Information Acquisition and Under-Diversification

Stijn Van Nieuwerburgh and Laura Veldkamp*

January 8, 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 a variety of 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: New York University, Stern School of Business, 44 West Fourth St., 9-190, New York, NY 10012; email:svnieuwe@stern.nyu.edu, Phone:212-998-0673, Fax: 212-995-4256. Veldkamp: New York University, Stern School of Business, 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.

This is a pre-copyedited, author-produced version of an article accepted for publication in The Review of Economic Studies following peer review. The version of record [Van Nieuwerburgh, Stijn, and Pierre-Olivier Weill. "Why Has House Price Dispersion Gone up?" *The Review of Economic Studies* 77, no. 4 (October 2010): 1567-1606] is available online at:
< https://doi.org/10.1111/j.1467-937X.2009.00583.x >.

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 rational-expectations investors should learn about. 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 could rationalize such concentration.

The standard utility function used in asymmetric information models of portfolio choice is one with constant absolute risk aversion (CARA). The standard measure of information flow in the information processing literature is the reduction in entropy.[2] When investors with CARA utility are charged per unit of entropy for their information, *any information choice* can be rationalized, and therefore the high degree of concentration in observed portfolios can be rationalized as well. The incentive to specialize and the desire to hold and learn about a diversified portfolio exactly offset each other.

[1]Recent 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.

[2]See Sims (2003), Moscarini (2004), and Mackowiak and Wiederholt (2006).

While the standard portfolio problem has the diversified portfolio as its unique solution, adding information choice breaks this uniqueness. A diversified portfolio is still one of the optimal portfolios. However, perturbing this problem by changing either the preferences or the learning technology does not recover this diversified portfolio. Rather, the unique optimal portfolio that emerges is typically a specialized portfolio.

Section 1 shows 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 exponential preferences, independent assets and independent signals. Information is not required to hold an asset, as in Merton (1987); rather it is a tool to reduce the conditional variance (the uncertainty) of the asset's payoff, as in Grossman and Stiglitz (1980). Investors can acquire more precise signals at a higher cost. An important question is how to model this cost function. We describe the physical learning process that underlies the standard entropy measure and an alternative measure which is linear in signal precision. We also propose and examine a third learning technology: a version of entropy-based learning that builds in decreasing returns to learning about a given asset. After the investor observes signals drawn from the distribution whose precision he has chosen, he solves a standard portfolio problem.

Section 2 explores optimal information acquisition choices. We begin with the indifference result in our baseline model. With CARA preferences and entropy-based learning costs, the solution pins down how much information to acquire, but prescribes no particular allocation of that information across assets. Replacing the entropy-based learning technology with a technology that delivers additional signal precision at a constant marginal cost delivers a unique solution. Investors want to learn about assets with high expected returns and assets they are very uncertain about. We call this learning strategy *broadening knowledge*. If we instead replace the CARA utility function with a mean-variance objective and keep the entropy-based cost, our investors choose to learn about assets with high expected returns and assets with less uncertain payoffs. Because these investors learn about assets they already know more about, we call this strategy *deepening knowledge*. Other combinations of technologies and preferences deliver information acquisition strategies in the same categories.

Section 3 characterizes optimal investment choices. In each version of this 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. Except in a knife-edge case, investors learn about one or a small set of assets. That ensures that the learning portfolio is concentrated. The total portfolio is comprised of a diversified component and a concentrated component, a pattern observed in the data (see e.g. Polkovnichenko (2004)). For an investor with no information, a diversified portfolio is optimal; our theory collapses to the standard model. As the precision of the investor's information increases, holding a perfectly diversified portfolio is still feasible, but no longer optimal.

Section 4 shows how to alter the model setup slightly to analytically solve the problem with CRRA preferences. With the linear capacity constraint, the results are like the CARA model. With the entropy constraint, results mimic the model with mean-variance preferences.

While our baseline model makes the point that concentrated portfolios can be rational, the alternative models are useful because they offer precise, testable predictions. Section 5 attempts to distinguish between the various models: It asks which learning technology the investor would prefer; it explores differences in observable model predictions, and it investigates which model produces predictions most in line with investment data.

Many questions about the efficient organization of the financial services and consulting industry could be explored with this framework: 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? How do investment research choices affect asset prices? This paper provides a building block by exploring how information and investment choice can be jointly modeled, what various modeling choices 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. Signal choices and portfolio choices in this setting are circular: What an investor wants to learn depends on what he expects he will invest in and what he wants to invest in depends on what he has actually learned. To ensure that beliefs and actions are consistent, we use backwards induction. We first solve the period 2 portfolio problem for arbitrary beliefs. Then, we substitute the optimal portfolio rule into the period 1 information choice problem. Figure 1 illustrates the sequence of events.

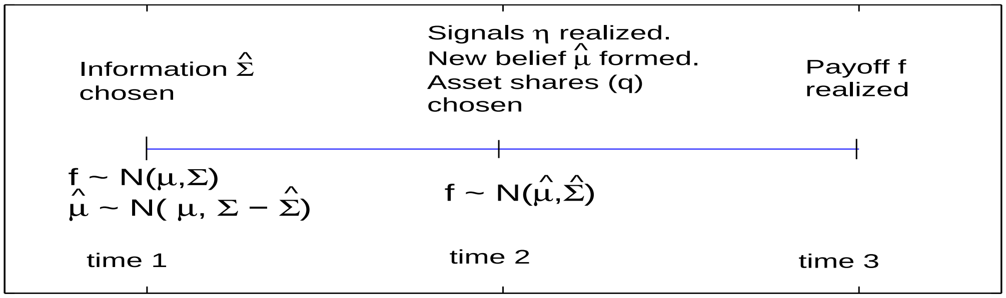

Figure 1: Sequence of events.

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 $\mu \sim N(f, \Sigma)$. 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 f. At time 2, the investor combines his signal $\eta \sim N(f, \Sigma_\eta)$ and his prior belief, using Bayes' law. Let $\hat{\mu}$ and $\widehat{\Sigma}$ be the posterior mean and variance of payoffs, conditional on all information known to the investor in period 2:

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

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

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

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 $\hat{\mu}$, $\widehat{\Sigma}$. Since the investor forms his portfolio after observing his signals, $\hat{\mu}$, $\widehat{\Sigma}$ 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

The paper considers three different commonly-used utility functions. We explore the first two, both variants of exponential, CARA preferences, in this section. Modeling CRRA preferences requires a slightly different setup, and is considered separately in section 4. In all three cases, investors have preferences over the wealth W they possess at the end of period 3 and a utility cost $c(K)$ (foregone leisure) they incur when they acquire a quantity of information K.

CARA (expected exponential) preferences With CARA utility and absolute risk-aversion coefficient ρ, the investor maximizes

$$U_1 = -E_1\left[E_2\left[\exp(-\rho W)\right]\right] - c(K) = -E_1\left[\exp(-\rho W)\right] - c(K). \tag{3}$$

The second equality follows from the law of iterated expectations.

Mean-variance preferences Investors have mean-variance utility with absolute risk aversion ρ:

$$U_1 = E_1\left[E_2[W] - \frac{\rho^2}{2} V_2[W]\right] - c(K). \tag{4}$$

Mean-variance utility routinely arises in settings where investors have CARA (exponential) utility and face normally distributed payoffs. If there were no information choice, then (3) and (4) would be equivalent because mean-variance utility is the log, a monotonically increasing transformation, of the objective function. Information choice introduces an extra expectations operator because the posterior beliefs $(\hat{\mu}, \widehat{\Sigma})$ are not known at time 1. The utility function in (4) is equivalent to

$$U_1 = -E_1\left[\log\left(E_2\left[\exp\left(-\rho W\right)\right]\right)\right] - c(K).$$

The log operator 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).

