NYU STERN
Salomon Center for the Study of Financial Institutions

Working Paper Series ASSET MANAGEMENT Research Group

INFORMATION ACQUISITION AND PORTFOLIO UNDER-DIVERSIFICATION

Stijn Van Nieuwerburgh
Laura Veldkamp

SC-AM-05-08

Information Acquisition and Portfolio Under-Diversification

Stijn Van Nieuwerburgh and Laura Veldkamp*

September 19, 2005

ABSTRACT

We develop a rational model of investors who choose which asset payoffs to acquire information about, before forming portfolios. Scale economies in information acquisition lead investors to specialize in learning about a set of highly-correlated assets. Knowing more about these assets makes them less risky and more desirable to hold. Benefits to specialization compete with benefits to diversification. The resulting asset portfolios appear under-diversified from the perspective of standard theory, but are optimal. In equilibrium, information is a strategic substitute because assets that many investors learn about have low expected returns. Increasing returns, combined with strategic substitutability leads ex-ante identical investors to specialize in different information, and hold different portfolios. Information choice rationalizes investing in a diversified fund and a set of highly-correlated assets, an allocation observed in the data but usually deemed anomalous.

*Stijn Van Nieuwerburgh is in the finance department at New York University, Stern School of Business, 44 West Fourth St., 9-190, New York, NY 10012 (email:svnieuwe@stern.nyu.edu). Laura Veldkamp is in the economics department at New York University, Stern School of Business, 44 West Fourth St., 7th floor, New York, NY 10012 (email:lveldkam@stern.nyu.edu). We thank Dave Backus, Bernard Dumas, Ned Elton, Diego Garcia, Fatih Guvenen, Bob Hall, Ron Kaniel, Hanno Lustig, Massimo Massa, Pascal Maenhout, Antii Petajisto, Urs Peyer, Matthew Pritsker, Tom Sargent, Chris Sims, Eric Van Wincoop, Pierre-Olivier Weill, seminar participants at NYU, Cornell, USC, UCLA, Stanford, UCSB, INSEAD, Ohio State, and the Federal Board of Governors, and conference participants at the 2005 Society for Economic Dynamics in Budapest, the 2005 CEPR Asset Pricing Institute in Gerzensee, the 2005 World Congress of the Econometric Society in London, and the 2005 European Finance Association in Moscow for helpful comments. **JEL classification**: G11, D83, G12, G14. **Keywords**: Portfolio theory, asymmetric information, information choice.

An important outstanding question in finance is why investor portfolios are so concentrated. This fact challenges standard portfolio theory, which stresses the benefits of diversification. However, when an investor can first choose what assets to learn about, scale economies in information generate gains to specialization. We propose a new framework for analyzing joint learning and portfolio choices where incentives to diversify and specialize compete. The model predicts the cross-section of information sets that investors will acquire in equilibrium and uses this result to bring portfolio theory and capital asset pricing model results closer to observed portfolios and prices.

Analyzing investment and information choices jointly is valuable because it is the feedback of one decision on the other that generates gains to specialization. When choosing information, investors can acquire noisy signals about future payoffs of many assets, or they can specialize and acquire precise information about a few assets. Choosing to learn more about an asset makes investors expect to hold more of it, because for an average signal realization, they prefer an asset they are better informed about. As expected asset holdings rise, returns to information increase; one signal applied to one share generates less profit than the same signal applied to many shares. Specialization arises because the more an investor holds of an asset, the more valuable it is to learn about that asset; but the more an investor learns about the asset, the more valuable that asset is to hold.

Investors have two means of reducing uncertainty and risk: They can learn about the payoffs of assets they will hold, or they can diversify. The optimal strategy does some of each. Investors hold some fraction of their assets in a diversified fund, about which they learn nothing, and hold the other fraction in a set of highly-correlated assets that they specialize in learning about. In equilibrium, ex-ante identical investors specialize in different risks. Assets that many investors learn about command a lower risk premium. This makes investors want to research assets that others are not researching, and hold portfolios unlike what the average investor holds.

The force behind specialization is a general one, increasing returns to information. Radner and Stiglitz (1984) showed that the marginal value of information increases as more of it is obtained. The idea that information about large investments is more valuable dates back to Wilson (1975). He found that information value is increasing in a firm's scale of operation. Applying these basic economic insights to an information and portfolio choice problem, produces a theory of what investors should specialize in and how specialization and diversification trade-off. Embedding this theory in a general equilibrium model tells us how investors' learning choices interact and how aggregate learning affects asset prices.

Starting with identical prior beliefs, investors can obtain additional signals about what the realizations of future asset payoffs will be. Information is not required to hold an asset (as in Merton (1987)); rather it is a tool to reduce the conditional variance of the asset's payoff. Because we focus on the choice of what to learn, rather than how much to learn, we endow investors with a fixed budget of signal precision to allocate across assets. This budget, which we call capacity, is quantified as an increase in the generalized precision of posterior beliefs about asset payoffs, relative to prior beliefs. After allocating capacity, the investor observes signals drawn from a distribution whose precision he has chosen. Conditional on these signals, he solves a standard CARA-normal portfolio problem. We examine the predictions of this model in both partial equilibrium and general equilibrium settings.

Section II analyzes a partial equilibrium model where the investor takes asset prices as given. When asset returns are uncorrelated, the investor chooses to learn about one asset. Because a piece of information is most profitable when it is applied to many shares, the investor allocates her capacity to the asset with the highest squared Sharpe ratio, the asset she is likely to hold most of. She invests in a diversified portfolio and adds to that a "learning portfolio" consisting of a single asset. When asset payoffs are correlated, the investor learns about a single risk factor instead of a single asset. Her "learning portfolio" contains assets in proportion to their covariance with the risk factor. In both cases, it is optimal for the investor with zero information capacity to hold a diversified portfolio; our theory collapses to the standard model. As the investor's information capacity increases, holding a perfectly diversified portfolio is still feasible, but no longer optimal.

Specialization arises because of the interaction of the information and portfolio choice problems, not because of the form of the information constraint. Even with a learning technology that exhibits decreasing marginal returns, specialization, though moderated, still persists. With a small amount of capacity, investors fully specialize in learning about one risk. Given sufficient information capacity, the investor will learn about more than one risk factor (section II.C). The increase and then decrease in the marginal value of information is similar to the more general learning results in Keppo, Moscarini and Smith (2005).

Section III investigates a general equilibrium model where a continuum of investors interact (as in Admati (1985)). Endogenous prices act as an additional source of information: they are a noisy signal of what other investors know. While investors still have an incentive to specialize in one risk factor, they also have an incentive to specialize in a different risk factor from the ones other investors are learning about: Learning is a strategic substitute. The reason is that assets that more investors learn about carry lower equilibrium returns, as their demand bids up the price. When ex-ante identical investors choose to learn about different risk factors, they end up holding different concentrated asset portfolios. Asset returns in our model are described by the CAPM that would hold if each investor had the average of all investors' signal precisions. By characterizing the aggregate allocation of capacity, we can determine what this heterogeneous-information CAPM predicts for the cross-section of asset prices. When an asset's value is correlated with large, high-return risk factors, its price should be higher than the standard CAPM predicts.

Recent empirical research confirms the predictions of our theory. Many individuals hold under-diversified portfolios of common stock, in addition to a well-diversified mutual fund. The median retail investor at a large on-line brokerage company holds only 2.6 stocks (Barber and Odean (2001)). These portfolios of directly-held equity not only contain too few stocks, but the stocks they contain are positively correlated (Goetzmann and Kumar (2003)). But directly-held equities are only 40% of the median household's portfolio; the remaining 60% is in stock and bond mutual funds (Polkovnichenko (2004)). Using Swedish data on investors' complete wealth portfolio, Massa and Simonov (2005) document similar facts. They rule out the explanation that this concentration optimally hedges labor income risk.

If investors concentrate their portfolios because they have informational advantages, then concentrated portfolios should outperform diversified ones. In contrast, if transaction costs or behavioral biases are responsible, then concentrated portfolios should offer no advantage. Ivkovic, Sialm and Weisbenner (2005) find that concentrated investors outperform diversified ones by as much as 3% per year. This excess return is even higher for investments in local stocks, where natural informational asymmetries are most likely to be present (see also Coval and Moskowitz (1999), Coval and Moskowitz (2001); Massa and Simonov (2005); Ivkovic and Weisbenner (2005)). Likewise, mutual funds with a higher concentration of assets by industry outperform diversified funds (Kacperczyk, Sialm and Zheng (2005)). If some investors have higher capacity than others, they should consistently earn higher returns. Indeed, the top 10% most successful investors do consistently earn higher excess returns (Coval, Hirshleifer and Shumway (2003)), as do institutional investors with degrees from more selective universities (Chevalier and Ellison (1999)). Finally, if asymmetric information exists in the market, then investors who learn from prices should outperform investors who buy and hold a market index. Using CRSP data (1927-2000), Biais, Bossaerts and Spatt (2004) show that price-contingent strategies generate annual returns (Sharpe ratios) that are 3% (16.5%) higher than the indexing strategy. These results highlight the economic importance of asymmetric information and help to rationalize the multi-billion dollar financial management industry.

Why is it relevant to consider information constraints when so much investment is professionally managed? While individuals can avoid processing information by paying a mutual fund manager, even fund managers must decide which stocks to follow, which reports to read and what research to do. The model could be reinterpreted as solving a fund manager's problem. This manager can earn excess returns from active portfolio management, even though markets are efficient. Section IV relates our results to traditional theories of active portfolio management (Treynor and Black (1973)). In the conclusion, we sketch a model with a market for portfolio management services. Our results could extend the one asset model of Garcia and Vanden (2005) to better understand the relationship between fund styles and expected returns.

Many theories in economics and finance have predictions that depend crucially on what information investors have. This information is usually treated as an endowment. By asking what information rational investors would want to acquire, predictions contingent on information sets can be turned into more general predictions. This paper provides a tractable framework and set of tools for analyzing optimal information choices and incorporating those choices into commonly-used models of portfolio composition and asset pricing.

I. Setup

This is a static model which we break up into 3-periods. In period 1, the investor chooses the variance of signals about asset payoffs. That choice is constrained by information capacity, which bounds the total precision of the signals, and by principal components analysis, which limits the linear combinations of signals the investor can choose and keeps the problem tractable. In period 2, the investor observes signals and then chooses what assets to purchase. In period 3, he receives the asset payoffs and realizes his utility. Signal choices and portfolio choices in this setting are circular: What an investor wants to learn depends on what he thinks he will invest in and what he wants to invest in depends on what he has learned. To ensure that beliefs and actions are consistent, we use backwards induction. We first solve the period 2 portfolio problem for arbitrary beliefs. Then, we substitute the optimal portfolio rule into the period 1 information optimization problem.

The vector of unknown asset payoffs $f \sim \mathcal{N}(\mu, \Sigma)$ is what the investor will devote capacity to learning about. After learning, the investor will have posterior beliefs about asset payoffs: $f \sim N(\hat{\mu}, \widehat{\Sigma})$. Let r be the risk-free return and q and p are Nx1 vectors of the number of shares the investor chooses to hold and the asset prices. Following Admati (1985), we call $f_i - rp_i$ the excess return on asset i. Investors have mean-variance utility with absolute risk aversion parameter ρ:

$$U = E\left[q'(f - pr) - \frac{\rho}{2} q' \widehat{\Sigma} q \;\middle|\; \mu \right]. \tag{1}$$

Mean-variance utility arises in settings where investors have negative exponential utility and face normally distributed payoffs.[1] It allows a tractable solution to an equilibrium model, while treating learned information and prior information as equivalent. This investor chooses information that maximizes his expected utility at the time when he must make his portfolio decision. When choosing what to learn, our investor asks himself, "When I invest, what information would I most like to know?"

