A Model-Free Assessment of the Importance of Subjective Beliefs for Asset Pricing*

Paymon Khorrami
Duke University

March 19, 2025

Abstract

Are belief dynamics or risks and risk attitudes more important for asset pricing? Allowing both, I use survey data combined with subjective-belief versions of stochastic discount factor (SDF) volatility bounds to shed new light on this classic question. I estimate lower bounds for the volatility of the SDF attributable to (i) risks relevant for investor marginal utility, versus (ii) subjective belief dynamics. The estimates suggest that risks, particularly long-term risks, make up at least half of SDF volatility. An example extrapolation model with a modest direct contribution of beliefs to SDF volatility (about 25%) can account for my estimates. This example also highlights the potential for a novel mechanism, *subjective risk*, which is the indirect impact beliefs have on asset prices through induced marginal utility volatility.

*E-mail: paymon.khorrami@duke.edu. Thanks to Spencer Kwon and participants at the Yale Junior Finance Workshop for comments.

1 Introduction

Are beliefs or risks more important for asset pricing? Despite a long line of research, answers remain elusive because the question is posed at a seemingly impenetrable level of generality. A common approach is to write down a fully-specified model (behavioral or rational) and test it against a set of important asset market data. One problem: the results from such models can often depend on assumptions that are hard to verify or measure. For example, while a subjective belief model may be able to account for many properties of asset prices, its success hinges not only on its belief dynamics but also on auxiliary assumptions on preferences, frictions, and heterogeneity. Similarly, many prominent rational asset-pricing models have been characterized as sensitive to "dark matter" that is difficult to measure with statistical accuracy (Chen, Dou and Kogan, 2024a).

This paper, instead, takes a "model-free" approach similar to Hansen and Jagannathan (1991) and Alvarez and Jermann (2005) but in a context where subjective beliefs may differ from rational expectations. Rather than specifying a particular model, I characterize some properties any asset-pricing model must have and in the process provide general estimates for the contributions of beliefs versus risks.

More specifically, I estimate the fraction of stochastic discount factor (SDF) volatility that originates from risks versus beliefs. Since SDFs summarize *all* asset prices, such model-free bounds and estimates should be useful to quantitative researchers seeking to align their models with data. Furthermore, the analysis provides a clear way to map financial survey data to key aspects of SDFs and other objects.[1] Finally, approaching the problem at some level of generality can help unify disparate results from various specific markets. That is, rather than ask how much beliefs matter for, say, stock market volatility, this paper asks how much beliefs matter for asset pricing in general.

The problem. Encoded in asset prices is the key fact that the SDF must be highly volatile (Hansen and Jagannathan, 1991). But what is not known is whether this volatility stems primarily from beliefs or risk. Ultimately, this identification problem is related to the fact that asset prices only encode the product of beliefs and risks and do not identify them separately (Harrison and Kreps, 1979). An Arrow-Debreu security for state x will have

[1] The analysis of financial surveys and their implications has been a fast-growing area of research (Chun, 2011; Amromin and Sharpe, 2014; Greenwood and Shleifer, 2014; Piazzesi et al., 2015; Crump et al., 2018; Xu, 2019; Wang, 2021; d'Arienzo, 2020; De La O and Myers, 2021; Adam et al., 2021; Nagel and Xu, 2023). I build, in particular, on the surveys used by Nagel and Xu (2023).

price

$$\text{Price}(x) = \text{Probability}(x) \times \underbrace{\text{Beliefs}(x) \times \text{MU}(x)}_{\text{SDF}(x)},$$

where $\text{Probability}(x)$ is the objective state probability, $\text{Beliefs}(x)$ denotes investors' belief distortion (ratio of subjective-to-objective probability), and $\text{MU}(x)$ is the representative investor's marginal utility in state x. It is this latter piece, $\text{MU}(x)$, that captures risk and risk attitudes. With enough data, perhaps the econometrician can figure out the state frequency $\text{Probability}(x)$. But beyond that, asset prices only identify the product $\text{Beliefs}(x) \times \text{MU}(x)$. To attribute SDF volatility to either component, asset prices do not suffice, and that is where financial surveys enter.

The literature now understands further that most of the (objective) SDF volatility comes from permanent shocks (Alvarez and Jermann, 2005), i.e., from the first term in a decomposition like

$$\text{SDF} = \text{Permanent} \times \text{Transitory}. \tag{1}$$

Does this additional knowledge help us separate $\text{Beliefs}(x)$ from $\text{MU}(x)$? Unfortunately not. The entirety of the belief distortion $\text{Beliefs}(x)$ enters into the permanent component, and so it cannot be distinguished from "long-term risks" to $\text{MU}(x)$ (Borovička, Hansen and Scheinkman, 2016). I show how, armed with financial survey data, one can separate long-term risks from beliefs.

Solution and results. It is relatively obvious that surveys provide *some* information about beliefs, but precisely what information and how should it be used? The core idea here is that survey data on expected returns are connected to risk by no-arbitrage. If investors are optimistic about an asset, they must at least perceive it to be risky in some way. This ties the survey responses to investor marginal utility (MU), which captures risk sensitivities and is *distinct from beliefs*. On the other hand, realized returns contain information about the volatility of the SDF (Hansen and Jagannathan, 1991), which is affected by both risks and beliefs. By comparing survey expected returns to objective expected returns, one can thus learn how much of SDF volatility comes from MU volatility, i.e., how much stems from risk. I formalize this idea and explore variations in the paper.

My main empirical evidence, building on the logic above, suggests that at least half of total SDF volatility stems from risk. The key reason is the simple observation that subjective equity premia are large, almost as large as their true unconditional counterparts.

What type of risk matters? I provide some direction by estimating a subjective version of the Alvarez and Jermann (2005) bounds. In particular, a decomposition like (1) exists for investor's MU as well as the SDF. I estimate that the permanent component of MU is responsible for at least 90% of total MU volatility. This cannot be attributed to beliefs, at least not directly, and it says that investors perceive substantial long-term risks.

Finally, I analyze several example models to illustrate and interpret the results. I consider a long-run risks model (Bansal and Yaron, 2004) and a growth extrapolation model that features "perceived long-run risks" (Collin-Dufresne et al., 2017; Nagel and Xu, 2022).[2] The latter is an interesting case study because both beliefs and risks are operational: investors fear a hypothetical long-run risk, even though they are wrong. Should SDF volatility stemming from this interaction between beliefs and preferences be classified as "belief-driven" or "risk-driven"? While not obvious a priori, my framework provides a clear answer. The calibrated growth extrapolation model is surprisingly consistent with my estimated volatility bounds, revealing as a by-product that beliefs can account for about 25% of SDF volatility. I also investigate what happens when the role for beliefs is magnified, by introducing sentiment shocks. In that case, the transitory component of the SDF is given a much larger role, standing against the existing empirical evidence. Thus, within this class of models, the direct role of beliefs cannot be too large. Instead, beliefs must matter largely indirectly by creating marginal utility volatility, i.e., perceived risks.

Related literature and other approaches. My approach relies on volatility bounds, similar to Alvarez and Jermann (2005) and Bakshi and Chabi-Yo (2012), that provide *some* information about SDFs, risks, and beliefs. If one wants more information, one needs to make more assumptions. For instance, assuming rational expectations and putting structure on objective dynamics allows the econometrician to extract (some aspects of) investor marginal utility (Bansal and Viswanathan, 1993; Hansen and Jagannathan, 1997; Jackwerth, 2000; Aıt-Sahalia and Lo, 2000; Rosenberg and Engle, 2002; Ghosh et al., 2017;

[2]Other risk-based models consistent with a volatile SDF that features a large permanent component include Beaudry and Portier (2004); Bloom (2009); Christiano et al. (2014); Fajgelbaum et al. (2017); Bianchi et al. (2018); Di Tella (2017); Bianchi et al. (2023).

At the same time, behavioral finance models with distorted beliefs can also produce similar implications. Some recent examples include those with "fundamentals extrapolation" (Barberis et al., 1998; Cecchetti et al., 2000; Lansing, 2006; Collin-Dufresne et al., 2017); "price extrapolation" (Branch and Evans, 2011; Barberis et al., 2015; Adam et al., 2016; Barberis et al., 2018; Jin and Sui, 2022); "diagnostic beliefs" (Bordalo et al., 2021, 2022); "learning from experiences" (Schraeder, 2016; Ehling et al., 2018; Malmendier et al., 2020; Nagel and Xu, 2022); "probability weighting" (Barberis and Huang, 2008); and "natural expectations" (Fuster et al., 2012). While I do not examine SDF volatility in all of these models, there are connections between the examples I study in Section 4 and the extrapolation models.

Beason and Schreindorfer, 2022; Chen et al., 2024b). An alternative is to put some structure on marginal utility (e.g., assumptions about preferences and risks) and treat beliefs as a "residual" needed to match asset market data (Chen, Hansen and Hansen, 2024c), but then cross-check the resulting beliefs against survey data (Ghosh and Roussellet, 2023).

While these approaches are useful, my goal is to see how much can be said without these additional assumptions. I do not take a stand on investor preferences. I make no assumptions about the underlying objective dynamics and the types of risks. Is there long-run growth risk; is uncertainty time-varying; does the economic state "jump" or evolve continuously? I also avoid some practical challenges like the high dimensionality and nonlinear mapping from states to the SDF dynamics. As long as we trust the survey data to proxy for investor beliefs, my procedure imposes minimal additional assumptions. The cost is that we will only be able to learn certain bounds on the volatilities of SDFs and marginal utilities, rather than their actual values.

A separate stream of behavioral finance literature is concerned with how belief dynamics may help resolve "excess volatility" puzzles, by combining present-value relations with survey data (Chen et al., 2013; De La O and Myers, 2021, 2024; De La O et al., 2023; Bordalo et al., 2019, 2024a,b). This paper differs by focusing more on excess returns rather than excess volatility, and examining the SDF rather than a particular asset's present-value relation.[3] However, there is a connection between my work and this literature. For example, if the market excess return is an important factor for the SDF, then its volatility is a key contributor to SDF volatility. In that case, subjective beliefs' role in SDF volatility is tied to their importance for stock market volatility and vice versa.

2 Theory: Subjective belief bounds

2.1 General environment

Let $\mathbb{P}$ denote the objective probability measure and $\mathbb{E}$ the associated expectation operator. I allow the marginal investor to possess *subjective beliefs* and use a distorted probability and expectation, which we will denote by $\tilde{\mathbb{P}}$ and $\tilde{\mathbb{E}}$. To maintain a minimal

[3] While a decomposition into "future cash flows" and "future returns" is interesting in its own right, it is the wrong framework to evaluate whether return premia are due to behavioral biases or risk. For this question, it is better to work directly with returns and return expectations. For instance, it is not clear whether or not subjective expectations about future cash flows, even if they are biased, capture risk premia or not. It is not even necessary for valuation ratios or expected returns to be time-varying; even if those objects are constant, one wonders what fraction of constant discounts/premia arise from beliefs versus risk.

amount of consistency for beliefs, I assume that $\mathbb{P}$ and $\tilde{\mathbb{P}}$ agree on null sets, so that their discrepancy can be modeled via the positive martingale B_t (likelihood ratio):[4]

$$B_t = \left(\frac{d\tilde{\mathbb{P}}}{d\mathbb{P}}\right)_t \tag{2}$$

In other words, the subjective probability can be defined by $\tilde{\mathbb{P}}(\mathcal{E}) := \mathbb{E}[\frac{B_t}{B_0}\mathbf{1}_{\mathcal{E}}]$ for any event $\mathcal{E}$ that resides in the time-t information set.

I will assume absence of arbitrage opportunities. (Because $\mathbb{P}$ and $\tilde{\mathbb{P}}$ are equivalent measures, there is no ambiguity to simply saying "arbitrage" without reference to any probability.) In that case, let $(S_t)_{t\geq 0}$ denote the *stochastic discount factor* (SDF) process that prices all assets, i.e., S is positive and for any gross return R_{t+1} we have

$$1 = \mathbb{E}_t\left[\frac{S_{t+1}}{S_t}R_{t+1}\right], \tag{3}$$

where $\mathbb{E}_t$ denotes the conditional expectation operator. At the same time, there is a subjective SDF $\tilde{S}_t$ that prices assets under investor beliefs, i.e., investors' Euler equation:

$$1 = \tilde{\mathbb{E}}_t\left[\frac{\tilde{S}_{t+1}}{\tilde{S}_t}R_{t+1}\right]. \tag{4}$$

The interpretation of equation (4) is that $\tilde{S}_t$ represents the *investors' marginal utility of wealth*—we will often refer to $\tilde{S}$ as marginal utility, in contrast to the objective SDF S. Without loss of any generality for my points, impose the normalization $S_0 = \tilde{S}_0 = B_0 = 1$. Then, comparing (3)-(4) and using (2), we have

$$S_t = \tilde{S}_t B_t. \tag{5}$$

In this sense, the SDF is the product of investor marginal utility and investor beliefs. Roughly speaking, the idea of this paper is to use survey data to learn about the volatilities of S and $\tilde{S}$ separately, thereby providing information about the belief distortion B.

In addition to the volatilities of S and $\tilde{S}$, information is contained in the volatilities of their permanent components. Under fairly general conditions, which I implicitly

[4]This assumption on subjective beliefs is quite general and captures almost all belief-based models in the behavioral finance literature. However, some models are ruled out, such as those where the law of iterated expectations (LIE) is violated (e.g., Fuster et al., 2010). Other models, like the fading memory model of Nagel and Xu (2022), may violate LIE when interpreted literally but can be written in an observationally equivalent form where LIE holds. Therefore, that model would be included in the general setup here.

assume, the SDF obeys a unique factorization into *permanent and transitory components* (Alvarez and Jermann, 2005; Hansen and Scheinkman, 2009; Qin and Linetsky, 2016, 2017):

$$S_t = G_t H_t, \tag{6}$$

where G is a trend-stationary process (transitory component) and H is a $\mathbb{P}$-martingale (permanent component). An analogous factorization holds for the investor marginal utility:

$$\tilde{S}_t = \tilde{G}_t \tilde{H}_t, \tag{7}$$

where $\tilde{G}$ is trend-stationary and $\tilde{H}$ is a $\tilde{\mathbb{P}}$-martingale.

Luckily, equation (5) then implies that

$$H_t = \tilde{H}_t B_t \tag{8}$$

and hence

$$G_t = \tilde{G}_t. \tag{9}$$

To see this, notice that $\mathbb{E}_0[H_t] = H_0 \mathbb{E}_0[\frac{B_t}{B_0}\frac{\tilde{H}_t}{\tilde{H}_0}] = H_0 \tilde{\mathbb{E}}_0[\frac{\tilde{H}_t}{\tilde{H}_0}] = H_0$, so that H and $\tilde{H}$ are indeed $\mathbb{P}$- and $\tilde{\mathbb{P}}$-martingales, respectively, when they satisfy (8). The interpretations of $\tilde{H}$ and H are different: $\tilde{H}$ captures the impact of permanent risks on investor marginal utility, while $H = \tilde{H}B$ captures the joint impact of permanent risks and belief distortions. The permanent components H and $\tilde{H}$ will be given special attention, in light of evidence from Alvarez and Jermann (2005) and Bakshi and Chabi-Yo (2012) that G carries a relatively minor contribution to overall SDF volatility. What is not yet agreed is whether the prominence of H is due to $\tilde{H}$ or B.

This entire theoretical setup has referenced a single investor belief and marginal utility, when in fact there is clearly heterogeneity in the real world and in the surveys that I will use empirically. Appendix B shows that the presence of heterogeneity can be accommodated in this setup. In particular, the common practice of averaging across survey respondents (the so-called consensus forecast) allows us to interpret B, $\tilde{S}$, and $\tilde{H}$ as the average belief, the belief-weighted average marginal utility, and its permanent component. This type of averaging is also consistent with theoretical models of belief

disagreement.[5]

Finally, I will make frequent reference to a few special returns in this theory section. First, I denote the *one-period risk-free rate* R_t^f. Since this is conditionally risk-free, we have from (3)-(4) that $R_t^f = (\mathbb{E}_t \frac{S_{t+1}}{S_t})^{-1} = (\tilde{\mathbb{E}}_t \frac{\tilde{S}_{t+1}}{\tilde{S}_t})^{-1}$. I will assume, as in most of the asset-pricing literature, that the risk-free rate is traded.

Second, I will make use of the *long bond return* R_{t+1}^∞, defined as the one-period holding return on an arbitrarily long-maturity discount bond:

$$R_{t+1}^\infty := \lim_{T\to\infty} \frac{\mathbb{E}_{t+1}[\frac{S_T}{S_{t+1}}]}{\mathbb{E}_t[\frac{S_T}{S_t}]}.$$

The long bond was originally studied in Kazemi (1992) (see also Martin and Ross, 2013, and Ross, 2015). In particular, it turns out that R_{t+1}^∞ identifies the transitory component of the SDF:

$$R_{t+1}^\infty = \lim_{T\to\infty} \frac{\mathbb{E}_{t+1}[\frac{S_T}{S_{t+1}}]}{\mathbb{E}_t[\frac{S_T}{S_t}]} = \lim_{T\to\infty} \frac{G_t}{G_{t+1}} \frac{\mathbb{E}_{t+1}[\frac{H_T}{H_{t+1}} G_T]}{\mathbb{E}_t[\frac{H_T}{H_t} G_T]} = \frac{G_t}{G_{t+1}}, \tag{10}$$

because the ratio of conditional expectations converges to 1 under mild conditions. Of course, an infinite-maturity discount bond does not exist, but we can proxy this with the longest-maturity bonds available.

Third, I will occasionally consider the *growth-optimal return* R_{t+1}^*. This is the portfolio that maximizes the growth rate of wealth, i.e., it solves $\max_R \mathbb{E}_t[\log(R)]$ subject to the pricing constraint (3). The solution is $R_{t+1}^* = \frac{S_t}{S_{t+1}}$ (Roll, 1973; Bansal and Lehmann, 1997). Thus, there is a well-known equivalence between extracting the SDF from the data and finding the growth-optimal portfolio. Because of the difficulties in identifying R^*, I will not rely on any one particular approach but rather present results from several.

2.2 Volatility bounds

I consider the following *entropy measures of volatility*. For any positive random variable X, define

$$L_t(X) := \log\big(\mathbb{E}_t[X]\big) - \mathbb{E}_t\big[\log(X)\big] \tag{11}$$

$$\tilde{L}_t(X) := \log\big(\tilde{\mathbb{E}}_t[X]\big) - \tilde{\mathbb{E}}_t\big[\log(X)\big]. \tag{12}$$

[5] For some examples, see Detemple and Murthy (1994), Zapatero (1998), Scheinkman and Xiong (2003), David (2008), Bhamra and Uppal (2014), Kasa et al. (2014), Atmaz and Basak (2018).

Entropy measures played an important role in Alvarez and Jermann (2005), but have also been used in various other asset-pricing contexts (Backus et al., 2011, 2014; Martin, 2017) By Jensen's inequality, $L_t(X) \geq 0$ and $\tilde{L}_t(X) \geq 0$. In the special case that X_{t+1} is conditionally lognormal, $L_t(X_{t+1}) = \frac{1}{2}\text{Var}_t[\log X_{t+1}]$, but in other cases entropy conveys more information about the left tail of a distribution, relative to variance.

Whereas $L_t(X)$ represents objective volatility, $\tilde{L}_t(X)$ represents subjective, or perceived, volatility. To make interpretable comparisons, I will sometimes invoke the following condition equating objective and subjective entropies:

Definition 1. X satisfies the *conditional entropy identity* (CEI) if $L_t(X_{t+1}) = \tilde{L}_t(X_{t+1})$.

When does a variable X_{t+1} to satisfy CEI? One important case when CEI holds is when the underlying joint dynamics are conditionally lognormal over sufficiently short time periods (i.e., no jumps).[6] Most models in macro-finance have this property. A second broader case is when investors hold accurate beliefs about higher moments beyond the mean (e.g., agents understand variances, skewness, kurtosis, etc., well). This automatically holds in the conditionally lognormal case just mentioned, but it may happen to be true even outside of conditional lognormality. In this paper, I will impose CEI only sometimes and only for $\Delta\tilde{S}_{t+1} := \frac{\tilde{S}_{t+1}}{\tilde{S}_t}$ and $\Delta\tilde{H}_{t+1} := \frac{\tilde{H}_{t+1}}{\tilde{H}_t}$, remaining agnostic about CEI on other variables, so in some sense the conditions I will assume are even weaker.