Agents with expected utility (3) only value information for its ability to increase expected portfolio returns. They see a total amount of uncertainty about asset payoffs Σ at time 1, know that it will all be resolved when payoffs are revealed at time 3, and are indifferent to the amount of risk faced at time 2, just after signals are revealed but before the investment portfolio is formed. With a preference for early resolution of uncertainty, as in (4), information is also valued for its ability to reduce portfolio risk. Investors with mean-variance utility choose information that maximizes the certainty-equivalent of expected utility, at the time when they make portfolio decisions. When choosing what to learn, such investors ask themselves, “When I invest, what information will I most want to know?" Alternatively, (4) would be the objective of a profit-maximizing portfolio manager who invests on behalf of clients with CARA utility (3). These clients would be willing to pay up to their certainty equivalent consumption for information services that learning $(\hat{\mu}, \widehat{\Sigma})$ represents. The portfolio manager therefore chooses information that maximizes the expected value of this certainty equivalent. Finally, mean-variance utility is a second-order Taylor approximation of CRRA utility around the expected level of wealth.[3]

1.3 Portfolio Allocation Choice

Given his posterior beliefs, the investor chooses the $N \times 1$ vector $q \equiv [q_1, \ldots, q_N]'$ 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 $p \equiv [p_1, \ldots, p_N]'$. In making that choice, he is subject to a budget constraint

$$W = W_0 r + q'(f - pr). \tag{5}$$

Following Admati (1985), we call $f_i - p_i r$ asset i's excess return and $q'(f - pr)$ the portfolio 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 problem is then exactly the same, if q, p and f are the quantity invested, price and payoff of the linear combinations of assets. Investors learn about and invest in the risk factors, just as if they were independent underlying assets.

1.4 Information Allocation Choice

There are two aspects of information costs: the function $c(K)$ and the mapping between the matrix of signal precisions and the scalar K, which we call *capacity*. The former aspect determines how much capacity in acquired. By a simple duality argument, we know that for every function $c(\cdot)$, there

[3] A 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.

is an equivalent endowment of capacity K^* that delivers the same portfolio predictions. Therefore, we 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 take that optimal level of K as given for the remainder of our analysis. The mapping between signal precisions and K is important for how investors form their portfolios. Therefore, we investigate three such mappings, which we call *learning technologies*.

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. Learning makes the conditional variance $\widehat{\Sigma}$ (uncertainty) lower than the unconditional variance Σ.

Independent signals We make a simplifying assumption that is not without loss of generality. Investors cannot obtain signals with correlated information about risks that are independent. Working with uncorrelated signals simplifies the problem greatly. Such an assumption has a straightforward interpretation: Investors do not choose a correlation structure for posterior beliefs. Instead, they only solve a signal precision allocation problem. They take the structure of risks in the world as given and decide how much to reduce each risk through information acquisition. Appendix A.1 shows that the key results survive with correlated assets and signals.

A no-forgetting constraint The first constraint is common to all learning technologies. The variance of each signal must be non-negative. Without this constraint, the investor could erase what he knows about one asset in order to obtain a more precise signal about another, without violating the capacity constraint. Ruling out increasing uncertainty implies that investors cannot choose to forget information. Using (2) and the independence of assets and signals, this implies that posterior variance can never exceed prior variance

$$\widehat{\Sigma}_{ii} \leq \Sigma_{ii} \qquad \forall i \tag{6}$$

Learning Technology #1: An Entropy-Based Cost 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.[4] Entropy measures the amount of uncertainty in a random variable. It is also used to measure the complexity of information transmitted. Following Sims (2003), we model the amount of information transmitted as the reduction in entropy achieved by conditioning on that additional information (mutual information).

This technology represents learning as a process of more and more refined searching. A capacity K is equivalent to a number of binary signals that partition states of the world.[5] A simple example is where a first signal tells the investor whether the payoff realization is above or below the median outcome. The second signal tells the investor, conditional on being in the top half or the bottom half, what quartile of the state space the outcome is in. In conjunction with the first two signals, the third reveals what eighth of the sample space the outcome is in, and so forth. Because the interpretation of each signal depends on its predecessors, this learning technology has the characteristics of more and more directed or refined search for an answer. This technology does not allow an investor to dynamically re-optimize his learning choice, based on signal realizations. Rather, imagine that in period 1, the investor tells the computer what asset information to download. In period 2, he reads the computer output written in binary code. When reading the binary code, the meaning of each 0 or 1 depends on the sequence of 0's and 1's that precede it.

The amount of capacity K an investor is endowed with limits how much his signals η can reduce payoff uncertainty: $|\Sigma|/|\widehat{\Sigma}| \leq K$. The work on information acquisition with one risky asset quantified information as the ratio of prior and posterior belief variances (Verrecchia (1982)). The more information a signal contains, the more the posterior variance of the asset falls below the prior variance, and the more information capacity is required to observe the signal.[6] Entropy-based capacity is a simple generalization of the same measure. In a multi-asset setting, capacity becomes the ratio of the *generalized* prior variance $|\Sigma|$ to the generalized posterior variance $|\widehat{\Sigma}|$, where the generalized variance is the determinant of the variance-covariance matrix. Because payoffs and signals are independent across assets, the determinants can be re-written as the product of the diagonal elements. Thus, a given $K > 1$ measures the product of the precisions of the investor's

[4] In 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 et al. (2002). It has been previously used in economics to model limited mental processing ability (Radner and Van Zandt (1999) and Sims (2003)) and in representative investor models in finance Peng (2004).

[5] See Cover and Thomas (1991) chapter 9.3 for a proof that the entropy of a random variable is an approximation of the number of binary signals needed to convey the same information.

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

posterior beliefs about each asset.

$$K = \frac{\prod_{i=1}^{N} \widehat{\Sigma}_{ii}^{-1}}{\prod_{i=1}^{N} \Sigma_{ii}^{-1}} \tag{7}$$

Learning Technology #2: A Linear Precision Cost An alternative learning technology is one where more information means a higher sum of signal precisions.

$$K = \sum_{i=1}^{N} (\widehat{\Sigma}_{ii}^{-1} - \Sigma_{ii}^{-1}) \tag{8}$$

With this constraint, learning takes the form of a sequence of independent draws. Each independent draw of a normally distributed signal with mean f and variance σ adds σ^{-1} to the precision of posterior beliefs. If each signal draw requires equal resources to acquire, then resources devoted to information acquisition will be a sum of signal precisions. Thus, the entropy technology represents a process of more and more refined searching while the linear technology models search as a sequence of independent explorations. As before, the investor must choose the set of signals he will draw before observing any signal realizations.

Learning Technology #3: Entropy with Decreasing Returns The final learning technology we explore represents a process of more and more refined searching, combined with some amount of randomness that can never be known. We call uncertainty that can never be resolved – not even with infinite capacity – un-learnable risk. Adding un-learnable risk is a way of generating decreasing returns to learning, which has intuitive appeal. If all risk were learnable and capacity approached infinity, the payoff variance of the optimally chosen portfolio would approach zero, an arbitrage would arise, and profit would become infinite. Un-learnable risk imposes a finite, maximum benefit to learning. Reducing an asset's learnable payoff variance to zero requires an unbounded amount of information capacity and yields only a finite benefit.

The cost function is formulated so that eliminating all learnable risk (reducing $\widehat{\Sigma}$ to $\alpha\Sigma$) requires infinite capacity, where α measures the fraction of risk that is unlearnable. When $\widehat{\Sigma} = \Sigma$, the investor is not learning anything, and no capacity is required. The resulting capacity constraint is:

$$K = \frac{|\Sigma - \alpha\Sigma|}{|\widehat{\Sigma} - \alpha\Sigma|} = \frac{\prod_{i=1}^{N} \left(\widehat{\Sigma}_{ii} - \alpha\Sigma_{ii}\right)^{-1}}{\prod_{i=1}^{N} \left(\Sigma_{ii} - \alpha\Sigma_{ii}\right)^{-1}}. \tag{9}$$

1.5 The Investor's Problem

Given a chosen level of capacity K, a solution to the model is: (1) A choice of $\widehat{\Sigma}$, the variance of posterior beliefs, to maximize utility (3 or 4), subject to the learning cost function (7, 8, or 9), the no-forgetting constraint (6), and rational expectations about q; (2) given a signal η about asset payoffs f, posterior means $\hat{\mu}$ and variances $\widehat{\Sigma}$ are formed according to Bayes' law (1) and (2); (3) A choice of portfolio q to maximize expected utility, conditional on the signal realizations.

