

Discussion Paper Series – CRC TR 224

Discussion Paper No. 337

Project C 02

Expectations Data in Asset Pricing

Klaus Adam ¹

Stefan Nagel ²

February 2022

¹ University of Mannheim, Department of Economics, Email : adam@uni-mannheim.de

² University of Chicago, Booth School of Business, Email: stefan.nagel@chicagobooth.edu

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
through CRC TR 224 is gratefully acknowledged.

Expectations Data in Asset Pricing

Klaus Adam*

University of Mannheim, CEPR and CESifo

Stefan Nagel†

*University of Chicago,
NBER, CEPR, and CESifo*

January 2022

Abstract

Asset prices reflect investors' subjective beliefs about future cash flows and prices. In this chapter, we review recent research on the formation of these beliefs and their role in asset pricing. Return expectations of individual and professional investors in surveys differ markedly from those implied by rational expectations models. Variation in subjective expectations of future cash flows and price levels appear to account for much of aggregate stock market volatility. Mapping the survey evidence into agent expectations in asset pricing models is complicated by measurement errors and belief heterogeneity. Recent efforts to build asset pricing models that match the survey evidence on subjective belief dynamics include various forms of learning about payout or price dynamics, extrapolative expectations, and diagnostic expectations. Challenges for future research include the exploration of subjective risk perceptions, aggregation of measured beliefs, and links between asset market expectations and the macroeconomy.

JEL codes: G12, G41, E71

Keywords: Investor beliefs, survey forecasts, return expectations, cash flow expectations, belief formation, asset price dynamics

*University of Mannheim, Department of Economics, L7, 3-5, 68131 Mannheim, Germany, e-mail: adam@uni-mannheim.de. Funding by the German Research Foundation (DFG) through CRC TR 224 (Project C02) is gratefully acknowledged.

†University of Chicago, Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637, e-mail: stefan.nagel@chicagobooth.edu

1 Introduction

Asset prices are inherently forward looking. The willingness to pay for an asset today depends on investors' expectations about the asset's future payouts and the future price at which the asset can be sold again. The market price of assets therefore reflects investors' price and payout expectations, as well as the risk-adjustments associated with discounting future payouts and sale prices.

Understanding the temporal behavior of asset prices thus requires understanding how price and payout expectations move over time. The traditional approach to this problem is to assume rational expectations (RE). Under RE, investors' expectations are objective in the sense that they reflect the true underlying law of motion that generates asset payoffs. With investor expectation tied down in this way, much of the large observed swings in asset prices over time are then attributed to changes in risk premia rather than changing expectations of future payouts and prices. Consequently, much of the asset pricing literature has focused on searching for specifications of preferences or technology that produce sufficiently volatile risk premia.

While analytically convenient, RE is a strong assumption. Whether or not it is a plausible one is ultimately an empirical question. A growing recent literature examines whether some of the empirical difficulties of RE asset pricing models could be addressed by allowing subjective beliefs of investors to deviate from RE.

Developing models in which investors price assets based on their subjective beliefs about future payouts and prices then requires taking a stand on how these subjective beliefs are formed. Without the tight link of beliefs to some underlying model of objective reality that RE entails, there are many possibilities. Reverse-engineering subjective beliefs from asset prices seems unattractive. Presumably, there are many different belief formation mechanisms that are observationally equivalent in terms of their predictions for asset price movements seen in historical asset price.

Expectations data therefore play an important role. Mechanisms of subjective belief formation in asset pricing models should not only produce empirically realistic asset price behavior, but they should also be plausible in light of observable data on investor expectations. With increasing availability of survey data, the study of investor expectations has become a very active area of research. In this chapter, we review this work. We start with a basic asset pricing framework that clarifies the role that different types of expectations play in asset pricing. We then discuss existing empirical evidence on the dynamics of investor expectations, followed by a review of work that aims to build asset pricing models with subjective beliefs that are consistent with

this empirical evidence. We conclude with some thoughts on the outlook for future research in this area.

2 A general asset pricing framework

We consider a general asset pricing setup allowing for heterogeneity across agents in their beliefs and preferences. The setup nests many structural asset pricing models as special cases. It allows us to illustrate how different assumptions about beliefs in these models affects asset pricing outcomes.

Let us consider a particular asset with a (potentially) stochastic payout stream D_t and let P_t denote the (ex-dividend) price of the asset in period $t \geq 0$. The asset may have payouts over a finite time horizon only, as is typically the case with bonds, or it can be infinitely lived, as is the case with stocks.

To price the asset, structural economic models must determine—at a minimum—the following three elements: (1) the stochastic process determining the set of marginal agents $\{\mathcal{M}_t\}_{t=0}^\infty$ pricing the asset in all periods and contingencies, (2) the one-step-ahead stochastic discount factor (SDF) $\{M_{t+1}^m\}_{t=0}^\infty$ that discounts period $t+1$ payouts into period t of at least one marginal agent, and (3) this marginal agent’s (subjective) probability measures \mathcal{P}_t^m , which provides in each period t the perceived probability distribution over the next period’s asset price P_{t+1} and payout D_{t+1} .

The probability measure \mathcal{P}_t^m is part of a probability space $(\Omega, \mathcal{S}, \mathcal{P}_t^m)$, where the space of outcomes Ω contains (among other things) the infinite sequence of price and payout outcomes $(P_0, D_0, P_1, D_1, \dots)$ and where \mathcal{S} is the sigma-algebra of all Borel subsets of Ω . The probability measure is a model primitive under subjective beliefs that needs to be specified by the modeler. In a setting with dynamically consistent beliefs, we have $\mathcal{P}_t^m = P^m$, even in the presence of learning. We allow here for belief specifications where the probability measure varies over time. Such time dependence arises, for instance, when agents forget about old data or face memory constraints.

Element (1) allows for the possibility that not all agents are marginal at all points time or in all contingencies. This is the case whenever agents are constrained in their portfolio choices. Whether or not an agent is constrained depends also on the agent’s beliefs. Optimistic agents, for instance, would perhaps prefer taking a levered position in the asset, but a leverage constraint may prevent them from doing so. Conversely, pessimistic agents would perhaps like to take a short position, but a short-selling constraint may prevent them from doing so. Belief heterogeneity thus interacts with portfolio constraints to determine the set of marginal agents \mathcal{M}_t .

Element (2), which captures marginal agents' stochastic discount factor, allows for heterogeneity in investor preferences and in the optimal consumption plans. The agents' optimal consumption plan also depends on the agent's beliefs.

Beliefs are captured by element (3). For asset pricing with subjective beliefs, this element is of particular importance. Unlike in RE models, the true law of motion that generates payouts does not pin down beliefs about future payouts and prices. From a theoretical viewpoint, this requires taking a stand on the mechanism by which investors form beliefs. Empirically, the lack of a tight link to the true law of motion raises the question whether an assumed belief specification is plausible. Expectations data from surveys can help answer this question.

Given the three elements introduced above, one can price the asset. Let m be an index for marginal agents $m \in \mathcal{M}_t$ and let E_t^m denote the expectation of marginal agent m , as determined by her beliefs \mathcal{P}_t^m . The asset price P_t then satisfies in each period $t \geq 0$ the first-order necessary condition for optimality of marginal agent m , i.e.,

$$P_t = E_t^m[M_{t+1}^m(P_{t+1} + D_{t+1})]. \quad (1)$$

Economic models differ in the way they determine who is marginal, in the way they model the marginal agents' discount factor M_{t+1}^m , and in the way they assign beliefs \mathcal{P}_t^m to marginal agents. Nevertheless, they have a common core in terms of equation (1), which we will use as the starting point of our discussion.

Without loss of generality, we can write marginal agent m 's SDF as

$$M_{t+1}^m = \delta_t^m \xi_{t+1}^m, \quad \text{where} \quad E_t^m[\xi_{t+1}^m] = 1 \quad (2)$$

and hence δ_t^m controls the conditional mean of the SDF under agent m beliefs while ξ_{t+1}^m captures the variation of the SDF across states of the world. The pricing equation (1) then becomes

$$P_t = \delta_t^m E_t^m[D_{t+1} + P_{t+1}] - \delta_t^m \text{cov}_t^m(D_{t+1} + P_{t+1}, \xi_{t+1}^m). \quad (3)$$

Suppose there also exists a risk-free asset with unit payoff at $t + 1$ that all agents that are marginal for the risky asset have access to.¹ Let $R_{f,t}$ denote the gross return of this risk-free asset. The pricing equation for this risk-free asset then implies

$$\frac{1}{R_{f,t}} = \delta_t^m \quad (4)$$

¹We also assume that all agents understand that this asset has a unit payoff.

for each agent m . This means that the agents must adjust their portfolios—by trading in the risk-free and/or the risky asset—such that the δ_t^m are equalized for all marginal agents m :

$$\delta_t^m = \delta_t. \quad (5)$$

Equation (3) thus simplifies further to

$$P_t = \delta_t E_t^m [D_{t+1} + P_{t+1}] + \delta_t \text{cov}_t^m (D_{t+1} + P_{t+1}, \xi_{t+1}^m). \quad (6)$$

The first term captures the agent’s subjective payoff and price expectations, discounted with the conditional mean of the SDF. The second term represents a subjective risk premium. Lower covariance of payoffs and prices with ξ_{t+1}^m implies a higher required risk premium and hence a lower price. Defining the gross return of the risky asset as $R_{t+1} = (P_{t+1} + D_{t+1})/P_t$, equations (4) and (6) yield an expression for the subjectively expected excess return

$$E_t^m [R_{t+1}] - R_{f,t} = -\text{cov}_t^m (R_{t+1}, \xi_{t+1}^m). \quad (7)$$

2.1 Rational expectations

The vast majority of asset pricing models in the literature assumes that investors hold RE. The RE assumption is stronger than the assumption of individual rationality in belief formation. Individual rationality implies that agents update subjective beliefs using Bayes’ rule and that they make optimal decisions given their subjective beliefs about variables beyond their control. RE additionally requires that all subjective distributions coincide with the objective distributions implied by the asset pricing model in equilibrium (Sargent (?)). With RE, the expectations that show up in our pricing equation (6) become

$$E_t^m [D_{t+1} + P_{t+1}] = E [D_{t+1} + P_{t+1} | \mathcal{J}_t] \text{ for all } t \geq 0, \quad (8)$$

where $E[\cdot | \mathcal{J}_t]$ denotes objective expectations given the information set \mathcal{J}_t available to agents at time t . More precisely, agents are endowed with knowledge on how to calculate the density of the payoff D_{t+1} conditional on \mathcal{J}_t , which means that they know the functional form and parameters of this density. They also know the function that maps agents’ beliefs about future payoffs into the equilibrium price, which allows them to form objective conditional expectations of P_{t+1} (Adam and Marcet (?)).