Period-2 investment problem Let $\hat{\mu}$ and $\widehat{\Sigma}$ be the mean and variance of payoffs, conditional on all information known to the investor in period 2. The investor chooses q to maximize

$$E[q'(f - pr) - \frac{\rho}{2} q' \widehat{\Sigma} q | \hat{\mu}, \mu]. \tag{2}$$

Updating Beliefs At time 1, the investor chooses how to allocate his information capacity by choosing a normal distribution from which he will draw an $N \times 1$ signal η about asset payoffs.[2] At time 2, the investor will combine his signal $\eta \sim N(f, \Sigma_\eta)$ and his prior belief $\mu \sim N(f, \Sigma)$, using Bayes' law. His posterior belief about the asset payoff f has a mean

$$\hat{\mu} \equiv E[f|\mu, \eta] = \left(\Sigma^{-1} + \Sigma_\eta^{-1}\right)^{-1} \left(\Sigma^{-1}\mu + \Sigma_\eta^{-1}\eta\right) \tag{3}$$

and a variance that is a harmonic mean of the prior and signal variances:

$$\widehat{\Sigma} \equiv V[f|\mu, \eta] = \left(\Sigma^{-1} + \Sigma_\eta^{-1}\right)^{-1}. \tag{4}$$

These are the conditional mean and variance that investors use to form their portfolios in period 2.

Since every signal variance Σ_η has a unique posterior belief variance $\widehat{\Sigma}$ associated with it, we can economize on notation and optimize over posterior belief variance $\widehat{\Sigma}$ directly. Prior (unconditional) variances and covariances are not random; they are given. Posterior (conditional) variances are also not random; they are choice variables that summarize the investor's optimal information decision.

Capacity Constraint There are 2 constraints governing how the investor can choose his signals. The first constraint is the *capacity constraint.* The work on information acquisition with one risky asset quantified information as the ratio of variances of prior and posterior beliefs (Verrecchia

[1] It is equivalent to $U = E\left[-\log\left(E\left[\exp\left(-\rho W\right)|\hat{\mu}, \mu\right]\right)|\mu\right]$, where $W = rW_0 + q'(f - pr)$ by the budget constraint and payoffs f are normally distributed. This formulation of utility is related to Epstein and Zin (1989) preference for early resolution of uncertainty and Hansen and Sargent (2004) models of risk-sensitive control.

[2] Choosing normally distributed signals η is optimal. When an objective is quadratic, normal distributions maximize the entropy over all distributions with a given variance (see Cover and Thomas (1991), Chapter 10). As we show below, the objective in our problem is quadratic: $E\left[\frac{1}{2}(\hat{\mu} - pr)'\widehat{\Sigma}^{-1}(\hat{\mu} - pr)|\mu\right]$.

(1982)). The more information a signal contains, the more the posterior variance of the asset falls below the prior variance, and the more information capacity is required to observe the signal. We generalize this metric to an multi-signal setting by calling capacity the ratio of the *generalized* prior variance to the generalized posterior variance, where generalized variance refers to the determinant of the variance-covariance matrix:

$$\frac{|\Sigma|}{|\widehat{\Sigma}|} \leq e^{2K} \tag{5}$$

The amount of capacity K bounds the reduction in uncertainty of payoffs due to the knowledge of the signal η.[3]

This capacity constraint is one possible description of a learning technology. We think it is a relevant constraint because it is a commonly-used distance measure in econometrics (a log likelihood ratio) and in statistics (a Kullback-Liebler distance[4]); it is equivalent to a bound on entropy reduction, which has a long history in information theory as a quantity measure for information (Shannon (1948)); it is a measure of information complexity (Cover and Thomas (1991)); it can be interpreted as a reduction in measurement error. It has been previously used in economics and finance (Sims (2003) and Peng (2004)) to model limited mental processing ability in a representative investor framework. That having been said, this particular formulation of the learning technology is not crucial for the results. Our capacity constraint is a way to describe a feasible set of learning possibilities that is rich enough to analyze the trade-off between diversification and specialization in learning. Section II.C considers an alternative learning technology, one with decreasing returns. The incentive to specialize persists, but is moderated. Finally, while studying the extensive margin of information acquisition is interesting, adding a cost for capacity won't change the nature of the capacity allocation decision. For every cost, there is an amount of capacity K that produces an identical result. In sum, because the increasing returns to specialization show up in the objective, through the endogenous portfolio choice, specialization is robust to changes in the learning technology.

No Negative Learning Constraint The second constraint is that the variance-covariance matrix of the signals must be positive semi-definite.

$$\Sigma_\eta \quad \text{positive semi-definite} \tag{6}$$

[3]To see the role of the signal, the capacity constraint can be restated as a bound on the precision Σ_η^{-1} of signals η: $|\Sigma_\eta^{-1}\Sigma + I| \leq e^{2K}$.

[4]In statistics, this distance is used as a measure of how difficult it is to distinguish one distribution from another.

Without this constraint, the investor could increase uncertainty about one variable in order to obtain a more precise signal about another, without violating the capacity constraint. Ruling out increasing uncertainty implies that investors cannot forget nor see signals with negative information content.

Learning About Correlated Risks When asset payoffs co-vary, learning about one asset's payoff is informative about others. To keep track of what is being learned about, we describe signals by how much information they contain about each principal component of asset payoffs. Studying principal component risks is a well-established idea in the portfolio literature (Ross (1976)). Principal components, or *risk factors*, are linear combinations of underlying assets that do not covary with each other. An eigen-decomposition splits the prior variance-covariance matrix Σ into a diagonal eigenvalue matrix Λ, and an eigenvector matrix Γ: $\Sigma = \Gamma\Lambda\Gamma'$. The Λ_i's are the variances of each risk factor i. The i^{th} column of Γ (denoted by Γ_i) gives the loadings of each asset on the i^{th} risk factor. Investors obtain signals about the payoffs of risk factors $f'\Gamma_i$. This does not rule out learning about many risks. It does rule out signals with correlated information about risks that are independent. For example, if oil price risk and rainfall were two independent risks, an investor could not observe their sum without acquiring a signal about each. Effectively, we assume that the investor cannot choose to observe signals about these zero-probability events.

Appendix B shows that the key results, specialization and strategic substitutability, hold for any given set of orthogonal risk factors. However, this particular decomposition simplifies the problem. Investors will have posterior beliefs with the same eigenvectors as their prior beliefs ($\widehat{\Sigma} = \Gamma\hat{\Lambda}\Gamma'$), but with lower weights $\hat{\Lambda}_i$ on some risks they chose to learn about. The decrease in risk factor variance $\Lambda_i - \hat{\Lambda}_i$ captures how much an investor learned about that risk.

The Investor's Problem

1. The investor chooses $\{\hat{\Lambda}_i\}_{i=1}^{N}$, the variance of posterior beliefs about each risk factor, to maximize (1), subject to the capacity constraint (5), the no-negative learning constraint (6), and rational expectations about q.

2. Given a signal η about asset payoffs f, posterior means $\hat{\mu}$ and variances $\widehat{\Sigma}$ are formed according to Bayes' law (3) and (4).

3. Investors choose optimal portfolios q to maximize (2).

The sequence of events is summarized in figure I.

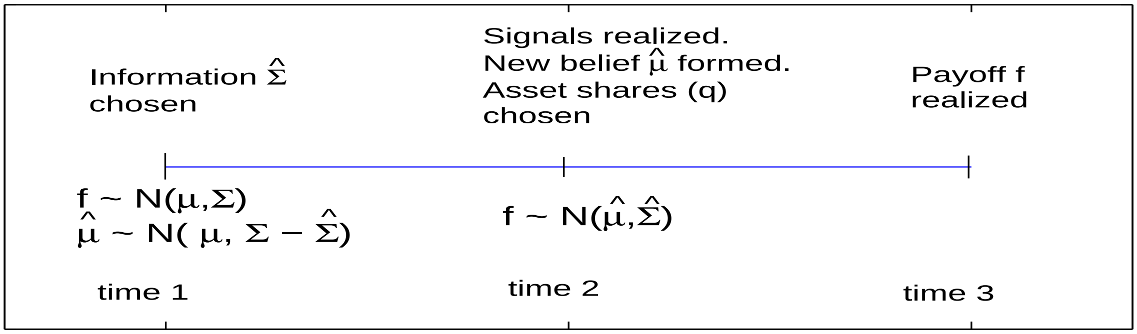

Figure 1. Sequence of events in partial equilibrium model

The optimal portfolio $q^\star$ is given by the first order condition of (2):

$$q^\star = \frac{1}{\rho}\widehat{\Sigma}^{-1}(\hat{\mu} - pr). \tag{7}$$

The model does not rule out short sales: $q^\star < 0$ when $\hat{\mu} - pr < 0$. Any remaining initial wealth is invested in the risk-free asset.

Substituting the optimal portfolio choice into (1) delivers the utility that results from having any beliefs $\hat{\mu}$, $\widehat{\Sigma}$ and investing optimally. The period-1 problem then amounts to choosing information (beliefs) optimally to maximize $U = E\left[\frac{1}{2}(\hat{\mu} - pr)'\widehat{\Sigma}^{-1}(\hat{\mu} - pr)|\mu\right]$, where posterior means are random: $(\hat{\mu} - pr) \sim N(E[\hat{\mu} - pr], V_{ER})$, and $V_{ER} \equiv Var[\hat{\mu} - pr|\mu]$. Taking the expectation of a non-central χ^2-distributed random variable, the time-1 problem is to maximize

$$max_{\widehat{\Sigma}} \;\; \frac{1}{2}Tr(\widehat{\Sigma}^{-1}V_{ER}) + \frac{1}{2}E[\hat{\mu} - pr]'\widehat{\Sigma}^{-1}E[\hat{\mu} - pr]. \tag{8}$$

II. Partial Equilibrium Results

A. Independent Assets

To gain intuition, it is helpful to first consider a simple case with N assets whose payoff variance-covariance matrix Σ is diagonal. Choosing signals with the same principal components as asset payoffs implies that signals are independent as well. The next section will generalize the problem to correlated assets.

When investors takes prices as given, $E[\hat{\mu} - pr] = \mu - pr$ and $V_{ER} = var[\hat{\mu}|\mu] = \Sigma - \widehat{\Sigma}$. Because the trace of a matrix is the sum of its eigenvalues, we can rewrite (8) as

$$max_{\{\widehat{\Sigma}_{11},\cdots,\widehat{\Sigma}_{NN}\}} \frac{1}{2}\{-N + \sum_{i=1}^{N} \frac{\Sigma_{ii}}{\widehat{\Sigma}_{ii}} + \sum_{i=1}^{N} \theta_i^2 \frac{\Sigma_{ii}}{\widehat{\Sigma}_{ii}}\}. \tag{9}$$

The constraints become

$$s.t. \quad \prod_{i=1}^{N} \frac{\Sigma_{ii}}{\widehat{\Sigma}_{ii}} = \exp(2K) \qquad and \qquad \frac{\Sigma_{ii}}{\widehat{\Sigma}_{ii}} \geq 1, \ \forall i$$

where θ_i^2 is the prior squared Sharpe ratio of asset i: $\theta_i^2 \equiv \frac{(\mu_i - p_i r)^2}{\Sigma_{ii}}$. The first constraint results from (5) and the fact that the determinant of a diagonal matrix is the product of the diagonal entries. The second constraint uses (6), (4) and the fact that a diagonal matrix is positive semi-definite if and only if all its elements are non-negative.

The key feature of the learning problem (9) is that it is convex in the variance $\widehat{\Sigma}_{ii}$. It is the convexity of the objective that delivers a corner solution. The corner solution is to reduce variance (increase precision) on the risk factor with the highest weight (θ_i^2) as much as possible.

Proposition 1. *The optimal information portfolio with N independent assets uses all capacity to learn about one asset, the asset with the highest squared Sharpe ratio* $\theta_i^2 = (\mu_i - p_i r)^2 \Sigma_{ii}^{-1}$.

Proof is in appendix A. For intuition, consider the problem of sequentially assigning units of capacity that can reduce the variance of an asset's payoff from Σ_{ii} to $\widehat{\Sigma}_{ii} = (1-\epsilon)\Sigma_{ii}$. The greatest utility gain is obtained by assigning the first unit of capacity to the asset with the highest value of $(\mu_i - p_i r)^2 \Sigma_{ii}^{-1}$. The value of assigning the next unit of capacity to asset i is then even greater: $(\mu_i - p_i r)^2 \widehat{\Sigma}_{ii}^{-1} > (\mu_i - p_i r)^2 \Sigma_{ii}^{-1}$. The value of assigning each subsequent unit of capacity to i rises higher and higher, while the value of assigning capacity to all other assets remains the same. Therefore, the optimal choice of posterior variance is $\widehat{\Sigma}_{ii} = e^{-2K}\Sigma_{ii}$, and $\widehat{\Sigma}_{jj} = \Sigma_{jj}$ for all $j \neq i$.

