t^j S_t^{-1} \left(\hat{\beta}_{1,t}^j S_t^{-1} p_t + Z_t \hat{\beta}_{2,t}^j \right)$$

The holdings data are available only every quarter, but the simulations use every month of characteristic data. Therefore, in the two months between estimation, the same estimated parameter values and $\hat{\epsilon}_t^j$ from the most recent quarter are used for each institution to obtain a demand function.

6.6 Estimation of Incumbent Asset Manager Demand

In this section, I estimate non-institutional demand, and use the sum of institutional demand and non-institutional demand as the incumbents' demand function. I refer to this incumbent demand function as the estimated incumbent demand, as opposed to the CARA utility based demand function above.

I define non-institutional demand, η_t, as just residual demand:

$$\eta_t = \iota - \left(\sum_j \mathrm{diag}(\upsilon_t^j) S_t^{-1} \left(\hat{\beta}_{1,t}^j S_t^{-1} p_t + Z_t \hat{\beta}_{2,t}^j + \hat{\epsilon}_t^j \right) \right)$$

The non-institutional demand function is estimated just as the institutional demand functions are estimated except with a slight change to the instrument and a standard two stage least squared approach is used because there is no censoring. Note that η_t contains classic measurement error due to the estimated institutional demand component, but because it is a dependent variable it is inconsequential.

The first stage is

$$S_t^{-1} p_t = a_1^\eta S_t^{-1} m_t^\eta + Z_t a_2^\eta + \upsilon_t^\eta$$

where m_t^η is the market equity of each asset if all institutional investors held equally weighted portfolios of assets in their investment universe. Notice that for the calculation of the instrument above, institution j's demand was excluded in calculating m_t^j, while in the calculation of m_t^η, no institutions are excluded. Denote $(\widehat{S_t^{-1} p_t})^\eta$ as the predicted dependent variable from this first stage regression.

The second stage is the same as above, except no censored approach is required. However, the coefficient on the price-to-size ratio is still constrained to be non-positive. The second stage yields estimated β_t^η for each period and residuals $\hat{\epsilon}_t^\eta$. Overall incumbent demand is defined as

$$x_{I,t} = \left(\sum_j \mathrm{diag}(v_t^j) S_t^{-1} \left(\hat{\beta}_{1,t}^j S_t^{-1} p_t + Z_t \hat{\beta}_{2,t}^j + \hat{\epsilon}_t^j \right) \right) + S_t^{-1} \left(\hat{\beta}_{1,t}^\eta S_t^{-1} p_t + Z_t \hat{\beta}_{2,t}^\eta + \hat{\epsilon}_t^\eta \right)$$

By design, at incumbent prices, this is a vector of ones. Thus demand equals supply.

CHAPTER 7

LEARNERS

This chapter describes the eleven different kinds of learners briefly, with more details in the appendix.

7.1 Market Indexer

A market indexer has portfolio weights proportional to p_t. Thus a market indexer has demand proportional to ι. I can therefore write $x_{L,t} = a_t \iota$ for some $a_t > 0$. Equilibrium is dictated by

$$\iota = \theta_{I,t} x_{I,t} + \theta_{L,t} a_t \iota$$

Simplifying this equilibrium equation yields

$$\iota = \frac{\theta_{I,t}}{1 - \theta_{L,t} a_t} x_{I,t}$$

This is identical to the incumbents-only equilibrium equation, except that $\theta_{I,t}$ is replaced by $\theta_{I,t}/(1-\theta_{L,t}a_t)$. Using the same logic from the incumbents-only equilibrium, it must be the case that investors choose

$$\frac{\theta_{I,t}}{1 - \theta_{L,t} a_t} = 1$$

which implies

$$\iota = x_{I,t}$$

Thus prices are just incumbents-only prices, and mispricing does not change. Note that investors are indifferent between incumbent and learner investment as long as $\theta_{I,t} = 1 - \theta_{L,t} a_t$. Therefore, I do not report the percent of investment captured by learners.

7.2 Bayesian Learner

Assume a learner knows everything, including Σ_t, but does not know π. Suppose the learner has a Bayesian prior given by:

$$\pi \sim N(\pi_0, \Xi_0^{-1})$$

The posterior is

$$\pi | Z_0, ..., Z_{T-1}, \tilde{r}_1, ..., \tilde{r}_T \sim N\left(\pi_T, \Xi_T^{-1}\right)$$

where

$$\pi_T = \left(\Xi_0 + \sum_{t=0}^{T-1} Z_t' \Sigma_t^{-1} Z_t\right)^{-1} \left(\Xi_0 \pi_0 + \sum_{t=0}^{T-1} Z_t' \Sigma_t^{-1} \tilde{r}_{t+1}\right)$$

$$\Xi_T = \left(\Xi_0 + \sum_{t=0}^{T-1} Z_t' \Sigma_t^{-1} Z_t\right)$$

The posterior predictive is

$$\tilde{r}_{T+1} | Z_0, ..., Z_{T-1}, Z_T, \tilde{r}_1, ..., \tilde{r}_T \sim N\left(Z_T \pi_T, \Sigma_T + Z_T \Xi_T^{-1} Z_T'\right)$$

Therefore, demand is

$$x_{L,T} = \frac{1}{\gamma} S_T^{-1} \left(\Sigma_T + Z_T \Xi_T^{-1} Z_T'\right)^{-1} \left(Z_T \pi_T - (1 + r_f) S_T^{-1} (p_T - \bar{p}_T)\right)$$

In the simulations, I consider the case where the prior is consistent with the truth, i.e. $\pi_0 = \pi$, as well as the case where the prior π_0 is a vector of zeros. I also set $\Xi_0^{-1} = \sigma^2 I$, where I consider a range of different σ values.

7.3 True DGP Maximum Likelihood Learners

The true DGP learner estimates model parameters from the exogenous dividend process and is thus immune to error from extrapolating based on past returns. Recall that by assumption

$$\tilde{r}_t = S_t^{-1}(d_{t+1} - (1 + r_f)\bar{p}_t) \sim N(\mu_t, \Sigma_t)$$

where $\bar{p}_t$ is fixed. These true DGP estimators observed both size and incumbents-only prices $\bar{p}_t$, and use these to estimate model variables from the process $S_t^{-1}(d_{t+1} - (1+r_f)\bar{p}_t)$. These learners use a maximum likelihood estimator to obtain estimates $\hat{\zeta}$, $\hat{\pi}$, and $\hat{\phi}$. These investors plug in these values in to get $\hat{\mu}_t$ and $\hat{\Sigma}_t$. True DGP learners have demand

$$x_{L,t} = \frac{1}{\gamma} S_t^{-1} \hat{\Sigma}_t^{-1} (\hat{\mu}_t - (1+r_f) S_t^{-1}(p_t - \bar{p}_t))$$

where

$$\hat{\mu}_t = Z_t \hat{\pi}, \quad \hat{\Sigma}_t = \hat{\Gamma}_t \hat{\Gamma}_t' + \hat{\zeta} I, \quad \hat{\Gamma}_t = Z_t \hat{\phi}$$

7.4 True DGP MAP Learners

These learners are identical to the previous learners, except the learners use MAP estimators instead of maximum likelihood. The MAP estimation specification is identical to the MAP estimation used above in Chapter 5. I consider a range of parameter precision parameter values, and report the results below.

7.5 True DGP Learners with Asymptotic Parameter Values

In the model simulations, there is only a finite amount of data to estimate the model for each learner. Considering only a finite sample fails to investigate the theoretical properties of the model when the sample period grows asymptotically.

This asymptotic learner simply uses the true parameter values. This investigates how the mispricing behaves when estimation error is eliminated and the true data generating process is used. To be precise, the demand function for this learner is

$$x_{L,t} = \frac{1}{\gamma} S_t^{-1} \Sigma_t^{-1} (\mu_t - (1+r_f) S_t^{-1}(p_t - \bar{p}_t))$$

where μ_t and Σ_t are the true conditional mean and covariance matrix.

7.6 Portfolio Optimizer

This learner's demand function is derived from choosing portfolio weights using the Brandt et al. (2009) method, as described below.

In both Brandt et al. (2009) and Kozak et al. (2020), the authors use excess returns r_t. For this paper, this should be replaced by scaled dollar excess returns,

$$r^s_{t+1} = S_t^{-1}(d_{t+1} - (1 + r_f)p_t)$$

Following Kozak et al. (2020), I define the $K + 2$ dimensional vector of portfolio returns F_{t+1} as

$$F_{t+1} = \hat{Z}_t' r^s_{t+1}$$

Let $\hat{\omega}$ and $\hat{\Omega}$ be the usual estimates of the mean and covariance of F_{t+1}, that is

$$\hat{\omega}_T = \frac{1}{T}\sum_{t=1}^{T} F_t \tag{7.1}$$

$$\hat{\Omega}_T = \frac{1}{T}\sum_{t=1}^{T}(F_t - \hat{\omega}_T)(F_t - \hat{\omega}_T)' \tag{7.2}$$

In Brandt et al. (2009), the authors assume that mean-variance efficient weights are linear in characteristics. So mean-variance weights can be written as

$$w_t = \hat{Z}_t b \tag{7.3}$$

The empirical Sharpe ratio of this portfolio is

$$\frac{\frac{1}{T}\sum_{t=1}^{T} w_t' r^s_{t+1}}{\sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(w_t' r^s_{t+1} - \frac{1}{T}\sum_{\tau=1}^{T} w_\tau' r^s_{\tau+1}\right)^2}} \tag{7.4}$$

By choosing b to maximize this empirical Sharpe ratio yields a version of the Brandt et al. (2009) estimator. Note for any positive scalar k, kb and b have the same empirical Sharpe ratio. Thus by including a constraint to the sum of the parameter b, the parameters can be pinned down.

Define the Brandt et al. (2009) estimator, denoted as b_T^* to be the b that maximizes the empirical Sharpe ratio in (7.4) above subject to equation (7.3) and

$$\iota' b = \iota' \hat{\Omega}_T^{-1} \hat{\omega}_T$$

which pins down the estimate as described above.

In Appendix C.2.10, I prove that this estimator has the following closed-form solution:

$$b_T^* = \hat{\Omega}_T^{-1} \hat{\omega}_T$$

I show in Appendix C.2.11 that with constant characteristics $\hat{Z}_t$, this estimator is consistent under the true data generating process. Plugging this estimator, b_T^*, into the linear demand function (6.1) yields the Brandt et al. (2009) learner demand:

$$x_{L,T} = \frac{1}{\gamma} S_T^{-1} \hat{Z}_T b_T^*$$

7.7 Portfolio Optimizer with Shrinking Learner

This learner uses a variation of the Kozak et al. (2020) model. The Kozak et al. (2020) estimator, $\hat{b}_T$, is estimated by solving

$$\hat{b}_T = \text{argmin}_b \ (\hat{\omega}_T - \hat{\Omega}_T b)' \hat{\Omega}_T^{-1} (\hat{\omega}_T - \hat{\Omega}_T b) + \lambda_1 ||b||_1 + \lambda_2 ||b||_2$$

where $||\cdot||_1$ and $||\cdot||_2$ are the L^1 and L^2 norms respectively, and λ_1 and λ_2 are penalty parameters that shrink the parameter values towards zero.

When the penalty parameters are set to zero, that is $\lambda_1 = \lambda_2 = 0$, this estimator yields

$$\hat{b}_T = b_T^* = \hat{\Omega}_T^{-1} \hat{\omega}_T$$

as shown in Appendix C.2.10.

For simplicity, I set $\lambda_1 = 0$, which allows me to estimate b with a simple closed-form formula as shown in Appendix C.2.10. I vary λ_2 as discussed in the results below. The demand function for this learner is

$$x_{L,T} = \frac{1}{\gamma} S_T^{-1} \hat{Z}_T \hat{b}_T$$

7.8 Asymptotic Portfolio Optimizer Learner

We can compute the conditional mean and covariance of portfolio returns as follows:

$$\mathbb{E}_t[F_{t+1}] = \hat{Z}_t' \hat{Z}_t \Pi, \quad \mathrm{Var}_t[F_{t+1}] = \zeta \hat{Z}_t' \hat{Z}_t + \hat{Z}_t' \hat{Z}_t \Phi \Phi' \hat{Z}_t' \hat{Z}_t$$

Consider a fixed set of the augmented characteristic matrices, $\hat{Z}_0$, ..., $\hat{Z}_{T-1}$. Consider generating a sample of scaled returns, each denoted as f_n, using the following process:

1. Select one of the indices, 0, ..., $T-1$, each with probability $1/T$. Let $\tau(n)$ be the chosen index. Let $\hat{Z}_{\tau(n)}$ denote the corresponding matrix.

2. Take a scaled return random portfolio draw, denoted as f_n, from the distribution

$$f_n = N\left(\hat{Z}_{\tau(n)}' \hat{Z}_{\tau(n)} \Pi, \zeta \hat{Z}_{\tau(n)}' \hat{Z}_{\tau(n)} + \hat{Z}_{\tau(n)}' \hat{Z}_{\tau(n)} \Phi \Phi' \hat{Z}_{\tau(n)}' \hat{Z}_{\tau(n)}\right)$$

Denote the random sample chosen using the above method as f_1, ..., $f_{\bar{n}}$. Analogous to equations (7.1) and (7.2), define

$$\hat{\omega}_{\bar{n},T} = \frac{1}{\bar{n}} \sum_{n=1}^{\bar{n}} f_n$$

$$\hat{\Omega}_{\bar{n},T} = \frac{1}{\bar{n}} \sum_{n=1}^{\bar{n}} (f_n - \hat{\omega}_{\bar{n},T})(f_n - \hat{\omega}_{\bar{n},T})'$$

In this case, the probability limit of these, as $\bar{n} \to \infty$, denoted with the symbol as $\underset{\bar{n}\to\infty}{\overset{P}{\to}}$, is

$$\hat{\omega}_{\bar{n},T} \underset{\bar{n}\to\infty}{\overset{P}{\to}} \frac{1}{T} \sum_{t=0}^{T-1} \zeta \hat{Z}_t' \hat{Z}_t + \hat{Z}_t' \hat{Z}_t \Phi \Phi' \hat{Z}_t' \hat{Z}_t$$

$$\hat{\Omega}_{\bar{n},T} \xrightarrow[\bar{n}\to\infty]{P} \frac{1}{T}\sum_{t=0}^{T-1} \hat{Z}_t'\hat{Z}_t\Pi$$

Thus define

$$b_{\bar{n},T} = \hat{\Omega}_{\bar{n},T}^{-1}\hat{\omega}_{\bar{n},T}$$

Denote the probability limit, as $\bar{n} \to \infty$, of $b_{\bar{n},T}$ as $b_{\infty,T}$. Then by the continuous mapping theorem, we have

$$b_{\infty,T} = \left(\sum_{t=0}^{T-1} \zeta\hat{Z}_t'\hat{Z}_t + \hat{Z}_t'\hat{Z}_t\Phi\Phi'\hat{Z}_t'\hat{Z}_t\right)^{-1}\left(\sum_{t=0}^{T-1}\hat{Z}_t'\hat{Z}_t\Pi\right)$$

Analogous to the above portfolio optimizers, the resulting demand function for this learner is

$$x_{L,T} = \frac{1}{\gamma}S_T^{-1}\hat{Z}_T b_{\infty,T}$$

7.9 DeMiguel et al. (2009) $1/N$ Learners

DeMiguel et al. (2009) argue that a $1/N$ portfolio avoids estimation error and performs well in terms of out-of-sample Sharpe ratios. Thus, if this claim is true, it gets closer to the mean-variance efficient frontier than more complex methods. A true $1/N$ portfolio has portfolio weights proportional to ι, which means that demand is proportional to $P_t^{-1}\iota$. This nonlinear demand is intractable, so I use a demand function proportional to $\bar{P}_t^{-1}\iota$. This means that demand invests an equal amount of money into each asset at incumbents-only prices. Demand is scaled such that the constraint $\theta_{L,t} \leq 1$ is nonbinding in the simulations below, as explained in Appendix C.2.12. In other words, $1/N$ demand is scaled so that there is no investor excess demand for this learner.

7.10 Random Forest Learners

The random forest learner has two estimation steps:

1. Fit a random forest using Z_0, ..., Z_{T-1} to predict $\tilde{r}_1$, ..., $\tilde{r}_T$. Using this random forest, let $\hat{\mu}_T$ be the random forest prediction of scaled incumbent returns using Z_T.

2. Fit a random forest that fits the covariance matrix, that is in the form of

$$\hat{\Sigma}_t = \varphi_t \varphi_t' + \hat{\zeta} I$$

where φ_t is a column vector and a random forest function of Z_t. This is described in Appendix C.2.13.

Using the random forest predictions of the mean and covariance matrix, the random forest learner demand is

$$x_{L,t} = \frac{1}{\gamma} S_t^{-1} \hat{\Sigma}_t^{-1} (\hat{\mu}_t - (1 + r_f) S_t^{-1} (p_t - \bar{p}_t))$$

In the simulations below, the depths of the trees in the mean and covariance random forests are two. Trees in a random forest typically only consider a random fraction of features (in this case exogenous characteristics), and thus approximately $\sqrt{K}$ exogenous characteristics are selected for each tree during training to fit the target variable.

7.11 Neural Network Learner

This learner is identical to the true DGP learner above, except it uses a general neural network as described in Appendix C.2.4. Only one hidden layer is used with 25 neurons.

CHAPTER 8

MODEL SIMULATION RESULTS

The rest of the paper is concerned with the simulation results of the model. As discussed above, the model cannot be solved in closed form, so I rely on parameter estimates and simulations to make inferences. For each simulation, I assume the period 1970 to the end of 1989 as the incumbents-only era, and 1990 to June 2014 as the combined era. I use actual exogenous characteristics Z_t from the data and parameter estimates described above to simulate actual dividends. The machine learner estimates a model, and invests accordingly based on all available information. I use 200 simulations in each of the results below. This number is constrained by computational resources. This may appear to be a low number of simulations, but results are averaged across simulations and across time within the simulation, which results in a large simulated sample size. The results across simulations are quite close to each other, indicating a high degree of reliability for these results.

In each of the simulation results below, I report statistics averaged over the combined era within simulation and across simulations. I report the average mispricing, ξ_t. I also report the average mispricing as a ratio of incumbents-only mispricing, $\xi_t/\bar{\xi}_t$, in percentage terms, where $\bar{\xi}_t$ is the mispricing if investors are constrained to only invest with incumbent asset managers. Refer to Table B.1 for definitions of these variables. I also report the maximum Sharpe ratio that can be achieved with an atomistic investor with monthly market Sharpe ratios of 0.12 (the monthly empirical market Sharpe ratio observed during this period). Finally, I report the percentage of all capital invested with learners as defined in Appendix C.2.2.

The results across simulations are shown in Table B.5. I show simulations results with variations of three of the learners—the Bayesian learner, true DGP with shrinking learner, and the portfolio optimizer with shrinking learner. Table B.6 shows simulation results for various values of the prior mean and prior precision matrix. Tables B.7 and B.8 shows the results, for CARA and estimated demand incumbents respectively, across various values of λ—the prior precision parameter. Table B.9 shows the simulation results for various λ_2 shrinking parameter values. I take the learner, across the various hyperparameter values, with the lowest mispricing to include with the main results, as displayed in Table B.5. The mispricing from this table, as a percent of incumbents-only mispricing, is shown graphically as bar plots in Figure A.7.

The simulations involving estimated demand learners are quite different than the other simulations. Recall from Chapter 6 that for each institutional demand function estimated in each period, there are three types of learners used in the simulations: 1) plain learners, 2) no residual learners, and 3) no residual and no investment universe restriction learners. The estimated demand simulations are conducted as follows. Consider one of the three learner types from a given institutional demand function estimation in a given period, and use this learner to solve for the equilibrium with both incumbent types. Repeating this exercise for every one of the three learner types for every institutional demand functions in every period yields mispricing outcome data for each incumbent type, learner type, period, and institution. The mispricing simulation results are then aggregated across institutions but within incumbent types, learner types, and periods with one of the three following methods:

1. Average: The average mispricing across institutions is calculated.

2. Highest Share: The mispricing of the institution with the highest fraction of investor capital is calculated.

3. Lowest Mispricing: The mispricing of the institution that achieves the lowest mispricing in equilibrium is calculated.

