

INFORMATION ACQUISITION AND UNDER-DIVERSIFICATION

Stijn Van Nieuwerburgh and Laura Veldkamp*
New York University, Stern School of Business and NBER

November 13, 2008

Abstract

If an investor wants to form a portfolio of risky assets and can exert effort to collect information on the future value of these assets before he invests, which assets should he learn about? The best assets to acquire information about are ones the investor expects to hold. But the assets the investor holds depend on the information he observes. We build a framework to solve jointly for investment and information choices, with general preferences and information cost functions. Although the optimal research strategies depend on preferences and costs, the main result is that the investor who can first collect information systematically deviates from holding a diversified portfolio. Information acquisition can rationalize investing in a diversified fund and a concentrated set of assets, an allocation often observed, but usually deemed anomalous.

*Van Nieuwerburgh: 44 West Fourth St., 9-190, New York, NY 10012; email:svnieuw@stern.nyu.edu, Phone:212-998-0673, Fax: 212-995-4256. Veldkamp: 44 West Fourth St., 7-180, New York, NY 10012 (email:lveldkam@stern.nyu.edu). We thank Andy Atkeson, Dave Backus, Hal Cole, Bernard Dumas, Ned Elton, Fatih Guvenen, Bob Hall, Ron Kaniel, Ralph Koijen, Hanno Lustig, Massimo Massa, Pascal Maenhout, Lin Peng, 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, the Federal Reserve Board, Chicago Fed, Kellogg, European Central Bank, Stockholm Institute for Financial Research, BI Oslo, Tilburg, and participants at the 2005 Gerzensee Summer AP Institute, the 2005 Econometric Society World Congress, the 2005 SED, EFA and FMA meetings, 2006 Econometric Society and AEA meetings, SEMSI workshop on model uncertainty, Caesarea conference on Capital Markets and Risk Management for helpful comments. A previous version of this paper was titled "Information Acquisition and Portfolio Under-Diversification." **JEL classification:** D83, D82, G11. **Keywords:** Information choice, portfolio theory, learning, asymmetric information.

The asset management industry is in the business of acquiring information and using that information to manage a portfolio of assets. While the rationality of asset portfolios and financial services is hotly contested, the debate has largely ignored the role of information. Little is known about what information rational-expectations investors should learn. Since the information learned determines which assets are invested in, understanding information acquisition is central to understanding investment behavior. This paper takes the first step in investigating this interplay between information choice and the investment problem.

Investment choices with given information and information choices with a single risky asset have each been studied before.¹ But analyzing the two choices jointly delivers new insights. Specifically, the feedback of one decision on the other can generate gains to specialization. When choosing information, investors can acquire noisy signals about future payoffs of many assets, or they can specialize and acquire more precise signals about fewer 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 benefit than the same signal applied to many shares. Specialization then 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. Standard investment theory is challenged by the high degree of concentration in observed portfolios. A joint learning-investment model can rationalize such concentration.

We explore how to jointly model investment and information choices. To make the logic as transparent as possible, we focus on a one-period, partial equilibrium model with independent assets and independent signals. Unlike Merton (1987), information is not required to hold an asset; rather it is a tool to reduce the conditional variance (the uncertainty) of the asset's payoff, as in Grossman and Stiglitz (1980). We endow an investor with a given amount of information and ask: How much should the investor learn about each asset? After the investor observes signals drawn from the distribution whose precision he has chosen, he solves a standard portfolio problem.

¹Recent work on the role of information in portfolio choice includes: Banerjee (2007), Maenhout (2004), Biais, Bossaerts and Spatt (2004), Bernhardt and Taub (2005), Kodres and Pritsker (2002), Albuquerque, Bauer and Schneider (2005), Wang (1993) and is reviewed by Brunnermeier (2001). Recent work on learning with one risky asset includes: Barlevy and Veronesi (2000), Bullard, Evans and Honkapohja (2005), Cagetti, Hansen, Sargent and Williams (2002), Peress (2004). Learning by a representative agent is modeled by Timmermann (1993), Sims (2003) and Peng (2004); these models cannot address portfolio allocation because a representative agent must hold the market portfolio for the market to clear. The most closely related work is Peress (2006), who studies how exogenous differences in portfolios affect information acquisition, Vayanos (2003), who studies the optimal organizational structure for transmitting investment information, and Brunnermeier, Gollier and Parker (2007) who consider the choice of optimism, which is the mean of a signal, in contrast to information allocation which is a choice of signal variance.

When portfolio choice is preceded by an information choice, the diversified portfolio prescribed by standard theory is no longer optimal. What form under-diversification takes depends on the convexity or concavity of the investor's constrained maximization problem in signal precision, which in turn depends on the investor's preferences and the information choice set. When the objective function is convex in signal precision, the investor learns only about one asset or risk factor. We call this *specialized learning*. When this objective is concave, an investor with sufficient capacity learns about multiple assets. We call this *generalized learning*. In both cases, the investor takes a larger position in the assets he learns about, on average, than standard portfolio theory dictates. Surprisingly, the convexity of the objective in signal precision is not determined by the investor's risk aversion. Instead, it is related to the investor's preference for early resolution of uncertainty. Information acquisition is a technology to resolve uncertainty earlier. It lets an investor learn sooner what he would otherwise observe after asset payoffs are realized. An investor may prefer early resolution of uncertainty either because he has non-expected utility or because when outcomes are known earlier, he can use that information to adjust his portfolio.

Exponential preferences that exhibit constant absolute risk aversion (CARA preferences) are commonly used in the literature on information acquisition and asymmetric information in finance. We show that such preferences deliver less of a preference for specialization than the constant relative risk aversion (CRRA) preferences used in much of the rest of the finance literature. Constant relative risk aversion means that absolute risk aversion fluctuates, depending on the investor's realized wealth. This fluctuation works to hedge the risk that specialized learning entails.

Section 1 constructs the investor's problem for a general class of preferences and learning technologies. It also describes two of the most commonly-used learning technologies. The *additive signal precision* technology is a continuous version of the following problem: For each asset, construct an urn filled with slips of paper, each with an independent noisy signal about the asset's payoff. Give the investor N draws and ask how many draws the investor would like to take from each urn. The *entropy* learning technology constrains the bits of information observed. In other words, the investor can download information about assets' payoffs for n seconds with a modem of 1 bit per second. How much information should he download about each asset?

Section 2 gives conditions under which the preferences and learning technology deliver generalized or specialized learning. To make the theoretical results more concrete, we work through four examples with specific utility functions and learning technologies. Example 1 shows that an investor with CARA preferences and an additive signal precision technology will generalize and learn about many assets. Example 2 shows that giving the CARA investor a preference for early

resolution of uncertainty makes specialization more desirable. Such an investor has mean-variance preferences, which are standard in the portfolio literature. Example 3 solves the problem with CRRA preferences. With the entropy technology, results mimic the mean-variance specialization results. Finally, example 4 returns to CARA expected utility with the entropy technology and proves that the investor becomes indifferent to any information allocation. The incentive to specialize and the desire to hold and learn about a diversified portfolio exactly offset each other. Thus, this model can rationalize almost any portfolio.

All the models dictate that an investor with no information should hold a diversified portfolio; our theory collapses to the standard portfolio choice model. As the precision of the investor's information increases, holding a perfectly diversified portfolio is still feasible, but no longer optimal. In each version of the model, the optimal portfolio has a diversified component that is the portfolio an investor without the capacity to learn would hold, plus a "learning portfolio" consisting of assets the investor learned about. Such a portfolio composition appears in the data (Polkovnichenko (2004)). With specialized learning, the learning portfolio is concentrated: it contains one asset when asset payoffs are independent or one risk factor when payoffs are correlated. With generalized learning, the investor learns about the riskiest assets. Except in a knife-edge case with symmetric assets, generalized learning also does not return the portfolio prescribed by standard theory.

While the merits of various preference specifications are well known, the choice of learning technology is less understood. Section 3 argues that the entropy technology has desirable properties. In particular, it is scale neutral; redefining what one share of an asset is does not change the feasible signal set. In addition, the entropy technology is more conducive to specialized learning. We cite evidence suggesting that specialized learning is responsible for investors' concentrated portfolios.

This paper is only one step in understanding the role of information choice in financial markets. Its framework can be used to understand the effect of information choice on asset prices. Doing so requires an equilibrium model. Section 2.5 discusses the results from such a model (which is solved in a separate technical appendix) and shows that a modified version of a CAPM emerges. Many questions about the efficient organization of the financial services and consulting industry could also be explored: What investments should analysts research? How should portfolio management services be priced? How many mutual funds should there be? What metrics reveal whether a fund manager is investing based on information or is earning higher returns by taking on high-risk investments? This paper provides a building block by exploring how information and investment choice can be jointly modeled, what various model features mean, and what predictions they deliver.

1 Setup

This is a static model which we break up into 3 periods. In period 1, the investor chooses the precision of signals about asset payoffs, subject to an increasing cost for more precise information. 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.

1.1 Defining Information Sets

The exogenous vector of unknown asset payoffs f is what the investor learns about. He is endowed with a prior belief that $f \sim N(\mu, \Sigma)$. 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 vector η about asset payoffs: $\eta = f + e_\eta$ where $e_\eta \sim N(0, \Sigma_\eta)$. The signal is noisy, but unbiased. At time 2, the investor combines his signal and his prior belief, using Bayes' law. Let $\hat{\mu}$ and $\hat{\Sigma}$ be the posterior mean and variance of payoffs, conditional on all information known to the investor in period 2:

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

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

$$\hat{\Sigma} \equiv V[f|\mu, \eta] = (\Sigma^{-1} + \Sigma_\eta^{-1})^{-1}. \quad (2)$$

We use $E_1[\cdot]$ and $V_1[\cdot]$ to denote the mean and variance conditional on prior beliefs alone, and we use $E_2[\cdot]$ and $V_2[\cdot]$ to denote the mean and variance conditional on information from priors and signals. That period-2 information is summarized by the moments $\hat{\mu}$, $\hat{\Sigma}$. Since the investor forms his portfolio after observing his signals, $\hat{\mu}$, $\hat{\Sigma}$ are the conditional mean and variance that govern the investor's portfolio choice. Likewise, we use U_1 and U_2 to denote expected utility, conditional on time-1 and time-2 information sets.

1.2 Preferences

We begin by considering a general form of utility function. It allows for both expected utility preferences and a preference for early/late resolution of uncertainty. In the next section, we will explore examples with specific function forms, such as CARA and CRRA utility. Investors have preferences over the wealth W they possess at time 3 that are described by continuous, twice

differentiable, increasing functions u_1 and u_2 :

$$U_1 = E_1 [u_1(E_2[u_2(W)])]. \quad (3)$$

1.3 Portfolio Allocation Choice

Given his posterior beliefs, the investor chooses the $N \times 1$ vector $q \equiv [q_1, \dots, q_N]'$ of quantities of each asset that he chooses to hold. The investor takes as given the risk-free return r and the $N \times 1$ vector of asset prices $p \equiv [p_1, \dots, p_N]'$. In making that choice, he is subject to a budget constraint

$$W = W_0 r + q'(f - pr). \quad (4)$$

Following Admati (1985), we call $f_i - p_i r$ asset i 's excess return and $q'(f - pr)$ the portfolio excess return.

Independent assets Without loss of generality, we consider independent assets: Σ is diagonal. For any set of correlated assets with full rank variance-covariance payoff matrix Σ , we can form principal components – linear combinations of these correlated assets such that the linear combinations are independent. Principal components are frequently used in the portfolio literature to represent risk factors such as business-cycle risk, industry-specific risk, and firm-specific risk (Ross (1976)). The solution to the correlated asset problem is then exactly the same, as long as we define q , p and f to be the quantity invested, price and payoff of the linear combinations of assets (risk factors). Investors learn about and invest in the risk factors, just as if they were independent underlying assets. Investing additional funds in one asset translates in to investing in one risk factor, which is a portfolio of assets with correlated payoffs. So while the solution is unchanged, its interpretation differs with correlated assets.

