Implications of Return Predictability across Horizons for Asset Pricing Models*

Carlo A. Favero[†]
Bocconi University & IGIER & CEPR

Fulvio Ortu[‡]
Bocconi University & IGIER

Andrea Tamoni[§]
LSE

Haoxi Yang[¶]
Nankai University

This version: March 24, 2017

*We thank for helpful comments and suggestions Geert Bekaert, Mike Chernov, Dimitris Papanikolaou, Nicola Pavoni, Cesare Robotti, Enrique Sentana, and Yuzhao Zhang (FMA discussant), seminar participants at Manchester Business School, UBC, Hebrew University, Collegio Carlo Alberto, Warwick Business School, BI Norwegian Business School, HKUST Business School, REIM-SWUFE, and ICEF - HSE, and conference participants at the 2015 University of York Asset Pricing Workshop, 2013 IFABS in Nottingham, 20th Annual Conference of the Multinational Finance Society and FMA 2013 Annual Meeting.

[†]Deutsche Bank Chair in Quantitative Finance and Asset Pricing, Bocconi University, Department of Finance, Milan, 20136, Italy. e-mail: carlo.favero@unibocconi.it.

[‡]Bocconi University, Department of Finance, Milan, 20136, Italy. e-mail: fulvio.ortu@unibocconi.it.

[§]London School of Economics, Department of Finance, London, WC2A 2AE, UK. e-mail: a.g.tamoni@lse.ac.uk

[¶]Nankai Univeristy, School of Finance, Tianjin, 300350, China, e-mail: haoxi.yang@outlook.com

Implications of Return Predictability across Horizons for Asset Pricing Models

Abstract

We use the evidence on predictability of returns at different horizons to discriminate among competing asset pricing models. Specifically, we employ predictors-based variance bounds, i.e. bounds on the variance of the Stochastic Discount Factors (SDFs) that price a given set of returns conditional on the information contained in a vector of return predictors. We show that return predictability delivers variance bounds that are much tighter than the classical, unconditional Hansen and Jagannathan (1991) bounds. We use the predictors-based bounds to discriminate among three leading classes of asset pricing models: rare disasters, long-run risks and external habit. We find that the rare disasters model of Nakamura, Steinsson, Barro, and Ursúa (2013) is the best performer since it satisfies rather comfortably the predictors-based bounds at all horizons. As for long-run risks, while the classical version of Bansal and Yaron (2004) is the model most challenged by the introduction of conditioning information since it struggles to meet the bounds at all horizons, the more general version of Schorfheide, Song, and Yaron (2016), which accounts for multiple volatility components, satisfies the 1- and 5-year bounds as long as the set of test assets includes only equities and T-Bills. The Campbell and Cochrane (1999) habit model lies somehow in the middle: it performs quite well at our longest 5-year horizon while it struggles at the 1-year horizon. Finally, when the set of test assets is augmented with Treasury Bonds, the only model that is able to satisfy the predictors-based bounds is the rare disasters model.

J.E.L. CLASSIFICATION NUMBERS: G12, E21, E32, E44

Keywords: return predictability, predictors-based bound, asset pricing models

1 Introduction

If there is valuable information for predicting stock and bond prices over time, and the more so the longer the horizon, how can we use this information to discriminate among competing asset pricing models?[1] In this paper we propose an empirical approach that, by exploiting the predictability of asset returns, permits either to tell asset pricing models apart, or to select the common behavior among apparently different models. The strategy uses sequentially two necessary conditions that must be satisfied by any asset pricing model. The first necessary condition says that the predictability of returns should disappear when returns are discounted by the model-implied Stochastic Discount Factors (SDF). The second condition requires the variance of the model-implied SDF to be larger than the variance dictated by the predictors-based variance bounds, i.e. bounds on the variance of those SDF that price a given set of returns conditional on the information contained in a vector of returns predictors.

We examine three leading classes of asset pricing models: external habit formation (Abel (1990), Campbell and Cochrane (1999)), rare disasters (see Rietz (1988), Barro (2006), and Tsai and Wachter (2015) for a recent review of disasters models), and long-run risks (Bansal and Yaron (2004), Bansal, Kiku, and Yaron (2012)). Within each class, we analyze an estimated asset pricing model representative of that class in order to evaluate the effect of estimation uncertainty on the moments of a model-implied SDF. In particular, we analyze the habit-formation model of Campbell and Cochrane (1999) as estimated by Aldrich and Gallant (2011). Within the long-run risks class we analyze the model of Bansal et al. (2012) as estimated by Bansal, Kiku, and Yaron (2016), and an extension of that model suggested and estimated by Schorfheide et al. (2016). Finally, with respect to the rare disasters class, we focus on the specification proposed and estimated by Nakamura et al. (2013) that introduce partial recoveries into the Barro-Rietz disasters model, and allows for disasters to unfold over several years.

The three models under scrutiny match closely both the historical unconditional annual real return on the risk-free bond and the equity market. Moreover, they all incorporate low frequency

[1] "*There is no way to predict the price of stocks and bonds over the next few days or weeks. But it is quite possible to foresee the broad course of these prices over longer periods, such as the next three to five years.*" Press release of the The Royal Swedish Academy of Sciences for the 2013 Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel

components that should make asset pricing puzzles less pronounced at longer horizons. Consistent with these statements, the conclusion drawn from the standard unconditional Hansen and Jagannathan (1991) bounds is not surprising: all models satisfy these unconditional bounds at medium and long horizons. To see if these conclusions are robust to the inclusion of information, we employ as stock market predictors the dividend-price and consumption-wealth ratios, and as bond returns predictors the yield spread and the short-term rate. Our procedure is as follows. We consider three alternative investment horizons of $h = 1, 4, 20$ quarters respectively, and two alternative sets of returns: a first set that includes the equity market plus the returns from rolling a 3-month Treasury Bill over the horizon h, and a second set which adds to the first set the returns from holding Treasury Bonds with constant maturities of 5, 7, 10, 20 and 30 years. We first test the predictability of discounted returns where discounting is based on the SDF of the model. We show that, for all horizons and sets of assets, the (null of) absence of predictability of *model-discounted* returns cannot be rejected for any of the models under scrutiny. Therefore we proceed to compute predictors-based bounds as in Gallant, Hansen, and Tauchen (1990) and Bekaert and Liu (2004).[2]

Three main conclusions emerge from our horse race. First, the rare disasters model is the best performer, since it satisfies the predictors-based bounds across all horizons and for both sets of assets. Second, when looking at the set of assets that includes only equities and T-Bills, both the habit and the classical Bansal et al. (2016) long-run risks models are rejected at the 1-year horizon. However, we are able to tell these two models apart at long horizons: when we focus on the predictors-based bound at the 5-year horizon constructed from equities, it is the habit model that satisfies the bounds, while the long-run risks model fails to meet the restriction imposed by the bounds. Third, within the long-run risks class of models, the novel specification of Schorfheide et al. (2016) with multiple volatility components constitutes a noticeable improvement upon the classical long-run risks model based on a single time-varying volatility. Indeed, the Schorfheide et al. (2016) model now satisfies the 1-year and 5-year predictor-based bounds constructed from equities and T-Bills, although it still fails the bounds when the Treasury Bonds

[2] In particular, we follow the duality-based approach of Bekaert and Liu (2004): this approach is robust to misspecification of the conditional mean and variance in returns, and is as tight as the Gallant et al. (1990) bound when the conditional moments are known.

are added to the test assets.

Importantly, our bounds provide a conservative estimation of the minimum variability required by any valid SDF. Our choice of predictors for stock and bond returns is parsimonious: we select only predictors that follow from accounting identities (see, e.g., the log-linear framework of Campbell and Shiller (1988a) and Lettau and Ludvigson (2001) for the stock market returns, and Campbell and Ammer (1993) for bond returns). Therefore, our list of potential return predictors is not exhaustive and one could expand the set of predictors to make our predictors-based bounds even tighter.

The fact that only the rare disasters model survives the predictors-based bounds constructed from equities, T-Bills, and Treasury Bonds, is consistent with Tsai (2015) who shows that a model where disasters are accompanied by period of high inflation can generate realistic implications both for the aggregate stock market and for the interaction between stock and bond markets. Also, our analysis highlights the importance of disasters that unfold over time but are transitory in nature. In particular, consistent with the international evidence documented by Gourio (2008a) and Nakamura et al. (2013), the predictors-based bounds favor a specification where consumption disasters unfold over a few years, and are followed by substantial recovery over the more classical approach where disasters are instantaneous and permanent (Rietz (1988), Barro (2006)) .