Once these three measures of mispricing are calculated for both incumbent types, all three learner types, and for all three aggregation methods in each period, the results are averaged across periods within these bins. This yields 18 mispricing results (two incumbent types $\times$ three learner types $\times$ three aggregation methods).

These 18 mispricing results for estimated demand learners are shown in Figure A.8. The results with CARA incumbents are in Panel A, and these results indicate that under CARA incumbents, mispricing is essentially unchanged. In Panel B, the results with estimated demand incumbents are displayed. Mispricing increases relative to incumbents on average and for the incumbents who capture the highest investment management market share. Mispricing decreases only very slightly relative to the incumbents-only equilibrium for the learners with the lowest possible mispricing.

Some of the main findings from these two types of simulations are discussed below.

8.1 Mispricing can Increase When Sophisticated Learner Assets Managers Enter the Market

Machine learning can increase mispricing rather than improving price efficiency. This can be seen in Panel B of Figure A.7. Both types of portfolio optimizer learners (with and without shrinking), simple $1/N$ learners, and neural network learners all increase mispricing relative to the incumbents-only equilibrium. While the example above explains how mispricing can increase when learners enter, the simulation shows that under an estimated data generating process that drives returns, reasonable investment strategies can exacerbate mispricing.

Panel B of Figure A.8 shows a similar result. Mispricing on average increases when estimated demand learners are placed in the model. While Sharpe ratios of entrants can be appealing to investors, these same entrants can lead to price distortions.

In academic papers, α's and Sharpe ratios are used as a common assessment of portfolios. However, these portfolios are only theoretical and do not actually change asset prices. This paper shows that while these portfolios may seem appealing when they are not actually in the market, in practice these investment strategies can actually distort prices.

8.2 Price Elasticites are Important Determinants in Mispricing Outcomes

As discussed above, CARA incumbents have a demand function with a slope with respect to prices that equals the typical CARA utility demand function. However, the estimated demand incumbents have a different slope with respect to prices. Panel A of Figure A.7 shows the results with CARA incumbents, and Panel B shows the results with estimated demand incumbents. These results indicate that with these learners, mispricing never increases with CARA incumbents. However, as discussed above, with estimated demand incumbents mispricing does sometimes increase. Thus the price elasticity of demand for incumbents plays a large role in mispricing outcomes, including in determining whether mispricing increases or decreases. Also, in eight of the eleven learners, mispricing is higher with estimated demand incumbents than with CARA incumbents. The market indexer learner has identical mispricing. The mispricing point estimates are only lower with estimated demand incumbents rather than CARA incumbents with two learners—the true DGP with shrinking learner and the random forest learner.

Figure A.8 echoes these results. Mispricing never increases with CARA incumbents. However, mispricing on average increases across estimated demand learners. Incumbents in this model represent actual market demand since the insertion of learners is counterfactual. Thus the way that the market reacts to changing prices strongly influences mispricing outcomes when new players enter the market.

8.3 Shrinking Can Eliminate More Mispricing Than Knowing the True Covariance Matrix

As discussed above, the Bayesian learners know the true covariance matrix of returns and only learn the conditional expectation of returns parameter π. The true DGP with shrinking learner however does not know any of the parameters and uses a simple shrinking MAP estimator. Strikingly, under CARA incumbents 42.5% of incumbents-only mispricing remains, while only 30.7% remains with true DGP MAP learner incumbents. Thus the learner that does not use the true covariance matrix but estimates the covariance matrix with a shrinking estimation strategy eliminates *more* mispricing than the Bayesian learner who knows the true covariance matrix.

This relatively counter-intuitive result is analogous to MacKinlay and Pástor (2000), who show that a shrinking estimator can outperform a strategy that uses the true covariance matrix. This paper builds on this idea, showing that shrinking also has better mispricing outcomes as well as portfolio performance outcomes.

8.4 Other Interesting Results

These results show that when the mix of assets changes through time, a portfolio based approach can struggle to eliminate all mispricing. Recall from equation (6.1) that learner demand with full knowledge of the parameters is

$$x_{L,t} = \frac{1}{\gamma} S_t^{-1} \Sigma_t^{-1} (\mu_t - (1 + r_f) S_t^{-1} (p_t - \bar{p}_t)) = \frac{1}{\gamma} S_t^{-1} \hat{Z}_t b_t$$

where

$$b_t = \frac{1}{\zeta} \Pi + c_t \Phi$$

where c_t varies through time as a function of the augmented characteristics. This equation is used to justify the portfolio optimizer learners. However, in Figure A.7, the asymptotic version of this learner shows that 15.4% of mispricing remains with CARA incumbents and 18.4% of mispricing remains under estimated demand incumbents. This occurs exactly because the true b_t depends on the mix of assets and asset characteristics, which changes through time. Thus while the asymptotic true DGP learner eliminates all mispricing, this asymptotic portfolio optimizer learner fails to eliminate all mispricing. Thus, the portfolio based approach fails to eliminate all mispricing because the mix of assets and asset characteristics changes through time.

While the neural network learner estimates the true data generating process, this estimation is done using many parameters. The true model has parameters π, ϕ, and ζ, which is a total of $2K + 1$ parameters. The neural network has 25 neurons in its single hidden layer, which means there are $50K + 53$ parameters. This causes a relatively extreme overfit problem.

Panel B of Figure A.8 shows that even learners who capture a high investment management market share can lead to large mispricing increases, even with rational optimizing investors. Although intuitively, one might assume that asset managers that smart optimizing investors allocate more money to will improve price efficiency, this model indicates that this logic is not sound. In many cases, mispricing actually increases fairly drastically under these asset managers.

CHAPTER 9
CONCLUSION

This paper presents a novel econometric estimator, as well as a variety of novel estimation and theoretical innovations. The key economic results and most important details are discussed below.

The key finding of this paper is that mispricing can *increase* when sophisticated asset managers enter the market, even with optimizing intelligent investors. This is caused by learners who can create appealing portfolios before they move prices, but struggle to adapt as the data generating process of returns changes due to their own demand. Mispricing increases when asset manager demand for an asset is insufficiently price reactive, leading to an overshooting effect. In other words, the asset is initially underpriced (overpriced) and the asset becomes overpriced (underpriced) through an aggressive investment position that fails to adapt quickly as prices change (insufficiently price reactive). Other asset positions in the new entrant portfolio make the overall portfolio appealing to investors, leading to additional capital allocation even as overall asset prices become more distorted. In model simulations, this increasing-mispricing scenario plays out. Thus, it is not only theoretically possible, but likely in some situations given estimated model parameters.

This paper illustrates that adapting as the data generating process of returns changes is fundamentally difficult. Canonical portfolio optimization methods, such as Brandt et al. (2009) and Kozak et al. (2020), tend to be insufficiently price reactive and therefore adapt poorly. Even learners who estimate models trying to learn the fundamental exogenous dividend process in the model adapt poorly when these learners overfit the true dividend process.

The demand function of the market without the entrant—referred to as the incumbent demand throughout the paper—plays a key role in mispricing outcomes. If incumbent demand reacts to changing prices just as mean-variance utility dictates, mispricing does not tend to increase. However, by using estimated demand functions for incumbents, mispricing can easily increase.

Also, new entrant asset managers who do not know the true model parameters but use shrinking—prior parameter restrictions—to estimate the model can actually eliminate *more* mispricing than Bayesian asset managers who know the true covariance matrix of returns. While MacKinlay and

Pástor (2000) show that forming portfolio weights with shrinking can actually outperform portfolios using the true covariance matrix, this paper gives an analogous result for asset mispricing.

This paper presents an interesting framework to consider parameter estimation, asset management, and mispricing outcomes. Promising areas of future research include allowing an endogenous dividend process which would allow interesting welfare and mispricing analysis, or allowing the relationship between dividends and characteristics to vary through time.

BIBLIOGRAPHY

Baba Yara, Fahiz, Brian H. Boyer, and Carter Davis, 2021, Evaluating the anomaly zookeepers, *working paper* .

Brandt, Michael W., Pedro Santa-Clara, and Rossen Valkanov, 2009, Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns, *The Review of Financial Studies* 22, 3411–3447.

Centers for Disease Control and Prevention, 2021, About antibiotic resistance, *CDC Website* online.

Cho, Thummim, 2020, Turning alphas into betas: Arbitrage and endogenous risk, *Journal of Financial Economics* 137, 550 – 570.

Cochrane, John H., 2011, Presidential address: Discount rates, *The Journal of Finance* 66, 1047–1108.

DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, Positive feedback investment strategies and destabilizing rational speculation., *The Journal of Finance* 45, 379.

DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal, 2009, Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy?, *Review of Financial Studies* 22, 1915 – 1953.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1 – 22.

Freyberger, Joachim, Andreas Neuhierl, and Michael Weber, 2020, Dissecting characteristics non-parametrically, *The Review of Financial Studies* 33, 2326–2377.

Gibbons, Michael R., Stephen A. Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 57, 1121.

Gu, Shihao, Bryan Kelly, and Dacheng Xiu, 2020, Empirical asset pricing via machine learning, *The Review of Financial Studies* 33, 2223–2273.

Hansen, Lars Peter, and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.

Hart, Oliver D., and David M. Kreps, 1986, Price destabilizing speculation., *Journal of Political Economy* 94, 927 – 952.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman, 2001, *The Elements of Statistical Learning*, Springer Series in Statistics (Springer New York Inc., New York, NY, USA).

Hong, Harrison, and Jeremy C. Stein, 1999, A unified theory of underreaction, momentum trading, and overreaction in asset markets., *The Journal of Finance* 54, 2143.

Israel, Ronen, Bryan Kelly, and Tobias Moskowitz, 2020, Can machines "learn" finance?, *Journal Of Investment Management* 18.

Kelly, Bryan, Seth Pruitt, and Yinan Su, 2019, Characteristics are covariances: A unified model of risk and return, *Journal of Financial Economics* 134, 501 – 524.

Koijen, Ralph S. J., and Motohiro Yogo, 2019, A demand system approach to asset pricing, *Journal of Political Economy* 127, 1475–1515.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2018, Interpreting factor models, *The Journal of Finance* 73, 1183–1223.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2020, Shrinking the cross section, *Journal of Financial Economics* 135, 271 – 292.

MacKinlay, A. Craig, and Ľuboš Pástor, 2000, Asset pricing models: Implications for expected returns and portfolio selection, *The Review of Financial Studies* 13, 883–916.

Martin, Ian, and Stefan Nagel, 2019, Market efficiency in the age of big data, *NBER Working Paper* No. w26586.

Mitchell, Tom M., 1997, *Machine Learning* (McGraw-Hill, New York).

Powell, James L, 1984, Least absolute deviations estimation for the censored regression model, *Journal of Econometrics* 25, 303 – 325.

Pástor, Ľuboš, 2000, Portfolio selection and asset pricing models, *Journal of Finance* 55, 179–223.

Pástor, Ľuboš, and Pietro Veronesi, 2009, Learning in financial markets, *Annual Review of Financial Economics* 1, 361–381.

Stein, Jeremy C., 1987, Informational externalities and welfare-reducing speculation., *Journal of Political Economy* 95, 1123.

Stein, Jeremy C., 2009, Presidential address: Sophisticated investors and market efficiency, *The Journal of Finance* 64, 1517–1548.

van Binsbergen, Jules H., and Christian C. Opp, 2019, Real anomalies, *Journal of Finance* 74, 1659 – 1706.

APPENDIX A

APPENDIX A: FIGURES

Figure A.1: Machine Learning and Investment Management Survey

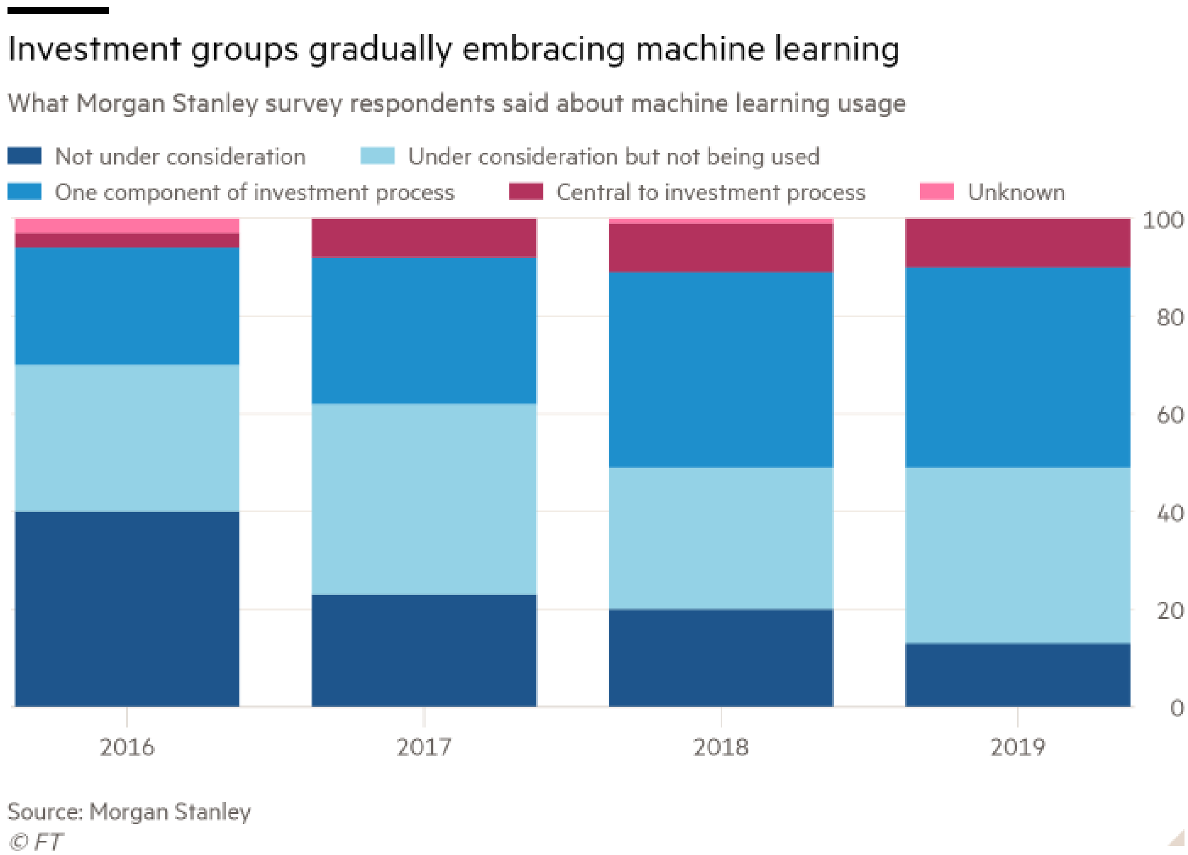

This plot shows the results from a Morgan Stanley survey of investment groups about their use of machine learning. Source: Robin Wigglesworth, Feb 10 2020, "Stockpickers turn to big data to arrest decline", *Financial Times*.

Figure A.2: Model Timeline

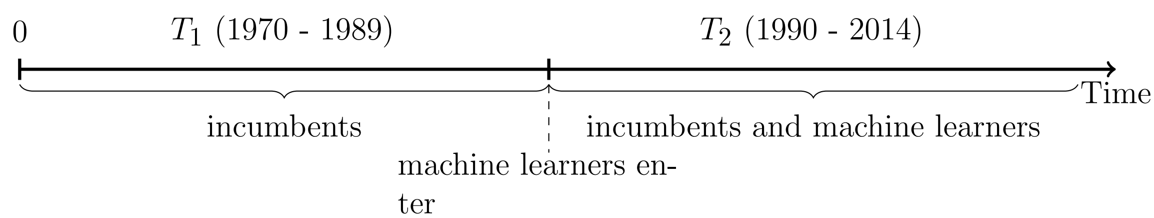

This diagram shows the basics dynamics of the model. The first era, with T_1 periods, has only incumbent asset managers. In the model simulations throughout the paper, I use firm monthly exogenous firm characteristics from 1970 - 1989 for this period. Then machine learner asset managers enter, and investors choose between investing with both incumbents and these learners. In the simulations, I use the period 1990 - 2014 for this combined era.

Figure A.3: Vector Plot of Equilibrium Example

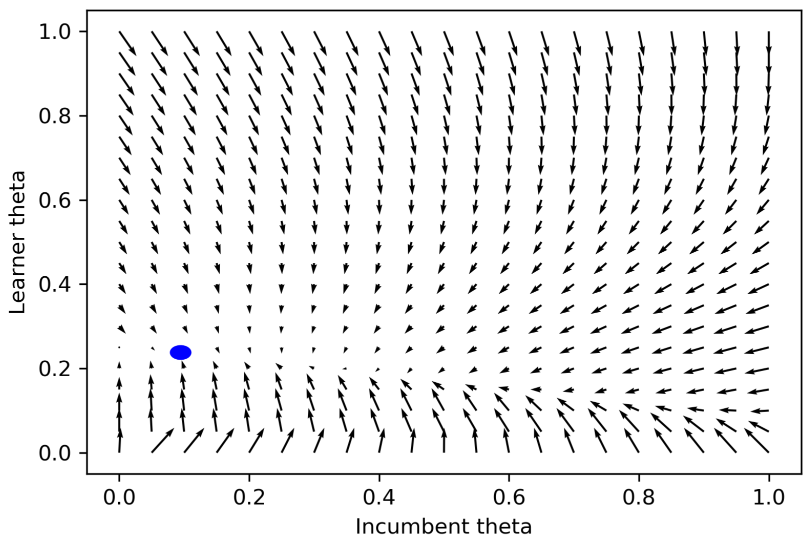

This plot gives an example of a vector plot of $f(\theta) - \theta$ for each $\theta = (\theta_I, \theta_L)$. The variable θ_I (Incumbent theta) represents the incumbent asset managers investment, while θ_L (Learner theta) represents the learner investment as defined in the paper. $f(\theta)$ is defined as $f(\theta) = h(g(\theta))$, where $g(\theta) = p$ is the price that solves the equilibrium equation given fund investment θ, and $h(p)$ is the optimal investment θ that investors make given prices p. Thus equilibrium is found by solving for θ^* such that $f(\theta^*) = \theta^*$. Thus this plot can be intuitively thought of how investors move to the equilibrium from any initial starting point, with the arrows showing the direction. The blue dot represents the equilibrium θ^*.

Figure A.4: Mispricing and Learner Information Ratio Along Path From Incumbents-Only Investment to Equilibrium

Panel A: Path From Incumbents-Only Investment to Equilibrium

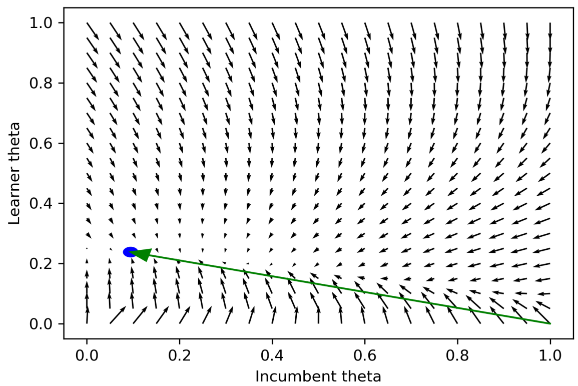

Panel B: Mispricing Along Path

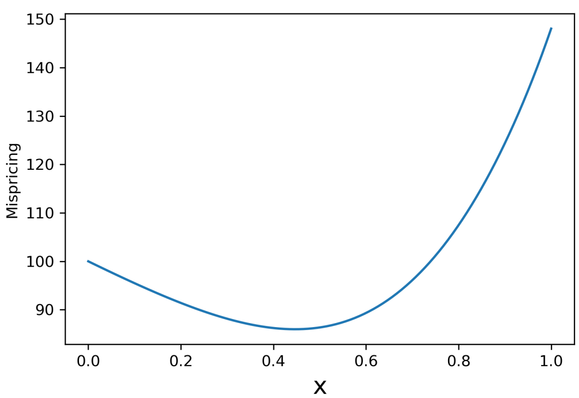

Panel C: Learner Information Ratio Along Path

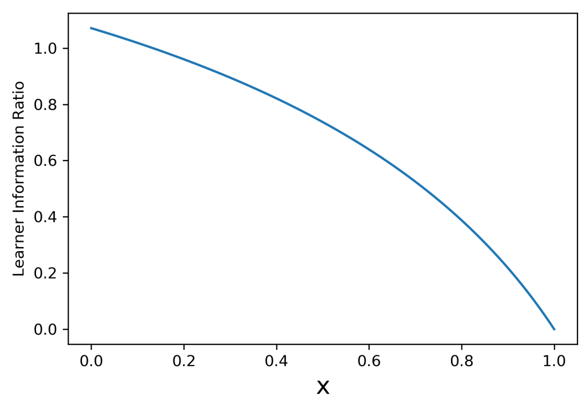

This figure is continued on the next page.