1.4 Information Allocation Choice

There are two aspects of information choice: how much information to acquire and how to allocate that information among assets. We call the quantity of information an investor observes his *capacity*, K . We take K as given and focus on the allocation problem.² The mapping between between signal

²By a simple duality argument, we know that for every choice problem with an information cost function $c(\cdot)$ such that $U = U_1 - c(K)$, there is an equivalent endowment of capacity K^* that delivers the same portfolio predictions. Therefore, we could assume that $c(\cdot)$ is increasing and is sufficiently convex to deliver an interior optimal level of K^* . This sufficient condition varies depending on the preferences. Then, we would take that optimal level of K^* as given for the remainder of our analysis.

precisions of N assets and K is what we call a *learning technology*.

Since every signal variance Σ_η has a unique posterior belief variance $\widehat{\Sigma}$ associated with it (see equation 2), we can economize on notation and optimize over posterior belief variance $\widehat{\Sigma}$ directly. The prior covariance matrix Σ is not random; it is given. The posterior (conditional) covariance matrix $\widehat{\Sigma}$, which measures investors' uncertainty about asset payoffs, is also not random; it is the choice variable that summarizes the investor's optimal information decision.³ Learning about asset i makes its conditional variance $\widehat{\Sigma}_{ii}$ lower than its unconditional variance Σ_{ii} . The learning technology is a continuous function κ that maps prior and posterior variances into a level of capacity:

$$\kappa(\Sigma, \widehat{\Sigma}) \leq K \tag{5}$$

Capacity allocation is not a sequential choice. An investor is not allowed to re-optimize his learning choice, based on signal realizations. Rather, we imagine that in period 1, the investor tells a computer what asset information to download. In period 2, he observes the computer output and decides how to invest.

Independent signals We make an assumption that simplifies the problem substantially: Investors cannot obtain signals with correlated information about asset (risk factor) payoffs that are independent. Such an assumption has a straightforward interpretation: Investors do not choose a correlation structure for posterior beliefs. Instead, they take the structure of risks in the world as given and decide how much to reduce each risk through information acquisition. An on-line technical appendix shows that the indifference result with CARA preferences and the corner solution with mean-variance preferences survive with correlated signals and correlated assets.⁴

A no-forgetting constraint The second constraint is that the variance of each signal must be non-negative. Without this constraint, the investor could choose to forget what he knows about one asset in order to obtain a more precise signal about another, without violating the capacity constraint. Using (2) and the independence of assets and signals, this implies that posterior variance can never exceed prior variance on any asset i

$$0 < \widehat{\Sigma}_{ii} \leq \Sigma_{ii}, \quad \text{or equivalently} \quad \infty > \widehat{\Sigma}_{ii}^{-1} \geq \Sigma_{ii}^{-1} \quad \forall i. \tag{6}$$

³The property that the posterior variance is a deterministic function of the prior variance and signal precision is a property of normal variables. This is not generally true for other probability distributions.

⁴The on-line appendix is posted at http://pages.stern.nyu.edu/~lveldkam/pdfs/portfolio_appdx.pdf .

1.5 The Investor's Problem

Given a chosen level of capacity K , a solution to the model is: A choice of $\hat{\Sigma}$, the variance of posterior beliefs, to maximize utility (3), subject to the information constraint (5), the no-forgetting constraint (6), and rational expectations about q ; given a signal η about asset payoffs f , posterior means $\hat{\mu}$ and variances $\hat{\Sigma}$ are formed according to Bayes' law (1 and 2); a choice of portfolio q to maximize expected utility, conditional on the signal realizations.

Figure 1 illustrates the sequence of events.




Figure 1: Sequence of events.

1.6 Two commonly-used learning technologies

In the results that follow, we will explore models of learning that employ one of two technologies.

Additive precision This learning technology is one where more information capacity means a higher *sum* of signal precisions.

$$K \geq \sum_{i=1}^N (\hat{\Sigma}_{ii}^{-1} - \Sigma_{ii}^{-1}) \quad (7)$$

With this constraint, learning takes the form of a sequence of independent draws. Each independent draw of a normally distributed signal η with mean f and variance σ^2 adds σ^{-2} to the precision of posterior beliefs. Therefore, constraining the sum of precisions is like constraining the number of independent signal draws.

Entropy The standard measure of information in information theory is entropy. It is frequently used in econometrics and statistics and has been used in economics to model limited information processing by individuals and to measure model uncertainty.⁵ Following Sims (2003), we model

⁵In econometrics, it is a log likelihood ratio. In statistics, it is a difference of the prior and posterior distributions' Kullback-Liebler distance from the truth. In robustness, it is interpreted as a reduction in measurement error (Cagetti

the amount of information transmitted as the reduction in entropy achieved by conditioning on that additional information. In other words, the amount of capacity \tilde{K} an investor is endowed with limits how much his signals η can reduce payoff uncertainty: $|\Sigma|/|\hat{\Sigma}| \leq \tilde{K}$. Because payoffs and signals are independent across assets, the determinants can be re-written as the product of the diagonal elements, $\tilde{K} \geq (\prod_{i=1}^N \hat{\Sigma}_{ii}^{-1})/(\prod_{i=1}^N \Sigma_{ii}^{-1})$. Thus, this learning technology is one where more information capacity means a higher *product* of signal precisions. It will become more convenient to work with an equivalent log form of this constraint where $K = \ln(\tilde{K})$:

$$K \geq \sum_{i=1}^N \ln(\hat{\Sigma}_{ii}^{-1} \Sigma_{ii}). \quad (8)$$

2 Results: Optimal Information Acquisition

If the investor's problem is concave in the precision of beliefs about all assets, he should learn about assets he initially knows less about. We call this "generalized learning." If the problem is convex in the precision of beliefs, a corner solution emerges and the investor learns about one asset. We call this "specialized learning." Finally, a knife-edge case arises when the objective and the constraint take the same form, leaving the investor indifferent between any capacity allocation. This section explores these three scenarios and offers examples of each.

The first order condition of the objective function with respect to q yields the optimal portfolio. Denote that optimal portfolio, which depends on the posterior mean and precision of asset payoffs, $q(\hat{\mu}, \hat{\Sigma}^{-1})$. Substituting the optimal portfolio into the budget constraint and then into the objective delivers the expected utility of having posterior precision $\hat{\Sigma}^{-1}$ and investing optimally. The period-1 Lagrangian problem is to choose $\{\hat{\Sigma}_{ii}^{-1}\}_{i=1}^N$ to maximize

$$\begin{aligned} \mathcal{L} \equiv & \int u_1 \left(\int u_2 \left(rW_0 + q(\hat{\mu}, \hat{\Sigma}^{-1})'(f - pr) \right) d\Phi_2(f|\hat{\mu}, \hat{\Sigma}) \right) d\Phi_1(\hat{\mu}|\mu, \Sigma - \hat{\Sigma}) \\ & + \psi \left(K - \kappa(\Sigma^{-1}, \hat{\Sigma}^{-1}) \right) + \sum_i \lambda_i (\hat{\Sigma}_{ii}^{-1} - \Sigma_{ii}^{-1}) \end{aligned} \quad (9)$$

where ψ and λ_i are Lagrange multipliers on the information capacity and no-forgetting constraints respectively, and Φ_1 and Φ_2 are normal cumulative density functions.

The key property of \mathcal{L} will be its convexity or concavity in $\hat{\Sigma}_{ii}^{-1}$, $\forall i$. This determines whether

et al. (2002)). It has been previously used in economics to model limited mental processing ability (Radner and Van Zandt (1999), Sims (2003), Moscarini (2004), and Mackowiak and Wiederholt (2008)) and in representative investor models in finance (Peng (2004)). See Sims (2003) for a discussion of the link between this information measure and the length of the binary code mentioned in the introduction.

investors want to become specialists (information problem has a corner solution) or generalists (information problem has an interior solution). In a multivariate optimization problem, a necessary condition for convexity of the objective is

$$\frac{\partial^2 \mathcal{L}}{\partial(\hat{\Sigma}_{ii}^{-1})^2} \geq 0 \quad \forall i. \quad (10)$$

The sufficient condition for convexity is that the Hessian matrix, containing the second derivatives and all the cross-partial derivatives, is positive semi-definite. This turns out to be a difficult condition to verify because with N assets, there are N sub-determinants of the Hessian that all must be non-negative. To make the problem tractable, each of our examples will transform the Lagrangian to make it additively separable in $\hat{\Sigma}_{ii}^{-1}$ for each i , while leaving its maximum the same. Additive separability implies that $\partial^2 \mathcal{L} / \partial \hat{\Sigma}_{ii}^{-1} \partial \hat{\Sigma}_{jj}^{-1} = 0$, $\forall i \neq j$, which makes the Hessian a diagonal matrix. When the Hessian is diagonal, (10) is a necessary and sufficient condition for convexity. Convexity of $-\mathcal{L}$ implies that the Lagrangian \mathcal{L} is concave.

Since $\hat{\Sigma}$ determines the distributions $\Phi_1(\hat{\mu}|\mu, \Sigma - \hat{\Sigma})$ and $\Phi_2(f|\hat{\mu}, \hat{\Sigma})$, we use the change of variables formula to replace $f = \mu + (\Sigma - \hat{\Sigma})^{1/2} z_1 + \hat{\Sigma}^{1/2} z_2$ and $\hat{\mu} = \mu + (\Sigma - \hat{\Sigma})^{1/2} z_1$, where z_1 and z_2 are standard normal variables. These, in turn enter (9) through wealth W (equation 4) and through the portfolio choice q . Then, the second partial derivative of the Lagrangian is

$$\frac{\partial^2 \mathcal{L}}{\partial(\hat{\Sigma}_{ii}^{-1})^2} = \int_{z_1} \frac{\partial^2 u_1}{\partial(\hat{\Sigma}_{ii}^{-1})^2} \partial \Phi(z_1) - \Psi \frac{\partial^2 \kappa}{\partial(\hat{\Sigma}_{ii}^{-1})^2} \quad (11)$$

where

$$\begin{aligned} \frac{\partial^2 u_1}{\partial(\hat{\Sigma}_{ii}^{-1})^2} &= \frac{\partial^2 u_1}{\partial U_2^2} \cdot \left(\int_{z_2} \frac{\partial u_2}{\partial W} \frac{\partial W}{\partial \hat{\Sigma}_{ii}^{-1}} d\Phi(z_2) \right)^2 \\ &+ \frac{\partial u_1}{\partial U_2} \cdot \left(\int_{z_2} \frac{\partial^2 u_2}{\partial W^2} \left[\frac{\partial W}{\partial \hat{\Sigma}_{ii}^{-1}} \right]^2 + \frac{\partial u_2}{\partial W} \frac{\partial^2 W}{\partial(\hat{\Sigma}_{ii}^{-1})^2} d\Phi(z_2) \right), \end{aligned} \quad (12)$$

where Φ without a subscript denotes a standard normal cumulative density function.

While this expression has many pieces, they are straightforward to interpret. The last term in (11) is the convexity or concavity of the learning technology in signal precision. The additive signal precision technology is linear and therefore has $\partial^2 \kappa / \partial(\hat{\Sigma}_{ii}^{-1})^2 = 0$. But the entropy learning technology is a sum of logs, which are concave. Since $\partial^2 \kappa / \partial(\hat{\Sigma}_{ii}^{-1})^2 < 0$, the second term in the Lagrangian is $-\psi \partial^2 \kappa / \partial(\hat{\Sigma}_{ii}^{-1})^2 > 0$, which makes the Lagrangian more convex. This means that

the entropy technology will be more conducive to specialization.

