

Do Heterogeneous Beliefs Matter for Asset Pricing?

Evan W. Anderson

University of North Carolina - Chapel Hill

Eric Ghysels

University of North Carolina - Chapel Hill

Jennifer L. Juergens

Arizona State University

We study how heterogeneous beliefs affect returns and examine whether they are a priced factor in traditional asset pricing models. To accomplish this task, we suggest new empirical measures based on the disagreement among analysts about expected earnings (short-term and long-term) and show they are good proxies. We first establish that the heterogeneity of beliefs matters for asset pricing and then turn our attention to estimating a structural model in which we use the forecasts of financial analysts to proxy for agents' beliefs. Finally, we investigate whether the amount of heterogeneity in analysts' forecasts can help explain asset pricing puzzles.

Economic agents differ in their beliefs, preferences, and endowments. Despite these differences and despite strong and persuasive arguments put forward for including heterogeneity in finance and macroeconomics, the representative agent paradigm is still the leading structural approach to asset pricing.¹ This has happened for many reasons. First, in many contexts it is difficult to derive testable predictions in asset pricing models with heterogeneous agents, though many researchers have made progress; for example, see Duffie and Constantinides (1996), Heaton and Lucas (1995), and Shefrin (2001). Second, there is a lack of tangible data that reflect heterogeneity. Most of the data available on the consumption, endowments, and beliefs of individual agents is sparse and of questionable quality. Third, many of the most tractable formulations of heterogeneous

We thank Stan Levine (previously of First Call) and First Call/Thomson for providing the earnings data used in this study. We express our appreciation to Campbell Harvey and anonymous referees for providing invaluable comments and suggestions. This article is a substantially revised version of an earlier article by the last two authors, Ghysels and Juergens (2001). We also like to thank Günter Franke, Raman Uppal and Fernando Zapatero for their comments as well as participants at the 2002 EFA and 2003 AFA meetings, and seminars at Arizona State University, College of William and Mary, Penn State University, the Securities and Exchange Commission, and Tulane University. Address correspondence to Jennifer L. Juergens: Department of Finance, W.P. Carey School of Business, Arizona State University, Tempe, AZ 85287, or email: jennifer.juergens@asu.edu

¹ For example see the arguments presented in Browning, Hansen, and Heckman (2000).

agent models are observationally equivalent to representative agent models. For example, when preferences Gorman aggregate (Gorman 1953) there is often no need to explicitly consider heterogeneous agents because there exists a representative agent, with a utility function of the same form as the agents, which can be used for asset pricing [see the early work of Wilson (1968) and Rubinstein (1974)].

In this article, we focus on the heterogeneity of beliefs and for simplicity let agents have the same preferences. The heterogeneity of beliefs captures how individual agents interpret or have access to differing information sets. Models with agents who have heterogeneous beliefs have been previously studied by Abel (1989), Basak (2000), Detemple and Murthy (1994), Li (1999), Shefrin (2001), Shefrin and Statman (1994), Williams (1977), and Zapatero (1998). One of our main innovations is to use the publicly stated forecasts of financial analysts to proxy for the beliefs of agents in an asset pricing model with optimizing agents. Financial analysts may not be a random sample from the general population so we discuss various specifications, some of which drive a wedge between the expectations of agents and the predictions of analysts. In our model, there is a representative agent whose beliefs are the composite of the beliefs of individual agents. Despite the existence of a representative agent, there is useful information in the beliefs of individual agents. This article (1) develops an empirically testable theoretical prediction that expected returns are impacted by heterogeneous beliefs; (2) uses a factor structure to demonstrate empirically that heterogeneity explains a portion of expected returns and volatility; and (3) estimates a consumption-based model that incorporates dispersion and biases in analysts' forecasts. In work that was undertaken and completed at about the same time as earlier drafts of this article, Shefrin (2001) studies closely related issues. Shefrin's approach is slightly more general, as we make specific assumptions for the empirical implementation of the model.

Before estimating our structural model, we take a reduced form approach and examine first whether the heterogeneity of beliefs is a priced factor in traditional asset pricing models. We suggest new empirical measures of heterogeneity. Unlike factors for size, book-to-market, and momentum that have emerged because of the evidence of empirical regularities, our hypothesis that dispersion is priced emerges from a theoretical model. We show factors constructed from the disagreement among analysts about expected (short-term and long-term) earnings have explanatory power in asset pricing models. Next we focus on return volatility. Since one might expect the same factors that predict returns also predict future volatility, we examine whether disagreement among analysts explains return volatility. Implementing methodology similar to Chan, Karceski, and Lakonishok (1999), we find simple models, particularly those that include only the

dispersion factors, produce lower forecast errors and have better sample properties than complex multi-factor models.

Although the explicit analysis of heterogeneity of beliefs is only timidly present in traditional asset pricing models, it is most often omnipresent in behavioral finance models.² The latter class of models emphasizes the psychology of individuals and its relevance for financial markets, particularly in the pricing of assets. The analysis in this article relates to the behavioral finance literature in several ways. The use of analyst forecasts prompts the question about biases in their predictions and recommendations [see Daniel, Hirshleifer, and Teoh (2002) for a comprehensive overview of analyst biases]. We estimate parameters which reflect the extent to which agents have the beliefs of financial analysts. More importantly, both psychological and purely rational theories of asset pricing generally imply that returns are predictable. However, most often behavioral finance models are not formulated explicitly in the context of pricing kernel models. One notable exception is Shefrin (2001) who analyzes the manner in which traders' errors (called sentiments) affect the pricing kernel. In this article, we estimate pricing kernels for models with heterogeneous agents whose beliefs are connected to those of analysts.

The remainder of the article is organized as follows. In Section 1 we discuss fundamental pricing equations in heterogeneous agent economies. Section 2 discusses the limited form of rationality that agents in our models have and shows that dispersion is priced. Section 3 describes the data used in this article and how the beliefs of individuals are measured. Section 4 examines the role of heterogeneity of beliefs, particularly dispersion of analyst forecasts, as a priced factor in traditional asset pricing models. This section also describes the contribution of the dispersion factor above fundamental factors to predict out-of-sample returns and volatility. Section 5 introduces the behavioral elements of amplification and bias into our model. Section 6 investigates whereas analyst forecasts can help explain asset pricing puzzles. Section 7 concludes.

1. Beliefs and the Pricing Kernel

There are a variety of ways to formulate models with agents who have heterogeneous beliefs. For example, see Detemple and Murthy (1994), Zapatero (1998), Basak (2000), and Shefrin (2001), among others. In this section, we describe a complete markets model in which agents have heterogeneous beliefs and identical power reward functions. The first subsection describes the pricing kernel, whereas the second subsection specifies how the distribution of beliefs is determined from data on beliefs about conditional means.

² It is impossible to cite the many articles in this area. See, however, the recent surveys and general discussions by Barberis and Thaler (2002) and Daniel, Hirshleifer, and Teoh (2002).

The third subsection describes how we approximate consumption forecasts. We conclude by showing how the return on an asset can be decomposed into a fundamental component that depends only on the true probabilities and a component which depends on the actual beliefs of agents.³

1.1 The pricing kernel

Consider a complete markets heterogeneous agent economy in which agents have different beliefs about the world. We assume the beliefs of any agent are absolutely continuous with respect to the true probabilities and vice versa. For now, we make no other assumptions about what agents believe or how their beliefs evolve. Agents may be Bayesians who optimally update their beliefs or they may have degenerate priors and refuse to take into account new information.

Assume there are I types of agents who have power reward functions and live an infinite number of periods. The lifetime utility function for agents of type i is

$$E_{i0} \sum_{t=0}^{\infty} \beta^t \frac{c_{it}^{1-\gamma}}{1-\gamma} \quad (1)$$

where β is a time discount factor, E_{it} is the expectation with respect to information available at time t , c_{it} is time t consumption, and $\gamma > 0$ is the coefficient of relative risk aversion. When $\gamma = 1$ we interpret the agents' reward functions as logarithmic. We allow agents in each period to either consume or invest in one of n different assets. We assume there are enough assets to make markets complete. Let the gross real return, including dividends, from holding asset p between periods t and $t + 1$ be $r_{pt} + 1$. For a risk-free asset this return is observed at time t . For a risky asset, the return is unknown at time t but observed at time $t + 1$. The budget constraint of a type i agent at time t is

$$c_{it} + \sum_{p=1}^n \varphi_{ipt} = w_{it} + \sum_{p=1}^n \varphi_{ipt-1} r_{pt} \quad (2)$$

where at time t , the amount invested in asset p is φ_{ipt} , and labor income is w_{it} . An agent's problem is to maximize lifetime utility, Equation (1), subject to the sequence of budget constraints given by Equation (2) for each date t . A constraint to rule out Ponzi schemes is also imposed. The agents choose

³ An earlier version of the decomposition in Section 1.4 that applied to the model of Detemple and Murthy (1994) appeared in Ghysels and Juergens (2001). That version only applied when agents have logarithmic utility and did not link the beliefs of agents to available data. The earlier version was written independently and simultaneously to the work by Shefrin (2001) who reported similar results and related them to the behavioral finance literature. The latter will be further discussed in Section 5. The more general version of the decomposition included in this article was written before we were aware of the work of Shefrin (2001).

consumption and the amount to invest in each asset at each date adapted to information that is known. At time t agents know the value of variables dated t and earlier. They also know the return on risk-free assets from time t to $t + 1$. It is straightforward to show that if the time t price of an asset is one then its random payoff r_{pt+1} at time $t + 1$ satisfies

$$1 = E_{it} \left[\beta \left(\frac{c_{it}}{c_{it+1}} \right)^\gamma r_{pt+1} \right]. \quad (3)$$

This is the fundamental equation for consumption-based asset pricing, and it is one way to write the Euler equation for a type i agent's optimization problem.⁴ This equation holds for any asset using the beliefs and consumption of any agent.

Let λ_{it} be the time t Pareto weight for agents of type i (a formal definition of Pareto weights appears in Appendix A). The Pareto weight measures the importance of type i agents and is related to their wealth. The Pareto weights are constructed so that they sum to one at every date. We will assume that a function of the Pareto weights are constant over time:

Assumption 1. (*Constant Distribution*). *The distribution of agents, as measured by,*

$$\sum_{i=1}^I (\lambda_{it})^{\frac{1}{\gamma}} \quad (4)$$

is constant over time.

This assumption holds exactly if the economy has reached an invariant distribution across agents and always holds exactly when $\gamma = 1$, even if an invariant distribution has not been reached. Note that the Pareto weight of any individual agent can evolve over time; however, the population distribution, as measured by Equation (4), is assumed time invariant. In data, Assumption 1 may not be exactly satisfied. (Indeed it is not exactly satisfied by the data we use to measure the Pareto weights in subsequent sections.) Nevertheless, Assumption 1 is a plausible simplifying approximation over short time intervals when (1) γ is close to one, (2) the economy is near an invariant distribution, and (3) the beliefs of agents are not too disparate. See Blume and Easley (1992) for examples of economies for which it would be unreasonable to assume a constant distribution. It is possible to do the theoretical analysis in this section without imposing Assumption 1. Shefrin (2001) takes this approach. Since we do not have reliable data on the short-term movements of the Pareto weights, it is convenient for our empirical analysis to impose Assumption 1.

⁴ The agent's labor income does not enter this equation and will play no direct role in this article.

Using our notion of Pareto weights, it is shown in Appendix A that the pricing Equation (3) can be written as

$$1 = \sum_{i=1}^I \lambda_{it} E_{it} \left[\beta \left(\frac{c_t}{c_{t+1}} \right)^{\gamma} r_{pt+1} \right]. \quad (5)$$

Equations similar to Equation (5) have been studied by Detemple and Murthy (1994), Basak (2000), and Shefrin (2001) and tells us that security prices are a weighted average, across agents, of the security prices that would obtain when all agents are identical. Equation (5) is in terms of aggregate consumption whereas Equation (3) is in terms of the consumption of a single agent. Equation (5) allows us to state the pricing equation in terms of fundamentals—aggregate consumption and asset returns—and the variables pertaining to individual agents which are embedded in λ_{it} and E_{it} .

1.2 The distribution of an individual's beliefs

In this article, we empirically estimate Equation (5) using the forecasts of analysts as proxies for beliefs of agents. For many of the variables, including asset returns, we face the following situation. We have an historical record of actual values, and we have implied forecasts of conditional expectations. To estimate Equation (5), we need more than this: we need the entire distribution of each agent's beliefs. In this section, we explain a method for approximating the entire distribution by tying the higher order moments of agents' beliefs to the true distribution.⁵

To make things concrete, let x be a vector of positive variables for which we have a time series record of observations. At time t , the actual value of x is denoted x_t . Assume we also have a time series record of conditional expectations of several different agents of the one period ahead value of x : $\mu_{xit+1|t} = E_{it}x_{t+1}$. We have no other direct information about the individual agent's perception of the distribution of x . Since our empirical work requires more information about the individuals' perceived distribution of x , we need to make an assumption about the distribution of an individual's beliefs, namely:⁶

Assumption 2. (Homotopy). *Let x_{t+1} be generated by the conditional cdf $\xi_{xt}(x_{t+1})$ and $\mu_{xt+1|t}$ be the true conditional mean of x_{t+1} given time t information. Let the conditional expectation of agent i be $\mu_{xit+1|t} = E_{it}x_{t+1}$.*

⁵ In an ideal world, we would not only have mean forecasts across analysts for each stock, but we would also have some measure of dispersion of each analysts' beliefs regarding the distribution of forecasts. Graham and Harvey (2003), using survey data of CFOs about the expected risk premium, are able to obtain a distribution of beliefs for each individual respondent. In the future, work along the lines of Graham and Harvey (2003) may eventually yield more data regarding the distribution of beliefs for individual stocks or portfolios.

⁶ We thank the referee for the suggestion to call the following assumption homotopy.

Then the conditional cdf of x_{t+1} given time t information from the point of view of agent i is

$$\xi_{ixt}(x_{t+1}) = \xi_{xt} \left(\frac{x_{t+1} * \mu_{xit+1|t}}{\mu_{xit+1|t}} \right).$$

Here the notation “ $*$ ” denotes element by element multiplication and the division is element by element.

The assumption completely specifies the distribution of individual beliefs in terms of the known conditional mean held by an agent, the unknown true conditional mean, and the unknown true distribution. This assumption entails that if agent i has a more optimistic view of the conditional mean than agent j then agent i is globally more optimistic than agent j in the sense that if $\mu_{xit+1|t} \geq \mu_{xjt+1|t}$ then for any x_{t+1} it is the case that $\xi_{ixt}(x_{t+1}) \leq \xi_{jxt}(x_{t+1})$. Similar to Assumption 1, it is possible to do the theoretical analysis in this section without imposing Assumption 2. Shefrin (2001) takes this approach. Since we do not have data on the distribution of an individual’s beliefs, it is convenient for our empirical analysis to impose Assumption 2.⁷

This specification has a number of appealing additional properties. Under this specification, if an agent has the correct beliefs about the conditional mean then he has the correct beliefs about the entire distribution. That is if $\mu_{xit+1|t} = \mu_{xt+1|t}$ then $\xi_{ixt}(x_{t+1}) = \xi_{xt}(x_{t+1})$, and from the Homotopy Assumption 2 agent i knows the complete distribution of x_{t+1} . This specification easily allows us to transform integrals against $\xi_{ixt}(x_{t+1})$ into integrals against $\xi_{xt}(x_{t+1})$. We have for any function f , when the integrals below exist,

$$\int f(x_{t+1}) d\xi_{ixt}(x_{t+1}) = \int f(x_{t+1}) d\xi_{xt} \left(\frac{x_{t+1} * \mu_{xit+1|t}}{\mu_{xit+1|t}} \right) \quad (6)$$

$$= \int f \left(\frac{y_{t+1} * \mu_{xit+1|t}}{\mu_{xit+1|t}} \right) d\xi_{xt}(y_{t+1}). \quad (7)$$

From this it follows that $\mu_{xit+1|t}$ really is the mean of i ’s beliefs since

$$\int x_{t+1} d\xi_{ixt}(x_{t+1}) = \left(\frac{\mu_{xit+1|t}}{\mu_{xit+1|t}} \right) * \int y_{t+1} d\xi_{xt}(y_{t+1}) = \mu_{xit+1|t}.$$

Appendix B shows that this specification has the appealing property that all agents correctly know the second and higher order conditional central moments of $\log(x)$.