Under RE, equation (7) then implies

$$E [R_{t+1} | \mathcal{J}_t] - R_{f,t} = -\text{cov} (R_{t+1}, \xi_{t+1}^m | \mathcal{J}_t), \quad (9)$$

which shows that the risk premium must be the same for all marginal agents.

The RE assumption is convenient in several ways. First, RE greatly simplifies asset pricing by removing the need to separately specify how agents form beliefs. There is no need to study subjective expectations data to understand the belief formation mechanism. Given a model of the economy, the model-consistency requirement of RE pins down agents beliefs.² Whether the assumed beliefs are empirically plausible is a different question.

Second, RE is convenient for econometric evaluation. While an econometrician outside the model may not be able to observe $E[\cdot|\mathcal{J}_t]$, simply because the econometrician's information set \mathcal{A}_t is smaller than the agents', the econometrician can make use of the fact that with $\mathcal{A}_t \subseteq \mathcal{J}_t$ the law of iterated expectations (LIE) holds, i.e., $E[E[\cdot|\mathcal{J}_t]|\mathcal{A}_t] = E[\cdot|\mathcal{A}_t]$. Taking conditional expectations of equation (9), the econometrician can approximate the risk-premium as³

$$\text{cov}(R_{t+1}, \xi_{t+1}^m | \mathcal{A}_t) = E[R_{t+1} | \mathcal{A}_t] - R_{f,t}. \quad (10)$$

Since these expectations are consistent with the underlying economy generating the data, a sufficiently large sample of empirical data will allow the econometrician to use empirical moments to approximate the population moments on both sides of this equation. The econometrician can then test statistically whether the equality implied by (10) holds.

While RE offers these convenient simplifications, the RE assumption is rather strong and generates empirically unattractive features. For instance, RE implies that expectations are homogeneous and conditionally unbiased. As we will discuss, these predictions are difficult to square with empirical evidence on beliefs. Therefore, we now turn to approaches that allow subjective beliefs to differ from RE. But this means giving up on some of the convenient properties of RE. Expectations data then becomes a crucial input for implementation and evaluation of these asset pricing models.

2.2 Subjective beliefs in a single-period setting

To highlight some key relationships in subjective beliefs models in the simplest possible setting, we start with a risky asset with a maturity of one period that pays a single cash flow of D_{t+1} in $t + 1$. The asset pricing equation (6)

²This assumes that the model does not allow for the presence of rational bubbles. In such cases, the RE assumption does not uniquely pin down expectations.

³To see this, note that $\text{cov}(R_{t+1}, \xi_{t+1}^m | \mathcal{J}_t) = E[R_{t+1}\xi_{t+1}^m | \mathcal{J}_t]$ because $E[\xi_{t+1}^m | \mathcal{J}_t] = 0$. Therefore, $E[\text{cov}(R_{t+1}, \xi_{t+1}^m | \mathcal{J}_t) | \mathcal{A}_t] = \text{cov}(R_{t+1}, \xi_{t+1}^m | \mathcal{A}_t)$.

then simplifies to⁴

$$P_t = \delta_t E_t^m[D_{t+1}] + \delta_t \text{cov}_t^m(D_{t+1}, \xi_{t+1}^m). \quad (11)$$

The first term captures the agent's subjective payoff expectations, discounted with the conditional mean of the SDF. The second term represents a subjective risk premium. From these simple equations, one can derive several equilibrium relations that must hold in asset-pricing models with subjective beliefs for one-period assets with a single payoff.

2.2.1 Homogeneous subjective beliefs

We first consider the case of belief homogeneity among marginal investors. Equation (11) tells us that homogeneity in $E_t^m[D_{t+1}]$ implies homogeneity in the subjectively required risk premium $-\text{cov}_t^m(D_{t+1}, \xi_{t+1}^m)$ among marginal investors. All marginal agents must thus have adjusted their portfolios and consequently their SDFs such that the subjectively required risk premia are the same.⁵

With pricing under investors' subjective beliefs, the interpretation of empirical data can then be very different compared to a setting where the econometrician assumes investors have RE. In particular, there is a difference between *subjective risk premia*, perceived by the agents pricing the assets, and *objective risk premia* extracted by an econometrician studying empirical data ex post. Consider an econometrician who uses data on realized returns $R_{t+1} = D_{t+1}/P_t$ and a statistical model to approximate $E[R_{t+1}|\mathcal{A}_t]$. Taking expectations of the return definition under the econometrician's and the agents' beliefs, and comparing these expectations, we obtain

$$E[R_{t+1}|\mathcal{A}_t] - R_{f,t} = E_t^m[R_{t+1}] - R_{f,t} + \frac{E[D_{t+1}|\mathcal{A}_t] - E_t^m[D_{t+1}]}{P_t}. \quad (12)$$

Therefore, if the econometrician observes, for example, a high objective risk premium $E[R_{t+1}|\mathcal{A}_t] - R_{f,t}$, this does not imply that agents necessarily demanded a high subjective risk premium $E_t^m[R_{t+1}] - R_{f,t}$ when they priced the asset. The high objective risk premium could, instead, be a manifestation of agent pessimism about the future payoffs, which gives rise to a positive belief wedge $E[D_{t+1}|\mathcal{A}_t] - E_t^m[D_{t+1}]$, as, e.g., in Cogley and Sargent (?). To disentangle the effects of risk aversion, perceived risk, and beliefs about payoffs on the objective risk premia, researchers thus need direct measurements

⁴This assumes that investors understand that $P_{t+1} = 0$, which is the case in the absence of pure bubbles.

⁵If the asset market is complete, agents' SDFs must be equal state by state.

of payoff expectations from survey data. Econometric analysis of asset price data alone cannot provide such a decomposition.

The pricing equation under subjective beliefs (11), in conjunction with the expression for expected returns (7), also reveals several fundamental properties of the relation between payoff and return expectations in a homogeneous beliefs equilibrium. While it may seem intuitive that more optimistic subjective expectations of terminal payoffs would also imply higher subjective return expectations, this is not true unless high $E_t^m[D_{t+1}]$ is also accompanied by higher perceived risk or risk aversion, and hence greater magnitude of the subjective risk premium term in (7): Optimistic payoff expectations at time t will lead to higher prices P_t , but, according to (7), not in itself to higher return expectations. Equilibrium requires that the asset will be priced such that $E_t^m[R_{t+1}] - R_{f,t}$ is equal to the risk premium that investors demand to hold the supply of the asset given their risk aversion and subjective perception of risk.

This is already a hint that it is not entirely straightforward to devise equilibrium models in which subjective return expectations vary a lot over time. Any such variation would have to be associated with time-variation in perceived risk and/or risk aversion.

2.2.2 Heterogeneous subjective beliefs

We now allow for heterogeneity in beliefs about D_{t+1} . Equation (11) then again delivers important restrictions. It tells us that heterogeneity in beliefs about payoffs must be accompanied by heterogeneity in subjective risk premia. For example, if we pick two marginal agents, A and B , with payoff expectations $E_t^A[D_{t+1}] > E_t^B[D_{t+1}]$, equation (11) tells us that agent A must also demand a higher subjective risk premium than B such that both agents can agree on the same price P_t . One way in which this can play out in equilibrium is that agent A would devote a larger share of her portfolio to the risky asset than B , which generates a higher required risk premium that then coincides with the subjectively perceived risk premium (Martin and Papadimitriou (?)).

If agents' differences in payout beliefs vary over time, their risky asset exposures, and consequently their subjective risk premia vary over time as well. This opens up a channel for time-varying beliefs about payouts to generate time-varying expectations of excess returns.

This also means that in a heterogeneous belief setting, with a cross-section of agents, there is a tight *cross-sectional* relationship between cash flow expectations and return expectations. Disagreement about $E_t^m[D_{t+1}]$ is reflected one for one in disagreement about future returns. This is very

different from the absence of a *time-series* relation between $E_t^m[D_{t+1}]$ and return expectations that we noted above in the homogeneous belief setting.

Since a heterogeneous beliefs equilibrium may require that agents hold heterogeneous portfolios, it may also happen that some agents' desired portfolio is not feasible due to portfolio constraints. Especially in models with risk-neutral agents, portfolio constraints play a crucial role for a heterogeneous-belief equilibrium to exist. Without such constraints, risk-neutral agents with heterogeneous payoff expectations would want to take infinitely sized bets against each other. Another way to see the non-existence issue is to recognize that the SDF in this case is conditionally deterministic, $M_{t+1}^m = \delta_t$, and hence the pricing equation (11) simplifies to $P_t = \delta_t E_t^m[D_{t+1}]$. Since the left-hand side is the same for all agents, the right-hand side must be the same as well. This leaves no room for differences in $E_t^m[D_{t+1}]$ between marginal agents. The only way for equilibrium to exist is that all but a subset of agents with homogeneous beliefs are not marginal because they are stuck at constraints, e.g. on leverage for optimists and short-selling for pessimists (Geanakoplos (?))

2.3 Subjective beliefs in a multi-period setting

We now turn to assets that offer a payout stream D_t over multiple periods, potentially extending to infinity. The key difference to the case with a single-period assets is that agents can now buy or sell the asset at market prices in intermediate periods. This causes expectations about future market prices to become relevant for equilibrium pricing outcomes.

Defining capital gains as

$$\beta_{t+1}^P \equiv P_{t+1}/P_t \quad (13)$$

and using equation (4), we can express subjectively expected excess returns as

$$E_t^m \left[\frac{P_{t+1} + D_{t+1}}{P_t} \right] - R_{f,t} = E_t^m[\beta_{t+1}^P] + \frac{E_t^m[D_{t+1}]}{P_t} - \frac{1}{\delta_t}. \quad (14)$$

Expected excess returns now have three components: (i) contributions from expected capital gains; (ii) contributions from next-period's payouts; and (iii) the risk-free interest rate, which depends on the expected value of the discount factor.