For some results, I will exploit an even weaker condition that allows us to bound time-series averages of conditional entropies.

Definition 2. X satisfies the *unconditional entropy bound* (UEB) if $L(X_{t+1}) \geq \mathbb{E}[\tilde{L}_t(X_{t+1})]$.

Clearly, if X satisfies CEI, then it satisfies UEB. This is because $L(X) = \mathbb{E}[L_t(X)] + L(\mathbb{E}_t[X]) \geq \mathbb{E}[L_t(X)]$, so if $L_t(X) = \tilde{L}_t(X)$, then we get UEB. But some variables can satisfy UEB even if CEI fails. One obvious case is if $L_t(X) \geq \tilde{L}_t(X)$, meaning that the subjective measure understates risk to X. This case is relevant because, empirically, survey measures of risk tend to be lower than actual risk. Another case is if $L(\mathbb{E}_t[X])$ is large, meaning that X has significant predictability. This case is relevant because actual returns do tend to have non-trivial predictability.

[6] One easy way to see this is that $L_t(X_{t+1}) = \frac{1}{2}\text{Var}_t[\log X_{t+1}]$ and $\tilde{L}_t(X_{t+1}) = \frac{1}{2}\tilde{\text{Var}}_t[\log X_{t+1}]$ under conditional lognormality. Then, combine this result with the fact that belief distortions about variance necessarily become negligible as the time interval shrinks. See Appendix A.6 for a formal proof.

Lemma 1. *The following conditional volatility bounds hold for all returns R_{t+1}:*

$$L_t\Big(\frac{S_{t+1}}{S_t}\Big) \geq \mathbb{E}_t\big[\log R_{t+1}\big] - \log R_t^f \tag{13}$$

$$\tilde{L}_t\Big(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}\Big) \geq \tilde{\mathbb{E}}_t\big[\log R_{t+1}\big] - \log R_t^f \tag{14}$$

and

$$L_t\Big(\frac{H_{t+1}}{H_t}\Big) \geq \mathbb{E}_t\big[\log R_{t+1} - \log R_{t+1}^{\infty}\big] \tag{15}$$

$$\tilde{L}_t\Big(\frac{\tilde{H}_{t+1}}{\tilde{H}_t}\Big) \geq \tilde{\mathbb{E}}_t\big[\log R_{t+1} - \log R_{t+1}^{\infty}\big] \tag{16}$$

The proofs of all theoretical results are contained in Appendix A.

The SDF bound (13) reminiscent of the Hansen and Jagannathan (1991), except in terms of entropy rather than variance. What is new is the marginal utility bound (14). Whereas the first object $L_t(\Delta S_{t+1})$ measures the true conditional volatility of $S = B\tilde{S}$, which jointly captures contributions from beliefs (B) and risks ($\tilde{S}$), the second object $\tilde{L}_t(\Delta\tilde{S}_{t+1})$ isolates the role of risks. If CEI holds, then $\tilde{L}_t(\Delta\tilde{S}_{t+1}) = L_t(\Delta\tilde{S}_{t+1})$, so that the perceived marginal utility volatility equals true marginal utility volatility. In that case, comparing the volatilities from (13)-(14) provides information about the importance of the belief distortion B, which is the wedge between S and $\tilde{S}$.

A similar discussion holds for the permanent components H and $\tilde{H}$. Whereas (15) measures the total volatility coming from beliefs (B) and long-run risks ($\tilde{H}$) together, the subjective bound (16) measures the perceived conditional volatility of the long-run risks ($\tilde{H}$) alone. The gap between them informs us about how important beliefs alone are to asset pricing.

At first glance, it is not obvious why $\tilde{\mathbb{E}}_t[\log R_{t+1}] - \log R_t^f$ and $\tilde{\mathbb{E}}_t[\log R_{t+1} - \log R_t^{\infty}]$ are informative about risks, rather than beliefs. Is it not possible that investors are just optimistic about an asset return R, which has nothing to do with their marginal utility $\tilde{S}$? No-arbitrage rules this out: if investors are optimistic about an asset, they must perceive it to have *some* risk. Since marginal utility $\tilde{S}$ summarizes investors' risk attitudes, this risk necessarily manifests in their marginal utility volatility.

One additional benefit of working under subjective beliefs is the direct availability of conditional expectations. In particular, survey data on subjective expected returns contain proxies for $\tilde{\mathbb{E}}_t[\log R_{t+1}]$ and $\tilde{\mathbb{E}}_t[\log R_{t+1}^{\infty}]$, which combined with Lemma 1 delivers time series of bounds for $\tilde{L}_t(\Delta\tilde{S}_{t+1})$ and $\tilde{L}_t(\Delta\tilde{H}_{t+1})$. By contrast, it is often more difficult to find a reasonable proxy for objective conditional expectations of returns (although I

will present a few candidates from the literature).

Building on the conditional results of Lemma 1, I also present unconditional bounds. For the objective probability, this is straightforward. On the other hand, the subjective bound does not generalize immediately to a useful unconditional version. Indeed, the *unconditional* subjective expectation $\tilde{\mathbb{E}}$ is not directly observable; taking time-series averages of *conditional* subjective expectations $\tilde{\mathbb{E}}_t$ will not yield an estimate of $\tilde{\mathbb{E}}$ unless beliefs are rational. We can make progress if we invoke UEB in Lemma 1 prior to taking time series averages.

Proposition 1. *The following unconditional volatility bounds hold for all returns R_{t+1}:*

$$L\Big(\frac{S_{t+1}}{S_t}R_t^f\Big) \geq \mathbb{E}\big[\log R_{t+1} - \log R_t^f\big] \tag{17}$$

$$L\Big(\frac{H_{t+1}}{H_t}\Big) \geq \mathbb{E}\big[\log R_{t+1} - \log R_{t+1}^{\infty}\big] \tag{18}$$

Additionally, if UEB holds for $\frac{\tilde{H}_{t+1}}{\tilde{H}_t}$ and $\frac{\tilde{S}_{t+1}}{\tilde{S}_t}$, then

$$L\Big(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}R_t^f\Big) \geq \mathbb{E}\Big[\tilde{\mathbb{E}}_t\big[\log R_{t+1} - \log R_t^f\big]\Big] \tag{19}$$

$$L\Big(\frac{\tilde{H}_{t+1}}{\tilde{H}_t}\Big) \geq \mathbb{E}\Big[\tilde{\mathbb{E}}_t\big[\log R_{t+1} - \log R_{t+1}^{\infty}\big]\Big] \tag{20}$$

Finally, while Lemma 1 and Proposition 1 hold in levels, we may want a ratio version in order to lower-bound the fraction of long-run SDF volatility that is due to marginal utility rather than beliefs. This approach requires more assumptions—namely the hypothesis that a return is "closer" to the objective growth-optimal return than a subjective counterpart (which holds, for example, if the chosen return is actually the growth-optimal return, but in other cases too). But the benefit is this approach permits a more intuitive scale as a ratio of volatilities.

Definition 3. For any return R, define its *distance to growth-optimal* and *distance to perceived growth-optimal*, respectively, by

$$\delta_t(R) := \mathbb{E}_t[\log R_{t+1}^*] - \mathbb{E}_t[\log R_{t+1}]$$
$$\tilde{\delta}_t(R) := \mathbb{E}_t[\log \tilde{R}_{t+1}^*] - \mathbb{E}_t[\log R_{t+1}].$$

Proposition 2. *Suppose UEB holds for $\frac{\tilde{H}_{t+1}}{\tilde{H}_t}$ and $\frac{\tilde{S}_{t+1}}{\tilde{S}_t}$. Let R be any return that is (on average) closer to growth-optimal than perceived growth-optimal, in the sense that $\mathbb{E}[\delta_t(R)] \leq \mathbb{E}[\tilde{\delta}_t(R)]$.*

Then, the following volatility bounds hold:

$$\frac{L(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}R_t^f)}{L(\frac{S_{t+1}}{S_t}R_t^f)} \geq \min\left\{1, \frac{\mathbb{E}\left[\tilde{\mathbb{E}}_t[\log R_{t+1} - \log R_t^f]\right]}{\mathbb{E}[\log R_{t+1} - \log R_t^f]}\right\} \tag{21}$$

$$\frac{L(\frac{\tilde{H}_{t+1}}{\tilde{H}_t})}{L(\frac{H_{t+1}}{H_t})} \geq \min\left\{1, \frac{\mathbb{E}\left[\tilde{\mathbb{E}}_t[\log R_{t+1} - \log R_{t+1}^{\infty}]\right]}{\mathbb{E}[\log R_{t+1} - \log R_{t+1}^{\infty}]}\right\} \tag{22}$$

The results so far constitute the core theoretical results of the paper, allowing me to compare risk-based volatility to total volatility and thus infer something about beliefs.

I conclude this section with a few auxiliary results, stated as corollaries, that measure the size of the permanent components of the SDF and marginal utility—the first one due to Alvarez and Jermann (2005), and the second one a novel subjective (and conditional) version. Whereas the Alvarez and Jermann (2005) bound quantifies the permanent component H, the subjective version below quantifies the subjective permanent component $\tilde{H}$, which is informative about how much risk investors perceive as coming from permanent sources. I will also report empirical estimates for these bounds.

Corollary 1.

$$\frac{L\left(\frac{H_{t+1}}{H_t}\right)}{L\left(\frac{S_{t+1}}{S_t}\right)} \geq \min\left\{1, \frac{\mathbb{E}\left[\log R_{t+1} - \log R_{t+1}^{\infty}\right]}{\mathbb{E}\left[\log R_{t+1} - \log R_t^f\right] + L(1/R_t^f)}\right\} \tag{23}$$

for any return R such that $\mathbb{E}[\log R_{t+1} - \log R_t^f] + L(1/R_t^f) > 0$.

Corollary 2.

$$\frac{\tilde{L}_t\left(\frac{\tilde{H}_{t+1}}{\tilde{H}_t}\right)}{\tilde{L}_t\left(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}\right)} \geq \min\left\{1, \max\left\{0, \frac{\tilde{\mathbb{E}}_t\left[\log R_{t+1} - \log R_{t+1}^{\infty}\right]}{\tilde{\mathbb{E}}_t\left[\log R_{t+1}\right] - \log R_t^f}\right\}\right\} \tag{24}$$

for any return R such that $\tilde{\mathbb{E}}_t \log R_{t+1} > \log R_t^f$.

3 Estimating the bounds

In the results to follow, I use the following return notations. Let R^m denote the aggregate stock market return. Given a positive number k, let y^k and R^k denote the (continuously-compounded) yield and holding return, respectively on a discount bond with maturity

k. For example, R^{10} denotes the return on a 10-year Treasury. The risk-free rate between t and $t+1$ will alternatively be written either as R^1_{t+1}, following the notation for other bonds, or as R^f_t, following the convention from the theory section. I assume a period length of one year, so the risk-free rate corresponds to the yield of a 1-year Treasury security.

3.1 Data and variable construction

The key data comes from multiple sources, summarized in Table 1.

	Frequency	Sample Period
A. Stocks		
Excess return	Monthly	1926:12–2020:12
Nagel-Xu	Quarterly	1972:06–1977:06[7] and 1987:03–2021:03
Livingston	Semi-annually	1952:05–2021:05
CFO	Quarterly	2000:08–2021:05
B. Treasuries		
Par yield, 1-year	Monthly	1953:04–2024:07
Par yield, 10-year	Monthly	1953:04–2024:07
Par yield, 30-year	Monthly	1961:06–2024:07
ZCB yield, 10-year	Monthly	1971:08–2024:07
ZCB yield, 30-year	Monthly	1985:11–2024:07
Return, 10-year	Monthly	1962:06–2024:07
Return, 30-year	Monthly	1962:06–2024:07
Return, index ($>$ 10-year)	Monthly	1954:04–2020:12
BCFF	Monthly	2001:01–2018:06

Table 1: Samples from key data sources

For the realized returns and yields, the data sources are standard. Realized returns on the aggregate value-weighted stock market are obtained from CRSP. I supplement this with a few volatility measures: I construct realized variance from daily return series obtained from Ken French's website; I use the CBOE VIX obtained from OptionMetrics; and I use the "SVIX" from Ian Martin's website (Martin, 2017). Realized Treasury yields and returns come from the St. Louis Fed (FRED), the Fed Board updates to Gürkaynak et al. (2007) (GSW), and CRSP. When there is overlapping data on par yields, I use the FRED yields as the first choice, followed by the GSW second, followed by the CRSP last. For zero-coupon yields, GSW is the only source. Appendix C.1 displays these various

[7] From 1972 to 1977, the data is annual only.

sources for comparison. For bond returns, CRSP is the only source. From the fixed-term indexes file, I obtain return series corresponding to an approximately 10-year and 30-year bond, respectively. I also obtain returns on a bond index tracking bonds with maturities above 10 years.

Survey-based US stock market expected returns come from Nagel and Xu (2023), who compile three sources: (i) the individual investor expectations from Nagel and Xu (2022), which agglomerate the UBS/Gallup survey, the Conference Board survey, and the Michigan Survey of Consumers, plus a few smaller surveys; (ii) the Livingston Survey from the Philadelphia Fed; and (iii) the Graham-Harvey CFO survey (Ben-David et al., 2013). From these sources, Nagel and Xu (2023) construct proxies for annual expected excess returns on the stock market (or S&P 500, depending on the source). See Nagel and Xu (2023) for more details on this data.

I convert expected stock returns $\tilde{\mathbb{E}}_t[R^m_{t+1}]$ to expected log returns $\tilde{\mathbb{E}}_t[\log R^m_{t+1}]$ via the identity $\tilde{\mathbb{E}}_t[\log R_{t+1}] = \log \tilde{\mathbb{E}}_t[R_{t+1}] - \tilde{L}_t(R_{t+1})$. While I do not have access to a measure of perceived entropy $\tilde{L}_t(R_{t+1})$, the lognormal approximation implies $\tilde{L}_t(R_{t+1}) \approx \frac{1}{2}\tilde{\text{Var}}_t[\log R_{t+1}]$. Nagel and Xu (2023) construct a measure of $\tilde{\text{Var}}_t[\log R^m_{t+1}]$ from the elicited 10th and 90th percentile return forecasts in the CFO survey. To extend this subjective variance measure to a longer history, I regress it on measures of uncertainty that are available further back in time. First, I regress the subjective variance $\tilde{\text{Var}}_t[\log R^m_{t+1}]$ onto the contemporaneous squared VIX, its one-month lag, and its trailing-12-month moving average. This regression generates the fitted value $\hat{y}_t = 0.012 + 0.52\text{VIX}^2_t + 0.005\text{VIX}^2_{t-1} + 0.071(\frac{1}{12}\sum_{j=0}^{11}\text{VIX}^2_{t-j})$, has an R-squared of 0.60, and its fit appears to be roughly a slightly smoothed version of the CFO subjective variance estimate.[8] I then backfill fitted values for months that pre-date the VIX by using the Nagel and Xu (2023) procedure of projecting VIX onto the news-implied volatility (NVIX) measure of Manela and Moreira (2017). Figure 12 in Appendix C.3 displays the CFO's reported subjective variance versus its fitted value from this procedure. The result of this is a time series of fitted and backfilled subjective variance estimates $\tilde{\text{Var}}_t[\log R^m_{t+1}]$ that I use to construct expected log market returns as $\tilde{\mathbb{E}}_t[\log R^m_{t+1}] = \log \tilde{\mathbb{E}}_t[R^m_{t+1}] - \frac{1}{2}\tilde{\text{Var}}_t[\log R^m_{t+1}]$. The result-

[8]The VIX is theoretically an appropriate regressor. Indeed, Result 3 of Martin (2017) shows that $\text{VIX}^2_{t\to T} = \frac{2}{T-t}L^*_t(R^m_{t\to T})$, where L^* is the risk-neutral entropy. To the extent that the risk-neutral and investors' subjective entropies roughly coincide at the one-year horizon (as they would in a conditionally log-normal model), the VIX is the ideal regressor. That said, the standard VIX is a one-month-ahead implied volatility, while the independent variable $\tilde{\text{Var}}_t[\log R^m_{t+1}]$ is a one-year ahead volatility. In ongoing work, I am constructing the one-year VIX from options to improve this step. To add some information about uncertainty at the one-year horizon, I have also tried adding the one-year squared SVIX (Martin, 2017) to this regression and found very little difference in the resulting fitted value of CFO perceived return volatility—in particular, for the sample where SVIX is available, I find an R-squared of 0.778 from including it, versus 0.774 without it.

ing expected excess log stock returns are depicted in Figure 1.

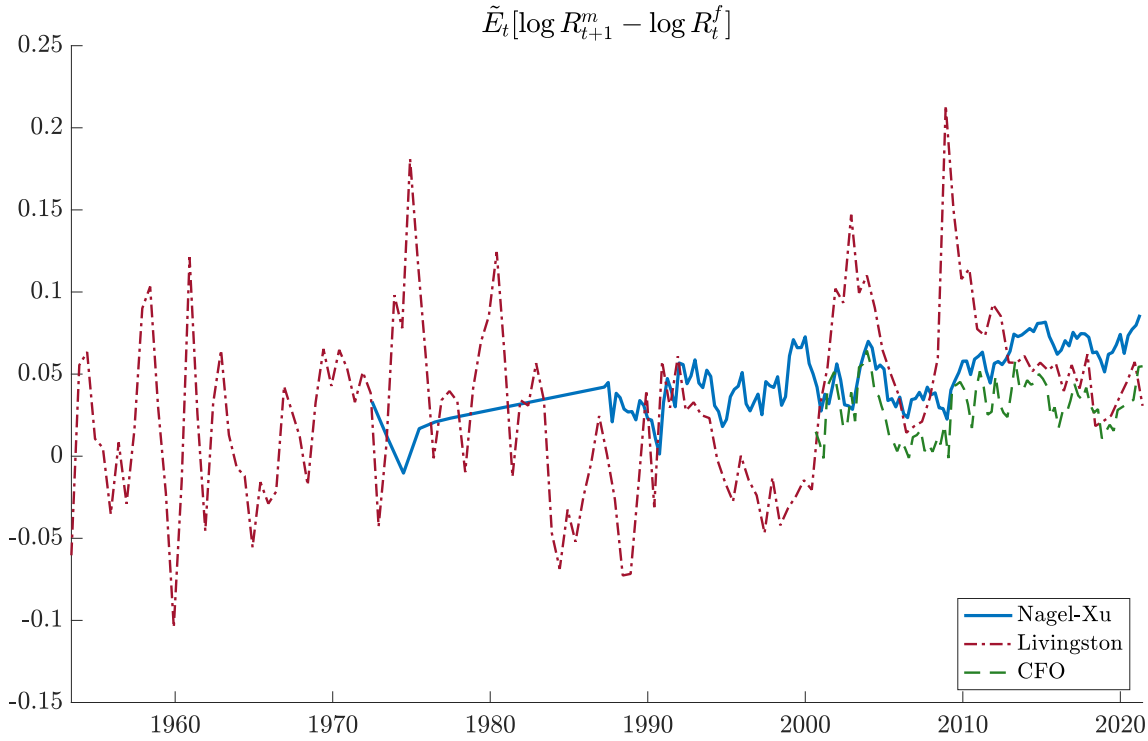


Figure 1: Expected one-year log stock returns in excess of the one-year Treasury rate. Expected log returns are obtained from expected arithmetic returns via a variance adjustment that is described in the text.