The value of learning about an asset is indexed by its squared Sharpe ratio $(\mu_i - p_i r)^2 \Sigma_{ii}^{-1}$. Another way to express the same quantity is as the product of two components: $(\mu_i - p_i r)$ and $(\mu_i - p_i r)/\Sigma_{ii}$, which is $\rho E[q_i]$ for an investor who has zero capacity. An investor wants to learn about an asset that has (i) high expected excess returns $(\mu_i - p_i r)$, and (ii) features prominently in his (expected) portfolio $E[q_i]$. The fact that an investor wants to invest all capacity in one asset comes from the anticipation of his future portfolio position $E[q]$. The more shares of an asset he expects to hold, the more valuable information about those shares is, and the higher the index value he assigns to learning about the asset. But, as he learns more about the asset, the amount he expects to hold $E[q_i] = (\mu_i - p_i r)/(\rho\widehat{\Sigma}_{ii})$ rises. As he learns, devoting capacity to the same asset becomes more and more valuable. This is the increasing return to learning.

How does this learning strategy affect the investor's portfolio? For the assets that the investor does not learn about, the number of shares does not change. For the asset he does learn about, the expected number of shares increases by $E[q^{learn}] = \frac{1}{\rho \Sigma_{ii}}(\mu_i - p_i r)(e^{2K} - 1)$. Call the portfolio of shares that the investor would hold if he had zero-capacity and could not learn, q^{div}. This is the benchmark portfolio predicted by the standard CARA-normal model. Since it contains no signals, it is not random: $E[q^{div}] = q^{div}$. The portfolio of an investor with positive capacity is the sum of q^{div} and the component due to learning, q^{learn}, (plus his position in the risk free asset).

Proposition 2. *As long as there is at least one asset i for which $(\mu_i - p_i r) \neq 0$, then when capacity rises, the expected fraction of the optimal portfolio consisting of fully-diversified assets $(|q^{div}|/(|q^{div}| + |E[q^{learn}]|)$ falls.*

Proof: As capacity (K) increases from zero, the zero-capacity portfolio q^{div} is, by definition, unchanged. As long as there is an asset s.t. $(\mu_i - p_i r) \neq 0$, then proposition 1 tells us that an investor will learn about an asset $i^\star$ s.t. $(\mu_{i^\star} - p_{i^\star} r) \neq 0$. The only quantity that changes in K is the expected amount of asset $i^\star$ held due to learning: $|E[q_{i^\star}^{learn}]| = \frac{1}{\rho \Sigma_{i^\star i^\star}} |\mu_{i^\star} - p_{i^\star} r|(e^{2K} - 1)$. Since $\mu_{i^\star} - p_{i^\star} r \neq 0$, $|E[q_{i^\star}^{learn}]|$ is strictly increasing in K. $\square$

Only expected portfolio holdings can be predicted. Since actual signal realizations and therefore posterior beliefs $\hat{\mu}$ are random variables, the true portfolio chosen in period 2 could be either larger or smaller in absolute value, than it would have been without the signal. But, for any given belief about payoffs $\hat{\mu}_i$, having more capacity to reduce the variance of that belief $\widehat{\Sigma}_{ii}$, makes the investor take a larger position in the asset $|q_i|$.

This result can be easily restated in terms of the more familiar value-weighted fraction of shares in the learning and diversified funds. As long as the expected excess return and price for the learning asset i are positive, then the expected value-weighted fraction of shares held in the diversified portfolio falls. This is the sense in which learning and diversification trade off.

Corollary 1. *An investor who optimally chooses a less diversified portfolio earns a higher expected return than an investor who chooses a more diversified portfolio.*

Proof in appendix C. Proposition 2 tells us that investors who have high information capacities K choose highly under-diversified portfolios. Such investors makes more informed investment choices and obtain a higher expected profit. The reason is that these investors achieve a higher correlation between asset payoffs and portfolio shares: they hold a long position when the asset is likely to have a high payoff and a short position when the asset payoff is likely to be low. This prediction is corroborated empirically by the findings of Ivkovic et al. (2005) and Kacperczyk et al. (2005). They show that investors with concentrated portfolios buy stocks have consistently higher returns than the stocks they sell, so that under-diversified portfolios significantly outperform diversified ones. This finding is present across wealth groups.

Data Example with Independent Assets

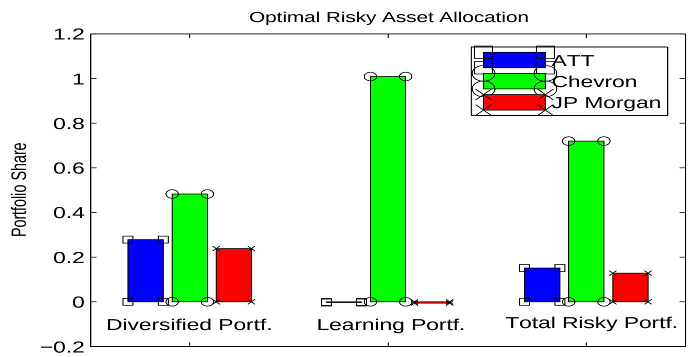

Figure 2. Under-Diversification and the Increasing Returns to Learning: Uncorrelated Assets.

We illustrate the portfolio composition with a simple numerical example. Figure 2 illustrates the case of three uncorrelated S&P 500 assets.[5] The monthly excess returns on AT&T, Chevron, and JP Morgan were nearly orthogonal in the sample period. Chevron had the highest Sharpe ratio (.58 annualized). When faced with the mean excess returns and the covariance matrix of returns of three assets, an investor with zero information capacity would hold an optimally diversified portfolio, consisting of 28% AT&T, 48% Chevron, and 24% JP Morgan ('diversified portfolio'). When given some information capacity, the investor specializes in learning about Chevron. ($K = .5$ here, which allows the investor to reduce the standard deviation of one asset by 39%.) The 'learning fund' is fully invested in Chevron. As a result, the total portfolio is under-diversified: 15% AT&T, 72% Chevron, and 13% JP Morgan.

B. Correlated assets

When assets are correlated, signals about individual asset payoffs are no longer principal components. Instead, principal components are linear combinations of asset payoffs with weights on each asset given by an eigenvector of Σ. Rather than reduce the risk of independent assets, investors now

[5]Monthly return data runs from November 1986 and December 2003 (206 observations). Excess returns are constructed by subtracting the return on a 1-month T-bill.

choose how to reduce the variance of these independent risk factors. The factors could represent risks such as business cycle risk, pharmaceutical industry risk, or idiosyncratic risk. The variance of each risk factor is given by its eigenvalue (Λ_i). After transforming assets into independent risk factors, the results for independent assets can be restated for the correlated assets case.

When an investor learns about principal components, his posterior belief variance $\widehat{\Sigma}$ has the same eigenvectors (Γ) as Σ. Therefore, the investor's choice is over the diagonal eigenvalue matrix $\hat{\Lambda}$, where $\widehat{\Sigma} = \widehat{\Gamma}\hat{\Lambda}\widehat{\Gamma}'$. The investor solves problem (9), where the prior squared Sharpe ratios θ_i^2 now refer to risk factors: $\theta_i^2 \equiv \frac{((\mu - pr)'\Gamma_i)^2}{\Lambda_i}$. The capacity constraint still takes the form $\prod_{i=1}^{N} \frac{\Lambda_i}{\hat{\Lambda}_i} = \exp(2K)$ because the determinant of $\Sigma\widehat{\Sigma}^{-1}$ is the product of its eigenvalues $\Lambda_i \hat{\Lambda}_i^{-1}$ and because Σ and $\widehat{\Sigma}$ share eigenvectors Γ. The no-negative learning constraint, which requires $\Sigma - \widehat{\Sigma}$ to be positive semi-definite, or equivalently that all its eigenvalues are non-negative becomes $\frac{\Lambda_i}{\hat{\Lambda}_i} \geq 1$.

Proposition 3. *The optimal information portfolio with N correlated assets uses all capacity to learn about one linear combination of asset payoffs. The linear combination coefficients are given by the eigenvector* Γ_i*, with the highest factor squared Sharpe ratio* $\theta_i^2 = ((\mu - pr)'\Gamma_i)^2 \Lambda_i^{-1}$.

The proof follows immediately from proposition 1 and the new definition of θ_i^2. There are two components of this result. The first component tells us how the investor initially ranks learning about each risk factor Γ_i. The second tells us that he specializes completely in whatever risk factor he ranks first. What direction an investor decides to learn in is determined by the magnitude of the expected return on the risk factor $\Gamma_i'(\mu - pr)$ and by ρ times the expected holding of that risk factor: $\rho\Gamma_i' E[q]$. The fact that the investor wants to devote all capacity to learning about one risk factor comes from increasing returns, and does not depend on the particular risk factor structure (see appendix B). As the investor learns more about risk factor i, the investor expects to hold more of that risk factor: $\Gamma_i' E[q]$ grows. As he expects to hold more of the risk factor, the value of learning more about it rises.

What does this result mean for portfolio allocation? The investor will hold shares of each asset given by $\frac{1}{\rho}(\Gamma\hat{\Lambda}\Gamma')^{-1}(\hat{\mu} - pr)$. Again, this portfolio can be decomposed into the diversified benchmark portfolio that an investor with no capacity would hold $q^{div} = \frac{1}{\rho}(\Gamma\Lambda\Gamma')^{-1}(\hat{\mu} - pr)$, and the number of extra shares of assets that will be held due to learning,

$$q^{learn} = \frac{e^{2K} - 1}{\rho\Lambda_i}\Gamma_i\Gamma_i'(\hat{\mu} - pr)$$

where i is the factor the investor optimally learns about. This learning portfolio puts more weight on assets in proportion to how correlated they are with the risk factor that the investor is learning about. Since the 'learning' assets are highly correlated with a common risk factor, they are also highly correlated with each other. As K grows, the expected weight on this highly-correlated component of the portfolio rises. As learning increases, diversification falls.

Data Example with Correlated Assets

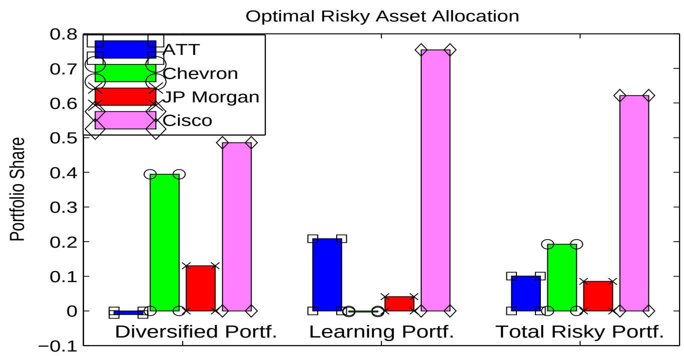

Figure 3. Under-Diversification and the Increasing Returns to Learning: Correlated Assets.

Figure 3 illustrates the case of correlated assets. It adds to the three uncorrelated assets described above a fourth asset, Cisco. Cisco has a low correlation with Chevron (-.008) and with JP Morgan (.068), but a high correlation with AT&T (.296). Cisco has a much higher Sharpe ratio than the other three firms. When offered these four assets, an investor with zero information capacity would hold an optimally diversified portfolio, consisting of -1% AT&T, 39% Chevron, 13% JP Morgan, and 49% Cisco ('diversification fund'). When given some information capacity (K is still .5 here), the investor learns about Cisco, the most valuable asset to learn about, but also about AT&T. The reason is that both Cisco and AT&T load positively on the most valuable risk factor (correlations .96 and .27 respectively). The 'learning fund' is invested for 75% in Cisco and 21% in AT&T. As a result of the specialization in learning, the total portfolio is under-diversified: 10% AT&T, 19% Chevron, 9% JP Morgan, and 62% Cisco. The new optimal portfolio has a variance (conditioning on past public information) that is 25% higher than the diversified portfolio variance; it is under-diversified.

C. Un-Learnable Risk and Decreasing Returns to Learning

In the previous results, investors never diversify their information because learning substitutes for diversification. As learning increases and risk falls, the value of diversification falls as well. With un-learnable risk, there is some risk that learning cannot eliminate, but diversification can. This risk revives some benefits to diversification and makes high-capacity investors learn about multiple risk factors. Thus, high capacity investors may hold more diversified portfolios than low capacity investors. This could explain why portfolios of individual investors tend to be less diversified than portfolios of institutional investors.