We solve the model by backwards induction. The first order condition of the objective function with respect to q yields the optimal portfolio . Substituting the optimal portfolio into the objective delivers the indirect utility of having any beliefs $\hat{\mu}$, $\widehat{\Sigma}$ and investing optimally. To solve the period-1 problem, we take an expectation over what will be known after signals are observed ($\hat{\mu}$) and choose signal precisions (equivalently $\widehat{\Sigma}$) to maximize expected utility.

2 Results: Optimal Information Acquisition

For each utility function, we describe the optimal learning strategy with entropy-based and with linear-precision learning technologies. In the interest of space, the decreasing returns entropy technology is analyzed only in the mean-variance preference model.

2.1 CARA (Expected Exponential) Utility

After substituting the budget constraint (5) into the objective (3), the first order condition yields the optimal portfolio of risky assets. Any remaining initial wealth is invested in the risk-free asset.

$$q^\star = \frac{1}{\rho}\widehat{\Sigma}^{-1}(\hat{\mu} - pr). \tag{10}$$

2.1.1 Optimal information with an entropy constraint

After substituting in the optimal portfolio (10) into the budget constraint (5), substituting for wealth W in the objective (3) and taking the time-1 expectation of $\hat{\mu}$, utility is

$$U_1 = -A\left(|\Sigma|/|\hat{\Sigma}|\right)^{-1/2} \tag{11}$$

where $A \equiv \exp(-1/2(\mu - pr)\Sigma^{-1}(\mu - pr))$ is positive and exogenous. Maximizing this expression is equivalent to maximizing $|\Sigma|/|\hat{\Sigma}|$. Recall that the entropy constraint is $|\Sigma|/|\hat{\Sigma}| = K$. The solution to the information allocation problem is indeterminate because the time-1 expected utility only depends on capacity K, not on how that capacity is allocated across assets. Equation (11) and the following result are derived in appendix A.2.

Proposition 1. *Given the capacity constraint in (7), an investor with CARA utility (3) is indifferent between any allocation of his capacity.*

This indifference result does not say that information has no value for a CARA utility investor. Indeed, U_1 is an linearly increasing function of K. The investor still values information because it allows him to choose assets that will have higher payoffs on average. The expected excess portfolio return achieved through learning is $E_1[q'(f-pr)] - E_1[q]'E_1[f-pr]$; learned information allows investors to profit by holding a portfolio that covaries with realized payoffs. They hold a large position in assets that are likely to have a high payoff and a small (or negative) position in assets that are likely to have low payoffs. But all allocations of capacity increase utility equally. Thus, entropy is the *exponential utility neutral* learning technology.

The intuition for specialization in the introduction was that the more the investor learned about an asset, the larger the position he expected to take in that asset and the more valuable it would be to learn about the asset that would now feature more prominently in his portfolio. That mechanism is still present in this problem. But it is exactly offset by the investor's concern for portfolio variance.

An investor who learns about one asset (specializes) will learn a lot about that asset. Learning more about an asset means that posterior beliefs $\hat{\mu}$ and therefore the portfolio based on those beliefs are very uncertain at time 1. Because the investor knows he is likely to hold a large (positive or negative) position in the asset he learned about, his portfolio payoffs are very uncertain. Even though much of this risk will be resolved at time 2 when the investor observes his signal and realizes $\hat{\mu}$, this offers no additional utility to an investor who is indifferent to the timing of the resolution of uncertainty. Thus, the investor regards specialization as a risky strategy because it generates lots of risk between time 1 and time 2. Even though it leaves less risk to be resolved between time 2 and time 3, after the portfolio has been formed, the higher total amount of risk to the investor's wealth just offsets the benefit of higher expected returns from specializing. This leaves the investor indifferent between any allocation of capacity.

The left panel of figure 2 illustrates the results for a two-asset example. The horizontal and vertical axes measure the posterior precision of asset 1 and asset 2 respectively. The dashed line denotes the minimum values for the posterior precisions, given by the no-forgetting constraint. With CARA utility, the time-1 expected utility U_1 depends on the product of posterior precisions, which is the convex solid line. Higher expected utility curves lie to the northeast. With entropy-based learning, the capacity constraint is also a product of precisions. The shaded area denotes the feasible learning choices for a given K. As suggested by the indifference result, the utility function lies directly on top of the information constraint. Hence, any allocation on the north-east ridge of the feasible set delivers the same expected utility.

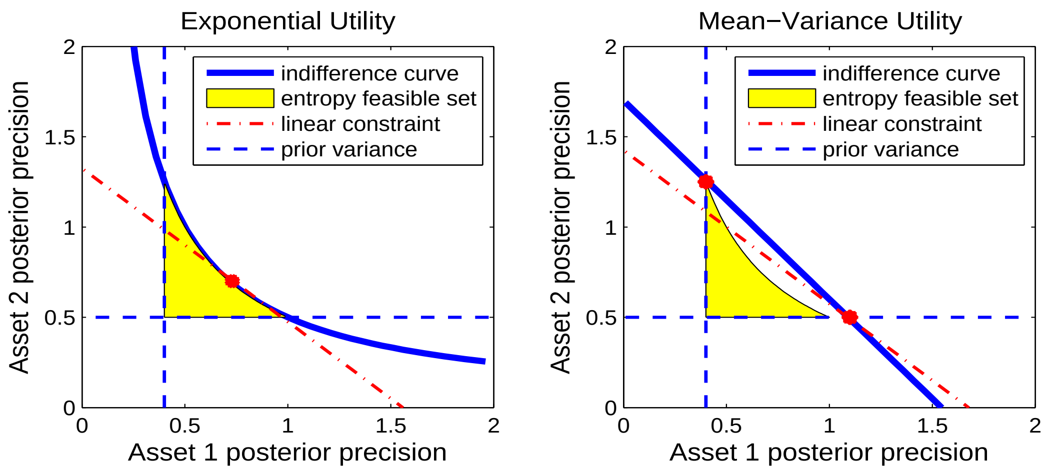

Figure 2: Indifference curves for and information constraints in a 2-asset example. Shaded area represents feasible information choices: posterior variances that satisfy the capacity and no-negative learning constraints. The two dots represent optimal information choices. In this example, the investor 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.

2.1.2 Optimal information with a linear precision cost

Linear precision is not exponential utility neutral. Among allocations of information with the same total precision, those that allocate more precision to assets with higher initial uncertainty achieve higher expected utility than others. Thus, changing the learning technology breaks the indifference result and results in an optimal portfolio that weights initially high-risk assets more than a diversified portfolio does. The reason is that doubling the precision of information about any asset increases utility (U_1) proportionately. But doubling a low precision adds less to the sum of all precisions and therefore requires less additional capacity than doubling a high precision. Therefore, it is more efficient for investors to learn more about low precision (high-uncertainty) assets. In contrast, with an entropy constraint, it is equally costly to double any precision, making the cost-benefit ratio of learning equal across all assets. Thus, investors are indifferent.

The previous result shows that the CARA utility objective U_1 takes the form of a square root of the product of posterior precisions (11). Maximizing a product (or its square root), subject to a sum constraint (8) yields an interior solution.

$$\widehat{\Sigma}^*_{ii} = min(\underline{\sigma}, \Sigma_{ii}) \quad \forall i \tag{12}$$

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. He *broadens his knowledge.* The number of assets that are being learned about is weakly increasing in capacity K. With sufficient capacity, he would set the posterior uncertainty about all assets equal.