Treasury Bond expected returns are constructed from yield forecasts in the Blue Chip Financial Forecasts (BCFF) survey. BCFF has consensus (i.e., averaged across respondents) yield forecasts for Treasuries with maturities 6 months, 1 year, 2 years, 5 years, 10 years, and 30 years.[9] I follow Kim and Orphanides (2012) in treating the BCFF forecasts as approximately mid-quarter forecasts, and so I use a weighted-average of the 3-, 4-, and 5-quarter ahead forecasts, with weights depending on the survey calendar month, to get a forecast at a one-year horizon. Since the BCFF surveys are taken nearer to the beginning of the month, I also shift the responses and align them with with previous-month yields and asset prices, which are measured at month-end (this shift does not materially affect any result). I then interpolate these yields to obtain a full term structure by linear interpolation for maturities up to 5 years and a fitted Nelson-Siegel model for maturities beyond 5 years (with the Nelson-Siegel model fitted to expectations data on maturities 2+ years). The result is a term structure of one-year-ahead expectations of the par yield curve. I then bootstrap the approximate expected continuously compounded

[9]From March 2002 through February 2006, there was no 30-year bond issued, so the survey records a forecast of the 20-year yield. I account for this in all results with both the survey expected returns and realized returns.

zero-coupon yield curve $(\tilde{\mathbb{E}}_t[y^k_{t+1}])_{k=1}^{30}$ from this expected par yield curve.[10] Appendix C.2 contains additional documentation of the transformation from the raw data (sparse par yield forecasts) to an interpolated par yield curve forecast to a bootstrapped zero-coupon yield curve forecast. Finally, I construct one-year subjective expected log returns via $\tilde{\mathbb{E}}_t[\log R^k_{t+1}] = ky^k_t - (k-1)\tilde{\mathbb{E}}_t[y^{k-1}_{t+1}]$ at the $k = 10$-year and $k = 30$-year maturities.

As an extension, I also project the bond expectations back and forward in time by regressing them on various yields and their lags.[11] I regress $\tilde{\mathbb{E}}_t y^9_{t+1}$ (resp., $\tilde{\mathbb{E}}_t y^{29}_{t+1}$) onto the time-$t$ 1-year and 10-year (30-year) par yields, as well as their one-month lagged values and a trailing-12-month average of their values. The in-sample R-squareds from these regressions are 0.9821 and 0.9612, respectively. I use this estimation to construct hypothetical values of $\tilde{\mathbb{E}}_t[\log R^{10}_{t+1}]$ and $\tilde{\mathbb{E}}_t[\log R^{30}_{t+1}]$ going back as far as I have actual yield data.

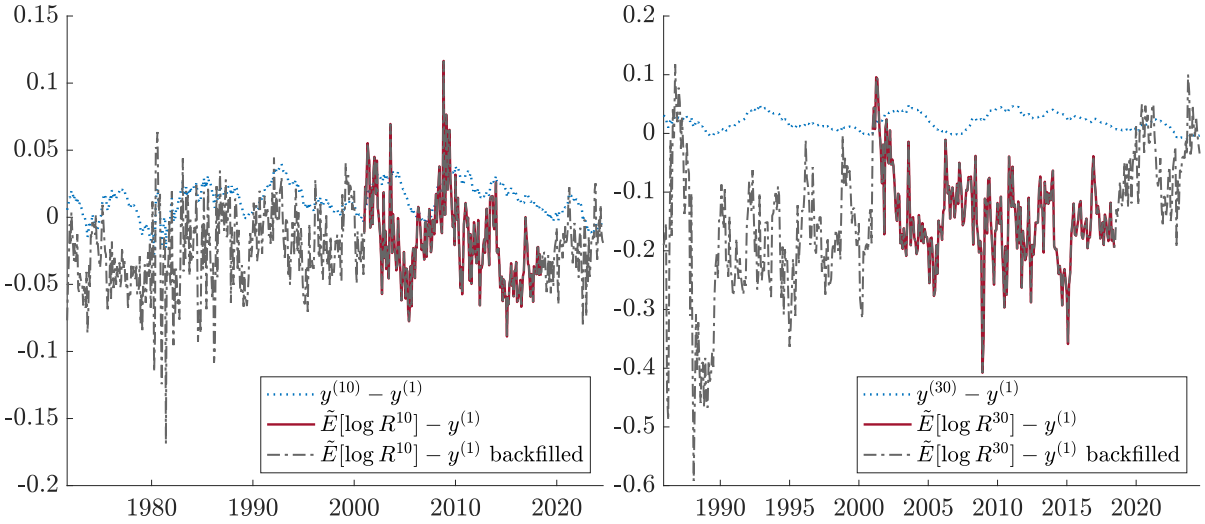


Figure 2: Expected log (zero-coupon) bond returns constructed using BCFF consensus forecasts of the one-year ahead yield curve. "Backfilled" refers to expected returns constructed from projecting yield forecasts onto the contemporaneous and lagged yield curve.

The resulting subjective expected excess log bond returns, as well as their extended backfilled values, are depicted in Figure 2. Notice from the figure that the subjective expected returns tend to be below the long-maturity yield. Since the yield measures the long-term expected return, this implies that survey respondents are pessimistic about bond returns in this sample. The situation is particularly extreme in 30-year bonds. As Figure 3 shows, this stems from pessimism about yields: survey respondents think

[10] Technically, this bootstrapping is a nonlinear transformation. I am effectively assuming a small amount of uncertainty in the par yield expectations to construct the zero-coupon yield expectations using this approximation.

[11] I am in the process of obtaining a longer time series of BCFF data and will de-emphasize this backfilling in a subsequent version.

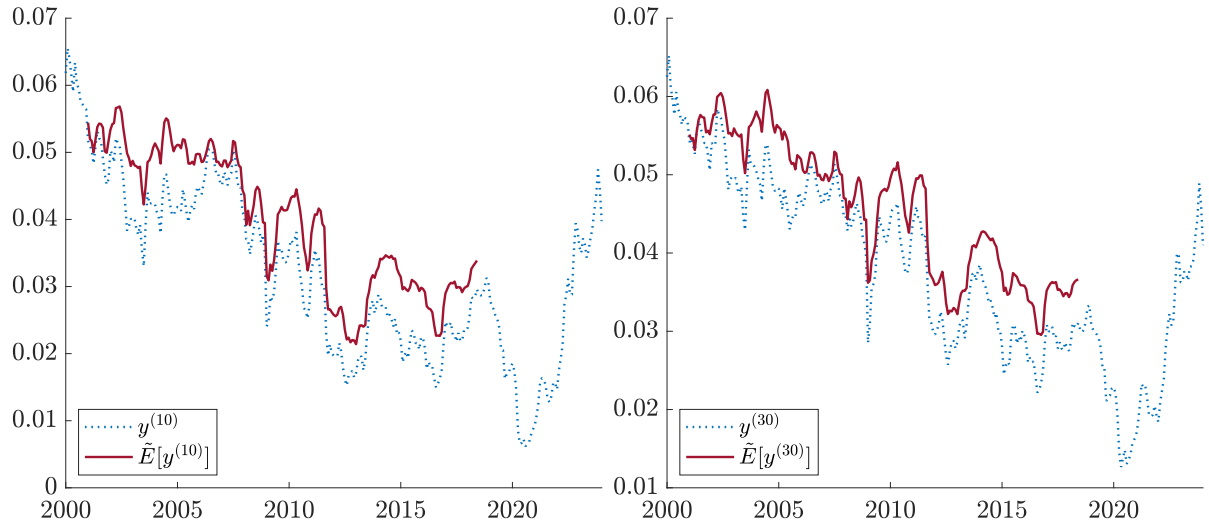


Figure 3: Par yields and the BCFF consensus one-year ahead forecasts of those par yields. All yields and yield forecasts are converted from coupon-equivalent to continuously compounded units.

yields will either stop falling or even rise during this period. The extremely pessimistic 30-year bond expected return is then the result of general pessimism about yields (Figure 3 shows that both the 10-year and 30-year yield forecasts are consistently about 0.5% above yield levels) combined with the significantly higher duration of the 30-year bond.

3.2 Results

To start, Table 2 reviews and refreshes the original Alvarez and Jermann (2005) results on the size of the permanent component in the SDF. Overall, the message from the point estimates is that at least 2/3 of the volatility of S stems from H. The reason: the equity premium (5.55%) is substantially larger than the term premium (between 1% and 2%, depending on the method), so that $\mathbb{E}[\log R^m - \log R^\infty]$ is likely to be nearly as large as the entire equity premium. Standard errors, shown in parentheses, are computed using a Newey and West (1987) window with 36 months of lags to account for the overlap in returns and persistence in spreads. The standard errors for the size of the permanent component are computed using the Delta Method. Based on these standard errors, the estimates appear to be reasonably precise.

As in the original paper, I use several methods for the long bond component. I use both holding period returns (Panel A) and yields (Panel B) as alternatives to measure expected returns. In a stationary and ergodic environment, the time-series average yield also measures a bond expected return; yields are much less volatile (though more persistent) than holding returns, which explains why Panel B features much tighter standard errors. I also examine both the 10-year and 30-year maturities as proxies for the long

bond. While 30-year bonds are preferable in principle, the discrepancy in results is negligible in practice.

	Equity Premium $E[\log R^m/R^1]$	Term Premium $E[\log R^{(k)}/R^1]$	Adjustment for Volatility of Short Rate $L(1/R^1)$	Size of Permanent Component $L(\Delta H)/L(\Delta S)$
A. Holding returns.				
$k = 10$ years	0.0555	0.0100	0.0005	0.8134
	(0.0161)	(0.0099)	(0.0001)	(0.1641)
$k = 30$ years		0.0091		0.8292
		(0.0143)		(0.2308)
$k > 10$ (index)		0.0125		0.7686
		(0.0087)		(0.1564)
B. Yields.				
$k = 10$ years (par)	0.0555	0.0097	0.0005	0.8185
	(0.0161)	(0.0016)	(0.0001)	(0.0566)
$k = 10$ years (zcb)		0.0129		0.7615
		(0.0022)		(0.0718)
$k = 30$ years (par)		0.0119		0.7791
		(0.0027)		(0.0692)
$k = 30$ years (zcb)		0.0195		0.6428
		(0.0031)		(0.1096)

Table 2: Objective Alvarez and Jermann (2005) volatility bound from Corollary 1. Panel A uses holding period returns to measure the term premium; each row uses a different proxy for R^∞. Panel B uses yields (par or zero-coupon bond yields) to measure the term premium. Newey-West asymptotic standard errors using 36 months of lags are shown in parentheses.

Next, I perform a similar analysis in the subjective measure. Table 3 presents estimates for the average size of the permanent component $\tilde{H}$ in marginal utility $\tilde{S}$. The estimates in the final column are the time-series average of the conditional subjective bound from Corollary 2. (Figure 4 below shows the time series of this conditional subjective bound, before taking averages.) As before, standard errors, shown in parentheses, are computed using a Newey and West (1987) window with 3 years of lags; since the equity surveys feature different sampling frequencies, this means a different number of period lags in each panel (Nagel-Xu and CFO data use 12 quarter lags; the Livingston survey use 6 semi-annual lags). The results, across the three equity surveys and four different proxies for the long bond expected return, broadly echo the message of Alvarez and Jermann (2005): investors perceive that the lion's share of marginal utility volatility (90% or more) emanates from permanent risks.

In my sample, I find a very negative subjective long bond term premium, particularly

	Subj. Equity Premium $E[\tilde{E}_t(\log R^m_{t+1}/R^1_{t+1})]$	Subj. Term Premium $E[\tilde{E}_t(\log R^{(k)}_{t+1}/R^1_{t+1})]$	Size of Subj. Permanent Component $E[\tilde{L}_t(\Delta\tilde{H}_{t+1})/\tilde{L}_t(\Delta\tilde{S}_{t+1})]$
A. Nagel-Xu.			
$k = 10$ years	0.0476	-0.0175	0.8817
	(0.0044)	(0.0075)	(0.0547)
$k = 30$ years		-0.1454	0.9740
		(0.0137)	(0.0234)
$k = 10$ (backfill)		-0.0226	0.9098
		(0.0035)	(0.0294)
$k = 30$ (backfill)		-0.1499	0.9737
		(0.0161)	(0.0139)
B. Livingston.			
$k = 10$ years	0.0300	-0.0175	0.9080
	(0.0081)	(0.0075)	(0.0383)
$k = 30$ years		-0.1454	0.9719
		(0.0137)	(0.0271)
$k = 10$ (backfill)		-0.0226	0.9288
		(0.0035)	(0.0215)
$k = 30$ (backfill)		-0.1499	0.9478
		(0.0161)	(0.0260)
C. CFO.			
$k = 10$ years	0.0297	-0.0175	0.8686
	(0.0038)	(0.0075)	(0.0561)
$k = 30$ years		-0.1454	0.9848
		(0.0137)	(0.0149)
$k = 10$ (backfill)		-0.0226	0.8816
		(0.0035)	(0.0498)
$k = 30$ (backfill)		-0.1499	0.9627
		(0.0161)	(0.0254)

Table 3: Subjective versions of the Alvarez and Jermann (2005) volatility bound from Corollary 2. Panels A, B, and C use the Nagel-Xu, Livingston, and CFO survey expectations for the market, respectively. The rows labeled "backfill" refer to BCFF forecasts fitted to the contemporaneous and lagged yield curve, and then projected to other periods in which the survey is not available. Newey-West asymptotic standard errors using 3 years of lags (i.e., 12 periods for Nagel-Xu, 6 periods for Livingston, 12 periods for CFO) are shown in parentheses.

for the 30-year bond. This is the source of the large estimated subjective permanent component. One caveat is that the BCFF sample period (2001–2018) only covers a period of low long-term yields, which survey respondents expected to revert. Thus, it is possible that the extremely low subjective term premia is a particular result of the sample period. The backfilled yield forecasts, which lead to very similar results in Table 3, may not

address this issue because they inherit the estimated in-sample pessimism from the BCFF survey.[12]

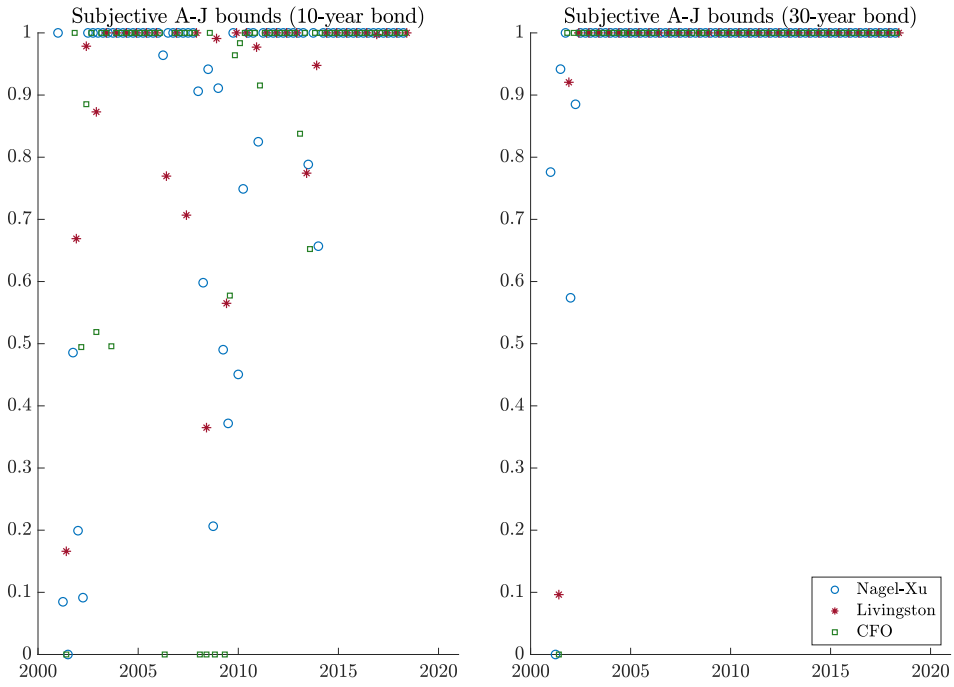


Figure 4: Time series of subjective versions of the Alvarez and Jermann (2005) volatility bound from Corollary 2.

Table 4 presents the main empirical result that risk accounts for at least half the volatility of the SDF. The table displays each unconditional mean in Proposition 1, followed by ratios of the subjective-to-objective bounds. Strictly speaking, the subjective-to-objective ratio may not be meaningful if the stock market is far from the growth-optimal portfolio. But as Proposition 2 makes clear, if the stock market is close to growth-optimal, the ratios provide lower bounds on the contribution of risk to various components of the SDF (i.e., to the risk-neutral density $R^f \Delta S$ or the permanent component ΔH, respectively). A weaker interpretation is also permitted: if S represents an SDF that prices the stock market, riskless rate, and long bond (which may be taken to be unique by projecting into this space of returns), then all I require for the subjective-to-objective ratios to be meaningful is that the stock market is close to growth-optimal *among this subset of assets*.

[12]For example, when projecting the 29-year zero-coupon yield forecast onto the 1-year and 30-year par yields, their one-month lags, and their trailing-12-month average, I estimate

$$\widehat{\tilde{\mathbb{E}}_t[y_{t+1}^{(29)}]} = -0.074 + 0.316 Y_t^{(1)} + 0.614 Y_t^{(30)} - 0.097 Y_{t-1}^{(1)} + 0.182 Y_{t-1}^{(30)} - 0.434 \Big(\frac{1}{12} \sum_{j=0}^{11} Y_{t-j}^{(1)} \Big) + 0.518 \Big(\frac{1}{12} \sum_{j=0}^{11} Y_{t-j}^{(30)} \Big)$$

with an R-squared of 0.9612. The negative estimated constant can be interpreted as estimated pessimism.

Standard errors assume the same Newey-West lag structure as before, with the standard errors for the volatility ratios computed using the Delta Method.

The extreme pessimism in long bond forecasts, particularly the 30-year bond, results in enormous estimates of $L(\Delta\tilde{H})/L(\Delta H)$ in the final column. As mentioned earlier, such a large volatility ratio estimate may be due partly to the specificity of the data sample; the very large standard errors suggest that the large estimate may also be due to the tremendous volatility in long-term bond expectations. If $L(\Delta\tilde{H})/L(\Delta H) > L(R^f\Delta\tilde{S})/L(R^f\Delta S)$ is true, even if the gap is not as large as the point estimates of Table 4 suggest, this means that the permanent component H is driven by risks to a greater extent than is the overall SDF S. This can happen if belief fluctuations partly manifest in transitory SDF variation, i.e., via G. The example models in Section 4 contain mechanisms of this type.

	Obj. Equity Premium $E[\log R^m/R^1]$	Subj. Equity Premium $E[\tilde{E}_t(\log R^m_{t+1}/R^1_{t+1})]$	Size of Risks in SDF $L(R^1\Delta\tilde{S})/L(R^1\Delta S)$	Obj. Equity minus Term Premium $E[\log R^m/R^{(k)}]$	Subj. Equity minus Term Premium $E[\tilde{E}_t(\log R^m_{t+1}/R^{(k)}_{t+1})]$	Size of Risks in Permanent Component $L(\Delta\tilde{H})/L(\Delta H)$
A. Nagel-Xu.						
$k = 10$ years	0.0555	0.0476	0.8567	0.0455	0.0651	1.4290
	(0.0161)	(0.0044)	(0.2449)	(0.0172)	(0.0081)	(0.4817)
$k = 30$ years				0.0464	0.1930	4.1568
				(0.0184)	(0.0138)	(1.6913)
B. Livingston.						
$k = 10$ years	0.0555	0.0300	0.5412	0.0455	0.0476	1.0443
	(0.0161)	(0.0081)	(0.1566)	(0.0172)	(0.0069)	(0.3165)
$k = 30$ years				0.0464	0.1755	3.7795
				(0.0184)	(0.0127)	(1.3275)
C. CFO.						
$k = 10$ years	0.0555	0.0297	0.5349	0.0455	0.0472	1.0367
	(0.0161)	(0.0038)	(0.1524)	(0.0172)	(0.0071)	(0.3365)
$k = 30$ years				0.0464	0.1751	3.7720
				(0.0184)	(0.0132)	(1.5321)

Table 4: Estimates of the importance of risks versus belief volatility, as in Proposition 1. Panels A, B, and C use the Nagel-Xu, Livingston, and CFO survey expectations for the market, respectively. Newey-West asymptotic standard errors using 3 years of lags (i.e., 12 periods for Nagel-Xu, 6 periods for Livingston, 12 periods for CFO) are shown in parentheses.