Using equation (7), we can replace the left-hand side of (14) with $-\text{cov}_t^m(\xi_{t+1}^m, R_{t+1})$ and solve for P_t , which yields

$$P_t = \frac{E_t^m[D_{t+1}]}{\frac{1}{\delta_t} - \text{cov}_t^m(\xi_{t+1}^m, R_{t+1}) - E_t^m[\beta_{t+1}^P]}. \quad (15)$$

Naturally, the equilibrium asset price depends positively on expected payouts, $E_t^m[D_{t+1}]$, and negatively on the subjective risk premium, $-\text{cov}_t^m(\xi_{t+1}^m, R_{t+1})$ and the risk-free rate $\frac{1}{\delta_t}$. In addition, expectations about future capital gains can now potentially affect the current price level in this multi-period setting.

Suppose, for example, that $E_t^m[D_{t+1}]$, δ_t , and the subjective risk premium $-\text{cov}_t^m(\xi_{t+1}^m, R_{t+1})$ stay unchanged, but investors revise upward their view of $E_t^m[\beta_{t+1}^P]$. Equation (15) tells us that P_t will rise in this case. Intuitively, since $\text{cov}_t^m(\xi_{t+1}^m, R_{t+1})$ and δ_t stay fixed, the total expected return $E_t^m[R_{t+1}]$ must stay fixed, too. For this unchanged $E_t^m[R_{t+1}]$ to be consistent with a rise in expected capital gains, P_t must rise sufficiently such that a decline in the expected payout yield $E_t^m[D_{t+1}]/P_t$ exactly offsets the rise in expected capital gains.

Clearly, there is a limit to how much capital gain expectations can rise in this example. As the dividend cannot be negative, it must be that $E_t^m[\beta_{t+1}^P] < E_t^m[R_{t+1}]$, otherwise an equilibrium does not exist. In general, however, $\text{cov}_t^m(\xi_{t+1}^m, R_{t+1})$ and δ_t , and hence $E_t^m[R_{t+1}]$, may also respond in equilibrium to a change in expected capital gains, ensuring existence. How much P_t moves in response to changing capital gains expectations, and how the subjective risk premium and the discount factor respond depends on the specifics of the model (preferences, technology, and belief dynamics).

While the multi-period setting thus opens room for capital gains expectations as a source of variation in asset prices aside from payout expectations, subjective risk premium, and the discount factor, this does not uncouple subjective expected excess returns from perceived risk and risk aversion. Equation (7) always holds, which makes subjective expected excess return to be exactly equal to $\text{cov}_t^m(\xi_{t+1}^m, R_{t+1})$ in equilibrium. Time series variation in expected capital gains $E_t^m[\beta_{t+1}^P]$ thus can only affect subjective excess return expectations if they also generate time series variation in perceived risk or risk aversion. Intuitively, when subjective expected excess returns are high (and perhaps higher than objectively justified), it must be that perceived risk and/or risk aversion is high, otherwise the high expected excess returns would be an unexploited investment opportunity, which can not be an equilibrium outcome.

2.3.1 Common knowledge

Whether price expectations play an independent role as a source of asset price variation in addition to payout expectations depends on the assumptions about how investors form beliefs. Specifically, the crucial question is whether it is common knowledge among agents (i.e., they know, and they know that other agents know, and they know that other agents know that other agents

know, ...) that each period t the asset price is determined according to equation (1) and that the discount factor is given by equation (5).

To simplify the exposition and to focus on the key differences to the setting with a single payoff, we consider a setting with risk-neutral marginal agents, which have the same constant SDF at all times:

$$M_{t+1}^m = \delta \in (0, 1) \text{ for all } t \geq 0. \quad (16)$$

Equipped with common knowledge, the marginal agent can iterate forward on the market-pricing equation (1) and express the equilibrium asset price as a generalized discounted sum of subjectively expected payouts:⁶

$$P_t = \delta E_t^m[D_{t+1}] + \delta^2 E_t^m[E_{t+1}^m[D_{t+2}]] + \delta^3 E_t^m[E_{t+1}^m[E_{t+2}^m[D_{t+3}]]] \dots \quad (17)$$

Equation (17) involves the marginal agent's (first-order) expectation about the payout in the next period and higher-order expectations of future payouts, i.e., expectations of future marginal agents' payout expectations, expectations of future marginal agents' expectations of future marginal agents' expectations, and so on.

Equation (17) implies that subjective price expectations in equation (1) reflect higher-order payout expectations, i.e.,

$$E_t^m[P_{t+1}] = \delta E_t^m[E_{t+1}^m[D_{t+2}]] + \delta^2 E_t^m[E_{t+1}^m[E_{t+2}^m[D_{t+3}]]] \dots \quad (18)$$

Without further assumptions on what agents know or believe about other agents' expectations, the higher-order payout expectations remain undetermined and the same holds true for the subjective price expectations $E_t^m[P_{t+1}]$. In particular, absent further assumptions, the marginal agents' *first-order expectations* about future payouts ($E_t^m[D_{t+j}]$, $j \geq 1$) generally fail to determine the marginal agents' *higher-order expectations* about future payouts and thus also fail to determine the subjective price expectations ($E_t^m[P_{t+1}]$), see Adam and Marcet (?).

When beliefs are homogeneous and dynamically consistent and it is common knowledge that agents share the same subjective payout expectations, the Law of Iterated Expectations (LIE) holds also *across* agents, despite expectations being subjective. For instance, we then have⁷

$$E_t^m[E_{t+1}^m[D_{t+2}]] = E_t^m[D_{t+2}].$$

⁶For the terminal price to disappear, it must also be common knowledge that the expected discounted terminal price is equal to zero under all agents' beliefs.

⁷Recall, $E_{t+1}^m[\cdot]$ denotes the expectations operator that uses the beliefs of an agent that is marginal in time $t+1$. This agent can differ from the one that is marginal in time t , whose expectations are given by $E_t^m[\cdot]$.

The general pricing equation (17) then again greatly simplifies to

$$P_t = \delta E_t^m [D_{t+1} + \delta D_{t+2} + \delta^2 D_{t+3} + \dots]. \quad (19)$$

Price fluctuations are now driven by fluctuations in the marginal agents' *subjective* (first-order) payout expectations. Under the stated assumptions, an asset-pricing model requires a specification of the dynamics of individual agents' own first-order payout expectations, but there is no need for a separate modeling of agents' higher-order expectations and price expectations, as common knowledge and the LIE implicitly determine them from first-order payout expectations.

An example of an asset pricing setup that fits into this framework is a Bayesian learning model with common priors and common knowledge in which agents learn about the properties of an exogenous payout process, possibly coupled with constraints on portfolios that lead to changes in the set of marginal agents in different contingencies.

However, once we entertain belief dynamics that deviate from Bayesian rationality, it is not obvious that subjective beliefs necessarily obey the LIE. For example, if agents update beliefs like a Bayesian, but their memory of past data fades over time, the LIE no longer holds (Nagel and Xu (?)). When agents form diagnostic expectations as in Bordalo et al. ((?)), whether the LIE holds depends on parameter values (see Bianchi et al. (?)). In this case, even with homogeneous beliefs and common knowledge, higher-order payout expectations cannot be eliminated from the pricing equation (17) without further assumptions.

Asset pricing then requires taking a stand on how agents' think about the beliefs of future marginal agents, i.e., how $E_t^m [E_{t+1}^m [D_{t+2}]]$ and other expectations in equation (17) are determined. One possibility, as in Nagel and Xu (?), is to assume that in period t agents form expectations about time $t + 1$ beliefs of agents in a rational Bayesian fashion and that future agents will do the same. For example, in the case of fading memory, this would imply that time t agents rationally anticipate that time $t + 1$ will have experienced some loss of memory of the data that is known to time t agents. This leads to a specific version of the pricing equation (17) with a chain of nested expectations.

Another possibility is to assume that agents at t believe that future agents' expectations are formed in a way that LIE applies. In the fading memory model, this would mean assuming that agents at t believe that for this point in time onwards, agents' memory will not fade further, and hence the LIE can be applied to future agents' expectations. Under either interpretation higher-order expectations collapse to first-order expectations. For example,

we get $E_t^m[E_{t+1}^m[D_{t+2}]] = E_t^m[D_{t+2}]$, so that the pricing equation collapses to one like in (19) where only first-order expectations of future payouts appear. Bordalo et al. (?) take this route with diagnostic expectations.

Our discussion showed that theory restricts higher-order expectations tightly in this homogeneous beliefs, common knowledge setting. We need to make an assumption whether agents are sophisticated about their behavioral limitation and anticipate that future selves will have them, too, or whether they naively anticipate that future agents won't have them.⁸ But once this choice is made, the law of motion for first-order expectations also pins down agents' higher-order expectations. Which of these assumptions is a better description of agents' expectations formation is ultimately an empirical question. In particular, it remains to be explored to what extent agents' first-order dividend expectations are in fact tightly linked to agents' expectations about the future asset price, as implied by these approaches.

The asset pricing literature also studied settings in which agents hold heterogeneous subjective first-order expectations about payouts and 'agree to disagree' about future payouts. Differences in payout beliefs can arise from differences in subjective prior beliefs (Harrison and Kreps (?)) or differences in the way incoming information is interpreted (Dumas, Uppal, Kurshev (?)).

With risk averse agents that do not face portfolio constraints, all agents are marginal at all times, even in a heterogeneous belief setting. In this case, if the LIE holds for individual beliefs, higher-order expectations of stochastically discounted payouts collapse to first-order expectations. As a consequence, heterogeneity affects asset prices through uncertainty about agents' future portfolio positions and hence their SDFs, but not through higher-order expectations (Martin and Papadimitriou (?)).

With risk-neutral agents, as in our illustrative example here, portfolio constraints or specific forms of market incompleteness are required to ensure existence of equilibrium. In this case, the identity of the marginal agent can change over time and higher-order payout expectations do not collapse to first-order expectations.