Our results show the importance of understanding how we can employ the information contained in a set of returns predictors. The dynamic asset pricing models under consideration are constructed from a mixture of assumptions about preferences (such as recursive utility or habit persistence) and exposure to fundamental shocks (such as long-run risks or rare disasters). If one were to look only at the unconditional bounds, the conclusion would be that the equity premium puzzle can be resolved as long as sufficient time-nonseparability is incorporated in the preferences. However, by accounting for the information in the predictors our results, while showing that time-nonseparable preferences are not the full story, emphasize the importance of the interplay between preferences and the dynamics of the state variables. At the light of this interplay, the predictors-based bounds are simultaneously able to (a) tell the habit model apart from the classical long-run risk model at the 5-year horizon when only equities and T-Bills are to be priced; (b) detect the common features across these two models, such as the low variance of their SDFs at the 1-year horizon; (c) let the rare disaster model of Nakamura et al. (2013) emerge as the model most difficult to rejected across all horizons and all sets of assets.

Our work is related to Cochrane and Hansen (1992), who have been the first to look systematically at Hansen and Jagannathan (1991) (HJ, henceforth) bounds across horizons to ascertain the relative performance of different asset pricing models. The present paper extends their analysis in several directions. First, we account explicitly for the effect of returns predictability at different horizons. Moreover, we highlight the interplay between preferences, dynamics and horizons in a wider variety of models. From this standpoint, in particular, we extend the comparative horizons analysis of Cochrane and Hansen (1992) to models that explicitly contain a low frequency component, such as the long-run risk and rare disaster models, and investigate if and how different specifications of this components can make asset pricing puzzles less pronounced at longer horizons. Finally, we account for estimation uncertainty in both the bounds and the model-implied SDFs.

Our work is also related to Kirby (1998), who provides an explicit link between linear predictability and the Hansen and Jagannathan (1991) bounds. Whereas Kirby (1998) investigates whether the ability of predictors to forecast a given set of return is correctly priced by *some* rational asset pricing model, in the sense that there exist SDFs that price correctly those dynamic strategies which condition on the predictors, our interest here is different: we want to exploit the informational content of *a given set of predictors* to investigate the potential of *a given asset pricing model* to price *a given set of returns*.[3]

Finally, our work shares some of the intuition of the recent literature which, using a decomposition of the model's dynamics into transient and permanent components, investigates the implications of these components for valuation (see Hansen and Scheinkman (2009) and Borovicka, Hansen, Hendricks, and Scheinkman (2011)). In particular we view the predictors-based bounds as a useful tool for understanding the high- and low-frequency components of such models. Our work is also related to the recent information-theoretic literature that uses

[3] De Roon and Szymanowska (2012) extend the approach by Kirby (1998) by analyzing the cross-sectional differences in return predictability and taking into account market frictions. The effect of cross-sectional information and frictions on the predictors-based bounds are interesting avenue for future research.

entropy bounds to restrict the admissible regions for the SDF and its components (see Bakshi and Chabi-Yo (2012) and Ghosh, Julliard, and Taylor (2011)). In particular our conclusions are in line with Backus, Chernov, and Zin (2014) who show that the entropy of a model should be sufficiently large to account for observed excess returns.

The rest of this paper is organized as follows. Sections 2 discusses the interplay between predictability, bounds and asset pricing models from the standpoint of the predictors-based variance bounds. Section 3 documents the existence of significant predictable variation in stock and bond returns and shows how conditioning information plays an important role in the construction of the bounds at different horizons. We then assess whether various SDF specifications are consistent with the predictors-based bounds. Section 4 addresses the question of which among the asset classes considered in the paper, stocks or bonds, is key to our results. It also investigates whether misspecification of the conditional moments would change the results. Section 5 concludes.

2 Predictability, Variance Bounds, and Asset Pricing

In this section we discuss how to derive implications of returns predictability across horizons for asset pricing models. We first derive predictors-based bounds on the volatility of SDFs. We proceed by discussing how the predictors-based bound can be used to discriminate empirically these asset pricing models as long as the predictability of model-discounted returns is rejected. Finally, we briefly review the main characteristics of the SDFs of the three main asset pricing models we employ in our empirical analysis, i.e. the Long-Run Risks model, the Habit model and the Rare Disasters model.

2.1 Predictability of Returns and Variance Bounds

Suppose that N assets are traded at a given time t, with returns vector R_{t+h}, where h is the investment horizon. We denote by μ the vector of unconditional mean returns, and by Σ their unconditional covariance matrix. We start by briefly reviewing the classical Hansen-Jagannathan (1991, HJ henceforth) bounds which produce testable restrictions based on μ and Σ on the set of SDFs that price the returns R_{t+h} unconditionally. Recall that an SDF that prices the returns R_{t+h} is a random variable m_{t+h} that satisfies the pricing equation

$$E\left(m_{t+h}R_{t+h}\right) = \mathbf{1} \tag{1}$$

where $\mathbf{1}$ is the unit vector. Letting then $E\left(m_{t+h}\right) = \nu$, consider a regression of the demeaned SDF on the demeaned returns, i.e.

$$m_{t+h} - \nu = \beta^T \left(R_{t+h} - \mu\right) + \epsilon_t$$

where the regression error ϵ_t is orthogonal to the regressors $R_{t+h} - \mu$. Exploiting the standard regression formula $\beta = \Sigma^{-1}E\left[\left(R_{t+h} - \mu\right)\left(m_{t+h} - \nu\right)\right]$ together with the pricing equation $E\left(m_{t+h}R_{t+h}\right) = \mathbf{1}$ and $E\left(m_{t+h}\right) = \nu$, one readily obtains

$$m_{t+h} = \nu + \left(\mathbf{1} - \nu\mu\right)^T \Sigma^{-1}\left(R_{t+h} - \mu\right) + \epsilon_t$$

Letting

$$m^{HJ}_{t+h} = \nu + \left(\mathbf{1} - \nu\mu\right)^T \Sigma^{-1}\left(R_{t+h} - \mu\right) \tag{2}$$

and substituting m^{HJ}_{t+h} into the pricing equation (1), simple algebra shows that $E\left(m^{HJ}_{t+h}R_{t+h}\right) = \mathbf{1}$, i.e. m^{HJ}_{t+h} is a valid SDF. More importantly, since $E\left(\epsilon_t^2\right) \geq 0$ then $Var\left(m_{t+h}\right) \geq Var\left(m^{HJ}_{t+h}\right)$. Therefore, m^{HJ}_{t+h} is the SDF with minimum variance among all the SDFs with mean ν. This leads to defining the quadratic form in ν

$$\sigma^2_{HJ}\left(\nu\right) \equiv Var\left(m^{HJ}_{t+h}\right) = \left(\mathbf{1} - \nu\mu\right)^T \Sigma^{-1}\left(\mathbf{1} - \nu\mu\right) \tag{3}$$

which represents the unconditional HJ bound. Given that the unconditional HJ bound depends only on $E\left(m_{t+h}\right)$ and observable moments of the vector of returns, the relationship $\sigma\left(m_{t+h}\right) \geq \sigma\left(m^{HJ}_{t+h}\right)$ defines a portion of the plane in the dimensions $E\left(m_{t+h}\right), \sigma\left(m_{t+h}\right)$, which takes the form of a cup-shaped region.

To illustrate how the region is constructed consider the typical case in which there is no real risk free interest rate and therefore no linear combination of the vector of asset payoffs used by the econometrician is identically equal to one, and consider equity market real returns plus the returns from rolling a three-month Treasury Bill over the holding period. The sample spans 1952Q2- 2012Q3. The unconditional lower bound to volatility in stochastic discount factor is plotted in Figure 1 for the horizons $h = 1, 4, 20$.

[Insert Figure 1 about here]

Next, we consider the case in which there exists a vector Z_t of variables that predict the returns R_{t+h}. More precisely, we let $\mu_t \equiv E\left[R_{t+h}\left|Z_t\right.\right]$, $\Sigma_t \equiv Var\left[R_{t+h}\left|Z_t\right.\right]$ and we assume that $Var(\mu_t) > 0$. We concentrate now our attention on the SDFs that price returns conditionally on the realizations of the predictors Z_t, that is

$$E\left(m_{t+h}R_{t+h}\left|Z_t\right.\right) = \mathbf{1} \tag{4}$$

Importantly, while the law of iterated expectations guarantees that any such SDF also prices the returns R_{t+h} unconditionally, in general an SDF that satisfies the unconditional pricing equation may fail to price correctly the returns when the information in the predictors Z_t is accounted for. In other words, the set of SDFs that price returns conditionally on Z_t is a subset of the larger set of SDFs that satisfy the unconditional pricing equation.