4 Example structural models

I provide three example models to help interpret the general theory above. All models are exchange economies but with differing endowment and belief structures. The first model features rational expectations about a time-varying long-run growth rate (Bansal and Yaron, 2004). The second model features distorted beliefs, formed using past growth extrapolation, about future growth (Collin-Dufresne et al., 2017; Nagel and Xu, 2022). In these two settings, there is a representative agent with recursive preferences a la Epstein and Zin (1989) and elasticity of intertemporal substitution (EIS) equal to one. These preferences are described by the utility recursion (Bellman equation)

$$U_t = (1-\beta)\log C_t + \beta \frac{\log \tilde{\mathbb{E}}_t[\exp((1-\gamma)U_{t+1})]}{1-\gamma}, \tag{25}$$

where $\tilde{\mathbb{E}}$ represents the agent's subjective expectation, γ is her risk aversion, and β is her discount factor. Finally, the third model also features growth extrapolation but with CRRA preferences for comparison. After presenting the details of the environments, Section 4.4 compares them across various measures of SDF volatility.

4.1 Long-run risks

In this canonical long-run risks example, the agent has rational expectations: $\tilde{\mathbb{E}} = \mathbb{E}$. There are two shocks which are jointly Normal, $W := (W^{(1)}, W^{(2)})' \sim N(0, I)$. Aggregate consumption follows

$$\log C_{t+1} = \log C_t + x_t + \sigma_c \left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right) \cdot W_{t+1},$$

and its *expected growth rate* x_t has AR(1) dynamics

$$x_{t+1} = (1-\rho)\bar{x} + \rho x_t + \sigma_x \left(\begin{smallmatrix} (1-\kappa^2)^{1/2} \\ \kappa \end{smallmatrix}\right) \cdot W_{t+1} \tag{26}$$

The parameter ρ captures the persistence of expected consumption, and κ controls the correlation between consumption and growth shocks. A typical specification sets $\kappa = 1$, but this flexibility will be helpful going forward.

Solving the model, the value function is given by[13]

$$U_t = \text{constant} + \log C_t + \frac{\beta}{1-\beta\rho} x_t$$

From this solution, we can write the SDF explicitly as

$$\begin{aligned} \frac{S_{t+1}}{S_t} &= \beta \frac{C_t}{C_{t+1}} \frac{\exp[(1-\gamma)U_{t+1}]}{\mathbb{E}_t[\exp[(1-\gamma)U_{t+1}]]} \\ &= \exp\Big\{ -r_t - \pi \cdot W_{t+1} - \frac{1}{2}\|\pi\|^2 \Big\}, \end{aligned} \tag{27}$$

where the one-period interest rate $r_t = \log R_t^f$ and market price of risk π are given by

$$r_t := -\log(\beta) + x_t + \frac{1}{2}\sigma_c^2 - \gamma\sigma_c^2 - \frac{(\gamma-1)\beta}{1-\beta\rho}(1-\kappa^2)^{1/2}\sigma_x\sigma_c \tag{28}$$

$$\pi := \gamma\sigma_c \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \frac{(\gamma-1)\beta}{1-\beta\rho}\sigma_x \begin{pmatrix} (1-\kappa^2)^{1/2} \\ \kappa \end{pmatrix}. \tag{29}$$

The permanent and transitory components of the SDF are given by[14]

$$\frac{H_{t+1}}{H_t} = \exp\Big\{ -\frac{1}{2}\Big\| \pi + \frac{\sigma_x}{1-\rho}\begin{pmatrix} (1-\kappa^2)^{1/2} \\ \kappa \end{pmatrix} \Big\|^2 - \Big(\pi + \frac{\sigma_x}{1-\rho}\begin{pmatrix} (1-\kappa^2)^{1/2} \\ \kappa \end{pmatrix} \Big) \cdot W_{t+1} \Big\}. \tag{30}$$

and

$$\frac{G_{t+1}}{G_t} = \exp\Big[\eta + \frac{1}{1-\rho}(x_{t+1} - x_t) \Big], \tag{31}$$

[13]Guess that the value function is given by the form $U_t = \log C_t + u_0 + u_x x_t$, and obtain the values of u_0 and u_x by substituting this guess into the Bellman equation (25) and using the method of undetermined coefficients. The "constant" term u_0 (not reported in the main text) is

$$u_0 = \frac{\beta}{1-\beta}\Big[(1-\rho)\bar{x}u_x + \frac{1}{2}(1-\gamma)\big([\sigma_c + (1-\kappa^2)^{1/2}\sigma_x u_x]^2 + \kappa^2\sigma_x^2 u_x^2 \big) \Big]$$

[14]Guess that the transitory piece is given by $G_t = e^{\eta t - \zeta(x_t - x_0)}$ for some constants (η, ζ), and then substitute the guess $S_t = G_t H_t$ into the SDF formula. This yields

$$\begin{aligned} \frac{H_{t+1}}{H_t} &= \exp(-\eta + \zeta(x_{t+1} - x_t)) \frac{S_{t+1}}{S_t} \\ &= \exp\Big\{ -\eta - r_t - \big(\pi - \zeta\sigma_x \begin{pmatrix} (1-\kappa^2)^{1/2} \\ \kappa \end{pmatrix}\big) \cdot W_{t+1} - \frac{1}{2}\|\pi\|^2 + \zeta(1-\rho)(\bar{x} - x_t) \Big\} \end{aligned}$$

Using the conjecture that H is a martingale allows us to solve for η and ζ with the method of undetermined coefficients. These values are then $\zeta = -(1-\rho)^{-1}$ and $\eta = \log(\beta) - \frac{1}{2}\sigma_c^2 + \gamma\sigma_c^2 + \frac{(\gamma-1)\beta}{1-\beta\rho}(1-\kappa^2)^{1/2}\sigma_x\sigma_c + \zeta(1-\rho)\bar{x} + \frac{1}{2}\zeta^2\sigma_x^2 - \zeta\sigma_x \begin{pmatrix} (1-\kappa^2)^{1/2} \\ \kappa \end{pmatrix} \cdot \pi$.

for some constant η (given in footnote 14).

This is a model that can generate a large amount of SDF volatility via a very prominent permanent component. Indeed, amount of conditional SDF volatility is

$$L_t\Big(\frac{S_{t+1}}{S_t}\Big) = \frac{1}{2}\|\pi\|^2$$

A typical calibration sets $\gamma = 10$, $\beta = 0.99$ (or higher), $\rho = 0.95$ (or higher), $\sigma_c = 0.02$, $\sigma_x = 0.002$, and $\kappa = 0.9$ (Bansal et al., 2012), in which case $L(\frac{S_{t+1}}{S_t}) = \frac{1}{2}(0.43)^2$, corresponding to an annual Sharpe ratio of 0.43. As argued by Hansen and Jagannathan (1991), this large conditional SDF volatility (i.e., large short-term risk prices) is an important target for asset-pricing models. This class of models obtains large short-term risk prices via a highly persistent growth rate (high ρ), combined with an investor who is both patient (high β) and risk averse (high γ)—see equation (29).

At the same time, the Alvarez and Jermann (2005) volatility ratio is given by

$$\frac{L(\frac{H_{t+1}}{H_t})}{L(\frac{S_{t+1}}{S_t})} = \frac{\big\|\pi + \frac{\sigma_x}{1-\rho}\big(\begin{smallmatrix}(1-\kappa^2)^{1/2}\\ \kappa\end{smallmatrix}\big)\big\|^2}{\|\pi\|^2 + \mathrm{Var}[r_t]}$$

Theoretically, one can show that that this expression is always greater than 1 if $\gamma \geq 1$.[15] Quantitatively, under the calibration just given, $L(\frac{H_{t+1}}{H_t})/L(\frac{S_{t+1}}{S_t}) = 1.18$, in large part because of small riskless rate volatility, $\mathrm{Var}[r_t] = 0.00004$. On the other hand, the transitory component of the SDF is relatively small, with

$$\frac{L(\frac{G_{t+1}}{G_t})}{L(\frac{S_{t+1}}{S_t})} = \frac{\frac{2}{1-\rho}\mathrm{Var}[r_t]}{\|\pi\|^2 + \mathrm{Var}[r_t]}$$

For instance, under the calibration given above, $L(\frac{G_{t+1}}{G_t})/L(\frac{S_{t+1}}{S_t}) = 0.01$.

While obvious, let us just mention that any volatility ratio of perceived-to-actual will be one in this rational expectations model. Because $B \equiv 1$, marginal utility coincides with the SDF, $\tilde{S} \equiv S$, and the perceived permanent component coincides with the actual

[15] The unconditional variance of r_t is computed as $\sigma_x^2/(1-\rho^2)$ using its AR(1) dynamics inherited from x_t. Then, using the fact that $\pi \geq 0$ when $\gamma \geq 1$, followed by the assumption $0 < \rho < 1$, we have $\|\pi + \frac{\sigma_x}{1-\rho}\big(\begin{smallmatrix}(1-\kappa^2)^{1/2}\\ \kappa\end{smallmatrix}\big)\|^2 \geq \|\pi\|^2 + (\frac{\sigma_x}{1-\rho})^2 \geq \|\pi\|^2 + \frac{\sigma_x^2}{1-\rho^2} = \|\pi\|^2 + \mathrm{Var}[r_t]$.

permanent component, $\tilde{H} \equiv H$. Therefore,

$$\frac{L(\frac{B_{t+1}}{B_t})}{L(\frac{S_{t+1}}{S_t})} = 0 \quad \text{and} \quad \frac{L(\frac{\tilde{S}_{t+1}}{\tilde{S}_t})}{L(\frac{S_{t+1}}{S_t})} = \frac{L(\frac{\tilde{H}_{t+1}}{\tilde{H}_t})}{L(\frac{H_{t+1}}{H_t})} = 1.$$

Working in reverse, the latter two ratios being near one can be interpreted as evidence of near-rational expectations. The magnitudes of various volatilities and volatility ratios, under the calibration described here, are consolidated at the end of this section in Table 5.

4.2 Growth extrapolation

Now, consider a model in which actual growth is IID but agents perceive a stochastic growth rate. It turns out, with the appropriate modeling of beliefs, we can just reinterpret the previous example as x representing the *subjective expected growth rate*, which remains constant at $\bar{x}$ under the rational expectation. This will be a model with long-run risks only in investors' heads.

Therefore, the solution will be identical to the actual long-run risks model but with a corresponding reinterpretation of the results. The expression in (27) now corresponds to investors' marginal utility $\tilde{S}_t$. Its permanent component $\tilde{H}_t$ is given by the expression in (30). In essence, the extrapolative growth model can be solved as a long-run risk model but then simulated by turning off long-run risks.

To justify these claims, consider the following enriched version of an extrapolative growth model. True consumption dynamics are

$$\log C_{t+1} = \log C_t + \bar{x} + \sigma_c \left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right) \cdot W_{t+1}.$$

While conditional mean growth is constant at $\bar{x}$, agents are not convinced of this fact, and they do not observe the objective shocks W either. Building on many extrapolation models, agents will learn about the mean consumption growth via *constant-gain learning*. I will also allow the learning process to be subject to *sentiment shocks*.

Mathematically, let x_t denote the agent's perceived expected growth rate, i.e., $x_t := \tilde{\mathbb{E}}_t[\log C_{t+1} - \log C_t]$. Then, the agent's perception of consumption dynamics are

$$\log C_{t+1} = \log C_t + x_t + \sigma_c \left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right) \cdot \tilde{W}_{t+1},$$

where $\tilde{W}$ represent perceived shocks, i.e., a normally-distributed variable which has

zero mean under the agent's belief, $\tilde{\mathbb{E}}_t[\tilde{W}_{t+1}] = 0$. Learning is modeled as follows. The posterior expected growth is given by

$$x_t = (1-\phi)\bar{x} + \phi z_t, \tag{32}$$

where z_t is a "sentiment" variable that follows

$$z_{t+1} - z_t = \underbrace{\lambda[\log C_{t+1} - \log C_t - z_t]}_{\text{constant gain learning}} + \underbrace{\left(\tfrac{\nu_1 - \lambda\sigma_c}{\nu_2}\right) \cdot \tilde{W}_{t+1}}_{\text{sentiment shocks}} \tag{33}$$

In words, sentiment increases with excess growth realizations $\log C_{t+1} - \log C_t - z_t$, at a fixed learning rate λ, and is also subject to additional shocks. After sentiment is determined, posterior expected growth is determined by a type of "model averaging" between the correct time-invariant growth $\bar{x}$ and the growth sentiment z_t. The parameter ϕ governs the degree of rationality: $\phi = 0$ is fully rational, while $\phi = 1$ is fully extrapolative. Alternatively, as shown in Nagel and Xu (2022), $1 - \phi$ captures the weight a constant-gain learner puts on his prior.

Combining equations (32)-(33) with the perceived consumption law of motion yields the belief evolution

$$x_{t+1} = \lambda(1-\phi)\bar{x} + (1 - \lambda(1-\phi))x_t + \phi\nu \cdot \tilde{W}_{t+1}, \tag{34}$$

where $\nu := (\nu_1, \nu_2)'$. From equation (34), we can see the analogy to the long-run risks model. By inspection, the dynamics in this equation are identical to those of equation (26), provided we set the prior dependence parameter to any interior value $\phi \in (0,1)$ and then put $\lambda = \frac{1-\rho}{1-\phi}$, $\nu_1 = \phi^{-1}(1-\kappa^2)^{1/2}\sigma_x$, and $\nu_2 = \phi^{-1}\kappa\sigma_x$.[16] In that case, the impact of λ is entirely through the following "perceived growth persistence" that I will use going forward:

$$\tilde{\rho} := 1 - (1-\phi)\lambda.$$

Notice also that setting $\phi < 1$ serves a technical convenience by keeping $\tilde{\rho} < 1$ so that x_t stationary is under the subjective probability.

Given this isomorphism, the long-run risks solution above fully characterizes the

[16]The special case without sentiment shocks ($\nu_1 = \lambda\sigma_c$ and $\nu_2 = 0$) would be like a long-run risk model with perfect correlation between level and growth shocks ($\kappa = 0$) and with a particular calibration of growth volatility ($\sigma_x = \phi\lambda\sigma_c$).

investor marginal utility $\tilde{S}_t$ and its permanent component $\tilde{H}_t$. These are given by

$$\frac{\tilde{S}_{t+1}}{\tilde{S}_t} = \exp\Big\{ -r_t - \tilde{\pi}\cdot\tilde{W}_{t+1} - \frac{1}{2}\|\tilde{\pi}\|^2\Big\} \tag{35}$$

$$\frac{\tilde{H}_{t+1}}{\tilde{H}_t} = \exp\Big\{ -\frac{1}{2}\Big\|\tilde{\pi} + \frac{\phi\nu}{1-\tilde{\rho}}\Big\|^2 - \Big(\tilde{\pi} + \frac{\phi\nu}{1-\tilde{\rho}}\Big)\cdot\tilde{W}_{t+1}\Big\} \tag{36}$$

where r_t and $\tilde{\pi}$ are the riskless rate and perceived risk prices, respectively,

$$\begin{aligned} r_t &:= -\log(\beta) + x_t + \frac{1}{2}\sigma_c^2 - \gamma\sigma_c^2 - \frac{(\gamma-1)\beta}{1-\beta\tilde{\rho}}\sigma_c\phi\nu_1 \\ \tilde{\pi} &:= \gamma\sigma_c\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right) + \frac{(\gamma-1)\beta}{1-\beta\tilde{\rho}}\phi\nu. \end{aligned}$$

Similarly, the transitory component of the SDF, which recall is identical under both the subjective and objective probability, is given by

$$\frac{G_{t+1}}{G_t} = \exp\Big[\tilde{\eta} + \frac{1}{1-\tilde{\rho}}(x_{t+1} - x_t)\Big],$$

for some constant $\tilde{\eta}$. This fully solves the model from the perspective of agent's subjective probability.

It remains to characterize the belief distortion B_t, in order to map this solution back into the objective probability. In this conditionally lognormal environment, all belief distortions are characterized by a mean distortion to the shocks. In other words, while W are the true zero-mean shocks, the agent perceives them to have non-zero mean, $\tilde{\mathbb{E}}_t[W_{t+1}] = \mu_t$, and instead views $\tilde{W}_{t+1} = W_{t+1} - \mu_t$ as the relevant disturbance. This means that

$$\frac{B_{t+1}}{B_t} = \exp\Big[\mu_t\cdot W_{t+1} - \frac{1}{2}\|\mu_t\|^2\Big] \tag{37}$$

for some μ_t. To identify this mean, compare the true and perceived consumption dynamics to see that

$$\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)\cdot\mu_t = \frac{x_t - \bar{x}}{\sigma_c}.$$

While this determines the first component of μ_t, the second component is not pinned down, because any value for it leads to the same perceived growth dynamics. Intuitively, pure sentiment shocks are not tied down by any objective dynamic. For theoretical

clarity, let us set

$$\binom{0}{1} \cdot \mu_t = 0,$$

so that all results are generated via mechanisms that are tied down. (This can be easily generalized, and allowing a non-zero and possibly time-varying second component of μ_t would work like irrational exuberance in this model.)

Combining equations (35)-(36) with the belief distortion (37), we have the objective SDF and its permanent component

$$\frac{S_{t+1}}{S_t} = \exp\Big\{ -r_t - (\tilde{\pi} - \mu_t) \cdot W_{t+1} - \frac{1}{2}\|\tilde{\pi} - \mu_t\|^2 \Big\} \tag{38}$$

$$\frac{H_{t+1}}{H_t} = \exp\Big\{ -\frac{1}{2}\Big\|\tilde{\pi} - \mu_t + \frac{\phi\nu}{1-\tilde{\rho}}\Big\|^2 - \Big(\tilde{\pi} - \mu_t - \frac{\phi\nu}{1-\tilde{\rho}}\Big) \cdot W_{t+1} \Big\} \tag{39}$$

Thus, objective risk prices are lowered by μ_t relative to perceived risk prices. For example, if agents are pessimistic about growth, then the first component of μ_t is negative, which raises risk prices.

Now, we may compute various volatility ratios, analogously to the previous model. For example, the annual average SDF volatility is

$$\mathbb{E}\Big[2L_t\Big(\frac{S_{t+1}}{S_t}\Big)\Big]^{1/2} = \Big(\mathbb{E}\|\tilde{\pi} - \mu_t\|^2\Big)^{1/2} = \Big(\|\tilde{\pi}\|^2 + \frac{\mathrm{Var}[r_t]}{\sigma_c^2}\Big)^{1/2}$$

Relative to the long-run risk model, extrapolation induces time-variation in objective risk pricing, which raises SDF volatility via the second term above. Note that $\mathrm{Var}[r_t] = \mathrm{Var}[x_t]$, i.e., the objective variance of agent's growth perceptions. The sizes of the permanent and transitory components are given by

$$\frac{L(\frac{H_{t+1}}{H_t})}{L(\frac{S_{t+1}}{S_t})} = \frac{\|\tilde{\pi} + \frac{\phi\nu}{1-\tilde{\rho}}\|^2 + \frac{1}{\sigma_c^2}\mathrm{Var}[r_t]}{\|\tilde{\pi}\|^2 + (1 + \frac{1}{\sigma_c^2})\mathrm{Var}[r_t]}$$

$$\frac{L(\frac{G_{t+1}}{G_t})}{L(\frac{S_{t+1}}{S_t})} = \frac{\frac{2(1-\tilde{\rho}-\frac{\phi\nu_1}{\sigma_c})}{(1-\tilde{\rho})^2}\mathrm{Var}[r_t]}{\|\tilde{\pi}\|^2 + (1 + \frac{1}{\sigma_c^2})\mathrm{Var}[r_t]}$$

Finally, the ratio of subjective-to-objective SDF volatility is given by

$$\frac{L(\frac{\tilde{S}_{t+1}}{\tilde{S}_t})}{L(\frac{S_{t+1}}{S_t})} = \frac{\|\tilde{\pi}\|^2 + (1 - \frac{\tilde{\pi}_1}{\sigma_c})^2\mathrm{Var}[r_t]}{\|\tilde{\pi}\|^2 + (1 + \frac{1}{\sigma_c^2})\mathrm{Var}[r_t]}$$

and its permanent counterpart by

$$\frac{L(\frac{\tilde{H}_{t+1}}{\tilde{H}_t})}{L(\frac{H_{t+1}}{H_t})} = \frac{\|\tilde{\pi} + \frac{\phi\nu}{1-\tilde{\rho}}\|^2 + (\tilde{\pi}_1 + \frac{\phi\nu_1}{1-\tilde{\rho}})^2 \frac{1}{\sigma_c^2}\text{Var}[r_t]}{\|\tilde{\pi} + \frac{\phi\nu}{1-\tilde{\rho}}\|^2 + \frac{1}{\sigma_c^2}\text{Var}[r_t]}$$

These are all easily computable analytically, because of the lognormality of the equilibrium solution.[17]

A standard calibration without sentiment shocks (i.e., pure extrapolation) sets $\gamma = 4$, $\beta = 0.99$, $\sigma_c = 0.02$, $\phi = 0.9$ (almost fully extrapolative), $\lambda = 0.075$ (interpretable as roughly 13 years of effective information retention), and $\nu_1 - \lambda\sigma_c = \nu_2 = 0$ (no sentiment shocks).[18] This calibration generates perceived conditional growth rate volatility $\phi\|\nu\| = 0.0013$ and a perceived persistence $\tilde{\rho} = 1 - (1-\phi)\lambda = 0.993$. Relative to the long-run risk model calibration, the extrapolative calibration features a lower short-run volatility of perceived growth but a higher persistence hence higher long-run volatility.