Figure A.4, continued

Panel D: Overpricing of Assets Along Path

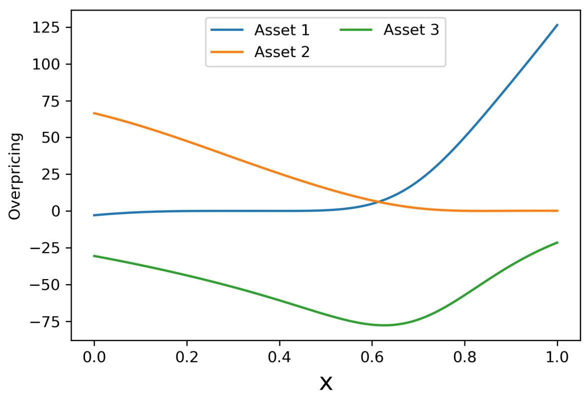

Panel E: Learner Information Ratio Component of Assets Along Path

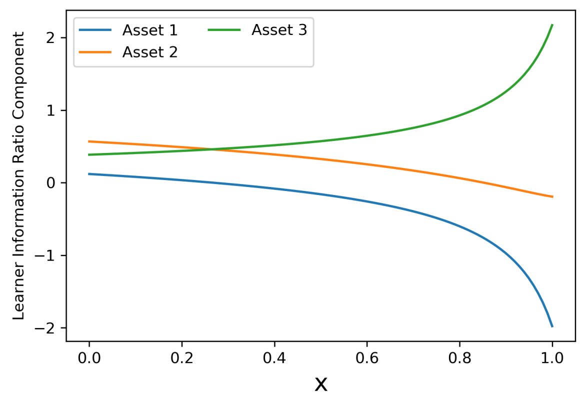

This plot gives an example of how mispricing can increase when learners are introduced. Panel A shows the same vector plot shown in Figure A.3, and the green arrow shows the linear path from incumbents-only investment, $\theta = (1, 0)$, to the equilibrium, $\theta^* = (\theta_I^*, \theta_L^*)$. The path is defined as $\theta = (1 - x + x\theta_I^*, x\theta_L^*)$ as x moves from 0 to 1. Panel B shows the overall mispricing as a percent of incumbents-only mispricing along this path, while Panel C shows the learner information ratio along this path. Panel D shows overpricing of each of the three assets along this path, while Panel E shows learner information ratio component of each of these three assets along this path. Overpricing and information ratio components are defined precisely in Appendix C.2.9.

Figure A.5: Dollar Mispricing and Utility Along Path From Incumbents-Only Investment to Equilibrium

Panel A: Dollar Mispricing Along Path From Incumbents-Only Investment to Equilibrium

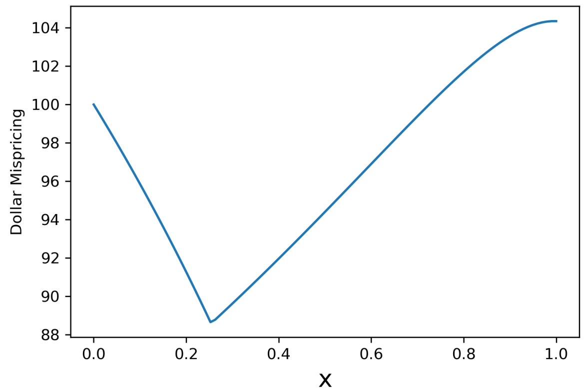

Panel B: Investor Utility

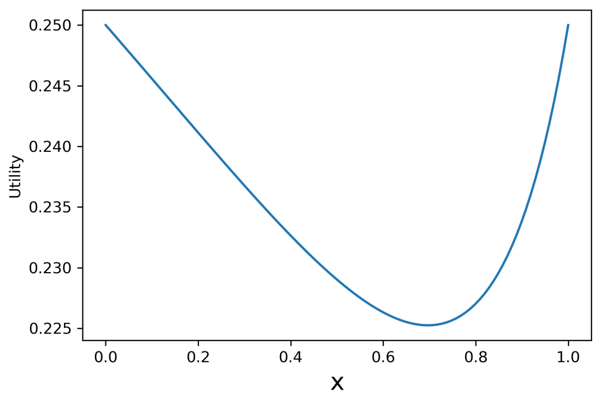

Panel C: Atomistic Unconstrained Investor Utility

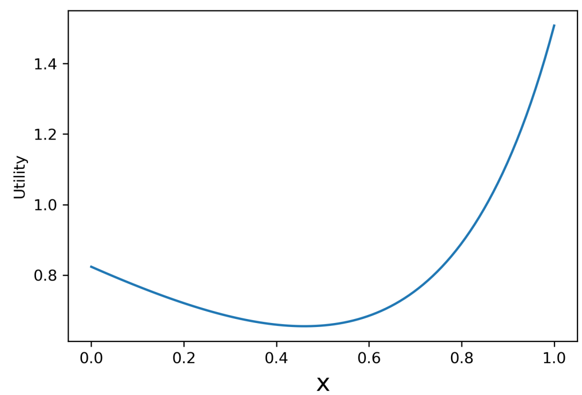

This figure shows some additional curves along the path from the incumbents-only investment equilibrium to the combined equilibrium as discussed in Figure A.4. Panel A shows dollar mispricing along this path, as a percent of total dollar mispricing in the incumbents-only equilibrium. Dollar mispricing is the sum of the absolute value of deviations of prices from the mean-variance efficient prices. Panel B shows investor utility along this path. Panel C shows the utility, along the path, of an atomistic investor that invests in individual assets—not asset managers—and has perfect knowledge of model parameters.

Figure A.6: Mispricing Target

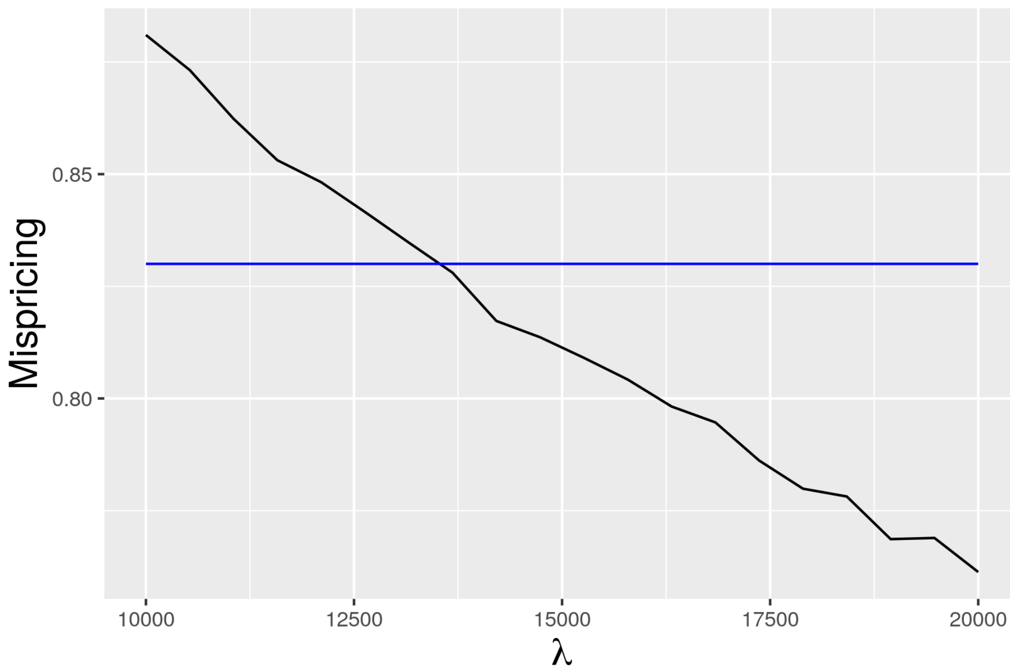

This plot shows how estimated average mispricing over the entire sample period, 1970 - June 2020, changes as the penalty prior, λ, in the maximum a posteriori (MAP) estimate changes. The target mispricing is 0.83 (see Table B.2). This target mispricing is achieved with a penalty parameter of about 13,125.

Figure A.7: Mispricing Results Summary

Panel A: Mispricing with CARA Incumbents

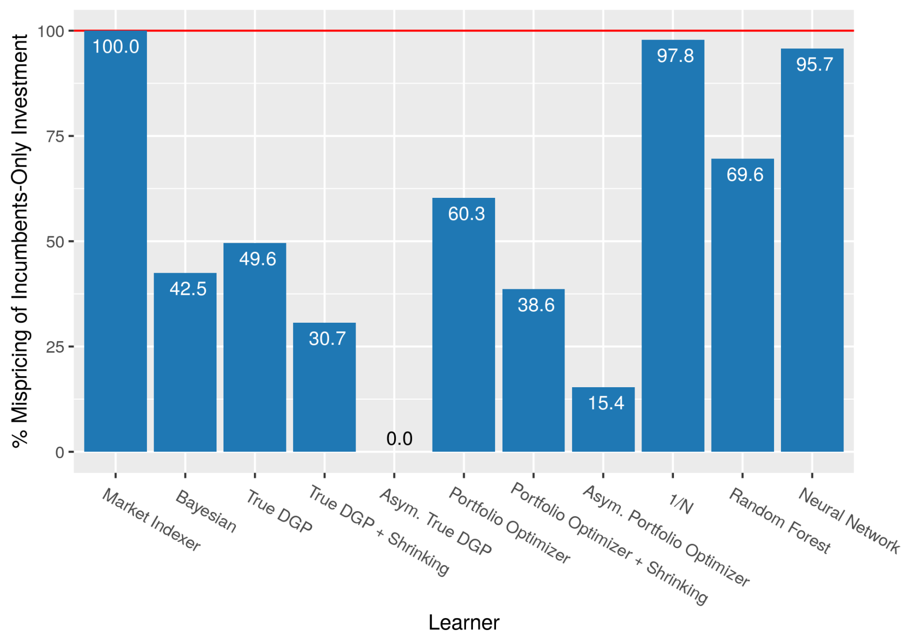

This figure is continued on the next page.

Figure A.7, continued

Panel B: Mispricing with Estimated Demand Incumbents

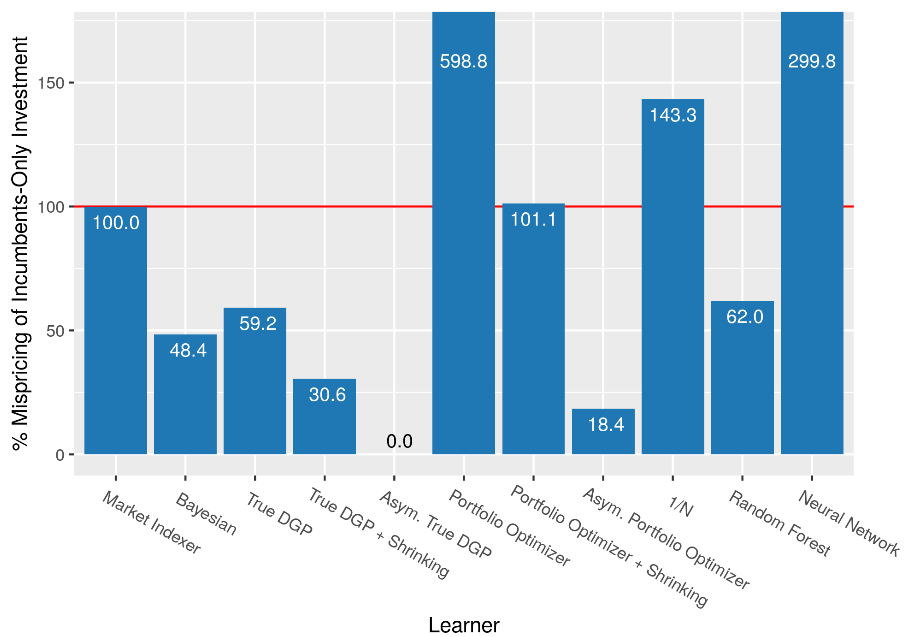

These two bar charts show how mispricing, as a percentage of incumbents-only mispricing, varies across different learner types and incumbent types. Panel A shows the results for CARA utility incumbent asset managers, while Panel B shows mispricing with estimated demand incumbents. The eleven learner asset managers are described in Chapter 7. Note that the plot in Panel B is cut off at the top; mispricing with portfolio optimizer learners and neural network learners mispricing is out of the plot region.

Figure A.8: Mispricing Results with Estimated Demand Learners

Panel A: Mispricing with CARA Incumbents

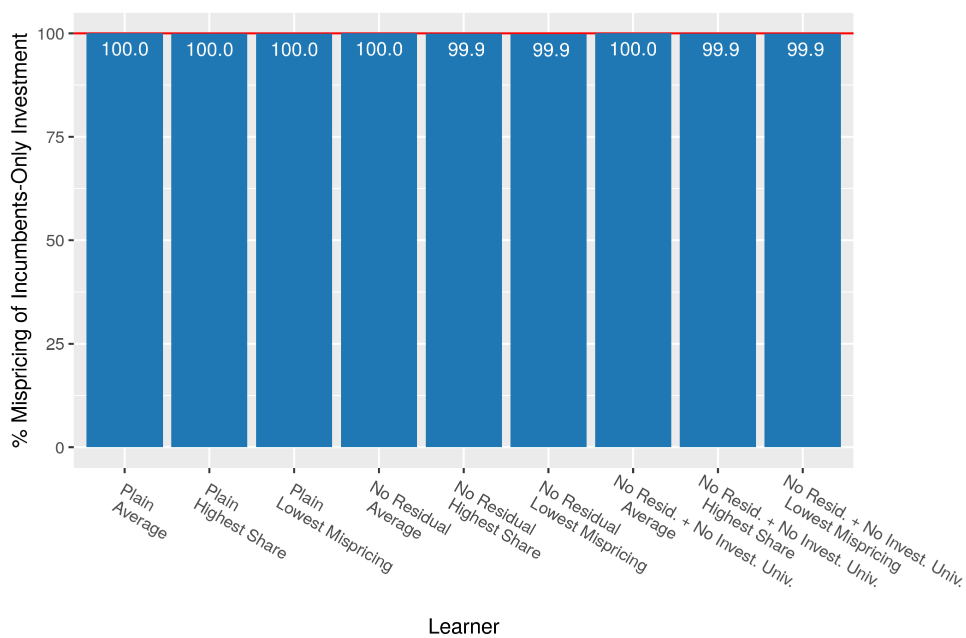

This figure is continued on the next page.

Figure A.8, continued

Panel B: Mispricing with Estimated Demand Incumbents

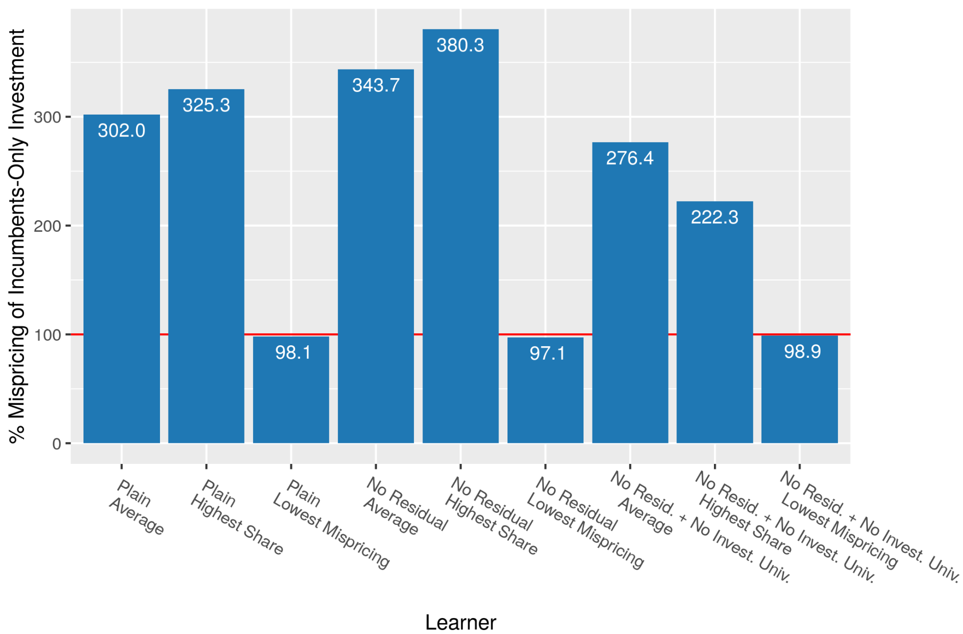

These two bar charts show how mispricing, as a percentage of incumbents-only mispricing, varies across different learner types and incumbent types. Panel A shows the results for CARA utility incumbent asset managers, while Panel B shows mispricing with estimated demand incumbents. There are nine total bars in each panel, which shows every combination of three learner types and three ways of aggregating mispricing. The three estimated demand learner types are described in Chapter 6: 1) plain learners, 2) no residual learners, and 3) no residual + no investment universe restriction learners. The three aggregating methods are as follows: 1) simple average mispricing across models and periods, 2) the average mispricing across periods with the learner that captures the highest market share of investor funds, and 3) the average mispricing across periods with the learner that generates the lowest possible mispricing.

APPENDIX B

APPENDIX B: TABLES

Table B.1: Definitions and Notation

This table shows the notation and definitions of some key variables in this paper. There are N assets, and d_t is the N dimensional column vector of dividends in period t. p_t is the N dimensional vector of prices at time t, and $P_t = \text{diag}(p_t)$. In other words, P_t is the $N \times N$ dimensional matrix with zeros on the off-diagonal and p_t on the diagonal. The top section of this table shows the standard defintions of excess returns, covariance matrices, market weights, market returns, market alpha, and mispricing, both in the dollar returns / demand share space and returns / portfolio weights space. The right-most column shows the implied relationship between the variables. The second part shows definitions of variance, excess returns, beta, alpha, and information ratio for a specific portfolio as described in the text in the center.

name	definition dollar returns / shares	definition returns / portfolio weights	implied relationship
excess returns	$r_{t+1} = d_{t+1} - (1 + r_{f,t})p_t$	$r^p_{t+1} = P_t^{-1}(d_{t+1} - (1 + r_{f,t})p_t)$	$r_{t+1} = P_t r^p_{t+1}$
covariance	$\Lambda_t = \text{Var}_t(r_{t+1})$	$\Lambda^p_t = \text{Var}_t(r^p_{t+1})$	$\Lambda_t = P_t \Lambda^p_t P_t$
market weights		$w_{M,t} = p_t/(\iota' p_t)$	
market returns	$r_{M,t+1} = \iota' r_{t+1}$	$r^p_{M,t+1} = w'_{M,t} r^p_{t+1}$	$r_{M,t+1} = (\iota' p_t) r^p_{M,t+1}$
beta	$\beta_t = \frac{\Lambda_t \iota}{\iota' \Lambda_t \iota}$	$\beta^p_t = \frac{\Lambda^p_t w_{M,t}}{w'_{M,t} \Lambda^p_t w_{M,t}}$	$\beta_t = P_t \beta^p_t / (\iota' p_t)$
alpha	$\alpha_t = \mathbb{E}_t[r_{t+1}] - \beta_t \mathbb{E}_t[r_{M,t+1}]$	$\alpha^p_t = \mathbb{E}_t[r^p_{t+1}] - \beta^p_t \mathbb{E}_t[r^p_{M,t+1}]$	$\alpha_t = P_t \alpha^p_t$
mispricing	$\xi_t = \sqrt{\alpha'_t \Lambda_t^{-1} \alpha_t}$	$\xi^p_t = \sqrt{(\alpha^p_t)'(\Lambda^p_t)^{-1}(\alpha^p_t)}$	$\xi_t = \xi^p_t$

This table is continued on the next page.