Returning to the left panel of Figure 2, the capacity constraint is linear in precision (dashed-dotted line). Utility is maximized with an interior solution that involves learning about and thus increasing the posterior precision of information about both assets. The investor learns comparatively more about asset 1, which he was initially most uncertain about.

2.2 Mean-Variance Utility Results

Mean-variance preferences in (4) only differ from CARA utility in how they treat risk at time 1. At time 2, when the portfolio choice is made, the utilities are equivalent. Therefore, the optimal portfolio choice is the one derived in equation (10). Substituting this choice into the period 1 utility function yields $U = E_1\left[\frac{1}{2}(\hat{\mu} - pr)'\widehat{\Sigma}^{-1}(\hat{\mu} - pr)\right]$, where posterior means $\hat{\mu}$ are normally distributed. Taking the expectation of a non-central χ^2-distributed random variable, the time-1 problem is to choose signals to maximize

$$max_{\widehat{\Sigma}} \; \frac{1}{2}Tr(\widehat{\Sigma}^{-1}V_1[\hat{\mu} - pr]) + \frac{1}{2}E_1[\hat{\mu} - pr]'\widehat{\Sigma}^{-1}E_1[\hat{\mu} - pr], \tag{13}$$

where $Tr(\cdot)$ 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

$$max_{\{\widehat{\Sigma}_{11},\cdots,\widehat{\Sigma}_{NN}\}} \; \frac{1}{2}\{-N + \sum_{i=1}^N \widehat{\Sigma}_{ii}^{-1}\Sigma_{ii}(1 + \theta_i^2)\}, \tag{14}$$

where $\theta_i^2 \equiv \frac{(\mu_i - p_i r)^2}{\Sigma_{ii}}$ is the squared 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.

2.2.1 Optimal information with an entropy constraint

The objective (14) is a weighted sum of posterior precisions. The entropy-based capacity constraint 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 and making the other as low as possible, maximizes the sum, while keeping the product low. Thus the result that follows has a simple mathematical intuition that is illustrated in the right panel of figure 2. It shows that the objective is linear in precision. The capacity constraint remains convex. The optimal choice is to devote all learning capacity to asset 2.

Proposition 2. *The optimal information acquisition strategy uses all capacity to learn about one asset, the asset with the highest learning index:* $\theta_i^2 = (\mu_i - p_i r)^2 \Sigma_{ii}^{-1}$.

Proof is in appendix A.3. For intuition, consider the problem of sequentially assigning units of capacity that reduce the variance of an asset's payoff from Σ_{ii} to $\widehat{\Sigma}_{ii}$. Devoting the first unit of capacity to the asset with the highest value of $(\mu_i - p_i r)^2 \Sigma_{ii}^{-1}$ achieves the greatest utility gain. The gain from assigning the next unit of capacity to asset i is then even greater because $(\mu_i - p_i r)^2 \widehat{\Sigma}_{ii}^{-1} > (\mu_i - p_i r)^2 \Sigma_{ii}^{-1}$. The value of assigning each subsequent unit of capacity to i rises higher and higher, while the value of assigning capacity to all other assets remains the same. Therefore, the optimal posterior variance is $\widehat{\Sigma}_{ii} = \Sigma_{ii}/K$, and $\widehat{\Sigma}_{jj} = \Sigma_{jj}$ for all $j \neq i$.

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

Timing of the resolution of uncertainty Another way of understanding the difference between the CARA and mean-variance problems is by seeing learning as a tool to resolve uncertainty sooner. Learning reduces $\widehat{\Sigma}$, which is risk that the investor faces at time 2. Since all risk is resolved at time 3, when asset payoffs are revealed, the only role for learning is to resolve that risk earlier (in period 2). The mean-variance investor is indifferent to risk borne between time 1 and time 2. He is only averse to the residual risk he faces after he learns when he forms his portfolio. Therefore, the mean-variance investor prefers specialization because it offers him high expected profits and he is not averse to the resulting high uncertainty about what he will learn at time 2.

In contrast, the CARA (expected utility) investor is as averse to risk that will be resolved at time 2 through learning as he is to risk that will be resolved at time 3 when asset payoffs are observed. Since specialization generates greater uncertainty about what will be learned, and this investor is averse to that risk, he does not strictly prefer a specialized portfolio.

2.2.2 Optimal information with a linear precision constraint

The investor's learning problem is to maximize the weighted sum of precisions in (14) subject to a constraint on the un-weighted sum of precisions in (8). As long as there is a unique $i^* = \text{argmax}_i \Sigma_{ii}(1+\theta_i^2)$, the optimal learning choice is to learn exclusively about asset i^*. The investor specializes in learning about an asset he is initially uncertain about (high Σ_{ii}) and has a high squared Sharpe ratio θ_i^2. When assets have similar levels of initial uncertainty and differ in their expected returns, this is the same solution as under entropy-based learning. But when assets have similar expected returns and differ in their initial uncertainty, the investor will specialize in a different asset. This is the case illustrated in Figure 2 (right panel).

2.2.3 Optimal information with decreasing returns to learning

With the entropy information constraint, investors always deepen and never broaden their knowledge. The reason they have no incentive to learn about diverse assets is that learning substitutes for diversification in reducing risk. As learning increases and risk falls, the value of diversification falls as well. With un-learnable risk, there is some risk that learning cannot eliminate, but diversification can. This risk revives benefits unique to diversification and makes high-capacity investors broaden their knowledge.

Proposition 3. *When there is un-learnable risk, the number of assets that the investor learns about is an increasing step function of capacity K.*

Proof is in appendix A.4. The investor chooses $\widehat{\Sigma}_{ii}$ to maximize (14), subject to the capacity constraint (9), with multiplier ξ, and the no-negative learning constraint (6) for each asset i, with multiplier ϕ_i. The reason for learning about additional assets can be seen by examining the first-order condition

$$(1+\theta_i^2)\frac{\Sigma_{ii}}{\widehat{\Sigma}_{ii}^2} = \xi K \frac{1}{\widehat{\Sigma}_{ii} - \alpha\Sigma_{ii}} - \phi_i. \tag{15}$$

The left side is the marginal benefit of reducing $\widehat{\Sigma}_{ii}$ by one unit. As the investor learns more and $\widehat{\Sigma}_{ii}$ decreases, the marginal benefit increases. Increasing returns to scale in learning are still present. However, the marginal cost of reducing $\widehat{\Sigma}_{ii}$ by one unit (the right side) is now also convex in $\widehat{\Sigma}_{ii}$. The difference of the two, the net marginal benefit, first increases until $\widehat{\Sigma}_{ii} = 2\alpha\Sigma_{ii}$, and then decreases. In the limit, as the investor gets closer to learning all the learnable risk ($\widehat{\Sigma}_{ii} \to \alpha\Sigma_{ii}$), the marginal cost approaches infinity, but the marginal benefit is finite. Therefore, there is some finite cutoff level of $\widehat{\Sigma}_{ii}^{-1}$ such that when the investor's information exceeds this precision, he begins to learn about a second asset.

3 Results: Optimal Portfolios

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

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

From here on, we will use the amount of under-diversification to mean the difference between the optimal portfolio with learning and this diversified portfolio. Note that since they do not depend on any signals, the diversified portfolio weights are not random: $E[q^{div}] = q^{div}$.

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

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

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

plus his position in the risk free asset.[7] Different learning technologies lead to a different composition

[7] This expression does not have a $\hat{\mu}$ term in it because it is the *unconditional expectation* of the learning portfolio.

of q^{learn} because they induce different choices of $\widehat{\Sigma}$. The following result, which shows how learning and diversification trade off, holds for all technologies.