Because of the relatively high persistence of perceived growth, the permanent component is even larger here than in the long-run risk calibration, $L(\frac{H_{t+1}}{H_t})/L(\frac{S_{t+1}}{S_t}) = 2.13$. It turns out the transitory component is also larger, i.e., $L(\frac{G_{t+1}}{G_t})/L(\frac{S_{t+1}}{S_t}) = 0.26$, because of negative correlation between G and H. In particular, optimism about perceived growth (high x_t) raises the transitory component of marginal utility, while simultaneously reducing the permanent component.

Finally, the two subjective-to-objective volatility ratios, which measure the amount of

[17] Every object of interest (e.g., S, H, $\tilde{S}$, $\tilde{H}$, G, B) in this model takes the form

$$\frac{M_{t+1}}{M_t} = \exp\left\{a + bx_t - [p + q\mu_t] \cdot W_{t+1} - \frac{1}{2}\|p + q\mu_t\|^2\right\}$$

for some scalars a, b, q and vector p. The dynamics of x_t under the objective measure are an AR(1) with modified persistence:

$$x_{t+1} - \bar{x} = \left(1 - \lambda(1-\phi) - \frac{\phi\nu_1}{\sigma_c}\right)(x_t - \bar{x}) + \phi\nu \cdot W_{t+1}$$

so that $\mathbb{E}[x_t] = \bar{x}$ and $\text{Var}[x_t] = \frac{\phi^2\|\nu\|^2}{1-(1-\lambda(1-\phi)-\frac{\phi\nu_1}{\sigma_c})^2}$. Similarly, $\mathbb{E}[\mu_t] = 0$ and $\text{Var}[(\begin{smallmatrix}1\\0\end{smallmatrix}) \cdot \mu_t] = \frac{1}{\sigma_c^2}\text{Var}[x_t]$. Therefore, for each object whose growth takes the form of M,

$$L\left(\frac{M_{t+1}}{M_t}\right) = \log\mathbb{E}\left[e^{a+bx_t}\right] - (a + b\bar{x}) + \frac{1}{2}\mathbb{E}\|p + q\mu_t\|^2 = \frac{1}{2}(b^2 + \frac{q^2}{\sigma_c^2})\text{Var}[x_t] + \frac{1}{2}\|p\|^2.$$

[18] These parameters largely follow Nagel and Xu (2022). They set $\lambda = 0.018$ at a quarterly frequency, which translates to approximately 4×0.018 annually. Furthermore, in their model $1 - \phi$ can be interpreted as the weight (in belief updating) that agents place on their prior. Their baseline model sets $\phi = 1$, but this would cause the agent to believe that $\tilde{x}_t$ follows a random walk, which causes some technical issues for unconditional subjective volatilities. Thus, I pick a value ϕ close to 1 but slightly below.

SDF volatility coming from risk, are obviously smaller in this distorted belief model than a rational expectations model. However, under the calibration given, $L(\frac{\tilde{S}_{t+1}}{\tilde{S}_t})/L(\frac{S_{t+1}}{S_t}) = 0.77$ and $L(\frac{\tilde{H}_{t+1}}{\tilde{H}_t})/L(\frac{H_{t+1}}{H_t}) = 0.91$, which are not far from their rational-expectations benchmarks of 1. Intuitively, although subjective beliefs about growth are non-trivial here, a large chunk of their effect comes through their interaction with investors' preferences (specifically the preference for early resolution of uncertainty). The result might be called *fear of perceived growth*, a mechanism that causes marginal utility and its permanent component to be highly sensitive to perceived growth shocks.

As with the long-run risk model, the magnitudes of various volatilities and volatility ratios for this model are summarized at the end of this section in Table 5.

4.3 Growth extrapolation (CRRA)

For comparison, I also consider a version of the previous extrapolative growth model under CRRA utility with risk aversion γ. The sentiment and belief processes are identical to the previous setting. But now, marginal utility growth is simply

$$\frac{\tilde{S}_{t+1}}{\tilde{S}_t} = \beta\Big(\frac{C_{t+1}}{C_t}\Big)^{-\gamma} = \exp\Big\{-r_t - \frac{1}{2}\|\tilde{\pi}\|^2 - \tilde{\pi}\cdot\tilde{W}_{t+1}\Big\},$$

where the riskless rate and perceived risk prices are the familiar expressions

$$\begin{aligned} r_t &:= -\log(\beta) + \gamma x_t - \frac{1}{2}(\gamma\sigma_c)^2 \\ \tilde{\pi} &:= \gamma\sigma_c\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right) \end{aligned}$$

Notice that the structure of the solution is very similar, with two key differences. First, perceived short-run risk prices $\tilde{\pi}$ do not inherit any growth rate exposure. Second, the growth rate volatility enters instead into magnified interest rate volatility. The result of these two differences will be (i) much less marginal utility volatility, hence much more of SDF volatility attributable to beliefs; and (ii) much more of SDF volatility coming from transitory fluctuations.

Following the same procedures as above, we can compute the perceived permanent component, the objective SDF, the objective permanent component, and the transitory

component as

$$\begin{aligned}
\frac{\tilde{H}_{t+1}}{\tilde{H}_t} &= \exp\Big\{-\frac{1}{2}\Big\|\tilde{\pi}+\frac{\gamma\phi\nu}{1-\tilde{\rho}}\Big\|^2-\Big(\tilde{\pi}+\frac{\gamma\phi\nu}{1-\tilde{\rho}}\Big)\cdot\tilde{W}_{t+1}\Big\}\\
\frac{S_{t+1}}{S_t} &= \exp\Big\{-r_t-\frac{1}{2}\|\tilde{\pi}-\mu_t\|^2-(\tilde{\pi}-\mu_t)\cdot W_{t+1}\Big\}\\
\frac{H_{t+1}}{H_t} &= \exp\Big\{-\frac{1}{2}\Big\|\tilde{\pi}-\mu_t+\frac{\gamma\phi\nu}{1-\tilde{\rho}}\Big\|^2-\Big(\tilde{\pi}-\mu_t+\frac{\gamma\phi\nu}{1-\tilde{\rho}}\Big)\cdot W_{t+1}\Big\}\\
\frac{G_{t+1}}{G_t} &= \exp\Big[\tilde{\eta}+\frac{\gamma}{1-\tilde{\rho}}(x_{t+1}-x_t)\Big],
\end{aligned}$$

for some constant $\tilde{\eta}$. As suggested above, notice that the transitory component is magnified by risk aversion γ, whereas this was not the case in the previous recursive utility version.

4.4 Comparing the models

Table 5 provides a comparison of the three models in terms of various volatilities and volatility ratios. Comparing the extrapolation models to the long-run risk model, one sees that more SDF volatility comes from beliefs and more of it is transitory. This situation is extreme in the CRRA model, whose transitory SDF volatility is far too high relative to the data (Alvarez and Jermann, 2005).

	Long-Run Risk	Extrapolation (EZ)	Extrapolation (CRRA)
Annual Avg SDF Volatility			
$E[2L_t(\Delta S_{t+1})]^{1/2}$	0.4265	0.3574	0.1948
Size of Permanent Component			
$L(\Delta H)/L(\Delta S)$	1.1786	2.1275	17.5983
Size of Transitory Component			
$L(\Delta G)/L(\Delta S)$	0.0090	0.2635	14.1140
Risk Component			
$L(\Delta\tilde{S})/L(\Delta S)$	1.0000	0.7736	0.1677
Permanent Risk Component			
$L(\Delta\tilde{H})/L(\Delta H)$	1.0000	0.9118	0.9831
Belief Component			
$L(\Delta B)/L(\Delta S)$	0.0000	0.2471	0.8270

Table 5: Entropies and entropy ratios across the three example models, in their benchmark calibrations. The "Long-Run Risk" model has $\gamma = 10$, $\beta = 0.99$, $\rho = 0.95$, $\sigma_c = 0.02$, $\sigma_x = 0.002$, $\kappa = 0.9$. The "Extrapolation (EZ)" model has $\gamma = 4$, $\beta = 0.99$, $\phi = 0.9$, $\lambda = 0.075$, $\sigma_c = 0.02$, and $\nu = (\lambda\sigma_c, 0)'$. "Extrapolation (CRRA)" uses the same parameters as Extrapolation (EZ), except with CRRA preferences.

The extrapolation model with recursive utility succeeds in generating high SDF volatility, a large permanent component, and a modest transitory component. In fact, the calibrated values of $L(\Delta\tilde{S})/L(\Delta S)$ and $L(\Delta\tilde{H})/L(\Delta H)$ are not too far from the empirical estimates in Table 4. Interestingly, only 24.71% of total SDF volatility comes from beliefs directly in this model. This is because the action in that model (e.g., the large value of $L(\Delta S)$) arises via an *interaction* between preferences and perceived growth dynamics. While shutting down the belief dynamics mitigates SDF volatility, so does eliminating the preference for early resolution of uncertainty. In general, when SDF volatility is created by interaction effects, it is not a priori clear how to attribute it; my framework provides a clear-cut answer to that question.

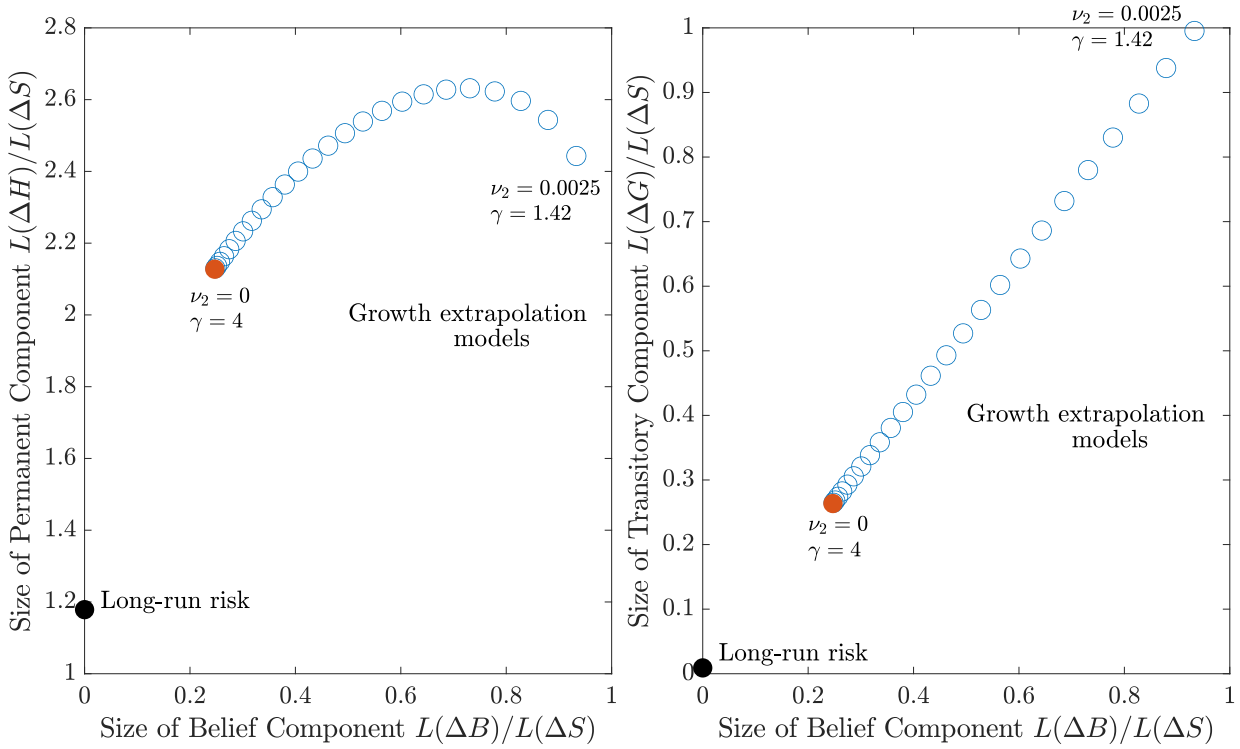


Figure 5: Comparison between the baseline long-run risk model (solid black dot), the baseline growth extrapolation model (solid red dot), and various growth extrapolation models with sentiment shocks (hollow blue circles), i.e., models with $\nu_2 > 0$. As ν_2 is increased, γ is also decreased to keep the measure of one-year average SDF volatility $\mathbb{E}[2L_t(\Delta S_{t+1})]^{1/2}$ unchanged at 0.3574 (the baseline value for the extrapolation model). The calibrations of the baseline models are the same as in Table 5.

In a final exercise, I also illustrate what happens when sentiment shocks are added the extrapolation model. For various models, Figure 5 displays results for the size of the permanent component $L(\Delta H)$, transitory component $L(\Delta G)$, and belief component $L(\Delta B)$, all scaled by the total volatility of the SDF $L(\Delta S)$. The extrapolation model is solved for different levels of ν_2 (the sensitivity of perceived growth to a shock orthogonal

to consumption). To keep these various settings on a level playing field, I recalibrate risk aversion γ to hold fixed the model-implied average SDF volatility $\mathbb{E}[2L_t(\Delta S_{t+1})]^{1/2}$; since sentiment shocks increase SDF volatility, this means that γ is reduced as ν_2 is increased. What the figure shows is that sentiment shocks raise the contribution of beliefs to SDF volatility, as expected, but they do so by increasing the size of both the permanent and transitory components of the SDF. At least in this class of models, there appears to be a tension: attributing large amounts of SDF volatility to beliefs necessitates a counterfactually transitory SDF.

5 Conclusion

Undoubtedly, asset markets are influenced jointly by risk, risk attitudes, and investors' subjective beliefs. For example, the popular narrative that long-term risk should be particularly important for asset pricing has accumulated a substantial amount of direct evidence.[19] On the other hand, pervasive violations of full-information rational expectations in macro-finance surveys suggests that belief dynamics must also matter.[20] In the theoretical literature, rigorous frameworks alternatively emphasizing risks or beliefs can account for a variety of facts in asset markets.

Against this backdrop, this paper seeks to quantify *how much* of asset market dynamics should be attributable to risk versus beliefs. I find that both risk and beliefs matter, with risk likely to be a more important factor. In particular, combining a volatility bounds approach (Hansen and Jagannathan, 1991; Alvarez and Jermann, 2005) with financial survey data in stocks and bonds, I estimate that at least 50% of SDF volatility must stem from risk. Some estimates for this risk-based fraction exceed 100%. I also show that a growth extrapolation model with preferences for early resolution of uncertainty, a la Nagel and Xu (2022), is remarkably consistent with my estimated volatility ratios. Within that particular model, beliefs are directly responsible for 25% of SDF volatility. That said, the model reveals an interesting possibility: beliefs can matter indi-

[19] For instance, consider the literature that argues news and news shocks matter for asset markets and the macroeconomy (McQueen and Roley, 1993; Francis and Ramey, 2005; Beaudry and Portier, 2006; Barsky and Sims, 2011; Schmitt-Grohé and Uribe, 2012; Kurmann and Otrok, 2013; Barsky et al., 2015; Leduc and Liu, 2016; Nakamura et al., 2017; Basu and Bundick, 2017; Schorfheide et al., 2018; Berger et al., 2020; Liu and Matthies, 2022).

[20] In addition to the belief dynamics uncovered in financial surveys (Chun, 2011; Amromin and Sharpe, 2014; Greenwood and Shleifer, 2014; Piazzesi et al., 2015; Crump et al., 2018; Xu, 2019; Wang, 2021; d'Arienzo, 2020; De La O and Myers, 2021; Nagel and Xu, 2023), similar phenomena have been found in macroeconomic surveys (Coibion and Gorodnichenko, 2015; Gennaioli et al., 2016; Bordalo et al., 2020; Bhandari et al., 2022; Bianchi et al., 2022; Farmer et al., 2023) and experimental settings (Coibion et al., 2018; Afrouzi et al., 2023).

rectly for SDF volatility by creating risk perceptions. Further investigation into this direct versus indirect role of beliefs would help sharpen our understanding bring nuance to the discussion of beliefs versus risks.

References

Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini, "Stock market volatility and learning," *The Journal of Finance*, 2016, *71* (1), 33–82.

_ , Dmitry Matveev, and Stefan Nagel, "Do survey expectations of stock returns reflect risk adjustments?," *Journal of Monetary Economics*, 2021, *117*, 723–740.

Afrouzi, Hassan, Spencer Y Kwon, Augustin Landier, Yueran Ma, and David Thesmar, "Overreaction in expectations: Evidence and theory," *The Quarterly Journal of Economics*, 2023, *138* (3), 1713–1764.

Aıt-Sahalia, Yacine and Andrew W Lo, "Nonparametric risk management and implied risk aversion," *Journal of Econometrics*, 2000, *94* (1-2), 9–51.

Alvarez, Fernando and Urban J Jermann, "Using asset prices to measure the persistence of the marginal utility of wealth," *Econometrica*, 2005, *73* (6), 1977–2016.

Amromin, Gene and Steven A Sharpe, "From the horse's mouth: Economic conditions and investor expectations of risk and return," *Management Science*, 2014, *60* (4), 845–866.

Atmaz, Adem and Suleyman Basak, "Belief dispersion in the stock market," *The Journal of Finance*, 2018, *73* (3), 1225–1279.

Backus, David, Mikhail Chernov, and Ian Martin, "Disasters implied by equity index options," *The Journal of Finance*, 2011, *66* (6), 1969–2012.

_ , _ , and Stanley Zin, "Sources of entropy in representative agent models," *The Journal of Finance*, 2014, *69* (1), 51–99.

Bakshi, Gurdip and Fousseni Chabi-Yo, "Variance bounds on the permanent and transitory components of stochastic discount factors," *Journal of Financial Economics*, 2012, *105* (1), 191–208.

Bansal, Ravi and Amir Yaron, "Risks for the long run: A potential resolution of asset pricing puzzles," *The Journal of Finance*, 2004, *59* (4), 1481–1509.

_ and Bruce N Lehmann, "Growth-optimal portfolio restrictions on asset pricing models," *Macroeconomic dynamics*, 1997, *1* (2), 333–354.

_ and Salim Viswanathan, "No arbitrage and arbitrage pricing: A new approach," *The Journal of Finance*, 1993, *48* (4), 1231–1262.

_ , Dana Kiku, and Amir Yaron, "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices," *Critical Finance Review*, 2012, *1* (1), 183–221.

Barberis, Nicholas and Ming Huang, "Stocks as lotteries: The implications of probability weighting for security prices," *American Economic Review*, 2008, *98* (5), 2066–2100.

_ , Andrei Shleifer, and Robert Vishny, "A model of investor sentiment," *Journal of Financial Economics*, 1998, *49* (3), 307–343.