Table B.1, continued

Let x_t denote asset demand for the N assets, so each $x_{i,t}$ represents the fraction of the total value of asset i held. If w_t represents the portfolio weights held by this same set of holdings, then $w_t = P_t x_t / (\iota' P_t x_t)$ assuming portfolio weights are rescaled so that the weights sum to one. The following section of this table gives quantities specific to the this portfolio, assuming this relationship between x_t and w_t.

name	definition dollar returns / shares	definition returns / portfolio weights	implied relationship
excess returns	$r_{x,t+1} = x_t' r_{t+1}$	$r^p_{x,t+1} = w_t' r^p_{t+1}$	$r_{x,t+1} = (\iota' P_t x_t) r^p_{x,t+1}$
variance	$\sigma^2_{x,t} = \mathrm{Var}_t(r_{x,t})$	$(\sigma^p_{x,t})^2 = \mathrm{Var}_t(r^p_{x,t})$	$\sigma^2_{x,t} = (\iota' P_t x_t)^2 (\sigma^p_{x,t})^2$
beta	$\beta_{x,t} = x_t' \beta_t$	$\beta^p_{x,t} = w_t' \beta^p_t$	$\beta_{x,t} = \frac{(\iota' P_t x_t) \beta^p_{t,x}}{\iota' p_t}$
alpha	$\alpha_{x,t} = \mathbb{E}_t[r_{x,t+1}] - \beta_{x,t} \mathbb{E}_t[r_{M,t+1}]$	$\alpha^p_{x,t} = \mathbb{E}_t[r^p_{x,t+1}] - \beta^p_{x,t} \mathbb{E}_t[r^p_{M,t+1}]$	$\alpha_{x,t} = (\iota' P_t x_t) \alpha^p_{x,t}$
Sharpe ratio	$\mathrm{SR}_{x,t} = \frac{\mathbb{E}_t[r_{x,t+1}]}{\sigma_{x,t}}$	$\mathrm{SR}^p_{x,t} = \frac{\mathbb{E}_t[r^p_{x,t+1}]}{\sigma^p_{x,t}}$	$\mathrm{SR}_{x,t} = \mathrm{SR}^p_{x,t}$
information ratio	$\mathrm{IR}_{x,t} = \frac{\alpha_{x,t}}{\sqrt{\sigma^2_{x,t} - \beta^2_{x,t} \iota' \Lambda_t \iota}}$	$\mathrm{IR}^p_{x,t} = \frac{\alpha^p_{x,t}}{\sqrt{(\sigma^p_{x,t})^2 - (\beta^p_{x,t})^2 w'_{M,t} \Lambda^p_t w_{M,t}}}$	$\mathrm{IR}_{x,t} = \mathrm{IR}^p_{x,t}$

Table B.2: Out-of-Sample Portfolio CAPM Regression

	Dependent variable:
	Portfolio Return
α	3.65*** t = 14.12
β	−0.09 t = −1.60
α/σ_ϵ	0.83
Observations	293
R^2	0.01
Note:	*p<0.1; **p<0.05; ***p<0.01

This is a standard Capital Asset Pricing Model (CAPM) regression of portfolio excess returns regressed on the market excess returns. The portfolio weights are formed using only prior information using the estimated Data Generating Process (DGP) econometric model outlined in the paper. I use maximum likelihood to estimate the parameters. I use the Freyberger et al. (2020) asset characteristics data, with the Kozak et al. (2020) characteristic transformations outlined in Appendix C.2.3. The data is divided into a training period, 1970 - 1989, and a testing period, 1990 - June 2014. In every month in the testing period, the entire sample period (training and testing) prior to the given month is used to estimate the parameters, and then portfolio weights are formed. The resulting portfolio excess returns during this testing period are then regressed on the CRSP market excess returns in order to estimate the standard alpha, beta, and information ratio during this period.

Table B.3: Description of Characteristics

Category	Name	Desc.
Payout	nop	dividend and net issuance to price ratio
	ldp	dividend price ratio
	o2p	alternative payout ratio
Issuance	d_so	log change in split-adjusted shares
	d_shrout	percentage change in shares outstanding
Beta	beta	historical CAPM beta
	beta_daily	alternative CAPM beta
Investment	d_ceq	% change in book value of equity
	investment	growth in total assets
	noa	net operating assets
	dpi2a	change in PP&E
	ivc	change in inventories
	at	total assets
Cash	roc	ratio of price and debt to cash
	c2d	cash flow to liabilities
	free_cf	cash flow to book equity
	c	cash and short-term assets to assets
Profitability	sat_adj	asset turnover by industry
	eps	earnings to shares
	ipm	pre-tax income to sales
	roe	return on equity
	roa	return on assets
	pm	profit margin
	d_dgm_dsales	changes in gross margin and sales
	roic	return on invested capital
	sales_g	sales growth
	pm_adj	profit margin by industry
	rna	return on net operating assets
	ato	net sales of lagged operating assets
	pcm	price-to-cost margin
	cto	capital turnover
	prof	profitability
	s2c	sales to cash
	sat	asset turnover

Category	Name	Desc.
Value	e2p	earnings to price
	a2me	assets to price
	beme	book-to-market
	beme_adj	book-to-market by industry
	s2p	sales to price
	debt2p	debt to price
	q	Tobin's Q
Trading	std_volume	std. dev. of volume regression
	total_vol	volatility
	idio_vol	idiosyncratic volatility
	lturnover	turnover
	ret_max	max daily return in previous month
	std_turn	std. dev. of turnover regression
	rel_to_high_price	high past year price to price
	dto	alternative turnover measure
	spread_mean	bid-ask spread
	suv	standard unexplained volume
	lme	market capitalization
	lme_adj	market capitalization by industry
Intangibles	oa	operating accruals
	aoa	absolute value of operating accruals
	tan	tangibility
	ol	operating leverage
Past Returns	cum_return_36_13	long-run reversals
	cum_return_12_2	momentum
	cum_return_6_2	recent momentum
	cum_return_12_7	old momentum
	cum_return_1_0	short term reversals

This table shows the 62 characteristics used in the out-of-sample performance test (see Table B.2). This dataset comes from Freyberger et al. (2020). See their paper for more details.

Table B.4: Description of Exogenous Characteristics

Category	Name	Desc.
Issuance	d_so	log change in split-adjusted shares
	d_shrout	percentage change in shares outstanding
Beta	beta	historical CAPM beta
	beta_daily	alternative CAPM beta
Investment	d_ceq	% change in book value of equity
	investment	growth in total assets
	noa	net operating assets
	dpi2a	change in PP&E
	ivc	change in inventories
	at	total assets
Cash	c2d	cash flow to liabilities
	free_cf	cash flow to book equity
	c	cash and short-term assets to assets
Profitability	sat_adj	asset turnover by industry
	eps	earnings to shares
	ipm	pre-tax income to sales
	roe	return on equity
	roa	return on assets
	pm	profit margin
	d_dgm_dsales	changes in gross margin and sales
	roic	return on invested capital
	sales_g	sales growth
	pm_adj	profit margin by industry
	rna	return on net operating assets
	ato	net sales of lagged operating assets
	pcm	price-to-cost margin
	cto	capital turnover
	prof	profitability
	s2c	sales to cash
	sat	asset turnover

Category	Name	Desc.
Trading	std_volume	std. dev. of volume regression
	total_vol	volatility
	idio_vol	idiosyncratic volatility
	lturnover	turnover
	std_turn	std. dev. of turnover regression
	dto	alternative turnover measure
	spread_mean	bid-ask spread
	suv	standard unexplained volume
Intangibles	oa	operating accruals
	aoa	absolute value of operating accruals
	tan	tangibility
	ol	operating leverage
Basic Accounting	book	book equity
	debt	numerator of debt2p
	earnings	numerator of e2p
	dividends	numerator of ldp
	net_payout	numerator of nop
	net_payout2	numerator of o2p
	sale	numerator of s2p
Book Ratios	debt2b	debt / book
	e2b	earnings / book
	div2b	dividends / book
	nob	net_payout / book
	o2b	net_payout2 / book
	s2b	sale / book

This table shows the 55 exogenous characteristics used in the simulations throughout the paper and which all come from the Freyberger et al. (2020) dataset. Note that the variables in the Basic Accounting and Book Ratios categories are added from Table B.3, and the following variables in that table have been dropped because they were a function of price: a2me, beme, beme_adj, debt2p, e2p, ldp, lme, lme_adj, rel_to_high_price, cum_return_12_2, cum_return_12_7, q, ret_max, cum_return_6_2, cum_return_1_0, cum_return_36_13, nop, o2p, s2p, roc.

Table B.5: Mispricing Summary Table

Incumbent	Learner	Mispricing	% Mispricing	$\sqrt{0.12^2 + \text{Misp}^2}$	% Learner Capital
CARA	Market Indexer	0.91	100	0.92	
CARA	Bayesian	0.39	42.47	0.40	73.81
CARA	True DGP	0.45	49.58	0.47	70.29
CARA	True DGP + Shrinking	0.28	30.67	0.30	97.75
CARA	Asym. True DGP	0	0	0.12	100
CARA	Portfolio Optimizer	0.55	60.35	0.56	59.86
CARA	Portfolio Optimizer + Shrinking	0.35	38.64	0.37	61.77
CARA	Asym. Portfolio Optimizer	0.14	15.39	0.18	64.02
CARA	1/N	0.90	97.82	0.90	16.01
CARA	Random Forest	0.64	69.62	0.65	90.67
CARA	Neural Network	0.88	95.74	0.88	33.08
ED	Market Indexer	0.91	100	0.92	
ED	Bayesian	0.44	48.40	0.46	9.54
ED	True DGP	0.54	59.17	0.55	12.10
ED	True DGP + Shrinking	0.28	30.57	0.30	24.88
ED	Asym. True DGP	0	0	0.12	100.00
ED	Portfolio Optimizer	5.45	598.82	5.45	41.82
ED	Portfolio Optimizer + Shrinking	0.93	101.14	0.93	57.12
ED	Asym. Portfolio Optimizer	0.16	18.38	0.20	55.75
ED	1/N	1.31	143.27	1.32	20.98
ED	Random Forest	0.57	62.00	0.58	77.02
ED	Neural Network	2.73	299.85	2.73	3.54

This table shows the results across the eleven types of learner asset managers and two types incumbent asset managers. The learner types are described in Chapter 7. The columns show, from left to right: incumbent asset manager type (ED stands for estimated demand); learner asset manager types; the average mispricing across time and simulations; the average ratio of mispricing to incumbents-only mispricing—reported as a percentage—averaged across time and simulations; the Sharpe ratio an (atomistic) individual trader could achieve by combining a market monthly Sharpe ratio of 0.12 with knowledge of asset mispricing; and the average percentage of capital in the market under learner management. Note that 0.12 is the average monthly market Sharpe ratio during the entire sample period used for these simulations.

Table B.6: Mispricing with Bayesian Learners

Incumbent	Correct Prior	σ	Mispricing	% Mispricing	$\sqrt{0.12^2 + \text{Misp}^2}$	% Learner Capital
CARA	No	0.01	0.40	43.55	0.41	74.06
CARA	No	0.05	0.40	43.78	0.42	71.10
CARA	No	0.1	0.40	44.39	0.42	73.25
CARA	No	0.5	0.41	44.95	0.43	72.70
CARA	Yes	0.01	0.39	42.47	0.40	73.81
CARA	Yes	0.05	0.40	44.34	0.42	72.44
CARA	Yes	0.1	0.41	44.68	0.42	71.95
CARA	Yes	0.5	0.40	44.17	0.42	71.90
ED	No	0.01	0.44	48.85	0.46	9.54
ED	No	0.05	0.46	50.73	0.48	9.92
ED	No	0.1	0.46	50.66	0.48	9.57
ED	No	0.5	0.46	50.82	0.48	11.94
ED	Yes	0.01	0.44	48.40	0.46	9.54
ED	Yes	0.05	0.46	50.48	0.47	9.94
ED	Yes	0.1	0.46	50.36	0.47	8.93
ED	Yes	0.5	0.46	50.86	0.48	9.94

This table shows mispricing with Bayesian learners as described in Chapter 7. The columns show, from left to right: incumbent asset manager type (ED stands for estimated demand); whether the prior is objectively correct, i.e. $\pi_0 = \pi$ (yes) or equals a vector of zeros (no); the value of σ in the prior precision matrix $\Xi_0^{-1} = \sigma^2 I$; the average mispricing across time and simulations; the average ratio of mispricing to incumbents-only mispricing—reported as a percentage—averaged across time and simulations; the Sharpe ratio an (atomistic) individual trader could achieve by combining a market monthly Sharpe ratio of 0.12 with knowledge of asset mispricing; and the average percentage of capital in the market under learner management. Note that 0.12 is the average monthly market Sharpe ratio during the entire sample period used for these simulations.

Table B.7: Mispricing with CARA Incumbents and True DGP + Shrinking Learners

$\lambda/100$	Mispricing	% Mispricing	$\sqrt{0.12^2 + \text{Misp}^2}$	% Learner Capital
0	0.45	49.58	0.47	70.29
0.5	0.34	37.87	0.36	88.05
1	0.33	35.96	0.35	91.17
1.5	0.32	34.82	0.34	93.24
2	0.31	34.35	0.34	94.65
2.5	0.30	32.85	0.32	94.34
3	0.30	33.36	0.33	94.95
3.5	0.30	33.19	0.33	96.23
4	0.30	32.89	0.32	96.64
4.5	0.31	33.78	0.33	96.77
5	0.28	30.67	0.30	97.75
10	0.29	31.85	0.31	98.95
15	0.31	34.28	0.33	99.43
20	0.32	35.56	0.35	99.75
25	0.35	38.26	0.37	99.81
30	0.36	39.55	0.38	99.85
35	0.37	40.36	0.39	99.86
40	0.39	42.84	0.41	99.89
45	0.40	43.26	0.41	99.91
50	0.41	44.33	0.42	99.94
75	0.45	49.45	0.47	99.92
100	0.49	53.19	0.50	99.81
125	0.51	56.22	0.53	99.52
150	0.54	59.07	0.55	98.99
175	0.56	61.53	0.58	98.34
200	0.57	62.37	0.58	98.13

This table shows mispricing with CARA incumbent managers as described in Chapter 5 and True DGP with Shrinking learners described in Chapter 7. Note that $\lambda = 0$ corresponds to the true DGP learner. The columns show, from left to right: the MAP penalty parameter λ divided by 100; the average mispricing across time and simulations; the average ratio of mispricing to incumbents-only mispricing—reported as a percentage—averaged across time and simulations; the Sharpe ratio an (atomistic) individual trader could achieve by combining a market monthly Sharpe ratio of 0.12 with knowledge of asset mispricing; and the average percentage of capital in the market under learner management. Note that 0.12 is the average monthly market Sharpe ratio during the entire sample period used for these simulations.

Table B.8: Mispricing with Estimated Demand Incumbents and True DGP + Shrinking Learners

$\lambda/100$	Mispricing	% Mispricing	$\sqrt{0.12^2 + \text{Misp}^2}$	% Learner Capital
0	0.54	59.17	0.55	12.10
0.01	0.52	56.77	0.53	12.12
0.02	0.50	54.98	0.51	13.21
0.03	0.49	53.66	0.50	14.89
0.04	0.48	52.83	0.49	11.97
0.05	0.48	52.34	0.49	12.12
0.06	0.47	51.25	0.48	13.64
0.07	0.46	50.78	0.48	14.45
0.08	0.45	49.75	0.47	14.68
0.09	0.45	49.59	0.47	15.35
0.10	0.45	49.23	0.46	16.26
1	0.33	36.58	0.35	22.12
2	0.30	32.92	0.32	24.25
3	0.29	31.59	0.31	26.69
4	0.28	31.12	0.31	23.73
5	0.28	30.57	0.30	24.88
6	0.28	31.03	0.31	22.84
7	0.28	31.12	0.31	20.75
8	0.29	31.81	0.31	21.53
9	0.29	32.28	0.32	21.43
10	0.30	32.63	0.32	20.79

This table shows mispricing with estimated demand incumbent managers as described in Chapter 6 and True DGP with Shrinking learners described in Chapter 7. Note that $\lambda = 0$ corresponds to the true DGP learner. The columns show, from left to right: the MAP penalty parameter λ divided by 100; the average mispricing across time and simulations; the average ratio of mispricing to incumbents-only mispricing—reported as a percentage—averaged across time and simulations; the Sharpe ratio an (atomistic) individual trader could achieve by combining a market monthly Sharpe ratio of 0.12 with knowledge of asset mispricing; and the average percentage of capital in the market under learner management. Note that 0.12 is the average monthly market Sharpe ratio during the entire sample period used for these simulations.

Table B.9: Mispricing with Portolio Optimizer + Shrinking Learners

Incumbent	λ_2	Mispricing	% Mispricing	$\sqrt{0.12^2 + \text{Misp}^2}$	% Learner Capital
CARA	0	0.55	60.35	0.56	59.86
CARA	1	0.41	45.66	0.43	60.86
CARA	2	0.39	42.63	0.40	60.02
CARA	3	0.37	41.19	0.39	62.40
CARA	4	0.36	40.34	0.38	60.45
CARA	5	0.36	40.24	0.38	63.57
CARA	6	0.36	39.33	0.38	57.19
CARA	7	0.35	39.20	0.37	61.59
CARA	8	0.35	38.85	0.37	61.16
CARA	9	0.35	39.02	0.37	61.04
CARA	10	0.35	38.64	0.37	61.77
CARA	11	0.35	38.96	0.37	63.22
CARA	12	0.36	39.92	0.38	64.26
CARA	13	0.36	39.55	0.38	59.83
CARA	14	0.36	39.50	0.38	62.26
CARA	15	0.36	40.18	0.38	59.40
CARA	16	0.36	40.06	0.38	63.26
CARA	17	0.36	39.92	0.38	63.06
CARA	18	0.36	40.16	0.38	56.37
CARA	19	0.37	41.47	0.39	59.72
CARA	20	0.38	41.94	0.40	61.36
CARA	25	0.39	42.84	0.40	64.52
CARA	30	0.41	44.97	0.42	59.59
CARA	40	0.43	47.21	0.44	62.71
CARA	50	0.45	49.85	0.47	62.21
ED	0	5.45	598.82	5.45	41.82
ED	1	3.03	338.22	3.04	55.21
ED	5	1.03	113.17	1.04	56.67
ED	25	0.95	104.08	0.96	56.57
ED	50	0.93	101.14	0.93	57.12
ED	100	0.93	101.42	0.93	65.76

This table shows mispricing with portfolio optimizer learners as described in Chapter 7. Note that $\lambda_2 = 0$ corresponds to the portfolio optimizer (Brandt et al. (2009)) learners without shrinking. This figure caption is continued on the next page.

Table B.9, continued

The columns show, from left to right: incumbent asset manager type (ED stands for estimated demand); the L^2 penalty parameter from Kozak et al. (2020) labeled here as λ_2; the average mispricing across time and simulations; the average ratio of mispricing to incumbents-only mispricing—reported as a percentage—averaged across time and simulations; the Sharpe ratio an (atomistic) individual trader could achieve by combining a market monthly Sharpe ratio of 0.12 with knowledge of asset mispricing; and the average percentage of capital in the market under learner management. Note that 0.12 is the average monthly market Sharpe ratio during the entire sample period used for these simulations.