Proposition 4. *As long as an investor learns about at least one asset $i^\star$ for which $(\mu_i - p_i r) \neq 0$, then when capacity rises, the expected fraction of the optimal portfolio consisting of fully-diversified assets $|q^{div}|/(|q^{div}| + |E[q^{learn}]|)$ falls.*

Proof: As capacity (K) increases, the zero-capacity portfolio q^{div} is unchanged. What changes is the expected holdings of each asset $i^\star$ that the investor learns about: $|E[q^{learn}_{i^\star}]| = \frac{1}{\rho \Sigma_{i^\star i^\star}} |\mu_{i^\star} - p_{i^\star} r|(K-1)$. Since $\mu_{i^\star} - p_{i^\star} r \neq 0$, it means that $\partial |E[q^{learn}_{i^\star}]|/\partial K > 0$. $\square$

This result can be restated in terms of the more familiar value-weighted fraction of shares in the learning and diversified funds. As long as the expected return and price for the learning asset i are positive, then the expected value-weighted fraction of shares in q^{div} falls as K rises.

To describe the characteristics of these optimal portfolios, we group the models discussed previously into four categories. The first category is a portfolio with *knowledge deepening*, where the investor chooses to learn more about assets he already know lots about. The second is a portfolio with *knowledge broadening*, where investors learn about assets whose payoffs they are initially uncertain about. Third, there are portfolios that result from the *diminishing returns* strategy where investors initially deepen and then broaden as capacity increases. Finally, the CARA utility model with entropy learning predicts indifference among all feasible learning strategies. The last category offers no clear predictions for information acquisition or portfolio holdings.

3.1 Portfolio with deeper knowledge

With mean-variance preferences and and entropy learning technology, the investor deepens his knowledge about the 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 $E[q^{learn}] = \frac{1}{\rho \Sigma_{ii}}(\mu_i - p_i r)(K-1)$. 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.

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 $\hat{\mu} - \mu$ term drops out in expectation.

3.2 Portfolio with broader knowledge

Broader learning arises in both models with linear capacity. Investors learn most about assets that have the most uncertain payoffs. Learning more about these high-risk assets makes the investor hold more of them, on average, in his portfolio. Therefore, broader knowledge tilts the optimal portfolio towards higher-risk assets. A mean-variance investor or a CARA investor with low capacity learns only about one asset (except in a knife-edge case). The resulting concentration of assets in the learning portfolio is similar to deeper-knowledge portfolios. The difference is that the broader investor would hold more of an asset that is less familiar to him, while the deeper investor would hold more of an asset that he was initially familiar with.

The strategy of investing aggressively in assets whose payoffs are very uncertain resembles what an over-confident investor might do. But instead of irrationally under-estimating uncertainty, our investors reduce uncertainty through learning. The informed investor uses his information to buy assets that are likely to have high payoffs and sell assets that are likely to have low payoffs ($cov(q, f)$ rises because learning makes $\hat{\mu}$ and f more correlated). Therefore, his portfolio returns should exceed those of the overconfident investor. The higher the investor's capacity, the more excess return he can earn.

There is a special case of the CARA model with linear capacity where learning results in diversified portfolios, on average – when prior variances on all assets are equal. If the investor uses his capacity to reduce uncertainty about each asset by an equal proportion, $E[q^{learn}]$ is proportional to q^{div}. While his realized portfolio still depends on the information he sees; his average (or expected) portfolio is the no-information, diversified one.

3.3 Portfolio with diminishing returns to learning

When an investor with mean-variance preferences and a diminishing returns learning technology has low capacity, he learns about only one asset. With more capacity, he learns about more assets. This pattern is similar to a CARA investor with a linear learning technology. But like the deep investors, diminishing returns investors initially choose more of assets with higher indices θ_i^2. This index is higher when expected returns are higher and initial uncertainty Σ_{ii} is lower.

As capacity increases, diversification falls, and then rises again. An investor with zero capacity holds only the diversified fund. An investor with infinite capacity holds a perfectly diversified learning fund. In the limit, he would eliminate all learnable risk, setting $\widehat{\Sigma} = \alpha\Sigma$. The learning fund would be $E[q^{learn}] = \frac{1}{\rho}\left(\frac{1}{\alpha} - 1\right)\Sigma^{-1}(\mu - pr)$. This is a scaled-up copy of the diversified fund (when α is the same across assets). In between the two perfectly diversified extremes, the investor with positive, finite capacity to learn deviates from the diversified portfolio by tilting his portfolio toward more familiar (and higher-return) assets.

4 An Information Choice Model with CRRA Preferences

While CARA preferences are commonly used in the asymmetric information literature because they are tractable, constant relative risk aversion (CRRA) preferences are much more commonly used in macroeconomics, asset pricing, and dynamic portfolio choice. This section shows how to analytically solve a CRRA portfolio problem with information choice. This change in utility from CARA to CRRA breaks the investor's indifference between deeper or broader knowledge, in favor of deepening, just like mean-variance preferences do.

In the continuous time portfolio literature, the canonical setting is the Merton (1987) problem. Since the question of how to value information that can be used for many periods is beyond the scope of this paper, we use a static version of that model.

4.1 Model setup

The investor's utility is a function of his end-of-period wealth and the amount of information he acquires. Utility over wealth exhibits constant relative risk aversion $\gamma > 1$.[8]

$$U_t = E_t\left[\frac{1}{1-\gamma}\left(W_{t+1}\right)^{1-\gamma}\right] - c(K) \tag{18}$$

In continuous time, an $N \times 1$ vector of asset values S_t has a stochastic process given by $dS_t = diag(S_t)(\mu dt + \Sigma^{1/2} dZ_t)$, 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: $dW_t = W_t\left[(r + q'(\mu - r))dt + q'\Sigma^{1/2} dZ_t\right]$. Because we analyze a static model, we approximate this continuous-time process with a discrete time process that has the same mean and variance as the continuous process.

$$W_{t+1} = W_t \exp\{r + q'(\mu - r) - \frac{1}{2}q'\Sigma q + q'\Sigma^{1/2} z_t\} \quad \text{where} \quad z_t = Z_{t+1} - Z_t \sim N(0, I_n). \tag{19}$$

[8] Appendix A.7 derives similar results for the case when $\gamma < 1$.

Information acquisition As before, the investor acquires information about the future value of the asset. The signal is normally distributed with mean z_t and a variance-covariance matrix that is chosen. As before, we assume that asset values and signals are independent across assets. Agents update using Bayes' law and form posterior means $\hat{\mu}$ and variances $\widehat{\Sigma}$ as in (1) and (2). We model the information choice as a choice of posterior variance $\widehat{\Sigma}$. 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 $\hat{\mu}$).

4.2 Results: The Portfolio Choice Problem

To solve the model, we substitute the wealth process into the utility function and compute expected utility for a given portfolio choice q. The investor chooses q to maximize that utility. The first-order condition characterizes the optimal portfolio, as a share of the investor's wealth:

$$q^* = \frac{1}{\gamma}\hat{\Sigma}^{-1}(\hat{\mu} - r). \tag{20}$$

This portfolio has the same form as the optimal portfolio with CARA preferences in section 2. Note that $(\hat{\mu} - r)$ is the expected return in excess of the risk free rate, while in the CARA problem the equivalent term was $(\hat{\mu} - pr)$, the expected payoff of the asset, minus its opportunity cost.