_ , Robin Greenwood, Lawrence Jin, and Andrei Shleifer, "X-CAPM: An extrapolative capital asset pricing model," *Journal of Financial Economics*, 2015, *115* (1), 1–24.

_ , _ , _ , and _ , "Extrapolation and bubbles," *Journal of Financial Economics*, 2018, *129* (2), 203–227.

Barsky, Robert B and Eric R Sims, "News shocks and business cycles," *Journal of Monetary Economics*, 2011, *58* (3), 273–289.

_ , Susanto Basu, and Keyoung Lee, "Whither news shocks?," *NBER Macroeconomics Annual*, 2015, *29* (1), 225–264.

Basu, Susanto and Brent Bundick, "Uncertainty shocks in a model of effective demand," *Econometrica*, 2017, *85* (3), 937–958.

Beason, Tyler and David Schreindorfer, "Dissecting the equity premium," *Journal of Political Economy*, 2022, *130* (8), 2203–2222.

Beaudry, Paul and Franck Portier, "An exploration into Pigou's theory of cycles," *Journal of Monetary Economics*, 2004, *51* (6), 1183–1216.

_ and _ , "Stock prices, news, and economic fluctuations," *American Economic Review*, 2006, *96* (4), 1293–1307.

Ben-David, Itzhak, John R Graham, and Campbell R Harvey, "Managerial miscalibration," *The Quarterly journal of economics*, 2013, *128* (4), 1547–1584.

Berger, David, Ian Dew-Becker, and Stefano Giglio, "Uncertainty shocks as second-moment news shocks," *The Review of Economic Studies*, 2020, *87* (1), 40–76.

Bhamra, Harjoat S and Raman Uppal, "Asset prices with heterogeneity in preferences and beliefs," *The Review of Financial Studies*, 2014, *27* (2), 519–580.

Bhandari, Anmol, Jaroslav Borovička, and Paul Ho, "Survey data and subjective beliefs in business cycle models," *Unpublished working paper*, 2022.

Bianchi, Francesco, Cosmin L Ilut, and Martin Schneider, "Uncertainty shocks, asset supply and pricing over the business cycle," *The Review of Economic Studies*, 2018, *85* (2), 810–854.

_ , Howard Kung, and Mikhail Tirskikh, "The origins and effects of macroeconomic uncertainty," *Quantitative Economics*, 2023, *14* (3), 855–896.

_ , Sydney C Ludvigson, and Sai Ma, "Belief distortions and macroeconomic fluctuations," *American Economic Review*, 2022, *112* (7), 2269–2315.

Bloom, Nicholas, "The impact of uncertainty shocks," *Econometrica*, 2009, *77* (3), 623–685.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Overreaction and diagnostic expectations in macroeconomics," *Journal of Economic Perspectives*, 2022, *36* (3), 223–244.

_ , _ , Rafael La Porta, and Andrei Shleifer, "Diagnostic expectations and stock returns," *The Journal of Finance*, 2019, *74* (6), 2839–2874.

_ , _ , _ , and _ , "Belief overreaction and stock market puzzles," *Journal of Political Economy*, 2024, *132* (5), 1450–1484.

_ , _ , Rafael La Porta, and Andrei Shleifer, "Finance Without (Exotic) Risk," 2024.

_ , _ , Spencer Yongwook Kwon, and Andrei Shleifer, "Diagnostic bubbles," *Journal of Financial Economics*, 2021, *141* (3), 1060–1077.

_ , _ , Yueran Ma, and Andrei Shleifer, "Overreaction in macroeconomic expectations," *American Economic Review*, 2020, *110* (9), 2748–2782.

Borovička, Jaroslav, Lars Peter Hansen, and José A Scheinkman, "Misspecified recovery," *The Journal of Finance*, 2016, *71* (6), 2493–2544.

Branch, William A and George W Evans, "Learning about risk and return: A simple model of bubbles and crashes," *American Economic Journal: Macroeconomics*, 2011, *3* (3), 159–191.

Cecchetti, Stephen G, Pok sang Lam, and Nelson C Mark, "Asset pricing with distorted beliefs: are equity returns too good to be true?," *American Economic Review*, 2000, *90* (4), 787–805.

Chen, Hui, Winston Wei Dou, and Leonid Kogan, "Measuring "Dark Matter" in Asset Pricing Models," *The Journal of Finance*, 2024, *79* (2), 843–902.

Chen, Long, Zhi Da, and Xinlei Zhao, "What drives stock price movements?," *The Review of Financial Studies*, 2013, *26* (4), 841–876.

Chen, Luyang, Markus Pelger, and Jason Zhu, "Deep learning in asset pricing," *Management Science*, 2024, *70* (2), 714–750.

Chen, Xiaohong, Lars Peter Hansen, and Peter G Hansen, "Robust inference for moment condition models without rational expectations," *Journal of Econometrics*, 2024, p. 105653.

Christiano, Lawrence J, Roberto Motto, and Massimo Rostagno, "Risk shocks," *American Economic Review*, 2014, *104* (1), 27–65.

Chun, Albert Lee, "Expectations, bond yields, and monetary policy," *The Review of Financial Studies*, 2011, *24* (1), 208–247.

Coibion, Olivier and Yuriy Gorodnichenko, "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review*, 2015, *105* (8), 2644–2678.

_ , _ , and Saten Kumar, "How do firms form their expectations? New survey evidence," *American Economic Review*, 2018, *108* (9), 2671–2713.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A Lochstoer, "Asset pricing when 'this time is different'," *The Review of Financial Studies*, 2017, *30* (2), 505–535.

Crump, Richard K, Stefano Eusepi, and Emanuel Moench, "The term structure of expectations and bond yields," 2018. Unpublished working paper.

d'Arienzo, Daniele, "Maturity increasing overreaction and bond market puzzles," 2020. Unpublished working paper.

David, Alexander, "Heterogeneous beliefs, speculation, and the equity premium," *The Journal of Finance*, 2008, *63* (1), 41–83.

Detemple, Jerome and Shashidhar Murthy, "Intertemporal asset pricing with heterogeneous beliefs," *Journal of Economic Theory*, 1994, *62* (2), 294–320.

Ehling, Paul, Alessandro Graniero, and Christian Heyerdahl-Larsen, "Asset prices and portfolio choice with learning from experience," *The Review of Economic Studies*, 2018, *85* (3), 1752–1780.

Epstein, Larry G and Stanley E Zin, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 1989, pp. 937–969.

Fajgelbaum, Pablo D, Edouard Schaal, and Mathieu Taschereau-Dumouchel, "Uncertainty traps," *The Quarterly Journal of Economics*, 2017, *132* (4), 1641–1692.

Farmer, Leland E, Emi Nakamura, and Jón Steinsson, "Learning About the Long Run," *Journal of Political Economy (forthcoming)*, 2023.

Francis, Neville and Valerie A Ramey, "Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited," *Journal of Monetary Economics*, 2005, *52* (8), 1379–1399.

Fuster, Andreas, Benjamin Hebert, and David Laibson, "Natural expectations, macroeconomic dynamics, and asset pricing," *NBER Macroeconomics Annual*, 2012, *26* (1), 1–48.

_ , David Laibson, and Brock Mendel, "Natural expectations and macroeconomic fluctuations," *Journal of Economic Perspectives*, 2010, *24* (4), 67–84.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer, "Expectations and investment," *NBER Macroeconomics Annual*, 2016, *30* (1), 379–431.

Ghosh, Anisha and Guillaume Roussellet, "Identifying beliefs from asset prices," 2023. Unpublished working paper.

_ , Christian Julliard, and Alex P Taylor, "What is the consumption-CAPM missing? An information-theoretic framework for the analysis of asset pricing models," *The Review of Financial Studies*, 2017, *30* (2), 442–504.

Greenwood, Robin and Andrei Shleifer, "Expectations of returns and expected returns," *The Review of Financial Studies*, 2014, *27* (3), 714–746.

Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright, "The US Treasury yield curve: 1961 to the present," *Journal of Monetary Economics*, 2007, *54* (8), 2291–2304.

Hansen, Lars Peter and José A Scheinkman, "Long-term risk: An operator approach," *Econometrica*, 2009, *77* (1), 177–234.

__ **and Ravi Jagannathan**, "Implications of security market data for models of dynamic economies," *Journal of Political Economy*, 1991, *99* (2), 225–262.

__ **and** __ , "Assessing specification errors in stochastic discount factor models," *The Journal of Finance*, 1997, *52* (2), 557–590.

Harrison, J Michael and David M Kreps, "Martingales and arbitrage in multiperiod securities markets," *Journal of Economic Theory*, 1979, *20* (3), 381–408.

Jackwerth, Jens Carsten, "Recovering risk aversion from option prices and realized returns," *The Review of Financial Studies*, 2000, *13* (2), 433–451.

Jin, Lawrence J and Pengfei Sui, "Asset pricing with return extrapolation," *Journal of Financial Economics*, 2022, *145* (2), 273–295.

Kasa, Kenneth, Todd B Walker, and Charles H Whiteman, "Heterogeneous beliefs and tests of present value models," *Review of Economic Studies*, 2014, *81* (3), 1137–1163.

Kazemi, Hossein B, "An intemporal model of asset prices in a Markov economy with a limiting stationary distribution," *The Review of Financial Studies*, 1992, *5* (1), 85–104.

Kim, Don H and Athanasios Orphanides, "Term structure estimation with survey data on interest rate forecasts," *Journal of Financial and Quantitative Analysis*, 2012, *47* (1), 241–272.

Kurmann, André and Christopher Otrok, "News shocks and the slope of the term structure of interest rates," *American Economic Review*, 2013, *103* (6), 2612–2632.

Lansing, Kevin J, "Lock-in of extrapolative expectations in an asset pricing model," *Macroeconomic Dynamics*, 2006, *10* (3), 317–348.

Leduc, Sylvain and Zheng Liu, "Uncertainty shocks are aggregate demand shocks," *Journal of Monetary Economics*, 2016, *82*, 20–35.

Liu, Yukun and Ben Matthies, "Long-Run Risk: Is It There?," *The Journal of Finance*, 2022, *77* (3), 1587–1633.

Malmendier, Ulrike, Demian Pouzo, and Victoria Vanasco, "Investor experiences and financial market dynamics," *Journal of Financial Economics*, 2020, *136* (3), 597–622.

Manela, Asaf and Alan Moreira, "News implied volatility and disaster concerns," *Journal of Financial Economics*, 2017, *123* (1), 137–162.

Martin, Ian, "What is the Expected Return on the Market?," *The Quarterly Journal of Economics*, 2017, *132* (1), 367–433.

__ **and Steve Ross**, "The long bond," 2013. London School of Economics and MIT Sloan Mimeo.

McQueen, Grant and V Vance Roley, "Stock prices, news, and business conditions," *The Review of Financial Studies*, 1993, *6* (3), 683–707.

Moreira, Alan and Tyler Muir, "Volatility-managed portfolios," *The Journal of Finance*, 2017, *72* (4), 1611–1644.

Nagel, Stefan and Zhengyang Xu, "Asset pricing with fading memory," *The Review of Financial Studies*, 2022, *35* (5), 2190–2245.

_ and _, "Dynamics of subjective risk premia," *Journal of Financial Economics*, 2023, *150* (2), 103713.

Nakamura, Emi, Dmitriy Sergeyev, and Jón Steinsson, "Growth-rate and uncertainty shocks in consumption: Cross-country evidence," *American Economic Journal: Macroeconomics*, 2017, *9* (1), 1–39.

Newey, Whitney K and Kenneth D West, "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica: Journal of the Econometric Society*, 1987, pp. 703–708.

O, Ricardo De La and Sean Myers, "Subjective cash flow and discount rate expectations," *The Journal of Finance*, 2021, *76* (3), 1339–1387.

_ and _, "Which Subjective Expectations Explain Asset Prices?," *The Review of Financial Studies*, 2024, *37* (6), 1929–1978.

_, Xiao Han, and Sean Myers, "The Return of Return Dominance: Decomposing the Cross-section of Prices," 2023.

Piazzesi, Monika, Juliana Salomao, and Martin Schneider, "Trend and Cycle in Bond Premia," 2015. Unpublished working paper.

Qin, Likuan and Vadim Linetsky, "Positive eigenfunctions of Markovian pricing operators: Hansen-Scheinkman factorization, Ross recovery, and long-term pricing," *Operations Research*, 2016, *64* (1), 99–117.

_ and _, "Long-term risk: A martingale approach," *Econometrica*, 2017, *85* (1), 299–312.

Roll, Richard, "Evidence on the "growth-optimum" model," *The Journal of Finance*, 1973, *28* (3), 551–566.

Rosenberg, Joshua V and Robert F Engle, "Empirical pricing kernels," *Journal of Financial Economics*, 2002, *64* (3), 341–372.

Ross, Steve, "The recovery theorem," *The Journal of Finance*, 2015, *70* (2), 615–648.

Scheinkman, Jose A and Wei Xiong, "Overconfidence and speculative bubbles," *Journal of Political Economy*, 2003, *111* (6), 1183–1220.

Schmitt-Grohé, Stephanie and Martin Uribe, "What's news in business cycles," *Econometrica*, 2012, *80* (6), 2733–2764.

Schorfheide, Frank, Dongho Song, and Amir Yaron, "Identifying long-run risks: A Bayesian mixed-frequency approach," *Econometrica*, 2018, *86* (2), 617–654.

Schraeder, Stefanie, "Information processing and non-Bayesian learning in financial markets," *Review of Finance*, 2016, *20* (2), 823–853.

Tella, Sebastian Di, "Uncertainty shocks and balance sheet recessions," *Journal of Political Economy*, 2017, *125* (6), 2038–2081.

Wang, Chen, "Under-and overreaction in yield curve expectations," 2021. Unpublished working paper.

Xu, Zhengyang, "Expectation Formation in the Treasury Bond Market," 2019. Unpublished working paper.

Zapatero, Fernando, "Effects of financial innovations on market volatility when beliefs are heterogeneous," *Journal of Economic Dynamics and Control*, 1998, *22* (4), 597–626.

Appendix

A Proofs

A.1 Proof of Lemma 1

We prove only the subjective versions; the objective bounds follow from those by taking $B \equiv 1$ (no belief distortion). Note by the pricing equation (4), the permanent-transitory factorization (7), the fact $\tilde{G} = G$, and the long bond result (10) that

$$\begin{aligned} 0 \leq \tilde{L}_t\Big(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}R_{t+1}\Big) &= -\tilde{\mathbb{E}}_t \log\Big(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}\Big) - \tilde{\mathbb{E}}_t \log R_{t+1} \\ &= -\tilde{\mathbb{E}}_t \log\Big(\frac{G_{t+1}}{G_t}\Big) - \tilde{\mathbb{E}}_t \log\Big(\frac{\tilde{H}_{t+1}}{\tilde{H}_t}\Big) - \tilde{\mathbb{E}}_t \log R_{t+1} \\ &= \tilde{\mathbb{E}}_t \log R^{\infty}_{t+1} + \tilde{L}_t\Big(\frac{\tilde{H}_{t+1}}{\tilde{H}_t}\Big) - \tilde{\mathbb{E}}_t \log R_{t+1} \end{aligned}$$

Rearranging leads to (16). To derive (14), start with the identity

$$\tilde{L}_t\Big(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}\Big) = \tilde{L}_t\Big(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}R^f_t\Big) = \tilde{\mathbb{E}}_t[\log R^{\infty}_{t+1}] - \log R^f_t + \tilde{L}_t\Big(\frac{\tilde{H}_{t+1}}{\tilde{H}_t}\Big) \tag{40}$$

and then use (16).

A.2 Proof of Proposition 1

To prove (20), use the UEB property to get $L(\frac{\tilde{H}_{t+1}}{\tilde{H}_t}) \geq \mathbb{E}[\tilde{L}_t(\frac{\tilde{H}_{t+1}}{\tilde{H}_t})]$. Combining this result with the bound in (16), we obtain (20).

The same argument with $\frac{\tilde{S}_{t+1}}{\tilde{S}_t}R^f_t$ in place of $\frac{\tilde{H}_{t+1}}{\tilde{H}_t}$ leads to $L(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}R^f_t) \geq \mathbb{E}[\tilde{L}_t(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}R^f_t)]$. Because R^f_t is conditionally risk-free, we have $\tilde{L}_t(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}R^f_t) = \tilde{L}_t(\frac{\tilde{S}_{t+1}}{\tilde{S}_t})$. Combining these results with the bound in (14), we obtain (19).

Finally, we can deduce (17)-(18) from the subjective versions by considering $B \equiv 1$ (i.e., no belief distortion), so that all objects and expectations with tildes become objective versions, and then applying the law of iterated expectations.

A.3 Proof of Proposition 2

First, make note of the following key result: the inequalities in (13) and (15) become equalities with $R_{t+1} = R^*_{t+1}$. Similarly, the subjective versions in (14) and (16) become equalities with $R_{t+1} = \tilde{R}^*_{t+1}$.

Now, let us characterize the "wedges" between the bounds in Proposition 1: For the objective bounds, we obtain

$$\begin{aligned} \omega(R) &:= L\Big(\frac{S_{t+1}}{S_t}R^f_t\Big) - \mathbb{E}[\log R_{t+1} - \log R^f_t] \\ &= \mathbb{E}[\log R^*_{t+1} - \log R^f_t] - \mathbb{E}[\log R_{t+1} - \log R^f_t] = \mathbb{E}[\delta_t(R)] \end{aligned}$$

For the subjective bounds, we also use the UEB assumption to obtain

$$\begin{aligned}\tilde{\omega}(R) &:= L\Big(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}R_t^f\Big) - \mathbb{E}\big[\tilde{\mathbb{E}}_t[\log R_{t+1} - \log R_t^f]\big] \\ &\geq \mathbb{E}\Big[\tilde{L}_t\Big(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}R_t^f\Big)\Big] - \mathbb{E}\big[\tilde{\mathbb{E}}_t[\log R_{t+1} - \log R_t^f]\big] \\ &= \mathbb{E}\big[\tilde{\mathbb{E}}_t[\log R_{t+1}^* - \log R_t^f]\big] - \mathbb{E}\big[\tilde{\mathbb{E}}_t[\log R_{t+1} - \log R_t^f]\big] = \mathbb{E}[\tilde{\delta}_t(R)]\end{aligned}$$

Using the assumption that R is closer to objective growth-optimal than subjective growth-optimal, i.e., $\mathbb{E}[\delta_t(R)] \leq \mathbb{E}[\tilde{\delta}_t(R)]$, we thus obtain

$$\omega(R) \leq \tilde{\omega}(R)$$

Therefore, compute

$$\frac{L(\frac{\tilde{S}_{t+1}}{\tilde{S}_t}R_t^f)}{L(\frac{S_{t+1}}{S_t}R_t^f)} = \frac{\mathbb{E}\big[\tilde{\mathbb{E}}_t[\log R_{t+1} - \log R_t^f]\big] + \tilde{\omega}(R)}{\mathbb{E}[\log R_{t+1} - \log R_t^f] + \omega(R)} \geq \frac{\mathbb{E}\big[\tilde{\mathbb{E}}_t[\log R_{t+1} - \log R_t^f]\big] + \omega(R)}{\mathbb{E}[\log R_{t+1} - \log R_t^f] + \omega(R)}$$

Note that this ratio is increasing (decreasing) in $\omega(R)$ if it is smaller (larger) than one. This leads to the result in (21). To obtain (22), we repeat a very similar argument to obtain

$$\frac{L(\frac{\tilde{H}_{t+1}}{\tilde{H}_t})}{L(\frac{H_{t+1}}{H_t})} \geq \frac{\mathbb{E}\big[\tilde{\mathbb{E}}_t[\log R_{t+1} - \log R_t^\infty]\big] + \omega_H(R)}{\mathbb{E}[\log R_{t+1} - \log R_t^\infty] + \omega_H(R)},$$

where $\omega_H(R) := L\Big(\frac{H_{t+1}}{H_t}\Big) - \mathbb{E}[\log R_{t+1} - \log R_t^\infty]$ is the corresponding wedge for the permanent component.

A.4 Proof of Corollary 1

See Alvarez and Jermann (2005), Proposition 2.