We fix now an SDF m_{t+h} that prices returns conditional on Z_t, we let $\nu_t \equiv E\left[m_{t+h}\left|Z_t\right.\right]$ (so that $E\left(\nu_t\right) = \nu$), and we want to characterize the SDF m^Z_{t+h} with minimum variance among all the SDFs that price returns conditionally and whose mean is ν. The solution to this problem is supplied by Gallant et al. (1990)(GHT henceforth), who show that m^Z_{t+h} is obtained, quite naturally, by replacing in (1) the unconditional moments with the conditional ones, i.e.

$$m^Z_{t+h} = \nu_t + \left(\mathbf{1} - \nu_t\mu_t\right)^T \Sigma_t^{-1}\left(R_{t+h} - \mu_t\right) \tag{5}$$

where μ_t and Σ_t are the vector of conditional expected returns and the conditional variance-covariance matrix, respectively.

The predictors-based bound $\sigma^2_Z(v)$ is then defined as the function that maps the mean of m^Z_{t+h}

into its variance. In fact, to make this relationship more transparent, Bekaert and Liu (2004) work on the expression of $Var\left(m_{t+h}^{Z}\right)$ to show that

$$\begin{aligned} \sigma_Z^2(v) \equiv Var\left(m_{t+h}^{Z}\right) = \; & E\left[\left(\mathbf{1} - \tfrac{\nu-b}{1-d}\mu_t\right)^T \left(\mu_t\mu_t^T + \Sigma_t\right)^{-1} \left(\mathbf{1} - \tfrac{\nu-b}{1-d}\mu_t\right)\right] \\ & - E\left[\left(\mathbf{1} - \tfrac{\nu-b}{1-d}\mu_t\right)^T \left(\mu_t\mu_t^T + \Sigma_t\right)^{-1} \mu_t\right] \end{aligned} \tag{6}$$

where $b = E\left[\mu_t^T \left(\mu_t\mu_t^T + \Sigma_t\right)^{-1} \mathbf{1}\right]$ and $d = E\left[\mu_t^T \left(\mu_t\mu_t^T + \Sigma_t\right)^{-1} \mu_t\right]$. Given v, therefore, $\sigma_Z^2(v)$ is a quadratic form in ν that yields the minimum variance across all SDFs that price the returns conditionally on the information in Z_t and whose unconditional mean is v. While this expression is more cumbersome than (2), yet it is readily seen that the quadratic form $\sigma_Z^2(v)$ lies above $\sigma_{HJ}^2(\nu)$ for any value of v. As mentioned above, in fact, the law of iterated expectation implies that any SDF which prices returns conditionally on Z_t does so unconditionally as well, hence $E\left(m_{t+h}^{Z} R_{t+h}\right) = \mathbf{1}$. Since m_{t+h}^{HJ} is the SDF with minimum variance among all the SDFs with mean ν, it follows immediately that $\sigma_Z^2(v) \equiv Var\left(m_{t+h}^{Z}\right) \geq Var\left(m_{t+h}^{HJ}\right) \equiv \sigma_{HJ}^2(\nu)$, an inequality that in fact holds strictly banning degenerate cases.

To illustrate the effect of predictability on the bounds we have reported in Figure 1 the unconditional bounds along with the conditional bounds based on two predictors, the dividend-price ratio and the consumption-wealth ratio. Interestingly the gap between the unconditional and conditional bounds increases with the holding period as a consequence of the fact that predictability increases with the horizon.

2.2 Predictors-based Bounds and Asset Pricing

We consider now the asset pricing modelling side of our argument. Denote with Ω_t the information set available to a (representative) agent that makes optimal consumption and portfolio decisions given the set of assets with returns R_{t+h}. In a nutshell, an asset pricing model is characterized by a couple $\left(X_t, m_{t+h}^{X}\right)$ where $X_t \subset \Omega_t$ denotes the set of state variables of the model and m_{t+h}^{X} denotes the SDF associated with the model. Ideally, the state variables should represent a sufficient statistic of the information set Ω_t, in the sense that the following Euler condition should hold

$$E\left(m_{t+h}^{X}R_{t+h}\middle|\,\Omega_t\right) = E\left(m_{t+h}^{X}R_{t+h}\middle|\,X_t\right) = \mathbf{1} \tag{7}$$

In words, if the state variables are a sufficient characterization of the information set Ω_t then the model SDF m_{t+h}^{X} should price not only all the managed portfolio that condition on the state variables X_t, but more generally the returns from managed portfolios that condition on all the available information.

Consider now the situation in which a set $Z_t \subset \Omega_t$ of observable predictors of the returns is available to the optimizing agent. The predictors are not necessarily a subset of the state variables: successful predictors, for instance, are derived by forward solution of linearization of returns or by the imposition of transversality conditions on the intertemporal budget constraint (see, e.g. Campbell and Shiller (1988a), Campbell and Mankiw (1989), and Lettau and Ludvigson (2001)). Still, since the predictors Z_t are in the information set Ω_t of the optimizing agent the law of iterated expectation applied to (7) imposes on any model-based SDF m_{t+h}^{X} the necessary condition

$$E\left(m_{t+h}^{X}R_{t+h}\middle|\,Z_t\right) = \mathbf{1} \tag{8}$$

i.e. that the predictability of returns disappears when returns are discounted by the model-implied SDF. Recalling now that the predictors-based bound $\sigma_Z^2(v)$ defined in equation (6) identifies a lower bound on the variance of any SDF that prices returns conditional on Z_t, when (8) cannot be empirically rejected, a further testable implication is supplied by the predictors-based bounds, i.e.

$$Var\left(m_{t+h}^{X}\right) \geq \sigma_Z^2(v)\ . \tag{9}$$

Our empirical strategy is to use the two conditions (8) and (9) to test different asset pricing models. Given an asset pricing model $\left(X_t, m_{t+h}^{X}\right)$ our first step consists in testing if the model-discounted returns are indeed unpredictable. Of course, if a model fails this test, then there is no point in proceeding to test the bound (9), since that model will never be able to price correctly the information in the predictors Z_t. On the other side, suppose that the SDFs of two competing asset pricing models both pass the test of unpredictability of model-discounted returns. The predictors-based bound can then be employed as a second-stage test to tell the two models apart. To exemplify this procedure, consider the two extreme cases in which a first model is based on logarithmic preferences, and the second model is a model with rare disasters a là Nakamura et al. (2013) (which we will discuss more in details later in the paper). Suppose moreover that only the market portfolio R_{t+h}^{MKT} is traded. Since the SDF associated with log preferences is $\left(R_{t+h}^{MKT}\right)^{-1}$, the model-discounted returns would be constantly equal to 1 and hence obviously unpredictable. In Section 3.3.1 we show that the Nakamura et al. (2013) rare disaster model also passes the test of unpredictability of model-discounted returns. However, the second stage test based on the predictors-based bound (9) allows us to tell the two models apart unambiguously. In fact, while as it is well documented in the literature and illustrated in our Figure 1, the log preference model fails even the weaker unconditional HJ bounds (and hence a fortiori the stricter predictors-based bounds) as we show in Section 3.3.2 the rare disasters model of Nakamura et al. (2013) passes comfortably the predictors-based bounds test. In sum, while the two models could not have been discriminated based only on the evidence of unpredictability of model-discounted returns, the predictors-based bounds allow us to reject the log preference model, while the Nakamura et al. (2013) rare disasters model cannot be rejected. Importantly, the rare disasters model is not rejected not simply because it adds volatility to the stochastic discount factor but because the mean *and* volatility of the SDF define a point in the space spanned by the state variables which is within the cup shaped region determined by the predictors-based bounds.

In the empirical part of the paper we use our sequential testing methodology based on equations (8) and (9) to discriminate among three leading asset pricing models. The first model is the Bansal et al. (2016) model of long-run risks, where the state variables are the first two conditional moments of log consumption growth g_t, that is $X_t = (x_t, \sigma_t^2)$, where the information generated in this case comes from the innovations in these two conditional moments, and the SDF takes the form

$$\ln\left(m_{t+h}^{X}\right) = A + Bg_{t+h} + Cr_{a,t+h} \tag{10}$$

where $r_{a,t+h}$ denotes the (continuously compounded) return on an asset that delivers a dividend equal to aggregate consumption, and A, B, C, are functions of the subjective discount factor, risk-aversion coefficient and intertemporal elasticity of substitution of the representative investor. We will also consider a more recent specification of the long-run risks model described in Schorfheide et al. (2016). This specification expands the set of state variables to account for three separate volatility components: one governing dynamics of the persistent cash flow growth component, $\sigma^2_{x,t}$, and the other two controlling temporally independent shocks to consumption and dividend volatility, $\sigma^2_{c,t}$ and $\sigma^2_{d,t}$.