Substituting q^* into (19) and (18) yields expected utility conditional on information $\widehat{\Sigma}$ and $\hat{\mu}$

$$U_2 = \frac{W_t^{1-\gamma}}{1-\gamma} e^{(1-\gamma)r} \exp\left\{ \frac{1-\gamma}{2\gamma}(\hat{\mu} - r)'\hat{\Sigma}^{-1}(\hat{\mu} - r) \right\}. \tag{21}$$

This is the payoff of the first stage decision problem if it results in beliefs $\hat{\mu}$ and $\hat{\Sigma}$. The investor does not choose the expected return $\hat{\mu}$, only the conditional variance of the return after learning $\hat{\Sigma}$. Each information choice results in a random $\hat{\mu} \sim N(\mu, \Sigma - \hat{\Sigma})$, where μ is the prior information about the mean return and Σ is its unconditional variance. Taking the time-1 expectation delivers the objective function an investor uses to make his information choice[9]

$$U_1 = \frac{W_t^{1-\gamma}}{1-\gamma} \frac{e^{(1-\gamma)r}\gamma^{1/2}}{|I + (\gamma - 1)\Sigma\hat{\Sigma}^{-1}|^{1/2}} \exp\left[\frac{1-\gamma}{2}(\mu - r)' \left(\hat{\Sigma} + (\gamma - 1)\Sigma\right)^{-1} (\mu - r) \right]. \tag{22}$$

Working with the log of this objective simplifies the calculations. To avoid taking the log of a

[9]See appendix A.5 for details of these calculations.

negative number, we maximize $-\log(-U_1)$, 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

$$U_1 = a + \frac{1}{2}\sum_{i=1}^{N} \log\left(1 + (\gamma - 1)\Sigma_{ii}\hat{\Sigma}_{ii}^{-1}\right) + \left(\frac{\gamma - 1}{2}\right)\sum_{i=1}^{N} \frac{(\mu_i - r)^2}{\hat{\Sigma}_{ii} + (\gamma - 1)\Sigma_{ii}}. \tag{23}$$

4.3 Results: Information Choice with an Entropy Constraint

With an entropy-based learning technology, there are increasing returns to devoting additional capacity to learning about a given asset. This makes deepening knowledge optimal, just like in the case with entropy and mean-variance preferences.

Proposition 5. *The optimal information acquisition strategy uses all capacity to learn about one asset, the asset with highest squared Sharpe ratio* $(\mu_i - r)^2/\Sigma_{ii}$.

Because the CRRA objective is not as simple as the CARA one, the proof (in appendix A.6) requires redefining the choice variables to be the amount of entropy capacity devoted to learning about each asset. The investor chooses $(K_1, \ldots, K_N) \geq 0$ where the choice variable measures the increase in precision: $\widehat{\Sigma}_{ii}^{-1} = e^{K_i}\Sigma_{ii}^{-1}$, subject to the constraint that $\sum_i K_i \leq \ln(K)$. Since this constraint is now linear in the choice variable, a corner solution arises if the objective is convex. The proof shows that indeed $\partial^2 EU/\partial K_i^2 > 0$. This implies that as the investor learns more about an asset, the marginal value of learning more rises. Furthermore, the asset that is most valuable to learn about is one with a high expected return and low initial uncertainty. Thus, the investor *deepens his knowledge*.

Why is the CRRA investor not indifferent like the CARA investor? As explained before, the CARA investor sees higher expected profits from specializing, but also exposes himself to more risk between time-1 and time-2. In the CARA problem, the investor cared about expected profit, minus one-half the portfolio variance. But the additional variance from specializing was exactly twice as large as the additional expected profits. Thus, the two concerns canceled out. The CRRA ($\gamma > 1$) investor cares about expected profit, minus $(\gamma-1)/(2\gamma)$ times the variance instead. Because $(\gamma-1)/(2\gamma) < 1/2$, for all finite $\gamma > 1$, this CRRA investor weights the drawback of higher variance between time 1 and time 2 less than the CARA investor. He therefore specializes in learning about one asset. In the limit, as $\gamma \to \infty$, the CRRA investor behaves the like the CARA investor and becomes indifferent between any learning strategies.

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 $\gamma > 1$. Rather, 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 gambles is lower makes him less averse to taking this large risk.

4.4 Results: Information Choice with a Linear Precision Constraint

The investor maximizes (23) subject to the information constraint (8), which is a constraint on the sum of precisions. Taking the partial derivative with respect to $\hat{\Sigma}_{ii}^{-1}$ yields the marginal value of information about asset i.

$$\frac{\partial U_1}{\partial \hat{\Sigma}_{ii}^{-1}} = \frac{(\gamma-1)^2}{2}\left[\frac{\Sigma_{ii}(1+(\gamma-1)\Sigma_{ii}\hat{\Sigma}_{ii}^{-1}) + (\mu_i - r)^2}{(1+(\gamma-1)\Sigma_{ii}\hat{\Sigma}_{ii}^{-1})^2}\right] \tag{24}$$

Because information precision $\hat{\Sigma}_{ii}^{-1}$ enters linearly in the numerator and quadratically in the denominator, the second derivative is negative ($\partial^2 U_1/\partial(\hat{\Sigma}_{ii}^{-1})^2 < 0$). Therefore, the first order condition characterizes the optimal information choice. Setting (33) equal for each asset i and imposing the capacity constraint determines a target level of information precision for each asset σ_i^* that depends on μ_i, and Σ_{ii}. The investor achieves this target precision, as long as it does not violate his no negative learning constraint: $\widehat{\Sigma}_{ii}^* = min(\sigma_i^*, \Sigma_{ii})$. With sufficient capacity, this investor may learn about several assets. The fact that the solution is an interior solution tells us that this problem, like the CARA problem with a linear constraint, induces the investor to *broaden his knowledge.*

Figure 3 shows that if the learning technology is linear, investors with sufficient capacity should learn about multiple assets. With entropy-based learning, they should specialize.

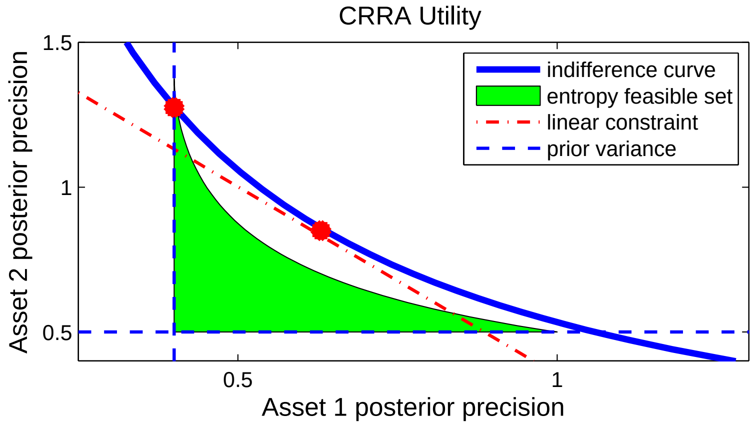

Figure 3: Indifference curves for and information constraints in a 2-asset example with CRRA utility. Shaded area represents feasible information choices: posterior variances that satisfy the capacity and no-negative learning constraints. The two dots represent optimal information choices.

4.5 Optimal Asset Portfolios with CRRA Preferences

The optimal portfolios with CRRA preferences have the same character as the CARA portfolios. 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 $E[q^{learn}] = (1/\gamma)(\widehat{\Sigma}^{-1} - \Sigma^{-1})(\mu - r)$. With the linear precision constraint, learning is generalized and q^{learn} is likely to load on multiple assets. With the entropy-based constraint, investors deepen their knowledge and q^{learn} is comprised of a single asset.

5 Comparing Learning Technologies

5.1 Utility comparison

One way to argue for one learning technology over another is to compare the welfare of the investor who uses each technology. This section shows that more informed investors prefer entropy-based learning to learning that is linear in precision because entropy-based learning uses existing information to interpret a new signal draw. With linear precision technology, the investor does not use past information when interpreting the new signal. This series of independent signal draws is preferred only when the investor has little initial information and little capacity.

With *CARA utility*, section 2.1 established that expected utility is linear in $(|\widehat{\Sigma}^{-1}|/|\Sigma^{-1}|)^{1/2}$. Since initial precision Σ^{-1} is exogenous, the learning technology that can increase $|\widehat{\Sigma}^{-1}|$ the most with an additional unit of K delivers the highest utility.

With entropy-based learning (7), each unit of capacity increases the determinant of posterior beliefs by $\partial|\widehat{\Sigma}^{-1}|/\partial K = |\Sigma^{-1}|$, no matter how it is allocated. With linear precision learning (8), investors learn about the lowest-precision (highest-variance) assets. Therefore, $\widehat{\Sigma}_{jj}^{-1} = \Sigma_{jj}^{-1} + K$ for $j = \text{argmax}_i \Sigma_{ii}$ and $\widehat{\Sigma}_{ii}^{-1} = \Sigma_{ii}^{-1}$ for all other assets. This learning strategy implies that $\partial|\widehat{\Sigma}^{-1}|/\partial K = |\Sigma^{-1}|\Sigma_{jj}$. Comparing the two expressions reveals that entropy learning increases utility more than linear learning if prior beliefs about asset j are sufficiently precise: $\Sigma_{jj} < 1$.