Adding un-learnable risk is also a way of generating decreasing returns to learning. Because this has intuitive appeal as an alternative learning technology, we explore its implications here. We find that a learning technology that embodies decreasing returns to specialization does not restore full diversification; it just moderates the extent of specialization. This reinforces the point that specialization is driven by increasing returns to information, not the form of the capacity constraint.

When all risk is learnable and capacity approaches infinity, the payoff variance of some portfolio approaches zero, an arbitrage arises, and profit becomes infinite. Un-learnable risk imposes a finite, maximum benefit to learning. To reduce an asset's learnable payoff variance to near zero costs an unbounded amount of information capacity and yields only a finite benefit. Therefore, learning an arbitrarily large amount about a single asset is never optimal.

In the modified model, the investor's preferences, the sequence of events, and the optimal period-2 portfolio remains unchanged. The period-1 choice of signal distributions is constrained by the fact that of the total variance in the prior beliefs Σ, $\alpha\Sigma$ is un-learnable, and only $(1-\alpha)\Sigma$ can be learned $(0 < \alpha < 1)$.[6] The new period-1 problem is to maximize (8) subject to a constraints on the reduction in entropy of the *learnable* component of asset payoffs. This constraint is formulated so that eliminating all learnable risk (reducing $\widehat{\Sigma}$ to $\alpha\Sigma$) requires infinite capacity. When $\widehat{\Sigma} = \Sigma$, the investor is not learning anything, and no capacity is required.

$$\frac{|\Sigma - \alpha\Sigma|}{|\widehat{\Sigma} - \alpha\Sigma|} \leq e^{2K} \tag{10}$$

The no-negative learning constraint is as before.

As before, the investor chooses the posterior variance $\hat{\Lambda}$ of given risk factors Γ to maximize (1) subject to (10) and no-negative learning (6). The first-order condition with respect to $\hat{\Lambda}_i$ describes an interior solution to the maximization problem:

$$(1+\theta_i^2)\frac{\Lambda_i}{-\hat{\Lambda}_i^2} = \xi\frac{1}{\hat{\Lambda}_i - \alpha\Lambda_i} - \phi_i, \tag{11}$$

where ξ is the Lagrange multiplier on (10), and ϕ_i is the Lagrange multiplier on the no-learning constraint. The left hand side is the marginal benefit of learning about risk factor i, the right hand side is the marginal cost. Once the variance of payoff i has been reduced sufficiently $(\hat{\Lambda}_i < 2\alpha\Lambda_i)$,

[6] For every result, except proposition 6, α can be a matrix where every element is $0 < \alpha_{ij} < 1$.

the second derivative reveals that the $\hat{\Lambda}_i$ that solves (11) is an interior maximum.

Proposition 4. *When there is un-learnable risk, the number of risk factors that the investor learns about is an increasing step function of K.*

Corollary 2. *When there is un-learnable risk and asset payoffs are independent, the number of assets held in the 'learning fund,' q^{learn}, is an increasing step function of capacity K.*

Proofs are in appendix D.

The reason for learning about additional assets can be seen by examining the marginal benefit and the marginal cost of learning. As the investor learns more and $\hat{\Lambda}_i$ decreases, the marginal benefit increases. Increasing returns to scale in learning are still present. However, the marginal cost is now also convex in $\hat{\Lambda}_i$. The difference of the two, the net marginal benefit, first increases until $\hat{\Lambda}_i = 2\alpha\Lambda_i$, and then decreases. In the limit, as the investor gets closer to learning all the learnable risk ($\hat{\Lambda}_i \to \alpha\Lambda_i$), the marginal cost approaches infinity. Therefore, there is some finite cutoff level of $\hat{\Lambda}_i$ such that when the investor reaches this level of learning for asset i, he begins to allocate some capacity to another risk factor. When assets are independent, allocating capacity to another risk factor means learning about another asset, and including the additional asset in the learning fund.

Proposition 5. *When there is un-learnable risk and there is some asset i with non-zero expected excess return $(\mu_i - p_i r) \neq 0$, then, as capacity rises, the fraction of the expected optimal portfolio consisting of fully-diversified assets $(|q^{div}|/(|q^{div}| + |E[q^{learn}]|)$ falls.*

Proof is in appendix E.

When the investor learns more about an asset, he expects to hold a larger position in that asset. Since the zero-capacity portfolio q^{div} does not change as capacity increases and more shares are held in the learning portfolio, the fraction of the expected portfolio that is diversified falls.

Proposition 6. *When there is un-learnable risk and capacity is infinite, the expected learning portfolio is fully diversified:* $\lim_{k\to\infty} E[q^{learn}] = \left(\frac{1}{\alpha} - 1\right) q^{div}$.

Proof: An investor with an infinite capacity would eliminate all learnable risk, setting $\hat{\Lambda} = \alpha\Lambda$, which implies $\widehat{\Sigma} = \alpha\Sigma$. In this limit, the learning fund is $E[q^{learn}] = \frac{1}{\rho}\left(\frac{1}{\alpha} - 1\right)\Sigma^{-1}(\mu - pr)$, a scaled-up copy of the diversified mutual fund. □

Putting the results together tells us that as capacity increases, diversification falls, and then rises again. An investor with zero capacity holds only the diversified fund. An investor with infinite capacity holds a perfectly diversified learning fund. In between the two perfectly diversified extremes, the investor with positive, finite capacity to learn is optimally under-diversified. So, an (individual) investor with low capacity would learn about fewer risk factors than an (institutional) investor with high capacity, and consequently hold a less diversified portfolio.

III. Equilibrium Information and Investment Choices

In general equilibrium, an investor must consider the information acquisition and investment strategies of other investors. Information is a strategic substitute in this setting: Investors want to learn about assets that others are not learning about. In equilibrium, this means that ex-ante identical investors will choose to observe different signals and will hold different assets. When all risk is learnable, the nature of the solution to the individuals problem does not change. After accounting for the actions that other investors will take and how these will affect asset prices, an investor chooses one risk factor and concentrates all his capacity on learning about that one factor. We begin by describing the modifications to the setup.

A. Setup

There is now a continuum of atomless investors, indexed by $j \in [0, 1]$. Preferences and payoffs are identical to the model described in section I. We continue to hold the risk-free rate fixed. Two additional assumptions are required. First, the per capita supply of the risky asset is $\bar{x} + x$, a constant plus a random $(n \times 1)$ vector with known mean and variance, and zero covariance across assets: $x \sim N(0, \sigma_x^2 I)$. This risky asset supply creates noise in the price level that prevents investors perfectly inferring the private information of others. Without it, there would be no private information, and no incentive to learn. We interpret this noise as liquidity shocks, life-cycle needs of traders, or errors that investors make when inverting prices.[7]

Second, we assume that noisy asset payoff signals are *independent.* Each investor adds his own noise when he interprets his information. An alternative mechanism would be a news agency that produces a noisy signal of the truth and transmits that same signal to all of investors. We revisit the idea of centralized markets for information processing in the conclusion.

Asset prices p are determined by market clearing: the sum of investors' demands for each asset equals its supply.

$$\int_0^1 \widehat{\Sigma}_j^{-1}(\hat{\mu}_j - pr)dj = \bar{x} + x \tag{12}$$

[7]See Biais et al. (2004) for an interpretation in terms of risky non-tradeable endowments.

B. Individual's Asset Allocation in Equilibrium

As before, we work backwards, starting with the optimal portfolio decision. In period 2, investors have three pieces of information that they aggregate to form their expectation of the assets' payoffs: their prior beliefs (common across investors), their signals (draws from distributions chosen in period 1), and the equilibrium asset price.

Proposition 7. *Asset prices are a linear function of the asset payoff and the unexpected component of asset supply.*

$$p = \frac{1}{r}(A + Bf + Cx)$$

This price is also a function of the posterior mean and variance of the 'average' investor:

$$p = \frac{1}{r}\left(\hat{\mu}_a - \rho\widehat{\Sigma}_a(\bar{x} + x)\right)$$

where the average posterior mean is $\hat{\mu}_a = \int_0^1 \hat{\mu}_j dj$ *and the 'average' posterior variance is a harmonic mean of all investors' variances* $\widehat{\Sigma}_a = \left(\int_0^1 \widehat{\Sigma}_j^{-1} dj\right)^{-1}$.

Proof is in appendix F, along with the formulas for A, B and C.

If prices take this form, then the mean and variance of the asset payoff, conditional on prices are $E[f|p] = B^{-1}(rp - A)$ and $V[f|p] = \sigma_x^2 B^{-1}CC'B^{-1'} \equiv \Sigma_p$. Denote its eigenvalue matrix by Λ_p. Then, posterior beliefs about the asset payoff f, conditional on prior beliefs $\mu \sim N(f, \Sigma)$, signals $\eta \sim N(f, \Sigma_\eta)$, and prices, can be expressed using standard Bayesian updating formulas. The posterior mean belief about f is

$$\hat{\mu} \equiv E[f|\mu, \eta, p] = \left(\Sigma^{-1} + \Sigma_\eta^{-1} + \Sigma_p^{-1}\right)^{-1}\left(\Sigma^{-1}\mu + \Sigma_\eta^{-1}\eta + \Sigma_p^{-1}B^{-1}(rp - A)\right) \quad (13)$$

with variance that is a harmonic mean of the three signal variances.

$$\hat{\Sigma} \equiv V[f|\mu, \eta, p] = \left(\Sigma^{-1} + \Sigma_\eta^{-1} + \Sigma_p^{-1}\right)^{-1}. \quad (14)$$

These are the conditional mean and variance that investors use to form their portfolios in period 2. As in the partial equilibrium problem, optimal portfolios are formed according to equation (7). Only the conditioning information changes.

C. Individual's Information Capacity Allocation in Equilibrium

In period 1, the investor chooses a covariance matrix for his posterior beliefs $\widehat{\Sigma}$, just as in the partial equilibrium problem. The difference is that at time 1, the time-2 expected excess return $(\hat{\mu} - pr)$ is now distributed $(\hat{\mu} - pr) \sim N\left((I-B)\mu - A, V_{ER}\right)$ where $V_{ER} \equiv \Sigma - \widehat{\Sigma} + B\Sigma B' + CC'\sigma_x^2 - 2\Sigma B'$. Using equation (8), the period-1 objective is:

$$max_{\widehat{\Sigma}} \left\{ \frac{1}{2} Tr(\widehat{\Sigma}^{-1} V_{ER}) + \frac{1}{2}((I-B)\mu - A)' \widehat{\Sigma}^{-1}((I-B)\mu - A) \right\}. \tag{15}$$

Just as in partial equilibrium, the choice of posterior covariance matrix $\widehat{\Sigma}$ is subject to two constraints. The constraints are formally the same as in section I, but require re-interpretation. The first constraint is that the total information the investor sees cannot reduce the generalized variance of payoffs by more than e^{2K}. Being a constraint on the distance between the posterior belief variance $\widehat{\Sigma}$ and the prior belief variance Σ, it assumes that investors use capacity to extract payoff relevant information both from private signals η and from prices. Some capacity must be devoted to price discovery; the remaining capacity can be optimally allocated to signals.[8] The second constraint is the equivalent of (6). This no-negative learning constraint prevents investors from forgetting information that is either contained in priors or in prices.

$$(\Sigma^{-1} + \Sigma_p^{-1})^{-1} - \widehat{\Sigma} \quad \text{positive semi-definite} \tag{16}$$

where $(\Sigma^{-1} + \Sigma_p^{-1})^{-1} = V[f|\mu, p]$ is what the conditional variance of asset payoffs would be if the investor observed priors and prices, but no private signals.

The sequence of events is as in the partial equilibrium problem, except that at time 2, prices p are revealed, in addition to private signals η.

As in partial equilibrium, learning about principal components of asset payoffs implies that prior and posterior variances have the same eigenvectors. This allows us to recast the problem in terms of eigenvalues. Recall the definition of the prior squared Sharpe ratio of risk factor i:

$$\theta_i^2 \equiv \frac{(E[\Gamma_i'(f - pr)])^2}{Var[\Gamma_i' f]} = \frac{(((I-B)\mu - A)'\Gamma_i)^2}{\Lambda_{ii}}. \tag{17}$$

[8] In the partial equilibrium problem the capacity constraint on signals was $|I + \Sigma\Sigma_\eta^{-1}| \leq e^{2K}$; in the general equilibrium setting it becomes $|I + \Sigma\Sigma_\eta^{-1} + \Sigma\Sigma_p^{-1}| \leq e^{2K}$.