The second model is the external habit model of Campbell and Cochrane (1999), where the state variable is log surplus consumption ratio s_t, so that in this case $X_t = s_t$, the information is now generated by the innovations in s_t, and the SDFs takes the form

$$\ln\left(m^X_{t+h}\right) = A' + B'\left(g_{t+h} + s_{t+h}\right)) \tag{11}$$

with A' and B' functions of the subjective discount factor and the risk aversion coefficient.

The last model is the Rare Disaster model of Nakamura et al. (2013), where the state variables are I_t, the indicator of disaster occurrence at time t, and z_t, the amount by which consumption differs from potential due to current and past disasters. Hence, in this case $X_t = (I_t, z_t)$, and the SDF takes the same functional form as in the long-run risk model.

3 Empirical Results

In this section we put the framework introduced in the previous section at work. The empirical investigation is articulated in three steps.

First, we review the predictability of raw returns within a standard, linear forecasting model (see Section 3.1). Given the evidence of significant predictability, we next compute the predictors-based bound $\sigma^2_Z(\nu)$ for different sets of assets and horizons, and compare them with the unconditional bounds (see Section 3.2). Importantly, we establish that the predictors employed in our linear forecasting model are indeed useful in sharpening the diagnostic efficacy of variance bounds with respect to the unconditional case.

Second, we compare the predictability of *raw* returns with the predictability of *discounted* returns (see Section 3.3.1). In particular we investigate the predictability of returns discounted using, respectively, the long run risk model, the external habit model and the rare disaster model. Interestingly, whereas for raw returns the null of no-predictability is rejected at medium and long-term horizon, when returns are duly discounted by all three models then the null of no-predictability cannot be rejected.

Third, since the three models under scrutiny pass the first-moment test of unpredictable discounted returns, we move to the natural second stage that compares the variability of model-implied SDFs against the predictors-based bounds (see Section 3.3.2). We show that, by fully exploiting the implications of information for the first conditional moment of returns across horizons, the predictors-based bounds provide a useful model-selection tool that permits either to set models apart, or to select the common behavior among apparently different models.

3.1 The predictive model

To conduct our empirical analysis, we consider three horizons and two alternative sets of assets for each horizon. Specifically, we concentrate on horizons of $h = 1, 4, 20$ quarters, and for each horizon we consider alternatively the following two sets of returns:

1. SET A: the equity market returns plus the returns from rolling a three-month Treasury Bill over the holding period h;[4]
2. SET B: the returns from SET A plus the returns from holding over the horizon h Treasury Bonds with constant maturities of 5, 7, 10, 20 and 30 years.

In particular the market return is the (gross) return on the value weighted portfolio of all stocks traded in the NYSE, the AMEX, and NASDAQ. All returns are gross and deflated using the consumer price index (CPI).[5]

[4] In alternative to the returns from rolling over the three-month T-bill, we have also considered the yield-to-maturity on a zero-coupon bond with maturity matching the horizon h. Results were basically unchanged and particularly so at the 1-year horizon. Since in all cases one needs to subtract the inflation realized over the given horizon h from the return on each strategy, for robustness we have also considered the case in which real yields on Treasury Inflation Protected Securities (TIPS) were used instead of nominal yields minus realized inflation. Once again results were unaffected.

[5] For a detailed description of data construction see the Appendix A.

These two sets correspond to a universe of equity and bond portfolios whose return properties are the subject of much scrutiny in the empirical asset pricing research. In particular SET A allows us to examine whether the equity premium puzzle (see Mehra and Prescott (1985)) can be explained once predictability is accounted for, while SET B will be informative about whether the puzzle extends beyond the market index, to include returns from government bonds, i.e. the term premium (see Fama and Bliss (1987) and Cochrane and Piazzesi (2005)).[6]

Table 1 presents full-sample statistics of the quarterly stock returns and 5-year constant maturity bond for the common sample period (1952Q2 to 2012Q4). Over this sample period, the mean nominal return on stocks was 11.45% per annum, the mean nominal return on bonds was 6.31% per annum, and the mean short-term interest rate - not shown in the table - was about 5.65% per annum. The standard deviation of stock nominal returns was 16.68% per annum, and the standard deviation of bond log returns was 5.77% per annum.

[Insert Table 1]

In contrast to the simple random walk view, stock and bond returns do seem predictable, and markedly more so the longer the horizon. To review this claim we use a typical specification that regresses rates of return on lagged predictors. In particular we consider the following linear predictive system:

$$R^i_{t+h} = \beta^i_{0,h} + \beta^i_{1,h} Z^i_t + u^i_{t+h} \tag{12}$$

where $i = S, B$ stands for stocks and bonds, respectively, and $Z_t = \left(Z^S_t, Z^B_t\right)$ denotes the vector of returns' predictors, potentially different for stock and bonds. As mentioned above, the holding period ranges from one quarter to five years, i.e. $h = 1, 4, 20$ quarters. The stock market return predictors are the dividend-price ratio, dp_t, and the consumption-wealth ratio, cay_t, i.e.

$$Z^S_t = [dp_t,\ cay_t]\ .$$

[6]The empirical term-structure literature has typically employed the CRSP unsmoothed Fama-Bliss zero-coupon yields. When the holding period is one year, the correlation between the returns on the 5-year constant-maturity Treasury bond and the Fama-Bliss 5-year zero-coupon is 0.989. We use constant maturity coupon bonds from CRSP instead of the Fama-Bliss data set because we want to study also the implication of long-horizon return predictability on the variance of a SDF. Whereas the Fama-Bliss data set comprises observations on zero-coupon bonds with maturities of 1, 2, 3, 4 and 5 years, the available maturities in the constant-maturity data set ranges from 5 to 30 years.

The bond returns predictors are the yield spread, spr_t, and the short-term US treasury-bill rates, y_t, i.e.

$$Z_t^B = [\text{spr}_t,\ y_t]\ .$$

The choice of dp_t as a stock market predictor is motivated by the present value logic, see Campbell and Shiller (1988a), while the choice of cay_t follows from a linearization of the accumulation equation for aggregate wealth in a representative agent economy, see Campbell and Mankiw (1989) and Lettau and Ludvigson (2001). They are both "noisy" predictors of future asset returns. The bond market predictors are motivated by a voluminous literature documenting the failure of the expectation hypotheses (see Fama and Bliss (1987) and Campbell and Shiller (1991)). Interestingly, the yield spread is the counterpart for the bond market of the dividend yield for stock market. According to the present value model, the dividend yield should either forecast dividend growth or returns, or both. Empirically, the dividend yield *does* forecast excess returns and not dividend growth. Analogously, the yield spread should forecast either short rates or excess bond returns, or both. Empirically, the yield spreads *do* forecast excess returns and not yield changes. Table 2-Panel A presents regressions of the real stock returns R^s_{t+h} onto Z^A_t. Although the $R^2 = 4.7\%$ at quarterly horizon is low, it then rises with the horizon, reaching a value of about 51%, at the 5 years horizon.[7] Each variable has an important impact on forecasting long horizon returns: using the price-dividend ratio as the sole forecasting variable, for instance, would lead to an R^2 of "only" 22% at the 5 years horizon. Table 2-Panel B presents regressions of the real returns from holding a 5-year maturity bond onto Z^B_t. The results show that our predictive system is able to capture fluctuations in bond returns at all horizons. These results are consistent with much of the recent empirical research on the predictability of stock and bond returns (see Campbell (1987), Fama and Bliss (1987), Fama and French (1989), Cochrane (2001; 2008), Lewellen (2004) and Viceira (2012) among others).