⁷ In this article, the vector x will consist of variables which are gross returns and hence nonnegative by construction. We make the additional assumption that $\mu_{xt+1|t}$ and $\mu_{xit+1|t}$ for all i are strictly positive. We only apply homotopy to the primitive gross returns for which we have data. Assumption 2 does not necessarily hold for constructed gross returns.

1.3 Beliefs about consumption growth

Our data does not provide direct measurements on analysts' perceived conditional means of consumption growth whereas our empirical work requires such information. We construct an approximation, $\mu_{git+1|t}$, to type i agents' mean forecast of the gross rate of aggregate consumption growth by tying the forecast to the forecasted market return. We state the approximation as an assumption.

Assumption 3. *Agent i 's belief about the conditional mean of aggregate consumption growth, $\mu_{git+1|t}$, satisfies*

$$\log \mu_{git+1|t} = q_t + \phi \log \mu_{mit+1|t} \quad (8)$$

where ϕ is a constant, q_t is an unknown function of current information, and $\mu_{mit+1|t}$ is agent i 's forecast of the market return.

We take the time-invariant constant ϕ and the time-varying function q_t to be the same for all agents. Note that $\mu_{git+1|t}$ is agent i 's forecast of aggregate consumption growth which is not necessarily equal to agent i 's forecast of his own consumption growth.

In our empirical work, we consider different values of ϕ . It is interesting to note that ϕ should be equal to one in a general equilibrium version of our model.

1.4 A decomposition

In this section, we decompose expected returns into a fundamental component and a heterogeneity component. We exploit our earlier assumptions to write the pricing Equation (5) as

$$1 = \sum_{i=1}^I \lambda_{it} E_{it} [\beta g_{t+1}^{-\gamma} r_{pt+1}] \quad (9)$$

$$= \sum_{i=1}^I \lambda_{it} E_t \left[\beta \left(g_{t+1} \frac{\mu_{git+1|t}}{\mu_{gt+1|t}} \right)^{-\gamma} r_{pt+1} \frac{\mu_{pit+1|t}}{\mu_{pt+1|t}} \right] \quad (10)$$

$$= \beta h_{pt} E_t g_{t+1}^{-\gamma} r_{pt+1} \quad (11)$$

where $g_{t+1} = c_{t+1}/c_t$ is the gross rate of aggregate consumption growth and

$$h_{pt} = \left[\sum_{j=1}^I \lambda_{jt} \left(\frac{\mu_{gjt+1|t}}{\mu_{gt+1|t}} \right)^{-\gamma} \left(\frac{\mu_{pjt+1|t}}{\mu_{pt+1|t}} \right) \right] \quad (12)$$

is the heterogeneity component of the pricing equation for asset p . Notice that h_{pt} varies across assets. Equation (10) follows from Equation (9) and the result in Equation (6) applied to conditional expectations. Equation (11) follows because the conditional mean forecasts are known at time t .

A simple conditional decomposition follows from rearranging terms in Equation (11) using the identity that $E_t xy = E_t x E_t y + \text{cov}_t(x, y)$. The expected return of asset p is therefore

$$E_t r_{pt+1} = \underbrace{\Lambda_t - \frac{\text{cov}_t(g_{t+1}^{-\gamma}, r_{pt+1})}{E_t g_{t+1}^{-\gamma}}}_{\text{Fundamental component}} + \underbrace{\Lambda_t \left(\frac{1 - h_{pt}}{h_{pt}} \right)}_{\text{Heterogeneity component}}$$

where we have defined

$$\Lambda_t = \frac{1}{\beta E_t g_{t+1}^{-\gamma}} \quad (13)$$

to be the risk-free rate if agents *agreed* on their forecasts of g for all assets. The actual risk-free rate, which we call r_{bt+1} , in our economy is different when agents disagree and is given by

$$r_{bt+1} = \frac{\Lambda_t}{h_{bt}}. \quad (14)$$

Even though all agents agree on forecasts of a risk-free return, there are pricing implications if agents disagree on forecasts of consumption growth.

This decomposition illustrates the effect of heterogeneity of beliefs on expected returns. When there is no disagreement, at all dates, $h_{pt} = 1$ for all t , and the heterogeneity component in the decomposition is zero. When there is heterogeneity, h_{pt} need not equal to one, and there can be pricing implications. The fundamental component of the return is not altered as heterogeneity is increased.⁸ We will revisit the decomposition of returns in Section 5.3.

We have derived the decomposition under Assumptions 1 and 2. It is possible to derive an analogous decomposition without requiring both assumptions. Shefrin (2001) takes this more general approach.

2. Limited Rational Expectations and Dispersion

For now, we assume agents have exactly the same beliefs as analysts. So, if there are I analysts, we assume there are I types of agents and agents of

⁸ This is true from the point of view of an econometrician faced with a fixed data set. In general, since heterogeneity can affect the return of assets it also can affect $\text{cov}_t(g_{t+1}^{-\gamma}, r_{pt+1})$ and hence alter the fundamental component of the return.

type i (for $i = 1, 2, \dots, I$) have the same beliefs as analyst i .⁹ The data used in this article provides us with measures of conditional mean forecasts of stock returns.¹⁰ We construct a conditional mean forecast for consumption growth as described at the end of Section 1.2. We do not have direct observations on the *true* conditional means of consumption growth and stock returns. To apply the Homotopy Assumption 2, we need to take a stand on these true conditional means. We propose a limited form of rational expectations in which the economy as a whole is correct on average about its expectations of consumption growth and stock returns. In this section, we assume the true conditional means of consumption growth and any stock p are¹¹

$$\mu_{gt+1|t} = \sum_{i=1}^I \lambda_{it} \mu_{git+1|t} \quad (15)$$

$$\mu_{pt+1|t} = \sum_{i=1}^I \lambda_{it} \mu_{pit+1|t}. \quad (16)$$

The Pareto weights are used to measure the importance of agents. This assumption is motivated from a desire to include the actual beliefs of agents in asset pricing models while making as few departures as possible from rational expectation models. As we explain at the end of this section, the economy as a whole will not have rational expectations about all variables in the economy and our model will include some features prevalent in the behavioral finance literature.

Recall that Equations (5) and (11) tell us security prices are a weighted average, across agents, of the security prices that would obtain when all agents are identical. The dispersion of beliefs about $g_{t+1}^{-\gamma} r_{pt+1}$ does not matter. Only the mean matters. However, the dispersion of beliefs about g_{t+1} and r_{pt+1} can matter. To see this, interpret the Pareto weights as probabilities and $\hat{E}_t$ and $\widehat{\text{cov}}_t$ as the expectation and covariance operators over these fictitious probabilities (so, for example, $\hat{E}_t \mu_{pit+1|t} = \sum_{i=1}^I \lambda_{it} \mu_{pit+1|t}$). We can decompose h_{pt} as

$$h_{pt} = \hat{E}_t \left[\left(\frac{\mu_{git+1|t}}{\mu_{gt+1|t}} \right)^{-\gamma} \right] \hat{E}_t \left(\frac{\mu_{pit+1|t}}{\mu_{pt+1|t}} \right) + \widehat{\text{cov}}_t \left[\left(\frac{\mu_{git+1|t}}{\mu_{gt+1|t}} \right)^{-\gamma}, \frac{\mu_{pit+1|t}}{\mu_{pt+1|t}} \right] \quad (17)$$

⁹ Throughout this article, unless explicitly said otherwise, we use “analyst” to refer to the fictitious analyst constructed from the views of all analysts in a particular brokerage house.

¹⁰ The computations of the conditional mean forecast of stock returns is described in Section 3.2.

¹¹ We assume that agents do not use the fact that the analysts are correct on average when formulating their beliefs.

$$= \hat{E}_t \left[\left(\frac{\mu_{gjt+1|t}}{\mu_{gt+1|t}} \right)^{-\gamma} \right] + \widehat{\text{cov}}_t \left[\left(\frac{\mu_{gjt+1|t}}{\mu_{gt+1|t}} \right)^{-\gamma}, \frac{\mu_{pj+1|t}}{\mu_{pt+1|t}} \right] \quad (18)$$

where the second line follows from Equation (16). If the co-dispersion in the last term increases then h_{pt} will increase if the first term remains constant. So, if the mean of $\mu_{gjt+1|t}^{-\gamma}$ and $\mu_{pj+1|t}$ remain constant h_{pt} can still increase. Likewise, depending upon the value of γ , it is possible that if the mean of $\mu_{gjt+1|t}$ remains constant, $\hat{E}_t \left[\left(\mu_{gjt+1|t} / \mu_{gt+1|t} \right)^{-\gamma} \right]$ can increase as dispersion in consumption growth increases.¹²

One could object to our formulation because if we assumed agents are correct on average about the conditional mean of $g_{t+1}^{-\gamma} r_{pt+1}$ then dispersion would not matter. Furthermore given we assumed in Equations (15) and (16) agents are correct on average about the conditional means of asset returns and consumption growth it might seem natural to assume they also are correct about the conditional mean of $g_{t+1}^{-\gamma} r_{pt+1}$. However this objection is not valid. If we assume agents are correct on average in their expectations of asset returns and consumption growth, then, in general, under Assumption 2, agents cannot have the correct expectations about $g_{t+1}^{-\gamma} r_{pt+1}$, when they *disagree* about the expectations of asset returns and consumption growth. This is what Equation (18) tells us.

Hence in our model, the economy as a whole, has a limited form of rational expectations. It has rational expectations about simple variables like asset returns and consumption growth but possibly irrational expectations about more complicated variables like $g_{t+1}^{-\gamma} r_{pt+1}$. This feature is also present in a number of behavioral finance models including the early work of Simon (1955) and Tversky and Kahneman (1974). Using this work, Hirshleifer (2001) argues that biases can be viewed as outgrowths of heuristic simplification. One can indeed argue, as in Equations (15) and (16), that agents are correct on average about the conditional means of asset returns and consumption, but may not be able to capture properly expectations of complex processes such as those that enter the stochastic discount factor. In reality agents' expectations of asset returns and consumption growth are probably not correct on average. If we allowed for this then including heterogeneous beliefs about asset returns and consumption growth would have two effects. One effect comes from the dispersion of beliefs, and another effect comes from errors in the conditional mean of beliefs about asset returns and consumption growth. In this section, we zero out the second effect and focus only on the dispersion effect. In Section 5, we allow agents (but not analysts) to have biased expectations of stock returns.

¹² Since consumption growth is an endogenous variable chosen by agents, an increase in dispersion of consumption is interpreted as increase in the dispersion of the information available to agents which in turn creates dispersion in their optimal choices.

3. Measuring Beliefs

Thus far, we studied the theoretical properties of models with heterogeneous beliefs; however, our ultimate goal is to empirically investigate the role played by the heterogeneity of beliefs. We follow two different empirical strategies. We first examine whether the heterogeneity of beliefs is a priced factor in traditional asset pricing models. This is a reduced form approach and requires the construction of factors that reflect heterogeneous beliefs. The second empirical strategy consists of estimating structural equations, based on the models presented in the previous sections, while being careful about measuring agents' actual beliefs.

In this section, we describe the data used for the various empirical exercises. Both empirical approaches require different data collections and constructions. We construct a factor specification for short-term earnings and long-term earnings growth forecasts (measured as the standard deviation of month-end forecasts) and examine the relevance of these factors in standard asset pricing models. Similarly, based on analyst earnings predictions we construct expected returns which are necessary to estimate fundamental pricing equations.

3.1 Constructing heterogeneity factors

We begin by acquiring Fama–French factors for a sample period that coincides with the heterogeneity factor. Monthly firm returns (denoted r for a generic firm), prices, and shares outstanding are collected from the Center for Research in Security Prices (CRSP) monthly tapes. Data to compute book value of equity is collected from the annual Compustat tapes. We also obtain the risk-free interest rate (r_b), monthly estimates of size, book-to-market, and momentum factors (SMB, HML, and UMD, respectively) as well as market returns (r_m).¹³ Tables 1 and 2 provide summary statistics for returns and explanatory variables. The average monthly firm and market excess returns of 1.23 and 1.16% are substantially larger than those computed by Fama and French (1993) of 0.67 and 0.43%, respectively.¹⁴ These results are not surprising given the time period of analysis (1991–1997) primarily occurred during up markets (market returns were positive in 61 out of 84 months). The average monthly size factor (SMB) is 0.14%, the book-to-market factor (HML) averages 0.47%, and the momentum factor (UMD) is 0.78% per month. These results are similar to those found in other research that use substantially longer time series.

¹³ This data along with a description of how the factors are formed is available on Ken French's Web site.

¹⁴ The average returns reported within are more closely aligned with those found by Griffin (2002) who also examines a later time period (1981–1995) than Daniel and Titman (1997) and Fama and French (1993).

Table 1
Descriptive statistics for Standard & Poor's 500 stocks (1991–1997)

Panel A: Summary statistics of S & P 500 stocks

	Mean	SD	Minimum	Q1	Median	Q3	Maximum
$r - r_b$	1.23	0.18	0.98	1.07	1.23	1.37	1.53
Market capitalization	6763	7732	657	2007	3602	6376	26147
BE/ME (%)	0.57	0.31	0.18	0.36	0.52	0.71	1.19
Disp. ST (\$)	0.20	0.41	0.00	0.03	0.08	0.21	8.64
Analysts ST	15	7	2	10	15	20	37
Disp. LT (%)	4.14	4.34	0.00	1.95	3.16	4.90	129.27
Analysts LT	12	5	2	8	12	15	33

Panel B: Comparison of high- and low-dispersion firm characteristics

	Short-term			Long-term		
	Low	High	<i>t</i> -statistics	Low	High	<i>t</i> -statistics
$r - r_b$	1.76	1.51	2.98	1.64	1.64	0.02
Market capitalization	8546	6686	4.02	8876	6392	5.39
BE/ME	0.49	0.66	−12.87	0.54	0.60	−4.84
Institutional ownership	0.58	0.60	−3.72	0.57	0.62	−6.93
Analysts	15	16	−3.07	13	13	−0.41

Panel C: Annual switching between high- and low-dispersion portfolios

	1992	1993	1994	1995	1996	1997
Disp. ST	39.60%	29.95%	31.83%	32.48%	31.79%	28.75%
Disp. LT	19.49%	17.61%	16.20%	12.05%	14.68%	14.02%

This table provides summary data for the Standard & Poor's (S & P) 500 Index firms between 1991 and 1997. Panel A provides summary statistics. Excess returns ($r - r_b$), market capitalization, a proxy for firm size, and the book-to-market ratio (BE/ME) are averaged cross-sectionally across 25 portfolios formed on size and book-to-market [Fama and French (1993)]. The dispersion of short-term (Disp. ST) and long-term earnings forecast (Disp. LT), and the number of analysts reporting short-term (Analysts ST) and long-term forecast (Analysts LT), are calculated annually for individual firms. Summary statistics are based on time series averages across the sample period. Panel B provides a comparison of average returns, size, book-to-market ratios, institutional ownership and the number of analysts between high and low-dispersion portfolios. Panel C examines the percentage of firms in each year that switch from high to low dispersion (or vice versa), where high-dispersion firms are those with average dispersion above the median. *P*-values are in parentheses.

We want to construct factors that pertain to heterogeneity in the same fashion. For the remainder of this subsection, we discuss our methodology.