With expected utility, $u_1(x) = x$. In that case, $\partial^2 U_1 / \partial U_2^2 = 0$, the first line of (12) drops out, and $\partial u_1 / \partial U_2 = 1$ on the second line. If the investor is risk-averse, then $\partial^2 u_2 / \partial W^2 < 0$. Since the squared term must be non-negative, (12) is negative whenever $E \left[\frac{\partial u_2}{\partial W} \frac{\partial^2 W}{\partial (\Sigma_{ii}^{-1})^2} \right] < 0$. Furthermore, with CARA or mean-variance utility, expected profit and W will turn out to be linear in signal precision. Therefore, $E \left[\frac{\partial^2 W}{\partial (\Sigma_{ii}^{-1})^2} \right] = 0$, which means that the expectation of the product is the covariance, $cov \left(\frac{\partial u_2}{\partial W}, \frac{\partial^2 W}{\partial (\Sigma_{ii}^{-1})^2} \right)$. In other words, the concavity of the Lagrangian with expected utility hinges on how marginal utility covaries with the increasing returns to learning. The rationale for specialization relies on an investor with an unexpectedly high or low signal taking an aggressive portfolio position to exploit his information. From the perspective of the investor at time 1, this is a risky thing to do but also offers higher expected profits. The question then is, in states where the investor learns information that prompts him to take bigger risks with more return ($\frac{\partial^2 W}{\partial (\Sigma_{ii}^{-1})^2}$ is high), is his marginal value of the additional wealth ($\frac{\partial u_2}{\partial W}$) low? If so, then specialization is not attractive because it pays off in states when the investor values additional wealth least. Conversely, the first condition for convexity is met (11 is positive) for expected utility investors when marginal utility covaries sufficiently positively with the increasing returns to learning.

2.1 Interior solutions

When (9) is a concave function, there is some target precision for each asset that the investor would like to achieve $\underline{\sigma}_i^{-1}$. For some assets, the investor may already have more precise prior information than what the target prescribes $\Sigma_{ii}^{-1} > \underline{\sigma}_i^{-1}$. But the no-forgetting constraint prevents the investor from un-learning that information. For assets whose payoffs are uncertain relative to the target level of uncertainty, the investor learns more about these assets and ends up holding more of them than he would have without any capacity to learn. We call this *generalized learning*.

Because one asset could always be much more valuable to learn about than the others, it is not possible to prove generally that an investor will always learn about multiple assets. One special case where the investor learns about all assets is when all assets are symmetric. Although this case is clearly not realistic, it is instructive because it shows that the breadth of assets that investors learn about is related to the asymmetry in the mean and variance of the assets' payoffs. The following proposition shows that if expected profits and the learning technology are symmetric for all assets and the investor initially knows the same amount about every asset, then he will learn about every asset, no matter how much capacity he has to learn.

Proposition 1. *Suppose that all assets are symmetric, meaning that $\mu_i = \mu_j$, $\Sigma_{ii} = \Sigma_{jj}$, $p_i = p_j$ $\forall i, j$, and the learning technology treats all assets symmetrically, meaning that $\partial\kappa/\partial\widehat{\Sigma}_{ii} = \partial\kappa/\partial\widehat{\Sigma}_{jj}$ whenever $\widehat{\Sigma}_{ii} = \widehat{\Sigma}_{jj}$, $\forall i, j$. Then, if the Lagrangian \mathcal{L} is strictly **concave** in signal precision $\widehat{\Sigma}_{ii}^{-1}$, $\forall i$, and the investor values capacity ($\partial\mathcal{L}/\partial K > 0$), there exists a solution to the investor's problem such that the investor learns an equal amount about all assets: $\widehat{\Sigma}_{ii}^{-1} = \widehat{\Sigma}_{jj}^{-1}$, $\forall i, j$.*

This and all further propositions are proven in appendix A.

When an investor is faced with this type of learning problem, but assets are asymmetric, the nature of his optimal portfolio depends on how much capacity he has to learn. With a small amount of capacity, the investor would learn only about the asset whose signal precision has the highest marginal value. Given a larger amount of capacity, the investor will learn about the payoffs of a larger number of assets to avoid the diminishing value to learning solely about one or a small number of assets. If the indirect utility of learning about each asset goes to zero with the quantity of information, then with sufficient capacity, the investor learns about all assets' payoffs. How much capacity is sufficient depends on how similar the initial marginal benefits to learning about each asset are.

Example 1: CARA preferences and additive precision technology Let $u_1(x) = x$ and $u_2(W) = -\exp(-\rho W)$. This is CARA expected utility with an absolute risk-aversion coefficient ρ . Using the law of iterated expectations, it can be rewritten as

$$U_1 = -E_1[\exp(-\rho W)]. \quad (13)$$

Substituting the budget constraint (4) into the objective (13), the first order condition yields the optimal portfolio of risky assets.

$$q = \frac{1}{\rho} \widehat{\Sigma}^{-1}(\hat{\mu} - pr). \quad (14)$$

Any remaining initial wealth is invested in the risk-free asset. After substituting the optimal portfolio (14) into the budget constraint (4), substituting for wealth W in the objective (13) and taking the time-1 expectation of $\hat{\mu}$ (see appendix A.2 for details), expected utility at time 1 is

$$U_1 = -A \left(|\Sigma|/|\widehat{\Sigma}| \right)^{-1/2} \quad (15)$$

where $A \equiv \exp(-1/2(\mu - pr)\Sigma^{-1}(\mu - pr))$ is positive and exogenous. Maximizing this expression is

equivalent to maximizing $|\widehat{\Sigma}^{-1}|/|\Sigma^{-1}| = \prod_i \widehat{\Sigma}_{ii}^{-1} / \prod_i \Sigma_{ii}^{-1}$, which is also equivalent to maximizing its log $\sum_i \ln(\widehat{\Sigma}_{ii}^{-1}/\Sigma_{ii}^{-1})$, which is additively separable. Recall that the additive technology (7) constrains the sum $\sum_i \widehat{\Sigma}_{ii}^{-1}$. Maximizing a concave function, subject to a sum constraint, yields an interior solution.

Corollary 1. *The solution is: $\widehat{\Sigma}_{ii}^{-1} = \max(\underline{\sigma}^{-1}, \Sigma_{ii}^{-1}) \forall i$ where $\underline{\sigma}$ solves $\sum_{i=1}^N \max(\underline{\sigma}^{-1}, \Sigma_{ii}^{-1}) = K$.*

In other words, the investor learns most about assets he is most uncertain about. The number of assets learned about is weakly increasing in capacity K . With sufficient capacity, he would set the posterior uncertainty about all assets equal. Figure 2 illustrates the expected utility and the feasible and optimal learning choices in a 2-asset example. The linear capacity constraint (dash-dotted line) intersects the indifference curve (solid line) at an interior point. This reveals that utility is maximized when the investor learns about both assets. The investor learns comparatively more about asset 1 because he was initially more uncertain about its payoffs.

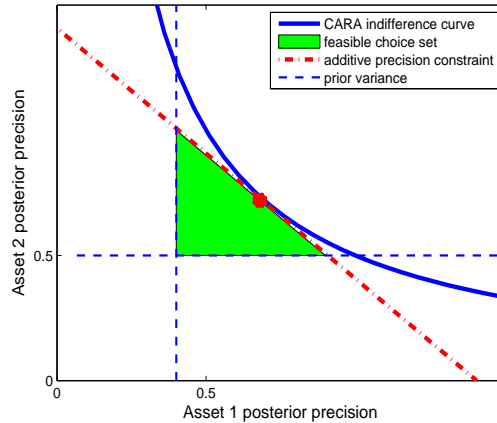


Figure 2: CARA Preferences and Additive Precision Technology: Indifference curve and information constraints. Shaded area represents feasible information choices, posterior variances that satisfy the capacity and no-negative learning constraints. The dot represents the optimal information choice. In this example, the investor is initially less uncertain about asset 2: The prior precision on asset 2 is 0.5, higher than the prior precision on asset 1 (0.4). The expected return on both assets is the same.

This is an example of a learning problem where the Lagrangian is concave. Recall that because this investor has expected utility, $\partial u_1 / \partial U_2 = 1$ and $\partial^2 u_1 / \partial U_2^2 = 0$, and because the learning technology is linear in $\widehat{\Sigma}_{ii}^{-1}$, $\partial^2 K / \partial (\widehat{\Sigma}_{ii}^{-1})^2 = 0$. Finally, $\partial^2 u_2 / \partial W^2 = -\rho^2 \exp(-\rho W) < 0$. Therefore, the reason that the sign of (11) is negative is that the covariance of the marginal utility of wealth

with the increasing returns to learning $cov\left(\frac{\partial u_2}{\partial W}, \frac{\partial^2 W}{\partial(\hat{\Sigma}_{ii}^{-1})^2}\right)$ is not too positive.⁶ In other words, the benefits from specialized learning do not accrue in states with high enough marginal utility to warrant the risk that specialized learning entails.

Portfolio with generalized learning Before we can assess how diversified portfolios are, we need a benchmark for comparison. The no-learning portfolio for CARA preferences, what standard theory would call a diversified portfolio, is

$$q^{div} = \frac{1}{\rho} \Sigma^{-1} (\mu - pr). \quad (16)$$

The optimal portfolio with learning is the sum of q^{div} and the component due to learning, q^{learn}

$$E[q^{learn}] = \frac{1}{\rho} (\hat{\Sigma}^{-1} - \Sigma^{-1}) (\mu - pr), \quad (17)$$

plus his position in the risk free asset.⁷ In the special case of perfect symmetry considered in proposition 1, the expected learning portfolio and the diversified portfolio are identical. From here on, we will use under-diversification to mean the difference between the optimal portfolio with learning q and q^{div} .

What does generalized learning with asymmetric assets imply for portfolio holdings? From (17), we see that learning more about asset i (higher $\hat{\Sigma}_{ii}^{-1}$) makes the investor expect to hold more of asset i (higher $E[q_i^{learn}]$), as long as the asset offers a positive prior expected return. Since generalized learning entails learning most about assets that have the most uncertain payoffs, such a learning strategy makes the investor hold more of these assets, on average, in his portfolio. Therefore, generalized learning tilts the optimal portfolio towards higher-risk assets.

The strategy of investing aggressively in assets whose payoffs are most uncertain resembles what an over-confident investor might do. But instead of irrationally under-estimating uncertainty, the generalist investor reduces uncertainty through learning. He uses his information to buy assets that are likely to have high payoffs and sell assets that are likely to have low payoffs: the time-1 covariance of q and f rises because learning makes $\hat{\mu}$ and f more correlated. Therefore, his portfolio

⁶The expectation of the product of the two terms in (12) is equal to their covariance because the product of their expectations is 0. This is because $E[\partial^2 W / \partial(\hat{\Sigma}_{ii}^{-1})^2] = 0$. To see this, note that $W = rW_0 + \frac{1}{\rho} (\hat{\mu} - pr)' \hat{\Sigma}^{-1} (f - pr)$. The period-1 expectation is $E[W] = rW_0 + \text{Trace}(\Sigma \hat{\Sigma}^{-1} - I) + (1/\rho) (\mu - pr)' \hat{\Sigma}^{-1} (\mu - pr)$. This is linear in the diagonal elements of $\hat{\Sigma}^{-1}$. Linearity ensures that the second derivative is zero.

⁷This expression does not have a $\hat{\mu}$ term in it because it is the *unconditional expectation* of the learning portfolio. The *realized* portfolio will depend on the observed signals and will therefore have a $\hat{\mu} - \mu$ term in it. Since beliefs are martingales, $E[\hat{\mu}] = \mu$ and the $\hat{\mu}$ term drops out in expectation.

returns should exceed those of the overconfident investor. The higher the investor's capacity, the more excess return he can earn through uncertainty reduction.