[Insert Table 2]

We conclude this section with three observations. First, our predictive model (12) is based on predictors that are not derived from a specific asset pricing model. In the language of

[7] Using excess returns yields similar results.

section 2.2, the information in the predictors Z_t is not required to be spanned by the state variables X_t. This is the interesting and empirically relevant case where, once a model passes the first-moment test of unpredictable *discounted* returns, a natural second stage test arises from the predictors-based bounds which, as we argued above, become legitimate lower bounds on the volatility of the SDF of the given asset pricing model. Second, we are interested in understanding the link between predictors, horizons and bounds, and not in finding the best predictive model. For this reason we consider a set of traditional and widely used predictors. Our list of potential return predictors is not exhaustive and, in this sense, our bounds provide a conservative estimation of the minimum variability required by any valid SDF. In fact one could expand the set of predictors to make our predictors-based bounds even tighter. For instance, the degree of predictability could be improved both at short- and long- horizons, by using the variance risk premium (see Bollerslev and Zhou, 2009) or the long-run past volatility (see Bandi et al., 2016), respectively. Future research might want to consider how to select the best (in terms of fit, measured by the R^2) subset of predictors to build even tighter bounds. Third, one might think that our conclusions are weakened by the Goyal and Welch (2003; 2008) results that return forecasts based on dividend yields and a number of other variables do not work out of sample. However in our analysis we only require *forecastable returns*. As shown by Cochrane (2008), the out-of-sample R^2 is important for the practical usefulness of return forecasts in forming aggressive real-time market-timing portfolios, but it is not a *test* of forecastable returns. Within our setting, this means that one can find bad out-of-sample performance even when the model actually posits that conditional returns do vary with predictors.[8]

3.2 Predictors-based bounds across horizons

It seems apparent from Table 2 that expected returns vary over time. To evaluate the ability of the predictors to sharpen variance bounds as a diagnostic tool, in this section we compare the predictors-based bounds to the classical, unconditional HJ bounds. We analyze predictive and unconditional bounds at different investment horizons. Whereas Cochrane and Hansen (1992) were the first to carry out this exercise on the unconditional variance bounds, our analysis

[8]For instance Cochrane (2008) sets up a null in which return forecasts account for all dividend-yield volatility, and finds out-of-sample performance as bad or worse than that in the data 40% of the time.

extends their results and highlights the interaction of conditioning information with the horizon dimension.

To compute the predictors-based bounds we use Eq. (6). To compute the first and second conditional moments of asset returns, μ_t and Σ_t, we use the linear predictive model in equation (12).[9] For simplicity, we assume the conditional covariance matrix for returns to be constant, and estimate it as the variance matrix of residuals in the forecasting regressions.

Figure 2 presents our results for SET A. We consider investment horizons of 1 quarter, 1 year, and 5 years. The shortest investment horizon coincides with the sampling interval of returns. In each panel we report the efficient bounds generated with conditioning information (solid lines) along with the unconditional HJ bounds (dashed lines) that make no use of conditioning information.[10] Similar to Cochrane and Hansen (1992), Figure 2 shows that the bottom of the mean standard deviation frontier shifts up and to the left as we increase the investment horizon. Importantly, the picture shows that the predictability across horizons documented in Table 2 translates into a tighter lower bound on the variance of the SDF. In particular Figure 2 shows that the predictor-based bounds are sharper relative to the unconditional ones. For instance, the minimum point of the frontier at the 1-year horizon (at the 5-year horizon, respectively) obtained using conditioning information is about 1.98 (1.79, respectively) times sharper than the unconditional lower bound, thereby substantiating the incremental value of conditioning information in asset pricing applications. The difference between the bounds with and without conditioning information across horizons reflects the predictability documented in Table 2.

Figure 3 presents the same analysis for SET B. In this case, the use of conditioning information yields a bound that is about 1.47 (1.34, respectively) times the unconditional HJ bound at the 1-year horizon (at the 5-year horizon, respectively). Upon comparing Figures 2 with 3, moreover, we observe that expanding the number of assets, i.e. moving from SET A to SET B, leads to a bound that is intrinsically tighter than the one obtained using only returns from SET

[9] We investigate the effect of potential model misspecification in our regressions on the construction of the bounds in Section 4.1.1.

[10] These bounds do not impose that the SDF is a strictly positive random variable. Computing the bounds imposing positivity requires a numerical search procedure; consistent with Hansen and Jagannathan (1991) we find that the bounds imposing positivity are nearly coincident with the simpler bounds in the portion of the parabola where the standard deviation is low, and depart from the simpler bounds only when the standard deviation is relatively high.

A.

[Insert Figures 2 and 3 about here]

Taken together Figures 2 and 3 impart two conclusions. First, these figures highlight the three effects that are at work simultaneously: the conditioning information embedded in moments of returns, the horizon at which this information becomes relevant and the set of assets available for investment. The tightening of the volatility bounds is the combination of these three forces simultaneously at work. Second, although in the predictive regressions the role of the information contained in the predictors becomes more apparent as we lengthen the investment horizon, the predictors-based bounds reveal the important role played by conditioning information even at short horizons.

3.3 Predictors-based bounds and asset pricing models

To illustrate how predictors-based bounds can be used to evaluate different asset pricing models, we select candidate models that embed different preferences and specify the long run and short run risk in distinct ways (see also Hansen (2012)). In particular, we investigate three leading classes of asset pricing models: the habit model, the long-run risks model, and the rare disasters model. We take as representative of the habit and rare disasters classes the Campbell and Cochrane (1999) model as estimated by Aldrich and Gallant (2011), and the model proposed and estimated by Nakamura et al. (2013), respectively.[11,12] With respect to the long-run risks

[11]We consider a specification of the habit model where the preference parameter b that determines the behavior of the risk-free rate is set to zero. Wachter (2006) allows b to differ from zero to match the upward-sloping yield curve for nominal Treasury bonds. Similarly Bansal and Shaliastovich (2013) propose a modified long-run risk framework that successfully match the observed bond yields. We leave for future research the sensitivity analysis of the parameter b on the volatility of the habit-implied SDF, and the analysis of long-run risk type of models that accounts for bond return predictability.

[12]Nakamura et al. (2013) estimate the probability of entering the disaster state using data for 24 countries. Although we use this probability estimated from a panel of countries, in our study the remaining country specific parameters refer to the US. Moreover Nakamura et al. (2013) allow for breaks both in the average growth rate of trend consumption and in the variance of the shock to the growth rate of trend consumption. In our simulations we fix these two values to their post 1973 and post 1946 values, respectively. We do so because our bounds are constructed using post-war data and to make the rare disaster model more comparable to the long-run and habit models which do not allow for trend breaks in consumption. Finally, a number of recent papers (see e.g. Gourio (2008b), Gabaix (2012), Wachter (2013a)) propose versions of the rare disasters model that employ time-varying probability and/or severity of disasters to explain the predictability and volatility of stock returns, among other anomalous features of asset returns. The effect of estimation uncertainty on the moments of a model-implied SDF can be evaluated within the Nakamura et al. (2013) framework but not in any of the above cited papers since they rely solely on calibration. This is why we concentrate on the Nakamura et al. (2013) instead of a framework with time-varying probability and severity of disasters.

model we consider two specifications, namely the models proposed and estimated by Bansal et al. (2016) and Schorfheide et al. (2016). The two models differ in their volatility specification. The first model is closer to the original specification of Bansal and Yaron (2004) in that it considers one common process driving the time-varying volatility in realized and expected consumption growths. The second model accommodates three separate volatility components: one governing dynamics of the expected consumption growth, and the other two controlling temporally independent shocks to consumption and dividend volatility.

We focus on estimated version of these models since we want to quantify the impact of parameter uncertainty on the mean and volatility of their SDFs. Deep parameters in all three classes of models are estimated using a long span of the data to better capture the overall low frequency variations in asset prices and macroeconomic data and to reduce the measurement errors that arise from seasonalities and other measurement problems (see e.g. Wilcox, 1992).[13,14] Finally, the three models have been solved by well established methods that facilitate the computation of the first and second unconditional moments of their SDFs. Specifically, for the habit and the rare disasters model, we solve numerically for a fixed point for the price of the consumption claim as a function of the state(s) of the economy. This method has been used by Campbell and Cochrane (1999), Nakamura et al. (2013), and Wachter (2013b), among others. For the long-run risks models, we follow Bansal and Yaron (2004) and use a linearized solution method based on the Campbell and Shiller (1988a) present-value relation.

Tables D1, D2, D3 and D4 report, for each model, the complete specification of the parameter values for preferences and exogenous dynamics, along with the standard errors of the estimated parameters.[15]

[13] Aldrich and Gallant (2011) use annual data from 1930 to 2008, Bansal et al. (2016) from 1930 till 2009, and Nakamura et al. (2013) from 1890 to 2006. Schorfheide et al. (2016) use a mixed frequency approach based on a sample that contains annual consumption growth data from 1930 to 1959 and monthly data from 1960:M1 to 2014:M12. We also observe that the rare disaster model is at an advantage since it can rely also on periods of time encompassing both World War I and the Great Influenza Epidemic of 1918–1920 to make inference on the dynamics of consumption.