3.1.1 Why analyst dispersion? We investigate whether the heterogeneity of beliefs is a priced factor in traditional asset pricing models. We argue that factors constructed from the disagreement among analysts about expected (short-term and long-term) earnings are good proxies for quantifying heterogeneous beliefs. Other proxies have been proposed. For example, as support for Harris and Raviv (1993) and Shalen (1993), Graham and Harvey (1996) find dispersion among newsletter “forecasts” is positively related to historical volatility, implied (or expected) volatility, and volume. Alternatively, Bessembinder, Chan, and Seguin (1996) examine the open interest on the Standard & Poor's (S & P) 500 Index futures as

Table 2
Descriptive statistics for explanatory factors

Panel A: Summary statistics of Fama–French, momentum, and dispersion factors, $N = 84$
Autocorrelations

	Mean	SD	<i>t</i> -statistics	Autocorrelations		
				1	2	12
$r_m - r_b$	1.16	3.09	3.46	−0.21	0.11	−0.06
SMB	0.14	2.68	0.47	0.10	−0.09	0.10
HML	0.47	2.44	1.75	0.21	−0.06	0.13
UMD	0.78	2.52	2.84	0.04	0.01	0.12
DISP	−0.28	1.44	−1.75	0.14	0.00	0.16
LTGDISP	0.23	2.27	0.96	0.09	−0.07	0.11

Panel B: Cross-sectional correlation coefficients: $N = 84$

	$r_m - r_b$	SMB	HML	UMD	DISP
SMB	0.0203 (0.8543)				
HML	−0.4420 (0.0001)	−0.3049 (0.0048)			
UMD	0.1662 (0.1308)	−0.2108 (0.0543)	−0.0136 (0.9023)		
DISP	0.1120 (0.3103)	0.3574 (0.0008)	0.1159 (0.2938)	−0.3546 (0.0009)	
LTGDISP	0.2495 (0.0221)	0.5052 (0.0001)	−0.3511 (0.0011)	−0.2263 (0.0385)	0.4913 (0.0001)

Summary data on the independent explanatory factors is provided in this table. The factors include the Fama and French (1993) risk factors: $r_m - r_b$ (the systematic risk factor of the excess market return over the risk-free rate), SMB (the risk factor for small firms over large firms), and HML (the risk factor for high book equity to market equity firms over low book-to-market firms). A factor to capture momentum, UMD, is also included. Risk factors to capture heterogeneity of beliefs, DISP (the risk factor between firms with high and low dispersion of earnings forecasts) and LTGDISP (the risk factor between firms with high and low dispersion of long-term earnings growth forecasts), are also presented. Panel A describes summary statistics (in returns) which are pooled monthly from January 1991 to December 1997. Autocorrelations through lag 12 are also included. Panel B provides pairwise Pearson correlation coefficients for the explanatory factors. *P*-values are in parentheses.

a measure of disagreement in opinion and link it to trading volume. Disagreement among analysts is a more appealing source to extract heterogeneity. Unlike other measures, such as short-sale constraints, breadth of ownership, open interest, and newsletter forecasts, dispersion measures are easily available for a large number of equities.

Diether, Malloy, and Scherbina (2002) examine short-term earnings forecast dispersion in predicting future stock returns and relate their findings to the short sale constraint literature, where increased disagreement among investors leads to lower future returns. Unlike their work, we explicitly construct a factor specification for short-term and long-term forecast dispersion, which allows us to measure the incremental contribution of dispersion in factor asset pricing models and sets the stage for out-of-sample predictive properties of these factors, which is detailed in Section 4. Before providing the details, we note that the idea to use the dispersion among analyst forecasts as a measure of heterogeneous beliefs is found in a number of articles. For instance, it has been shown that contemporaneous dispersion of analysts' short-term forecasts is significantly and positively related to volatility [see, for example, Ajinkya and Gift (1985), L'Her and Suret (1996), and Lobo and Tung (1998)].

Our analysis extends the existing literature of the role of analyst dispersion on returns by distinguishing fundamental factor contributions from factors for the heterogeneity of beliefs and by explicitly testing a decomposition of asset pricing returns and volatility that ties into theoretical models. Moreover, we focus more explicitly on return and volatility dynamics by examining out-of-sample performance. This safeguards us from data mining issues, as we construct a pricing factor that captures the heterogeneity of beliefs and test its predictive ability for out-of-sample returns and volatility.¹⁵

Although dispersion of analyst forecasts is a reasonable measure of heterogeneous beliefs, there are several issues, including biases and timing issues that should be addressed. There is an extensive literature that examines the value of analysts' earnings forecasts [(Daniel, Hirshleifer, and Teoh (2002))], and substantial evidence exists that analysts are biased in their forecasts. Analysts are generally optimistic about annual and longer-term forecasts [(De Bondt and Thaler (1985, 1987, 1990), LaPorta (1996), Dechow and Sloan (1997), Capstaff, Paudyal, and Rees (1998), Brown (2001))], but there is some evidence that as the forecast period declines, analysts become slightly pessimistic in their forecasts [O'Brien (1988), Brown (2001), Matsumoto (2002)]. Moreover, analysts tend to overreact to positive news and underreact to negative news [Easterwood and Nutt (1999)]. There are many reasons for this persistent bias, including agency issues [Rajan and Servaes (1997), Michaely and Womack (1999)] and behavioral biases [Barberis, Shleifer, and Vishny (1998), Daniel, Hirshleifer, and Subrahmanyam (1998)]. It also appears that investors do not fully correct for biases.¹⁶ Although we do not dispute this evidence, we should note that our analysis is based on portfolios rather than individual stocks, an issue we will revisit in Section 6.

Whatever the verdict about biases, we focus on *dispersion* among analysts. In the event forecasts are biased, we take the dispersion or deviation from the consensus forecast.¹⁷ Stocks with little disagreement, though perhaps biased in the mean, are treated differently from stocks where analysts have widely different views. However, focusing on dispersion is not entirely immune to bias issues.

For example, Lin and McNichols (1998) argue that analysts are more likely to discontinue coverage of firms about which they have unfavorable

¹⁵ Qu, Starks, and Yan (2003) construct factors for both short-term earnings forecast dispersion which impacts returns through differing private information as well as the volatility of that dispersion which captures uncertainty about firm fundamentals. However, they argue that rather than being "new" factors, dispersion captures elements of other well-documented, but often less easily explained, risk factors such as size and book-to-market.

¹⁶ Womack (1996) documents price drift measurable up to six months following a recommendation change.

¹⁷ Given the documented biases that exist in analyst forecasts, we will explicitly test this hypothesis. This material is deferred to Section 6 as we only deal with dispersion here.

opinions. Therefore, analysts that predict extremely negative earnings forecasts may disappear from the sample, which may reduce the “true” dispersion of analysts’ beliefs. Moreover, analysts are also under no obligation to update forecasts, even as new information arrives. The potential problem of staleness in forecasts and dispersion arises. However, we cannot determine from our data whether analysts insufficiently update their forecasts or if the forecasts are not updated because the beliefs of the analysts remain unchanged. We err on the side that analysts do not update because their beliefs remain unchanged and use all available forecasts in the construction of dispersion.¹⁸

A final note regarding sample selection is in order. Prior studies that examine the empirical implications of fundamental factors on returns and volatility have primarily analyzed relatively long time horizons for estimation and use large cross-sectional samples of firms [see, for example, Chan, Karceski, and Lakonishok (1999), Fama and French (1993), and Griffin (2002)]. However, because of the constraints on the availability of our measures of dispersion, long time series of data cannot be constructed. We examine the S & P 500 Index between 1991 and 1997 because many small firms had little or no analyst coverage in the early First Call data, and adequate analyst coverage is needed to create meaningful dispersion measures. Although the results obtained may be criticized as being sample specific given the intentional design of the sample and the short time period for analysis, we formulate tests similar to those of previous studies [specifically, Section 4.2 of Fama and French (1993)] for comparison purposes.

3.1.2 Analyst forecasts, dispersion, and factor specifications. We produce two measures of heterogeneous beliefs from the First Call/Thomson databases: dispersion of short-term (one-year ahead) dollar earnings forecasts and dispersion of long-term (five-year ahead) earnings growth rate forecasts, measured as the standard deviation of forecasts. Only the last available dispersion measure in each month is used. The average monthly level of short-term (long-term) forecast dispersion is \$0.20 (4.14%), with approximately 15 (19) analysts furnishing short-term (long-term) forecasts per firm in the S & P 500 Index, which is substantially higher than the average of three analysts per firm for the entire First Call database.

Panel B of Table 1 examines the differences in firm characteristics between high and low-dispersion firms. Similar to Diether, Malloy, and Scherbina (2002) and Qu, Starks, and Yan (2003), we find a negative relation between dispersion and returns. Confirming Crombez (2001) and

¹⁸ We believe this is reasonable considering that one metric used in ranking analysts is the timeliness and accuracy of forecasts. Stickel (1992) shows compensation is directly linked to rankings, thus forecasts should be updated as beliefs about valuation change.

others, we find larger firms tend to have less disagreement about earnings expectations, as do growth [low book-to-market ratio (BE/ME)] firms. Interestingly, we find that there is greater institutional participation and a larger number of analysts for high-dispersion firms, which is directly in contrast to Crombez (2001) and the short-sale restriction literature (Chen, Hong, and Stein 2000), which predicts that as disagreement increases, the breadth of ownership should decrease.

In Figure 1, we display average levels of dispersion. As noted, analyst optimism tends to decline as the fiscal year draws to a close. Although we do not actually make a comparison to realized earnings, we find some evidence that forecasts decline throughout the fiscal year for short-term forecasts. However, we do not find any evidence of seasonal patterns in long-term forecasts.

We also examine the incidence of “switching” from low- to high-dispersion firms across months. Panel C of Table 1 displays that for short-term dispersion approximately 30% of firms switch from high dispersion to low dispersion or vice versa in a given year. (Monthly switching is slightly less.) Approximately 15% of firms characterized by long-term dispersion switch per year.

In order to be comparable to the market, size, book-to-market, and momentum factors, we construct dispersion factors for both the short-term and long-term forecasts. Dispersion factors include all 500 firms in the S & P 500 Index, which are updated annually to reflect changes to the index. In each month, the dispersion is ranked, and the median value is used as the breakpoint to categorize high and low-dispersion forecasts.¹⁹ Value-weighted returns are then calculated each month for high- and low-dispersion portfolios. A zero investment strategy is realized as we purchase high-dispersion stocks and short low-dispersion stocks [see, for example, Chan, Karceski, and Lakonishok (1999) for the construction of mimicking portfolios]. The two factors examined are DISP and LTGDISP for short-term and long-term forecasts, respectively. Merton (1987) posits that dispersion represents a compensation to investors for the idiosyncratic risk from holding undiversified portfolios. Since dispersion implies higher variation in earnings streams, stocks with high dispersion should earn larger future returns.²⁰

As reported in Table 2, the average monthly return for the short-term dispersion factor, DISP, is -0.28% , whereas the average return for LTGDISP is 0.23% per month. Hence, LTGDISP appears to be consistent with the Merton hypothesis that dispersion proxies for idiosyncratic risk, whereas DISP is not. The market factor and UMD are highly

¹⁹ Womack (1996) documents price drift measurable up to six months following a recommendation change.

²⁰ Other breakpoints, including quintiles and deciles were examined. The results are robust to various specifications of high and low dispersion.




Figure 1
Average level of dispersion throughout the fiscal year

The top figure displays the average monthly dispersion of short-term (one-year ahead) analyst earnings forecasts in dollars for Standard & Poor's (S & P) 500 Index firms. The bottom figure shows the average monthly long-term (five-year average growth rate of earnings) dispersion. Average dispersion is computed as the equal-weighted average of S & P 500 firm's dispersion in each month of a firm's fiscal year (where month 12 is the final month in the fiscal year). Results are reported on an annual basis as well as averaged across the full sample (1991–1997).

statistically significant, and HML and DISP are marginally significant at the 10% level; however, the average returns for the SMB and LTGDISP factors are insignificantly different from zero. Table 2 also provides Pearson correlation coefficients for relations between the explanatory

factors. Unlike Fama and French (1993), we do not observe a significant relation between the market and the size factors; however, HML is significantly negatively correlated with both $r_m - r_b$ and SMB, which is consistent with the findings of Fama and French. Although positively correlated with both the market factor and HML, DISP is significantly positively correlated at the 1% level only with SMB, indicating that dispersion may be more important for smaller firms. LTGDISP is positively correlated with $r_m - r_b$, SMB, and DISP, and negatively correlated with HML. Both DISP and LTGDISP are significantly negatively correlated with UMD. Returns for both dispersion factors tend to be higher when small cap stocks outperform large cap stocks; however, we observe DISP (LTGDISP) returns are higher (lower) when value stocks outperform growth stocks. Overall, these findings suggest traditional factors may be encapsulating, as anticipated, some of the same risk measures dispersion factors have captured.

3.2 Construction of expected returns

In order to construct expected return forecasts to be used in the aforementioned structural models, we collect fiscal short-term earnings estimates, long-term earnings growth forecasts, and investment recommendations from First Call/Thomson databases. Brokerage houses are not required to follow a schedule, such as that for macroeconomic forecasts, when updating their investment recommendations or earnings forecasts. Therefore, there is an unevenness when forecasts occur in a given month. For example, an analyst may update an earnings forecast for IBM three times during July 1995 and then will not update again until October 1995. We take two steps to correct this. First, for each firm, we select only the last forecast in a given month to be the month-end forecast, regardless of where it may fall in the month. If there is no forecast within a given month, we carry forward the previous month's forecast until the analyst issues the next new forecast.

Our study focuses on the S & P 500 firms included in the Index between 1991 and 1997 (updated annually). Although the S & P 500 firms are among the largest firms listed in US markets, brokerage houses are under no obligation to cover any particular firm. Thus, although some of the brokerage houses cover at least a subsample of S & P 500 firms between 1991 and 1997, we examine six bulge-bracket brokers that covered identical monthly subsets of S & P 500 firms between April 1994 and December 1997 (45 monthly observations).²¹ Unlike the factor composition which includes all 500 firms in the Index, the number of firms that jointly

²¹ The six brokerage firms that we include are Credit Suisse First Boston, Lehman Brothers, Paine Webber, Prudential Bache, Salomon Brothers, and Smith Barney.

spans the six brokerage houses in each month ranges from 50 to 83 firms, with the mean (median) number of firms as 68 (66).

Analysts do not forecast future returns per se. Recent work by Brav, Lehavy, and Michaely (2003) employs analysts' price target estimates to approximate *ex ante* expected returns; however, price target data are unavailable during the time period we examine. Instead, given our data limitations, we rely on analysts' earnings forecasts and investment recommendations. In order to extrapolate return forecasts from analysts, we implement a modified constant dividend growth model.²² Although this model relies on expected future dividends to back out the market capitalization rate (expected future returns), there are several issues that make this model difficult to use. First, not all firms pay dividends. Second, although we have today's dividend in dollars, analysts do not predict future dividends or the expected growth of those dividends. Third, firms have the option to suspend dividend payments in times of financial trouble. Thus, although our measure of the forecasted returns is not exactly that intended by the constant dividend growth model, we believe that it is a fairly accurate approximation since all firms report earnings and most firms are covered by at least one brokerage firm that provides earnings forecasts. Thus, we substitute the next fiscal year's earnings forecast in for the expected future dividend and substitute the long-term earnings growth forecast for the growth rate of dividends.

Assumption 4. *At time t , brokerage house i 's belief about the expected gross return on stock p between periods t and $t + 1$ is*

$$\bar{\mu}_{pit+1|t} = \left[1 + \frac{\Delta_{ipt}}{Q_{pt}} + \delta_{ipt} \right]^{\frac{1}{12}}$$

where for stock p and brokerage house i , Δ_{ipt} is the one-year ahead fiscal year net income per share forecast, Q_{pt} is the current-month price, and δ_{ipt} is the five-year earnings growth forecast.

Δ and δ change when a brokerage house updates its short-term or long-term earnings forecasts; however, the expected return (or market capitalization rate) is updated at the end of each month. The beliefs of brokerage house i are the combined beliefs of all analysts at brokerage house i . In our data set for a given brokerage house i there is at most one analyst who covers each stock.²³ In subsequent sections, we often refer to the

²² The justification of using a modification of the constant dividend discount model can be found in Brav, Lehavy, and Michaely (2003); Crombez (2001); Gebhardt, Lee, and Swaminathan (2001); and Guay, Kothari, and Shu (2003); among others.

²³ One reason for this is that customers might become confused if different analysts within a brokerage gave different predictions for the same stock.

beliefs of brokerage house i as being the beliefs of a fictitious analyst that we refer to as analyst i .

Three series of returns are constructed based on short-term dispersion, long-term dispersion, and historical volatility (based on a 60-month rolling sample of historical standard deviation of price returns). In each month, the firms are ranked from lowest to highest for each of the three series. For example, in April 1994, the percentage short-term dispersion (broker-weighted average dispersion divided by month-end price) is computed for each firm. Rather than equally weighting each brokerage firm, we rank-weight each brokerage firm in each year according to the ranking of their analysts in the *Institutional Investor*. All American Research Team poll in the prior year. The weights are normalized to sum to 1 in each year. These percentage dispersions are ranked, and we separate firms into “high” or “low” dispersion based on where the firm falls relative to the median of the series.²⁴ Firms are then size-weighted within each category each month.