To insure tractability, models where agents agree-to-disagree assume identical conditioning information (even if information signals are interpreted differently by agents) and common knowledge of expectations. These assumptions imply that agents possess the same information about their own and other agents' beliefs, so that higher-order expectations do not depend on the identity of the marginal agent (e.g. Harrison and Kreps (?)). The

⁸O'Donoghue and Rabin (?) discuss a similar choice between sophistication and naivete about the behavior of future selves in the context of self-control problems.

general pricing equation (17) then simplifies to

$$P_t = \delta E_t^m[D_{t+1}] + \delta^2(E[E_{t+1}^m[D_{t+2}]|\mathcal{J}_t] + \delta^3(E[E_{t+2}^m[D_{t+3}]|\mathcal{J}_t] + \dots), \quad (20)$$

where $E[.|\mathcal{J}_t]$ denotes agents' common and rational expectation based on time t information. Asset prices now depend on the marginal agents' subjective one-step-ahead payout expectations and on rational expectations of future marginal agents' first-order payout expectations.

Harrison and Kreps (?) show that the equilibrium asset price can then exceed (in the presence of short-sale constraints) the subjectively expected discounted value of future payout of all agents in the economy, i.e., we have

$$P_t \geq E_t^i[\delta D_{t+1} + \delta^2 D_{t+2} + \delta^3 D_{t+3} + \dots]$$

for all investors i , unlike in the case with common subjective payout expectations in equation (19).

2.3.2 Lack of common knowledge

Common knowledge about a high-dimensional object such as all other investors' (first-order) payout expectations is a somewhat implausible starting point in terms of descriptive realism. How could investors possibly be sure about how other investors form beliefs about payoffs and about what other investors believe, etc.?

If we abandon the assumption that (first-order) payout expectations are common knowledge, asset pricing looks quite different. The higher-order payout expectations in (17) are then no longer determined by first-order payout expectations (or rational expectations thereof).

One approach to asset pricing without common knowledge is to explicitly model the (high-dimensional) process of higher-order payout expectations. An alternative approach is to directly model the process for subjective (first-order) capital gains expectations that appear in equation (15), which sidesteps the need to specify how higher-order expectations are formed. Under lack of common knowledge, the latter approach is consistent with individual rationality and rational belief formation (Adam and Marcet (?), Adam et al. (?)). From a modeling perspective, there are some advantages of this approach. First-order capital gains expectations are a considerably more tractable object than higher-order payout expectations. In addition, while higher-order payout expectations are very hard to observe empirically⁹, first-order capital gains expectations are regularly included in investor surveys.

⁹One exception is recent work by Coibion et al. (?) that provides survey data on higher-order macroeconomic expectations of firm managers.

Therefore, subjective capital gain expectations in asset pricing models can be disciplined by survey data.

Models with subjective beliefs about prices can generate strong belief-based amplification of asset price volatility. As we highlighted in our discussion of equation (15), a rise in capital gain expectations can generate a rise in asset prices. Therefore, if subjective capital gain expectations are positively influenced by observed past capital gains, as suggested by survey data that we discuss in the next section, then past price increases generate optimism about future capital gains and thus a further rise in asset prices. The resulting propagation over time allows models without common knowledge about payout beliefs to replicate the observed large volatility of stock prices, even in a setting with standard time-separable utility functions (Adam, Marcet and Nicolini (?)).

This dynamic feedback from past price changes to future prices is absent in setups in which agents hold subjective beliefs only about exogenous objects, e.g., exogenous payouts D_t that are independent of agents' beliefs. This absence of feedback effects makes it considerably harder to generate the empirically observed high volatility of stock prices.

Abandoning the assumption of common knowledge of (first-order) payout expectations provides economic models with additional degrees of freedom in specifying subjective capital gains expectations. While these can be disciplined with the help of investor survey data, the specification of subjective capital gains beliefs, and their relation to beliefs about payouts, is also subject to a set of restrictions generated from theory.

First, subjective beliefs must be consistent with agents' own optimality conditions, so as to have a well-defined agent problem. Adam and Marcet (?) refer to such beliefs as internally rational beliefs. In some special cases, internal rationality implies that capital gain beliefs cannot be specified independently of payout beliefs. For instance, when agents are risk-neutral and *know* to be marginal at all times, then the individual optimality condition, $P_t = \beta E_t^m [P_{t+1} + D_{t+1}]$, holds in all periods and all contingencies. If the LIE holds for individual agent beliefs, the agent can then forward-iterate on her *own* first-order condition to arrive at the asset pricing equation (19), without relying on common knowledge assumptions. Specifying price beliefs that differ from the ones implied by equation (19) would then lead to a situation where agents' first-order conditions are violated. More generally, however, when agents are risk-averse, subjective price beliefs cease to be determined by first-order payout beliefs, even when agents know to be marginal at all times (Adam, Marcet and Beutel (?)). This is the case because the stochastic discount factor is then endogenous to price beliefs, unlike in the special case with risk-neutrality.

Second, independently specifying subjective beliefs for returns and payouts, instead of specifying them for capital gains and payouts sharpens the non-existence problem we discussed following equation (15). To ensure existence, $\frac{1}{\delta_t} - \text{cov}_t^m(\xi_{t+1}^m, R_{t+1})$ may have to adjust. However, using equation (7) we see that fixing return expectations directly fixes $\frac{1}{\delta_t} - \text{cov}_t^m(\xi_{t+1}^m, R_{t+1})$, so this adjustment mechanism is not available. For this reason, one should formulate beliefs as subjective probability distributions over prices and payouts and refrain from formulating them over returns and payouts.

Third, since equation (7) ties subjective expected excess returns to a covariance with the SDF, the subjective expected excess returns can only change if subjectively perceived risk and/or risk aversion changes. Such changes in risk can arise when the SDF responds endogenously to a change in capital gains expectations.

3 Empirical dynamics of investor expectations

The asset pricing models with subjective beliefs that we sketched in the previous section require assumptions about belief formation. Unlike in RE models, the objective law of motion of the variables driving payoffs and the SDF does not pin down investors' subjective beliefs. Researchers must therefore make additional assumptions about how agents form beliefs. These assumptions in turn should be informed by empirical evidence on the dynamics of investors' subjective expectations about future prices or returns, asset cash flows, future interest rates, and beliefs about risks and higher moments. We now provide a brief overview of existing empirical evidence on the dynamics of investor expectations. The appendix at the end of this chapter lists the data sources for most of the empirical studies we discuss in this section.

3.1 Return and price expectations

Most empirical studies of investor subjective beliefs have focused on expectations of future returns and prices, and especially on expectations of returns and price levels of aggregate stock market indices.

For individual retail investors, expectations of stock market returns over the next year appear strongly related to past returns that these investors have experienced. Using the UBS/Gallup survey, Vissing-Jorgensen (?) finds a positive relation between expected returns and the (self-reported) returns of investors' own portfolios in the past. Using the same survey, Malmendier and Nagel (?) show that at the cohort level, subjective expectations of future

stock market returns are positively related to a weighted average of the lifetime stock market returns experienced by an individual's birth-cohort.

As we discussed in the previous section of this chapter, such a cross-sectional relationship between investors' experienced past returns and their expected returns indicates either that subjective expected cash flow expectations are positively related to past returns, or that, in absence of common knowledge, price expectations decoupled from cash flow expectations are related to experienced past returns.

From an equilibrium asset pricing perspective, the aggregate dynamics of subjective return expectations over time are more important than these cross-sectional relationships. Yet, it is not obvious that the apparent extrapolation from past returns in the cross-section of survey respondents also translates into a positive time-series relation between past returns when expectations are averaged across individuals. As we discussed in the earlier sections, if investors are generally optimistic about future cash-flows or the level of future prices, this generates high current prices and does not necessarily lead to high subjective expected returns.

Empirically, however, past returns and return expectations are positively related at the aggregate level, too. Using the UBS/Gallup survey and several other surveys, Greenwood and Shleifer (?) find a positive relation between the average return expectation of individuals and returns over the most recent 12 months and the log price-dividend ratio of the stock market portfolio.

Total return expectations can change over time due to movements in the dividend yield or the expected capital gain. Similarly, a change in the total expected return can reflect changes in the risk-free rate or changes in the subjective risk premium perceived by investors. To relate the expectations data to asset pricing models, it is important to understand the behavior of these different components.

Adam, Marcet, and Beutel (?) focus on capital gains expectations. They show, using econometric tests that account for small-sample biases, that there is a positive relation between expected stock market capital gains and the price-dividend ratio and that this is inconsistent with the RE hypothesis. Using the UBS/Gallup survey, the Yale/ICF survey of individual investors, and the Graham-Harvey survey of Chief Financial Officers, they find that the relation between the price-dividend ratio and future realized capital gains is considerably more negative than the relation between the price-dividend ratio and survey expectations of capital gains at all forecast horizons, ranging from one to ten years ahead.

Other work focuses on the subjective risk premium. A positive relationship between the price-dividend ratio and total subjective expected returns does not necessarily imply that the subjective risk premium is positively re-

lated to the price-dividend ratio. To understand whether perceived risk or risk aversion is changing with the price-dividend ratio one must isolate the subjective risk premium component. Bacchetta, Mertens and Van Wincoop (?) and Nagel and Xu (?) show, when return expectations are measured in excess of Treasury yields, i.e., as a subjective risk premium, they exhibit only a weak positive relationship with variables like the price-dividend ratio that capture slow-moving valuation cycles in the stock market. The reason is that interest rates tend to be procyclical so that part of the positive relationship between the price-dividend ratio and subjectively expected total returns is due to the risk-free rate. In contrast, as Nagel and Xu (?) document, the relationship of total expected returns with the past 12-month return works mainly through the subjective risk premium channel.

While it is not entirely clear to what extent individual investors' subjective excess return expectations are procyclical, it is clear that they fail to be countercyclical. Irrespective of whether they are expressed in terms of total returns, excess returns, or capital gains, the dynamics of subjective expectations do not match the countercyclical dynamics implied by predictive regressions of stock returns on valuation ratios and by RE models that are reverse-engineered to fit this predictive regression evidence. As a consequence, individual investors' forecast errors are strongly countercyclical and thus not unpredictable, as the RE hypothesis would imply.

To map the expectations evidence from investor surveys into asset pricing models, one must also take a stand on who the individual investors in these surveys represent. Is the average belief of individual investors a good approximation of the subjective expectations of a representative agent? Then a representative agent model that targets the average individual investor belief will provide a good match with the survey data. Or do investor groups that are excluded from individual investor surveys have systematically different expectations? Then perhaps a heterogeneous-agent approach, or a representative agent approach that targets an aggregate of beliefs across these different investor groups, would be more suitable.