[14] For our chosen specification of the habit model and of the long-run risk with one common volatility, calibrated parameters are available from Campbell and Cochrane (1999) and Bansal et al. (2012), respectively. We verify that our conclusions do not change if one computes the moments of the model-implied SDF m_{t+h}^{X} using calibrated parameters instead of estimated values. These results are available upon request.

[15] Aldrich and Gallant (2011) and Nakamura et al. (2013) present for each parameter the posterior mean and posterior standard deviation. We refer to the posterior mean of each parameter as our point estimate for that parameter.

[Insert Tables D1, D2, D3 and D4 about here]

3.3.1 Predictability of asset returns and discounted returns

Assessing predictability of returns and discounted returns is relevant as it can be argued that when returns discounted using a model-based pricing kernel are predictable, then that model is rejected. A model reflects a view in which the agents' information set is summarized by a specific set of state variables which can be used to build a discount factor that generates discounted gross returns that randomly fluctuate around the unity value. Predictability of discounted returns using any set of variables as predictors implies model rejection either because of the wrong specification of the pricing kernel (i.e. of the wrong transformation of the state variables) or because of some omitted relevant information in the choice of the state variables.

To evaluate predictability of discounted returns, i.e. returns discounted by the SDF of the external habit, the long-run risks and the rare disasters models, we need to proxy for the (unobserved variables in) the SDF of each model. For the rare disasters model and for the long-run risks specification of Bansal et al. (2016) we use Eq. (10). Although the two models share the same functional form, the risk aversion and intertemporal elasticity of substitution parameters entering the SDF are set to the estimated values reported in Bansal et al. (2016), for the long-run risks case, and to the estimated values reported in Nakamura et al. (2013), for the rare disasters case. When using expression (10), we proxy the unobserved return on total wealth by either the market return on the portfolio of all stocks traded in the NYSE, AMEX, and NASDAQ or by the return constructed from the wealth-consumption ratio, as described in Lustig, Van Nieuwerburgh, and Verdelhan (2013). Conclusions do not change by using either one of the two proxies, with results being slightly stronger when we use the return obtained from the wealth-consumption ratio. In the following we report the conservative results obtained by proxying the return on total wealth with the return on the stock market.[16]

For the long-run risks specification of Schorfheide et al. (2016), we replace the one-period

[16]The results obtained when the returns is computed using the consumption-wealth ratio are available upon requests.

log return on total wealth, $r_{a,t+1}$, in Eq. (10) with

$$r_{a,t+1} = \kappa_0 + \kappa_1 pc_{t+1} - pc_t + g_{t+1}$$

where the log price-consumption ratio is given by $pc_t = A_0 + A_1 x_t + A_{2,c}\sigma^2_{c,t} + A_{2,x}\sigma^2_{x,t}$, and the loadings A's are functions of the deep structural parameters given in Table D2. The unobserved latent expected consumption growth, x_t, its time varying variance, $\sigma^2_{x,t}$, and the variance of realized consumption, $\sigma^2_{c,t}$, are proxied with the median filtered values estimated in Schorfheide et al. (2016).[17]

Finally, for the habit case, see Eq. (11), we need to proxy for the surplus consumption ratio s_t. Recall that Campbell and Cochrane (1999) model s_t as an autoregressive, heteroskedastic process which is perfectly (conditionally) correlated with consumption growth innovations, ε^c_t, i.e.

$$\begin{aligned} s_t &= (1-\phi)\bar{s} + \phi s_{t-1} + \lambda_t \varepsilon^c_t \ , \\ \lambda_t &= \frac{1}{\bar{S}}\sqrt{1 - 2 * (s_t - \bar{s})} - 1 \ \ , \text{ if } s_t \leq s_{max} \text{ and zero otherwise} \end{aligned}$$

where the parameters, $\phi, \bar{s}, \bar{S}$ and s_{max} are set to the estimated values reported by Aldrich and Gallant (2011). The innovation term, ε^c_t, is the shock to consumption growth, and following Campbell and Cochrane (1999) we use demeaned values for real nondurables and services consumption log growth from the NIPA tables.[18]

With the SDF of each model at hand, we compute the F-statistics from regressions of discounted stock and bond returns and compare them with the critical values obtained from the bootstrapped empirical distribution of the F-test.[19] Specifically, we run the following discounted

[17] We are grateful to Dongho Song for providing us with these filtered series.

[18] Because the starting point of s_t is not specified, the process is started at its unconditional mean, $\bar{s}$, at the beginning of the consumption growth sample, 1947Q2. Given that our main analysis starts in 1952Q2, the level of s_t is not sensitive the choice of the initial condition.

[19] We bootstrap the empirical distribution of the F-test for the predictive regressions of raw stock and bond returns at different horizons (thus, our empirical distribution of the F-test is asset- and horizon-specific). We compute empirical p-values from a bootstrap experiment to account for poor small-sample properties of the asymptotic F-stats in long-horizon regressions, and to allow for the high persistence of the forecasting variables. The details of the bootstrap algorithm are provided in Appendix B.

returns regression:

$$m_{t,t+h}R^{i}_{t,t+h} = \alpha_{i,h} + \beta_{i,h}Z^{i}_{t} + \varepsilon_{i,t+h}$$

where $i = S, B$ denotes the asset class, $m_{t,t+h} = \prod_{i=1}^{h} m_{t+i}$ is the h-period stochastic discount factor based on the model-implied single period SDFs, $R^{i}_{t,t+h}$ is the return over the holding period h, and Z^{S}_{t} and Z^{B}_{t} are the vectors of stock and bond market predictors, respectively. We then compare the recorded F-statistics from these discounted stock and bond returns regressions with the critical values from the relevant (asset- and horizon-specific) bootstrapped empirical distribution of the F-test, see Table B1 in Appendix B. If the F-statistics is higher than the critical value, then we reject the null of no predictability and, thus, the validity of a given asset pricing model. Conversely, if the null of no predictability of returns duly discounted by the model-implied SDF is not rejected, then we use the predictors-based bounds to assess the performance of the models at long and short horizons and to discriminate among them.

Table 3 displays the results for the long-run risk, the external habit and rare disaster models when the test assets are the Equity Market (Panel A) and a 5-Year Maturity Bond (Panel B), respectively. For reference, we also report the F-test from the predictive regressions that use raw returns. We observe that, with raw returns, the F-stat firmly rejects the null of no predictability with a 95 percent, or larger, confidence interval. However, when we discount returns using the model-implied discount factors, the F-stat falls well below the critical values of the bootstrapped empirical distribution, as documented by the p-values reported in parentheses. In all, empirically, there is no horizon at which we can reject the null of no predictability for discounted returns, and this holds true for all three models. We obtain analogous results when we consider as test assets the returns from rolling over the Treasury bill and from holding bonds with constant maturities greater than five years.[20]

[Insert Table 3 about here]

Before concluding this section, two remarks must be made. First, the results so far show that the predictability of returns disappears when returns are discounted by the model-implied SDF. However the necessary condition in Eq. (8) also states that no mis-pricing should occur, i.e.

[20] Results are available from the authors upon request.

$E[m_{t,t+h} R^i_{t,t+h}] = 1$. In untabulated results, we find that we fail to reject the null that average discounted returns are equal to one for all three asset pricing models under scrutiny. Second, although an analysis of the size and power properties of our F-test based procedure is not an objective of this paper, we are still able to show that the condition of absence of predictability is in fact rejected in cases in which a rejection is the expected outcome. To see this, consider the simplest possible consumption-based asset pricing model, i.e. the model with a representative consumer with CRRA utility and whose endowment/consumption process exhibits i.i.d. growth. The column labeled "CRRA" shows that in this case the F-statistics from discounted returns regressions does reject the null of no predictability soundly, as one definitely expects.

3.3.2 Model-implied SDFs and predictors-based bounds

In this section we assess the three asset pricing models under scrutiny using the predictors-based bounds as a diagnostic tool. Before quantifying the distance between the model-implied SDF and the predictors-based bounds from a statistical perspective, we use Figures 2 and 3 to provide some intuition about the behavior of the model-implied SDF and their potential ability to satisfy the predictors-based bounds.