2.2 Corner solutions

When (9) is convex in $\widehat{\Sigma}_{ii}^{-1}$, there are increasing returns to learning about a single asset. As the following proposition shows, increasing returns induces the investor to specialize by learning about only one asset.

Proposition 2. *If the Lagrangian is increasing in capacity ($\partial\mathcal{L}/\partial K > 0$) and (weakly) **convex** in $\widehat{\Sigma}^{-1}$, then there exists a solution where the investor allocates all his capacity to one asset: $\widehat{\Sigma}_{ii}^{-1} > \Sigma_{ii}^{-1}$ for one asset j and $\widehat{\Sigma}_{jj}^{-1} = \Sigma_{jj}^{-1}$, $\forall j \neq i$.*

The following two examples show how either a preference for early resolution of uncertainty or CRRA preferences can generate a convex Lagrangian.

Example 2: Mean-variance preferences and entropy learning technology Let $u_1(x) = E_1[-\ln(-x)]$ and $u_2(W) = -E_2[\exp(-\rho W)]$. Then, using the fact that asset payoffs are normally distributed, we can rewrite investors' utility as a "mean-variance" utility function:

$$U_1 = E_1 \left[\rho E_2[W] - \frac{\rho^2}{2} V_2[W] \right]. \quad (18)$$

If there were no information choice, then (13) and (18) would be equivalent because u_2 is the same in both cases. The log operator in u_1 induces preference for early resolution of uncertainty, as in Epstein and Zin (1989). This objective function has been used in one-asset information choice problems by Wilson (1975) and Peress (2006).⁸ It can arise naturally in problems where a risk-neutral profit-maximizing portfolio manager invests on behalf of clients with CARA utility (13). Instead of having one investor who has different risk preferences in each period, this problem would introduce two agents with differing risk aversions, each with expected utility.⁹

⁸This preference for early resolution of uncertainty cannot be induced by adding a time-2 consumption decision alone, as in Spence and Zeckhauser (1972). There is a decision already being taken at time 2 that makes investors value information, just like the consumption choice would. Suppose investors can consume c_2 at the investment date and c_3 when asset payoffs are realized. If preferences are defined over $rc_2 + c_3$, where r is the rate of time preference, the solution will be identical. The earlier consumption choice simply takes the place of investing in the riskless asset.

⁹Suppose that a portfolio manager can extract all surplus from an investor. The investor would pay the manager an amount τ that would make him indifferent between a portfolio with payoff mean $\hat{\mu}$ and variance $\widehat{\Sigma}$ and a no-fee portfolio he would form based on his prior information: $-\exp(-(-\rho E[W|\hat{\mu}, \widehat{\Sigma}] - \tau + (\rho^2/2)V[W|\hat{\mu}, \widehat{\Sigma}])) = -\exp(-(-\rho E[W|\mu, \Sigma] + (\rho^2/2)V[W|\mu, \Sigma]))$. Since $E[W|\mu, \Sigma]$ and $V[W|\mu, \Sigma]$ are exogenous, maximizing the expected portfolio management fee entails maximizing the expected value of $\rho E[W|\hat{\mu}, \widehat{\Sigma}] - (\rho^2/2)V[W|\hat{\mu}, \widehat{\Sigma}]$, which is the same as (18).

Since u_2 is the same as in example 1, the optimal portfolio choice is also the same. Substituting (14) into the period 1 utility function yields $U_1 = E_1 \left[\frac{1}{2}(\hat{\mu} - pr)' \widehat{\Sigma}^{-1} (\hat{\mu} - pr) \right] + rW_0$, where posterior means $\hat{\mu}$ are normally distributed. Taking the expectation of this non-central χ^2 -distributed random variable, the time-1 problem is to choose signals to maximize

$$U_1 = \frac{1}{2} \text{Trace}(\widehat{\Sigma}^{-1} V_1 [\hat{\mu} - pr]) + \frac{1}{2} E_1 [\hat{\mu} - pr]' \widehat{\Sigma}^{-1} E_1 [\hat{\mu} - pr] + rW_0, \quad (19)$$

where $\text{Trace}(\cdot)$ stands for the trace of a matrix. The only unknown variable at time 1 is $\hat{\mu}$. Therefore, $E_1[\hat{\mu} - pr] = \mu - pr$ and $V_1[\hat{\mu} - pr|\mu] = \Sigma - \widehat{\Sigma}$. Because the trace of a matrix is the sum of its diagonal elements, we can rewrite the objective as

$$U_1 = \frac{1}{2} \left\{ -N + \sum_{i=1}^N \widehat{\Sigma}_{ii}^{-1} \Sigma_{ii} (1 + \theta_i^2) \right\} + rW_0, \quad (20)$$

where $\theta_i^2 \equiv \frac{(\mu_i - p_i r)^2}{\Sigma_{ii}}$ is the squared excess return per unit of prior variance, or the prior *squared Sharpe ratio* of asset i . The higher this Sharpe ratio, the more valuable it is to learn about asset i .

The objective (20) is a weighted sum of posterior precisions, plus constants that do not affect the investor's choice. The entropy-based capacity constraint (8) bounds the product of those precisions. Maximizing a sum subject to a product constraint yields a corner solution. Choosing a high precision of the signal with the highest linear weight θ_i^2 and keeping the others as low as possible, maximizes the sum, while keeping the product low. While proposition 2 established that the investor should learn about only one asset, the following corollary proves that the asset learned about is the one with the highest squared Sharpe ratio.

Corollary 2. *The optimal information acquisition strategy uses all capacity to learn about the asset with the highest squared Sharpe ratio: $\theta_i^2 = (\mu_i - p_i r)^2 \Sigma_{ii}^{-1}$.*

This result is illustrated in figure 3. While the objective is linear in precision, the learning technology constraint is convex. As before, asset 2 has a higher prior precision, but the same expected return as asset 1. So its squared Sharpe ratio is higher. Therefore, the optimal choice is to devote all learning capacity to asset 2. Recall that in example 1, the investor learned about both assets.

Returning to the Lagrangian second order condition, note that since U_1 is a weighted sum of signal precisions, $\partial^2 U_1 / \partial (\widehat{\Sigma}_{ii}^{-1})^2 = 0$. For the entropy constraint, $\partial^2 K / \partial (\widehat{\Sigma}_{ii}^{-1})^2 = -\widehat{\Sigma}_{ii}^2 > 0$. Since this term enters (11) negatively, it ensures that $\partial^2 \mathcal{L} / \partial (\widehat{\Sigma}_{ii}^{-1})^2 > 0, \forall i$. Recall that when the

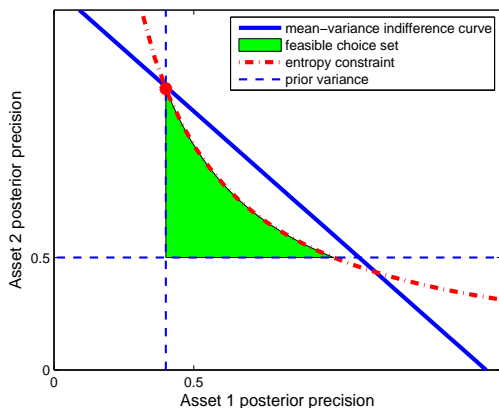


Figure 3: Mean-variance Preferences and Entropy Technology: Indifference curve and information constraints. Shaded area represents feasible information choices, posterior variances that satisfy the capacity and no-negative learning constraints. The dot represents the optimal information choice. In this example, the investor is initially less uncertain about asset 2: The prior precision on asset 2 is 0.5, higher than the prior precision on asset 1 (0.4). The expected return on both assets is the same.

objective and constraint are both additively separable, $\partial^2 \mathcal{L} / \partial (\hat{\Sigma}_{ii}^{-1})^2 > 0$, $\forall i$ is a sufficient condition for convexity. Therefore, the mean-variance problem with the entropy learning technology generates specialization because its Lagrangian is convex in signal precision.

Example 3: CRRA preferences and entropy learning technology Consider a static version of Merton (1987), the canonical model in the continuous-time portfolio choice literature. Let $u_1(x) = x$ and $u_2(W) = \frac{1}{1-\gamma} (W)^{1-\gamma}$, where $\gamma > 1$ is constant relative risk aversion.¹⁰

$$U_1 = E_1 \left[\frac{W^{1-\gamma}}{1-\gamma} \right] \quad (21)$$

Such a model entails a minor change to the payoff distribution set out in section 1 to keep it tractable. In continuous time, an $N \times 1$ vector of asset values S_t has a stochastic process given by $dS_t = \text{diag}(S_t)(\mu dt + \Sigma^{1/2} dZ_t)$, where r is the risk-free rate, μ is the expected rate of drift for the asset payoff, Σ is the variance-covariance matrix of payoffs and Z_t is a vector of standard Brownian motions. Therefore, an investor who holds a portfolio q of these risky assets faces a stochastic process for wealth: $dW_t = W_t [(r + q'(\mu - r))dt + q'\Sigma^{1/2}dZ_t]$. Because we analyze a static model, we approximate this continuous-time process with a discrete time process that has the same mean

¹⁰The on-line technical appendix derives similar results for the case when $\gamma < 1$ and for log preferences ($\gamma \rightarrow 1$).

and variance as the continuous process.

$$W_{t+1} = W_t \exp\left\{r + q'(\mu - r) - \frac{1}{2}q'\Sigma q + q'\Sigma^{1/2}z_t\right\} \quad \text{where} \quad z_t = Z_{t+1} - Z_t \sim N(0, I_n). \quad (22)$$

As before, we assume that asset values and signals are independent across assets. Each signal is normally distributed with mean z_t . Agents update using Bayes' law and form posterior means $\hat{\mu}$ and variances $\hat{\Sigma}$ as in (1) and (2). As before, we model the information choice as a choice of posterior variance $\hat{\Sigma}$. Information choice cannot affect the unconditional drift of the asset price. Rather, the *conditional* drift $\hat{\mu}$ summarizes the expected changes in the asset price between t and $t+1$. Information is moving some changes from the unexpected component (embedded in $\Sigma^{1/2}z_{t+1}$) to the expected component (embedded in $\hat{\mu}$). Thus, conditional on the time-2 information set, the process for wealth is

$$W = W_0 \exp\left\{r + q'(\hat{\mu} - r) - \frac{1}{2}q'\hat{\Sigma}q + q'\hat{\Sigma}^{1/2}\tilde{z}\right\} \quad \text{where} \quad \tilde{z} \sim N(0, I_n). \quad (23)$$

The first-order condition of the period-2 problem characterizes the optimal portfolio, as a share of the investor's wealth:

$$q = \frac{1}{\gamma}\hat{\Sigma}^{-1}(\hat{\mu} - r). \quad (24)$$

Substituting q into (22) and (21) yields expected utility conditional on information $\hat{\Sigma}$ and $\hat{\mu}$: $U_2 = W_0^{1-\gamma}/(1-\gamma)e^{(1-\gamma)r} \exp\left\{(1-\gamma)/(2\gamma)(\hat{\mu} - r)'\hat{\Sigma}^{-1}(\hat{\mu} - r)\right\}$.