Broadly, the group not covered in the individual investor surveys are professional investors. The available evidence on the dynamic properties of professional investor expectations is mixed. Andonov and Rauh (?) find that pension funds tend to extrapolate past returns: those with higher past performance expect higher risk premia on risky assets. Yet, this is again a cross-sectional result. Using the Yale/ICF survey of U.S. institutional investors, Bacchetta, Mertens and Van Wincoop (?) find that subjective expected excess returns are acyclical. Wu (?) uses return expectations aggregated from analyst price targets and forecasts from the Livingston survey and finds countercyclical (see, also, Wang (?)). Dahlquist and Ibert (?)

examine year stock equity premium expectations of asset managers, CFOs, and professional forecasters. They find countercyclical subjective risk premia for expectations at a one-year horizon. They also find that one-year expectations are not countercyclical enough relative to the RE benchmark, because forecast errors are still countercyclical, although for professional forecasters the forecast error predictability is not statistically significant. For 10-year expectations, they find countercyclical only for asset manager expectations, but not the others. What is not clear yet is how much of the countercyclical-ity of professionals' subjective risk premia is driven by a contrarian effect of recent past returns that are correlated with valuation ratios, i.e., whether it is the mirror image of individuals' apparent extrapolation from recent past returns, or whether it is due to lower-frequency variation.

Overall, the cyclical properties of individual investor subjective expected (excess) returns deviate sharply from the countercyclical-ity implied by ex-post predictive regressions and RE models. Professionals' expectations may be closer to the RE benchmark, but quantitatively it's not clear at this point whether the countercyclical-ity is strong enough, and the variation at the right frequency, to be consistent with RE. In any case, if market equilibrium reflects an average of beliefs that gives both groups substantial weight, the idea that countercyclical risk premia are a main driver of asset price booms and busts cannot be reconciled with the evidence on individual investor expectations.

3.2 Cash flow expectations

Return expectations alone carry only limited information about the link between subjective expectations of investors and asset prices. For example, asset prices could fluctuate wildly in response to volatile subjective expectations of future cash-flow growth or future prices, but, at the same time, subjective return expectations could be constant. Studying return expectations in this case would not reveal the extent to which asset prices are driven by subjective belief dynamics. Data on cash-flow expectations can provide another important piece of the picture how subjective belief dynamics generate asset price volatility.

The number of existing research studying directly subjective cash flow expectations of investors is relatively small, compared to the number of papers that examine return expectations. Moreover, the available evidence at this point is based exclusively on data from surveys of professional forecasters and from aggregated firm-level earnings or dividend forecasts of equity analysts. Existing individual investor surveys do not ask respondents for cash flow forecasts. The cash flow expectations also tend to be relatively crude in the sense that they are provided only for a few forecast horizons rather than

a full term structure.

Chen, Da, and Zhao (?) and De La O and Myers (?) use aggregated equity analyst earnings forecasts to measure subjective cash flow expectations. While their methods differ, they come to broadly similar conclusions. Chen et al. (?) work with a valuation model with a constant discount rate (implied cost of capital). They show that changes in these earnings forecasts go a long way toward explaining the observed movements in aggregate stock prices. De La O and Myers (?) work with a log-linearized approximate present value identity framework which allows decomposing the variance of the log price-dividend (or price-earnings) ratio into the covariance of these valuation ratios with subjective expected dividend (or earnings) growth and the covariance with subjective expected returns. This is the equivalent to the variance decomposition in Campbell and Shiller (?), but under subjective expectations instead of objective expectations implied by predictive regressions. De La O and Myers (?) find that about two thirds of the variation in the price-earnings ratio is attributable to subjective earnings growth expectations. Using a shorter sample of analyst dividend forecasts, they find that almost all variation in the price-dividend ratio is explained by subjective dividend growth expectations.

Perhaps surprisingly, De La O and Myers (?) find that variation in expectations of short-term cash flow growth seem to explain much of the variation in the price-earnings ratio. In contrast, Bordalo et al. (?), working with a similar Campbell-Shiller framework and similar data, come to the conclusion that long-term expectations are an important source of variation in the aggregate stock price level. The fact that likely reconciles these seemingly conflicting findings is that a substantial amount of variation in the price-earnings ratio comes from earnings rather than prices. For example, during the financial crisis, aggregate earnings plunged dramatically and, unlike the price-dividend ratio, the price-earnings ratio spiked up. In the depth of the crisis, analysts expected a strong reversal of this earnings drop, resulting in rise in forecasted earnings growth that coincided with a high price-earnings ratio. A lot of the short-term movements of earnings that affect the price-earnings ratio are largely offset, in terms of valuation implications, by a predictable near-term reversal of these earnings shocks. For this reason, the sources of variation in the price-earnings ratio (that De La O and Myers examine) are quite different from the sources of variation in the price level (that Bordalo et al. examine).¹⁰ For asset pricing, the variation of the price level is the main question of interest, not the effect of short-term earnings

¹⁰This is also why, to focus on movements of the price level, Shiller (?) uses a 10-year trailing moving average of earnings rather than current earnings to form a valuation ratio.

dynamics on the price-earnings ratio.

Furthermore, Bordalo et al. (?) find that variation in subjective long-term expectations is the main source of predictable forecast errors and predictable returns. Relatedly, Nagel and Xu (?) find their ‘experienced dividend growth’ variable—a slow moving exponentially weighted average of past aggregate dividend growth that predicts stock market excess returns—is positively related to analysts’ long-term aggregate earnings forecasts in a way that is consistent with subjective cash flow expectations as the source of asset price variation and predictable returns. Overall, for understanding the wedge between rational expectations forecasts and subjective cash flow growth expectations, the long-term component of these expectations seems to be important.

In summary, the evidence on subjective cash flow growth expectations is broadly consistent with the view that asset price fluctuations reflect, to a large extent, variation in investors’ subjective cash flow growth expectations. This fits well with the evidence from subjective return expectations that countercyclical movements in subjective risk premia cannot be the reason why asset prices are volatile. Taken together, the evidence suggest that high current prices, for example, are not associated with low subjective expectations of future returns, but rather with expectations of high future cash flows and high future prices. That said, the available data on cash flow expectations at this point is still rather limited. While it is reasonable to assume that analysts’ and investors’ forecasts may be closely related, it is unclear to what extent analyst forecasts are representative of investors’ forecasts. Additional data on cash flow expectations of market participants would thus be valuable to analyze this issue further.

3.3 Interest rate expectations

In a multi-period setting, the current price of an asset reflects not only expectations of future prices and cash flows, but also expectations of future discount factors of marginal agents. For default-free bonds with certain cash flows, uncertainty about future discount factors is the only uncertainty that matters for pricing. To see this, consider a zero-coupon bond with a sure payoff of \$1 at time $t + 2$. Iterating once on (1), we obtain

$$P_t = E_t^m [M_{t+1}^m E_{t+1}^m M_{t+2}^m]. \quad (21)$$

Since $1/R_{f,t} = E_t^m[M_{t+1}^m]$, we obtain

$$P_t = \frac{1}{R_{f,t}} E_{t+1}^m \left[\frac{1}{R_{f,t+1}} \right]. \quad (22)$$

Thus, the price of the bond reflects expectations of future short-term interest rates (or yields).

Analogous to stock return predictability with valuation ratios such as the dividend-price ratio, there is evidence in the bond pricing literature that returns on long-term bonds in excess of short-term interest rates are predictable with the spread of yields between long-term and short-term bonds. Rational expectations models explain this predictability of excess returns with time-varying risk or risk aversion. However, just as for stocks, bond return predictability could also be the consequence of predictable forecast errors due to deviations from rational expectations.

Using bond yield expectations of professional forecasters to measure subjective interest-rate expectations, Froot (?) finds that expectational errors contribute substantially to excess return predictability of U.S. Treasury bonds. Cieslak (?) and Piazzesi, Salomao, Schneider (?) find similar results in data that includes more recent decades. In particular, Piazzesi et al. find that subjective expected excess returns implied by the yield forecasts of professional forecasters are substantially less volatile than the forecasts from predictive regressions and do not show much cyclical variation. Thus, similar to stocks, much of the bond price variation that is associated with predictable future returns seems to be driven by subjective expectations, yet in this case not payoff expectations but discount factor or interest rate expectations.

3.4 Subjective risk perceptions

Our discussion of empirical work on expectations data in asset pricing so far focused on first moments. But subjective perceptions of second and higher moments are also relevant for asset pricing. For example, if there is time-variation in subjective risk premia, the economic reason for this variation may be that subjective perceptions of risk are time-varying. Moreover, empirically observed time-variation in risk premia on options and other derivatives—such as the variance risk premium, for example—could potentially reflect predictable forecast errors for asset return variances and higher moments rather than a subjective risk premium that investors priced in *ex ante*. In other words, similar to the case of the time-varying equity premium that we mostly focused on so far, subjective belief dynamics could be the driver of empirically observed time-varying risk premia. Data on subjective risk perceptions of investors can help disentangle these competing explanations,

The available evidence is rather limited at this point. Lochstoer and Muir (?) recently took a first look at this. They use survey data from the Graham and Harvey CFO survey in which respondents are asked to state the 10th and 90th percentile of stock returns over the next year and the Yale/ICF sur-

vey in which respondents provide subjective probabilities of a stock market crash over the next 6 months. Lochstoer and Muir (?) find that investors' subjective stock market risk perceptions seems to be slowly moving, with initial underreaction to volatility shocks and subsequent delayed overreaction. They show that these dynamics in forecast errors provide a potential explanation of the empirical dynamic response of the variance risk premium to volatility shocks. Claims that provide insurance against future volatility appear underpriced immediately following a rise in volatility, which matches the initial underreaction in survey data, and then overpriced later on, which is consistent with the delayed overreaction of subjective risk perceptions.

The elicitation of entire subjective distributions has been introduced recently in other contexts, e.g., in the Survey of Primary Dealers and the Survey of Market Participants of the Federal Reserve Bank of New York. The Survey of Professional Forecasters, as discussed in **Chapter 18, "Inference on Probabilistic Surveys..."**, also provides density forecasts of macro variables. Eliciting similar density forecasts in the context of investor survey would allow for a further and more detailed analysis of the role of risk perceptions.