APPENDIX C

APPENDIX C: FURTHER DETAILS AND PROOFS

C.1 Lemmas and Proofs

C.1.1 Lemma 1

If $x \sim N(\mu, \Sigma)$, ϕ and ψ are constant vectors, and A is a symmetric positive definite matrix, then

$$\mathbb{E}[(x+\phi)'A(x+\psi)] = \mathrm{tr}(A\Sigma) + \mu' A(\mu+\phi+\psi) + \phi' A\psi$$

$$\mathrm{Var}[(x+\phi)'A(x+\psi)] = 2\mathrm{tr}[(A\Sigma)^2] + 4\mu' A\Sigma A(\mu+\phi+\psi) + (\phi+\psi)' A\Sigma A(\phi+\psi)$$

For the proof, see Baba Yara et al. (2021).

C.1.2 Proof of Proposition 1.1

We set out to prove that

$$E[\alpha_T'\Lambda^{-1}\alpha_T] = \left(\frac{T}{\kappa+T}\right)^2 \left(\frac{N-1}{T}\right)$$

$$\mathrm{Var}[\alpha_T'\Lambda^{-1}\alpha_T] = \left(\frac{T}{\kappa+T}\right)^4 \left(\frac{2(N-1)}{T^2}\right)$$

While the variance equation above is not stated in the text, it is useful for proposition 1.2 below.

Expected returns are

$$\mathbb{E}_T[r_{T+1}] = \mathbb{E}_T[d_{T+1} - (1+r_f)p_T] = \mu - (1+r_f)p_T$$

$$= \mu - (\Lambda_0^{-1} + T\Lambda^{-1})^{-1}\left(\Lambda_0^{-1}\mu_0 + T\Lambda^{-1}\bar{d}_T\right) + \gamma\left((\Lambda_0^{-1} + T\Lambda^{-1})^{-1} + \Lambda\right)\iota$$

Asset β is

$$\beta_T = \frac{\Lambda\iota}{\iota'\Lambda\iota}$$

And the expected market return is

$$\mathbb{E}_T[\iota' r_{T+1}] = \iota'\mu - \iota'(\Lambda_0^{-1} + T\Lambda^{-1})^{-1}\left(\Lambda_0^{-1}\mu_0 + T\Lambda^{-1}\bar{d}_T\right) + \gamma\iota'\left((\Lambda_0^{-1} + T\Lambda^{-1})^{-1} + \Lambda\right)\iota$$

Thus the α_T is

$$\alpha_T = \mathbb{E}_T[r_{T+1}] - \beta_T\mathbb{E}_T[\iota' r_{T+1}]$$

$$= \left(I - \frac{\Lambda\iota\iota'}{\iota'\Lambda\iota}\right)\left(\mu - (\Lambda_0^{-1} + T\Lambda^{-1})^{-1}\left(\Lambda_0^{-1}\mu_0 + T\Lambda^{-1}\bar{d}_T\right) + \gamma\left((\Lambda_0^{-1} + T\Lambda^{-1})^{-1} + \Lambda\right)\iota\right)$$

$$= \left(I - \frac{\Lambda\iota\iota'}{\iota'\Lambda\iota}\right)\left(\mu - (\Lambda_0^{-1} + T\Lambda^{-1})^{-1}\Lambda_0^{-1}\mu_0 - (\Lambda_0^{-1} + T\Lambda^{-1})^{-1}T\Lambda^{-1}\bar{d}_T + \gamma\left((\Lambda_0^{-1} + T\Lambda^{-1})^{-1} + \Lambda\right)\iota\right)$$

If we plug in $\Lambda_0 = \kappa^{-1}\Lambda$ and $\mu_0 = \mu$, this simplifies to

$$\alpha_T = \left(I - \frac{\Lambda\iota\iota'}{\iota'\Lambda\iota}\right)\left(\frac{T}{\kappa+T}\mu - \frac{T}{\kappa+T}\bar{d}_T + \gamma\frac{1+\kappa+T}{\kappa+T}\Lambda\iota\right)$$

$$= A\bar{d}_T + b$$

where

$$A = -\frac{T}{\kappa+T}\left(I - \frac{\Lambda\iota\iota'}{\iota'\Lambda\iota}\right)$$

$$b = \left(I - \frac{\Lambda\iota\iota'}{\iota'\Lambda\iota}\right)\left(\frac{T}{\kappa+T}\mu + \gamma\frac{1+\kappa+T}{\kappa+T}\Lambda\iota\right)$$

Thus

$$\alpha_T'\Lambda^{-1}\alpha_T = (A\bar{d}_T + b)'\Lambda^{-1}(A\bar{d}_T + b)$$

Note that

$$\bar{d}_T \sim N\left(\mu, \frac{1}{T}\Lambda\right)$$

Thus

$$A\bar{d}_T \sim N\left(A\mu, \frac{1}{T}A\Lambda A'\right)$$

Thus, from Lemma 1 above,

$$E[\alpha_T'\Lambda^{-1}\alpha_T] = \frac{1}{T}\mathrm{tr}\left(\Lambda^{-1}A\Lambda A'\right) + \mu'A'\Lambda^{-1}(A\mu+2b) + b'\Lambda^{-1}b$$

$$\mathrm{Var}[\alpha_T'\Lambda^{-1}\alpha_T] = \frac{2}{T^2}\mathrm{tr}\left[\left(\Lambda^{-1}A\Lambda A'\right)^2\right] + \frac{4}{T}\mu'A'\Lambda^{-1}A\Lambda A'\Lambda^{-1}(A\mu+2b) + \frac{4}{T}b'\Lambda^{-1}A\Lambda A'\Lambda^{-1}b$$

We can calculate

$$A\Lambda A' = \left(\frac{T}{\kappa+T}\right)^2\left(I - \frac{\Lambda\iota\iota'}{\iota'\Lambda\iota}\right)\Lambda\left(I - \frac{\Lambda\iota\iota'}{\iota'\Lambda\iota}\right)'$$

$$= \left(\frac{T}{\kappa+T}\right)^2\left(\Lambda - \frac{\Lambda\iota\iota'\Lambda}{\iota'\Lambda\iota}\right)\left(I - \frac{\iota\iota'\Lambda}{\iota'\Lambda\iota}\right)$$

$$= \left(\frac{T}{\kappa+T}\right)^2\left(\Lambda - \frac{\Lambda\iota\iota'\Lambda}{\iota'\Lambda\iota}\right)$$

Thus

$$\Lambda^{-1}A\Lambda A' = \left(\frac{T}{\kappa+T}\right)^2\left(I - \frac{\iota\iota'\Lambda}{\iota'\Lambda\iota}\right)$$

Using this, we can calculate

$$\mathrm{tr}\left(\Lambda^{-1}A\Lambda A\right) = \left(\frac{T}{\kappa+T}\right)^2\mathrm{tr}\left(I - \frac{\iota\iota'\Lambda}{\iota'\Lambda\iota}\right) = \left(\frac{T}{\kappa+T}\right)^2\left(N - \frac{tr(\iota\iota'\Lambda)}{\iota'\Lambda\iota}\right)$$

$$= \left(\frac{T}{\kappa+T}\right)^2\left(N - \frac{\iota'\Lambda\iota}{\iota'\Lambda\iota}\right) = \left(\frac{T}{\kappa+T}\right)^2(N-1)$$

We can also calculate

$$A'\Lambda^{-1} = -\left(\frac{T}{\kappa+T}\right)\left(I - \frac{\Lambda\iota\iota'}{\iota'\Lambda\iota}\right)'\Lambda^{-1} = -\left(\frac{T}{\kappa+T}\right)\left(\Lambda^{-1} - \frac{\iota\iota'}{\iota'\Lambda\iota}\right)$$

Thus

$$A'\Lambda^{-1}A = \left(\frac{T}{\kappa+T}\right)^2\left(\Lambda^{-1} - \frac{\iota\iota'}{\iota'\Lambda\iota}\right)\left(I - \frac{\Lambda\iota\iota'}{\iota'\Lambda\iota}\right)$$

$$= \left(\frac{T}{\kappa+T}\right)^2\left(\Lambda^{-1} - \frac{\iota\iota'}{\iota'\Lambda\iota}\right)$$

Thus

$$\mu' A' \Lambda^{-1} A \mu = \left(\frac{T}{\kappa + T} \right)^2 \left(\mu' \Lambda^{-1} \mu - \frac{(\iota' \mu)^2}{\iota' \Lambda \iota} \right)$$

Similarly, we can calculate

$$\Lambda^{-1} b = \left(\Lambda^{-1} - \frac{\iota \iota'}{\iota' \Lambda \iota} \right) \left(\frac{T}{\kappa + T} \mu + \gamma \frac{1 + \kappa + T}{\kappa + T} \Lambda \iota \right)$$

First

$$A' \Lambda^{-1} b = -\frac{T}{\kappa + T} \left(I - \frac{\iota \iota' \Lambda}{\iota' \Lambda \iota} \right) \left(\Lambda^{-1} - \frac{\iota \iota'}{\iota' \Lambda \iota} \right) \left(\frac{T}{\kappa + T} \mu + \gamma \frac{1 + \kappa + T}{\kappa + T} \Lambda \iota \right)$$

$$= -\frac{T}{\kappa + T} \left(\Lambda^{-1} - \frac{\iota \iota'}{\iota' \Lambda \iota} \right) \left(\frac{T}{\kappa + T} \mu + \gamma \frac{1 + \kappa + T}{\kappa + T} \Lambda \iota \right)$$

$$= -\left(\frac{T}{\kappa + T} \right)^2 \left(\Lambda^{-1} \mu - \frac{\iota \iota' \mu}{\iota' \Lambda \iota} \right)$$

Thus

$$\mu' A' \Lambda^{-1} b = -\left(\frac{T}{\kappa + T} \right)^2 \left(\mu' \Lambda^{-1} \mu - \frac{(\iota' \mu)^2}{\iota' \Lambda \iota} \right)$$

We calculate

$$b' \Lambda^{-1} b = \left(\frac{T}{\kappa + T} \mu' + \gamma \frac{1 + \kappa + T}{\kappa + T} \iota' \Lambda \right) \left(I - \frac{\iota \iota' \Lambda}{\iota' \Lambda \iota} \right) \Lambda^{-1} \left(I - \frac{\Lambda \iota \iota'}{\iota' \Lambda \iota} \right) \left(\frac{T}{\kappa + T} \mu + \gamma \frac{1 + \kappa + T}{\kappa + T} \Lambda \iota \right)$$

$$= \left(\frac{T}{\kappa + T} \mu' + \gamma \frac{1 + \kappa + T}{\kappa + T} \iota' \Lambda \right) \left(\Lambda^{-1} - \frac{\iota \iota'}{\iota' \Lambda \iota} \right) \left(\frac{T}{\kappa + T} \mu + \gamma \frac{1 + \kappa + T}{\kappa + T} \Lambda \iota \right)$$

$$= \left(\frac{T}{\kappa + T} \right) \left(\frac{T}{\kappa + T} \mu' + \gamma \frac{1 + \kappa + T}{\kappa + T} \iota' \Lambda \right) \left(\Lambda^{-1} \mu - \frac{\iota \iota' \mu}{\iota' \Lambda \iota} \right)$$

$$= \left(\frac{T}{\kappa + T} \right)^2 \left(\mu' \Lambda^{-1} \mu - \frac{(\iota' \mu)^2}{\iota' \Lambda \iota} \right)$$

Thus we have

$$E[\alpha_T' \Lambda^{-1} \alpha_T] = \left(\frac{T}{\kappa + T} \right)^2 \left(\frac{N - 1}{T} \right) \tag{C.1}$$

Once again calculating:

$$A'\Lambda^{-1}A\Lambda A'\Lambda^{-1}A = \left(\frac{T}{\kappa+T}\right)^4\left(\Lambda^{-1}-\frac{\iota\iota'}{\iota'\Lambda\iota}\right)\Lambda\left(\Lambda^{-1}-\frac{\iota\iota'}{\iota'\Lambda\iota}\right)$$

$$= \left(\frac{T}{\kappa+T}\right)^4\left(\Lambda^{-1}-\frac{\iota\iota'}{\iota'\Lambda\iota}\right)$$

Which means that

$$\mu'A'\Lambda^{-1}A\Lambda A'\Lambda^{-1}A\mu = \left(\frac{T}{\kappa+T}\right)^4\left(\mu'\Lambda^{-1}\mu-\frac{(\iota'\mu)^2}{\iota'\Lambda\iota}\right)$$

We can also calculate:

$$A'\Lambda^{-1}A\Lambda A'\Lambda^{-1}b = -\left(\frac{T}{\kappa+T}\right)^4\left(\Lambda^{-1}-\frac{\iota\iota'}{\iota'\Lambda\iota}\right)\Lambda\left(\Lambda^{-1}\mu-\frac{\iota\iota'\mu}{\iota'\Lambda\iota}\right)$$

$$= -\left(\frac{T}{\kappa+T}\right)^4\left(\Lambda^{-1}-\frac{\iota\iota'}{\iota'\Lambda\iota}\right)\mu$$

premultiplying by μ' yields

$$\mu'A'\Lambda^{-1}A\Lambda A'\Lambda^{-1}b = -\left(\frac{T}{\kappa+T}\right)^4\left(\mu'\Lambda^{-1}\mu-\frac{(\iota'\mu)^2}{\iota'\Lambda\iota}\right)$$

We can also calculate

$$b'\Lambda^{-1}A\Lambda A'\Lambda^{-1}b = \left(\frac{T}{\kappa+T}\right)^4\left(\mu'\Lambda^{-1}-\frac{\mu'\iota\iota'}{\iota'\Lambda\iota}\right)\Lambda\left(\Lambda^{-1}\mu-\frac{\iota\iota'\mu}{\iota'\Lambda\iota}\right)$$

$$= \left(\frac{T}{\kappa+T}\right)^4\left(\mu'\Lambda^{-1}\mu-\frac{(\iota'\mu)^2}{\iota'\Lambda\iota}\right)$$

This means that

$$\mathrm{Var}[\alpha_T'\Lambda^{-1}\alpha_T] = \frac{2}{T^2}\mathrm{tr}\left[\left(\Lambda^{-1}A\Lambda A'\right)^2\right]$$

We can calculate

$$(\Lambda^{-1}A\Lambda A')^2 = \left(\frac{T}{\kappa+T}\right)^4 \left(I - \frac{\iota\iota'\Lambda}{\iota'\Lambda\iota}\right)\left(I - \frac{\iota\iota'\Lambda}{\iota'\Lambda\iota}\right)$$

$$= (\Lambda^{-1}A\Lambda A')^2 = \left(\frac{T}{\kappa+T}\right)^4 \left(I - \frac{\iota\iota'\Lambda}{\iota'\Lambda\iota}\right)$$

Thus we can calculate

$$\mathrm{tr}[(\Lambda^{-1}A\Lambda A')^2] = \left(\frac{T}{\kappa+T}\right)^4 (N-1)$$

Thus

$$\mathrm{Var}[\alpha_T'\Lambda^{-1}\alpha_T] = \left(\frac{T}{\kappa+T}\right)^4 \left(\frac{2(N-1)}{T^2}\right)$$

C.1.3 Proof of Proposition 1.2

We can easily see that:

$$\lim_{\substack{N/T=c\\ N,T\to\infty}} E[\alpha_T'\Lambda^{-1}\alpha_T] = \lim_{\substack{N/T=c\\ N,T\to\infty}} \left(\frac{T}{\kappa+T}\right)^2 \left(\frac{N-1}{T}\right) = c$$

and

$$\lim_{\substack{N/T=c\\ N,T\to\infty}} \mathrm{Var}[\alpha_T'\Lambda^{-1}\alpha_T] = \lim_{\substack{N/T=c\\ N,T\to\infty}} \left(\frac{T}{\kappa+T}\right)^4 \left(\frac{2(N-1)}{T^2}\right) = 0$$

Which means that

$$\operatorname*{plim}_{\substack{N/T=c\\ N,T\to\infty}} \alpha_T'\Lambda^{-1}\alpha_T = c$$

C.1.4 Proof of Proposition 2.1

From equation (3.1), we have

$$p_t = \frac{1}{1+r_f}\left(\mathbb{E}_t[d_{t+1}] - \gamma\Lambda\iota\right)$$

which means we can write

$$\mu - (1+r_f)p = \gamma\Lambda\iota \tag{C.2}$$

Recall that equilibrium is given by

$$\iota = \frac{1}{\gamma} Z_t b_t$$

This is the same as

$$\iota = \frac{1}{\gamma}(Z_I + P_t Z_S) b_t$$

which is equivalent to

$$\gamma\iota - Z_I b_t = P_t Z_S b_t \tag{C.3}$$

Since $Z_S b_t$ does not have any elements that exactly equal zero, equation (C.2) holds if and only if

$$U_t Z_S b_t - (1 + r_f) P_t Z_S b_t - \gamma\lambda Z_S b_t = 0$$

where $U_t = \text{diag}(\mu_t)$ and $\lambda = \text{diag}(\Lambda\iota)$. We can use equation (C.3) to rewrite this as:

$$U_t Z_S b_t - (1 + r_f)\gamma\iota + (1 + r_f) Z_I b_t - \gamma\lambda Z_S b_t = 0$$

Define

$$l_t = U_t Z_S b_t - (1 + r_f)\gamma\iota + (1 + r_f) Z_I b_t - \gamma\lambda Z_S b_t = -(1 + r_f)\gamma\iota + \tilde{Z} b_t \tag{C.4}$$

Thus the price equals the $\alpha_t = 0$ price of equation (3.1) if and only if

$$l_t' \Lambda l_t = 0$$

since Λ is positive definite.