Of course, this comparison assumes that units of capacity are directly comparable across learning technologies. Suppose instead that one unit of linear capacity $\tilde{K}$ had an opportunity cost of $C > 0$ units of entropy capacity ($\Delta\tilde{K} = -\Delta K/C$). The additional utility from an extra unit of K used for linear learning would be $\partial|\widehat{\Sigma}^{-1}|/\partial K = |\Sigma^{-1}|\Sigma_{jj}/C$. Entropy learning delivers more marginal utility whenever $\Sigma_{jj} < C$, leaving the qualitative conclusion unchanged.

A *mean-variance objective* (14) delivers the same conclusion. It is a constant plus a weighted sum of signal precisions. With entropy-based learning $\partial U/\partial K = (1+\theta_l)^2$, where l is the asset with the highest squared Sharpe ratio θ_l^2. With linear learning, the investor increases expected utility by $\Sigma_{jj}(1+\theta_j)^2$ per unit of capacity, where j is the asset that maximizes this expected marginal utility of information. If the asset that the investor chooses to learn about is the same under both technologies ($j = l$), then the investor prefers entropy learning if $\Sigma_{jj} < 1$. Such an investor prefers entropy-based signals that bisect the space of outcomes. Their prior knowledge is precise enough that doubling its precision is better than adding one additional unit of precision.

In sum, well-informed investors, perhaps mutual fund managers, would prefer to learn with an entropy-based learning technology. These investors have a precise prior to draw on; they can use that existing information to direct their search for additional information efficiently using an entropy-based algorithm. Meanwhile, some individual investors might be sufficiently poorly informed that they would prefer the linear learning technology.

5.2 Using Data to Select a Model of Information Acquisition

Recent empirical research supports the prediction that information is responsible for investors' concentrated portfolios. The median retail investor at a large on-line brokerage company holds too few stocks, and the stocks they contain are positively correlated (Goetzmann and Kumar (2003)). But directly-held equities are only 40% of the median household's portfolio; the remaining 60% is in more diversified assets (Polkovnichenko (2004)). This is consistent with the theory's prediction that the total portfolio has diversified and specialized components. Using Swedish data on investors' personal characteristics and complete wealth portfolio, Massa and Simonov (2005) rule out the explanation that this concentration optimally hedges labor income risk. Finally, direct evidence about information acquisition also supports the connection with portfolio concentration. Guiso and Jappelli (2006) examine survey data on the time customers of a leading Italian bank spend acquiring financial information. Those who spend more time on information collection invest more in individual stocks and relatively less in diversified mutual funds. While this evidence supports the hypothesis that information acquisition can explain portfolio concentration, two other predictions distinguish deeper from broader learning.

First, deeper investors concentrate their portfolios more when their capacity increases, while broader investors 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 deepen their knowledge. Ivkovic, Sialm and Weisbenner (2005) find that concentrated investors outperform diversified ones by as much as 3% per year.[10] 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 deepen their knowledge 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)). In contrast, broadening one's knowledge 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 firms study mostly US assets and EU investment managers study more EU assets tells us that financial research is building on initial information advantages, not diversifying away from them. This suggests that financial research is more like entropy-based learning, a process of refining one's search, than linear precision learning, a process of independent sampling.

6 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

[10] See also Coval and Moskowitz (2001), Massa and Simonov (2005) and Ivkovic and Weisbenner (2005).

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 hasn't 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, relative to the convexity of the information constraint. With standard CARA preferences and the entropy-based information constraint typically used in information theory, the objective and the constraint are proportional. That means that any information acquisition strategy can be justified and with sufficient information-gathering, any portfolio can be rationalized. In other settings, there are endogenous increasing returns to information that make the objective more convex. This induces the investor to deepen his knowledge about 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?" If a manager were to sell information, information resale would undermine profits. To avoid resale, they would manage investors' portfolios directly. 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. While this is a non-trivial new problem, understanding an individual's information choice problem is a necessary first step.

This framework could be used to explain patterns of real investments or international trade as well. For example, returns to specialization could explain why small differences in distance have large effects on how much countries trade. Countries that trade more should have more incentive to know their trading partner's next-period productivity, in order to better forecast their terms of trade. But, countries that know more about each other and better anticipate terms-of-trade shocks get higher benefits from trade and trade more. Thus, more trade leads to more information acquisition, which leads to more trade. Countries specialize in learning about and trading with a few nearby countries if nearby countries are the ones they initially expected to trade more with. Because of the incentive to specialize, small differences in initial benefits from trade can have large effects on realized trade volumes.

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.

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

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.

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

Grossman, Sanford and Joeseph 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," 2006. Northwestern University Working Paper.

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.

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

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 Arbitrary Risk Factor Structure

Rather than reconsider every combination of preferences and technologies, we work through one example of the correlated assets and correlated signals model. We show that the specialization result for the mean-variance investor with an entropy learning technology survives.

For tractability, the main text assumes that investors learn about independent risks independently. In the case of correlated assets (a non-diagonal Σ), this means that signals are about the payoffs of the principal components $f'\Gamma$, where Γ comes from the decomposition $\Sigma = \Gamma'\Lambda\Gamma$. In other words, investors can learn about N orthogonal risk factors. The first such risk factor (principal component) could be interpreted as "the market." So, investors are allowed to learn about the market. The same is true for all other risk factors, which could represent industry-specific risk factors or company-specific risk factors. When asset payoffs are independent, the matrix Γ is the identity matrix, so that investors are learning about asset payoffs, not linear combinations of asset payoffs.

Let the posterior covariance matrix be $\widehat{\Sigma} = \widehat{\Gamma}'\hat{\Sigma}\widehat{\Gamma}$, where $\widehat{\Gamma}$ is an *arbitrary but fixed* posterior eigenvector matrix of $\widehat{\Sigma}$. In particular, $\widehat{\Gamma} \neq \Gamma$. The objective function (13) then becomes:

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

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

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

Under the assumption $\widehat{\Gamma} = \Gamma$, we recover the learning index written in the main text: $\tau_l = \Lambda_l + ((\mu - pr)'\Gamma_l)^2$. But, whatever the exact specification of the learning index is, the investor always ranks the risk factors according to their learning index τ_l, and chooses to specialize in the risk factor with the highest τ_l. The reason is that the objective function is still linear in the $\{\hat{\Sigma}_l^{-1}\}$. Since the entropy constraint is still a constraint on the product of the posterior eigenvalues, it is a bound on $\prod_i \hat{\Sigma}_i^{-1}$. Thus, the problem is still maximizing a sum subject to a product constraint, which delivers a corner solution.

A.2 Proof of Proposition 1

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 (10) 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 lognormal 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 lognormal 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

$$E[e^{z'Fz+G'z+H}] = |I - 2\tilde{\Sigma}F|^{-1/2}\exp[\frac{1}{2}G'(I - 2\tilde{\Sigma}F)^{-1}\tilde{\Sigma}G + H] \tag{25}$$

Applying this formula yields

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

After combining terms, expected utility simplifies to

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

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 $\left(|\Sigma|/|\widehat{\Sigma}|\right)^{1/2}$.