4 Mapping survey expectations into asset pricing models

The mapping between investors' beliefs in an asset pricing model and the expectations data described in the previous section is often difficult to establish. In particular, the subjective expectations or distributions elicited in surveys may not always provide information about the beliefs of the relevant set of investors inside the model. The remainder of this section discusses the main issues that arise, focusing on the elicitation of expected values. Much of the discussion also applies to surveys that elicit other moments or probability distributions.

We use $\mathcal{E}^i[\cdot]$ to denote individual i 's expectation measured in the survey, which could be distinct from the expectation $E^i[\cdot]$ that the individual truly holds. We focus mostly on household survey data and thus abstract from the strategic considerations and career concerns that may distort reported expectations of professional forecasters. There exists a substantial body of literature dealing with these concern, which is reviewed by Marinovic, Ottaviani, and Sorensen (?).

4.1 Are survey expectations risk adjusted?

When respondents are asked to report an expected value or a probability assessment in a survey, the intention of the survey administrators is to elicit individuals' assessment of physical probabilities, i.e., their assessment of empirical frequencies that are not distorted by risk preference effects. In line with this intention, researchers typically interpret beliefs elicited in surveys as physical measure beliefs.

However, it is at least a theoretical possibility that responses risk preference effects could distort individuals responses to expectations questions in surveys. For example, an individual who is highly risk averse might put more weight on "bad" outcomes in high marginal utility states, reporting more pessimistic expectations than warranted under her subjective assessment of physical probabilities, as also discussed in **Chapter 26, "Looking Ahead to Research Enhancing Measurement of Expectations"**.

Along these lines, Cochrane (?) suggests that individuals might report expectations under the risk-neutral measure and that this could help explain the large wedges between survey expectations and investor expectations implied by rational expectations asset pricing models. This risk-neutral expectations hypothesis states that when individual i reports an expected value in a survey, say the expectation of an asset return R_{t+1} , then this expectation incorporates a risk-adjustment based on the individual's SDF

$$\mathcal{E}_t^i[R_{t+1}] = E_t^i \left[\frac{M_{t+1}^i}{E_t^i[M_{t+1}^i]} R_{t+1} \right] \quad (23)$$

where the ratio pre-multiplying R_{t+1} inside the expectations operator transforms the physical probability of future states, which enter the computation of the expectation $E^i[\cdot]$ into a risk-neutral, or marginal-utility weighted probability. Under this hypothesis, holding $E_t^i[R_{t+1}]$ fixed, greater risk aversion or greater risk of bad outcomes, would induce more pessimistic reported expectations $\mathcal{E}_t^i[R_{t+1}]$.

Whether individuals report physical expectations, risk-neutral expectations, or otherwise risk-adjusted expectations is an empirical question. Adam, Matveev, and Nagel (?) examine the evidence for stock market return expectations from various surveys of individual and professional investors. They first note that if investors have access to trading in the asset that delivers the return R_{t+1} and the risk-free asset with return $R_{f,t}$, then, for such a marginal investor m , their first-order conditions imply

$$E_t^m[M_{t+1}^m R_{t+1}] = 1, \quad E_t^m[M_{t+1}^m R_{f,t}] = 1, \quad (24)$$

which, together with (23) for $i = m$ implies that the risk-neutral expectation of asset returns equals the risk-free rate:

$$\mathcal{E}_t^m[R_{t+1}] = R_{f,t}. \quad (25)$$

This is a testable hypothesis and Adam et al. (?) show that it is strongly rejected in all survey data sets they examine. On average, reported return expectations are much higher than risk-free rate proxies.

A less extreme version of the risk-adjustment hypothesis would imply that individuals reported expectations that are tilted pessimistically towards risk-neutral expectations, but not necessarily all the way. Adam et al. (?) do not find empirical support for such pessimistic tilts. Comparing survey stock market return expectations to subsequently realized stock market returns, they find that unconditionally, survey expectations are close to unbiased, not pessimistically biased.

Overall, the evidence does not provide support for the notion that individuals report risk-adjusted return expectations in surveys. We therefore proceed under the assumption that when individuals respond to expectations questions in a survey, they are reporting $E_t^i[\cdot]$ without risk-adjustments.

4.2 Measurement error and cognitive uncertainty

Even if individuals are attempting to report their expectation under the physical measure, it is still not necessarily true that $\mathcal{E}_t^m[R_{t+1}]$ can be interpreted as a direct measurement of the physical expectation. The expectation that respondents provide in a survey may not be the expectation that individuals would truly hold if they had the time to reflect more carefully on their response and if they had to make decisions based on these expectations in a high-stakes environment. As a consequence, there may be an error component, ε_{t+1}^i , in measured expectations of individual i

$$\mathcal{E}_t^i[R_{t+1}] = E_t^i[R_{t+1}] + \varepsilon_{t+1}^i. \quad (26)$$

For example, as with any variable elicited in a survey, there is the possibility that survey expectations are subject to measurement error. Such measurement error could be the consequence of misunderstanding of the survey question or insufficient deliberation before providing a response; see also the discussion in section 2.4 of Chapter 19, “Expectations data in structural microeconomic models.”

The existing evidence indicates that survey return expectations contain useful information about individuals’ $E_t^i[R_{t+1}]$, but at the same time the

error ε_{t+1}^i is not negligible. In data from the UBS/Gallup survey, Vissing-Jorgensen (?) finds a positive cross-sectional correlation between individual retail investors' expected stock returns and the percentage of their portfolio that they report to hold in stocks. Adam et al. (?) show that the cross-sectional dispersion in expected returns correlates with trading volume over time, i.e., periods with high measured disagreement are periods with more active stock trading, which suggests that beliefs dispersion in surveys is unlikely driven by measurement error alone.

However, using survey data on stock return expectations combined with administrative data on portfolio holdings, Ameriks et al. (?) and Giglio et al. (?) show that the portfolio share of stocks in individuals' portfolios is substantially less sensitive to individuals' stock market return expectations than implied by standard portfolio choice models under plausible values for relative risk aversion.

Measurement error seems to be part of the reason for this low sensitivity. Using instrumental variable techniques that assume measurement error is uncorrelated across different survey questions that elicit return expectations and perceived probability distributions of stock market returns, Ameriks et al. (?) and Giglio et al. (?) find a stronger sensitivity of portfolio shares to expected returns, but it is still weaker than implied by standard models. There seem to be components of ε_{t+1}^i that are common across survey questions and which may not be interpretable as classical measurement error.

One possibility for the low sensitivity, suggested by Drerup, Enke, and von Gaudecker (?), is that individuals have a lack of confidence in their own stated beliefs. Along these lines, Enke and Graeber (?) show that if individuals perceive cognitive uncertainty about what the optimal action is, they may behave as if they shrink probabilities toward a cognitive default. However, it remains unclear whether this shrinkage only affects actions or also the expectations reported in surveys. Would individuals facing cognitive uncertainty respond with these shrunk probabilities or would they report their subjective assessment prior to shrinkage? The fact that Giglio et al. (?) find the sensitivity of portfolio shares to expectations to be stronger for investors who are active, confident in their beliefs, and pay attention is consistent with agents reporting their expectations prior to shrinkage.

Low sensitivity of decisions to expectations at the individual investor level does not necessarily imply low sensitivity at the aggregate level. The identification of sensitivity in Giglio et al. (?) rests on cross-sectional differences in expectations and actions between individuals. It is possible that the effects of measurement error, cognitive uncertainty, and other frictions largely cancel out when expectations and actions are aggregated. The fact, documented in Greenwood and Shleifer (?), that individuals' stock market

return expectations are strongly correlated with aggregate flows into equity mutual funds would be consistent with this latter interpretation.

4.3 Heterogeneity and beliefs aggregation

Heterogeneity of reported expectations is a pervasive feature of survey data.¹¹ This raises the question how researchers should deal with this dispersion when they want to map survey expectations into the expectations of agents in an asset pricing model.

For tractability reasons, many asset pricing models are set up as representative agent models. So which individual beliefs in the survey data should be mapped into the representative agent's beliefs? A typical approach is to take an equally-weighted mean or median of some observed set of expectations. Even if the survey data captured expectations of all investors in the economy, this approach would be subject to some approximation error. Jouini and Napp (?) construct a representative agent in an economy with heterogeneous beliefs and heterogeneous risk tolerance. They show that the representative agent's beliefs are a risk tolerance-weighted average of individual agent beliefs (belief dispersion also has an effect on the representative agent's discount factor and hence the risk-free rate).

In practice, without comprehensive data on risk tolerance and beliefs of the investor population, it is difficult to empirically implement such a weighting scheme. Moreover, additional complications may come into play. Participation constraints keep some individuals out of the market. Lack of attention may render some investors effectively non-participating for certain time intervals. It remains to be seen whether weighting schemes based on observable proxies for risk tolerance and likelihood of market participation can improve the fit between representative-agent asset pricing models and survey data.

For aggregating professional forecasters beliefs is also potentially important to consider that professional forecasters are not directly investing themselves, but their forecast may influence investment decisions of professional investors (who pay for these forecasts). How much influence they have on investors may differ between forecasters and it may depend on their past forecast performance. Accordingly, Buraschi, Piatti, and Whelan (?) construct an aggregated subjective bond risk premium measure that gives more weight to professional forecasters that were more accurate in the past.

¹¹See **Chapter XYZ, The Term Structure of Expectations**, which provides evidence on disagreement about output growth, inflation and interest rates over various forecast horizons.

An alternative approach is to move away from representative agent models to explicitly specify belief heterogeneity. David (?), for example, calibrates an asset pricing model with heterogeneous beliefs to earnings forecast data from the Survey of Professional Forecasters. Likewise, in models in which belief heterogeneity is tied to some observable agent characteristics, researchers can aim for a more detailed comparison of model implied beliefs and survey data beyond broad measures of dispersion. For example, in the model of Collin-Dufresne, Johannes, and Lochstoer (?) agents' beliefs are heterogeneous between age cohorts and they can be compared with cohort-aggregated survey data. Another observable dimension is professional vs. individual investors. Given the heterogeneity in return expectations that we discussed in Section 3.1, it may make sense to consider models that specify different belief dynamics for these two groups of participants.

Clearly, aggregating survey expectations within groups such as age-cohorts or among professionals and individuals, one again encounters the issue of how to weigh the empirically heterogeneous beliefs within each group to approximate the beliefs of each group's representative agent. Moreover, due to the technical difficulty of solving models with heterogeneous agents, representative agent models will likely continue to play an important role. For these reasons, the need to implement some aggregation scheme for survey expectations data is difficult to avoid.