A.5 Proof of Corollary 2

Throughout, suppose R is a return such that $\tilde{\mathbb{E}}_t \log R_{t+1} > \log R_t^f$. From (40), we have

$$\frac{\tilde{L}_t(\frac{\tilde{H}_{t+1}}{\tilde{H}_t})}{\tilde{L}_t(\frac{\tilde{S}_{t+1}}{\tilde{S}_t})} = \frac{\tilde{L}_t(\frac{\tilde{H}_{t+1}}{\tilde{H}_t})}{\tilde{\mathbb{E}}_t \log R_{t+1}^\infty - \log R_t^f + \tilde{L}_t(\frac{\tilde{H}_{t+1}}{\tilde{H}_t})}$$

Note that this ratio is increasing (decreasing) in $\tilde{L}_t(\frac{\tilde{H}_{t+1}}{\tilde{H}_t})$ if it is smaller (larger) than one. So if $\tilde{\mathbb{E}}_t \log R_{t+1}^\infty > \log R_t^f$, we use (16) to get

$$1 \geq \frac{\tilde{L}_t(\frac{\tilde{H}_{t+1}}{\tilde{H}_t})}{\tilde{L}_t(\frac{\tilde{S}_{t+1}}{\tilde{S}_t})} \geq \frac{\tilde{\mathbb{E}}_t[\log R_{t+1} - \log R_{t+1}^\infty]}{\tilde{\mathbb{E}}_t \log R_{t+1} - \log R_t^f}$$

On the other hand, if $\tilde{\mathbb{E}}_t \log R_{t+1}^\infty < \log R_t^f$, we use (16) to get

$$1 \leq \frac{\tilde{L}_t(\frac{\tilde{H}_{t+1}}{\tilde{H}_t})}{\tilde{L}_t(\frac{\tilde{S}_{t+1}}{\tilde{S}_t})} \leq \frac{\tilde{\mathbb{E}}_t[\log R_{t+1} - \log R_{t+1}^\infty]}{\tilde{\mathbb{E}}_t \log R_{t+1} - \log R_t^f}$$

Combine these results with the non-negativity of entropies to obtain the result.

A.6 CEI holds in continuous-time Brownian environments

The Conditional Entropy Identity (CEI) assumption says that all higher moments beyond the mean are perceived equally by the subjective and objective probability. In this section, I show, as claimed by Footnote 6, that the CEI automatically holds in continuous-time Brownian environments. (Hence, the weaker assumption of UEB also holds.) I model the dynamics of a variable $\log X_t$ via an Itô process and prove in such case that

$$\lim_{\Delta \to 0} \Delta^{-1} \tilde{L}_t(X_{t+\Delta}) = \lim_{\Delta \to 0} \Delta^{-1} L_t(X_{t+\Delta}).$$

i.e., local versions of the entropy operator coincide.

Suppose

$$d \log X_t = \alpha_t dt + \nu_t \cdot dW_t,$$

where W is a Brownian motion under $\mathbb{P}$. By Girsanov's theorem, the dynamics under the subjective measure are

$$d \log X_t = \tilde{\alpha}_t dt + \nu_t \cdot d\tilde{W}_t,$$

where $\tilde{W}$ is a Brownian motion under $\tilde{\mathbb{P}}$ and $\tilde{\alpha}_t$ is another adapted process. Compute

$$L_t(X_{t+\Delta}) = \log(\mathbb{E}_t^{\nu}[\exp(\int_t^{t+\Delta} (\alpha_s + \frac{1}{2}|\nu_s|^2)ds)]) - \int_t^{t+\Delta} \alpha_s ds$$
$$\tilde{L}_t(X_{t+\Delta}) = \log(\tilde{\mathbb{E}}_t^{\nu}[\exp(\int_t^{t+\Delta} (\tilde{\alpha}_s + \frac{1}{2}|\nu_s|^2)ds)]) - \int_t^{t+\Delta} \tilde{\alpha}_s ds,$$

where $\mathbb{E}^{\nu}$ and $\tilde{\mathbb{E}}^{\nu}$ are alternative expectation operators. Dividing both metrics by Δ and taking $\Delta \to 0$, and using L'Hôpital's rule, yields

$$\lim_{\Delta \to 0} \Delta^{-1} L_t(X_{t+\Delta}) = \alpha_t + \frac{1}{2}|\nu_t|^2 - \alpha_t = \frac{1}{2}|\nu_t|^2 = \tilde{\alpha}_t + \frac{1}{2}|\nu_t|^2 - \tilde{\alpha}_t = \lim_{\Delta \to 0} \Delta^{-1} \tilde{L}_t(X_{t+\Delta})$$

Thus, the two metrics coincide locally.

B Heterogeneity and aggregation

In reality, survey data do not contain a single measure of beliefs, but rather an entire cross-section. A common practice is to compute summary statistics like the cross-sectional average belief about an asset return. If we use such an average belief, then whose marginal utility are we bounding via Lemma 1?

Suppose there are $i \in \{1, \ldots, N\}$ individuals making return forecasts. All the analysis from the previous sections goes through, i-by-i. In particular, each of their beliefs can be characterized by a $\mathbb{P}$-martingale B_t^i, and each of them have a marginal utility process $\tilde{S}_t^i$. Assuming perfect risk-sharing for the shocks that are spanned by asset payoffs, we can treat $\tilde{S}_t^i$ as the projection of true individual marginal utility onto the return space. In that case, we have the generalization of formula (5) to

$$S_t = \tilde{S}_t^i B_t^i, \quad \forall i. \tag{41}$$

Following the permanent-transitory factorization arguments as above, we then also have the generalization of formula (8),

$$H_t = \tilde{H}_t^i B_t^i, \quad \forall i, \tag{42}$$

where $\tilde{H}^i$ is the permanent component of $\tilde{S}^i$, i.e., $\tilde{H}^i$ is a martingale under $\tilde{\mathbb{P}}^i$. All agents share the same transitory component $\tilde{G}^i = G$.

Averaging beliefs amounts to the following aggregation procedure. Consider B defined by

$$\frac{B_{t+1}}{B_t} := \frac{1}{N} \sum_{i=1}^{N} \frac{B_{t+1}^i}{B_t^i}. \tag{43}$$

Then, M is a $\mathbb{P}$-martingale as in the initial analysis. Computing the cross-sectional average belief is equivalent to using M:

$$\frac{1}{N} \sum_{i=1}^{N} \tilde{\mathbb{E}}_t^i[X_{t+1}] = \frac{1}{N} \sum_{i=1}^{N} \mathbb{E}_t[\frac{B_{t+1}^i}{B_t^i} X_{t+1}] = \mathbb{E}_t[\frac{B_{t+1}}{B_t} X_{t+1}] =: \tilde{\mathbb{E}}_t[X_{t+1}].$$

In that case, to satisfy the original relation (8), we must define

$$\tilde{S}_t := \frac{1}{N} \sum_{i=1}^{N} \left(\frac{B_t^i}{B_t}\right) \tilde{S}_t^i. \tag{44}$$

and

$$\tilde{H}_t := \frac{1}{N} \sum_{i=1}^{N} \left(\frac{B_t^i}{B_t}\right) \tilde{H}_t^i. \tag{45}$$

One can think of $\tilde{S}$ and $\tilde{H}$ as the "aggregated marginal utility" and its permanent component, and equations (44)-(45) shows that these are really belief-weighted-averages of individual variables. Furthermore, it is straightforward to check that $\tilde{H}$ is a $\tilde{\mathbb{P}}$-martingale, where $\tilde{\mathbb{P}}$ is defined via M (i.e., through the averaging procedure). With this definition, the same bounds on $\tilde{S}$ and $\tilde{H}$ hold as the

previous analysis (e.g., Lemma 1 and subsequent results), i.e., we have

$$\tilde{L}_t\Bigg(\underbrace{\frac{1}{N}\sum_{i=1}^{N}\Big(\frac{B^i_{t+1}/B^i_t}{B_{t+1}/B_t}\Big)\frac{\tilde{S}^i_{t+1}}{\tilde{S}^i_t}}_{=\tilde{S}_{t+1}/\tilde{S}_t}\Bigg) \geq \underbrace{\frac{1}{N}\sum_{i=1}^{N}\tilde{\mathbb{E}}^i_t}_{:=\tilde{\mathbb{E}}_t}\big[\log R_{t+1} - \log R^f_t\big] \tag{46}$$

$$\tilde{L}_t\Bigg(\underbrace{\frac{1}{N}\sum_{i=1}^{N}\Big(\frac{B^i_{t+1}/B^i_t}{B_{t+1}/B_t}\Big)\frac{\tilde{H}^i_{t+1}}{\tilde{H}^i_t}}_{=\tilde{H}_{t+1}/\tilde{H}_t}\Bigg) \geq \underbrace{\frac{1}{N}\sum_{i=1}^{N}\tilde{\mathbb{E}}^i_t}_{:=\tilde{\mathbb{E}}_t}\big[\log R_{t+1} - \log R^\infty_{t+1}\big] \tag{47}$$

Bounds (46)-(47) show that we may average the survey expected returns, but we must interpret this average belief as a lower bound on the volatility of *belief-weighted-averages* of individual-level variables like $\tilde{S}^i$ and $\tilde{H}^i$. Interpreting the bounds introduces some nuances in a heterogeneous-agent world.[21]

[21] An alternative to averaging is to simply use the entire cross-section of the survey. In particular, the following individual-specific bounds hold:

$$\tilde{L}^i_t\Big(\frac{\tilde{S}^i_{t+1}}{\tilde{S}^i_t}\Big) \geq \tilde{\mathbb{E}}^i_t\big[\log R_{t+1} - \log R^f_t\big] \tag{48}$$

$$\tilde{L}^i_t\Big(\frac{\tilde{H}^i_{t+1}}{\tilde{H}^i_t}\Big) \geq \tilde{\mathbb{E}}^i_t\big[\log R_{t+1} - \log R^\infty_{t+1}\big], \tag{49}$$

where $\tilde{L}^i_t$ is defined analogously to $\tilde{L}_t$. One can measure the right-hand-side via an individual-specific survey response. If an analogous version of CEI holds for $\tilde{S}^i_{t+1}/\tilde{S}^i_t$ and $\tilde{H}^i_{t+1}/\tilde{H}^i_t$, then these bounds in conjunction with (13) and (15) can be used to compare the marginal utility and belief distortion volatilities *for a cross-section of investors*.

C More details on data

C.1 Yield data

Figure 6 displays the very tight relationship between the three data sources I use for par yields. The actual yield series I use takes the FRED yields, then fills any missing values with GSW yields, then fills the remaining missing values with the CRSP yields.

Figure 7 displays par and zero-coupon bond (ZCB) yields. Some results (Table 2) use separately par and ZCB yields as proxies for long-term bond expected returns, because par yields are available for a longer history (and importantly contain more of the period before 1980 when yields were not secularly declining).

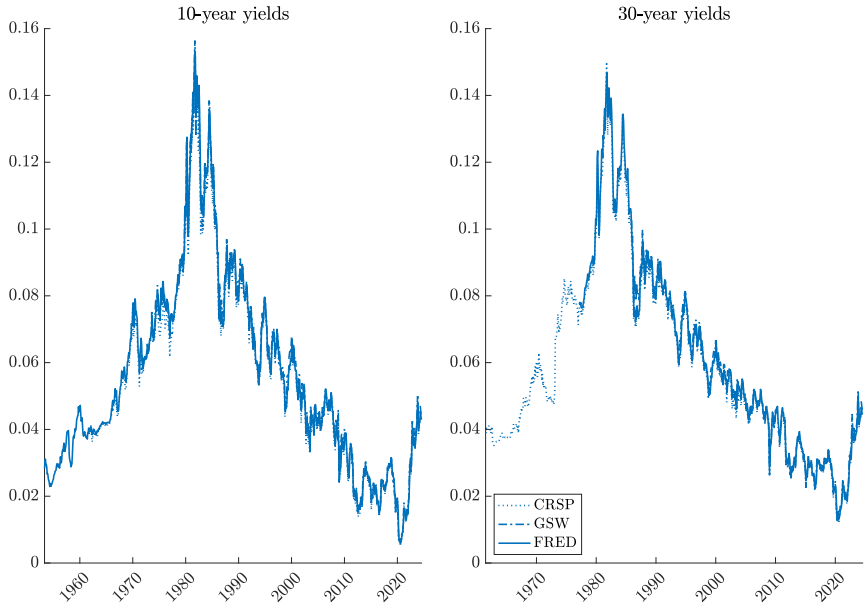


Figure 6: Actual par yields from the three data sources used in the paper. Yields are annualized and in continuously compounded units (the par yields are converted to cc from coupon-equivalent units).

C.2 Bond survey interpolation and bootstrapping

To transform the BCFF survey data (which contains one-year forecasts of 6-month, 1-year, 2-year, 5-year, 10-year, and 30-year par yields) into a full zero-coupon yield curve forecast, I (i) interpolate the par forecasts and then (ii) "bootstrap" a zero-coupon forecast curve.

The interpolation step is a combination of linear interpolation (for maturities ≤ 5 years) and a fitted Nelson-Siegel model (for maturities > 5 years). The Nelson-Siegel model says that the cc par yields of maturity n should take the form

$$y(n) = \beta_0 + \beta_1 \frac{1-\exp(-n\theta)}{n\theta} + \beta_2 \Big[\frac{1-\exp(-n\theta)}{n\theta} - \exp(-n\theta) \Big]. \tag{50}$$

The parameters $(\beta_0, \beta_1, \beta_2, \theta)$ in equation (50) are estimated at every date (thus time-varying) to fit the 2-year, 5-year, 10-year, and 30-year par yield forecasts. I estimate this model by computing the fitting three model-implied slopes: between $n = 5$ and $n = 2$; between $n = 10$ and $n = 5$; and between $n = 30$ and $n = 10$. I compute these three slopes both model and survey, after first

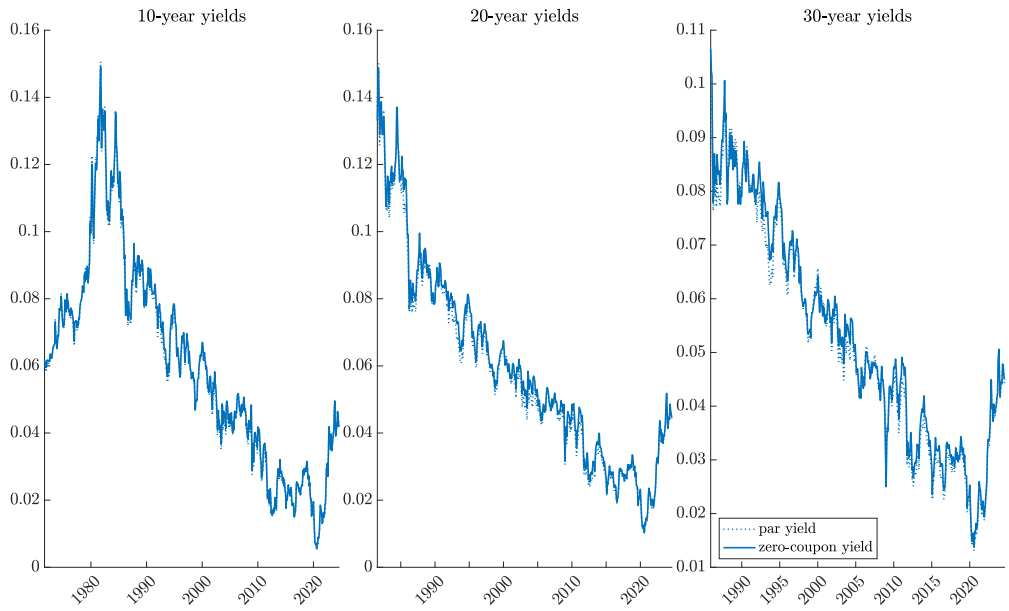


Figure 7: Actual par and zero-coupon yields from GSW. Yields are annualized and in continuously compounded units (the par yields are converted to cc from coupon-equivalent units).

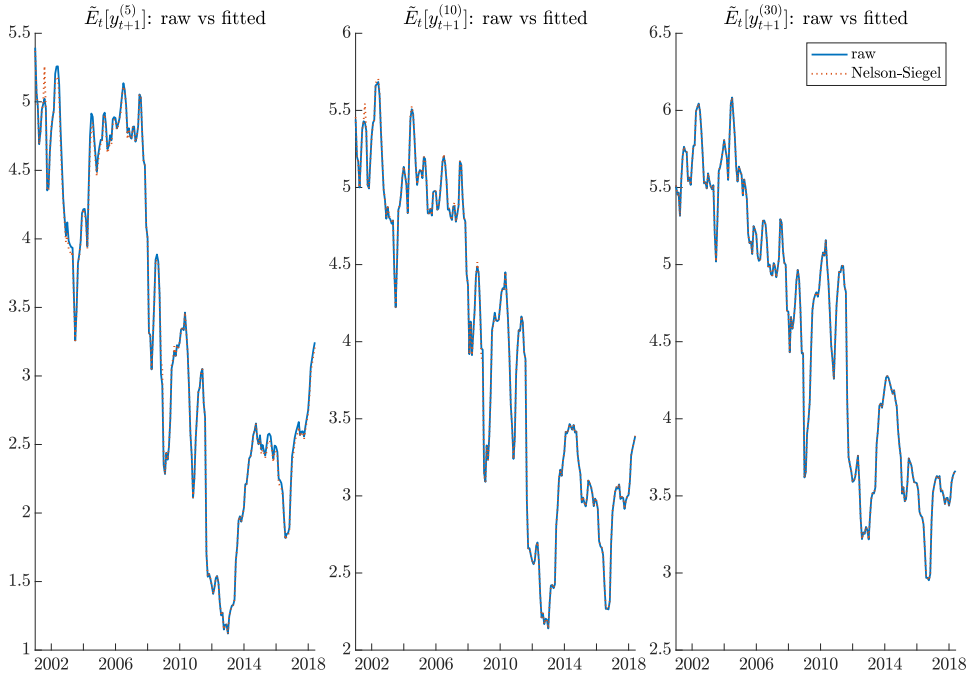


Figure 8: BCFF consensus expectations of one-year ahead par yields ("raw", solid lines) as well as values from a Nelson-Siegel model fitted to the one-year ahead forecasts (dotted lines). All yield expectations are converted from coupon-equivalent to continuously compounded units.

converting the survey yield forecasts to cc units. Since the slopes difference out β_0, this procedure estimates $(\beta_1, \beta_2, \theta)$ to fit these three slopes. I weight the three slopes by the number of maturity years enveloped in the two maturities (i.e., weights of 3, 5, and 20). I then finally pick β_0 to match exactly the 30-year forecast level, since the long-maturity bonds are the most important in my analysis. Figure 8 displays the resulting Nelson-Siegel fit to the key par yield forecasts. Figure 9

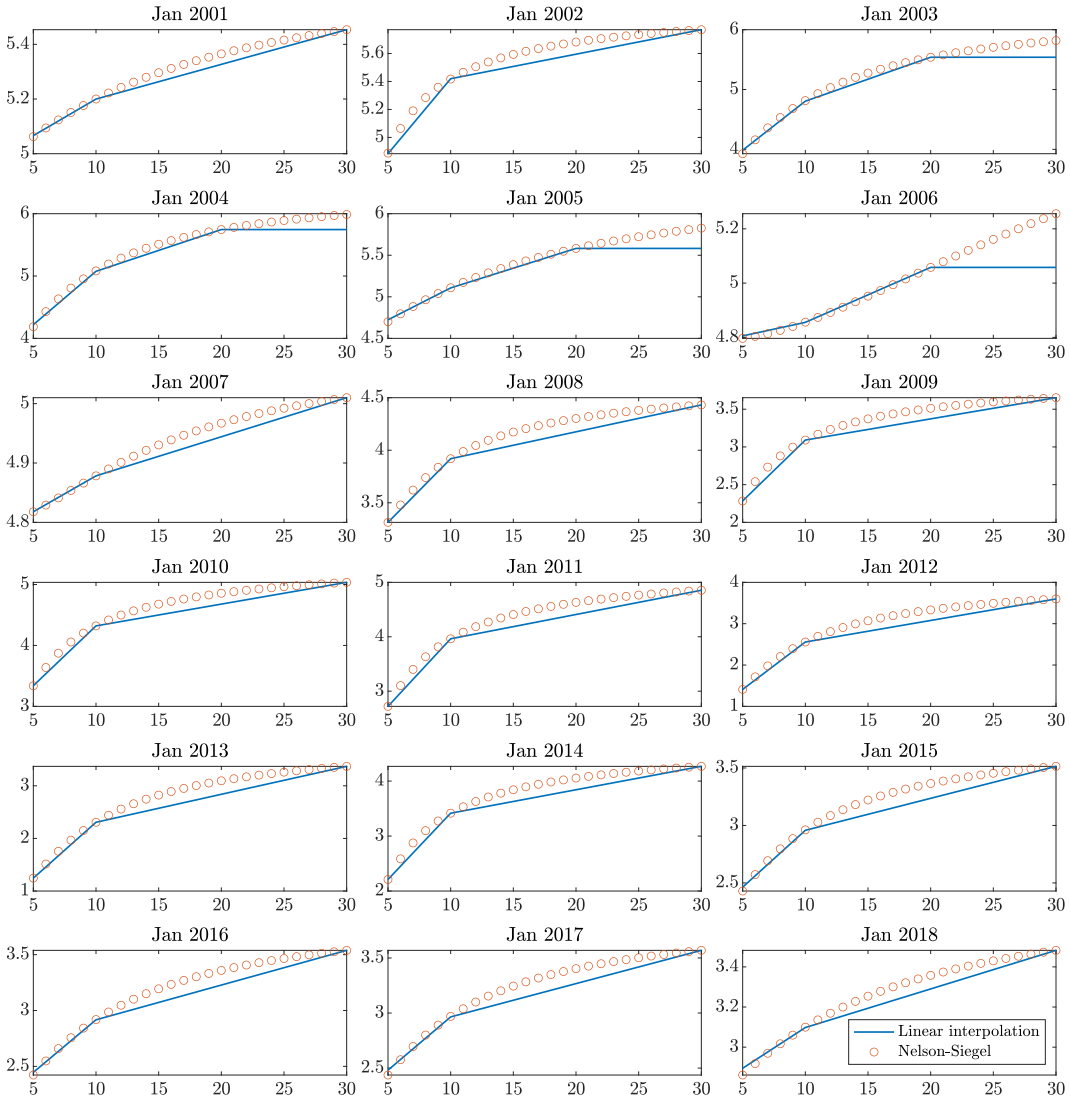


Figure 9: BCFF consensus expectations of the one-year par yield curve (constructed via linear interpolation of different maturities versus interpolation via fitting a Nelson-Siegel model). All yield expectations are converted from coupon-equivalent to continuously compounded units.

contrasts the Nelson-Siegel fit to a linear interpolation in January of each year: the key distinction is the introduction of concavity at the long end, which is typically a feature of the yield curve.