Figures 2 and 3 display the predictors-based bounds and the SDFs generated by the three competing classes of models for different horizons and for different sets of test assets. For the long-run risks, we report with a triangle the value of the SDF for the Bansal et al. (2016) specification with one stochastic volatility process, and with a star the value for the Schorfheide et al. (2016) specification with three volatility components. Figures 2 and 3 display the predictors-based bound along with the classical HJ bound, for SET A and SET B respectively. In all cases, the mean and the standard deviation of the SDFs reported in the figures represent population values obtained using estimated parameters.[21]

A first observation that emerges clearly from Figures 2 and 3 concerns the importance of

[21] In this section we abstract from parameter and small-sample uncertainty. Thus, to compute the mean and the volatility of the SDF of each model we use the dynamics of consumption growth and of the state variables posited by that model, and we simulate 600,000 monthly observations (50,000 years) of the model-implied SDF for the long-run risks and habit models, and 50,000 annual observations for the rare disaster model. From this long time series we then calculate the unconditional moments of the corresponding SDF. Using a single simulation run to infer the population values for the entities of interest is consistent with, among others, the approach of Campbell and Cochrane (1999) and Beeler and Campbell (2009).

jointly considering conditioning information and horizons for the equity premium puzzle. Figure 2 shows that the SDFs of all models satisfy the unconditional HJ bounds at the 1-year horizon. This is not surprising for the Aldrich and Gallant (2011) habit and the Bansal et al. (2016) long-run risks models since in these two cases the estimated parameters are close to the calibrated ones provided by Campbell and Cochrane (1999) and Bansal et al. (2012), respectively. In turn, these two latter models are calibrated to offer conformity with the historical unconditional annual real return on the risk-free bond and the equity market. However, the conclusions are different when we incorporate conditioning information. In this case the classical long-run risks model with one time-varying volatility process struggles to meet the bounds, whereas the specification with multiple volatility components satisfy the bounds comfortably. The habit model meets the restriction, since its SDF overlaps with the predictors-based bound. As expected, using SET B and hence expanding the set of assets exacerbates these results: the habit and classical long-run risks models fall largely below the bounds, whereas the long-run risks specification that accommodates three volatility processes lies exactly on top of the predictors-based bound (see Figure 3). It is important to stress that if we considered the 1-year bounds with no conditioning information, we would have concluded that the equity premium puzzle can be resolved as long as sufficient time-nonseparability is incorporated in preferences. However, the predictors-based bounds highlight that what really matters is the interaction between preferences and the dynamics of the state variables.

The figures shed also some light on the long-horizon behavior of the SDFs from the competing models. This is of particular interest given the recent theoretical and empirical research in macro-finance that has highlighted the importance of capturing low frequency components for an asset pricing model to be successful. One would hope for these low frequency components to make asset pricing puzzles less pronounced at longer horizons. With the visual aid of Figure 2 we can see that this statement is fulfilled when no conditioning information is incorporated: the 5-year unconditional HJ bound is satisfied with good margin by all models. However, this conclusion changes significantly at the light of the predictors-based bound at the 5-year horizon. In this case the habit and rare disaster models satisfy the predictors-based bound with a good margin, while the long-run risk model finds it onerous to satisfy the bounds at such a long horizon (see

Panel A in Figures 2 and 3).

The above results point to an interesting fact about the role played by preferences and state dynamics. Although the rare disasters and long-run risks models share the same preferences for early resolution of uncertainty, the two model-implied SDFs have very different behaviors. This can be explained by the different ways the long-run risks and the rare disasters models decompose consumption. Both models assume that the level of log consumption includes a deterministic trend and a stochastic trend. In the long-run risk model, in particular, the growth rate of the stochastic trend captures expected consumption growth and contains (i) a persistent component, (ii) long-run variation in volatility. In the rare disaster model, on the other hand, the growth rate of the stochastic trend follows a jump process and captures potential consumption. Differently from the long-run risk, the rare disaster model of Nakamura et al. (2013) incorporates also a transitory component in the log consumption level: this component is labeled disaster gap, and allows for partial recoveries after disasters. It is the interaction between time-nonseparability in preferences and these state dynamics that drives the different ability of these two models to satisfy the predictors-based bounds. To reinforce this point we consider in Figure 4 two specifications of the rare disaster model. Both specification have the same recursive preferences, but they allow for different disaster dynamics. In particular we compare the model-implied SDF of Nakamura et al. (2013) (triangles), with the SDF of a model with permanent, one-period disasters of the type analyzed in Barro (2006) (stars). The SDF from the permanent disasters model is below the predictors-based bound (solid line) at the one-year horizon independently from the test assets used, it satisfies the predictors-based bound at long-horizons using SET A, while it struggles when we incorporate bond returns. On the other hand, the SDF from the model that allows for partial recoveries after disasters that unfold over multiple years meets comfortably the predictors based bound, with an SDF well above the predictors-based bound across all horizons and for all test assets.

[Insert Figure 4 about here]

The evidence so far points to the ability of the predictors-based bounds to discriminate across models. We next quantify the robustness of these conclusions when we account for estimation uncertainty.

3.3.3 How uncertain is the distance from the bounds?

In this section we test the volatility restrictions imposed by the predictors-based bounds. In particular, we evaluate whether the difference between the estimated model-implied variance of the SDF, $Var\left(m_t^X\right)$, and the estimated predictors-based bound, $\sigma_Z^2(v)$, is large in a statistical sense. To properly compare the moments of the model-implied SDF with the variance bounds we must account for two sources of uncertainty. First, the computation of the mean and variance of a model-implied SDF relies on the estimates of the exogenous state dynamics, and hence it reflects the uncertainty in the parameters describing these dynamics. Second, the volatility bounds are estimated from the data, and hence they must reflect the uncertainty surrounding the linear predictive model used to compute the conditional moments of returns, see Eq. (12). Hereafter, we account for these sources of uncertainty to obtain the finite sample distribution of the difference $\Delta = Var\left(m_t^X\right) - \sigma_Z^2(v)$ (see Cecchetti, Lam, and Mark (1994) and Burnside (1994) for a related approach). Intuitively, for the three models under scrutiny we draw the parameters for the state dynamics using the values given in Tables D1, D2, D3 and D4. We draw the coefficients of each asset return predictors analogously. Given these parameters, for each model we simulate an SDF of length equal to our dataset, i.e. 742 months. Finally, we compute the model-implied variance of the SDFs, the predictors-based bounds, and their difference. We repeat this exercise 10000 times. Further details are provided in Appendix C.

The results are summarized in Table 4, for SET A and SET B. We start by looking at the rightmost block of results. These results reflect the estimation uncertainty in both the predictors-based bounds and the moments of the model-implied SDF. The first conclusion that emerges from Table 4 is that independently from the horizon and test assets considered, the SDF of the rare disaster model satisfies comfortably the predictors-based bound even after accounting for estimation uncertainty. When we consider the habit and the more classical Bansal et al. (2016) long-run risks models, and we focus on the predictors-based bound at the 5-year horizon constructed from SET A, it is the habit model that satisfies the bounds, while the long-run risks model fails to meet the restriction imposed by the bounds at the 5% confidence interval. At the 1- and 4-quarter horizons both the habit and long-run risks models are rejected at the 5% confidence level. Interestingly, within the long-run risks class of models, the novel specification of Schorfheide et al. (2016) with multiple volatility components constitutes a noticeable improvement upon the classical long-run risks model based on a single time-varying volatility. Indeed, the Schorfheide et al. (2016) model now satisfies the 1-year and 5-year predictor-based bounds constructed from SET A, although it still fails the the bounds constructed from SET B.

[Insert Table 4 about here]

The remaining two blocks of Table 4 report a "decomposition" of the uncertainty in the comparison of $Var\left(m_t^X\right)$ with $\sigma_Z^2(v)$.[22] Following Cecchetti et al. (1994), we think of this uncertainty as arising from three basic sources. Given the expected value of the SDF, v, there is uncertainty in both the location of the bound, $\sigma_Z^2(v)$, and the standard deviation implied by the model, $Var\left(m_t^X\right)$. In addition, there is uncertainty induced by the fact that the mean of the SDF for the model, v, must be estimated. The leftmost block reports the estimated distance for a fixed mean of the SDF. The middle block of the table reports the uncertainty in Δ that arises solely from randomness in v. If no uncertainty in the mean of the SDF is considered then all models meets the volatility restriction at the 1- and 5-year horizons, for both SET A and SET B, with a 5% confidence level (see leftmost panel). Thus, uncertainty surrounding the volatility bounds is sufficient to make the SDFs satisfy the restrictions *if* one is fully confident about the location of the mean of the SDF. However, once the uncertainty about the mean of the SDF is accounted for, the conclusions are reversed (see the middle block) and, when using SET B, all models fail to meet the volatility restrictions with the sole exception of the rare disaster model. As one would expect, re-introducing uncertainty about the bounds helps the model: moving from the middle to the rightmost block shows that the Δ becomes closer to positive values. However, the uncertainty on the bounds is not large enough to undo the uncertainty in the mean of the SDF. Hence, by comparing the distance of a model-implied SDF across the different columns in Table 4, it is clear that the main source of uncertainty lies in the fact that v must be estimated.