C.1.5 Proof of Proposition 2.2

Define

$$\hat{d}_T = \frac{1}{T}\sum_{t=1}^{T} d_t, \quad \hat{p}_{T-1} = \frac{1}{T}\sum_{t=0}^{T-1} p_t,$$

$$\hat{r}_T = \frac{1}{T}\sum_{t=1}^{T} r_t = \frac{1}{T}\sum_{t=1}^{T}(d_t - (1 + r_f)p_{t-1}) = \hat{d}_T - (1 + r_f)\hat{p}_{T-1}$$

Based on our assumption about the distribution of dividends, it must be the case that

$$\hat{d}_T \sim N\left(\mu, \frac{1}{T}\Lambda\right)$$

Recall that

$$b_T \sim N\left((\bar{Z}'\Lambda\bar{Z})^{-1}\bar{Z}'(\mu - (1+r_f)p), \frac{1}{T}(\bar{Z}'\Lambda\bar{Z})^{-1}\right)$$

We can rewrite the wedge from equation (C.4) as

$$l_t = -(1+r_f)\gamma\iota + ((1+r_f) + Z_I(U_t + \gamma\lambda)Z_S)b_t$$

Define

$$\tilde{Z} = (1+r_f) + Z_I(U_t + \gamma\lambda)Z_S$$

Thus we can write

$$l_t = -(1+r_f)\gamma\iota + \tilde{Z}b_t$$

Thus

$$l_t \sim N\left(\mathbb{E}[l_t], \mathrm{Var}(l_t)\right)$$

where

$$\mathbb{E}[l_t] = -(1+r_f)\gamma\iota - \tilde{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\bar{Z}'((1+r_f)p - \mu)$$

$$\mathrm{Var}(l_t) = \frac{1}{T}\tilde{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\tilde{Z}'$$

Thus by Lemma 1 above, we can write:

$$\mathbb{E}[l_t'\Lambda l_t] = \mathrm{tr}\left(\Lambda\frac{1}{T}\tilde{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\tilde{Z}'\right) + (1+r_f)^2\gamma^2\iota'\Lambda\iota$$

$$+((1+r_f)p-\mu)'\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\tilde{Z}'\Lambda\tilde{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\bar{Z}'((1+r_f)p-\mu)$$

$$+2(1+r_f)\gamma((1+r_f)p-\mu)'\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\tilde{Z}'\Lambda\iota$$

$$= \frac{1}{T}\text{tr}\left(\tilde{Z}'\Lambda\tilde{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\right) + (1+r_f)^2\gamma^2\iota'\Lambda\iota$$

$$+\text{tr}\left(\bar{Z}'((1+r_f)p-\mu)((1+r_f)p-\mu)'\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\tilde{Z}'\Lambda\tilde{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\right)$$

$$+2(1+r_f)\gamma\ \text{tr}\left(\tilde{Z}'\Lambda\iota((1+r_f)p-\mu)'\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\right)$$

$$= \frac{K}{T}\bar{\phi} + N^2(1+r_f)^2\gamma^2\bar{\chi} + K\bar{\psi} + 2(1+r_f)\gamma K\bar{\omega} \tag{C.5}$$

where

- ϕ_i is the element in the i^{th} row and i^{th} column of $\tilde{Z}'\Lambda\tilde{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}$.
- $\chi_{i,j}$ is the element in the i^{th} row and j^{th} column of Λ.
- ψ_i is the element in the i^{th} row and i^{th} column of $\bar{Z}'((1+r_f)p-\mu)((1+r_f)p-\mu)'\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\tilde{Z}'\Lambda\tilde{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}$.
- ω_i is the element in the i^{th} row and i^{th} column of $\tilde{Z}'\Lambda\iota((1+r_f)p-\mu)'\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}$.

and

$$\bar{\phi} = \frac{1}{K}\sum_{i=1}^{K}\phi_i,\ \ \bar{\chi} = \frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N}\chi_{i,j},\ \ \bar{\psi} = \frac{1}{K}\sum_{i=1}^{K}\psi_i,\ \ \bar{\omega} = \frac{1}{K}\sum_{i=1}^{K}\omega_i$$

C.1.6 Proof of Proposition 2.3

If the historical price equals the zero α price of equation (3.1), then $(1+r_f)p = \mu - \gamma\Lambda\iota$. This of course means that $(1+r_f)P = U - \gamma\lambda$. Thus

$$\tilde{Z} = (1+r_f)Z_I + (U-\gamma\lambda)Z_S = (1+r_f)(Z_I + PZ_S) = (1+r_f)\bar{Z}$$

Plugging this into (C.5) above yields

$$\mathbb{E}[l_t'\Lambda l_t] = \frac{1}{T}\text{tr}\left(\tilde{Z}'\Lambda\tilde{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\right) + (1+r_f)^2\gamma^2\iota'\Lambda\iota$$

$$+\text{tr}\left(\bar{Z}'((1+r_f)p-\mu)((1+r_f)p-\mu)'\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\tilde{Z}'\Lambda\tilde{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\right)$$

$$+2(1+r_f)\gamma\ \text{tr}\left(\tilde{Z}'\Lambda\iota((1+r_f)p-\mu)'\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\right)$$

$$= \frac{1}{T}(1+r_f)^2 \mathrm{tr}\left(\bar{Z}'\Lambda\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\right) + (1+r_f)^2\gamma^2\iota'\Lambda\iota$$

$$+(1+r_f)^2\gamma^2 \mathrm{tr}\left(\bar{Z}'\Lambda\iota\iota'\Lambda\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\bar{Z}'\Lambda\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\right)$$

$$-2(1+r_f)^2\gamma^2 \ \mathrm{tr}\left(\bar{Z}'\Lambda\iota\iota'\Lambda\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\right)$$

$$= \frac{K}{T}(1+r_f)^2 + (1+r_f)^2\gamma^2\left(N^2\bar{\chi} - \mathrm{tr}\left(\bar{Z}'\Lambda\iota\iota'\Lambda\bar{Z}(\bar{Z}'\Lambda\bar{Z})^{-1}\right)\right)$$

$$= \frac{K}{T}(1+r_f)^2 + (1+r_f)^2\gamma^2\left(N^2\bar{\chi} - \mathrm{tr}\left(\bar{Z}'\Lambda\iota\iota'\Lambda\bar{Z}LL'\right)\right)$$

where $LL' = (\bar{Z}'\Lambda\bar{Z})^{-1}$ is the cholesky decomposition of $(\bar{Z}'\Lambda\bar{Z})^{-1}$. Thus

$$\mathbb{E}[l_t'\Lambda l_t] = \frac{K}{T}(1+r_f)^2 + (1+r_f)^2\gamma^2\left(N^2\bar{\chi} - \mathrm{tr}\left(L'\bar{Z}'\Lambda\iota\iota'\Lambda\bar{Z}L\right)\right)$$

$$= \frac{K}{T}(1+r_f)^2 + (1+r_f)^2\gamma^2\left(N^2\bar{\chi} - K\bar{g}\right)$$

where the vector G is defined as

$$G = L'\bar{Z}'\Lambda\iota$$

and

$$\bar{g} = \frac{1}{K}\sum_{i=1}^{K} G_i^2$$

C.2 Additional Details

C.2.1 Theoretical Model Details

This appendix section has four parts: 1) a derivation of CARA incumbent demand, 2) a description of the incumbents-only equilibrium, 3) a description of demand and the combined equilibrium, and 4) a proof that a combined equilibrium exists.

CARA Incumbent Demand

Incumbent asset managers know that individual investors have CARA utility. Let x_t^j denote asset manager demand for any given investor that asks for investment. Let p_t be an N dimensional vector of prices, let r_f be a risk free rate, and let d_{t+1} be an N dimensional vector of dividends with a multivariate normal distribution as follows:

$$d_{t+1} \sim N(\mathbb{E}_t[d_{t+1}], \Lambda_t)$$

Dollar excess returns are $d_{t+1}-(1+r_f)p_t$, but incumbent asset manager j believes that expected excess dollar returns are $\mathbb{E}_t[d_{t+1}-(1+r_f)p_t+\delta_t^j]$ instead of $\mathbb{E}_t[d_{t+1}-(1+r_f)p_t]$.

Asset managers, endeavoring to provide the best portfolio for investors, choose x_t^j to maximize CARA utility demand:

$$x_t^j = \text{argmax}_{x_t^j} E_t[-e^{-\gamma^j (x_t^j)'(d_{t+1}-(1+r_f)p_t+(1+r_f)W_{0,t}^j+\delta_t^j)}]$$

where $W_{0,t}^j$ is the initial wealth of asset manager j in period t. Note that γ^j is allowed to differ from investor risk aversion coefficient γ.

Since d_{t+1} has a conditional multivariate normal distribution, we derive the classic CARA utility / mean-variance result:

$$x_t^j = \text{argmax}_{x_t^j} (x_t^j)'\mathbb{E}_t[d_{t+1}-(1+r_f)p_t+\delta_t] - \frac{\gamma^j}{2}(x_t^j)'\Lambda_t x_t^j$$

The first order condition is

$$\mathbb{E}_t[d_{t+1}-(1+r_f)p_t+\delta_t^j] - \gamma^j \Lambda_t x_t^j = 0$$

Thus

$$x_t^j = \frac{1}{\gamma^j}(\Lambda_t)^{-1}\,\mathbb{E}_t[d_{t+1}-(1+r_f)p_t+\delta_t^j]$$

Incumbents-only Equilibrium Details

Let incumbents be indexed by $j \in [0,\bar{j}]$, and let investors be indexed by $n \in [0,\bar{n}]$. Incumbent asset managers know that investors have CARA utility. Incumbent asset managers endeavor to deliver this return. Let x_t^j be the N dimensional vector of asset demand for asset manager j, as described in Appendix C.2.1. The cost of the portfolio for an investor for asset manager j's portfolio is $(x_t^j)'p_t$, which is just the portfolio cost. I assume perfect competition leads incumbents to charge fees of zero. Although there is heterogeneity of incumbent managers, for each manager j, there is enough identical managers to j that compete away all management fees. The payouts investors receive are just portfolio payouts: $(x_t^j)'d_{t+1}$.

Investors can invest with incumbent asset manager portfolios or the risk free rate. Let A_n be a Lebesgue measurable subset of incumbent asset managers $[0, \bar{j}]$ that investor n chooses to invest with. Define

$$\mathbb{1}_{A_n}(x) = \begin{cases} 1 & \text{if } x \in A_n \\ 0 & \text{otherwise} \end{cases}$$

And define $x_{I,t}$, γ^J, and δ_t as

$$x_{I,t} = \frac{1}{\bar{j}} \int_{[0,\bar{j}]} x_t^j dj$$

$$\gamma^J = \left(\frac{1}{\bar{j}} \int_{[0,\bar{j}]} \frac{1}{\gamma^j} dj \right)^{-1}$$

$$\delta_t = \frac{1}{\bar{j}} \int_{[0,\bar{j}]} \frac{\gamma^J}{\gamma^j} \delta_t^j dj$$

With these previous three definitions, $x_{I,t}$ can be written as

$$x_{I,t} = \frac{1}{\gamma^J} (\Lambda_t)^{-1} \mathbb{E}_t[d_{t+1} - (1 + r_f)p_t + \delta_t] \tag{C.6}$$

Aggregate asset demand is defined as

$$y_t = \int_{[0,\bar{n}]} \int_{[0,\bar{j}]} x_t^j \mathbb{1}_{A_n}(j) dj dn \tag{C.7}$$

This can be intuitively thought of as the sum of demand of the incumbent asset manager that are chosen by investors.

With these definitions, I make the following two assumptions:

Assumption 1: Incumbent Even Spreading

Let A_n by any Lebesgue measurable subset of $[0, \bar{j}]$ with measure u_n, then assume

$$\gamma^J = \left(\frac{1}{u_n} \int_{[0,\bar{j}]} \frac{1}{\gamma^j} \mathbb{1}_{A_n}(j) dj \right)^{-1}$$

$$\delta_t = \frac{1}{u_n} \int_{[0,\bar{j}]} \frac{\gamma^J}{\gamma^j} \delta_t^j \mathbb{1}_{A_n}(j) dj$$

This assumption assumes that investors are "evenly spread" such that no matter what subset of incumbents an investor chooses, the portfolio returns will be the same. Mathematically, this assumption yields the result that for any any Lebesgue measurable subset A_n of $[0, \bar{j}]$ with measure u_n

$$x_{I,t} = \frac{1}{u_n} \int_{[0,\bar{j}]} x_t^j \mathbb{1}_{A_n}(j) dj$$

Assumption 2: Average Belief Wedge is Zero

More formally, $\frac{1}{N} \sum_{i=1}^{N} \delta_{i,t} = 0$ where $\delta_{i,t}$ is the i^{th} element of δ_t. Let ι be an N dimensional vector of ones. This assumption is identical to assuming that $\iota' \delta_t = 0$. Thus while CARA incumbents make mistakes, the average mistake is zero.

Assumption 1 gives us the result that the subset of incumbent managers that a given investor chooses does not affect the portfolio, we can rewrite demand without loss of generality as

$$y_t = \int_{[0,\bar{n}]} \int_{[0,\theta_{I,t}^n]} x_t^j dj dn$$

where $0 \leq \theta_{I,t}^n \leq \bar{j}$ is the scalar representing the mass of incumbent managers that investor n choose to as their asset managers. Since investors are identical, we can just write $\theta_{I,t}$ instead of $\theta_{I,t}^n$, and we can rewrite aggregate demand, using Assumption 1, as

$$y_t = \theta_{I,t} \bar{n} x_{I,t} \tag{C.8}$$

Investors choose $\theta_{I,t}$ to maximize their CARA utility with risk aversion coefficient γ as follows:

$$\theta^*_{I,t} = \text{argmax}_{\theta_{I,t}} \mathbb{E}_t[-e^{-\gamma\theta_{I,t}\bar{n}x'_{I,t}(d_{t+1}-(1+r_f)p_t)}]$$

such that

$$0 \leq \theta_{I,t} \leq \bar{j}$$

Consider the unconstrained problem (ignoring the constraint that $0 \leq \theta_{I,t} \leq \bar{j}$). Since the incumbent portfolio return is normally distributed, this problem is the same as

$$\theta^*_{I,t} = \text{argmax}_{\theta_{I,t}} \theta_{I,t}\bar{n}x'_{I,t}(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t) - \frac{\gamma}{2}\theta^2_{I,t}\bar{n}^2 x'_{I,t}\Lambda_t x_{I,t}$$

This yields the first order condition:

$$\bar{n}x'_{I,t}(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t) - \gamma\theta_{I,t}\bar{n}^2 x'_{I,t}\Lambda_t x_{I,t} = 0$$

Solving for $\theta_{I,t}$ means that

$$\theta_{I,t} = \frac{x'_{I,t}(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t)}{\gamma\bar{n}x'_{I,t}\Lambda_t x_{I,t}}$$

Plugging in $x_{I,t}$ from equation (C.6) yields

$$\theta_{I,t} = \frac{\gamma^J(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t + \delta_t)'(\Lambda_t)^{-1}(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t)}{\gamma\bar{n}(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t + \delta_t)'(\Lambda_t)^{-1}(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t + \delta_t)}$$

$$= \left(\frac{\gamma^J}{\gamma\bar{n}}\right)\frac{(\mathbb{E}_t[d_{t+1}]-(1+r_f)p_t+\delta_t)'(\Lambda_t)^{-1}(\mathbb{E}_t[d_{t+1}]-(1+r_f)p_t)}{(\mathbb{E}_t[d_{t+1}]-(1+r_f)p_t+\delta_t)'(\Lambda_t)^{-1}(\mathbb{E}_t[d_{t+1}]-(1+r_f)p_t)+(\mathbb{E}_t[d_{t+1}]-(1+r_f)p_t+\delta_t)'(\Lambda_t)^{-1}\delta_t} \quad \text{(C.9)}$$

As discussed in the body of the paper, supply of each asset is normalized to unity. This enables the interpretation of prices p_t to be the market equity of the asset for stocks. Thus equilibrium is defined as

$$\iota = y_t$$

Plugging in equations (C.8) and (C.6), this can be rewritten as

$$\iota = \frac{\theta_{I,t}\bar{n}}{\gamma^J}\left(\Lambda_t\right)^{-1}\mathbb{E}_t[d_{t+1} - (1+r_f)p_t + \delta_t]$$

which implies

$$\frac{\gamma^J}{\theta_{I,t}\bar{n}}\iota = \left(\Lambda_t\right)^{-1}\mathbb{E}_t[d_{t+1} - (1+r_f)p_t + \delta_t]$$

Thus

$$\left(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t + \delta_t\right)'\left(\Lambda_t\right)^{-1}\delta_t = \frac{\gamma^J}{\theta_{I,t}\bar{n}}\iota'\delta_t = 0$$

where the assumption $\iota'\delta_t = 0$ is used above. Thus plugging in this result into the denominator of equation (C.9) yields

$$\theta_{I,t} = \left(\frac{\gamma^J}{\gamma\bar{n}}\right)\frac{\left(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t + \delta_t\right)'\left(\Lambda_t\right)^{-1}\left(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t\right)}{\left(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t + \delta_t\right)'\left(\Lambda_t\right)^{-1}\left(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t\right)} = \frac{\gamma^J}{\gamma\bar{n}}$$

Recalling the constraint $0 \leq \theta_{I,t} \leq \bar{j}$, thus the unconstrained problem has the same solution to the constrained problem as long as $\frac{\gamma^J}{\gamma\bar{n}} \leq \bar{j}$. If this holds, then demand is

$$y_t = \frac{\gamma^J}{\gamma^2\bar{n}}\left(\Lambda_t\right)^{-1}\left(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t + \delta_t\right)$$

Thus any demand function of this form can be replicated by normalizing the variables $\bar{j} = 1$, $\bar{n} = 1$, and $\gamma^J = \gamma$, and changing γ appropriately. Thus in order to simplify this equation, I make the following assumption:

Assumption 3: Parameter Normalization

Assume $\bar{j} = 1$, $\bar{n} = 1$, and $\gamma^J = \gamma$.

Given these normalizations, we have

$$y_t = \frac{1}{\gamma}\left(\Lambda_t\right)^{-1}\left(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t + \delta_t\right)$$

or simply

$$y_t = x_{I,t}$$

This assumption also means that

$$\theta_{I,t} = 1$$

with this normalization, $\theta_{I,t}$ can be interpreted as the fraction of asset managers that investors choose to invest with in equilibrium. In this incumbents-only equilibrium, $\theta_{I,t} = 1$, which means each investor fully invests with incumbents, or other words that they invest with all incumbent asset managers.

Thus the CARA-incumbents-only equilibrium is the price p_t that solves

$$\iota = \frac{1}{\gamma} (\Lambda_t)^{-1} (\mathbb{E}_t[d_{t+1}] - (1 + r_f)p_t + \delta_t) \tag{C.10}$$

Combined Equilibrium

In this section, the same notation and assumptions from Appendix C.2.1. are used. Here, y_t is aggregate risky asset demand with learners as well, not just incumbents.

Define learner demand for learner j to be

$$x^j_{L,t} = \frac{1}{\gamma}(a^j_{L,t} + J^j_{L,t}(p_t - \bar{p}_t))$$

where $a^j_{L,t}$ and $J^j_{L,t}$ are $N \times 1$ and $N \times N$ matrices respectively. There is a unit mass of learners in this model, so that if investors demand, there is exactly enough learner asset managers to completely replace incumbent asset managers.

Define

$$a_{L,t} = \int_{[0,1]} a^j_{L,t} dj$$

$$J_{L,t} = \int_{[0,1]} J^j_{L,t} dj$$

$$x_{L,t} = \frac{1}{\gamma} (a_{L,t} + J_{L,t}(p_t - \bar{p}_t))$$

Aggregate demand in the full model is

$$y_t = \int_{[0,1]} \int_{[0,1]} x^j_{I,t} \mathbb{1}_{A_n}(j) dj dn + \int_{[0,1]} \int_{[0,1]} x^j_{L,t} \mathbb{1}_{B_n}(j) dj dn$$

where B_n is the Lebesgue measurable subset of the set of learners, $[0,1]$, that investor n chooses to invest with in equilibrium. We assume that investors choose A_n and B_n to maximize CARA utility with risk aversion coefficient γ. The N dimensional vector $x^j_{L,t}$ is the j^{th} learner's demand function. Define

$$x_{L,t} = x^j_{L,t}$$

The following assumption is analogous to Assumption 1 above:

Assumption 4: Learner Even Spreading

Assume that for every Lebesgue measurable set B_n, that is a subset of $[0,1]$, with measure u_n,

$$a_{L,t} = \frac{1}{u_n} \int_{[0,1]} a^j_{L,t} \mathbb{1}_{B_n}(j) dj$$

$$J_{L,t} = \frac{1}{u_n} \int_{[0,1]} J^j_{L,t} \mathbb{1}_{B_n}(j) dj$$

This assumption, along with Assumption 1, and the fact that investors are identical allows demand to be rewritten as

$$y_t = \theta_{I,t} x_{I,t} + \theta_{L,t} x_{L,t}$$

where $\theta_{I,t}$ and $\theta_{L,t}$ are chosen to maximize CARA utility with the constraints

$$0 \leq \theta_{I,t} \leq 1, \quad 0 \leq \theta_{L,t} \leq 1$$

where the constraints are due to the fact that there is a unit mass of incumbents and learners each.

Define $q_{I,t}(p_t) = \gamma x_{I,t}(p_t)$ and $q_{L,t}(p_t) = \gamma x_{L,t}(p_t)$. Thus these are not functions of γ. Define X_t and Q_t to be the $N \times 2$ dimensional matrices

$$X_t = [x_{I,t}(p_t) \;\; x_{L,t}(p_t)] \text{ and } Q_t = [q_{I,t}(p_t) \;\; q_{L,t}(p_t)]$$

Investors choose how much to invest with each type of asset manager, $\theta_t = (\theta_{I,t}, \theta_{L,t})'$, in order to maximize CARA utility:

$$\mathbb{E}_t[-e^{-\gamma\left(\theta' X_t'\left(d_{t+1}-(1+r_f)p_t\right)\right)}]$$

subject to

$$0 \leq \theta_{I,t} \leq 1, \quad 0 \leq \theta_{L,t} \leq 1$$

This translates to solving the problem:

$$\theta_t' X_t'(\mathbb{E}_t[d_{t+1}] - (1 + r_{f,t})p_t^*) - \frac{\gamma}{2}\theta_t' X_t' \Lambda_t X_t \theta_t$$

subject to

$$0 \leq \theta_{I,t} \leq 1, \quad 0 \leq \theta_{L,t} \leq 1$$

This solution to optimizing this objective function is the same as optimizing

$$\theta_t' Q_t'(\mathbb{E}_t[d_{t+1}] - (1 + r_{f,t})p_t^*) - \frac{1}{2}\theta_t' Q_t' \Lambda_t Q_t \theta_t$$

subject to

$$0 \leq \theta_{I,t} \leq 1, \quad 0 \leq \theta_{L,t} \leq 1$$

which is not a function of γ.