5 Models of expectations formation

A growing body of work develops asset pricing models in which investors' subjective beliefs deviate from RE. Many of these papers aim not only to reproduce the key empirical properties of asset prices, but also the stylized facts about investor beliefs in survey data that we presented in Section 3. The belief specifications differ, but many involve some form of learning where investors use observed data to form expectations about future payouts or prices. Most use a homogeneous-beliefs setup, but we also discuss a few papers that explore the effects of belief heterogeneity.

5.1 Learning about payouts

Learning about dividend payouts can generate volatile asset prices and high risk premia. Belief revisions about the parameters of the payout process contribute to variations in expected payouts and hence to asset price volatility, unlike in full information RE models where the parameters governing the payout process are assumed to be known. Risk premia can be high because

parameter uncertainty contributes to perceived consumption uncertainty in a way that covaries with payouts. Both features help achieve a better fit with asset price data. The belief dynamics also help match data on the dynamics of subjective payout expectations.

Early work in this literature, e.g., Timmermann (?), shows that excess volatility and return predictability can emerge from learning about the payout process. This literature follows Kreps (?) in using an anticipated-utility framework in which the agents pricing assets ignore posterior uncertainty about payout process parameters and the fact that their beliefs will be revised in the future. This means that the additional subjective uncertainty (relative to the case of RE) that investors face in these models is not priced and therefore does not contribute to risk premia.

In Collin-Dufresne et al. (?), Bayesian investors learn about the mean of an i.i.d. log endowment growth process

$$\Delta d_t = \mu + \varepsilon_t, \quad \eta_t \sim N(0, \sigma^2) \quad (27)$$

and they price a consumption claim that has the endowment $D_t = \exp(d_t)$ as payout. As they show, switching to the assumption that investors fully take into account posterior uncertainty when they price assets actually has only very small effects when investors have constant relative risk aversion (CRRA) utility. In other words, for the CRRA case, anticipated utility is a good approximation (see, also, Cogley and Sargent (?)). However, when investors have Epstein-Zin utility, the result can be drastically different. With Epstein-Zin utility, investors demand a large risk premium for uncertainty about long-run endowment growth. As long as investors have not seen enough data to have precise beliefs about long-run growth, risk premia can therefore be very high.

When the parameters of the endowment process are time-invariant and investors prior beliefs reflect this, Bayesian learning has the perhaps unrealistic implication that the learning effects disappear in the long run: asymptotically, the model predictions approach the predictions of an RE model.

Models with perpetual learning avoid this outcome. Perpetual learning can arise for a number of reasons. For example, if investors believe, based on their prior, that there is time-variation in the payout process parameters, this may lead them to discount observations in the distant past as seem of little relevance for current parameter values. For example, if μ in equation (27) is not constant, but instead follows a random walk, a Bayesian investor's posterior mean would be an exponentially-weighted average, with fixed weights, of past endowment growth rates. As a consequence, the posterior mean would continue to drift forever and subjective uncertainty about μ would never disappear.

In Nagel and Xu (?) learning is perpetual because agents have slowly fading memory of past growth rate observations. This also gives rise to persistently high uncertainty about endowment growth. Nagel and Xu (?) also show that the pricing implications of the fading memory model are similar to one in which investors have full memory and they believe that μ follows a random walk, but only if the true μ is in fact constant. The wedge between investors' time-varying posterior mean of μ and the true constant μ are needed for the model to produce excess return predictability and return expectations forecast error predictability that is in line with the data. Using aggregate earnings growth as a proxy for payout growth, Nagel and Xu (?) also show that analysts' long-run earnings growth forecast errors are predictable with an exponentially-weighted average of past observed payout growth observations, which is in line model predictions and consistent with the evidence discussed in Section 3.2 that long-term payout growth expectations appear to be an important source of variation in the aggregate stock price level.

Models in which investors are uncertain about long-run growth can have a built-in fragility. It is important to keep this in mind, because this fragility may not be apparent in log-linearized or numerically solved versions of these models. To illustrate this issue, consider first the simplest case with risk-neutrality as in equation (19), where we had

$$P_t = E_t^m \left[\lim_{T \rightarrow \infty} \sum_{j=1}^T \delta^j D_{t+j} \right]. \quad (28)$$

Suppose log dividends are generated as $d_t = d_{t-1} + \mu + \eta_{t+1}$ where $\eta_t > 0$ is i.i.d. with $E[\eta_t] = 0$ and $\mu > 0$ is an unknown parameter. If investors' prior beliefs about μ assign arbitrarily small but positive probability mass to growth rates above δ^{-1} , then the asset price (28) diverges to infinity. When prior beliefs about μ are bounded strictly below δ^{-1} , arbitrarily small prior mass sufficiently close to δ^{-1} gives rise to arbitrarily large (albeit finite) equilibrium price levels. This shows how a small amount of uncertainty can make a large difference for the asset price predictions.

Geweke (?) shows that the blow-up problem becomes even starker with CRRA preferences. When the endowment is log-normal as in (27) with uncertainty about μ , then expected utility ceases to exist under conjugate Bayesian prior beliefs, unless the intertemporal substitution elasticity is exactly equal to one or unless one restricts the support of prior beliefs.¹² Weitzmann

¹²Further results about the existence of present value relationships under Bayesian learning are derived in Pesaran, Pettenuzzo and Timmermann (?) and Adam and Marcet (?).

(?) shows the restrictions in prior beliefs can then dominate the model’s asset pricing implications. The divergence problem also arises with Epstein-Zin preferences, whenever the inter-temporal elasticity of substitution differs from a value of exactly one (Collin-Dufresne et al. (?); Nagel and Xu (?)).

As a result, Collin-Dufresne, Johannes and Lochstoer (?) truncate the state space to insure finite outcomes. Pastor and Veronesi (?, ?) assume that uncertainty about growth rates disappears after some time period T , which can be stochastic. Nagel and Xu (?) specify informative prior beliefs that pull long-run growth expectations toward a prior mean. In their perpetual learning model, the persistent pull toward this prior mean applies to beliefs of all future agents, which is sufficient to ensure finite valuations in with Epstein-Zin utility even if the intertemporal elasticity of substitution is different from one.

There are a number of alternative views about the economic relevance of the blow-up problem. One view is that it would be simply unreasonable for investors to entertain the possibility of long-run growth rates that imply very high valuations or valuations that are extremely sensitive to growth rate uncertainty. An alternative view is that it will be very difficult to empirically discipline present value models using survey data, as the asset pricing implications can be driven by small amounts of uncertainty that are empirically hard to determine. Yet another view is that it is ultimately just an empirical question whether or not measured payout expectations from survey data are able to explain asset price behavior.

5.2 Learning about prices

A number of asset pricing models introduce learning directly about price behavior, as discussed in section 2.3.2. Price learning can generate stronger endogenous propagation of fundamental disturbances than payout learning, due to the ‘self-referential’ nature of price learning: price beliefs affect price outcomes and price outcome future revisions in price beliefs. This feedback loop, which is absent when learning is about exogenous fundamentals, say payouts, can generate plausible asset price volatility, even with time-separable preference specifications. Models of price learning can also explain the strong positive comovement between recent past returns and investors’ return expectations observed in survey data, which is an empirical fact that models featuring payout learning struggle to explain. Models specifying learning about prices rely on the one-step ahead asset pricing equation (1), rather than the discounted sum formula (28), therefore also do not face the kind of large sensitivity of the asset price predications to small changes in subjective price beliefs

Early models of learning about prices, e.g., Timmermann (?), studied learning about the *level* of next period’s asset price. Such learning specifications generated only modest amounts of additional asset price volatility, despite the feedback between price beliefs and price outcomes. The reason is that the strength of the feedback tends to be weak under price *level* learning. This is so because the expected *future* price level affects the *current* price level approximately one-to-one, see equation (3). In the vicinity of rational price beliefs, outcomes and beliefs will thus move virtually in the same way. The weak divergence between outcomes and beliefs causes learning-induced belief revisions to be weak, so that learning adds little asset price volatility.

As a result, subsequent models studied learning about capital gains, i.e., learning about the *change* in the price level from one period to the next. As should be clear from equation (15), realized asset prices and thus realized capital gains are rather sensitive to revisions in capital gain expectations $E_t^m[\beta_{t+1}^P]$, as $\beta_t^m = 1/R_{f,t}$ tends to be close to one. Adam, Marcet and Nicolini (?) show how many forms of learning about capital gains impart momentum and long-horizon mean-reversion into asset price dynamics, thereby generating large and persistent swings in the price-dividend ratio. This allows simple asset pricing models with time-separable preferences to generate realistic amounts of stock price volatility, even if they struggle to fully replicate the equity premium.

Learning-induced variations in subjective capital gain expectations also generate positive comovement of the price-dividend ratio and subjective expected excess returns. As we discussed in Section 2.3.2, such time-variation in subjective expected excess returns requires that the SDF adjusts in response to a change capital gains expectations. In Adam, Marcet, and Beutel (?), in a setting with CRRA utility and with wage income as a source of wealth outside of the stock market, an agent optimistic about future capital gains anticipates that a greater share of future wealth is exposed to stock market risk. As a consequence, perceived risk rises. Specifically, under the agent’s subjective beliefs, the SDF is more volatile when with capital gains expectations are higher. In equilibrium, optimism about capital gains therefore produces not only a high price-dividend ratio, but also a high subjective risk premium. Price growth expectations in their model are a function of an exponentially-weighted average of past price growth, similar to models of fading memory or perceived parameter drift.

Jin and Sui (?) pursue an alternative approach that retains the assumption of common knowledge. In their model, investors forecast future prices, but they implicitly form payout growth expectations that justify their expectation of the future price such that a discounted sum valuation like (28) holds under subjective beliefs. This brings the model quite close to mod-

els of payout learning. It also allows the model to match the evidence on time-varying subjective cash flow expectations that we discussed in Section 3.2.

Overall, models of payout learning and price learning share a common thrust in that they explain cycles in asset price valuations with waves of optimism and pessimism about future levels of stock prices. They differ in what data agents use to form these expectations and how agents reason about the justification for these beliefs about future prices. In the end, elements from both classes of models may be needed.