The problem can be written as:

$$max_{\{\hat{\Lambda}_1,\cdots,\hat{\Lambda}_N\}} \frac{1}{2}\{-N + \sum_{i=1}^{N} X_i \frac{\Lambda_i}{\hat{\Lambda}_i} + \sum_{i=1}^{N} \theta_i^2 \frac{\Lambda_i}{\hat{\Lambda}_i}\}. \tag{18}$$

$$s.t. \quad \prod_{i=1}^{N} \frac{\Lambda_i}{\hat{\Lambda}_i} = \exp(2K)$$

$$\frac{\Lambda_i}{\hat{\Lambda}_i} \geq 1 + \frac{\Lambda_i}{\Lambda_{pi}}, \ \forall i$$

where X_i measures the magnitude of the exploitable pricing errors in risk factor i. If an investor becomes informed, his valuation of an asset, based on his private information, will deviate from the realized price. We call this deviation an 'exploitable pricing error'. X_i measures the period-1 expected squared pricing error.

Appendix H shows that exploitable pricing errors depend on how much asset prices are affected by true payoffs (fundamentals) and by asset supply shocks:

$$X_i = (1 - \Lambda_{B,i})^2 + \Lambda_{C,i}^2 \sigma_x^2, \tag{19}$$

where Λ_{Bi} and Λ_{Ci} are the i^{th} eigenvalues of B and C in proposition 7.

The first term shows that pricing errors increase when prices are less reflective of true payoffs. Λ_{Bi} captures the covariance of the i^{th} risk factor's price with its true payoff. When Λ_{Bi} is small, a low covariance makes exploitable pricing errors X_i large. For example, if a well-informed investor sees a low price and knows that the true payoffs are likely to be high, he can exploit this by buying the asset. The uninformed investor, on the other hand, knows little about the true payoff and cannot exploit this difference.

The second term shows that pricing errors increase when prices are more reflective of supply shocks. Λ_{Ci} is the weight of the i^{th} risk factor's supply shock on the factor's price. When Λ_{Ci} is high, supply shocks create noise in prices that is exploitable by a well informed investor. For example, if such an investor sees a low price and knows it is due to a high supply shock, he can exploit this by buying the asset. The uninformed investor, on the other hand, attributes this low price to fundamentals. In sum, risk factors wit a higher X_i are more valuable to learn about, because informed investors can make more profit.

Proposition 8. *In general equilibrium with a continuum of investors, each investor's optimal information portfolio uses all capacity to learn about one linear combination of asset payoffs. The linear combination weights are given by the eigenvector* Γ_i *associated with the highest value of the* *squared Sharpe ratio plus exploitable pricing error:* $\theta_i^2 + X_i$.

Proof in appendix G. Just as in partial equilibrium, each investor continues to specialize in learning about one risk factor. Again, the reason is that the objective function is convex in $\hat{\Lambda}_i$. The learning index $\theta_i^2 + X_i$ determines which risk factor the investor learns about. The most valuable risk factor to learn about has (i) a high expected return $\Gamma_i' E[f - pr]$, (ii) a large expected portfolio share $\Gamma_i' E[q]$, and (iii) a large exploitable pricing error. The size of exploitable pricing errors is determined by the fraction of investors who learn about risk factor i. This is a new effect that shows up in general equilibrium only.

D. Aggregate Information Portfolios and Asset Prices

The previous section characterized the optimal information and asset allocation for an individual investor. This section describes how these choices aggregate across investors.

Aggregate Information Allocation Learning is a strategic substitute: The more precise the posterior beliefs of the average investor about the payoff of risk factor i, the less valuable it is to learn about. Let Ψ be the average of investors' signal precision matrices $\Psi = \int_0^1 \Sigma_{\eta j}^{-1} dj$, where $\Sigma_{\eta j}$ is the variance-covariance matrix of the signals that investor j observes. Let $\Lambda_{\Psi i}$ be the eigenvalue of Ψ corresponding to the i^{th} risk factor.

Proposition 9. *There is strategic substitutability in learning:* $(X_i + \theta_i^2)$ *is a strictly decreasing, monotonic function of* $\Lambda_{\Psi i}$.

Proof is in appendix H. When other investors learn more about a risk factor, the size of its exploitable pricing errors and its expected return fall.

In equilibrium, ex-ante identical investors get the same expected utility from learning about any of the risk factors that the economy learns about. The substitutability result is crucial to preventing diversification in general equilibrium. Because of substitutability, investors want to learn about risk factors that other investors are not learning about. If learning were a strategic complement, investors would want to specialize in learning about the same risk factor. This would make information symmetric: All investors would face the same payoff variance. As a result, they would all want to hold a lot of the risk factor they learn about. Since markets must clear, and investors each have an equal benefit of holding the asset they specialized in, they end up holding an equal share of the market in expectation. Full diversification would arise.

Proposition 10. *The number of risk factors that the economy learns about is weakly increasing in the economy's aggregate capacity* K.

Proof in appendix I. How much investors learn about an asset is summarized by the aggregate precision of beliefs $\widehat{\Sigma}_a^{-1}$. Manipulating the price in proposition 7 tells us that as long as assets are in positive net supply ($\bar{x} > 0$), the increase in information about an asset (fall in $\widehat{\Sigma}_a$) will cause its expected return to fall:

$$E[f - pr] = \rho \widehat{\Sigma}_a \bar{x}. \tag{20}$$

A reduction in expected return makes assets less valuable to learn about (lower θ_i^2 in proposition 8). When more investors learn about a factor, the expected return on the assets that load heavily on that factor falls. This makes that factor less desirable to learn about.

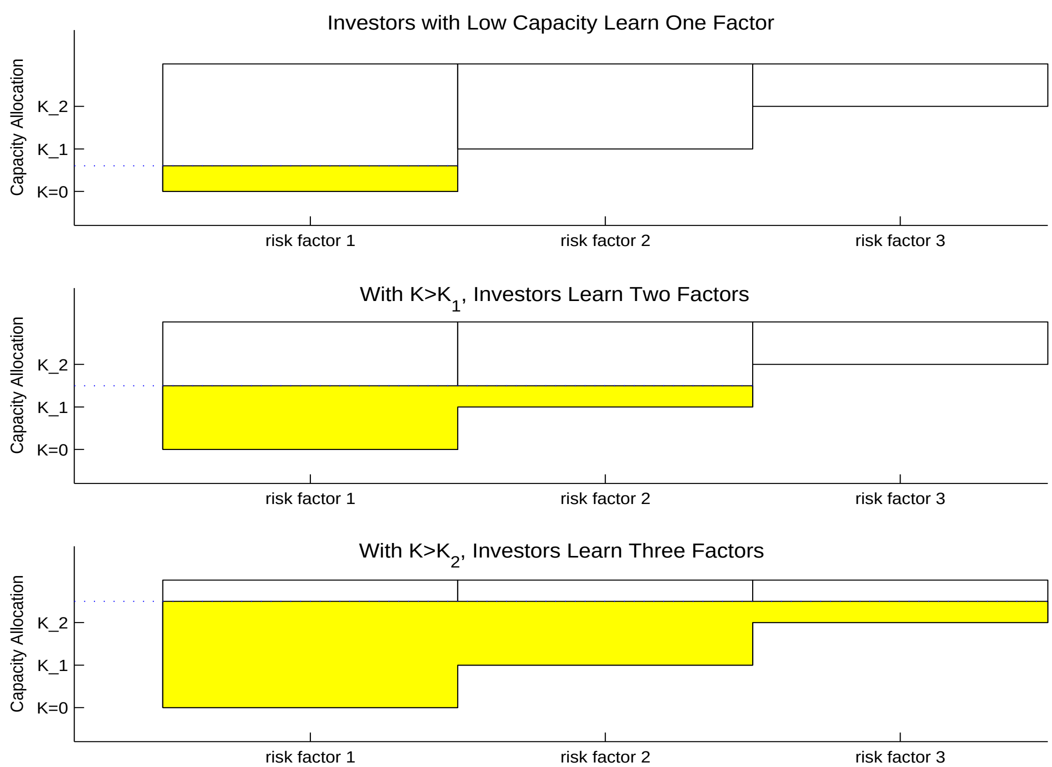

Figure 4. Aggregate allocation of information capacity for low, medium, and high levels of aggregate capacity K.

The equilibrium information allocations follow a cutoff rule. For intuition's sake, consider constructing the Nash equilibrium by letting investors sequentially choose how to allocate their capacity. The first investor learns about the risk factor that is most valuable when no other learning takes place. This is risk factor 1 in figure 4, the risk factor with the highest learning index $(\theta_i^2 + \rho^2 \sigma_x^2 \Lambda_i)$. This is the same risk factor as the one the partial equilibrium investor would learn about for $\sigma_x^2 = 0$. Subsequent investors will continue to allocate their capacity to factor 1 until the value of learning about risk factor 1 has dropped sufficiently that it equals the value of learning about the next most valuable risk factor 2. This cutoff is when capacity $K = K_1$ in figure 4. Then, some investors will find it beneficial to learn about risk factor 2. The proportions of investors that learn about 1 and about 2 is such that all investors remain indifferent. The reason that the first risk factor becomes gradually less valuable to learn about is that, as more investors become informed about it, $\hat{\Lambda}_{a1}^{-1}$ increases, thereby reducing both the exploitable pricing errors X_1 (proposition 9) and the equilibrium squared Sharpe ratio θ_1^2, the latter through a decrease in expected return (equation 20). Subsequent investors will continue to allocate capacity to these two risk factors, until all investors become indifferent between learning about 1, 2 and some third risk factor (where $K = K_2$ in figure 4). This process continues until all capacity is allocated. This type of result is referred to as 'water-filling' in the information theory literature. Because of strategic substitutability, the equilibrium aggregate information allocation is unique.

Asset Holdings in Equilibrium The cross-section of asset holdings is fully pinned down by the cross-section of information allocation. The mapping is as described in proposition 2. Each investor holds a diversified portfolio, plus a learning portfolio. The diversified portfolio now depends on information contained in prices. The learning portfolio contains assets in proportion to the one risk factor the investor learns about.

Atomless Investors and Limits to Arbitrage We assumed that there is a continuum of atomless investors, who by definition, cannot impact asset prices. This matters for equilibrium learning strategies because it makes the returns to learning unbounded. As our investor learns more about an asset, he can take larger and larger positions in that asset to fully exploit what he has learned, without worrying about his information being revealed through the price level. In contrast, an investor that is large in the market will move the asset price level when he trades. If he tries to exploit very precise information by taking large asset positions, his impact on the market price will partially reveal what he knows. This diminishes the value of his information and re-introduces decreasing returns to learning about a single risk factor. In figure 4, the investor is filling a bin on his own. For example, his capacity may exceed cutoff K_1.

Similar to the case where some risk is not learnable (section II.C), giving investors some mass in the market will make them want to specialize for low levels of capacity, but broaden their learning to multiple factors as capacity increases. In order to analyze a setting where large capacity investors interact, we would need to model strategic investors who consider the effect of their own learning on the price level (see Bernhardt and Taub (2005) or Pritsker (2005)). This question is beyond the scope of the current paper. In the conclusion, we return to the idea of modeling large portfolio managers.

Giffen Assets The fact that investors choose to learn about assets with large exploitable pricing errors means that assets are unlikely to be Giffen goods. In Admati (1985), investors' demand rises with the price of assets whose prices are very informative. Bhattacharya and Spiegel (1991) show that, with insufficient noise in prices (low X_i in our setting), a no-trade problem arises and markets break down. But when no trade in an asset is possible, information about the asset has no value. Less extremely, when good (bad) news about an asset's payoff makes its price rise (fall) dramatically, the value of such information is low (low X_i). Thus, rational learning should eliminate precisely the situations where such market anomalies are likely to arise.