This problem can be easily solved numerically by checking all possible solutions given the constraints.

Existence of an Equilibrium

This section uses the notation from Appendix C.2.1 above. Both types of incumbents, CARA and estimated demand, have demand that is linear in prices. Thus, $q_{I,t}$ can be written as

$$q_{I,t} = a_{I,t} + J_{I,t}(p_t - \bar{p}_t)$$

where $a_{I,t}$ and $J_{I,t}$ the relevant intercept $N \times 1$ column matrix and $N \times N$ slope matrix. As discussed above, $x_{I,t} = q_{I,t}/\gamma$. Define $E = [\epsilon_1, 1] \times [\epsilon_2, 1]$ where either $\epsilon_1 > 0$ or $\epsilon_2 > 0$ but $\epsilon_1 < 1$ and $\epsilon_2 < 1$. Equilibrium is defined as the $\theta_t = (\theta_{I,t}, \theta_{L,t})' \in E$ that solves the above investor problem and the price p_t that solves

$$\iota = \frac{\theta_{I,t}}{\gamma} \left(a_{I,t} + J_{I,t}(p_t - \bar{p}_t)\right) + \frac{\theta_{L,t}}{\gamma} \left(a_{L,t} + J_{L,t}(p_t - \bar{p}_t)\right) \tag{C.11}$$

Note that θ_t is restricted to E. Assume for any price p_t that $a_{I,t} + J_{I,t}(p_t - \bar{p}_t)$ and $a_{L,t} + J_{L,t}(p_t - \bar{p}_t)$ are linearly independent, and that for any $\theta_t \in E$, $\theta_{I,t} J_{I,t} + \theta_{L,t} J_{L,t}$ is invertible. Then an equilibrium, as defined above, exists.

Proof

Define

$$g_t(\theta_t) = \bar{p}_t + \left(\theta_{I,t} J_{I,t} + \theta_{L,t} J_{L,t}\right)^{-1} \left(\gamma\iota - \theta_{I,t} a_{I,t} - \theta_{L,t} a_{L,t}\right)$$

Thus for any $\theta_t \in E$, $g_t(\theta_t) = p_t$ is the price that solves equation (C.11). Define

$$h_t(p_t) = \text{argmax}_{\theta_t \in E} \;\; \theta_t' Q_t' (\mathbb{E}_t[d_{t+1}] - (1 + r_{f,t}) p_t) - \frac{1}{2} \theta_t' Q_t' \Lambda_t Q_t \theta_t$$

Because $a_{I,t} + J_{I,t}(p_t - \bar{p}_t)$ and $a_{L,t} + J_{L,t}(p_t - \bar{p}_t)$ are linearly independent, Q has full column rank. Thus $Q' \Lambda_t Q$ is positive definite. This means the above objective function is strictly concave, which means the argument maximum above is unique. Since the objective function above is also continuous, by Claude Berge's maximum theorem, h_t is continuous.

Define $f_t(\theta_t) = h_t(g_t(\theta))$. Since h_t and g_t are continuous on their domains, f_t is continuous on E. If I show that there exists a $\theta_t \in E$ such that $f_t(\theta_t) = \theta_t$, then an equilibrium exists. Since f_t maps from E to E and is continuous, such an equilibrium exists by Brouwer's fixed-point theorem.

C.2.2 Learner Capital Percent

This section of the appendix shows how the learner capital percentage is calculated.

While $\theta_{L,t}$ is the fraction of learners that investors invest with, then

$$\theta_{L,t} x'_{L,t} p_t$$

The percentage of total market capital is

$$\theta_{L,t} \frac{x'_{L,t} p_t}{\iota' p_t}$$

Occasionally, large negative positions skew the average capital negatively with demand functions without prior parameter restrictions (no shrinking). In order to drop these outliers, the tables report learner capital percentage as

$$\text{\% Learner Capital} = \frac{1}{T} \sum_{t=1}^{T} \theta_{L,t} \text{min} \left(\max \left(\frac{x'_{L,t} p_t}{\iota' p_t}, 0 \right), 1 \right)$$

C.2.3 Transformation of Characteristics

Denote $\tilde{z}_t^k$ as the k^{th} characteristic of N assets at time t. Thus $\tilde{z}_t^k$ is an N dimensional column vector, which could contain, for example, the dollar book value of assets known at time t for each asset. Let percentile($\cdot$) denote the function that converts each element of the input vector to percentiles between 0 and 1. This papers uses the following transformation, used in both Kozak et al. (2020) and Kelly et al. (2019):

$$z_t^k = \text{percentile}(\tilde{z}_t^k) - 0.5\iota$$

Thus the resulting $N \times K$ dimensional matrix Z_t of K characteristics of the N assets is filled with these transformed columns z_t^k.

For more discussion on why this transformation is important see Kozak et al. (2020) and Kelly et al. (2019).

C.2.4 General Neural Network

In the body of the paper, it is assumed as in Koijen and Yogo (2019) that

$$\mu_t = Z_t \pi$$

$$\Sigma_t = \Gamma_t \Gamma_t' + \zeta I, \quad \Gamma_t = Z_t \phi$$

where ζ is a scalar, π and ϕ are $K \times 1$ vectors.

A general neural network, with activation functions $a(\cdot)$ and H hidden layers can be written as the series of equations

$$L_{1,t} = a(Z_t W_1), \ \ L_{2,t} = a(L_1 W_2 + b_2), \ \ L_{3,t} = a(L_2 W_3 + b_3), \ \ ...,$$

$$L_{H,t} = a(L_{H-1,t} W_H + b_H), \ \ L_{H+1,t} = a(L_H W_{H+1,t} + b_{H+1})$$

where $L_{h,t}$, W_h, and b_h are the layers, weights, and biases respectively. Note that there is no bias in the first equation because Z_t has a column of ones, making a bias term redundant. $L_{H+1,t}$ is the output layer. Assume that $L_{H+1,t}$ is $N \times (M+1)$ dimensional, where $M \geq 1$. Define μ_t to be the first column of $L_{H+1,t}$ and Γ_t to be the last M columns of $L_{H+1,t}$.

If there are zero hidden layers and $M = 1$, then this describes perfectly the simple linear model used in the paper. However, this neural network generalization gives greater flexibility in the followings ways:

1. The mean and covariance matrices can be highly nonlinear functions of the characteristics Z_t

2. Γ_t is allowed to be $N \times M$ dimensional, instead of just $N \times 1$ dimensional, which allows for greater covariance matrix flexibility.

C.2.5 Consistency of Main Econometric Model

Define

$$\mu_t = Z_t \pi$$

$$\Sigma_t = \Gamma_t \Gamma_t' + \zeta I, \quad \Gamma_t = Z_t \phi$$

$$\psi = \begin{bmatrix} \pi \\ \phi \\ \zeta \end{bmatrix}$$

Assume $y_1, ..., y_T$ is a sample of i.i.d. random vectors that are distributed according to a multivariate normal distribution as follows:

$$y_t \sim N(\mu_t, \Sigma_t)$$

Assume Z_t is has full column rank for all t. Let $l(\psi|y_t, Z_t)$ be the log of the multivariate normal probability density function parameterized as shown above. Assume a compact parameter space Ψ, such that for all $(\pi', \phi', \zeta)' \in \Psi$, $\zeta > 0$. Let ψ^0 denote the true data generating process parameters. Then the maximum likelihood estimator $\hat{\Psi}_T$ is defined as

$$\hat{\Psi}_T = \text{argmax}_{\psi \in \Psi} \sum_{t=1}^{T} l(\psi|y_t, Z_t)$$

Then this estimator $\hat{\Psi}_T$ is a consistent estimator, i.e. $\hat{\Psi}_T$ converges in probability to ψ^0 as $T \to \infty$.

proof

The parameter space is compact, and the probability density function is obviously continuous. Since the sample is i.i.d., if parameter identification is shown, then this estimator is consistent by the classic maximum likelihood estimator theorem.

Let $\psi_1 = (\pi_1', \phi_1', \zeta_1)' \in \Psi$ and $\psi_2 = (\pi_2', \phi_2', \zeta_2)' \in \Psi$ such that

$$l(\psi_1|y_t, Z_t) = l(\psi_2|y_t, Z_t)$$

for any y_t and matrix Z_t with full column rank. Multivariate normal distributions are uniquely identified by their mean and covariance parameter matrices. Thus the above condition holds if and only if

$$Z_t\pi_1 = Z_t\pi_2, \quad Z_t\phi_1\phi_1' Z_t' + \zeta_1 I = Z_t\phi_2\phi_2' Z_t' + \zeta_2 I$$

$$\iff \pi_1 = \pi_2, \quad (Z_t' Z_t)^{-1}(\zeta_1 - \zeta_2) + (\phi_1\phi_1' - \phi_2\phi_2') = 0$$

where the second equation above holds for any full column rank Z_t if and only if $\zeta_1 = \zeta_2$ and $\phi_1 = \phi_2$. Thus the parameters are identified and the estimator is consistent.

C.2.6 Matrix Identities for Econometric Method

Recall that in the paper, it is assumed

$$\mu_t = Z_t\pi$$

$$\Sigma_t = \Gamma_t\Gamma_t' + \zeta I_N, \quad \Gamma_t = Z_t\phi$$

where ζ is a scalar, π and ϕ are $K \times 1$ vectors. Note that the identity matrix I was replaced with I_N to denote that it is an $N \times N$ identity matrix. In this section, Γ_t is assumed to be a more general $N \times M$ dimensional, rather than the simple $N \times 1$ dimensional as assumed by Koijen and Yogo (2019). See Appendix for C.4 for details about a more general representation.

In order to calculate the maximum likelihood fucntion, Σ_t^{-1} and the determinant of Σ_t, denoted as $|\Sigma_t|$ need to be calculated efficiently, even with large N. This section shows formulas that can do this.

By the Woodbury Matrix Identity,

$$\Sigma_t^{-1} = \frac{1}{\zeta}\left(I_N - \Gamma_t(\zeta I_M + \Gamma_t'\Gamma_t)^{-1}\Gamma_t'\right)$$

Note that if $M = 1$, then $\zeta I_M + \Gamma_t'\Gamma_t$ is a scalar. If M is much smaller than N, then $(\zeta I_M + \Gamma_t'\Gamma_t)^{-1}$ is much easier to numerically calculate than Σ_t^{-1}.

By Sylvester's matrix identity,

$$|\Sigma_t| = \zeta^N \left| I_M + \frac{1}{\zeta}\Gamma_t'\Gamma_t \right|$$

Note that if $M = 1$, $I_M + \frac{1}{\zeta}\Gamma_t'\Gamma_t$ is a scalar. If M is much smaller than N, then the determinant of $I_M + \frac{1}{\zeta}\Gamma_t'\Gamma_t$ is much easier to numerically calculate than the determinant of Σ_t.

C.2.7 Calculating the Value of the Market

In this appendix section, we prove that formula given in the paper:

$$\iota' p_t = \frac{\mathbb{E}_t[r_{M,t+1}]}{\gamma \mathrm{Var}_t[r_{M,t+1}]}$$

This section uses the identities and definitions found in Table B.1 heavily. Recall from equation (C.10) that the CARA-incumbents-only equilibrium equation is

$$\iota = \frac{1}{\gamma} (\Lambda_t)^{-1} \left(\mathbb{E}_t[d_{t+1}] - (1 + r_f)p_t + \delta_t\right)$$

Multiplying both sides by $\gamma\iota'\Lambda_t$ yields

$$\gamma\iota'\Lambda_t\iota = \iota' \left(\mathbb{E}_t[d_{t+1}] - (1 + r_f)p_t + \delta_t\right)$$

Using the assumption that $\iota'\delta_t = 0$ from Appendix C.2.1, this equation can be simplified to

$$\gamma\iota'\Lambda_t\iota = \iota' \left(\mathbb{E}_t[d_{t+1}] - (1 + r_f)p_t\right) \tag{C.12}$$

Using the simple identities from Table B.1, we can rewrite

$$\iota'\Lambda_t\iota = \iota'(P_t\Lambda_tP_t)\iota = p_t'\Lambda_tp_t = (\iota'p_t)^2 w_{M,t}'\Lambda w_{M,t} = (\iota'p_t)^2 \mathrm{Var}_t(r_{M,t+1}) \tag{C.13}$$

and

$$\iota' \left(\mathbb{E}_t[d_{t+1}] - (1+r_f)p_t\right) = \iota'\mathbb{E}_t[r_{t+1}] = \iota' P_t P_t^{-1}\mathbb{E}_t[r_{t+1}] = p_t'\mathbb{E}_t[r_{t+1}]$$

$$= (\iota' p_t) w_{M,t}'\mathbb{E}_t[r_{t+1}] = (\iota' p_t)\mathbb{E}_t[r_{M,t+1}] \tag{C.14}$$

Plugging in equations (C.13) and (C.14) into equation (C.12), we get

$$\gamma(\iota' p_t)^2 \mathrm{Var}_t(r_{M,t+1}) = (\iota' p_t)\mathbb{E}_t[r_{M,t+1}]$$

Solving for the total value of the asset market, $\iota' p_t$, yields

$$\iota' p_t = \frac{\mathbb{E}_t[r_{M,t+1}]}{\gamma \mathrm{Var}_t[r_{M,t+1}]}$$

C.2.8 Details of Example of Increasing Mispricing

In this section I give the parameters used to generate the example that shows how mispricing can increase in Section 4.

In this example, $\Gamma = (0.1, 0.1, 0.1)'$ and $\zeta = 0.01$. The size vector $s = (1, 2, 3)'$. Like in the paper, $S = \mathrm{diag}(s)$, and $\Lambda = S\Sigma S$ where $\Sigma = \Gamma\Gamma' + \zeta I$. Also $\gamma = 1$ and $r_f = 0$. Incumbent prices are $\bar{p} = (1, 2, 3)'$ and incumbent alpha are $\bar{\alpha} = (0.03, -0.18, 0.15)'$. Fixing these variables pins down expected dividends, $\mathbb{E}[d]$, with the equation

$$\mathbb{E}[d] = \gamma \Lambda_d \iota + (1 + r_f)\bar{p} + \bar{P}\alpha$$

where $\bar{P} = \mathrm{diag}(\bar{p})$.

I assume that incumbent demand and learner demand is linear in price:

$$x_I(p) = \frac{1}{\gamma}(a_I + \mathrm{diag}(b_I)(p - \bar{p})) \text{ and } x_L(p) = \frac{1}{\gamma}(a_L + \mathrm{diag}(b_L)(p - \bar{p}))$$

where x_I and x_L are incumbent and learner demand respectively, and a_I, b_I, a_L, and b_L are 3×1 column vectors. I assume that $b_I = (-1, -35, -20)'$ and $b_L = (-2, -35, -20)'$.

Fixing this slope for incumbents pins down a_I with the equation $a_I = \gamma\iota$. I set the learner intercept such that demand at incumbents-only prices is mean-variance optimal:

$$a_L = (\Lambda)^{-1}(\mathbb{E}[d] - (1 + r_f)\bar{p})$$

C.2.9 Overpricing and Information Ratio Components

In this section, I define overpricing and learner information ratio asset components as shown in Panels D and E in Figure A.4.

For asset demand x, the information ratio of the portfolio is

$$\text{IR}_x = \frac{x'\alpha}{\sqrt{x'\Lambda x - (x'\beta)^2\iota'\Lambda\iota}}$$

The asset specific component of this quantity, l_i^x, is defined as

$$l_i^x = \frac{x_i\alpha_i}{\sqrt{x'\Lambda x - (x'\beta)^2\iota'\Lambda\iota}}$$

thus

$$\text{IR}_x = \sum_{i=1}^{N} l_i^x$$

Let x^* be the (unscaled) mean-variance optimal demand:

$$x^* = (\Lambda)^{-1}\,\mathbb{E}[r]$$

where r is the vector of excess dollar returns.

Then mispricing is

$$\xi = \sqrt{(\alpha)'(\Lambda_d)^{-1}(\alpha)} = \frac{(x^*)'\alpha}{\sqrt{(x^*)'\Lambda(x^*) - ((x^*)'\beta)^2\iota'\Lambda\iota}} = \sum_{i=1}^{N} l_i^{x^*}$$

Overpricing for asset i, denoted as ξ_i, is

$$\xi_i = \frac{-\text{sign}(\alpha_i)\left(l_i^{x^*}\right)^2 \xi}{\sum_{n=1}^{N}\left(l_n^{x^*}\right)^2}$$

Thus

$$\xi = \sum_{i=1}^{N} |\xi_i|$$

C.2.10 Brandt et al. (2009) and Kozak et al. (2020) Estimator

This section of the appendix defines the Brandt et al. (2009) estimator b_T^* and the Kozak et al. (2020) estimator $\hat{b}_T$, and shows they are equivalent given some assumptions.

In both Brandt et al. (2009) and Kozak et al. (2020), the authors use excess returns r_t. For this paper, this should be replaced by scaled dollar excess returns,

$$r_{t+1}^s = S_t^{-1}(d_{t+1} - (1 + r_f)p_t)$$

Following Kozak et al. (2020), I define the $K + 2$ dimensional vector of portfolio returns as

$$F_{t+1} = \hat{Z}_t' r_{t+1}^s$$

Let $\hat{\omega}$ and $\hat{\Omega}$ be the usual estimates of the mean and covariance of F_{t+1}, that is

$$\hat{\omega}_T = \frac{1}{T}\sum_{t=1}^{T} F_t$$

$$\hat{\Omega}_T = \frac{1}{T}\sum_{t=1}^{T}(F_t - \hat{\omega}_T)(F_t - \hat{\omega}_T)'$$

The Kozak et al. (2020) estimate $\hat{b}_T$ is found by

$$\hat{b}_T = \text{argmin}_b \; (\hat{\omega}_T - \hat{\Omega}_T b)'\hat{\Omega}_T^{-1}(\hat{\omega}_T - \hat{\Omega}_T b) + \lambda_1||b||_1 + \lambda_2||b||_2$$

For now, ignore the shrinking parameters by setting $\lambda_1 = \lambda_2 = 0$. Then

$$\hat{b}_T = \hat{\Omega}_T^{-1}\hat{\omega}_T \tag{C.15}$$

In Brandt et al. (2009), the authors assume that mean-variance efficient weights are linear in characteristics. So mean-variance weights can be written as

$$w_t = \hat{Z}_t b \tag{C.16}$$

The empirical Sharpe ratio of this portfolio is

$$\frac{\frac{1}{T}\sum_{t=1}^{T} w_t' r_{t+1}^s}{\sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(w_t' r_{t+1}^s - \frac{1}{T}\sum_{\tau=1}^{T} w_\tau' r_{\tau+1}^s\right)^2}}$$

Plugging in equation (C.16) yields

$$\frac{\frac{1}{T}\sum_{t=1}^{T} b'\hat{Z}_t' r_{t+1}^s}{\sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(b'\hat{Z}_t' r_{t+1}^s - \frac{1}{T}\sum_{\tau=1}^{T} b'\hat{Z}_\tau' r_{\tau+1}^s\right)^2}}$$

$$= \frac{b'\left(\frac{1}{T}\sum_{t=1}^{T}\hat{Z}_t' r_{t+1}^s\right)}{\sqrt{b'\left(\frac{1}{T}\sum_{t=1}^{T}\left(\hat{Z}_t' r_{t+1}^s - \frac{1}{T}\sum_{\tau=1}^{T}\hat{Z}_\tau' r_{\tau+1}^s\right)\left(\hat{Z}_t' r_{t+1}^s - \frac{1}{T}\sum_{\tau=1}^{T}\hat{Z}_\tau' r_{\tau+1}^s\right)'\right) b}}$$

$$= \frac{b'\hat{\omega}_T}{\sqrt{b'\hat{\Omega}_T b}}$$

Brandt et al. (2009) suggest that b be selected to maximize this quantity. Note that they also fix weights on a benchmark portfolio, and this part is ignored here. Note that if b maximizes the above quantity kb will also maximize the above quantity for any positive scalar k. Thus, in order to get a unique solution, we solve

$$\max_b \ \frac{b'\hat{\omega}_T}{\sqrt{b'\hat{\Omega}_T b}}$$

subject to

$$\iota'\hat{\Omega}_T^{-1}\hat{\omega}_T = \iota' b \tag{C.17}$$

where this choice of the sum of b is chosen for convenience, as will be made more clear below.