5.3 Learning biases

The models we discussed so far employ forms of belief formation that can broadly be motivated by Bayesian learning, albeit with some additional tweaks, such as fading memory or a prior belief that parameters are drifting and hence data far in the past has become irrelevant. A different strand of the literature starts instead with an assumption, motivated by psychological experimental evidence, that agents use certain heuristics in belief formation. Reliance on these heuristics generates biases in updating of beliefs in response to incoming information.

One heuristic that seems to be particularly suitable for explaining cyclical asset price behavior is the representativeness heuristic of Kahnemann and Tversky (?). An early example is the model presented by Barberis, Shleifer and Vishny (?) where a string of positive or negative earnings change leads an investor to view this repeat performance as representative and adopt a forecasting model that extrapolates the past performance too far into the future. In this model, this extrapolation has no grounding in reality as the true process does not have any persistence.

In the diagnostic expectations approach of Bordalo, Gennaioli and Shleifer (?), agents apply the representativeness heuristic in a way that is closer to Bayesian learning. In their setup, agents observe a signal that has predictive information for the variable that they want to forecast. In line with the representativeness heuristic, however, they overweight outcomes that have become more likely in light of recently incoming information. As a consequence, agents exaggerate true predictability patterns in the data. Bordalo, Gennaioli, La Porta and Shleifer (?) apply this approach to stock valuation and show that it can explain a number of interesting financial market regularities, for instance, why the returns on stocks with the most optimistic analyst long-term earnings growth forecasts are lower than those on stocks with the most pessimistic forecasts.

One tension in the diagnostic expectations approach is that agents observe

(and overreact to) objective news, that is, the realization of variables relative to an undistorted Bayesian forecast from the previous period. And the news reaction is added to the previous period's Bayesian forecast. Thus, while their actual forecast in the previous period was distorted by the representativeness heuristic, agents are able, in the current period, to calculate the innovation relative to the previous period's Bayesian forecast and adjust this Bayesian forecast in the direction of the innovation. This raises the question where the knowledge of the Bayesian forecast comes from.

Rather than assuming that agents exaggerate true predictability patterns, or see predictability where there is none, an alternative approach to biased learning is to assume that agents use simplified models that do not capture the full complexity of the true predictability that exists in the data. Fuster, Hebert, and Laibson (?) call this approach 'natural expectations'. For example, the true data-generating process may be a high-order ARMA process, while agents estimate a forecasting model that allows only for a small number of lags. In this case, agents may fail to fully perceive the degree of mean reversion inherent in fundamental dynamics, which can also lead to overreaction that generates asset price cycles with procyclical optimism in investors' subjective beliefs.

The difference between these biased-learning approaches and Bayesian learning is smaller than it may seem. Updating that deviates from Bayesian learning with priors grounded in objective reality can often be rationalized as Bayesian learning under a particular subjective prior belief. To illustrate, consider the natural expectations approach and suppose the true data-generating process is an AR(5), but the agent estimates a simplified AR(1) model to construct forecasts. The agent's approach could be rationalized by giving the agent a dogmatic subjective prior that the data-generating process is an AR(1). Given this prior, the agent's updating of beliefs is in accordance with Bayes' law. In this sense, there often exists an equivalent Bayesian belief formulation, so that it becomes difficult to identify whether updating bias arises due to a deviation from Bayes' law or simply due to an application of Bayes' law under a particular subjective prior. The distinction between both views is more of a philosophical nature. What matters for model predictions is the sequence of subjective beliefs that updating gives rise to.

5.4 Heterogeneity

Motivated by the fact that survey measures of expectations typically display a large degree of heterogeneity, a number of papers study asset pricing setups in which agents hold different beliefs. Giglio, Maggiori, Stroebel and Utkus (?) show that belief heterogeneity has predictive power for the composi-

tion of individual portfolios and for trading behavior. Belief heterogeneity also appears relevant for understanding the large volume of assets traded on financial markets, as arguably these volumes are difficult to square with risk-sharing motives alone. Heterogeneity also allows studying the redistributive effects associated with subjective belief dynamics.

In Barberis, Greenwood, Jin and Shleifer (?) some investors form price growth expectations by extrapolating from past price growth, while other investors have rational expectations. The extrapolators are assumed to hold implicit subjective payout expectations that support their beliefs about future prices under a discounted sum valuation as in equation (28). Following a string of price increases, extrapolators are therefore effectively more optimistic about future payouts than rational investors and therefore hold a greater share of the outstanding supply of stocks. Due to their resulting greater exposure to stock market risk, they demand a higher subjective risk premium, which allows high past price growth to coincide with optimistic return expectations of the extrapolators in equilibrium. In contrast, rational investors have countercyclical return expectations. One potential interpretation is that the extrapolators represent individual investors while the rational investors represent more sophisticated professional investors. For tractability in the presence of heterogeneity, investors in the model have CARA preferences. But this has the consequence that the model does not produce realistic asset pricing predictions on a number of dimensions, such as the equity premium, the volatility of the price-dividend ratio, and the long-run behavior in the presence of economic growth.

Collin-Dufresne, Johannes and Lochstoer (?) explore heterogeneity of a different kind. Motivated by the empirical evidence in Malmendier and Nagel (?) and Malmendier and Nagel (?) that investors learn from life-time experiences, they assume that overlapping cohorts of investors learn about the dividend process from the dividend history observed during their life times. The model does well on a number of standard asset pricing moments and in matching the learning-from-experience evidence in microdata and surveys, but the dynamics of asset prices and beliefs are somewhat difficult to evaluate quantitatively as there are only two overlapping cohorts and subjective risk premia jump every 20 years when there is a generational shift.

Adam, Marcet, Merkel and Beutel (?) consider a model in which agents are heterogeneous in their tendency to extrapolate past capital gains. The heterogeneity is motivated by the empirical observation that the capital gain expectations of investors with more years of stock market experience react significantly weaker to past capital gains than the expectations of less experienced investors. Working with the one-step ahead asset pricing equation (1) for all investors allows the use of CRRA preferences without losing tractabil-

ity. The model produces quantitatively realistic asset price dynamics, can replicate patterns of trading volume and generates significant amounts of wealth redistribution over stock price boom and bust cycles.

Overall, the literature that incorporates belief heterogeneity in asset pricing in a way that is disciplined by survey evidence is still in its infancy. In many settings, lack of tractability makes it difficult to entertain empirically realistic belief heterogeneity and, at the same time, produce quantitatively plausible asset price behavior.

6 Future Research Directions

During the past decade, researchers have made substantial progress in linking investor expectations data with asset pricing theory. We conclude by highlighting several areas in which further advances would be desirable:

- We need more evidence on the links between expectations and investor portfolio decisions. Does low sensitivity of actions to expectations at the individual level translate into low sensitivity at more aggregated levels (e.g., cohorts, investor category, market-wide)? Which are the investors influenced by professional forecasters' expectations, and how strong is the influence?
- Subjective risk perceptions may be as important as the perceived first moments of returns and payoffs. So far, however, the available empirical evidence on the dynamics of subjective risk perceptions is rather limited. We need more work that explores how investors form beliefs about asset risks and how these risk perceptions are linked to the subjective risk premia that they demand to hold risky assets.
- There is substantial heterogeneity in the subjective beliefs of different groups of market participants, e.g., between professional forecasters and individual investors. Is it important for asset pricing to account for this heterogeneity? How should the belief formation be modeled for these groups?
- More generally, it would be desirable to make progress on the question of how to best aggregate the heterogeneous expectations of different investors, with the objective of identifying the marginal agents' beliefs. While this will prove difficult because survey responses are contaminated by response errors and tend to be measured infrequently, it is an important question that needs to be tackled.

- Subjective belief dynamics may not only be a major source of asset price volatility, but they may also play an important role as a source of macroeconomic fluctuations. Expectations about asset payoffs should be linked to the expectations that shape the decisions of consumers and firms about investment, production, and consumption. Therefore, there is much to be learned from integrating asset pricing based on subjective belief dynamics into macroeconomic models. Relatedly, expectations of financial intermediaries may affect credit supply, macroeconomic outcomes, and asset prices. Studying intermediaries' expectations, as done recently in Ma et al. (?), may provide new insights for macro-finance models with intermediary sectors.
- Analysis of policy in models with subjective beliefs brings up special challenges. On the theory side, the challenge is modeling how subjective beliefs react to changes in the policy environment, e.g., as in Adam and Woodford (?). This may in turn provide new insights about asset price reactions to policy announcements. On the empirical side, we need more evidence on how expectations change in response to policy interventions. Such evidence can be obtained through information experiments, e.g., as in Coibion et al. (?) that induce changes in participants' perceptions about actual policy. Chapter 5, "Survey Experiments on Economic Expectations" provides an overview of this line of research. One open question is whether participants' responses in survey experiments could differ from real-world responses because the experiment forces individuals to pay attention to stimuli that they would have ignored in a real-world setting.

Appendix

Table 1: Investor Survey Data Sets

Survey	Population	Repository
<i>Panel A: Stock market return or capital gain expectations</i>		
UBS/Gallup	Individuals	Roper Center ¹
Yale/ICF	Wealthy individuals	Yale ICF ²
Yale/ICF	Institutional investors	Yale ICF ²
Michigan Survey of Consumers	Individuals	UM Survey Research Center ³
Graham-Harvey CFO	Financial managers	FRB of Richmond ⁴
Livingston	Professional forecasters	FRB of Philadelphia ⁵
<i>Panel B: Stock market cash flow expectations</i>		
IBES	Equity Analysts	WRDS ⁶
Survey of Professional Forecasters	Professional forecasters	FRB of Philadelphia ⁷
<i>Panel C: Interest rate expectations</i>		
Survey of Professional Forecasters	Professional forecasters	FRB of Philadelphia ⁷
Bluechip Financial Forecasts	Professional forecasters	Wolters Kluwer ⁸

¹<https://ropercenter.cornell.edu>

²<https://som.yale.edu/centers/international-center-for-finance/data>

³<https://data.sca.isr.umich.edu>

⁴<https://www.richmondfed.org/cfosurvey>

⁵<https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/livingston-survey>

⁶<https://wrds-www.wharton.upenn.edu>

⁷<https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/survey-of-professional-forecasters>

⁸<https://www.wolterskluwer.com/en/solutions/vitallaw-law-firms/blue-chip>