This is the payoff of the first stage decision problem, if it results in beliefs $\hat{\mu}$ and $\hat{\Sigma}$. But each choice of $\hat{\Sigma}$ results in a random $\hat{\mu} \sim N(\mu, \Sigma - \hat{\Sigma})$. Taking the time-1 expectation over $\hat{\mu}$ delivers the objective function an investor uses to make his information choice.¹¹ Working with the log of this objective simplifies the calculations. To avoid taking the log of a negative number, we maximize $-\log(-U)$, collecting the constants in a new term (a). Since all the matrices in the problem are diagonal, the objective can be expressed as sums of the matrix diagonal

$$U_1 = a + \frac{1}{2} \sum_{i=1}^N \log\left(1 + (\gamma - 1)\Sigma_{ii}\hat{\Sigma}_{ii}^{-1}\right) + \left(\frac{\gamma - 1}{2}\right) \sum_{i=1}^N \frac{(\mu_i - r)^2}{\hat{\Sigma}_{ii} + (\gamma - 1)\Sigma_{ii}}. \quad (25)$$

As figure 4 illustrates, utility is no longer linear in each $\hat{\Sigma}_{ii}^{-1}$, as it was with the mean-variance objective. With an entropy-based learning technology, there are increasing returns to devoting

¹¹See appendix A.5 for details of these calculations.

additional capacity to learning about a given asset. Furthermore, the asset that is most valuable to learn about is one with a high expected return and low initial uncertainty.

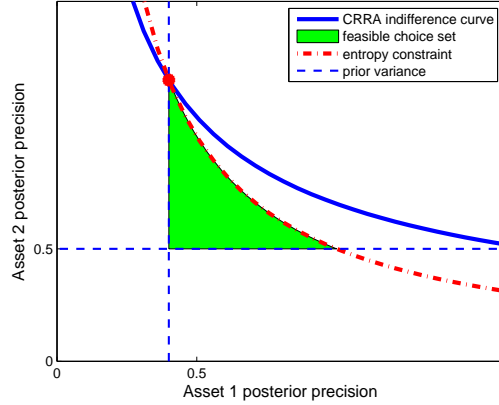


Figure 4: CRRA Preferences and Entropy Technology: Indifference curve and information constraints. Shaded area represents feasible information choices, posterior variances that satisfy the capacity and no-negative learning constraints. The dot represents the optimal information choice. In this example, the investor is initially less uncertain about asset 2: The prior precision on asset 2 is 0.5, higher than the prior precision on asset 1 (0.4). The expected return on both assets is the same.

Corollary 3. *The optimal information acquisition strategy uses all capacity to learn about one asset, the asset with highest squared Sharpe ratio $(\mu_i - r)^2 / \Sigma_{ii}$.*

The proof shows that Lagrangian problem is convex in signal precision. This is not obvious because (25) is concave in $\widehat{\Sigma}_{ii}^{-1}$. In principle, a concave objective minus a concave constraint could yield either a concave or convex Lagrangian function. But changing the choice variable to $K_i = \ln(\widehat{\Sigma}_{ii}^{-1} \Sigma_{ii})$ leaves the constraint linear in K_i and the objective convex in K_i . Since the objective is additively separable in K_i , $\partial^2 \mathcal{L} / \partial K_i^2 > 0$, $\forall i$ is sufficient for convexity of \mathcal{L} .

One source of convexity is the fact that expected wealth is convex in signal precision. In the CARA model with normally distributed shocks, expected wealth was linear in precision, which meant that $E \left[\frac{\partial u_2}{\partial W} \frac{\partial^2 W}{\partial (\widehat{\Sigma}_{ii}^{-1})^2} \right]$ in (12) was just the covariance of the two terms $\frac{\partial u_2}{\partial W}$ and $\frac{\partial^2 W}{\partial (\widehat{\Sigma}_{ii}^{-1})^2}$. Now, with CRRA preferences and log-normal innovations, expected wealth is $E_2[W] = W_0 \exp[r + (1/\gamma)(\widehat{\mu} - r)' \widehat{\Sigma}^{-1} (\widehat{\mu} - r)]$, an increasing exponential function of $\widehat{\Sigma}_{ii}^{-1}$, which is convex. This means that $E \left[\frac{\partial u_2}{\partial W} \frac{\partial^2 W}{\partial (\widehat{\Sigma}_{ii}^{-1})^2} \right]$ is the same covariance plus a positive number. The additional positive term in the second derivative makes the Lagrangian more convex.

Portfolio with specialized learning With mean-variance preferences and entropy learning technology, the investor specializes in learning about one highest-information-value asset for all levels of capacity. 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 $E[q_i^{learn}] = \frac{1}{\rho \Sigma_{ii}}(\mu_i - p_i r)(K - 1)$. The optimal portfolios with CRRA preferences have the same character. They are a sum of a diversified component, $q^{div} = (1/\gamma)\Sigma^{-1}(\mu - r)$ and a component that weights on the assets learned about $E[q_i^{learn}] = (1/\gamma)(\widehat{\Sigma}_{ii}^{-1} - \Sigma_{ii}^{-1})(\mu - r)$. Investors specialize and q^{learn} is comprised of a single asset. In both cases, since the learning portfolio is always comprised of one asset, and more capacity causes the investor to hold more of that asset, more capacity optimally results in a less diversified portfolio.

An investor engaged in generalized learning might also learn only about one asset, if assets are asymmetric and his capacity is sufficiently low (section 2.1). The resulting concentration of assets in the learning portfolio is similar to that in specialized portfolios. The difference is that the specialized investor would hold more of an asset that is more familiar to him, while the generalized-learning investor would hold more of an asset that he was initially more uncertain about.

The feedback between learning and investing The intuition for specialization laid out in the introduction was that investors want to learn about assets they expect to hold a lot of and then expect to hold more of assets that they choose to learn about. In every one of the examples explored so far, the second part of this argument holds true: The expression for the optimal portfolio q is multiplied by $\widehat{\Sigma}^{-1}$. This means that investors take larger positions in assets they choose to learn about. As long as the asset has a positive expected excess return ($\mu - pr > 0$ or $\mu - r > 0$ for CRRA), learning more about an asset makes an investor expect to hold more of it ($\partial E[q_i]/\partial \widehat{\Sigma}_{ii}^{-1} > 0$).

The first part of the feedback effect is that expecting to hold more of an asset should make it more valuable to learn about: $E[q_i]$ should enter positively in the marginal value of signal precision $\partial \mathcal{L}/\partial \widehat{\Sigma}_{ii}^{-1}$. In examples 2 and 3, the marginal value of learning about an asset is indexed by its squared Sharpe ratio. Another way to express the same quantity is as the product of the expected excess return, risk aversion and the expected quantity invested by an investor with zero capacity. This zero-capacity portfolio is q^{div} . For the mean-variance investor, $\theta_i^2 = (\mu_i - p_i r)\rho E[q^{div}]$. With CRRA, this squared Sharpe ratio is $\theta_i^2 = (\mu_i - r)\gamma E[q^{div}]$. Each investor wants to learn about an asset that has (i) high expected excess returns ($\mu_i - p_i r$), and (ii) features prominently in the portfolio he expects to hold, before accounting for new information.

This feedback effect generates convexity in the Lagrangian problem because if a higher expected

investment in asset i increases the marginal value of learning about i , then $\partial^2 \mathcal{L} / \partial \widehat{\Sigma}_{ii}^{-1} \partial E[q] > 0$. If learning more about an asset increases the expected investment, then $\partial E[q] / \partial \widehat{\Sigma}_{ii}^{-1} > 0$. Using the chain rule, this implies that $\partial^2 \mathcal{L} / \partial (\widehat{\Sigma}_{ii}^{-1})^2 > 0$, which is convexity in the investor's objective.

In example 1, the feedback effect breaks down. In that setting, the additive precision learning technology makes low $E[q]$ assets more valuable to learn about. Since $\partial E[q] / \partial \widehat{\Sigma}_{ii}^{-1} > 0$, low $E[q]$ assets are ones with low information precision. Adding a marginal increase to the precision of information about such an asset results in a proportionately larger increase in precision and a proportionately larger increase in utility. That makes the high-uncertainty, low $E[q]$ assets more valuable to learn about ($\partial^2 \mathcal{L} / \partial \widehat{\Sigma}_{ii}^{-1} \partial E[q] < 0$). By the same chain rule argument, this implies that $\partial^2 \mathcal{L} / \partial (\widehat{\Sigma}_{ii}^{-1})^2 < 0$, which makes the Lagrangian concave.

2.3 Indifference results

Indifference arises when the investor's expected utility (first term in 9) collapses to a function of $\kappa(\Sigma, \widehat{\Sigma})$. In other words, if every $\widehat{\Sigma}$ term can be substituted out by K , then every allocation of information precision that causes the information constraint to bind delivers equal expected utility. Expected utility may still be increasing in K , so the investor prefers more information to less, but he is indifferent about its allocation among assets.

While the possibility that the indirect utility takes exactly the same form as the constraint is a knife-edge case, there is an important class of problems that has exactly this property.

Example 5: CARA preferences and entropy learning technology Equation (15) in example 1 shows that the CARA utility objective U_1 takes the form of a square root of $|\Sigma|/|\widehat{\Sigma}|$. Recall also that the entropy constraint is $|\Sigma|/|\widehat{\Sigma}| = K$. The solution to the information allocation problem is therefore indeterminate because the time-1 expected utility only depends on capacity K , not on how that capacity is allocated across assets. Figure 5 shows how the indifference curve lies on top of the entropy constraint in a 2-asset example. In other words, entropy is the *exponential utility neutral* learning technology.

Proposition 3. *Given the entropy capacity constraint in (8), an investor with CARA utility (13) is indifferent between any allocation of his capacity.*

Indifference can arise with either concave or a convex utility. In this case, utility is concave in $\widehat{\Sigma}_{ii}^{-1}$ (see example 1). The interior solution associated with a concave problem could re-emerge if we "perturbed" the learning technology to make some risks slightly easier to learn about than others.

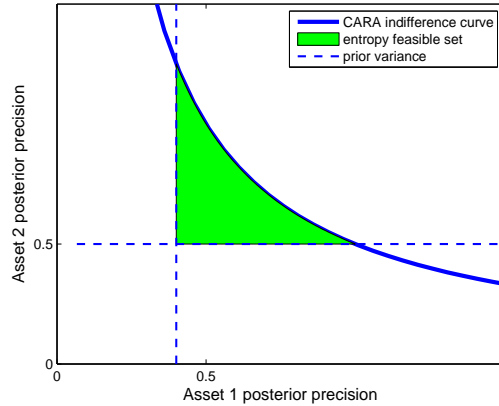


Figure 5: CARA Preferences and Entropy Technology: Indifference curve and information constraints. Shaded area represents feasible information choices, posterior variances that satisfy the capacity and no-negative learning constraints. Any capacity allocation along the north-east frontier of the feasible set achieves the highest possible expected utility. Every one of these allocations is a solution to the investor’s information choice problem.

For example, if the technology bounded the product of precisions each raised to the exponent ϵ_i , arbitrarily close to one ($K = \prod_i \hat{\Sigma}_{ii}^{-\epsilon_i}$), the investor would no longer be indifferent.

2.4 Comparing CARA to CRRA and mean-variance utility

First, why is the mean-variance investor not indifferent like the CARA investor, when faced with the same entropy learning technology? One way to understand the difference between the CARA and mean-variance problems is by seeing learning as a tool to resolve uncertainty sooner. Learning reduces $\hat{\Sigma}$, which is the residual risk the investor faces between time 2 and time 3. Since all risk is resolved at time 3, when asset payoffs are revealed, the only role for learning is to resolve that risk earlier (at time 2). Because the mean-variance investor has a preference for early resolution of uncertainty, he is less averse to the risk of learning useless information (a risk that will be resolved at time 2) than he is to the risk of earning low payoffs (a risk that will be resolved at time 3). Specialization appeals to this investor because it makes him more uncertain about what he will learn; but at time 2, he is less uncertain about the payoffs of his optimal portfolio. The latter is what he values most. In contrast, the CARA (expected utility) investor is just as averse to both risks. Therefore, he does *not* prefer a learning strategy that reduces portfolio payoff uncertainty by engaging in a more risky learning strategy.