Let b_T^* be the Brandt et al. (2009) vector b that is the solution to this problem. The first order condition to this problem is

$$\frac{\hat{\omega}_T}{\sqrt{(b_T^*)'\hat{\Omega}_T b_T^*}} - \frac{(b_T^*)'\hat{\omega}_T\hat{\Omega}_T b_T^*}{\left((b_T^*)'\hat{\Omega}_T b_T^*\right)^{\frac{3}{2}}} = 0$$

Solving for b_T^* yields

$$b_T^* = \frac{1}{(b_T^*)'\hat{\omega}_T}\left((b_T^*)'\hat{\Omega}_T b_T^*\right)\hat{\Omega}_T^{-1}\hat{\omega}_T \tag{C.18}$$

Premultiplying both sides with ι' yields

$$\frac{1}{(b_T^*)'\hat{\omega}_T}\left((b_T^*)'\hat{\Omega}_T b_T^*\right)\hat{\iota}'\Omega_T^{-1}\hat{\omega}_T = \iota' b_T^* = \iota'\Omega_T^{-1}\hat{\omega}_T$$

where the constraint in equation (C.17) is used. This implies that

$$\frac{1}{(b_T^*)'\hat{\omega}_T}\left((b_T^*)'\hat{\Omega}_T b_T^*\right) = 1$$

Plugging this into equation (C.18) yields

$$b_T^* = \hat{\Omega}_T^{-1}\hat{\omega}_T$$

Thus the Brandt et al. (2009) estimate b_T^* that maximizes the Sharpe ratio is the same as the Kozak et al. (2020) estimate $\hat{b}_T$ with no shrinking from the penalty parameters. That is $b_T^* = \hat{b}_T$.

C.2.11 Brandt et al. (2009) and Kozak et al. (2020) Estimator Consistency Given Data Generating Process

The demand function for a machine learner who knows all the parameters is

$$\frac{1}{\gamma} S_t^{-1} \Sigma_t^{-1} (\mu_t - (1 + r_f) S_t^{-1} (p_t - \bar{p}_t))$$

$$= \frac{1}{\gamma} S_t^{-1} \left[\frac{1}{\zeta} (\mu_t - (1 + r_f) S_t^{-1} (p_t - \bar{p}_t)) + c_t \Gamma_t \right]$$

where

$$c_t = -\frac{\Gamma_t' (\mu_t - (1 + r_f) S_t^{-1} (p_t - \bar{p}_t))}{\zeta (\zeta + \Gamma_t' \Gamma_t)}$$

Consider the case where c_t is a constant c, like Koijen and Yogo (2019). Define the following:

$$\hat{Z}_t = \begin{bmatrix} S_t^{-1} p_t & S_t^{-1} \bar{p}_t & Z_t \end{bmatrix}, \quad \Pi = \begin{bmatrix} -(1 + r_f) \\ 1 + r_f \\ \pi \end{bmatrix}, \quad \Phi = \begin{bmatrix} 0 \\ 0 \\ \phi \end{bmatrix}, \quad b = \frac{1}{\zeta} \Pi + c \Phi$$

Then machine learner demand, with perfect knowledge of the parameters, is

$$\frac{1}{\gamma} S_t^{-1} \hat{Z}_t b$$

The purpose of this section is to prove that under some assumptions, that $\hat{b}_T$ from Kozak et al. (2020) and Brandt et al. (2009) described precisely the in the previous Appendix section is a consistent estimator of $b = \Pi/\zeta + c\Phi$. That is, the goal is to show that

$$\hat{b}_T \xrightarrow{\mathbb{P}} b$$

where $\xrightarrow{\mathbb{P}}$ is convergence in probability as T goes to infinity.

Theorem Assume that $\hat{Z}_t = \hat{Z}$ is constant, denoted just as $\hat{Z}$. Assume the same data generating process for r_t^s as described in the paper. Then

$$\hat{b}_T \xrightarrow{\mathbb{P}} b$$

Proof:

The variable c_t can be written as

$$c_t = -\frac{\Gamma_t'(\mu_t - (1+r_f)S_t^{-1}(p_t - \bar{p}_t))}{\zeta(\zeta + \Gamma_t'\Gamma_t)}$$

$$= -\frac{\Phi'\hat{Z}_t'\hat{Z}_t\Pi}{\zeta(\zeta + \Phi'\hat{Z}_t'\hat{Z}_t\Phi)}$$

$$= -\frac{\Phi'\hat{Z}'\hat{Z}\Pi}{\zeta(\zeta + \Phi'\hat{Z}'\hat{Z}\Phi)}$$

Thus c_t is in fact constant under these assumptions. It is denoted simply as c.

Directly calculating

$$\mathbb{E}_t[F_{t+1}] = \mathbb{E}_t[\hat{Z}_t' r_{t+1}^s] = \hat{Z}_t'\hat{Z}_t\Pi$$

Thus

$$\mathbb{E}[F_t] = \hat{Z}'\hat{Z}\Pi$$

$$\mathrm{Var}_t[F_{t+1}] = \hat{Z}_t'\mathrm{Var}_t[r_{t+1}^s]\hat{Z}_t = \hat{Z}_t'\Sigma_t\hat{Z}_t = \hat{Z}_t'\left(\zeta I + \Gamma_t\Gamma_t'\right)\hat{Z}_t$$

$$= \hat{Z}_t'\left(\zeta I + \hat{Z}_t\Phi\Phi'\hat{Z}_t'\right)\hat{Z}_t = \zeta\hat{Z}_t'\hat{Z}_t + \hat{Z}_t'\hat{Z}_t\Phi\Phi'\hat{Z}_t'\hat{Z}_t$$

Thus

$$\mathrm{Var}[F_t] = \zeta\hat{Z}'\hat{Z} + \hat{Z}'\hat{Z}\Phi\Phi'\hat{Z}'\hat{Z}$$

By the Woodbury Matrix Identity, we have

$$\left(\mathrm{Var}[F_t]\right)^{-1}\mathbb{E}[F_t] = \left(\zeta\hat{Z}'\hat{Z} + \hat{Z}'\hat{Z}\Phi\Phi'\hat{Z}'\hat{Z}\right)^{-1}(\hat{Z}'\hat{Z}\Pi)$$

$$= \left(\frac{1}{\zeta}(\hat{Z}'\hat{Z})^{-1} - \frac{1}{\zeta}(\hat{Z}'\hat{Z})^{-1}\hat{Z}'\hat{Z}\Phi\left(\zeta + \Phi'\hat{Z}'\hat{Z}(\hat{Z}'\hat{Z})^{-1}\hat{Z}'\hat{Z}\Phi\right)^{-1}\Phi'\hat{Z}'\hat{Z}(\hat{Z}'\hat{Z})^{-1}\right)(\hat{Z}'\hat{Z}\Pi)$$

$$= \frac{1}{\zeta}\Pi - \frac{1}{\zeta}\Phi\left(\zeta + \Phi'\hat{Z}'\hat{Z}\Phi\right)^{-1}\Phi'\hat{Z}'\hat{Z}\Pi$$

$$= \frac{1}{\zeta}\Pi + c\Phi = b$$

Thus by the weak law of large numbers and the continuous mapping theorem

$$\hat{b}_T \xrightarrow{\mathbb{P}} b$$

Thus the estimators $\hat{b}_T$ and b_T^* from Kozak et al. (2020) and Brandt et al. (2009) as described in the previous section is a consistent estimator of b under these assumptions.

C.2.12 1/N Demand Rescaling

The $1/N$ learner demand, as set out in the paper, is

$$\frac{\chi_t}{\gamma}\bar{P}_t^{-1}\iota$$

where χ_t is set so to inflate $1/N$ learner demand to have sufficient demand to invest in the full fraction of every asset if investors demand. Thus

$$\chi_t = \gamma \max(\bar{p}_t)$$

Thus the smallest value of $\frac{\chi_t}{\gamma}\bar{P}_t^{-1}\iota$ is one, which equals supply of that asset.

C.2.13 Random Forest Estimator

This section presupposed the reader is familiar with regression trees and regression tree terminology. For a basic reference, see Hastie et al. (2001).

Stack the data as follows

$$Z = \begin{bmatrix} Z_0 \\ Z_1 \\ \vdots \\ Z_{T-1} \end{bmatrix}, \quad R = \begin{bmatrix} r_1 \\ r_2 \\ \vdots \\ r_T \end{bmatrix}$$

In order to fit an estimate of the conditional mean, $\mathbb{E}_t[r_{t+1}]$, I fit a typical random forest regression using Z to predict R, which yields $\hat{\mu}_t$ for each t.

A estimate of the covariance matrix is the other component of estimated MVE weights. In order to do this, I set

$$\hat{\Sigma}_t = \frac{1}{S}\sum_{s=1}^{S}(g^s(Z_t)g^s(Z_t)' + \zeta^s I)$$

where and $\varphi^s(Z_t)$ is an $N \times 1$ dimensional output of the s^{th} regression tree in the forest represented by φ^s, and ζ^s is the estimate of ζ from the s^{th} tree in the forest.

We fit the regression tree, φ^s, as described below. In order to simplify notation, I drop s and simply write φ. Let $\varphi_i = \varphi_i(Z_t)$ be the element of $\varphi(Z_t)$ corresponding to the i^{th} element (asset) of $\varphi(Z_t)$. Let $Z_{i,t}$ be the i^{th} row of Z_t, and let $Z_{i,k,t}$ correspond to the i^{th} row and k^{th} column of Z_t.

Consider a given tree represented by φ, with $\bar{L}$ leaves (terminal nodes), denoted as

$$\varphi_i(Z_t) = \sum_{l=1}^{\bar{L}} \hat{\varphi}_l \mathbb{1}(Z_{i,t} \in C_l)$$

where C_l is the partition of the data represented by leaf l, and $\hat{\varphi}_l$ represented the constant scalar associated with partition l.

I describe below how a given leaf l, can be "grown" into a decision node with subsequent left leaf L and right leaf R. Since fitting a regression tree is the process of recursively growing leaves into decision nodes with two subsequent leaves, the following process can be recursively applied to fit an entire tree.

For each of the K columns of Z_t, consider a range of candidate split values $c_1^l(k), c_2^l(k), ..., c_M^l(k)$. Consider a given split value $c_j^l(k)$. For each candidate split, define the N dimensional vectors $\varphi_{-l,t}$,

$\iota_{L,t}$, and $\iota_{R,t}$, with corresponding i^{th} elements denoted as $\varphi_{i,-l,t}$, $\iota_{i,L,t}$, and $\iota_{i,R,t}$, as

$$\varphi_{i,-l,t} = \sum_{j=1,j\neq l}^{L} \hat{\varphi}_j \mathbb{1}(Z_{i,t} \in C_j)$$

$$\iota_{i,L,t} = \mathbb{1}(Z_{i,t} \in C_l \text{ and } Z_{i,k,t} \leq c_j^l(k))$$

$$\iota_{i,R,t} = \mathbb{1}(Z_{i,t} \in C_l \text{ and } Z_{i,k,t} > c_j^l(k))$$

Define

$$\bar{\varphi}_t = \varphi_{-l,t} + \varphi_L \iota_{L,t} + \varphi_R \iota_{R,t}$$

where φ_L and φ_R are scalars chosen for this given candidate split. If this candidate split is chosen among all the other candidates, then leaf l will become a decision node with left leaf L and right leaf R. We'll set $\hat{\varphi}_L$ to the fitted value of φ_L as described below. Similarly, we'll set $\hat{\varphi}_R$ to the fitted value of φ_R as described below. Also, for any given $Z_{i,t}$, we can define C_L and C_R such that

$$Z_{i,t} \in C_L \text{ if and only if } Z_{i,t} \in C_l \text{ and } Z_{i,k,t} \leq c_j^l(k)$$

$$Z_{i,t} \in C_R \text{ if and only if } Z_{i,t} \in C_l \text{ and } Z_{i,k,t} > c_j^l(k)$$

This occurs only if this candidate split is chosen among all candidate splits according to the criteria below.

The covariance matrix for this candidate split is

$$\Sigma_t = \bar{\varphi}_t \bar{\varphi}_t' + \zeta I$$

Define

$$\tilde{\Sigma}_t = y_t y_t'$$

where $y_t = r_{t+1} - \hat{\mu}_t$.

For each candidate split, we choose φ_L, φ_R, and ζ to minimize

$$\sum_t \sum_i \sum_j (\tilde{\Sigma}_{i,j,t} - \Sigma_{i,j,t})^2$$

where $\Sigma_{i,j,t}$ is the element of Σ_t in the i^{th} row of the j^{th} column. $\tilde{\Sigma}_{i,j,t}$ is similar.

This is the same as minimizing

$$\sum_t \left|\left|\tilde{\Sigma}_t - \Sigma_t\right|\right|_F^2$$

where $||\cdot||_F$ is the Frobenius matrix norm. Then we can write

$$\sum_t \left|\left|\tilde{\Sigma}_t - \Sigma_t\right|\right|_F^2 = \sum_t (y_t'y_t)^2 + (\bar{\varphi}_t'\bar{\varphi}_t)^2 + \zeta^2 N_t - 2(y_t'\bar{\varphi}_t)^2 - 2\zeta y_t'y_t + 2\zeta(\bar{\varphi}_t'\bar{\varphi}_t)$$

After the entire tree is trained, ζ is estimated, as described below. Thus ζ is a nuisance parameter. Notice above that the first order condition of ζ is

$$\sum_t 2\zeta N_t - 2y_t'y_t + 2\bar{\varphi}_t'\bar{\varphi}_t = 0$$

Thus solving for ζ yields

$$\zeta = \frac{1}{\sum_t N_t} \sum_t y_t'y_t - \bar{\varphi}_t'\bar{\varphi}_t \tag{C.19}$$

Plugging this in eliminates the nuisance parameters, and thus φ_L and φ_R are found by minimizing

$$\begin{aligned}\sum_t \Bigg(&(y_t'y_t)^2 + (\bar{\varphi}_t'\bar{\varphi}_t)^2 + \left(\frac{1}{\sum_\tau N_\tau}\sum_\tau x_\tau' x_\tau - \bar{\varphi}_\tau'\bar{\varphi}_\tau\right)^2 N_t \\ &-2(y_t'\bar{\varphi}_t)^2 + 2(\bar{\varphi}_t'\bar{\varphi}_t - y_t'y_t)\left(\frac{1}{\sum_\tau N_\tau}\sum_\tau x_\tau' x_\tau - \bar{\varphi}_\tau'\bar{\varphi}_\tau\right)\Bigg)\end{aligned} \tag{C.20}$$

This is a relatively simple optimization problem because the function is smooth, the derivatives are relatively simple to calculate, and it's only a two dimensional problem. The partial derivatives of f with respect to φ_L and φ_R, which can be used to solve this problem numerically, are found below.

After φ_L and φ_R are chosen by the above minimization process, the values are plugged back into the objective in equation (C.20) for each candidate split. The candidate split which reaches the lowest value of quantity (C.20) is selected as the new split. For example, assume that $c_j^l(k)$, with its associated φ_L and φ_R values achieves the lowest value when plugged into quantity (C.20) above. As described above, then leaf l will become a decision node with left leaf L and right leaf R. We'll set $\hat{\varphi}_L = \varphi_L$. Similarly, we'll set $\hat{\varphi}_R = \varphi_R$. Just as described above, for any given $Z_{i,t}$, we can define C_L and C_R such that

$$Z_{i,t} \in C_L \text{ if and only if } Z_{i,t} \in C_l \text{ and } Z_{i,k,t} \leq c_j^l(k)$$

$$Z_{i,t} \in C_R \text{ if and only if } Z_{i,t} \in C_l \text{ and } Z_{i,k,t} > c_j^l(k)$$

When the tree is completely trained, ζ is set for the tree, using equation (C.19), such that

$$\zeta = \frac{1}{\sum_t N_t} \sum_t y_t' y_t - \varphi_t' \varphi_t \tag{C.21}$$

In summary, for a range of split values on different columns of Z, φ_L and φ_R are chosen by minimizing the objective function (C.20). The candidate with the lowest fitted values of (C.20) is selected as the new split, and the leaf l is transformed into a decision node with left leaf L—corresponding to the fitted value φ_L—and right leaf R—corresponding to the fitted value φ_R. Finally, when the tree is trained by repeating this process recursively to the desired depth, ζ is set according to equation (C.21).

Random Forest Objective Function Derivatives

Let the function f below denote the objective function described above in equation (C.20):

$$f(\varphi_L, \varphi_R) = \sum_t \Bigg((y_t' y_t)^2 + (\bar{\varphi}_t' \bar{\varphi}_t)^2 + \left(\frac{1}{\sum_\tau N_\tau} \sum_\tau x_\tau' x_\tau - \bar{\varphi}_\tau' \bar{\varphi}_\tau \right)^2 N_t$$

$$-2(y_t' \bar{\varphi}_t)^2 + 2(\bar{\varphi}_t' \bar{\varphi}_t - y_t' y_t) \left(\frac{1}{\sum_\tau N_\tau} \sum_\tau x_\tau' x_\tau - \bar{\varphi}_\tau' \bar{\varphi}_\tau \right) \Bigg)$$

Also

$$\bar{\varphi}_t = \varphi_{-l,t} + \varphi_L \iota_{L,t} + \varphi_R \iota_{R,t}$$

Define $N_{L,t} = \iota'_{L,t}\iota_{L,t}$, which is the number of ones in $\iota_{L,t}$. In other words, this is the sample size of the training data in the candidate left leaf. Similarly, let $N_{R,t} = \iota'_{R,t}\iota_{R,t}$. We can write

$$\bar{\varphi}'_t\bar{\varphi}_t = \varphi'_{-l,t}\varphi_{-l,t} + \varphi_L N_{L,t} + \varphi_R N_{R,t}$$

Using this, then we can write

$$\frac{\partial f}{\partial \varphi_L} = \sum_t \left(2N_{L,t}(\bar{\varphi}'_t\bar{\varphi}_t) - 2N_t \left(\frac{1}{\sum_\tau N_\tau} \sum_\tau N_{L,\tau} \right) \left(\frac{1}{\sum_\tau N_\tau} \sum_\tau x'_\tau x_\tau - \bar{\varphi}'_\tau \bar{\varphi}_\tau \right) \right.$$

$$\left. -4(y'_t\bar{\varphi}_t)(y'_t\iota_{L,t}) + 2N_{L,t} \left(\frac{1}{\sum_\tau N_\tau} \sum_\tau x'_\tau x_\tau - \bar{\varphi}'_\tau \bar{\varphi}_\tau \right) - 2(\bar{\varphi}'_t\bar{\varphi}_t - y'_t y_t) \left(\frac{1}{\sum_\tau N_\tau} \sum_\tau N_{L,\tau} \right) \right)$$

And similarly

$$\frac{\partial f}{\partial \varphi_R} = \sum_t \left(2N_{R,t}(\bar{\varphi}'_t\bar{\varphi}_t) - 2N_t \left(\frac{1}{\sum_\tau N_\tau} \sum_\tau N_{R,\tau} \right) \left(\frac{1}{\sum_\tau N_\tau} \sum_\tau x'_\tau x_\tau - \bar{\varphi}'_\tau \bar{\varphi}_\tau \right) \right.$$

$$\left. -4(y'_t\bar{\varphi}_t)(y'_t\iota_{R,t}) + 2N_{R,t} \left(\frac{1}{\sum_\tau N_\tau} \sum_\tau x'_\tau x_\tau - \bar{\varphi}'_\tau \bar{\varphi}_\tau \right) - 2(\bar{\varphi}'_t\bar{\varphi}_t - y'_t y_t) \left(\frac{1}{\sum_\tau N_\tau} \sum_\tau N_{R,\tau} \right) \right)$$