Note the objective collapses down to maximizing e^K. This tells us that, while the agent prefers more capacity to less (thus, the capacity budget constraint binds), every allocation of capacity yields equal expected utility. □

A.3 Proof of Proposition 2

Consider the trade-off between learning about any two assets i and j. Learning means that the agent can increase the posterior precision of his beliefs about asset i by x and of asset j by K/x, where $x \in [0, K]$ and K is capacity. The second precision is K/x because the capacity constraint is a constraint on the product of the posterior variances. Thus it is also a product constraint on precisions. Substituting in the constraint, the objective is

$$(\mu_i - p_i r)^2\Sigma_{ii}^{-1}x + (\mu_j - p_j r)^2\Sigma_{jj}^{-1}K/x$$

Its second order condition is $2(\mu_j - p_j r)^2\Sigma_{jj}^{-1}/x^3$. Since the second order condition is positive for all $x > 0$, the problem must have a corner solution. Next, we compare corner solutions. Devoting all capacity to asset i produces utility $(\mu_i - p_i r)^2\Sigma_{ii}^{-1}K + (\mu_j - p_j r)^2\Sigma_{jj}^{-1}$. Devoting all capacity to asset j produces utility $(\mu_i - p_i r)^2\Sigma_{ii}^{-1} + (\mu_j - p_j r)^2\Sigma_{jj}^{-1}K$. The first utility is greater than the second whenever $(\mu_i - p_i r)^2\Sigma_{ii}^{-1} > (\mu_j - p_j r)^2\Sigma_{jj}^{-1}$. If i is the asset with the highest learning index $(\mu_i - p_i r)^2\Sigma_{ii}^{-1}$, then there is no profitable deviation from learning about any other asset. □

A.4 Proof of Proposition 3

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

Independent assets From the proof of Proposition 4, we know that a non-zero quantity of an asset is held in the learning fund whenever the investor learns about the asset and the expected excess return is not equal to zero. Getting a signal from a continuous distribution that implies a zero excess return is a zero probability event. Since asset payoffs are independent, each risk factor corresponds to one and only one asset. Proposition 3 shows that when capacity increases, the number of risk factors learned about rises. Thus the number of assets learned about rises, and the number of different assets held in the learning fund rises. $\square$

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 to an investor who has posterior beliefs $(\hat{\mu}, \widehat{\Sigma})$.

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

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

$$U_2 = \frac{1}{1-\gamma}W_t^{1-\gamma}\exp\left\{(1-\gamma)\left[(r+q'(\hat{\mu}-r))-\frac{1}{2}q'\hat{\Sigma}q+\frac{1}{2}(1-\gamma)q'\hat{\Sigma}q\right]\right\} \tag{26}$$

$$U_2 = \frac{1}{1-\gamma}W_t^{1-\gamma}\exp\left\{(1-\gamma)\left[r+q'(\hat{\mu}-r)-\frac{1}{2}\gamma q'\hat{\Sigma}q\right]\right\}. \tag{27}$$

Given this expected utility, the investor chooses q to maximize that 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 (20).

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 $(\hat{\mu}-r-m) \sim N(0, \Sigma-\widehat{\Sigma})$, 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 (25) for a the expectation of a Wishart variable

$$U_1 = \frac{W_t^{1-\gamma}}{1-\gamma} e^{(1-\gamma)r} |I - 2(\Sigma - \widehat{\Sigma}) \frac{1-\gamma}{2\gamma} \widehat{\Sigma}^{-1}|^{-1/2} \times$$

$$\exp\left\{ \frac{(1-\gamma)^2}{2\gamma^2} (\mu - r)' \widehat{\Sigma}^{-1} \left(I - 2(\Sigma - \widehat{\Sigma}) \frac{1-\gamma}{2\gamma} \widehat{\Sigma}^{-1} \right)^{-1} (\Sigma - \widehat{\Sigma}) \widehat{\Sigma}^{-1} (\mu - r) + \frac{1-\gamma}{2\gamma} (\mu - r)' \widehat{\Sigma}^{-1} (\mu - r) \right\}.$$

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

Combining terms and rearranging this expression yields (22).

A.6 Proof of proposition 5

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 $(K_1, \ldots, K_N) \geq 0$ 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 (23) is

$$\mathcal{L} = a + \frac{1}{2} \sum_i \log(1 + (\gamma - 1)e^{K_i}) + \frac{\gamma - 1}{2} \sum_i \frac{(\mu_i - r)^2}{\Sigma_{ii}(e^{-K_i} + \gamma - 1)} + \xi(\ln(K) - \sum_i K_i) + \phi_i K_i \tag{28}$$

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

$$\frac{\partial EU}{\partial K_i} = \frac{\gamma - 1}{2} \left(\frac{(1 + (\mu_i - r)^2 \Sigma_{ii}^{-1}) e^{-K_i} + \gamma - 1}{(e^{-K_i} + \gamma - 1)^2} \right) - \xi + \phi_i = 0. \tag{29}$$

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

$$\frac{\partial^2 EU}{\partial K_i^2} = \frac{\gamma - 1}{2} e^{-K_i} \left(\frac{(1 + (\mu_i - r)\Sigma_{ii}^{-1}) e^{-K_i} + \gamma - 1 + (\gamma - 1)(\mu_i - r)^2 \Sigma_{ii}^{-1}}{(e^{-K_i} + \gamma - 1)^3} \right). \tag{30}$$

This expression is positive because we assumed that $\gamma > 1$. The positive second derivative reveals that the solution to the information choice problem is a corner solution. 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 (28). 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 is 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 Results for CRRA < 1

Expected utility is still given by (22). Working with the log of this objective simplifies the calculations. As before, we collect 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

$$EU = a + \frac{-1}{2}\sum_{i=1}^{N}\log\left(1+(\gamma-1)\Sigma_{ii}\hat{\Sigma}_{ii}^{-1}\right) + \frac{\gamma-1}{2}\sum_{i=1}^{N}\frac{(\mu_i-r)^2}{\hat{\Sigma}_{ii}+(\gamma-1)\Sigma_{ii}}. \tag{31}$$

Solution with entropy costs We redefine the choice variable as in (23). The objective and the first order condition are identical, except with opposite sign. The second derivative is

$$\frac{\partial^2 EU}{\partial K_i^2} = \frac{1-\gamma}{2}e^{-K_i}\left(\frac{(1+(\mu_i-r)\Sigma_{ii}^{-1})e^{-K_i}+\gamma-1+(\gamma-1)(\mu_i-r)^2\Sigma_{ii}^{-1}}{(e^{-K_i}+\gamma-1)^3}\right). \tag{32}$$

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

Canceling out the $(e^{-K_i}+\gamma-1)$ term leaves and expression that is positive for all assets. The positive results means there are increasing returns to learning about one asset. Thus, as long as the asset under investigation has an expected return that exceeds the risk-free rate, investors want to deepen their knowledge of that asset.

Solution with linear precision costs

$$\frac{\partial EU}{\partial \hat{\Sigma}_{ii}^{-1}} = \frac{1-\gamma}{2}\left[\frac{\Sigma_{ii}(1+(\gamma-1)\Sigma_{ii}\hat{\Sigma}_{ii}^{-1})+(\mu_i-r)^2}{(1+(\gamma-1)\Sigma_{ii}\hat{\Sigma}_{ii}^{-1})^2}\right] \tag{33}$$

The second derivative of this expression is

$$\frac{\partial^2 EU}{\partial \hat{\Sigma}_{ii}^{-2}} = \frac{-(1-\gamma)^2\Sigma_{ii}}{2}\left[\frac{\Sigma_{ii}((1-\gamma)\Sigma_{ii}\hat{\Sigma}_{ii}^{-1}-1)-2(\mu_i-r)^2}{(1+(\gamma-1)\Sigma_{ii}\hat{\Sigma}_{ii}^{-1})^3}\right] \tag{34}$$

The sign of this expression depends on the coefficient of relative risk aversion, the squared Sharpe ratio $\Sigma_{ii}^{-1}(\mu_i-r)^2$, and the amount of capacity the investor has because it determines $\Sigma_{ii}\hat{\Sigma}_{ii}^{-1}$. For very low capacity and very high capacity, this second derivative is positive. That means there are increasing returns to learning about one asset. Thus, investors deepen their knowledge. For an intermediate level of capacity, they broaden their knowledge.