We also use these weights to proxy for the Pareto weights of agents.

Assumption 5. *The number of types of agents, I , is equal to the number of brokerage houses. The Pareto weight of agents of type i is determined by the ranking of analysts who work in brokerage house i .*

Throughout this article, we link agents of type i with brokerage house i . Assumption 5 says that brokerage houses with higher quality analysts correspond to agents who have more total wealth.²⁵

4. The Dispersion of Beliefs as a Factor

Although the market, size, and book-to-market factors are designed to capture the component of expected returns and volatility related to fundamental factors, a major focus of this article is whether a factor that captures the heterogeneity of investors’ beliefs improves upon either the factor models or the predictive ability of these models for out-of-sample returns or volatility.

In Section 1.4, we observed that expected returns can be decomposed into two components: one component determined by traditional factors and the other component dependent on the heterogeneity of beliefs. The first empirical assessment of heterogeneity will exploit this decomposition. In particular, consider Hansen and Richard (1987) and subsequent work, and suppose one projects the unobserved stochastic discount factor

²⁴ Medians, rather than upper and lower quartile rankings, are used because of the limited size of the sample in each month.

²⁵ Ideally, the Pareto weights should be based on consumption or wealth shares of agents. Unfortunately data on assets under management are not readily obtainable nor comparable among the various institutions we include in our study.

onto a set of asset prices. Namely, we replace s_{t+1} in the fundamental pricing equation by the projection $s_{t+1}^f = \text{Proj}(s_{t+1}|f_t)$, onto a set of “traditional factors,” f_t , that can be viewed as proxies for marginal utility growth. Besides the “traditional factors,” that is the Fama–French factors, we use a dispersion factor that relates to heterogeneity. We proceed in two steps, namely first project returns onto the Fama–French factors and then project the pricing error onto the dispersion factor. This two-step strategy allows the traditional factors to explain as much as possible the dependence of the stochastic discount factor on heterogeneity.

In this first empirical assessment, we do not only limit our attention to excess returns, but also examine return volatility. One might expect the same factors that explain returns also predict return volatility. Chan, Karceski, and Lakonishok (1999) employ various factor models and compare each model’s forecasts of future volatility and covariance.²⁶ They find that the most simple models, in particular, the one-factor market model, have smaller forecast errors than more complex multi-factor models in predicting variances, indicating (1) factors predicting the mean return do not predict volatility and (2) higher-dimensional models may overfit the data. When we include a factor for dispersion of beliefs in our estimates of out-of-sample volatility, the forecast error in volatility models is substantially reduced, confirming again the theoretical predictions of the decomposition put forward.

Before moving on, it is worth pointing out that our analysis is based on conventional asset pricing models where no other frictions are present apart from the heterogeneity of beliefs. The *reduced form* approach taken in this section does not preclude alternative interpretations, many which rely on some type of market friction.²⁷ For example, since Miller (1977) there have been attempts to resolve the issue of short sales regulations and divergence of opinions. When there are short sale constraints on assets, an increase in the divergence of opinions leads to lower subsequent returns. Chen, Hong, and Stein (2000) use the breadth of ownership as a measure of dispersion and find that decreases in breadth of ownership (an increase in divergence of opinion) leads to lower subsequent returns. In a separate line of research, Easley, Hvidkjaer, and O’Hara (2000) utilize market microstructure principles and create a measure of private information. They find investors need to be compensated for the risk of holding firms with greater information asymmetries between investors as the

²⁶ The factors they examine include the Fama–French factors (Fama and French 1993), technical or past-return factors [Chan, Jegadeesh, and Lakonishok (1996), Jegadeesh and Titman (1993)], macroeconomic factors [Fama and Gibbons (1984)], principal component or statistical factors [Connor and Korajczyk (1988)], and the return on the market portfolio.

²⁷ When we move to the empirical investigation of structural asset pricing models our analysis will be more discriminatory with regards to such alternative interpretations.

probability of private information increases.²⁸ These models suggest that neither short sale constraints nor private information asymmetries can be diversified away and are likely reflected in measures of dispersion of beliefs. They represent a form of systematic risk; therefore, investors should be compensated for that risk.

4.1 Construction of factor models, forecasts of returns, and volatility

Consider the return process, which we shall denote r_t , and the linear projection:

$$r_{t+1} = E_t^L[r_{t+1}|f_t] + \eta_t^f = d'f_t + \eta_t^f \quad (19)$$

where E^L stands for the linear projection and assume the Fama–French set of factors are being used. As noted before, such a specification can be extended to include other factors as well, namely those pertaining to heterogeneous beliefs. The main motivation for the empirical work reported in this section is to find empirical factors, h_t , that are related to the residual variation of Equation (19) namely:

$$\eta_t^f = d'h_t + \eta_t. \quad (20)$$

In time series regressions of factor models, the slopes and R^2 s are direct evidence of whether pricing factors capture common variation in stock returns. More specifically, in the fixed parameter case with excess returns the intercept provides a measure of the pricing error of each model whereas the slope coefficients give the factor loadings of each stock i on each of the factors $r_m - r_b$, SMB, HML, UMD, DISP, and LTGDISP.

For the purpose of empirical testing we construct 15 factor models that incorporate various combinations of the six factors. Model 1 is a standard one-factor capital asset pricing model (CAPM) model that uses the excess return on the value-weighted market index. Model 2 supplements the excess market return with size and book-to-market factors (Fama and French 1993). Model 3 represents the Carhart (1997) four-factor model, namely Fama–French, plus a factor which captures momentum. Models 4 and 5 examine the excess market return in conjunction with the short-term and long-term dispersion factors, respectively. Models 6 and 8 combine the three-factor model with short-term and long-term dispersion, whereas Models 7 and 9 augment Models 6 and 8 with the momentum factor. Models 10 and 11 examine the individual power of the dispersion factors, whereas Models 12 and 14 combine the dispersion factors with size and book-to-market and Models 13 and 15 supplement with momentum.

²⁸ The latter is also closely related to the work of Amihud and Mendelson (1986) and numerous subsequent articles that try to test whether less liquid assets have a premium and are held by a different clientele, which have longer investment horizons. Here heterogeneity pertains to investment horizons, not beliefs, and therefore is a separate issue.

Models 10–15 exclude the presence of the market factor, which Fama and French show to be significantly related to expected returns, in order to isolate the impact of the other factors.²⁹ In Section 4.2, we examine the in sample predictive power of the dispersion factor on excess returns.

We next create estimations of equity returns and volatility employing the methodology similar to that used in Chan, Karceski, and Lakonishok (1999). Using a rolling sample of return data over the previous 60 months as our estimation period, we obtain the estimated sensitivities, the mean and variance of each factor, the covariances between factors, as well as residuals for return estimates and squared residuals for volatility estimates. The factor model described in Chan, Karceski, and Lakonishok (1999) is utilized to obtain forecasts of monthly returns and volatility. Monthly out-of-sample returns and volatility are constructed using a 12-month rolling sample starting in the month following the factor measurements. Out-of-sample tests are presented in Section 4.3.

4.2 Is dispersion of beliefs a priced factor?

We examine the role of dispersion as a factor in various factor models using individual security returns as well as portfolio returns based on size and book-to-market ratios (see Fama and French (1993) for portfolio formation details). Table 3 contains estimates of in-sample time-series regressions for a value-weighted portfolio of all stocks in the S & P 500 Index, constructed in each of the 84 months of the sample period. Only the nine factor models which contain the market factor are presented in Table 3. Although the dispersion factors are highly significant in Models 10–15, we do not include these models in Table 3 because of the very large pricing errors.³⁰ We obtain the intercept (the pricing error of the model), coefficients for each factor and relevant *t*-statistics, and adjusted R^2 s for each of the time-series regressions.

Both measures of the dispersion factor are highly statistically significant, and the coefficients are positively related to the S & P 500 Index returns as summerized in Table 3. The market factor is the dominant factor, although each of the other factors is generally statistically significant. The size factor is positively and significantly related to excess returns; however, the

²⁹ Although the CAPM is the standard benchmark for asset pricing models, standard asset pricing models based on the CAPM do not do very well at predicting returns or volatility out of sample. Therefore, we include models which exclude the market in an attempt to determine whether nonmarket based factors have superior predictive ability for either returns or volatility.

³⁰ Fama-MacBeth regressions are also performed. Although the parameter estimates and their statistical significance are quantitatively similar to those presented in Table 3, the fit of the models is substantially reduced (adjusted R^2 s are approximately 13%). Results are available upon request. Individual portfolios based on size and book-to-market breakpoints as well as the aggregate of these portfolios are also examined. In the aggregate, only the market factor is significant in the models 1–9, whereas models 10–15 indicate that dispersion factors are positively and significantly related to portfolio excess returns and HML is negatively and significantly related to portfolio excess returns. Moreover, time series regressions that utilize the original Fama-French factors provide results similar to those found in Fama and French (1993), implying that the results obtained are not likely sample specific and dispersion may be a priced risk factor for other firms. Individual portfolio results are also available upon request.

Table 3
Regressions of excess returns on explanatory factors

Model	Intercept	$r_m - r_b$	SMB	HML	UMD	DISP	LTGDISP	Adjusted R^2
1	0.06 (0.73)	1.03 (40.42)						0.9516
2	-0.07 (-0.88)	1.08 (42.04)	0.10 (3.71)	0.14 (4.06)				0.9614
3	0.01 (0.14)	1.10 (47.98)	0.08 (3.10)	0.14 (4.60)	-0.13 (-4.94)			0.9701
4	0.15 (1.99)	1.02 (45.92)				0.26 (5.39)		0.9639
5	0.06 (0.74)	1.01 (40.68)					0.12 (3.36)	0.9570
6	0.04 (0.45)	1.05 (42.63)	0.06 (1.90)	0.10 (2.83)		0.19 (3.69)		0.9667
7	0.06 (0.86)	1.08 (45.83)	0.05 (1.94)	0.11 (3.54)	-0.10 (-3.92)	0.12 (2.37)		0.9718
8	-0.08 (-1.07)	1.06 (43.20)	0.06 (1.90)	0.15 (4.73)			0.13 (3.39)	0.9659
9	-0.01 (-0.09)	1.08 (48.07)	0.05 (1.69)	0.15 (5.12)	-0.11 (-4.47)		0.09 (2.79)	0.9725

This table provides time-series regression coefficients of excess returns on explanatory factors for a value-weighted portfolio of the S & P 500 firms from January 1991 to December 1997 ($N = 84$). The explanatory factors include the excess market factor, the size factor, the book-to-market factor, the momentum factor, the short-term dispersion factor and the long-term dispersion factor. The general form of the model follows from equations (19) and (20). t -statistics in parentheses and adjusted R^2 's are reported for each model. Model 1 is a one-factor model that uses the excess return on the value-weighted market index, which corresponds to the standard capital asset pricing model (CAPM). Model 2 supplements the excess market return with size and book-to-market factors (Fama and French 1993). Model 3 represents the Carhart (1997) four-factor model, which includes the market, size, and book-to-market factors, plus a factor which captures momentum. Models 4 and 5 examine the excess market return in conjunction with the short-term and long-term dispersion factors, respectively. Models 6 and 8 combine the three-factor model with short-term and long-term dispersion, whereas Models 7 and 9 augment Models 6 and 8 with the momentum factor.

magnitude and significance of the size factor decreases when either of the dispersion factors are included. Moreover, in unreported tests of individual portfolios, we observe that dispersion is significantly positive for firms in the bottom two quintiles, indicating that dispersion as a factor may have even greater significance for smaller firms not in the S & P 500 Index. HML is significantly and positively related to excess returns, whereas the momentum factor is significantly and negatively related to excess returns. On average, dispersion captures nine to 26 basis points of excess return, depending on the model employed. Using the adjusted R^2 as a diagnostic, the model with the best fit is Model 9, which includes the four-factor model, plus the long-term dispersion factor (adjusted $R^2 = 0.9725$), followed by Model 7 which incorporates the short-term dispersion factor (adjusted $R^2 = 0.9718$). Generally, the pricing errors of the models are small and insignificantly different from zero. Models with the lowest pricing errors are the four-factor model and the four-factor model with long-term dispersion.³¹

Evident from the factor model analysis, a factor for dispersion cannot fully compensate for the explanatory power of fundamental factors, but the inclusion of this factor does improve the fit of the factor models, especially for small firms. In the next section, we create estimates of asset returns and volatility from the 15 factor models. We then determine which models best predict out-of-sample returns and return variances.

4.3 Using factor models to estimate returns and volatility of S & P 500 firms

Does the choice of factor model have implications for portfolio management? We investigate this issue by examining the out-of-sample predictive ability of traditional factors and the heterogeneity factor for returns and variance. We estimate 60-month rolling regressions beginning every month from January 1991 through the remaining sample period, using both individual securities as well as portfolios formed on size and book-to-market.³² As noted in Griffin (2002), we estimate the time-series regressions without intercept terms because this has been shown to generate more accurate estimates [Fama and French (1997), Simin (2000)]. We then implement the methodology detailed above to calculate both return estimates [similar to that detailed in Griffin (2002)] and variance estimates [Chan, Karceski, and Lakonishok (1999)].

³¹ A two-step regression analysis is also performed. Residuals from excess returns projected onto traditional factors are projected onto the dispersion factors. We find that dispersion is significantly related to the residuals and explain between 3% and 25% of the pricing error. Results are available upon request.

³² Unlike the analysis performed in the previous subsections, we use individual firm data rather than a value-weighted portfolio for estimating returns and volatilities out-of-sample. Because of the limited First Call data, we do not have complete 60-month rolling samples in the early years. Since the initial First Call data on both short-term and long-term forecasts began in earnest in March 1989, we have only 21 months of data for the initial month of the sample (January 1991). Although this could lead to problems with estimation of returns and variances, we do not find any statistical difference between those observations with the full five-year estimation period and those with reduced estimation periods.

We construct forecasts of stock returns using Fama and French (1997) methodology for the 15 factor models for both individual firms and portfolio returns, where portfolios are formed on size and book-to-market factors. There is a general consensus that expected returns are notoriously difficult to predict. Overall, none of the models does a particularly good job of predicting out-of-sample returns, and we do not find substantial deviations in the predictive power of a variety of models. More parsimonious models are slightly better, but forecasts of returns from all models underestimate actual returns. The results are not surprising given the strong market for most of the actualization period. Forecasts from models that excluded the market had marginally improved fit, as measured by adjusted R^2 . Similar results are obtained when portfolios of the S & P 500 stocks are examined, indicating that factor models do not perform well in estimating out-of-sample returns, regardless of their ability to explain in-sample returns, as shown in Section 4.2.³³

In Table 4, we examine whether the addition of dispersion factors in conjunction with models of fundamental factors improves the predictability of out-of-sample return variances. We implement the methodology detailed in Chan, Karceski, and Lakonishok (1999) for estimating return variances for individual securities variances. Their results indicate that the

Table 4
Performance of variance forecasting models for individual firms

Model	Correlation	Slope	Adjusted R^2
1	0.5403	0.66	0.2990
2	0.4203	0.63	0.2267
3	0.5339	0.59	0.2851
4	0.5404	0.66	0.2987
5	0.5388	0.66	0.2968
6	0.4170	0.60	0.2208
7	0.5367	0.61	0.2880
8	0.4172	0.59	0.2236
9	0.5337	0.61	0.2848
10	0.5457	0.68	0.3051
11	0.4874	0.66	0.3033
12	0.5358	0.65	0.2927
13	0.5358	0.62	0.2870
14	0.5332	0.63	0.2898
15	0.5347	0.61	0.2859

This table displays forecasts of monthly return variances (in %) generated from 15 models based on a rolling 60-month window of data for each stock. Reported are the Pearson correlation coefficients, the slope coefficient in the regression of realizations on forecasts and the adjusted R^2 values for those regressions. Model 1 contains the market factor. Model 2 is the Fama and French (1993) three-factor model, whereas Model 3 is the three-factor model plus the momentum factor. Model 4 (5) is computed from the market factor plus short-term (long-term) dispersion. Model 6 (8) expands on 4 (5) by adding size and book-to-market factors, whereas Model 7 (9) expands on 6 (8) with the inclusion of the momentum factor. Model 10 (11) isolates the impact of the short-term (long-term) dispersion factor. Models 12, 13, 14, and 15 exclude the impact of the market factor from models 6, 7, 8, and 9, respectively.