E. Cross-Section of Asset Returns

An APT Representation of Asset Prices Our theory justifies the arbitrage-free pricing theory practice of using the principal components of the asset payoff matrix as priced risk factors (Ross (1976)). We can rewrite the risk premium on an asset i as the sum of its loading on each principal component k times the equilibrium risk premium of that principal component:

$$E[f_i - rp_i] = \sum_{k=1}^{n} \Gamma_{ik} \left(\Gamma_k' E[f - rp] \right)$$

The equilibrium risk premium of factor k can be rewritten, using equation (20) as:

$$\Gamma_k' E[f - rp] \;=\; \rho \hat{\Lambda}_{ak} (\Gamma_k' \bar{x}). \tag{21}$$

The equilibrium risk premium depends on (i) the risk aversion of the economy ρ, (ii) the supply of the risk factor $\Gamma_k' \bar{x}$, and most importantly (iii) the weight $\hat{\Lambda}_{ak}$, the k^{th} eigenvalue of aggregate variance matrix $\widehat{\Sigma}_a$. This weight measures how much the economy (the average investor) learns about risk factor k. A risk factor that the economy does not learn about has weight $\hat{\Lambda}_{ak} = \Lambda_k$. A risk factor that the economy learns about has a weight $\hat{\Lambda}_{ak} < \Lambda_k$. In other words, as more investors learn about risk factor k, $\hat{\Lambda}_{ak}$ decreases, its risk premium falls. An asset that loads heavily on those risk factors has a low risk premium. Our theory predicts which risk factors are learned about in equilibrium (see figure 4).

A CAPM Representation of Asset Prices The equilibrium asset prices and returns are equivalent to the prices and returns that would arise in a representative investor economy. That representative investor is endowed with the belief that payoffs f are normally distributed with mean $E_a[f]$ and covariance $\widehat{\Sigma}_a$, the heterogeneously informed investors' arithmetic average mean and harmonic average covariance (see equations 25 and 24 in appendix F). In our model with heterogenous information and partially revealing prices, a version of the Capital Asset Pricing Model holds.

Proposition 11. *If the market payoff is defined as* $f_m = \sum_{k=1}^{N}(\bar{x} + x_k)f_k$, *the market return is* $r_m = \frac{f_m}{\sum_{k=1}^{N}(\bar{x}+x_k)p_k}$, *and the return on an asset* i *is* $r_i = \frac{f_i}{p_i}$, *then the equilibrium price of asset* i *can be expressed as*

$$p_i = \frac{1}{r}\left(E_a[f_i] - \rho\, Cov_a[f_i, f_m]\right). \tag{22}$$

The equilibrium return is

$$E_a[r_i] - r = \frac{Cov_a[r_i, r_m]}{Var_a[r_m]}(E_a[r_m] - r) \equiv \beta_a^i(E_a[r_m] - r). \tag{23}$$

The proposition, similar to Lintner (1969), states that the equilibrium expected return on a security is proportional to its beta and to the market price of risk expressed in beta units of a representative investor. Without a theory of information acquisition, this pricing relationship is not testable. The information of the representative investor used in equations (22) and (23) cannot be observed by an econometrician or deduced from prices. Our contribution is to predict the information set of the representative investor.

Incorporating our results into the CAPM (equation 23) can explain why a public-information based CAPM under-prices large assets. In their seminal paper, Fama and French (1992) show that large firms offer lower average returns than small firms *for a given beta.* The standard CAPM fails to explain the cross-section of size portfolio returns because the beta for large (small) firms is 'too high' ('too low') to account for the return difference. This beta is based on public (prior) information. When investors can learn, the true risk of an asset depends on its 'learning beta' β_a^i, which is based on public information and private information investors have chosen to learn. Combining equations (17), (20), and proposition 8, the value of learning about risk factor i is given by $\rho^2\left(\frac{\hat{\Lambda}_i^a}{\Lambda_i}\right)(\Gamma_i'\bar{x})^2\hat{\Lambda}_i^a + X_i$. Learning value is increasing in the size of the risk factor $(\Gamma_i'\bar{x})$. If large assets load heavily on these large risk factors, the representative investor will be well-informed about large assets. Our findings suggest that any assets that load heavily on the largest principal components should have returns that are lower than the standard CAPM predicts.

IV. Active Portfolio Management

While the paper's original motivation was the composition of individual investor portfolios, the model also dictates optimal allocations of research and financial resources for institutional investors. It is a theory of active portfolio management. Through the lens of our theory, we see a specialized fund, such as a hedge fund or 'alpha-fund,' as an optimally under-diversified component of an institution's portfolio. Their investment strategy is to hold assets along one risk dimension in order to exploit the increasing returns to learning.

Optimal portfolio management is a long-standing issue in the mutual fund literature. The seminal paper by Treynor and Black (1973) departs from the efficient markets hypothesis by assuming that individual portfolio managers can exploit mis-pricing to make abnormal returns. Portfolio managers can analyze only a limited number of securities. For each such security k, they estimate the alpha $\alpha_k = r_k - r - \beta_i'(r^{div} - r) - \varepsilon_k$, where r^{div} represents diversified portfolio returns and ε_k is idiosyncratic risk, with variance $\sigma^2(\varepsilon_k)$. The optimal portfolio tilts away from the diversified one, towards securities with a high 'information ratio': $\alpha_k/\sigma^2(\varepsilon_k)$.[9]

Treynor and Black (1973) and our paper both recognize the fundamental trade-off between diversification and specialization. However, the theories differ along several dimensions. First, ours is an equilibrium pricing model of active portfolio management. There is no irrational mis-pricing. A Treynor and Black (1973) regression in our model will produce α's that capture *public* information already impounded in prices. If a portfolio manager purchased stocks with a positive (public) information ratio, his stocks would have prices that were depressed by privately informed investors' bad news. Our theory suggests another notion of α: Investors demand different risk premia for the same asset because they have an individual-specific α, arising from private information.

Second, instead of allowing investors to analyze a fixed set of securities, we examine the choice of what to learn. In both models, investors who learn about an asset's α want to take a large position in that asset.[10] But the feedback mechanism, where taking that large position makes an investor want to learn more about the asset, is unique to our setting.

[9] By defining q^{div} as the zero-capacity portfolio, we avoid a non-uniqueness problem of TB's portfolio decomposition. To understand the non-uniqueness, suppose that the optimal diversified portfolio contains shares of asset 1 and 2 in the ratio of 1 to 2. The market (asset supply) is 2 shares of each asset. The asset supply can be decomposed into one share of the diversified portfolio, plus one share of asset 1 in the learning portfolio, or alternatively into 2 shares of the diversified portfolio and two shares sold short of asset 2.

[10] Our model is formally equivalent to one where investors learn about α.

V. Conclusion

Most theories of portfolio allocation and asset pricing take investors' information as given. Investigating optimal information choice has the potential to yield valuable insights into many portfolio and asset pricing puzzles. When investors can choose what information to acquire, given a fixed information capacity, they optimally devote all capacity to learning about the payoffs (or alphas) of a set of highly correlated assets. Even when some risk is not learnable, and there are diminishing returns to learning, investors still specialize in a small number of risks. Since risk-averse investors prefer to take larger positions in assets they are better-informed about, high-capacity investors hold larger 'learning portfolios', causing their total portfolio to be less diversified. In equilibrium, investors specialize in different assets from other investors. Ex-ante identical investors may optimally hold different portfolios.

The model has new cross-sectional asset pricing predictions. Assets that many investors learn about command lower risk premia, than standard asset pricing models predict. These assets are ones that co-vary with the largest principal components.

While this model has focussed on a static information allocation problem, it could be extended along many dimensions. The quantity of information could be endogenized with a capacity production function. A model of dynamic information choice, incorporating recent advances in the dynamics of information value (Bernhardt and Taub (2005)), could be used to explain the persistence and turnover of investor portfolio holdings and time variation in expected asset returns. Finally, analyzing a market for information capacity could answer questions about the organization of the portfolio management industry.

A natural question to pose in this setting is: "Why can't an investor delegate his portfolio management to someone who processes information for many investors?" If a manager were to sell information, information resale would undermine their profits. To avoid this problem, they would manage investors' portfolios directly. If information capacity were costly, then managers would maximize profit by each specializing in a different risk factor. Whether an investor's portfolio will also be concentrated hinges on how portfolio managers set prices. We conjecture that a competitive equilibrium would have fund managers charge a non-linear fee. Such quantity discounts induce more investment in a fund. The additional investment reduces the fund manager's per-share cost and allows him to compete linear-price suppliers out of the market. Quantity discounts make investing small amounts in many funds costly. Competitive pricing of portfolio management services forces investors to internalize increasing returns to specialization; optimal under-diversification reappears.

A theory of information choice in financial markets is vital to understanding or justifying the active portfolio management industry. While a formal analysis of the industry and its effect on investor portfolios is left for future work, understanding an individual's information choice problem is a necessary first step.

References

Admati, Anat, "A Noisy Rational Expectations Equilibrium for Multi-Asset Securities Markets," *Econometrica*, 1985, *53(3)*, 629–57.

Barber, Brad and Terrance Odean, "The Internet and the Investor," *Journal of Economic Perspectives*, 2001, *15(1)*, 41–54.

Bernhardt, Dan and Bart Taub, "Strategic Information Flows In Stock Markets," 2005. Working Paper.

Bhattacharya, Utpal and Matthew Spiegel, "Insiders, Outsiders, and Market Breakdowns," *Review of Financial Studies*, 1991, *4(2)*, 255–282.

Biais, Bruno, Peter Bossaerts, and Chester Spatt, "Equilibrium Asset Pricing Under Heterogenous Information," 2004. Working Paper.

Chevalier, Judith and Glenn Ellison, "Are Some Mutual Fund Managers Better than Others? Cross-Sectional Patterns in Behavior and Performance," *Journal of Finance*, 1999, *54(3)*, 875–899.

Coval, Joshua and Tobias Moskowitz, "Home Bias at Home: Local Equity Preference in Domestic Portfolios," *Journal of Finance*, 1999, *54*, 1–54.

___ and ___ , "The Geography of Investment: Informed Trading and Asset Prices," *Journal of Political Economy*, 2001, *109(4)*, 811–41.

___ , David Hirshleifer, and Tyler Shumway, "Can Individual Investors Beat the Market?," 2003. Dice Center Working Paper 2003-5.

Cover, Thomas and Joy Thomas, *Elements of Information Theory*, first ed., John Wiley and Sons, New York, New York, 1991.

Epstein, Larry and Stanley Zin, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 1989, *57*, 937–69.

Garcia, Diego and Joel Vanden, "Information Acquisition and Mutual Funds," 2005. Working Paper.

Goetzmann, William and Alok Kumar, "Diversification Decisions of Individual Investors and Asset Prices," 2003. Working Paper.

Hansen, Lars and Thomas Sargent, *Recursive Models of Dynamic Linear Economies*, Monograph, 2004.

Ivkovic, Zoran and Scott Weisbenner, "Local Does as Local Is: Information Content and the Geography of Individual Investors' Common Stock Investments," *Journal of Finance*, 2005, *60*, 267–306.

——, **Clemens Sialm, and Scott Weisbenner**, "Portfolio Concentration and the Performance of Individual Investors," 2005. Working Paper.

Kacperczyk, Marcin, Clemens Sialm, and Lu Zheng, "On the Industry Concentration of Actively Managed Equity Mutual Funds," *Journal of Finance*, 2005, *Forthcoming.*

Keppo, Jussi, Giuseppe Moscarini, and Lones Smith, "The Demand For Information: More Heat Than Light," 2005. NBER Working Paper.

Lintner, John, "The Aggregation of Investors'Diverse Judgements and Preferences in Purely Competitive Markets," *Journal of Financial and Quantitative Analysis*, 1969, *4(4)*, 347–400.

Massa, Massimo and Andrei Simonov, "Hedging, Familiarity and Portfolio Choice," *The Review of Financial Studies*, 2005, *Forthcoming.*

Merton, Robert, "A Simple Model of Capital Market Equilibrium with Incomplete Information," *Journal of Finance*, 1987, *42(3)*, 483–510.

Peng, Lin, "Learning with Information Capacity Constraints," *Journal of Financial and Quantitative Analysis*, 2004, *40(2).*

Polkovnichenko, Valery, "Household Portfolio Diversification: A Case for Rank Dependent Preferences," *The Review of Financial Studies*, 2004, *Forthcoming.*

Pritsker, Matthew, "A Fully-Rational Liquidity-Based Theory of IPO Underpricing and Underperformance," 2005. Working Paper.