Next, we compare the CARA to the CRRA investor. The CRRA investor has absolute risk

aversion that changes, depending on his expected level of wealth. These changes in his aversion to fixed gambles hedge some of the risk that is borne between time-1 and time-2. When the investor gets information that he cannot profit much from (e.g. he learns that the expected excess return is near zero), he does not deviate much from the diversified portfolio. Because his expected wealth is lower in these states, his aversion to large gambles is high. The extra utility the investor gets from taking a conservative portfolio position when his risk aversion is high offsets some of the loss of expected return. Conversely, when a CRRA investor gets a signal that he can profit from greatly, he takes a large position in an asset, making his portfolio riskier. The fact that the investor's expected wealth is high in these states and therefore his aversion to fixed-size gambles is lower helps to hedge his risk. It is not the case that the CRRA investor is less risk averse than the CARA investor. These results hold even when absolute risk aversion $\rho > 0$ is less than the absolute risk aversion implied by $\gamma > 1$.

This logic rests on changes in risk aversion, which show up in the second derivative of the Lagrangian (second line of 12), in the term $\frac{\partial^2 u_2}{\partial W^2} \left[\frac{\partial W}{\partial \hat{\Sigma}_{ii}^{-1}} \right]^2$. When the investor observes a particularly valuable signal, additional precision contributes greatly to the expected portfolio return: $\frac{\partial W}{\partial \hat{\Sigma}_{ii}^{-1}}$ is high. At the same time, a higher expected wealth makes absolute risk aversion fall, which makes $\frac{\partial^2 u_2}{\partial W^2}$ less negative. In other words, the two terms rise and fall together.¹² Positive covariance makes the expected value of their product of the two terms less negative, and makes the second derivative of the Lagrangian (11) more positive for CRRA preferences than it is for CARA preferences.

2.5 Equilibrium Effects

An important question is whether equilibrium effects might undo under-diversification. Specifically, if all investors want to specialize in learning about the same asset, they cannot then all hold more of that asset because the market has to clear. While a systematic investigation of equilibrium models is beyond the scope of this paper, the on-line technical appendix studies the model in example 2, with mean-variance preferences and entropy learning. It shows that specialization persists and investors' equilibrium portfolios are not diversified. The reason is that there is *strategic substitutability* in information choice. With sufficiently high capacity, investors choose to learn about different assets in equilibrium. This leads them to hold different, under-diversified portfolios. The intuition for this is that investors want to buy assets that others do not have high demand for because such assets

¹²To see the positive covariance formally, note that both terms are increasing in wealth $\partial^3 u_2 / \partial W^3 = (1 + \gamma) / \gamma W_{t+1}^{-2-\gamma} > 0$ and $\partial W / \partial \hat{\Sigma}_{ii}^{-1} = W_0 (\hat{\mu}_i - r)^2 \exp(r + (\hat{\mu} - r)' \hat{\Sigma}^{-1} (\hat{\mu} - r))$, which is a positive constant times expected wealth.

are less expensive. If an investor wants to buy assets that others are not buying, then they should learn information that will cause their asset demand to be different. That induces them to learn information that other investors do not know. Hence, specialization survives.

3 Comparing Learning Technologies

Scale neutrality A desirable property for a learning technology is that the definition of what constitutes one share of an asset should not change the feasible set of signals. The entropy-based learning technology exhibits this property, while the additive technology does not.¹³

Consider the following example: Take a share of an asset with payoff $f \sim N(\mu, \sigma)$ and split it into 2 shares of a new asset, each with payoff $f/2$. The new asset has payoffs with $1/2$ the standard deviation and $1/4$ the variance. The prior precision of information about its payoff is $4\sigma^{-1}$.

The additive learning technology allows the investor to add K units of precision to his information. If he adds K units of precision to his information about the new asset, the new precision is $4\sigma^{-1} + K$. This implies that the old asset has payoff precision $\sigma^{-1} + K/4$, after the investor learns. If the investor added K units of precision to his information about the old asset, the posterior precision would be $\sigma^{-1} + K$. The two posterior precisions are different. Thus, changing what constitutes a share of an asset changes the precision of the information an investor can acquire.

In contrast, the entropy learning technology allows the investor to multiply the precision of that information by K . Increasing the precision of information about the new asset by a factor of K will result in the precision of information about the old asset also being K times higher. This is the scale neutrality of entropy.

The entropy technology in (8) is not the only scale neutral learning technology. Consider a version of entropy-based learning where a fraction α of risk is unlearnable: $\tilde{K} = |\Sigma - \alpha\Sigma|/|\hat{\Sigma} - \alpha\Sigma|$. Resolving all learnable risk ($\hat{\Sigma}_{ii} = \alpha\Sigma_{ii}$) requires infinite capacity. Adding un-learnable risk is a way of generating decreasing returns to learning, which has intuitive appeal. It is scale neutral because dividing all the variances by any positive scalar x results in the same constraint.

Using Data to Select a Model of Information Acquisition Recent empirical research supports the prediction that information is responsible for investors' concentrated portfolios.¹⁴ Some

¹³Thanks to Michael Woodford for suggesting this argument.

¹⁴Using Swedish data on investors' personal characteristics and complete wealth portfolio, Massa and Simonov (2005) find that portfolio concentration does not hedge labor income risk. Guiso and Jappelli (2006) examine survey data and find that those who spend more time on information collection invest more in individual stocks and relatively less in diversified mutual funds. Finally, Polkovnichenko (2004) shows that directly-held equities are only 40% of the

of these findings can distinguish specialized from generalized learning.

First, specialized investors concentrate their portfolios more when their capacity increases, while generalists do the opposite. We cannot observe capacity in the data, but one good proxy for it is portfolio returns. Better-informed investors earn higher portfolio returns because they buy more assets that are likely to have high returns. The data reveal a positive correlation between expected returns and portfolio concentration, suggesting that investors specialize. Ivkovic, Sialm and Weisbenner (2005) find that concentrated investors outperform diversified ones by as much as 3% per year.¹⁵ Kacperczyk, Sialm and Zheng (2005) show that funds with above-median industry concentration yield an average return that is 1.1% per year higher than those with below-median concentration.

Second, investors who specialize hold more assets initially familiar to them, on average. Such a theory can simultaneously explain home bias, local bias and the tilt in investors' portfolios toward the sector they work in (see Van Nieuwerburgh and Veldkamp (2005) also for a quantitative evaluation of the theory). In contrast, generalized learning dictates that investors learn about unfamiliar assets. That would undo initial information advantages and reduce portfolio bias imparted by differences in initial information. The fact that US investment managers research mostly US assets and EU investment managers research mostly EU assets suggests that financial research is building on initial information advantages, not diversifying away from them. This supports the notion that financial research is more like entropy-based learning, a process of refining one's search, than additive precision learning, a process of independent sampling.

4 Conclusion

The assumption in most portfolio models, that investors cannot acquire information before investing, is not innocuous. When investors can choose what asset payoffs to learn about, the most basic prediction of portfolio theory, that investors should diversify, is overturned. This paper shows how to solve jointly for optimal information acquisition and portfolio allocation. The main message is that when investors can choose what information to acquire before they invest, they may invest in portfolios that would be sub-optimal for an investor who has not learned. From the point of view of standard portfolio theory, these portfolio might be deemed anomalous or irrational.

A key ingredient in the analysis is the convexity of the objective function in signal precision,

median household's portfolio; the remaining 60% is in more diversified assets. This is consistent with the theory's prediction that the optimal portfolio has a diversified component and a learning component.

¹⁵See also Coval and Moskowitz (2001), Massa and Simonov (2005) and Ivkovic and Weisbenner (2005).

relative to the convexity of the information constraint. With CARA preferences and the entropy-based information constraint, the objective and the constraint are proportional. That means that any information acquisition strategy can be justified and nearly any portfolio can be rationalized. But this is a knife-edge case. In other settings, with either a preference for early resolution of uncertainty or an ability to hedge learning risk, there are endogenous increasing returns to information that make the objective more convex. This induces the investor to specialize in assets he is initially well-informed about. Investors with more capacity learn more and therefore hold more of the asset they learn about. Other combinations of preferences and technologies generate decreasing returns. In these settings, the investor broadens his knowledge by acquiring information on assets he is most uncertain about. Such an investor holds a portfolio that deviates substantially from the diversified portfolio, by over-weighting high-risk assets.

One of the objections to an information-based theory of portfolio choice is that information is unobservable and therefore the model is not testable. A theory of information choice circumvents this problem by predicting investors' information sets, and therefore their portfolio holdings, on the basis of observable features of assets. It links observable variables to observable portfolios. Although we can't know what investors have learned directly, we can know which assets they would want to learn more or less about. Thus, each of these models presented offers testable predictions that can be compared with data to determine the most relevant model. Armed with the empirically-relevant model of financial information acquisition, we can better understand the financial services industry. For example, interpreted as a fund manager's research allocation problem, the model would predict the patterns of manager expertise studied by Dasgupta and Prat (2006) and Koijen (2007).

Another 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?" With entropy-based information capacity, managers maximize mean-variance-based profit by each specializing in one risk factor. Whether an investor's portfolio will also be concentrated hinges on how portfolio managers set fees. We conjecture that in a competitive equilibrium, fund managers offer quantity discounts, to induce more investment in their fund. Additional investment reduces the fund manager's per-share cost and allows him to compete linear-fee suppliers out of the market. Quantity discounts make investing in many funds costly. Competitive pricing of portfolio management services forces investors to internalize increasing returns to specialization; optimal under-diversification reappears. In future research on portfolio delegation, we plan to formalize this conjecture. The individual's information choice problem is a necessary first step in this research agenda.

References

- Admati, Anat**, “A Noisy Rational Expectations Equilibrium for Multi-Asset Securities Markets,” *Econometrica*, 1985, *53*(3), 629–57.
- Albuquerque, Rui, Greg Bauer, and Martin Schneider**, “International Equity Flows and Returns: A Quantitative Equilibrium Approach,” 2005. Working Paper.
- Banerjee, Snehal**, “Learning from Prices and Dispersion in Beliefs,” 2007. Northwestern University working paper.
- Barlevy, Gadi and Pietro Veronesi**, “Information Acquisition in Financial Markets,” *Review of Economic Studies*, 2000, *67*, 79–90.
- Bernhardt, Dan and Bart Taub**, “Strategic Information Flows In Stock Markets,” 2005. Working Paper.
- Biais, Bruno, Peter Bossaerts, and Chester Spatt**, “Equilibrium Asset Pricing Under Heterogenous Information,” 2004. Working Paper.
- Brunnermeier, Markus**, *Asset Pricing under Asymmetric Information: Bubbles, Crashes, Technical Analysis and Herding*, first ed., Oxford University Press, 2001.
- , **Christian Gollier, and Jonathan Parker**, “Optimal Beliefs, Asset Prices and the Preferences for Skewed Returns,” *American Economic Review, Papers and Proceedings*, 2007, *97*(2), 159–165.
- Bullard, James, George Evans, and Seppo Honkapohja**, “Near-Rational Exuberance,” 2005. Working Paper.
- Cagetti, Marco, Lars Hansen, Thomas Sargent, and Noah Williams**, “Robustness and Pricing with Uncertain Growth,” *Review of Financial Studies*, 2002, *15*(2), 363–404.
- Coval, Joshua and Tobias Moskowitz**, “The Geography of Investment: Informed Trading and Asset Prices,” *Journal of Political Economy*, 2001, *109*(4), 811–41.
- Dasgupta, Amil and Andrea Prat**, “Financial Market Equilibrium with Career Concerns,” *Theoretical Economics*, 2006, *1*(1), 67–93.
- 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.
- Grossman, Sanford and Joseph Stiglitz**, “On the Impossibility of Informationally Efficient Markets,” *American Economic Review*, 1980, *70*(3), 393–408.
- Guiso, Luigi and Tullio Jappelli**, “Information Acquisition and Portfolio Performance,” 2006. CEPR Working Paper #5901.
- 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, *60(4)*, 1983–2011.
- Kodres, Laura and Matthew Pritsker**, “A Rational Expectations Model of Financial Contagion,” *Journal of Finance*, 2002, *57*, 769–99.
- Koijen, Ralph**, “The Cross-section of Managerial Ability and Risk Preferences,” 2007. Working Paper.
- Mackowiak, Bartosz and Mirko Wiederholt**, “Optimal Sticky Prices Under Rational Inattention,” *American Economic Review*, 2008, *forthcoming*.
- Maenhout, Pascal**, “Robust Portfolio Rules and Asset Pricing,” *The Review of Financial Studies*, 2004, *17(4)*, 951–983.
- 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.
- Moscarini, Giuseppe**, “Limited Information Capacity As A Source of Inertia,” *Journal of Economic Dynamics and Control*, 2004, *28(10)*, 2003–35.
- Peng, Lin**, “Learning with Information Capacity Constraints,” *Journal of Financial and Quantitative Analysis*, 2004, *40(2)*.
- Peress, Joel**, “Wealth, Information Acquisition and Portfolio Choice,” *The Review of Financial Studies*, 2004, *17(3)*, 879–914.
- , “The Tradeoff between Risk Sharing and Information Production in Financial Markets,” 2006. INSEAD Working Paper.
- Polkovnichenko, Valery**, “Household Portfolio Diversification: A Case for Rank Dependent Preferences,” *The Review of Financial Studies*, 2004, *Forthcoming*.
- Radner, Roy and Timothy Van Zandt**, “Real-Time Decentralized Information Processing and Returns to Scale,” *Economic Theory*, 1999, *17 (3)*, 545–575.
- Ross, Stephen**, “The Arbitrage Theory of Capital Asset Pricing,” *Journal of Economic Theory*, 1976, *13*, 341–360.
- Sims, Christopher**, “Implications of Rational Inattention,” *Journal of Monetary Economics*, 2003, *50(3)*, 665–90.
- Spence, Michael and Richard Zeckhauser**, “The Effect of the Timing of Consumption Decisions and the Resolution of Lotteries on the Choice of Lotteries,” *Econometrica*, 1972, *40*, 401–401.