³³ Results from estimation of out-of-sample returns are available upon request.

estimates from more parsimonious models (historical market variance and historical security variance) have the lowest forecast errors and the highest correlations with realized variances.

In predicting out-of-sample volatility, Model 10, where short-term dispersion is the only factor is the best estimate of individual security return variance. When the factor for short-term dispersion is the sole estimator, the slope and the correlation between the estimate and the realized variance is maximized (0.5457). The model also has the highest adjusted R^2 of the models tested (0.3051). The market model (Model 1) and the market accompanied by short-term dispersion (Model 4) are also good models for predicting volatility. Similar to the findings of Chan, Karceski, and Lakonishok (1999), more parsimonious models have better prediction ability than the complex multifactor models. Models that contain the Fama–French three-factor specification tend to have the largest standard deviation, but also the lowest forecast averages, indicating that these models underestimate individual firm volatility. These models also perform the most poorly in prediction as given by their correlation coefficients and the fit of each model. We observe the short-term dispersion model arrives at better forecasts of individual stock variance than even the market. These results may occur because unlike the other factors that rely upon historical data, earnings forecasts are forward-looking expectations, which are significantly related to both future returns and volatility of firms.

5. Behavioral Expectations

The heterogeneity of beliefs and risk tolerance are basic premises of the behavioral decision-making literature. The emphasis in behavioral finance is on how psychological *biases* affect investor behavior and prices. Investors' subjective return distributions are at times too high (too optimistic or overconfident) or too low (too pessimistic). These behavioral features have prompted a substantial debate about many key issues ranging from predictability of asset returns and plausible explanations based on investor biases, regarding momentum and long-run reversals, equity premium and risk free rate puzzles, to name just a few.³⁴

In this section, we introduce the behavioral elements of amplification and bias into the model described in Section 1. Amplification magnifies the differences among analyst forecasts. Biases allow the beliefs of analysts to be systematically different from the beliefs of agents. Effectively biases adjusts the conditional means of analyst forecasts whereas amplification adjusts the variance (and higher moments) of forecasts

³⁴ An abundance of references should be cited here which appear in a systematic, organized way in the recent surveys of Barberis and Thaler (2002) and Daniel, Hirshleifer, and Teoh (2002). A detailed overview of biases in analyst forecasts and recommendations is provided in Section 3.1.1.

across analysts. The first subsection pertains to amplifying beliefs whereas the second subsection pertains to biases. The final subsection relates our models to Shefrin's work on sentiments and pricing kernels.

5.1 Amplifying beliefs

The diversity of analyst forecasts used in our empirical work may not fully represent the diversity of beliefs in the economy as a whole. Financial analysts receive similar training, work in institutions with similar goals, and interact in a community of other financial analysts. It is likely the views of financial analysts will underrepresent the views of agents in the economy as a whole since many other agents will have much different experiences and points of view.³⁵

Another reason the diversity of analyst forecasts may not fully represent the diversity of beliefs in the economy is that it is perhaps in the interest of analysts to understate their beliefs. Analysts may face pressure to state forecasts that are not too far out of line with analysts at other institutions. It may be optimal for analysts to understate their beliefs to perform well in their job. When an analyst's constructed return forecast indicates that he believes that the return on an asset is going to be 0.01 higher than the market, he may really mean that he believes that it is going to be 0.05 higher but in case he is wrong and the asset underperforms the market he hedges his estimate. If he reported 0.05 and that turns out to be wrong, he may lose credibility in the future if all other analysts were much less optimistic. If 0.05 is right and he reports 0.01, he may still get credit and be rewarded if he had the highest forecast among analysts.

Therefore, the heterogeneity of beliefs present in our sample may under/over represent the heterogeneity of beliefs in the economy. In this subsection, we continue to assume that the diversity of analyst forecasts correctly measures the direction of agents' beliefs but not their magnitudes. Consequently, we introduce a parameter θ that can amplify or dampen heterogeneity by amplifying or dampening the stated beliefs of financial analysts.³⁶ We now let the ratio of beliefs of type i agents of the conditional means of consumption growth and stock returns to the true conditional means be

$$\frac{\mu_{git+1|t}}{\mu_{gt+1|t}} = \frac{\left(\bar{\mu}_{git+1|t}\right)^\theta}{\sum_{k=1}^I \lambda_{kt} \left(\bar{\mu}_{gkt+1|t}\right)^\theta} \quad (21)$$

³⁵ De Bondt (1993), using primarily experimental settings, observes that individual investors tend to follow trends, whereas expert investors mean revert. In contrast, [LaPorta (1996)] finds excessive optimism (pessimism) following good (bad) announcements when he proxies for the beliefs of naive investors with analyst forecasts.

³⁶ When we say "stated beliefs" we mean the return forecasts constructed (using Assumption 4) from earnings forecasts.

$$\frac{\mu_{pit+1|t}}{\mu_{pt+1|t}} = \frac{(\bar{\mu}_{pit+1|t})^\theta}{\sum_{k=1}^I \lambda_{kt} (\bar{\mu}_{pkt+1|t})^\theta} \quad (22)$$

where $\bar{\mu}_{git+1|t}$ is the conditional expectation of consumption growth stated by financial analyst i and $\bar{\mu}_{pit+1|t}$ is the conditional expectation of the return on stock p stated by financial analyst i .³⁷ Note that we still link the beliefs of type i agents to those of analyst i , but we no longer require them to be equal to the stated by beliefs of analyst i . Our previous analysis is a special case in which θ is assumed to be one. When θ is greater than one then the heterogeneity among analysts is amplified, and when it is less than one but greater than zero it is dampened.

Following our discussion at the beginning of this section, there are two interpretations of amplification. The first interpretation is that financial analysts really believe their stated forecasts, $\bar{\mu}_{gt+1|t}$ and $\bar{\mu}_{pt+1|t}$ for all p , but to obtain the beliefs of agents we need to amplify the forecasts of analysts. The second interpretation is that financial analysts understate their true beliefs and by amplifying their beliefs we can recover their true beliefs which also are the beliefs of agents. For our analysis, it does not matter which interpretation is correct, though the language we use in the rest of the article follows the second interpretation. Regardless of the interpretation, Equations (21) and (22) can be viewed as maintaining limited rational expectations in which the true conditional expectations of consumption growth and stock returns are the weighted means of the beliefs of agents

$$\sum_{i=1}^I \lambda_{it} \frac{\mu_{git+1|t}}{\mu_{gt+1|t}} = 1$$

$$\sum_{i=1}^I \lambda_{it} \frac{\mu_{pit+1|t}}{\mu_{pt+1|t}} = 1$$

which are equal to the weighted means of the amplified beliefs stated by analysts.

Including θ gives us a convenient way to test for the importance of including analyst forecasts. The formula for h_{pt} is still given by Equation (12), though the specifications of $\mu_{git+1|t}/\mu_{gt+1|t}$ and $\mu_{pit+1|t}/\mu_{pt+1|t}$ have changed. When θ is zero the forecasts of analysts have no effect on asset pricing since $h_{pt} = 1$. If estimates of θ turn out to be significantly different from zero then there is evidence that analyst forecasts do matter for asset

³⁷ As mentioned earlier we use “analyst” to refer to the fictitious analyst constructed from the views of all analysts in a particular brokerage house.

pricing. In Section 6, we will estimate θ . If the beliefs of analysts are symmetric then the sign of θ does not affect the value of h_{pt} .³⁸

For our empirical analysis, we only need the ratio of beliefs given in Equations (21) and (22). We do not need to determine $\mu_{git+1|t}$ and $\mu_{pit+1|t}$. If one took $\mu_{git+1|t}$ and $\mu_{pit+1|t}$ to be the numerators of the right hand side of Equations (21) and (22) then the mean of beliefs across analysts can be unreasonably large when θ is large. This mean can easily be reduced to sensible levels without changing the ratios by dividing the beliefs of analysts by any number which is constant across analysts at a given date. The number can vary across dates. See Section 6.1 for more discussion and an example.

5.2 Incorporating biases

The starting point of behavioral models is that the economy as a whole features biases. In previous sections, we have assumed that the beliefs of financial analysts are unbiased proxies for the beliefs of agents. In this section we allow for biases. We use data on the buy/sell recommendations of analysts to approximate the difference of the beliefs of analysts and agents. When a financial analyst strongly recommends buying a stock we view this as signifying that the analyst views the future prospects of the stock more favorably than agents. Since our models require the beliefs of agents, we adjust the return forecasts of financial analysts by analyst recommendations. We continue to maintain the assumption that financial analysts are correct [where their beliefs are given by Equations (21) and (22)] on average about actual expected returns.

Financial analysts provide recommendations on many different stocks. Their recommendations are classified into five categories: strong sell, sell, hold, buy, and strong buy. We assign the following point structure for financial analyst i 's recommendation of stock j

$$R_{ij} = \begin{cases} -2 & \text{if strong sell} \\ -1 & \text{if sell} \\ 0 & \text{if hold} \\ 1 & \text{if buy} \\ 2 & \text{if strong buy} \end{cases}$$

so that a higher number indicates a stronger positive recommendation. Our analysis uses a single measure of the recommendations of *all* financial

³⁸ The notion of symmetry used here is as follows: Let τ_t be a nonnegative constant that can vary over time and be arbitrarily chosen at any particular date. Beliefs are symmetric if for any ϱ the total Pareto weight on agents who believe the conditional mean of a variable is ϱ is equal to the total Pareto weight of agents who believe the conditional mean is τ_t/ϱ .

analysts about a portfolio of stocks where we have the recommendations of several investment houses for each stock in the portfolio.³⁹ We construct this measure for portfolio p at time t as

$$b_{pt+1|t} = \sum_{i=1}^I \sum_{j=1}^m \lambda_i \vartheta_j R_{ij}$$

when there are m stocks in the portfolio and I financial analysts. Here λ_i is a measure of the importance of financial analyst i , and ϑ_j is a measure of the market share of stock j . We require $\sum_{i=1}^I \lambda_i = 1$, $\sum_{j=1}^m \vartheta_j = 1$, $\lambda_i \geq 0$, and $\vartheta_j \geq 0$.

Many researchers have studied the bias of financial analysts with respect to *ex post* realizations. In this article, we assume that the return forecasts of financial analysts are not biased but agents' beliefs can be biased. Although some evidence suggests that financial analysts provide biased predictions of reality [see Daniel, Hirshleifer, and Teoh (2002) for a comprehensive overview of analyst biases], most of the emphasis in the literature concerns individual stocks. For the portfolios discussed in this article, it is plausible that financial analysts are not biased. In Table 5, we see that the difference between the unconditional return forecasts and the actual market return is very small. For other assets, the differences are slightly larger but still arguably small.

Table 5
Comparison of actuals, forecasts and recommendations (April 1994 to December 1997)

	Actual returns		Return forecasts		Recommendations	
	Mean	Standard error	Mean	Standard error	Mean	Standard error.
Market	1.0198	0.0333	1.0205	0.0062	0.8118	0.0052
High volatility portfolio	1.0178	0.0454	1.0039	0.0095	0.9055	0.0876
High short-term dispersion portfolio	1.0002	0.0391	0.9895	0.0222	0.7473	0.0690
High long-term dispersion portfolio	1.0194	0.0397	1.0031	0.0093	0.8596	0.0909
Nominal risk-free bond	1.0040	0.0004	—	—	—	—

This table provides summary statistics for the actual gross returns, the average forecasted gross returns, and the analyst recommendations for the portfolios used in our GMM estimations. The first return used is from April to May of 1994 and last return used is from November to December of 1997. All returns are adjusted for inflation with the CPI. At each date, the average forecasted return is computed by weighting the forecasts of each investment house by their importance. The means reported in the table are the means over all dates in our sample. The recommendations were computed as follows. A single analyst recommendation of a single stock is assigned -2 if the analyst strongly recommends selling the stock, -1 if the analyst recommends selling the stock, 0 if the analyst recommends holding the stock, 1 if the analyst recommends buying the stock, and 2 if the analyst strongly recommends buying the stock. Recommendations for the portfolios are constructed by taking a weighted average across analysts (where the weights are the importance of the analysts) and a weighted average across the stocks in the portfolio (where the weights are determined by the market share of the stocks). The market recommendations reported are only for the subset of stocks for which we have data—the returns forecasts and the actual returns are for a larger group of stocks. We also do not have any data on the return forecasts or the recommendations of the nominal risk-free bond used in our estimations.

³⁹ The group of investment houses is held fixed across time and portfolios. See Section 3.2 for more details.

In this article, biases are situations in which agents and analysts disagree. Since analysts' return forecasts are assumed to be correct on average when there is a wedge between the beliefs of agents and analysts, agents have biased beliefs of reality. However, we do not allow for biases in all variables. We assume that there are no biases between financial analysts and individuals, on average, across all stocks at any date. Thus in our model, individuals and financial analysts have the same beliefs about the performance of the market portfolio.⁴⁰ We also assume that individuals and financial analysts have exactly the same beliefs about nominal risk-free bonds, whose returns are known with certainty in nominal terms. To the extent, individuals and analysts disagree about future inflation this is a poor assumption. In addition, for convenience we assume that there are no biases in the beliefs about consumption growth. For all other assets and portfolios, there can be biases.

Although we assume analysts' return forecasts are unbiased, the language that analysts use to state recommendations may be biased. It is well known that analysts on average recommend buying stocks. For our sample, in Table 5 we see that the average recommendation across all analysts of the market is 0.8118. This average is only across stocks for which we have recommendations as we do not have recommendations for all stocks in the market. Since we assume on average there are no biases, we need to adjust for favorable recommendations. We let ζ be the adjustment factor. We will consider models in which ζ is set to 0.8118 and, for comparison purposes, models in which it is set to 0.9. We take a recommendation of ζ for a portfolio to indicate that agents should hold their position. In the rest of this article, when we say a financial analyst recommends holding a stock we mean that the recommendation for the portfolio is ζ . Any recommendation above (below) this signifies a buy (sell) recommendation. The difference between the recommendation and ζ measures the strength of the buy (sell) recommendation.

Our approach views the recommendations of analysts as comments on the *current* beliefs of agents. When analysts strongly recommend agents buy a stock, analysts think agents are too pessimistic about the stock. Equivalently, analysts think if agents had the correct beliefs then the stock is underpriced. However, agents are optimizing and under their beliefs the stock is correctly priced. This interpretation can be problematic if agents adjust their beliefs upon hearing the recommendations of analysts, and analysts do not immediately revise their recommendations. For example, if upon hearing a change in recommendations from hold (ζ) to buy (an average level of recommendations greater than ζ), agents update their beliefs,

⁴⁰ Analysts' stated recommendations for the market vary over time. One could use this to allow for biases in the market return. Our approach assumes that there are never any biases in the market return. This is partially justified because as we see from Table 5 the stated market recommendations have a very low standard deviation over time. From a practical standpoint another reason we choose to assume there are no biases in the market is that we do not have recommendations for all stocks in the market.

so that their beliefs are identical to the beliefs of analysts then the price of the stock should rise so that the stock is correctly priced under the updated beliefs. After the rise, the stock is no longer an attractive buy, and analysts should revise their recommendations to be ζ . There is some evidence which bears on this issue. In practice, analyst recommendations are generally long term in nature, whereas prices are continually updated as individual beliefs change. There is evidence [Womack (1996), Michaely and Womack (1999)] that agents adjust their beliefs upon hearing the recommendations of analysts. However, as Michaely and Womack (1999) report, the adjustments are insufficient to bring their expectations in line with analysts.