Next, I bootstrap the zero-coupon forecasts as follows. A par yield Y_N^p (in coupon-equivalent units) is the yield such that an N-year bond paying that yield as its semi-annual coupon would trade at par, i.e.,

$$1 = \frac{Y_N^p}{2} \sum_{n=1}^{2N-1} \exp\Big(-\frac{n}{2} y_{n/2}\Big) + \Big(1 + \frac{Y_N^p}{2}\Big) \exp\Big(-N y_N\Big), \tag{51}$$

where y_j is the cc zero-coupon yield for maturity j. Thus, one can think of these par yields as the relevant discount rate for a coupon bond. One can solve for the implied N-period zero-coupon yield

as

$$y_N = -\frac{1}{N}\log\left[\frac{1-\frac{Y_N^p}{2}\sum_{n=1}^{2N-1}\exp\left(-\frac{n}{2}y_{n/2}\right)}{1+\frac{Y_N^p}{2}}\right] \tag{52}$$

I recursively obtain $(y_{n/2})_{n=1}^{60}$ in this way. The only nuance is that I treat Y_N^p as the par yield *forecast* and obtain the zero-coupon yield *forecast* y_N, which effectively ignores the nonlinearities in this transformation.

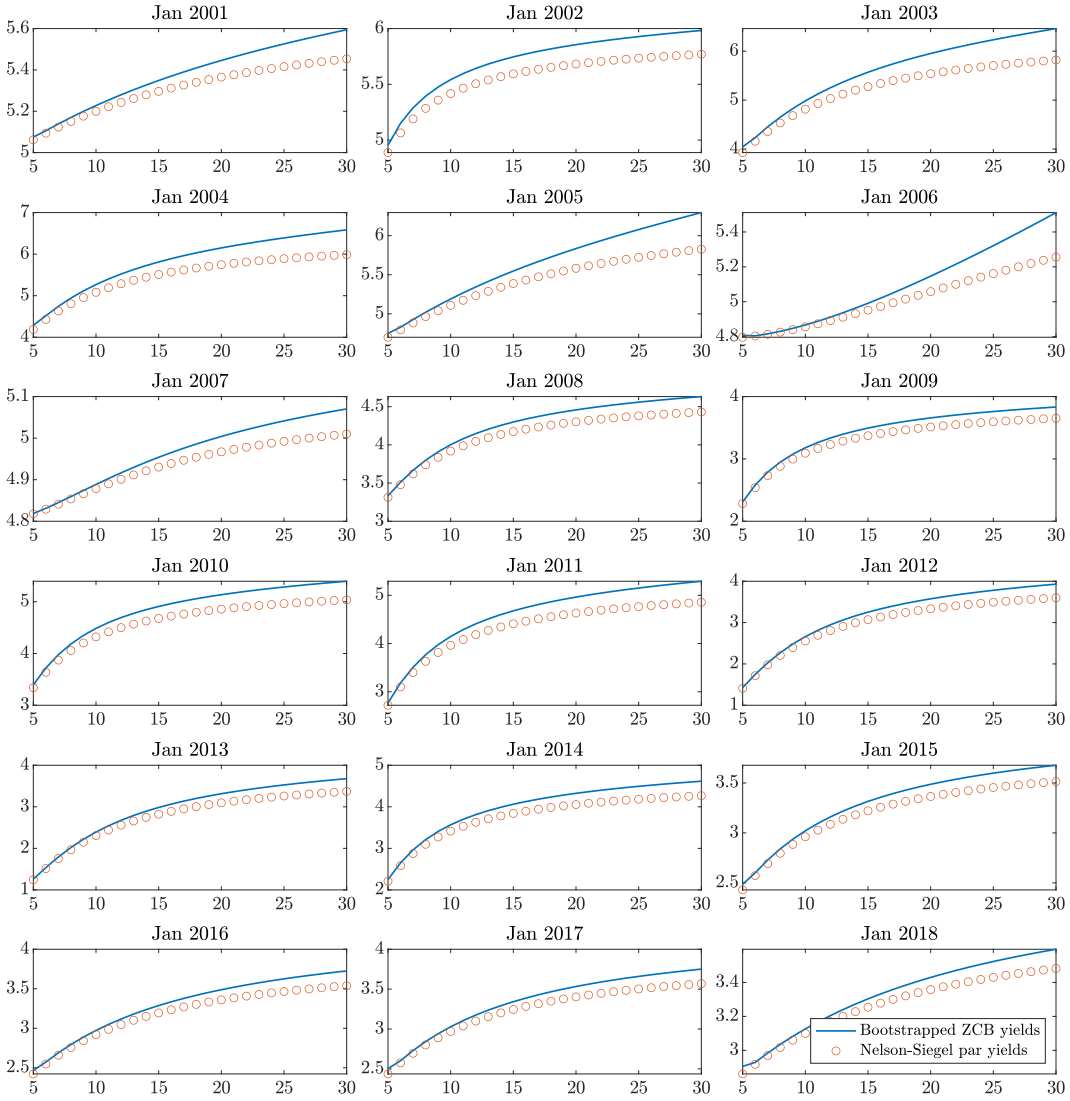


Figure 10: BCFF consensus expectations of the one-year ahead par yield curve (constructed via fitting a Nelson-Siegel model) and the one-year ahead zero-coupon yield curve (constructed via boostrapping from the Nelson-Siegel par yield curve expectation). All yield expectations are converted from coupon-equivalent to continuously compounded units.

Figure 10 shows yield curve forecasts (par vs zero-coupon) at the same January dates as above. Notice that the zero-coupon yield forecasts are higher than the par forecasts, mainly because zero-coupon bonds have higher durations due to their absence of coupons. Finally, Figure 11 displays yield forecast time series and shows how the various transformations affect these forecasts. The left

panel shows linearly interpolated par forecasts. The middle panel shows the Nelson-Siegel fitted par forecasts. The right panel shows the bootstrapped zero-coupon yield forecasts.

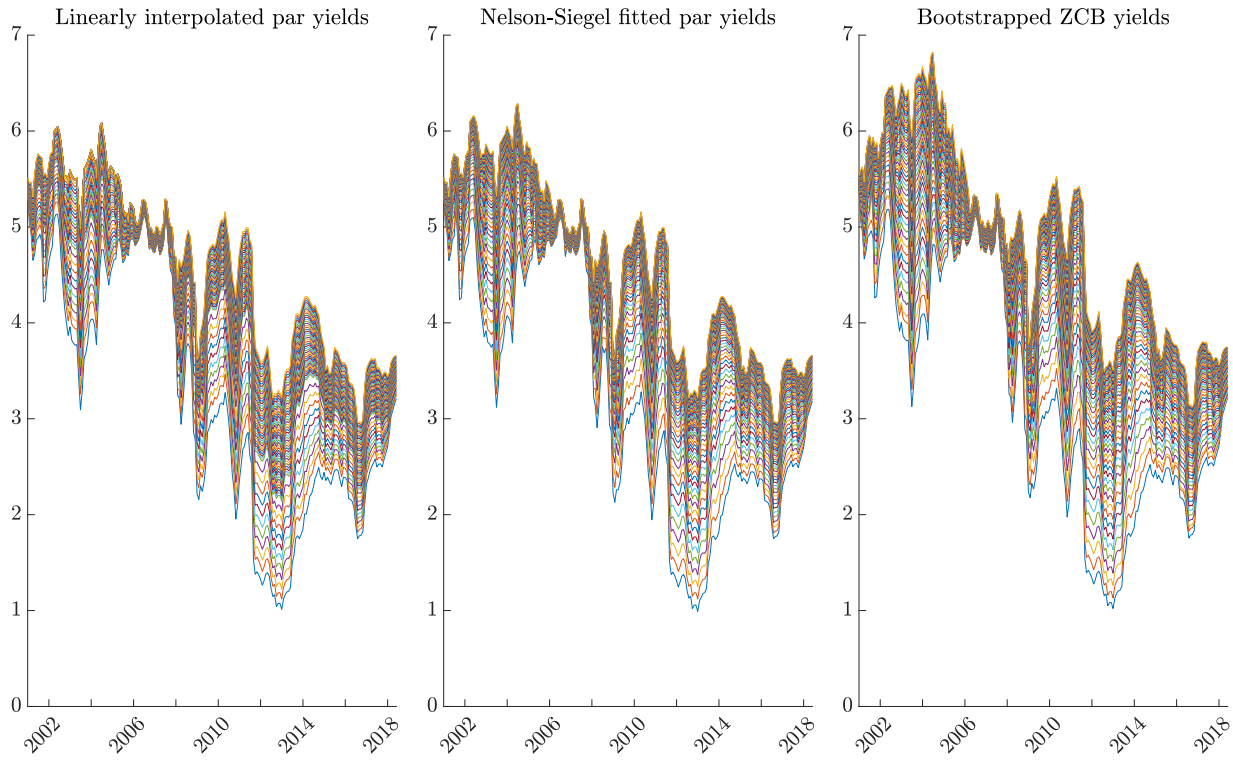


Figure 11: BCFF consensus expectations of the one-year ahead yield curve. The first panel displays expectations of the par yield curve, constructed by linearly interpolating between maturities. The second panel displays expectations of the par yield curve, constructed by fitting Nelson-Siegel model to the expectations. The third panel displays expectations of the zero-coupon yield curve, constructed by bootstrapping from the Nelson-Siegel par curve. All yield expectations are converted from coupon-equivalent to continuously compounded units.

C.3 Fitted volatility forecast

To construct subjective expected log returns, I need to perform a variance adjustment to the surveys, which contain expected arithmetic returns. My starting point is the CFO survey, whose participants report a 90th and 10th percentile for the future stock return. From these answers, Nagel and Xu (2022) construct an implied forecast one-year-ahead market return variance. This is plotted in Figure 12 in the dotted blue line. To apply this variance correction further back in time, I also fit a model by regressing CFO forecast variance onto the contemporaneous squared VIX, its one-month lag, and its trailing 12-month average value. In an alternative specification, I also include the contemporaneous squared SVIX. These fitted values are also displayed in Figure 12.

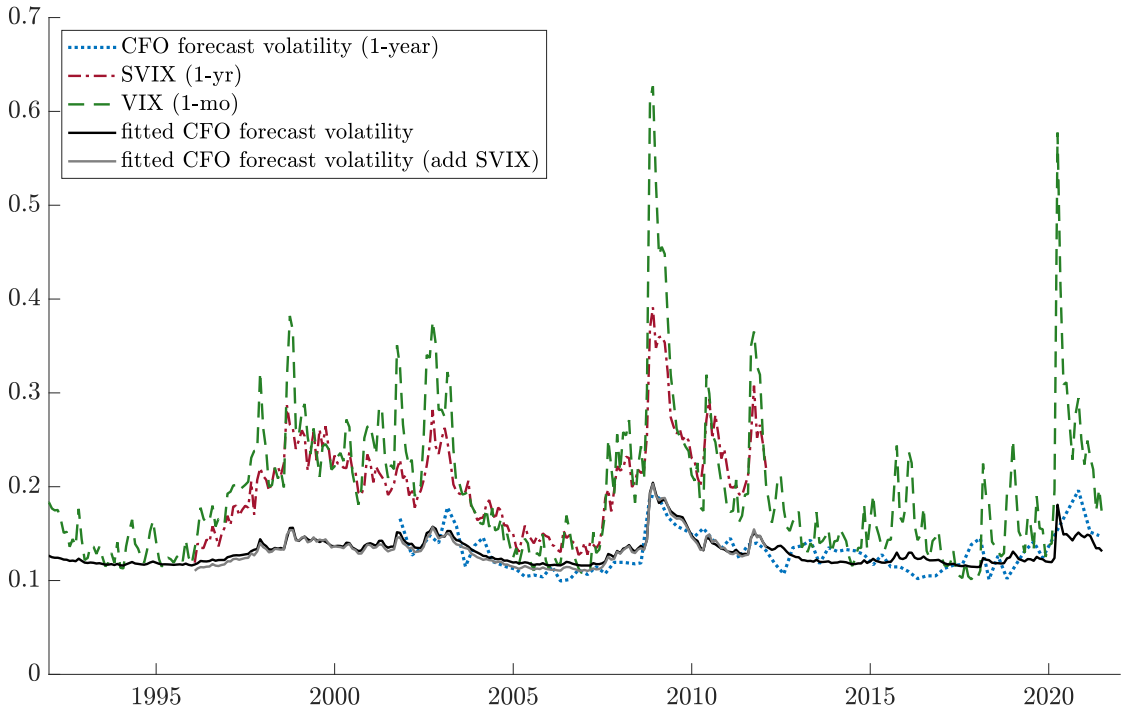


Figure 12: Fitted volatility from a regression of CFO's subjective variance forecasts on the squared VIX (and lags) and squared SVIX.

C.4 Growth-optimal portfolio and objective variance

To construct a proxy for the growth-optimal portfolio, I consider a time-varying mix between the stock market and risk-free rate (i.e., market timing). The growth-optimal problem in this context is

$$\max_{\theta_t} \mathbb{E}_t\left[\log(\theta_t R^m_{t+\Delta} + (1-\theta_t)R^f_t)\right] \tag{53}$$

for some re-investment time interval Δ. (Note also that R^f_t here denotes the risk-free rate between t and $t+\Delta$.) Using the lognormal approximation (or second-order approximation to the log), problem (53) becomes

$$\max_{\theta_t} \mathbb{E}_t\left[\theta_t R^m_{t+\Delta} + (1-\theta_t)R^f_t\right] - \frac{1}{2}\theta_t^2 \mathrm{Var}_t\left[R^m_{t+\Delta}\right], \tag{54}$$

which has the standard mean-variance optimal solution

$$\theta^*_t = \frac{\mathbb{E}_t\left[R^m_{t+\Delta} - R^f_t\right]}{\mathrm{Var}_t\left[R^m_{t+\Delta}\right]} \tag{55}$$

Implementation of (55) requires proxies for the objective conditional expected excess return and objective conditional variance.

I make one of three sets of assumptions.

(i) First, I will use the theory and evidence from Martin (2017) to construct expected stock returns with option-implied variance (SVIX). In particular, I assume (c.f., equation 15 of Martin, 2017)

$$\mathbb{E}_t[R^m_{t+\Delta} - R^f_t] = \Delta R^f_t \mathrm{SVIX}^2_t, \tag{56}$$

where SVIX_t^2 represents the squared SVIX with the same horizon as the period length.

I assume further that risk-neutral and objective conditional variances coincide, analogous to the CEI condition 1, so that (c.f., equation 13 of Martin, 2017)

$$\text{Var}_t[R^m_{t+\Delta}] = \Delta(R^f_t)^2\text{SVIX}^2_t. \tag{57}$$

Under (56)-(57), we have $\theta^*_t = 1/R^f_t \approx 1$, so that the market is approximately growth-optimal as is effectively assumed in Table 4.

(ii) Second, I keep assumption (56) about expected returns but refrain from assuming anything about the gap between risk-neutral and objective variance. Instead, I run a forecasting model for monthly realized variance. Let $v_t := 252 \times \left(\frac{1}{N_m}\sum_{i=1}^{N_m}(R^m_{t-i/252})^2 - (\frac{1}{N_m}\sum_{i=1}^{N_m}R^m_{t-i/252})^2\right)$ denote the trailing month realized variance (annualized) constructed from daily returns (N_m is the number of trading days in the month; I use $N_m = 21$). I propose an AR(1) for monthly variance:

$$v_{t+1/12} = \nu + \varrho v_t + \epsilon_{t+1/12}. \tag{58}$$

Estimating this model on the aggregate stock market realized variance from 1926:08–2024:05 delivers $\hat{\nu} = 0.013$, $\hat{\varrho} = 0.532$, and an R-squared of 0.2827.[22] Using the one-month variance forecasts $\widehat{\text{Var}}_t[R^m_{t+1/12}] = \frac{1}{12}(\hat{\nu} + \hat{\varrho}v_t)$ in conjunction with the SVIX estimate for expected returns in (56), I then construct θ^*_t from (55), i.e.,

$$\theta^*_t = \frac{R^f_t\text{SVIX}^2_t}{\hat{\nu} + \hat{\varrho}v_t} \tag{59}$$

When piecing together these monthly investment weights, I form an annual return via

$$R^*_{t+1} = \prod_{i=0}^{11}\left(\theta^*_{t+i/12}R^m_{t+(i+1)/12} + (1-\theta^*_{t+i/12})R^f_{t+i/12}\right) \tag{60}$$

I construct this growth-optimal proxy.

(iii) Third, I keep the variance forecasting exercise from (ii), but I drop the assumption that arithmetic returns are predictable (i.e., assume constant objective expected returns). In that case, equation (59) is modified to

$$\theta^*_t = \frac{\widehat{\mathbb{E}}[R^m - R^f]}{\hat{\nu} + \hat{\varrho}v_t}, \tag{61}$$

where $\widehat{\mathbb{E}}[R^m - R^f]$ denotes the time-series average of monthly market excess returns. Here, the idea is closer to the "volatility-managed portfolios" model of Moreira and Muir (2017), where higher variance forecasts signal an optimal deleveraging from the stock market. I construct this as an alternative growth-optimal proxy.

[22] I have also run several alternative specifications and found trivial gains in forecasting power. I have augmented the regression with more lags of monthly realized variance. I have tried adding last-month squared-VIX as a regressor. And I have run a GARCH(1,1) model on the monthly demeaned returns. At most, the R-squared rises by 0.01, so I stick with the simple AR(1) model (58).