- Timmermann, Allan**, “How Learning in Financial Markets Gnerates Excess Volatility and Predictability in Stock Prices,” *Quarterly Journal of Economics*, 1993, *108*, 1135–1145.
- Van Nieuwerburgh, Stijn and Laura Veldkamp**, “Information Immobility and the Home Bias Puzzle,” 2005. Working Paper.
- Vayanos, Dimitri**, “The Decentralization of Information Processing in the Presence of Interactions,” *Review of Economic Studies*, 2003, *70*, 667–695.
- Wang, Jiang**, “A Model of Intertemporal Asset Prices Under Asymmetric Information,” *Review of Economic Studies*, 1993, *60*, 249–282.
- Wilson, Robert**, “Informational Economies of Scale,” *Bell Journal of Economics*, 1975, *6*, 184–95.

A Appendix

A.1 Proof of proposition 1 (generalized learning)

By symmetry, $\partial\mathcal{L}/\partial\widehat{\Sigma}_{ii}^{-1} = \partial\mathcal{L}/\partial\widehat{\Sigma}_{jj}^{-1}$ for all assets i, j , as long as $\widehat{\Sigma}_{ii} = \widehat{\Sigma}_{jj}$. Suppose that the solution involves learning more about some asset i than about another asset j : $\widehat{\Sigma}_{ii}^{-1} > \widehat{\Sigma}_{jj}^{-1}$. Then, by concavity of the objective $\partial\mathcal{L}/\partial\widehat{\Sigma}_{ii}^{-1} < \partial\mathcal{L}/\partial\widehat{\Sigma}_{jj}^{-1}$. Since the first-order condition for an optimum is $\partial\mathcal{L}/\partial\widehat{\Sigma}_{ii}^{-1} = 0 \forall i$, the inequality implies that $\partial\mathcal{L}/\partial\widehat{\Sigma}_{jj}^{-1} > 0$. Therefore, the candidate solution does not satisfy the first-order condition in $\widehat{\Sigma}_{jj}^{-1}$ and therefore cannot be an optimum.

A.2 Proof of corollary 1

Step 1: Derive expected utility. CARA utility, after substituting in the budget constraint $W = rW_0 + q'(f - pr)$ is

$$U_1 = -E_1[\exp(-\rho(rW_0 + q'(f - pr)))].$$

Since r and W_0 are not choice variables and are multiplicative constants, we can drop these without changing the optimization problem. Substituting in the optimal portfolio in (14) and canceling out ρ in the denominator and numerator yields

$$U_1 = -E_1[\exp(-(\hat{\mu} - pr)' \widehat{\Sigma}^{-1} (f - pr))].$$

Using the law of iterated expectations, we take expectations in two steps. First, we compute an expectation over f , conditional on all the information known at time 2. Payoffs f have a mean of $\hat{\mu}$ and a variance of $\widehat{\Sigma}$. Using the formula for the mean of a log-normal to compute the expectation over f yields

$$U_1 = -E_1[\exp(-\frac{1}{2}(\hat{\mu} - pr)' \widehat{\Sigma}^{-1} (\hat{\mu} - pr))].$$

The second expectation is taken over the unknown posterior belief $\hat{\mu}$ at time 1. $\hat{\mu}$ has mean μ and variance $\Sigma - \widehat{\Sigma}$. However, this is not a log-normal variable anymore because $\hat{\mu}$ enters as a square. This is a Wishart variable. To compute its expectation, it is useful to rewrite the objective function in terms of the mean zero variable $\hat{\mu} - \mu$.

$$U_1 = -E[\exp\{-\frac{1}{2}[(\hat{\mu} - \mu)' \widehat{\Sigma}^{-1} (\hat{\mu} - \mu) + 2(\mu - pr)' \widehat{\Sigma}^{-1} (\hat{\mu} - \mu) + (\mu - pr)' \widehat{\Sigma}^{-1} (\mu - pr)]\}].$$

Then, we apply the formula for an expectation of a Wishart (Brunnermeier (2001), p.64). If $z \sim N(0, \tilde{\Sigma})$, then

$$E[e^{z'Fz + G'z + H}] = |I - 2\tilde{\Sigma}F|^{-1/2} \exp[\frac{1}{2}G'(I - 2\tilde{\Sigma}F)^{-1}\tilde{\Sigma}G + H] \quad (26)$$

Applying this formula yields

$$\begin{aligned} U_1 &= -|I - 2(\Sigma - \widehat{\Sigma})(-\frac{1}{2}\widehat{\Sigma}^{-1})|^{-1/2} \exp\{\frac{1}{2}(\mu - pr)' \widehat{\Sigma}^{-1} (I - 2(\Sigma - \widehat{\Sigma})(-\frac{1}{2}\widehat{\Sigma}^{-1}))^{-1} \\ &\quad \times (\Sigma - \widehat{\Sigma}) \widehat{\Sigma}^{-1} (\mu - pr) - \frac{1}{2}(\mu - pr)' \widehat{\Sigma}^{-1} (\mu - pr)\} \\ &= -|I + (\Sigma \widehat{\Sigma}^{-1} - I)|^{-1/2} \exp\{\frac{1}{2}(\mu - pr)' \widehat{\Sigma}^{-1} \left((I + (\Sigma \widehat{\Sigma}^{-1} - I))^{-1} (\Sigma \widehat{\Sigma}^{-1} - I) - I \right) (\mu - pr) \} \end{aligned}$$

After combining terms, expected utility simplifies to

$$U_1 = - \left(|\Sigma| / |\widehat{\Sigma}| \right)^{-1/2} \exp(-\frac{1}{2}(\mu - pr)' \Sigma^{-1} (\mu - pr)). \quad (27)$$

Step 2: Solve for optimal precision allocation. Since (27) is a strictly increasing function of $|\widehat{\Sigma}|^{-1} = \prod_i \widehat{\Sigma}_{ii}^{-1}$ and all the other terms are comprised solely of exogenous variables, maximizing (27) is equivalent to maximizing $\prod_i \widehat{\Sigma}_{ii}^{-1}$, which in turn is equivalent to maximizing the natural log of that objective $\sum_i \ln(\widehat{\Sigma}_{ii}^{-1})$. Thus the investor's problem is equivalent to maximizing the Lagrangian

$$\max_{\{\widehat{\Sigma}_{ii}^{-1}\}_i} \sum_i \ln(\widehat{\Sigma}_{ii}^{-1}) + \psi \left(K - \sum_i \widehat{\Sigma}_{ii}^{-1} \right) + \sum_i \lambda_i (\widehat{\Sigma}_{ii}^{-1} - \Sigma_{ii}^{-1}). \quad (28)$$

Note that since this is a concave objective plus linear constraints that the objective is strictly concave over the entire domain of $\widehat{\Sigma}_{ii}^{-1}$. Therefore, the first order condition characterizes a global optimum. That first order condition is

$$\widehat{\Sigma}_{ii} = \psi + \lambda_i \quad (29)$$

The complementary slackness condition tell us that $\lambda_i (\widehat{\Sigma}_{ii}^{-1} - \Sigma_{ii}^{-1}) = 0$. This means that $\widehat{\Sigma}_{ii} = \psi$ whenever $\widehat{\Sigma}_{ii}^{-1} > \Sigma_{ii}^{-1}$. Thus, $\widehat{\Sigma}_{ii}^{-1} = \max(\psi^{-1}, \Sigma_{ii}^{-1})$. Finally, it must be that $\underline{\sigma} = \psi^{-1}$ because (27) is strictly increasing in signal precision. Thus, the capacity constraint must bind. The capacity constraint binding implies that $\sum_i \max(\psi^{-1}, \Sigma_{ii}^{-1}) = K$, which delivers the solution in the corollary.

A.3 Proof of proposition 2 (specialization)

By definition, for any convex function \mathcal{L} and $\alpha \in (0, 1)$, $\alpha\mathcal{L}(x_1) + (1 - \alpha)\mathcal{L}(x_2) \geq \mathcal{L}(\alpha x_1 + (1 - \alpha)x_2)$, for any arguments x_1 and x_2 in the domain of \mathcal{L} . Consider two arguments of the Lagrangian, each of which is a posterior precision matrix that uses all capacity to increase the precision of one asset. Let $S(i)$ be a matrix that has every entry identical to Σ^{-1} , except for the i, i entry. Then, any choice of $\widehat{\Sigma}^{-1}$ that uses capacity to learn about two or more assets can be expressed as a linear combination of S matrices.

Suppose the investor learns about two assets i and j , then $\widehat{\Sigma}^{-1} = \alpha S(i) + (1 - \alpha)S(j)$. Since $\mathcal{L}(\alpha S(i) + (1 - \alpha)S(j)) \leq \alpha\mathcal{L}(S(i)) + (1 - \alpha)\mathcal{L}(S(j))$, it must be that either $\mathcal{L}(\alpha S(i) + (1 - \alpha)S(j)) \leq \mathcal{L}(S(i))$ or $\mathcal{L}(\alpha S(i) + (1 - \alpha)S(j)) \leq \mathcal{L}(S(j))$. Thus, for any interior solution $\widehat{\Sigma}^{-1}$, either $S(i)$ or $S(j)$ achieves an equal or greater value of the objective \mathcal{L} .