We make an assumption about how the beliefs of agents are tied to the beliefs of the analysts:

Assumption 6. *The ratio of beliefs of type i agents about the conditional mean of consumption growth and the conditional mean of the return on stock p to the true conditional means are*

$$\frac{\mu_{git+1|t}}{\mu_{gt+1|t}} = \frac{(\bar{\mu}_{git+1|t})^\theta}{\sum_{k=1}^I \lambda_{kt} (\bar{\mu}_{gkt+1|t})^\theta} \quad (23)$$

$$\frac{\mu_{pit+1|t}}{\mu_{pt+1|t}} = \left[\frac{(\bar{\mu}_{pit+1|t})^\theta}{\sum_{k=1}^I \lambda_{kt} (\bar{\mu}_{pkt+1|t})^\theta} \right] \exp[-d(b_{pt+1|t} - \zeta)] \quad (24)$$

where $\bar{\mu}_{git+1|t}$ is the conditional expectation of consumption growth stated by financial analyst i and $\bar{\mu}_{pit+1|t}$ is the conditional expectation of the return on stock p stated by financial analyst i .

Appendix C justifies the specification for $\mu_{pit+1|t}/\mu_{pt+1|t}$ in Assumption 6 using arguments based on optimizing behavior. We do not adjust analyst i 's beliefs by analyst i 's recommendation to obtain the beliefs of type i agents but rather adjust analyst i 's beliefs by the average level of recommendations across all analysts. The constant d is a parameter which is constant across stocks and measures the extent to which recommendations signify differences in the beliefs of analysts and agents. If there are no differences then $d = 0$, and there is no informational content in analyst recommendations. Our previous analysis is a special case in which d is assumed to be zero. When $d > 0$, if analysts strongly recommend buying a stock then agents are pessimistic about the stock relative to analysts. In this case $b_{pt+1|t} > \zeta$ and

$$\frac{\mu_{pit+1|t}}{\mu_{pt+1|t}} < \left[\frac{(\bar{\mu}_{pit+1|t})^\theta}{\sum_{k=1}^I \lambda_{kt} (\bar{\mu}_{pkt+1|t})^\theta} \right]. \quad (25)$$

Here the right hand side can be interpreted as the ratio of beliefs of analyst i about the conditional mean of the return on stock p to the true conditional mean. When $d > 0$, if analysts strongly recommend selling a stock then agents are optimistic about the stock relative to analysts. In this case, $b_{pt+1|t} < \zeta$ and the inequality in Equation (25) is reversed. We estimate d in later sections and find it to be positive, though usually not significantly different from zero. The parameter θ is included to allow for amplification, and ζ is included to allow for biased language in recommendation statements.

We continue to assume a limited form of rational expectations in which the average amplified beliefs of analysts are rational. Because of Assumption 6, agents on average are correct about consumption growth but possibly incorrect on average about some assets when $d \neq 0$ since

$$\sum_{i=1}^I \lambda_{it} \frac{\mu_{pit+1|t}}{\mu_{pt+1|t}} = \exp[-d(b_{pt+1|t} - \zeta)]. \quad (26)$$

Analysts are still correct on average about all stocks. In Section 2, we imposed Equations (15) and (16) which stated that agents are correct on average about the conditional means of asset returns and consumption growth. It was noted that this restriction does not imply we have nothing to say about the distributions of other variables since we do not assume that weighted average forecasts of all variables are correct. Thus, even in our earlier models the economy as a whole can be biased in some variables. The analysis in this section effectively broadens the scope of potential biases to include mean stock returns.

5.3 Behaviorally based decompositions

In section 1.4, we decomposed the expected return of an asset into fundamental and heterogeneity components. When there is no disagreement, at date t , $h_{pt} = 1$, and the heterogeneity component is zero. With our biased beliefs specification in the previous section the pricing equation is again given by Equation (11), and h_{pt} is given by Equation (12). Now h_{pt} includes components related to bias in conditional mean stock returns, in addition to heterogeneity components, since the beliefs of agents about conditional mean stock returns can be biased. Even if all agents have identical beliefs h_{pt} need not equal one if beliefs about conditional mean stock returns are biased. The decomposition in Section 1.4 still applies with the new specification of beliefs.

Shefrin (2001) formulates a similar decomposition in a different way. Shefrin (2001) shows that the log-pricing kernel can be decomposed into two stochastic processes, one pertaining to fundamentals and the other pertaining to sentiments where the latter refers to the errors in the expectations of a stand-in representative agent. Prices are efficient whenever the sentiment component is uniformly zero. Shefrin's sentiment component can include the effects of heterogeneity, biases, and irrational beliefs. The decomposition set forth in

Shefrin (2001) is easiest to state in terms of notation similar to that used in our Appendix A as it is driven by a likelihood ratio $\pi_R(\omega^v|\omega^s)/\pi(\omega^v|\omega^s)$. The ratio involves on the one hand $\pi(\omega^v|\omega^s)$, which is the true probability, based on information available at time s that the history of states will actually be ω^v at time v . On the other hand, it involves at time s , the representative agent beliefs $\pi_R(\omega^v|\omega^s)$, which are a weighted average of the beliefs of individual agents. Whenever the two probabilities $\pi_R(\omega^v|\omega^s)$ and $\pi(\omega^v|\omega^s)$ coincide, the representative agent holds correct beliefs, and the sentiment component in the Shefrin (2001) decomposition disappears. In any other situation, there is a sentiment component present in the pricing kernel because of biased aggregate beliefs. Shefrin's approach is an elegant way to characterize the importance of sentiments on asset returns.⁴¹

6. GMM Estimation of Models with Heterogeneity and Biases

In this section, we estimate structural models of the stochastic discount factor described in earlier sections. In Section 6.1, we estimate the heterogeneous agent model without biases and in Section 6.2 we include biases. Our results suggest that both heterogeneity and bias matter for asset pricing.

The pricing Equation (11) implies the unconditional moment condition

$$E[\beta g_{t+1}^{-\gamma} h_{pt} r_{pt+1} - 1] = 0 \quad (27)$$

holds for any asset. We estimate a model with five assets using GMM (Hansen 1982) without any instruments.⁴² The five assets are the market, a nominal risk-free bond, a portfolio of stocks with a high degree of past volatility, a portfolio of stocks with a high degree of dispersion among analysts short-term forecasts and a portfolio of stocks with a high degree of dispersion among analysts long-term forecasts. We assume there is no disagreement about the return on the nominal risk-free bond, b , so that $\mu_{bit} + 1|_t = \mu_{bt} + 1|_t$ for all i . All asset returns and consumption growth are in real terms, and we use the monthly CPI for inflation. We measure consumption with the sum of monthly services and nondurables.

In Section 6.1, for any asset p , h_{pt} is formed using the amplified beliefs presented in Equations (21) and (22). In Section 6.2, for any asset p , h_{pt} is

⁴¹ This heterogeneity/bias component (h_{pt}) in our work captures all deviations from fundamentals and plays the same role as the sentiment component in Shefrin (2001). We prefer the label heterogeneity/bias component to sentiment component because sentiment has a connotation that agents are irrational and somehow are using a mechanical procedure to form their beliefs. As Shefrin states in a book about behavioral finance: "A consistent theme in this book is that sentiment is the reflection of heuristic-driven bias," Shefrin (2000, page 53). The heterogeneity/bias component is present because of heterogeneity and bias whether or not they are heuristically driven.

⁴² We use the fixed weighting matrix $W = (E r_t r_t')^{-1}$ where r_t is the vector of gross returns on the five assets used in our estimation. The expectation is approximated with the sample average. Hansen and Jagannathan (1997) show that this weighting matrix has a number of appealing properties. While not all of those properties hold in the setting of this article, it is a convenient fixed weighting matrix.

formed using the amplified beliefs with biases presented in Equations (23) and (24). Both sections set

$$\log \bar{\mu}_{git+1|t} = \bar{q}_t + \phi \log \mu_{mit+1|t} \quad (28)$$

for analysts. Note that when forming h_{pt} the value of $\bar{q}_t$ cancels and that $\bar{q}_t$ need not equal q_t . (Recall that q_t was defined in Assumption 3.) The formula for h_{pt} involves the current Pareto weights of agents, and we approximate for the Pareto weights using Assumption 5. We update the Pareto weights once a year.⁴³

6.1 Heterogeneity

In this subsection, we discuss estimates of various combinations of θ , β , and γ in the consumption-based model without biases. Each panel of Table 6 presents estimates of the same 12 models except that panel A fixes ϕ at one and panel B fixes ϕ at 0.05. In models one through three, we estimate only θ for several different values of γ . In models four through six we estimate versions of the model in which all agents have the same beliefs. These models are obtained by setting $\theta = 0$ and consequently do not use data on analyst forecasts.⁴⁴ In models seven through nine, we set $\theta = 1$ so that the beliefs of agents are unamplified. In models 10 through 12, we estimate the amplification of beliefs, θ , along with various combinations of other parameters.

If financial agents were a random sample from the population and accurately reported their forecasts, we would expect that models in which $\theta = 1$ would perform better than models which ignored analyst forecasts and set $\theta = 0$. Our results confirm this because models seven, eight, and nine which set $\theta = 1$ do better than the corresponding models four, five, and six which set $\theta = 0$. Though, the gains from increasing θ from zero to one are extremely small. For example, consider the reduction in the GMM objective from going from model seven, which is the best model which restricts θ to be zero, to model nine, which is the best model which restricts θ to be one. The GMM objective is reduced from 0.3692 to 0.3618 in panel A and from 0.3692 to 0.3690 in panel B.

To understand why the gains are small, it is helpful to study the magnitudes of the differences in forecasts presented in Figure 2. In this figure, we graph the ratio of the most optimistic return forecast to the most pessimistic return forecast across investment houses. Although there is a lot of dispersion in forecasts in the high short-term portfolio, there is less dispersion in market forecasts. Since the dispersion in market

⁴³ Our theory tell us that the Pareto weights should vary from month to month if agents have different beliefs. Holding the Pareto weights fixed during a year is an approximation.

⁴⁴ When $\theta = 0$, the value of φ does not matter so the estimates presented for models four through six in panel A are identical to the estimates in panel B.

Table 6
Estimates of the consumption-based mode 1

Model	Panel A: $\phi = 1$				Panel B: $\phi = 0.05$			
	θ	β	γ	Objective	θ	β	γ	Objective
1	88.70 (35.98)	0.999	1	0.5206	476.05 (250.85)	0.999	1	0.5412
2	40.70 (14.21)	0.999	5	0.5188	180.88 (87.27)	0.999	5	0.5326
3	28.91 (9.77)	0.999	10	0.5068	125.85 (54.99)	0.999	10	0.5234
4	0	1.005 (0.030)	5	0.6038	0	1.005 (0.030)	5	0.6038
5	0	0.999	447.07 (235.41)	0.3698	0	0.999	447.07 (235.41)	0.3698
6	0	0.966 (0.3831)	457.29 (318.57)	0.3692	0	0.966 (0.3831)	457.29 (318.57)	0.3692
7	1	1.005 (0.030)	5	0.6036	1	1.005 (0.030)	5	0.6038
8	1	0.999	445.20 (235.51)	0.3626	1	0.999	447.10 (235.41)	0.3694
9	1	0.953 (0.3956)	459.70 (321.30)	0.3618	1	0.966 (0.3832)	457.31 (318.60)	0.3690
10	40.63 (14.32)	1.003 (0.004)	5	0.5096	179.89 (88.48)	1.009 (0.005)	5	0.5326
11	3.68 (2.01)	0.999	229.87 (229.87)	0.3188	18.93 (9.88)	0.999	439.54 (238.51)	0.3110
12	-4.60 (1.97)	0.688 (0.530)	497.00 (414.47)	0.2932	19.01 (9.54)	0.956 (0.414)	453.64 (343.07)	0.3104

This table displays GMM estimates of θ , β , and γ in the consumption-based model with five assets—the market, a nominal risk-free bond, and three portfolios. The portfolios consist of stocks with a high degree of past volatility, a high degree of long-term dispersion, and a high degree of short-term dispersion. The standard errors of the parameters are listed to the right of the variables in parenthesis. If there is no standard error present then the variable was fixed and not estimated. In this case the value of the variable in the estimate column is the value at which it is fixed. Objective is the value of the GMM criterion. In panel A we fix ϕ at one and in panel B we fix ϕ at 0.05.




Figure 2
Dispersion of beliefs

This figure plots a measure of the dispersion of beliefs for several different values of θ . Each figure displays the ratio of most optimistic forecast to the least optimistic forecast across houses for a given asset: $\{\max_i \mu_{pit+l|t} / [\min_i \mu_{pit+l|t}]\} = \left\{ \frac{\max_i (\bar{\mu}_{pit+l|t}^\theta)}{[\min_i (\bar{\mu}_{pit+l|t}^\theta)]} \right\}$. In each row, the left figure displays the ratio for the market portfolio and the right figure displays the ratio for the high short-term dispersion portfolio. In row one $\theta = 1$, in row two $\theta = 10$, in row three $\theta = 50$, and in row four $\theta = -4$.

forecasts is low the dispersion in consumption forecasts is low because of Assumption 3. If there is very little dispersion in consumption forecasts then h_{pt} is close to one unless there is much dispersion in asset p . In the extreme case in which there is no dispersion in consumption forecasts then h_{pt} is exactly one no matter how much dispersion there is in asset p . Hence, from the decomposition in Section 1.4, heterogeneity will have a small effect.

It might be the case that the dispersion in analyst forecasts has the right direction but the wrong magnitude. To investigate this, in models 10

through 12 we estimate θ and let the data tell us its appropriate values. When we estimate θ , we find that it is usually barely significantly different from zero. Because of our small data set and our choice to not use the optimal GMM weighting matrix, it is not surprising that the standard errors of our parameter estimates are large. In most of our models, all of the parameters are usually just barely significant.

We find that estimates of θ vary dramatically from -4.60 to 476.05 . Part of the reason for the large differences in estimates is that there is a trade-off between γ and θ . In models one through three, we see that if γ is fixed at a low value then the optimal estimate of θ is large. If γ is fixed at a high value then the optimal estimate of θ is small. This trade-off is also apparent in models 10 through 12. In model 10 in which γ is fixed at five, the optimal estimates of θ are much larger than that when γ is estimated in models 11 and 12. This happens because the estimated values of γ are much larger than five.

At first glance one may think that our estimates of θ range from being unreasonably negative to being unreasonably large and positive. If we consider what negative and large positive θ 's imply for beliefs most of our estimates are perhaps reasonable. The plots in Figure 2 suggest that amplifications of 10, -4 , and maybe even 50 are plausible for the market. When $\theta = 10$, the ratio of the most optimistic forecast to the least optimistic forecast of the gross market (monthly) return varies over time from about 1.005 to 1.014. When $\theta = -4$, it varies between about 1.002 and 1.006. When $\theta = 50$, it varies between about 1.03 and about 1.08. For the market, 1.08 is high but perhaps plausible. Looking at the high short-term dispersion portfolio, we see that an amplification of 10 entails a ratio that can be as high as 1.70. This is large but perhaps plausible since by construction this portfolio includes stocks that analysts disagree about. However, for the short-term dispersion portfolio an amplification of 50 is implausible: it implies that the ratio of the most optimistic forecast to the most pessimistic monthly forecast varies over time between about 2.0 and 14.0.

To understand the implications of amplification, consider an example in which $\theta = 10$, the most optimistic analyst's unamplified forecast of the gross market return is 1.011, and the most pessimistic analyst's unamplified forecast is 1.010. If we let the beliefs of agents be given by the numerator of the right hand side of Equation (22) then the beliefs of agents who correspond to the optimistic and pessimistic analysts are 1.1156 and 1.1046, respectively, which are, perhaps, unreasonably high. However, note that because of our specification of h_{pt} all that matters for our analysis are ratios of the form $\left(\bar{\mu}_{mit+1|t}\right)^\theta / \sum_{k=1}^I \lambda_{kt} \left(\bar{\mu}_{mkt+1|t}\right)^\theta$. This ratio is unaffected if we replace $\bar{\mu}_{mit+1|t}$ for all i with $\bar{\mu}_{mit+1|t}/\eta_t$ for any constant η_t . We could set $\eta_t = 1.1$ and let the beliefs of the agents who correspond to the optimistic and pessimistic analysts be 1.0142 and 1.0042. These beliefs are plausible and

would not change the parameter estimates in our consumption-based model. Note that we are free to divide by different constants at different dates and our empirical results would not change.