Radner, Roy and Joseph Stiglitz, "A Nonconcavity in the Value of Information," in M. Boyer and R.E. Kihlstrom, eds., *Bayesian Models in Economic Theory*, Elsevier Science Publishers B.V., 1984.

Ross, Stephen, "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic Theory*, 1976, *13*, 341–360.

Shannon, Claude E., "Mathematical Theory of Communication," *Bell System Technology Journal*, 1948, *27*, 379–423 and 623–656.

Sims, Christopher, "Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50(3)*, 665–90.

Treynor, Jack and Fischer Black, "How To Use Security Analysis to Improve Portfolio Selection," *The Journal of Business*, 1973, *46(1)*, 66–86.

Verrecchia, Robert, "Information Acquisition in a Noisy Rational Expectations Economy," *Econometrica*, 1982, *50(6)*, 1415–1430.

Wilson, Robert, "Informational Economies of Scale," *Bell Journal of Economics*, 1975, *6*, 184–95.

A. Appendix

A. Proof of Proposition 1

Consider a deviation from this solution that would allocate some capacity to another asset j, s.t. $\widehat{\Sigma}_{jj} = (1-\epsilon)\Sigma_{jj}$. Keeping total capacity constant implies that $\widehat{\Sigma}_{ii}$ must be increased by a factor of $1/(1-\epsilon)$. This deviation produces a net utility change

$$(\mu_j - p_j r)^2 \Sigma_{jj}^{-1} \left((1-\epsilon)-1\right) + (\mu_i - p_i r)^2 \Sigma_{ii}^{-1} \left(1-(1-\epsilon)\right)$$

Since i is the asset for which $(\mu_i - p_i r)^2 \Sigma_{ii}^{-1} > (\mu_j - p_j r)^2 \Sigma_{jj}^{-1}$, for all $j \neq i$, the net utility change from the deviation is negative. □

B. Arbitrary Risk Factor Structure

For tractability, the main text assumes that investors learn about the principal components of $\Sigma = \Gamma'\Lambda\Gamma$, $f'\Gamma$. In this appendix we show that the results of specialization does not hinge on this assumption. Let the posterior covariance matrix be $\widehat{\Sigma} = \widehat{\Gamma}'\hat{\Lambda}\widehat{\Gamma}$, where $\widehat{\Gamma}$ is an *arbitrary but fixed* posterior eigenvector matrix of $\widehat{\Sigma}$. In particular, $\widehat{\Gamma} \neq \Gamma$. The objective function (8) in the partial equilibrium model ($V_{ER} = \Sigma$) then becomes:

$$\frac{1}{2} Tr\left(\widehat{\Gamma}\hat{\Lambda}^{-1}\widehat{\Gamma}'\Sigma - I\right) + \frac{1}{2}(\mu - pr)'\widehat{\Gamma}\hat{\Lambda}^{-1}\widehat{\Gamma}'(\mu - pr).$$

One can write the ii element of the matrix $\widehat{\Sigma}^{-1}\Sigma$ as: $\left(\widehat{\Sigma}^{-1}\Sigma\right)_{ii} = \sum_{j=1}^{N}\sum_{l=1}^{N} \widehat{\Gamma}_{il}\hat{\Lambda}_l^{-1}\widehat{\Gamma}_{lj}\Sigma_{ji}$. Since the trace of a matrix is the sum of its diagonal elements, the objective function can be written as a sum of the $\{\hat{\Lambda}_l^{-1}\}$ each weighted by a scalar τ_l: $\frac{1}{2}\sum_{l=1}^{N}\tau_l\hat{\Lambda}_l^{-1} - \frac{N}{2}$, where τ_l is the learning index of the l^{th} risk factor:

$$\tau_l = \sum_{i=1}^{N}\sum_{j=1}^{N}\widehat{\Gamma}_{il}\left(\widehat{\Gamma}_{lj}\Sigma_{ji} + \widehat{\Gamma}_{jl}(\mu_i - p_i r)(\mu_j - p_j r)\right)$$

Under the assumption $\widehat{\Gamma} = \Gamma$, we recover the learning index written in the main text: $\tau_l = \Lambda_l + ((\mu - pr)'\Gamma_l)^2$. But, whatever the exact specification of the learning index is, the investor always ranks the risk factors according to their learning index τ_l, and chooses to specialize in the risk factor with the highest τ_l. The reason is that the objective function is still linear in the $\{\hat{\Lambda}_l^{-1}\}$, or convex in the choice variables $\{\hat{\Lambda}_l\}$. The first main result of specialization is not affected by the exact structure of risks that the investor learns about. This argument applies equally to sections II.A and B, i.e. Σ diagonal or not, as this only changes τ_l.

The same argument holds for the general equilibrium model of section III ($V_{ER} = Q - \widehat{\Sigma}$, where $Q = \Sigma + B\Sigma B' + CC'\sigma_x^2 - 2\Sigma B'$). Just replace the terms Σ_{ji} by Q_{ji} and $\mu - pr$ by $(I-B)\mu - A$ in the learning index. The definitions of the price coefficients A, B, and C are in the proof of proposition 7. The second main result of strategic substitutability also does not depend on the exact structure of risks. First, note that the proof of the equilibrium price in appendix F does not depend on the risk structure. Second, the effect of others' learning choices enter through A and B in the learning index τ_l. For any given risk, the value of learning about that risk declines the more others know about it.

The same argument can be made starting from some arbitrary but fixed *signal* eigenvector matrix $\widetilde{\Gamma} \neq \Gamma$ ($\Sigma_\eta = \widetilde{\Gamma}\Lambda_\eta\widetilde{\Gamma}'$), rather than an arbitrary *posterior* eigenvector matrix $\widehat{\Gamma} \neq \Gamma$. Because $\widehat{\Sigma}^{-1} = \Sigma^{-1} + \Sigma_\eta^{-1}$, $Tr(\widehat{\Sigma}^{-1}\Sigma - I) = Tr(\Sigma_\eta^{-1}\Sigma)$. The algebra (and the definition of the learning index) is then identical, simply replacing $\widehat{\Gamma}$ by $\widetilde{\Gamma}$ and $\hat{\Lambda}_l$ by $\Lambda_{\eta l}$. So, the objective function is a linear function of the inverse signal eigenvalues $\{\Lambda_{\eta l}^{-1}\}$, or convex in the $\{\Lambda_{\eta l}\}$, so that specialization arises again.

C. Proof of Corollary 1

Proposition 2 shows that an investor optimally chooses a portfolio with a low level of diversification, meaning a low $(|q^{div}|/(|q^{div}|+|q^{learn}|)$, if and only if he has a higher information capacity. What remains to be shown is that a higher information capacity entails a higher expected profit: $E[q'(f - rp)]$.

The portfolio weights q can be decomposed into q^{div}, the zero-capacity portfolio and $q_i^{learn} = \frac{1}{\rho\Sigma_{ii}}(\hat{\mu}_i - p_i r)(e^{2K} - 1)$. The profit from the diversified portfolio $E[q^{div'}(f - rp)]$ does not vary in the information capacity K. The profit from the learning portfolio is $E\left[\frac{1}{\rho\Sigma_{ii}}(\hat{\mu}_i - p_i r)(e^{2K} - 1)(f_i - rp_i)\right]$. This is increasing in K if $E[(\hat{\mu}_i - p_i r)(f_i - rp_i)] > 0$. Since the difference between f_i and $\hat{\mu}_i$ is a mean-zero, orthogonal expectation error,

$$E[(\hat{\mu}_i - p_i r)(f_i - rp_i)] = E[(\hat{\mu}_i - p_i r)^2] + 0 > 0.$$

□

D. Proof of Proposition 4

Proof: An investor learns about a risk factor whenever the marginal benefit of allocating the first increment of capacity to that risk factor $(1+\theta_i^2)\frac{\Lambda_i}{-\hat{\Lambda}_i^2}$ exceeds its marginal cost: $\xi\frac{1}{\hat{\Lambda}_i - \alpha\Lambda_i} - \phi_i$. K enters this inequality only through the Lagrange multiplier ξ, the shadow cost of capacity. When an investor learns about asset i, the no-negative learning constraint is no longer binding and $\phi_i = 0$. For each risk factor i, there is a cutoff value $\xi_i^\star = \frac{\Lambda_i}{\hat{\Lambda}_i}(1 - \alpha\frac{\Lambda_i}{\hat{\Lambda}_i})\left(1+\theta_i^2\right)$ where marginal benefit and cost are equal. For all $\xi < \xi_i^\star$, the marginal benefit is greater than the marginal cost and the investor will learn about risk factor i. We know from the proof of proposition 5 that $\partial\xi/\partial K \leq 0$. Therefore, the number of factors i for which $\xi < \xi_i^\star$ must be an increasing step function in K. □

Proof of Corollary 2

Proof: From the proof of proposition 2, we know that a non-zero quantity of an asset is held in the learning fund whenever the investor learns about the asset and the expected excess return is not equal to zero. Getting a signal from a continuous distribution that implies a zero excess return is a zero probability event. Since asset payoffs are independent, each risk factor corresponds to one and only one asset. Proposition 4 shows that when capacity increases, the number of risk factors learned about rises. Thus the number of assets learned about rises, and the number of different assets held in the learning fund rises. □

E. Proof of Proposition 5

Proof: The diversified portfolio q^{div} is what the investor would hold with zero capacity. It does not change as capacity rises. When K rises, the absolute value of $E[q^{learn}] = \frac{1}{\rho}(\widehat{\Sigma}^{-1} - \Sigma^{-1})(\mu - pr)$ is affected only through $\widehat{\Sigma}$. How K affects $\widehat{\Sigma}$ can be seen in the first-order condition; it enters through the Lagrange multiplier ξ. Solving for $(\frac{\Lambda_i}{\hat{\Lambda}_i} - \alpha \frac{\Lambda_i}{\hat{\Lambda}_i}^2)$ from equation (11) and substituting it into the capacity constraint (10) yields an expression for the multiplier $N \log(\xi) = \sum_{i=1}^{N} \left(2 \log(\frac{\Lambda_i}{\hat{\Lambda}_i}) + \log(1 + \theta_i^2 + \phi_i) \right) + \log(1 - \alpha) - 2K$, which is decreasing in K. Applying the implicit function theorem to the first-order condition (11), we get $\frac{\partial \frac{\Lambda_i}{\hat{\Lambda}_i}}{\partial K} \geq 0$, for every risk factor i, with strict inequality for those risk factors that are learned about. Since $\{\frac{\Lambda_i}{\hat{\Lambda}_i}\}$ are the eigenvalues of $\widehat{\Sigma}^{-1}\Sigma$, and Σ is a constant, $\frac{\partial \frac{\Lambda_i}{\hat{\Lambda}_i}}{\partial K} \geq 0$ implies that each element of the eigenvalue matrix $\hat{\Lambda}^{-1}$ of $\widehat{\Sigma}^{-1} = \Gamma \hat{\Lambda}^{-1} \Gamma'$, weakly rises with K. As a result, $\frac{\partial |E[q^{learn}]|}{\partial K} \geq 0$, with strict inequality for the risk factors that the investor learns about. □

F. Proof of Proposition 7

From Admati (1985), we know that equilibrium price takes the form $rp = A + Bf + Cx$ where

$$\begin{aligned} A &= \left(\Sigma^{-1} + \frac{1}{\rho^2 \sigma_x^2} \Psi' \Psi + \Psi \right)^{-1} \left(\Sigma^{-1} \mu - \rho \bar{x} \right), \\ B &= \left(\Sigma^{-1} + \frac{1}{\rho^2 \sigma_x^2} \Psi' \Psi + \Psi \right)^{-1} \left(\Psi + \frac{1}{\rho^2 \sigma_x^2} \Psi' \Psi \right), \\ C &= -\left(\Sigma^{-1} + \frac{1}{\rho^2 \sigma_x^2} \Psi' \Psi + \Psi \right)^{-1} \left(\rho I + \frac{1}{\rho \sigma_x^2} \Psi' \right). \end{aligned}$$

Here, Ψ is the average of investors' signal precision matrices $\Psi = \int_0^1 \Sigma_{\eta j}^{-1} dj$, and $\Sigma_{\eta j}$ is the variance-covariance matrix of the signals that investor j observes.[11]

Using (14), note that $\left(\Sigma^{-1} + \frac{1}{\rho^2 \sigma_x^2} \Psi' \Psi + \Psi \right)^{-1} = \hat{\Sigma}_a$, the posterior variance for an investor with the

[11] The Lebesgue integral may not be well defined when $\{\eta_j\}$ are processes of independent random variables for a continuum of investors j, because realizations may not be measurable with respect to the joint space of parameters and samples. Also, the sample function giving each investor's individual shock may not be Lebesgue measurable, and thus the fraction of investors associated with each shock may not be well defined. Making independence compatible with joint measurability requires defining an enriched probability space, where the one-way Fubini property holds. Then the exact law of large numbers is restored. See Hammond and Sun (2003), and Duffie and Sun (2004) for recent solutions.

average of all investors' posterior precisions:

$$\hat{\Sigma}_a \equiv \left(\int_0^1 \hat{\Sigma}_j^{-1} dj \right)^{-1} \tag{24}$$

Note also that $\Sigma_p \equiv \sigma_x^2 B^{-1} CC' B^{-1'} = (\frac{1}{\rho^2 \sigma_x^2} \Psi' \Psi)^{-1}$.