Similarly, suppose the investor learns about 3 assets i, j and l . Then, we can construct two posterior precision matrices, one matrix $S(i)$ that differs from Σ^{-1} only in the (i, i) entry and another matrix $S(j, l)$ that differs from Σ^{-1} only in the j, j and l, l entries, such that $\widehat{\Sigma}^{-1} = \alpha S(i) + (1 - \alpha)S(j, l)$. By the same convexity argument as above, either $S(i)$ or $S(j, l)$ must be weakly preferred to $\widehat{\Sigma}^{-1}$. Furthermore, we know from the previous step that any choice $S(j, l)$ is weakly dominated by $S(j)$ or $S(l)$. By induction, for any $\widehat{\Sigma}$ that differs from Σ^{-1} in more than one diagonal element, we can construct a matrix S that differs from Σ^{-1} in only one diagonal element that achieves at least as high a value of the objective \mathcal{L} . We know that a maximum exists because \mathcal{L} is assumed to be a continuous function on a compact set, $[0, \infty]$. Therefore, some matrix $S(i)$ that differs from Σ^{-1} in only one diagonal element i, i must maximize the Lagrangian function \mathcal{L} .

A.4 Proof of Corollary 2

Step 1: Show that the Lagrangian is convex To do this, we first redefine the choice variable. The investor chooses $(K_1, \dots, K_N) \geq 0$ where the choice variable measures the log increase in precision: $K_i = \ln(\widehat{\Sigma}_{ii}^{-1}/\Sigma_{ii}^{-1})$, subject to the constraint that $\sum_i K_i \leq \ln(K)$. The Lagrangian problem corresponding to the objective in (20) is

$$\mathcal{L} = \frac{1}{2} \left\{ -N + \sum_{i=1}^N e^{K_i} (1 + \theta_i^2) \right\} + rW_0 + \lambda(\ln(K) - \sum_i K_i) + \sum_{i=1}^N \phi_i k_i, \quad (30)$$

where λ is the Lagrange multiplier on the capacity constraint and ϕ is the multiplier on the no-forgetting constraint.

The first derivative of \mathcal{L} is

$$\frac{\partial \mathcal{L}}{\partial K_i} = \frac{1}{2} e^{K_i} (1 + \theta_i^2) - \lambda + \phi_i. \quad (31)$$

The second derivative is

$$\frac{\partial^2 \mathcal{L}}{\partial K_i^2} = \frac{1}{2} e^{K_i} (1 + \theta_i^2) > 0. \quad (32)$$

Thus, the problem has positive second partial derivatives in all the choice variables. The sufficient condition for convexity is that the Hessian of the problem is positive semi-definite. Note that \mathcal{L} is additively separable in the K_i 's. Therefore, all the cross-partial derivatives are zero. Since the Hessian of \mathcal{L} is a diagonal matrix, it is positive definite if and only if all its diagonal elements are positive. Since the diagonal elements of the Hessian are the second partial derivatives $\frac{\partial^2 \mathcal{L}}{\partial K_i^2}$, and we have shown those to be positive, the function \mathcal{L} is convex.

Step 2: Show which asset the investor specializes in. By proposition 2, we know that the investor optimally learns about only one asset. We determine which asset is optimal to learn about by comparing the expected utility of each corner solution. The additional utility of devoting the total capacity K to asset i is $(e^K - 1)(1 + \theta_i^2)$. This is maximized by choosing $i = \text{argmax}_j \theta_j^2$. Thus, the investor maximizes his objective by learning exclusively about the asset i with the highest squared Sharpe ratio θ_i^2 .

A.5 Solving the model with CRRA preferences

To solve the model, we substitute the wealth process into the utility function and compute expected utility for a given portfolio choice q for an investor who has posterior beliefs $(\hat{\mu}, \hat{\Sigma})$.

$$U_2 = \frac{1}{1-\gamma} W_t^{1-\gamma} \exp \left\{ (1-\gamma) \left[(r + q'(\hat{\mu} - r)) - \frac{1}{2} q' \hat{\Sigma} q \right] \right\} E_t \left[\exp((1-\gamma) q' \hat{\Sigma}^{1/2} z_t) \right]$$

We take the expectation of the last term, a lognormal variable, and then rearrange terms:

$$U_2 = \frac{1}{1-\gamma} W_t^{1-\gamma} \exp \left\{ (1-\gamma) \left[(r + q'(\hat{\mu} - r)) - \frac{1}{2} q' \hat{\Sigma} q + \frac{1}{2} (1-\gamma) q' \hat{\Sigma} q \right] \right\} \quad (33)$$

$$U_2 = \frac{1}{1-\gamma} W_t^{1-\gamma} \exp \left\{ (1-\gamma) \left[r + q'(\hat{\mu} - r) - \frac{1}{2} \gamma q' \hat{\Sigma} q \right] \right\}. \quad (34)$$

The investor chooses q to maximize this expected utility. The objective is concave in q so that the first-order condition characterizes the optimal portfolio. That first-order condition yields the portfolio in (24).

The next step is to compute expected utility with the optimal portfolio and given beliefs about the mean and variance of the next period's asset value. Substituting this portfolio into the utility function yields expected utility conditional on information Σ and μ .

$$\begin{aligned} U_2 &= \frac{W_t^{1-\gamma}}{1-\gamma} e^{(1-\gamma)r} \exp \left\{ \frac{1-\gamma}{\gamma} \left[(\hat{\mu} - r)' \hat{\Sigma}^{-1} (\hat{\mu} - r) - \frac{1}{2} (\hat{\mu} - r)' \hat{\Sigma}^{-1} (\hat{\mu} - r) \right] \right\} \\ &= \frac{1}{1-\gamma} W_t^{1-\gamma} e^{(1-\gamma)r} \exp \left\{ \frac{1-\gamma}{2\gamma} (\hat{\mu} - r)' \hat{\Sigma}^{-1} (\hat{\mu} - r) \right\}. \end{aligned}$$

It is useful to rewrite this problem in terms of the mean-zero random variable $(\hat{\mu} - r - m) \sim N(0, \Sigma - \hat{\Sigma})$, where $m = \mu - r$.

$$U_2 = \frac{1}{1-\gamma} W_t^{1-\gamma} e^{(1-\gamma)r} \exp \left\{ \frac{1-\gamma}{2\gamma} \left[(\hat{\mu} - r - m)' \hat{\Sigma}^{-1} (\hat{\mu} - r - m) + 2m' \hat{\Sigma}^{-1} (\hat{\mu} - r - m) + m' \hat{\Sigma}^{-1} m \right] \right\}$$

Then, we can apply formula (26) for the expectation of a Wishart variable

$$\begin{aligned}
U_1 &= \frac{W_t^{1-\gamma}}{1-\gamma} e^{(1-\gamma)r} |I - 2(\Sigma - \widehat{\Sigma}) \frac{1-\gamma}{2\gamma} \widehat{\Sigma}^{-1}|^{-1/2} \times \\
&\exp \left\{ \frac{(1-\gamma)^2}{2\gamma^2} (\mu - r)' \widehat{\Sigma}^{-1} \left(I - 2(\Sigma - \widehat{\Sigma}) \frac{1-\gamma}{2\gamma} \widehat{\Sigma}^{-1} \right)^{-1} (\Sigma - \widehat{\Sigma}) \widehat{\Sigma}^{-1} (\mu - r) + \frac{1-\gamma}{2\gamma} (\mu - r)' \widehat{\Sigma}^{-1} (\mu - r) \right\}. \\
&= \frac{(e^r W_t)^{1-\gamma}}{1-\gamma} |I - \frac{1-\gamma}{\gamma} \Sigma \widehat{\Sigma}^{-1}|^{-1/2} \exp \left\{ \frac{(1-\gamma)}{2\gamma} (\mu - r)' \widehat{\Sigma}^{-1} \left((1-\gamma)(I - (1-\gamma)\Sigma \widehat{\Sigma}^{-1})^{-1} (\Sigma \widehat{\Sigma}^{-1} - I) + I \right) (\mu - r) \right\}.
\end{aligned}$$

Combining terms, rearranging and taking the log of this expression yields (25).

Note that taking the log of a Lagrangian does not change whether it has an interior maximum or not. If a function has a unique interior maximum, then any strictly increasing function of that function must also have an interior maximum. If this were not true, then there would exist two points a in the interior and b on the boundary such that $f(a) > f(b)$ but $g(f(a)) \leq g(f(b))$, which would make g not a strictly increasing function. Likewise, if a function attains its maximum on the boundary of its domain, then any increasing function of that function must also attain its maximum on the boundary.

A.6 Proof of Corollary 3

To show that increasing returns arise, we redefine the choice variables to be the amount of entropy capacity devoted to learning about each asset. The investor chooses $(K_1, \dots, K_N) \geq 0$ where the choice variable measures the increase in precision: $\widehat{\Sigma}_{ii}^{-1} = e^{K_i} \Sigma_{ii}^{-1}$, subject to the constraint that $\sum_i K_i \leq \ln(K)$. The Lagrangian problem corresponding to the objective in (25) is

$$\mathcal{L} = a + \frac{1}{2} \sum_i \log(1 + (\gamma - 1)e^{K_i}) + \frac{\gamma - 1}{2} \sum_i \frac{(\mu_i - r)^2}{\Sigma_{ii}(e^{-K_i} + \gamma - 1)} + \xi(\ln(K) - \sum_i K_i) + \phi_i K_i \quad (35)$$

where ξ and ϕ are the Lagrange multipliers on the capacity constraint and the no-forgetting constraint. The first order condition of this problem is

$$\frac{\partial U_1}{\partial K_i} = \frac{\gamma - 1}{2} \left(\frac{(1 + (\mu_i - r)^2 \Sigma_{ii}^{-1}) e^{-K_i} + \gamma - 1}{(e^{-K_i} + \gamma - 1)^2} \right) - \xi + \phi_i = 0. \quad (36)$$

But this first order condition does not characterize the solution to the problem because the second order condition is not satisfied. The second derivative is

$$\frac{\partial^2 U_1}{\partial K_i^2} = \frac{\gamma - 1}{2} e^{-K_i} \left(\frac{(1 + (\mu_i - r)^2 \Sigma_{ii}^{-1}) e^{-K_i} + \gamma - 1 + (\gamma - 1)(\mu_i - r)^2 \Sigma_{ii}^{-1}}{(e^{-K_i} + \gamma - 1)^3} \right). \quad (37)$$

This expression is positive because we assumed that $\gamma > 1$. Since the Lagrangian is additively separable in each Σ_{ii}^{-1} and has positive second partial derivatives, it is a convex function. Therefore, by proposition 2, we know that the investor will learn nothing about all assets except one and devote all of his capacity to learning about (specializing in) one asset.

To determine which asset the investor specializes in, we look at (35). Note that the first two terms are the same for any asset the investor might specialize in because specialization involves setting $K_i = \ln(K)$. Likewise, the last two terms are zero, by the Kuhn-Tucker conditions. Therefore, the investor chooses to specialize in the asset that delivers the largest increase in $(\mu_i - r)^2 / (\Sigma_{ii}(e^{-K_i} + \gamma - 1))$. Choosing to specialize in asset i means that $1/(e^{-K_i} + \gamma - 1)$ is larger than it would be if the investor did not learn about i . Therefore, the investor maximizes his utility by choosing the asset with the largest value of $(\mu_i - r)^2 / (\Sigma_{ii})$ and devoting his capacity to that asset.

A.7 Proof of Proposition 3

Expected utility with CARA utility is given by (27). Since the exponential term is comprised of all variables the investors takes as given, and an exponential is positive (unless every single asset has zero expected excess return), then maximizing this objective is equivalent to maximizing $-\left(|\Sigma|/|\widehat{\Sigma}|\right)^{-1/2}$ or maximizing $\left(|\Sigma|/|\widehat{\Sigma}|\right)^{1/2}$.

Recall that the entropy learning technology (8) is $|\Sigma|/|\widehat{\Sigma}| \leq K$. The objective then becomes to maximize $K^{1/2}$, subject to a constraint on K . This tells us that, while the agent prefers more capacity to less (the capacity budget constraint binds), every allocation of a given amount of capacity yields equal expected utility.