6.2 Biases

Table 7 provides estimates of the consumption-based model when there are biases. We estimate 12 different combinations of d , θ , β , and γ for two different values of ζ and ϕ . The assets are the same as in the previous subsection, though the 12 models do not correspond to the 12 models estimated in Table 6. In models one through four, ϕ is fixed at one, and ζ is fixed at 0.8118. In models five and six, ϕ is fixed at 0.05, and ζ is fixed at 0.8118. In models seven through ten, ϕ is fixed at one, and ζ is fixed at 0.9. In models 11 and 12, ϕ is fixed at 0.05, and ζ is fixed at 0.9.

When we include biases, we find that optimal estimates of d are remarkably stable and range between 0.04 and 0.06. If $d = 0.05$ and $\zeta = 0.8118$ then when analysts on average state a buy recommendation, analysts are about 1% more optimistic than individuals about returns on the asset since $\exp[-d(1 - \zeta)] \approx .99$. If $d = 0.05$ and $\zeta = 0.8118$ then when analysts on average state a hold recommendation, analysts are about 4% more pessimistic than individuals about returns on the asset since $\exp[-d(0 - \zeta)] \approx 1.04$. We find both of these numbers reasonable.

When $\zeta = 0.8118$, $\phi = 1$ and we estimate d , θ , β , and γ the value of the GMM criterion is 0.1849, which is much lower than it was when we set $d = 0$ and estimated θ , β , and γ in model 12 of Table 6. Both dispersion and bias seem to be important. In model three, when we do not allow for dispersion

Table 7
Estimates of the consumption-based model with biases

Model	d	θ	β	γ	Objective
1	0.040 (0.031)	0	0.999	5	0.5312
2	0.048 (0.076)	43.90 (13.77)	1.003 (0.006)	5	0.4038
3	0.048 (0.037)	0	0.879 (0.478)	498.59 (354.02)	0.2501
4	0.047 (0.196)	-4.51 (0.3084)	0.674 (0.549)	509.04 (362.96)	0.1849
5	0.059 (0.031)	243.60 (55.60)	0.991 (0.034)	5	0.4014
6	0.044 (0.787)	17.91 (36.66)	0.928 (0.782)	473.54 (537.54)	0.2120
7	0.068 (0.034)	0	0.999	5	0.4210
8	0.058 (0.277)	26.68 (38.46)	1.004 (0.011)	5	0.4070
9	0.051 (0.037)	0	0.981 (0.400)	443.11 (328.51)	0.2282
10	0.047 (0.167)	2.45 (0.42)	0.922 (0.440)	447.65 (333.60)	0.2256
11	0.060 (0.085)	114.85 (101.06)	1.005 (0.022)	5	0.4109
12	0.053 (0.163)	0.00 (7.44)	0.993 (0.393)	442.18 (329.06)	0.2282

This table displays GMM estimates of d , θ , β , and γ in the consumption-based model with biases. The assets are the market, a nominal risk-free bond, and three portfolios. The portfolios consist of stocks with a high degree of past volatility, a high degree of long-term dispersion, and a high degree of short-term dispersion. The standard errors of the parameters are listed to the right of the variables in parenthesis. If there is no standard error present then the variable was fixed and not estimated. In this case the value of the variable in the estimate column is the value at which it is fixed. Objective is the value of the GMM criterion. In models one through four, ϕ is fixed at one, and ζ is fixed at 0.8118. In models five and six, ϕ is fixed at 0.05, and ζ is fixed at 0.8118. In models seven through ten, ϕ is fixed at one, and ζ is fixed at 0.9. In models 11 and 12, ϕ is fixed at 0.05, and ζ is fixed at 0.9.

and estimate d , β , and γ the model does not do as well. In models nine and ten, we see that when ζ is fixed at 0.9, this is not always the case. In these models, dispersion adds very little over bias which suggest in some circumstances biases may be more important than dispersion. However, since model four does better than model 10, we believe our evidence shows both dispersion and bias matter. Though our estimates of θ and d are usually not significantly different from zero.

It is well known from Hansen and Singleton (1982), Mehra and Prescott (1985), Hansen and Jagannathan (1991), Kocherlakota (1996) and many others, that in consumption-based models optimal estimates of γ tend to be unreasonably large when θ is restricted to be zero.⁴⁵ Many researchers consider any estimate of γ larger than five to be unreasonable. Moreover consumption-based models have a difficult time accounting for the difference between returns on stocks and returns on risk-free bonds. This is referred to as the equity premium puzzle. Can including analyst forecasts help resolve the equity premium puzzle? The standard model has trouble accounting for the equity premium because the variance of consumption growth and the correlation of consumption growth with stock returns is low. In the models presented in this article, the agents' perceived variance of unconditional consumption growth increases, and the agents perceived unconditional correlation with stock returns is affected. However, our results indicate that our model can not satisfactorily account for the equity premium. In addition, although there is a trade-off between θ and γ , the estimates of γ are still too large. Interestingly, when γ and θ are jointly estimated the estimates of γ can be larger than that when θ is fixed at zero. One possible reason for this is that when $\theta \neq 0$, γ affects the value of h_{pr} .

7. Conclusions

The contribution of this article is to empirically implement dynamic general equilibrium models with heterogeneous agents, being careful about measuring agents' actual beliefs. We examine whether the heterogeneity of beliefs is priced and show factors constructed from the disagreement among analysts about expected (short-term and long-term) earnings is a risk factor affecting both expected returns and volatility. Although previous research has documented positive and significant association between dispersion, returns, and volatility, almost none of these tests have examined how dispersion affects asset pricing models.⁴⁶ Second, we establish the out-of-sample properties of dispersion in predicting volatility and returns. We suggest that the disagree-

⁴⁵ There have been many proposed solutions to the equity premium puzzle including Campbell and Cochrane (1999) and Bansal and Yaron (2000). See Benartzi and Thaler (1995) and Barberis, Huang, and Santos (2001) for proposed behavioral solutions.

⁴⁶ The exception is Easley, Hvidkjaer, and O'Hara (2000), which uses a measure of private information and its relation to return in Fama-French type models.

ment among analysts about expected earnings is a good proxy that measures the heterogeneity of beliefs and has attributes that correspond to the theoretical predictions. We show that dispersion of earnings forecasts is a priced factor in traditional factor asset pricing models and is a good predictor of return volatility in out-of-sample tests. Our results indicate dispersion has better out-of-sample properties for individual firms than for portfolios of stocks, but the general findings about the importance of a factor that captures the heterogeneity of beliefs still remains.

Next, having established that the heterogeneity of beliefs matters for asset pricing we turn our attention to estimating a structural model. We develop a model in which agents have the correct beliefs about expected consumption growth but possibly incorrect beliefs about all other variables including the higher moments of consumption growth. We link the beliefs of agents about expected stock returns to the stated beliefs of analysts allowing the beliefs to differ according to the average level of recommendations. Since there is a representative agent whose beliefs are a weighted beliefs of analysts, the impact of heterogeneity is channelled through its effect on average beliefs. As in Shefrin (2001) the mean beliefs may not equal the true beliefs. When they do not, heterogeneity/bias matters. We estimate the discount factor β and risk aversion parameter γ , and the determinants of agents' beliefs via analyst forecasts. Our results suggest that including analyst forecasts can improve the performance of asset pricing models. However, the estimates of the parameters which determine amplification and bias are usually not significant.

It is common to ignore the heterogeneity of beliefs in empirical asset pricing models. Implementing such models can be troublesome, and the available data can be of low quality. Our article is an attempt to address these issues, suggesting factor specifications and suggesting how to empirically incorporate the heterogeneity of beliefs in stochastic discount fundamental asset pricing equations. The results in our article suggest that heterogeneity is a "missing factor" when it comes to predicting returns and volatility. The results also suggest that traditional asset pricing models have improved empirical fits when heterogeneity is taken into account.

Appendix A: Pricing with heterogeneous agents

This appendix derives Equation (5). We use more formal notation than in the text. Let ω_t be the state of the economy at time t which captures both exogenous and endogenous information. We take the number of possible values of ω_t for a given t to be finite, though our approach applies more generally. Let $\omega^t = (\omega_0, \dots, \omega_t)$ be the history of states from time zero to time t . Let $\pi_i(\omega^v | \omega^s)$ be the probability, based on information available in the history ω^s , that agents of type i believe that the history of states will be ω^v at time v . Let $\pi(\omega^v | \omega^s)$ be the true probability, based on information available in the history ω^s , that the history of states will actually be ω^v at time v . If the history of states is ω^t at time t then let $c_t(\omega^t)$ be aggregate consumption, $c_{it}(\omega^t)$ be the consumption of agent i , $\lambda_{it}(\omega^t)$ be the Pareto weight of agent i , and $\alpha_{it}(\omega^t)$ be the fraction of aggregate consumption that agent i consumes. Also consider an asset which has a time t price of one and pays $r_{pt+1}(\omega^{t+1})$ at time $t+1$.

Let there be I agents who have identical power reward functions with the same discount factor. To compute an equilibrium, we can solve the following Pareto optimal problem: maximize

$$\sum_{i=1}^I \sum_{t=0}^{\infty} \sum_{\omega^t} \lambda_{i0}(\omega^0) \pi_i(\omega^t | \omega^0) \beta^t \frac{[c_{it}(\omega^t)]^{1-\gamma}}{1-\gamma}$$

by choice of consumption allocations for all dates and histories subject to

$$\sum_{i=1}^I c_{it}(\omega^t) \leq c_t(\omega^t)$$

$$c_{it}(\omega^t) \geq 0$$

for all dates and histories. $\{\lambda_{i0}(\omega^0)\}_{i=1}^I$ are the initial Pareto weights which we construct to sum to one.

The solution to this problem entails that the optimal one period consumption allocation rules are

$$c_{it}(\omega^t) = \alpha_{it}(\omega^t) c_t(\omega^t) \quad (\text{A1})$$

where

$$\alpha_{it}(\omega^t) = \frac{[\lambda_{i0}(\omega^0) \pi_i(\omega^t | \omega^0)]^{\frac{1}{\gamma}}}{\sum_{j=1}^I [\lambda_{j0}(\omega^0) \pi_j(\omega^t | \omega^0)]^{\frac{1}{\gamma}}} \quad (\text{A2})$$

$$= \frac{[\lambda_{it}(\omega^t)]^{\frac{1}{\gamma}}}{\sum_{j=1}^I [\lambda_{jt}(\omega^t)]^{\frac{1}{\gamma}}} \quad (\text{A3})$$

can be written so that they only depend on the current Pareto weights, $\lambda_{it}(\omega^t)$. The current Pareto weights are defined so that the right hand side of Equation (A3) equals the right hand side of Equation (A2). Under this definition the Pareto weights evolve as

$$\lambda_{it+1}(\omega^{t+1}) = \frac{\lambda_{it}(\omega^t) \pi_i(\omega^{t+1} | \omega^t)}{\sum_{j=1}^I \lambda_{jt}(\omega^t) \pi_j(\omega^{t+1} | \omega^t)}. \quad (\text{A4})$$

The denominator normalizes the term so that the time t Pareto weights sum to one. If agents have the same beliefs then the Pareto weights will be constant over time. The pricing Equation (3) becomes

$$1 = \sum_{\omega^{t+1}} \beta \pi_i(\omega^{t+1} | \omega^t) \left[\frac{c_{it}(\omega^t)}{c_{it+1}(\omega^{t+1})} \right]^\gamma r_{pt+1}(\omega^{t+1}).$$

Substituting in first from Equations (A1) and (A3), second then from Equation (A4), and then rearranging allows us to obtain

$$\begin{aligned}
 1 &= \sum_{\omega^{t+1}} \beta \pi_i(\omega^{t+1} | \omega^t) \left[\frac{\lambda_{it}(\omega^t)}{\lambda_{it+1}(\omega^{t+1})} \right] \Omega_t(\omega^{t+1}) \left[\frac{c_t(\omega^t)}{c_{t+1}(\omega^{t+1})} \right]^\gamma r_{pt+1}(\omega^{t+1}) \\
 &= \sum_{\omega^{t+1}} \beta \left[\sum_{j=1}^I \lambda_{jt}(\omega^t) \pi_j(\omega^{t+1} | \omega^t) \right] \Omega_t(\omega^{t+1}) \left[\frac{c_t(\omega^t)}{c_{t+1}(\omega^{t+1})} \right]^\gamma r_{pt+1}(\omega^{t+1}) \\
 &= \sum_{j=1}^I \lambda_{jt}(\omega^t) \sum_{\omega^{t+1}} \pi_j(\omega^{t+1} | \omega^t) \beta \Omega_t(\omega^{t+1}) \left[\frac{c_t(\omega^t)}{c_{t+1}(\omega^{t+1})} \right]^\gamma r_{pt+1}(\omega^{t+1})
 \end{aligned}$$

where we define

$$\Omega_t(\omega^{t+1}) = \left[\frac{\sum_{j=1}^I [\lambda_{jt+1}(\omega^{t+1})]^\frac{1}{\gamma}}{\sum_{j=1}^I [\lambda_{jt}(\omega^t)]^\frac{1}{\gamma}} \right]^\gamma$$

to reflect the change in the distribution of wealth from periods t to $t + 1$. Notice that if agents have logarithmic reward functions ($\gamma = 1$) then $\Omega_t(\omega^{t+1})$ always equals one. Also if the distribution of wealth is the same in histories ω^t and ω^{t+1} then $\Omega_t(\omega^{t+1})$ will equal one. We will assume the distribution of wealth, as measured by,

$$\sum_{j=1}^I [\lambda_{jt}(\omega^t)]^\frac{1}{\gamma}$$

is constant and consequently our pricing Equation becomes

$$1 = \sum_{j=1}^I \lambda_{jt}(\omega^t) \sum_{\omega^{t+1}} \pi_j(\omega^{t+1} | \omega^t) \beta \left[\frac{c_t(\omega^t)}{c_{t+1}(\omega^{t+1})} \right]^\gamma r_{pt+1}(\omega^{t+1}). \quad (\text{A5})$$

Equation (A5) is a formal version of Equation (5) in the text.

Appendix B: Correct higher order log conditional central moments

We show that Assumption 2 implies that all agents correctly know the second and higher order conditional central moments of $\log(x)$. To see this, let x be a scalar and note that conditional expected value of $\log x_{t+1}$ from the point of view of agent i is

$$\begin{aligned}
 \mu_{xit+1|t}^* &= \int \log x_{t+1} d\zeta_{it}(x_{t+1}) \\
 &= \int \log x_{t+1} d\zeta_{it} \left(\frac{x_{t+1} * \mu_{xit+1|t}}{\mu_{xit+1|t}} \right) \\
 &= \int (y_{t+1} + \log \mu_{xit+1|t} - \log \mu_{xt+1|t}) d\zeta_{it}(y_{t+1}) \\
 &= \log \mu_{xit+1|t} - \log \mu_{xt+1|t} + \mu_{xit+1|t}^*
 \end{aligned}$$

where $\mu_{xit+1|t}^*$ is the true conditional mean of $\log x_{t+1}$. The $n > 1$ conditional central moment of $\log x_{t+1}$ from the point of view of agent i is

$$\begin{aligned}
 & \int \left(\log x_{t+1} - \mu_{xit+1|t}^* \right)^n d\zeta_{ixt}(x_{t+1}) \\
 &= \int \left(\log x_{t+1} - \mu_{xit+1|t}^* \right)^n d\zeta_{xt} \left(\frac{x_{t+1} * \mu_{xt+1|t}}{\mu_{xit+1|t}} \right) \\
 &= \int \left(y_{t+1} + \log \mu_{xit+1|t} - \log \mu_{xt+1|t} - \mu_{xit+1|t}^* \right)^n d\zeta_{xt}(y_{t+1}) \\
 &= \int (y_{t+1} - \mu_{xt+1|t}^*)^n d\zeta_{xt}(y_{t+1})
 \end{aligned}$$

which is the equal to the actual n th conditional central moment of $\log x_{t+1}$. This argument can be extended to the case in which x is a vector.