Then, the price equation can be rewritten as

$$rp = \widehat{\Sigma}_a (\Sigma^{-1}\mu + \Psi f + \Sigma_p^{-1}(f - \rho \Psi^{-1} x) - \rho(\bar{x} - x)$$

Simple algebra reveals that $(f - \rho\Psi^{-1}x) = B^{-1}(rp - A)$, the unbiased signal that investors observe from the price level. From equation 3, we note that the first three terms are equal to the posterior mean of the 'average' investor's beliefs:

$$\hat{\mu}_a \equiv E_a[f_i] \equiv \int_0^1 \hat{\mu}_j dj \tag{25}$$

Thus,

$$rp = \hat{\mu}_a - \rho \widehat{\Sigma}_a (\bar{x} + x). \tag{26}$$

The price level is increasing in the posterior belief of the average investor about the mean payoff, and decreasing in risk aversion, the amount of risk the average investor bears, and the supply of the asset.□

G. Proof of Proposition 8

As long as $\widehat{\Sigma}^{-1}$ and V_{ER} have the same eigenvectors as Σ, then the proof of the proposition follows immediately from the proof of proposition 4, where $E[f - pr]$ is now based on prior beliefs ($E[f - pr] = (I - B)\mu - A$), instead of on $(\mu - pr)$. Sums, products and inverses of matrices with identical eigenvectors preserve those eigenvectors. This tells us that Ψ can be rewritten as $\Psi = \int_0^1 \Gamma^{-1'} \Lambda_{\eta j}^{-1} \Gamma^{-1} dj$. Since eigenvector matrices have the property that $\Gamma^{-1} = \Gamma'$, and defining $\Lambda_{\eta a}^{-1} = \int_0^1 \Lambda_{\eta j}^{-1} dj$, this is equivalent to $\Psi = \Gamma \Lambda_{\eta a} \Gamma'$. Because Σ_p, $\hat{\Sigma}_a$, B, and C are result from a combination of sums, products and inverses of Σ and Ψ (see appendix F), all have eigenvectors Γ. □

H. Proof of Proposition 9

Proof: X_i and θ_i^2 are both decreasing in $\Lambda_{\psi i}$. Thus, their sum is decreasing. We start by deriving the expression for X_i. The first part of the objective is $Tr\left(\widehat{\Sigma}^{-1} V_{ER}\right)$, which we rewrite as $Tr\left(\widehat{\Sigma}^{-1}\Sigma\Sigma^{-1}(V_{ER} + \widehat{\Sigma} - \widehat{\Sigma})\right)$. This is $Tr\left(\widehat{\Sigma}^{-1}\Sigma\Sigma^{-1}(V_{ER} + \widehat{\Sigma}) - I\right)$ or $Tr\left(\widehat{\Sigma}^{-1}\Sigma\Sigma^{-1}(V_{ER} + \widehat{\Sigma})\right) - N$. The trace is the product of the eigenvalues. Let $\frac{\Lambda_i}{\hat{\Lambda}_i}$, be the ratio of the precision of the posterior to the precision of the prior, i.e. it is the i^{th} eigenvalue of $\widehat{\Sigma}^{-1}\Sigma$: $\frac{\Lambda_i}{\hat{\Lambda}_i} \equiv \hat{\Lambda}_i^{-1}\Lambda_i$. Let X_i be the i^{th} eigenvalue of $\Sigma^{-1}(V_{ER} + \widehat{\Sigma})$. Then the i^{th} eigenvalue of the matrix inside the trace is $\frac{\Lambda_i}{\hat{\Lambda}_i} X_i$, and $Tr\left(\widehat{\Sigma}^{-1}\Sigma\Sigma^{-1}(V_{ER} + \widehat{\Sigma})\right) = \sum_{i=1}^N X_i \frac{\Lambda_i}{\hat{\Lambda}_i}$. This is because Σ, $\widehat{\Sigma}$,

B and C all share the same eigenvectors Γ.

The expression for X_i, the i^{th} eigenvalue of $\Sigma^{-1}(V_{ER}+\widehat{\Sigma})$, is $X_i = \Lambda_i^{-1}\left[\Lambda_i\left(1+\Lambda_{Bi}^2-2\Lambda_{Bi}\right)+\Lambda_{Ci}^2\sigma_x^2\right]$, where Λ_{Bi} and Λ_{Ci} are the i^{th} eigenvalue of B and C respectively. Using the definition of B and $\hat{\Sigma}_a$, we can rewrite B as $I-\hat{\Sigma}_a\Sigma^{-1}$, which has eigenvalues $\Lambda_{Bi} = 1-\hat{\Lambda}_{ai}\Lambda_i^{-1}$, where Λ_{ai} is the i^{th} eigenvalue of $\hat{\Sigma}_a$. Also, using the definitions of B and C, we have $C = -\rho B\Psi^{-1}$, and hence $CC'\sigma_x^2 = B\left(\frac{1}{\rho^2\sigma_x^2}\Psi'\Psi\right)^{-1}B'$, which in turn equals $B\Sigma_p B'$. The i^{th} eigenvalues of the matrix $\Sigma^{-1}CC'\sigma_x^2$, $\Lambda_i^{-1}\Lambda_{Ci}^2\sigma_x^2$, are thus equal to $\Lambda_{ii}^{-1}\Lambda_{Bi}^2\Lambda_{pi} = \Lambda_i^{-1}(1-\hat{\Lambda}_{ai}\Lambda_i^{-1})^2\Lambda_{pi}$. Now we can rewrite X_i as:

$$X_i = \left(\frac{\hat{\Lambda}_{ai}}{\Lambda_i}\right)^2 + \left(\frac{\Lambda_{pi}}{\Lambda_i}\right)\left(1-\frac{\hat{\Lambda}_{ai}}{\Lambda_i}\right)^2.$$

An important property of X_i is that it is decreasing in the average signal precision of risk factor i, $\Lambda_{\Psi i}$, the i^{th} eigenvector of Ψ. To ease the burden of notation, define $a \equiv \frac{1}{\rho^2\sigma_x^2}$, $g \equiv \Lambda_{ii}^{-1}$, and $x \equiv \Lambda_{\Psi i}$. To show strict substitutability is to show $\frac{\partial X_i}{\partial x} < 0$. We first recall that $\Lambda_{pi}^{-1} = ax^2$ and $\hat{\Lambda}_{ai}^{-1} = g+ax^2+x$. We can rewrite X_i in our new notation as:

$$\begin{aligned} X_i &= g^2(g+ax^2+x)^{-2} + ga^{-1}x^{-2}(ax^2+x)^2(g+ax^2+x)^{-2}, \\ &= g(g+ax^2+x)^{-2}[g+a^{-1}+2x+ax^2]. \end{aligned}$$

Taking a partial derivative with respect to x, we get:

$$\begin{aligned} \frac{\partial X_i}{\partial x} &= -2g(g+ax^2+x)^{-3}[(g+a^{-1}+2x+ax^2)(2ax+1)-(g+ax^2+x)(ax+1)], \\ &= -2g(g+ax^2+x)^{-3}[a^2x^3+3ax^2+(3+ag)x+a^{-1}]. \end{aligned}$$

The partial derivative is strictly negative because $g > 0$, $a > 0$, $x > 0$, and hence the term in parentheses and the term in brackets are strictly positive. Using L'Hôpital's rule, it is easy to show that $lim_{x\to 0}X_i = 1+a^{-1}g^{-1}$, which equals $1+\rho^2\sigma_x^2\Lambda_i$. Because of the new source of risk induced by noisy asset supply (σ_x^2), X_i is strictly greater than 1 when nobody learns about risk factor i ($x = \Lambda_{\Psi i} = 0$). Note that this is consistent with $X_i = 1$ in partial equilibrium, where prices we taken as given ($\sigma_x^2 = 0$).

We conclude by showing that $\theta_i^2 = \frac{\left(((I-B)\mu-A)'\Gamma_i\right)^2}{\Lambda_i}$ is decreasing in $\Lambda_{\psi i}$. The denominator Λ_{ii} is exogenous. Using the formulas for A and B in appendix F, the expected return is $((I-B)\mu-A) = \rho\widehat{\Sigma}^a\bar{x}$. Thus, $((I-B)\mu-A)'\Gamma_i = \rho\hat{\Lambda}_i^a(\Gamma_i'\bar{x})$. Since, $\hat{\Lambda}_i^a = (\Lambda_i^{-1}+\frac{1}{\rho^s\sigma_x^2}\Lambda_{\psi i}^2+\Lambda_{\psi i})^{-1}$, the expected return and its square are decreasing in $\Lambda_{\psi i}$. $\square$

I. Proof of Proposition 10

From proposition 8, we know that investors always allocate their capacity to the asset with the highest value of $(\Gamma_i'E[f-pr])^2(\Lambda_i)^{-1}+X_i$. Begin by ordering risk factors by their learning index values when $K = 0$, such that $(\Gamma_i' E[f-pr])^2(\Lambda_i)^{-1}+X_i \geq (\Gamma_{i+1}' E[f-pr])^2(\Lambda_{i+1})^{-1}+X_{i+1}$. For small levels of K, capacity is allocated only to the most valuable risk factor, say risk factor 1, and to additional risk factors, only if their initial learning index value is equal to that of factor 1: $(\Gamma_1' E[f-pr])^2(\tilde{\Lambda}_1)^{-1}+X_1 = (\Gamma_i' E[f-pr])^2(\tilde{\Lambda}_i)^{-1}+X_i$. Is there some level of capacity K_j such that these two index levels are equal? For any non-zero index, there must be. As $K \rightarrow \infty$, precision of beliefs about asset 1 becomes infinite: $\Psi_{11} \rightarrow \infty$. Equation (20), shows that that $(\Gamma_1' E[f-pr])^2 \rightarrow 0$. Furthermore $X_i \rightarrow 0$ because there are no more exploitable pricing errors when beliefs are infinitely precise. Since index values are non-negative, there is some K_j for each asset j s.t. $\forall K > K_j$, investors learn about risk factor j. $\square$

J. Proof of Proposition 11

We can rewrite equation (26) for each asset $i \in \{1, 2, \cdots, N\}$ separately:

$$p_i = \frac{1}{r}\left(\hat{\mu}_a^i - \rho \sum_{k=1}^{N} Cov_a[f_i, f_k](\bar{x}+x_k)\right) = \frac{1}{r}\left(\hat{\mu}_a^i - \rho\, Cov_a[f_i, \sum_{k=1}^{N}(\bar{x}+x_k)f_k]\right)$$

where $Cov_a[f_i, f_k]$ denotes the (i,k) element of $\hat{\Sigma}_a$. Using the definition of f_m stated in the proposition, we obtain the first equation mentioned in the proposition: $p_i = \frac{1}{r}\left(E_a[f_i] - \rho\, Cov_a[f_i, f_m]\right)$. To rewrite this equilibrium price function in terms of returns divide both sides by the price. Denote the return on security i by $r^i \equiv \frac{f_i}{p_i}$. Simple manipulation leads to:

$$E_a[r_i] - r = \rho\, Cov_a[r_i, f_m]. \tag{27}$$

This is true for each asset i, and hence also for asset m:

$$E_a[r_m] - r = \rho\, p_m Cov_a[r_m, r_m]. \tag{28}$$

Solving (28) for the risk aversion coefficient ρ, and substituting it into (27), we get the second equation in the proposition.$\square$