Appendix C: Justification of Assumption 6

This appendix justifies the specification of $\mu_{pit+1|t}/\mu_{pt+1|t}$ in Assumption 6. Let the ratio of the beliefs of financial analyst i about the conditional mean of stock p to the true conditional mean be

$$\frac{(\bar{\mu}_{pit+1|t})^\theta}{\sum_{k=1}^I \lambda_{kt} (\bar{\mu}_{pkt+1|t})^\theta} \quad (C1)$$

and let Equation (23) hold. Recall that a type i agent's utility function from time t onward is

$$E_{it} \sum_{s=t}^{\infty} \beta^s \frac{\left(w_{is} + \sum_{p=1}^n \varphi_{ips-1} r_{ps} - \sum_{p=1}^n \varphi_{ips} \right)^{1-\gamma}}{1-\gamma}$$

where the term in parenthesis is equal to c_{is} . At time t the derivative of this utility function with respect to holdings of asset p between periods t and $t+1$ is

$$D_{ipt} = \beta^t c_{it}^{-\gamma} E_{it} \left[\beta \left(\frac{c_{it}}{c_{it+1}} \right)^\gamma r_{pt+1} - 1 \right] \quad (C2)$$

$$= \beta^t c_{it}^{-\gamma} E_t \left[\left(\frac{c_{it} \mu_{gt+1|t}}{c_{it+1} \mu_{git+1|t}} \right)^\gamma \frac{r_{pt+1} \mu_{pit+1|t}}{\mu_{pt+1|t}} - 1 \right] \quad (C3)$$

where we use Assumption 2 to replace E_{it} in Equation (C2) with E_t . Since agents maximize their utility, given their beliefs, D_{ipt} will always equal zero. Moreover, the derivative of a type i agent's lifetime utility function using the beliefs of financial analyst i is

$$\bar{D}_{ipt} = \beta^t c_{it}^{-\gamma} \bar{E}_{it} \left[\beta \left(\frac{c_{it}}{c_{it+1}} \right)^\gamma r_{pt+1} - 1 \right] \quad (C4)$$

$$= \beta^t c_{it}^{-\gamma} E_t \left[\left(\frac{c_{it} \mu_{gt+1|t}}{c_{it+1} \mu_{git+1|t}} \right)^\gamma \frac{r_{pt+1} (\bar{\mu}_{pit+1|t})^\theta}{\sum_{k=1}^I \lambda_{kt} (\bar{\mu}_{pkt+1|t})^\theta} - 1 \right] \quad (C5)$$

where $\bar{E}_{it}$ denotes the expectation with respect to analyst i 's beliefs. We have used the fact that there are no biases in consumption growth. $\bar{D}_{ipt}$ may or may not equal zero. We assume that

$$\frac{\bar{D}_{ipt}}{\beta^t c_{it}^{-\gamma}} = \exp[d(b_{pt+1|t} - \zeta)] - 1 \quad (C6)$$

where d is a positive constant. Notice that the right hand side does not depend on i . When financial analysts on average recommend holding stock p , (so that $b_{pt+1|t} = \zeta$) we view that as indicating analysts believe agents are optimizing under analysts' beliefs and hence $\bar{D}_{ipt} = \bar{D}_{ipt} = 0$. When financial analysts on average recommend buying a stock ($b_{pt+1|t} > \zeta$), analysts believe that agents could obtain more utility by purchasing more of stock p , and hence $\bar{D}_{ipt} > 0 = D_{ipt}$. When financial analysts on average recommend selling a stock ($b_{pt+1|t} < \zeta$), we view that as indicating analysts believe agents could obtain more utility by selling stock p and hence $\bar{D}_{ipt} < 0 = D_{ipt}$. Rearranging Equation (C3) using $D_{ipt} = 0$ and rearranging Equation (C5) using Equation (C6) yields

$$\frac{\mu_{pit+1|t}}{\mu_{pt+1|t}} = \left[E_t \left(\frac{c_{it} \mu_{git+1|t}}{c_{it+1} \mu_{git+1|t}} \right)^\gamma r_{pt+1} \right]^{-1} \quad (C7)$$

$$\frac{(\bar{\mu}_{pit+1|t})^\theta}{\sum_{k=1}^I \lambda_{kt} (\bar{\mu}_{pkt+1|t})^\theta} = \left[E_t \left(\frac{c_{it} \mu_{git+1|t}}{c_{it+1} \mu_{git+1|t}} \right)^\gamma r_{pt+1} \right]^{-1} \exp[d(b_{pt+1|t} - \zeta)]. \quad (C8)$$

Dividing Equation (C7) by Equation (C8) and rearranging terms yields the formula given in Assumption 6 of the text.

References

- Abel, A. B., 1989, "Asset Prices under Heterogeneous Beliefs: Implications for the Equity Premium," mimeo, University of Pennsylvania.
- Ajinkya, B. B., and M. J. Gift, 1985, "Dispersion of Financial Analysts' Earnings Forecasts and the (Option Model) Implied Standard Deviations of Stock Returns," *Journal of Finance*, 40, 1353–1365.
- Amihud, Y., and H. Mendelson, 1986, "Asset Pricing and the Bid-ask Spread," *Journal of Financial Economics*, 15, 223–249.
- Bansal, R., and A. Yaron, 2000, "Risks for the long run: A potential resolution of asset pricing puzzles," mimeo, Duke University.
- Barberis, N., M. Huang, and T. Santos, 2001, "Prospect Theory and Asset Prices," *Quarterly Journal of Economics*, 116, 1–53.
- Barberis, N., A. Shleifer, and R. Vishny, 1998, "A Model of Investor Sentiment," *Journal of Financial Economics*, 49, 307–343.
- Barberis, N., and R. Thaler, 2002, "A Survey of Behavioral Finance," NBER working paper, #9222.
- Basak, S., 2000, "A Model of Dynamic Equilibrium Asset Pricing with Heterogeneous Beliefs and Extraneous Risk," *Journal of Economic Dynamics and Control*, 24, 63–95.
- Benartzi, S., and R. H. Thaler, 1995, "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics*, 110, 73–92.
- Bessembinder, H., K. Chan, and P. Seguin, 1996, "An Empirical Examination of Information, Differences in Opinion and Trading Activity," *Journal of Financial Economics*, 11, 255–274.
- Blume, L., and D. Easley, 1992, "Evolution and Market Behavior," *Journal of Economic Theory*, 9, 9–40.
- Brav, A., R. Lehavy, and R. Michaely, 2003, "Using Expectations to Test Asset Pricing Models," mimeo, Duke University.

- Brown, L., 2001, "A Temporal Analysis of Earnings Surprise: Profits Versus Losses," *Journal of Accounting Research*, 39, 221–241.
- Browning, M., L. P. Hansen, and J. J. Heckman, 2000, "Micro Data and General Equilibrium Models," in J. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, North Holland, Amsterdam.
- Campbell, J. Y., and J. H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.
- Capstaff, J., K. Paudyal, and W. Rees, 1998, "Analysts' Forecasts of German Firms' Earnings: A Comparative Analysis Abstract," *Journal of International Financial Management and Accounting*, 9, 83–116.
- Carhart, M., 1997, "On Persistence in Mutual Fund Performance," *Journal of Finance*, 52, 57–82.
- Chan, L. K. C., N. Jegadeesh, and J. Lakonishok, 1996, "Momentum Strategies," *Journal of Finance*, 51, 1681–1713.
- Chan, L. K. C., J. Karceski, and J. Lakonishok, 1999, "On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model," *Review of Financial Studies*, 12, 937–974.
- Chen, J., H. Hong, and J. Stein, 2000, "Breadth of Ownership and Stock Returns," Discussion paper, Stanford University.
- Connor, G., and R. Korajczyk, 1988, "Risk and Return in an Equilibrium APT: Application of a New Test Methodology," *Journal of Financial Economics*, 21, 255–289.
- Crombez, J., 2001, "Updating Expected Returns Based on Consensus Forecasts," mimeo, Ghent University.
- Daniel, K., D. Hirshleifer, and A. Subrahmanyam, 1998, "Investor Psychology and Security Market Under- and Over-Reactions," *Journal of Finance*, 53, 1839–1886.
- Daniel, K., D. Hirshleifer, and S. Teoh, 2002, "Investor Psychology in Capital Markets: Evidence and Policy Implications," *Journal of Monetary Economics*, 49, 139–209.
- Daniel, K., and S. Titman, 1997, "Evidence on the Characteristics of Cross-Sectional Variation in Stock Returns," *Journal of Finance*, 52, 1–33.
- De Bondt, W. F. M., 1993, "Betting on trends: Intuitive forecasts of financial risk and return," *International Journal of Forecasting*, 9, 295–435.
- De Bondt, W.F.M., and R. Thaler, 1985, "Does the Stock Market Overreact?" *Journal of Finance*, 40, 793–808.
- De Bondt, W.F.M., and R. Thaler, 1987, "Further Evidence on Investor Overreaction and Stock Market Seasonality," *Journal of Finance*, 42, 557–581.
- De Bondt, W.F.M., and R. Thaler, 1990, "Do Security Analysts Overreact?" *American Economic Review*, 80, 52–57.
- Dechow, P., and R. Sloan, 1997, "Returns to Contrarian Investment Strategies: Tests of Naïve Expectations Hypothesis," *Journal of Financial Economics*, 41, 3–27.
- Detemple, J., and S. Murthy, 1994, "Intertemporal Asset Pricing with Heterogeneous Information," *Journal of Economic Theory*, 62, 294–320.
- Diether, K. B., C. J. Malloy, and A. Scherbina, 2002, "Differences of Opinion and the Cross-Section of Stock Returns," *Journal of Finance*, 57, 2113–2141.
- Duffie, D., and G. M. Constantinides, 1996, "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104, 219–240.
- Easley, D., S. Hvidkjaer, and M. O'Hara, 2000, "Is Information Risk a Determinant in Asset Returns?" Discussion Paper, Cornell University.
- Easterwood, J., and S. Nutt, 1999, "Inefficiency in Analysts Earnings Forecasts: Systematic Misreaction or Systematic Optimism?" *Journal of Finance*, 54, 1777–1797.

- Fama, E., and K. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.
- Fama, E., and K. French, 1997, "Industry Costs of Equity," *Journal of Financial Economics*, 43, 153–193.
- Fama, E., and M. Gibbons, 1984, "A Comparison of Inflation Forecasts," *Journal of Monetary Economics*, 13, 327–348.
- Gebhardt, W., C. Lee, and B. Swaminathan, 2001, "A Test of Efficiency of a Given Portfolio," *Journal of Accounting Research*, 39, 135–176.
- Ghysels, E., and J. L. Juergens, 2001, "Stock Market Fundamentals and Heterogeneity of Beliefs: Tests Based on a Decomposition of Returns and Volatility," http://papers.ssrn.com/sol3/papers.cfm?abstract_id=301836 (accessed September 28, 2002).
- Gorman, W., 1953, "Community Preference Fields," *Econometrica*, 21, 63–80.
- Graham, J., and C. Harvey, 1996, "Market Timing Ability and Volatility Implied in Investment Newsletters' Asset Allocation Recommendations," *Journal of Financial Economics*, 42, 397–421.
- Graham, J., and C. Harvey, 2003, "Expectations of Equity Risk Premia Volatility and Asymmetry," mimeo, Duke University.
- Griffin, J. M., 2002, "Are the Fama and French Factors Global or Country Specific?" *Review of Financial Studies*, 15, 783–803.
- Guay, W., S. P. Kothari, and S. Shu, 2003, "Properties of Implied Cost of Capital Using Analysts' Forecasts," mimeo, MIT Sloan School of Management.
- Hansen, L. P., 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054.
- Hansen, L. P., and R. Jagannathan, 1991, "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99, 225–262.
- Hansen, L. P., and R. Jagannathan, 1997, "Assessing Specification Errors in Stochastic Discount Factor Models," *Journal of Finance*, 52, 557–590.
- Hansen, L. P., and S. F. Richard, 1987, "The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models," *Econometrica*, 55, 587–613.
- Hansen, L. P., and K. J. Singleton, 1982, "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectation Models," *Econometrica*, 50, 1269–1286, Errata: *Econometrica*, 52, 267–268.
- Harris, M., and A. Raviv, 1993, "Differences of Opinion Make a Horse Race," *Review of Financial Studies*, 6, 473–506.
- Heaton, J., and D. Lucas, 1995, "The Importance of Investor Heterogeneity and Financial Market Imperfections for the Behavior of Asset Prices," *Carnegie-Rochester Series on Public Policy*, 42, 1–32.
- Hirshleifer, D., 2001, "Investor Psychology and Asset Pricing," *Journal of Finance*, 64, 1533–1597.
- Jegadeesh, N., and S. Titman, 1993, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," *Journal of Finance*, 48, 65–91.
- Kocherlakota, N. R., 1996, "The Equity Premium: It's Still a Puzzle," *Journal of Economic Literature*, 34, 42–71.
- LaPorta, R., 1996, "Expectations in the Cross-section of Stock Returns," *Journal of Finance*, 51, 1715–1742.
- L'Her, J.-F., and J.-M. Suret, 1996, "Consensus, Dispersion and Security Prices," *Contemporary Accounting Research*, 13, 209–228.
- Li, T., 1999, "Heterogeneous Beliefs, Asset Prices, and Volatility in a Pure Exchange Economy," mimeo, Washington University.

- Lin, H., and M. McNichols, 1998, "Underwriting Relationships and Analysts' Earnings Forecasts and Investment Recommendations," *Journal of Accounting and Economics*, 25, 101–127.
- Lobo, G., and S. Tung, 1998, "Financial Analysts' Earnings Forecast Dispersion and Intraday Stock Price Variability Around Quarterly Earnings Announcements," Discussion paper, Syracuse University.
- Matsumoto, D. A., 2002, "Management's Incentives to Avoid Negative Earnings Surprises," *The Accounting Review*, 77, 483–514.
- Mehra, R., and E. C. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161.
- Merton, R. C., 1987, "A Simple Model of Capital Market Equilibrium with Incomplete Information," *Journal of Finance*, 42, 483–510.
- Michaely, R., and K. Womack, 1999, "Conflict of Interest and the Credibility of Underwriter Analyst Recommendations," *Review of Financial Studies*, 12, 653–686.
- Miller, E., 1977, "Risk, Uncertainty, and Divergence of Opinion," *Journal of Finance*, 32, 1151–1168.
- O'Brien, P., 1988, "Analysts' Forecasts as Earnings Expectations," *Journal of Accounting and Economics*, 10, 53–83.
- Qu, S., L. Starks, and H. Yan, 2003, "Risk Dispersion of Analyst Forecasts and Stock Returns," mimeo, University of Texas at Austin.
- Rajan, R., and H. Servaes, 1997, "Analyst Following of Initial Public Offerings," *Journal of Finance*, 52, 507–530.
- Rubinstein, M., 1974, "An Aggregation Theorem for Securities Markets," *Journal of Financial Economics*, 1, 225–44.
- Shalen, C., 1993, "Volume, Volatility and the Dispersion of Beliefs," *Review of Financial Studies*, 6, 405–434.
- Shefrin, H., 2001, "On Kernels and Sentiment," http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2888258 (accessed June 5, 2002)
- Shefrin, H., 2000, *Beyond Greed and Fear: Understanding Behavioral Finance and the Psychology of Investing*. Harvard Business School Press, Boston.
- Shefrin, H., and M. Statman, 1994, "Behavioral Capital Asset Pricing Theory," *The Journal of Financial and Quantitative Analysis*, 29, 323–349.
- Simin, T., 2000, "The Predictive Performance of Asset Pricing Models," Discussion paper, Pennsylvania State University.
- Simon, H. A., 1955, "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, 69, 99–118.
- Stickel, S., 1992, "Reputation and Performance Among Security Analysts," *Journal of Finance*, 47, 1811–1836.
- Tversky, A., and D. Kahneman, 1974, "Judgment under Uncertainty: Heuristics and Biases," *Science*, 185, 1124–1131.
- Williams, J., 1977, "Capital Asset Prices with Heterogeneous Beliefs," *Journal of Financial Economics*, 5, 219–239.
- Wilson, R., 1968, "The Theory of Syndicates," *Econometrica*, 36, 119–132.
- Womack, K. L., 1996, "Do Brokerage Analysts' Recommendations Have Investment Value?" *Journal of Finance*, 51, 137–67.
- Zapatero, F., 1998, "Effects of Financial Innovations on Market Volatility when Beliefs are Heterogeneous," *Journal of Economic Dynamics and Control*, 22, 597–626.

Copyright of Review of Financial Studies is the property of Society for Financial Studies and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.