Finally, the table also shows that the decrease in distance due to the uncertainty in the mean of the SDF is more severe at longer horizons: Focusing on SET B, at the 1-year horizon, e.g., the distribution of the distance for the habit model shifts to the left from 0.349 (5% quantile

[22] The presence of nonzero covariances between moments of the SDF implies that entries in the table need not add up to other elements in the same row. That's why we abuse of the term "decomposition".

value) to −0.158; this is a much smaller shift than the one occurring at the 5-year horizon from 2.178 to −0.496. Although Figure 3 - Panel B suggests that the habit model has the ability to satisfy the bounds at long horizons, we find that the effect of the uncertainty in the mean SDF at long-horizon is large enough to make the habit-implied SDF fail when using SET B.

In sum, this section shows that by incorporating conditioning information from a well-established set of stock and bond predictors the predictors-based bounds are a useful tool to assess the performance of candidate asset pricing models at multiple horizons. It is noteworthy that each asset pricing model parametrization approximates quite reasonably the (annual) unconditional equity premium and the real risk-free return, while simultaneously calibrating closely to the first two moments of consumption growth. The rare disaster model handily meets the predictor-based bounds across horizons and asset classes, even after accounting for estimation uncertainty. When we look at the habit model and the classical version of the long-run risks with a single common source of stochastic volatility (see Bansal et al. (2016)), both models fail to meet the restrictions imposed by the predictor-based bounds at the 1-year horizon, with the standard deviation of their SDFs never approaching the bounds. The version of the long-run risks that accommodates multiple volatility components (see Schorfheide et al. (2016)) represents an improvement in the sense that it satisfies the 1-year bound constructed from SET A. At long horizons, finally, while the classical long-run risks model with one time-varying volatility process generates an SDF not volatile enough independently from the test assets considered, the habit model and a modified long-run risks with multiple volatility components satisfy the predictors-based bound when using SET A but they both fail when we consider our largest set of equity and bond returns.

3.4 Historical versus model-implied predictability

In this section we compare the ability of an asset pricing model to satisfy the predictors-based bound with the ability of the model to reproduce the return predictability observed in the data. The following empirical evidence concentrates exclusively on the stock market and uses the log price-dividend ratio as the sole predictor.

Figure 6 displays results for the predictive regressions of future gross real returns over the

1, 3, and 5 years horizons. The sample statistics considered are the R^2 values obtained from the regressions. For model-implied predictive regressions the figure displays the median R^2 with associated 95% confidence interval, obtained from 1000 regressions run on simulated data. Finally, the red triangles correspond to statistics computed from actual U.S. data.

[**Insert Figure 6 about here**]

The key finding is that all models perform well in terms of the univariate excess return predictability regressions. Specifically, for all horizons the model-based predictive intervals contain the R^2s obtained from the actual data.

Thus all these three classes of models can match the return predictability observed in the data when (sample and parameters) uncertainty is taken into proper account. This evidence, in turn, reinforces the findings that return predictability can be used to sharpen the variance bounds of SDFs and that the predictors-based bounds emerge as a criteria more effective than model-based predictability to discriminate among asset pricing models.

4 Extensions

In Section 3 we have shown that incorporating predictability of asset returns does make the variance bounds tighter and hence it imposes a harder yardstick on asset pricing models that deliver unpredictable discounted returns. In this Section we first answer the question of which asset class, stocks or bonds, contributes the most to the sharpening of unconditional variance bounds exhibited in the previous section. We also analyze the robustness of our results to the possibility of misspecification of the model for the conditional moments of returns.

4.1 Stock-based versus bonds-based variance bounds

To check the relative importance of different asset classes for predictability, and hence for sharpening the bounds, we consider the following experiment. We build the variance bounds according to two different scenarios, each one imposing different restrictions on the predictive system in (12). In the first case (restriction I) the restriction that stock returns are unpredictable is imposed. In the second case (restriction II), it is instead the returns from the strategy rolling over

Treasury bills that are restricted to be unpredictable.

Figure 5 displays the predictors-based bound obtained when we impose restriction I (dashed red line with circles) and restriction II (dashed green line with triangles). To compare the results with those obtained in the previous section, the Figure reports also our unrestricted predictors-based bound (solid black line) and the HJ variance bounds (dashed violet line) Panel A and Panel B display the results for SET A and SET B, respectively.

[Insert Figure 5 about here]

From Figure 5 we can draw two main conclusions. First, it is the predictability in stock returns that really tightens the variance bounds, particularly so at the long horizon (see Panel A). For instance, under Restriction I the minimum point of the frontier for the volatility of SDF at the 1-year horizon (at the 5-year horizon, respectively) based on the returns in SET A is only 0.63 (0.62, respectively) times the minimum point of the predictors-based bound obtained when predictability is unrestricted across asset classes. Second, by comparing Panel A with Panel B, it is apparent that the additional tightening brought about by stock return predictability is less effective as we expand the set of asset.[23] One might then wonder if it is the predictability of bond returns the key to the results of SET B. In unreported evidence (available from the authors upon request) we show that this is not the case: the additional tightening due to the predictability of treasury government bond is rather marginal.[24]

Summing up, the shape of the predictors-based variance bounds on SDFs essentially depends on the model we choose for predicting stock returns, whereas the predictability of bond returns plays a rather marginal role.

4.1.1 Predictability, model mis-specification and variance bounds

We conclude this section by investigating the performance of the bounds along two further dimensions: robustness and efficiency. Recall that the results presented so far are obtained

[23] Using SET B and imposing restriction I, the minimum point of the predictors-based bound is 0.725 at the 1-year horizon and 2.034 at the 5-year horizon. These points are close to the minimum point of the predictors-based bound obtained when predictability is unrestricted, namely 0.806 and 2.306.

[24] In particular even when we shut down the predictability of all the treasury government bonds simultaneously, the minimum value of the variance bound at the 5-year horizon is still 0.93 times that of the benchmark case (i.e. the variance bound generated with conditioning information and unrestricted predictability, see red solid line).

under the assumptions of a time-invariant variance-covariance matrix for returns and a linear model for their conditional means. To investigate possible mis-specification of the conditional moments and the efficiency of the bound we plot in Figure 7 alternative implementations of the variance bounds. Specifically, along with the predictors-based bound obtained following Bekaert and Liu (2004) (BL), in this figure we plot the bounds obtained following alternatively Gallant et al. (1990) (GHT) and Ferson and Siegel (2003, 2009) (FS).

[**Insert Figure 7 about here**]

Bekaert and Liu (2004) show that their bound, obtained by maximizing the Sharpe ratio over all returns obtained from portfolios that condition on Z_t and that cost 1 on average (see Section 2), must be a parabola under the null of correct moments specification. Figure 7 shows that in our case we obtain a smooth parabola indeed. The figure, moreover, shows that the GHT bound and the BL bound are virtually on top of each another, i.e. there is no duality gap. This suggests that the BL bound closely approximates the efficient use of conditioning information. Overall the three alternative implementations of the variance bounds that incorporate information from the predictors Z_t generate similar bounds with no visible misspecification. The FS is the lowest bound, this is readily understood by observing that the FS bound collects all those payoffs that are generated by trading strategies that reflect the information available at time t, and that have unit price almost surely equal to one, and not just on average as for the BL case.[25] Although the FS approach yields the most conservative bound, the differences between the three approaches would not change our conclusions. This evidence suggests that misspecification of the conditional moments does not seem to be a driver of our results.

[25]More formally, the FS bound (see Ferson and Siegel (2003)) is defined as

$$\sigma^2_{FS}(v) = \nu^2 \sup_{R^w_{t+h} \in \mathcal{R}^{FS}} \left(\frac{E(R^w_{t+h}) - \nu^{-1}}{Var\left(R^w_{t+h}\right)} \right)^2$$

where

$$\mathcal{R}^{FS} = \left\{ R^w_{t+h} \in \mathcal{R}^Z \mid w_t' e = 1 \; almost \; surely \right\}$$

i.e. the FS variance bound follows from maximizing the Sharpe ratio over the set of returns from portfolios that, while conditioning on Z_t, are required to have unit price almost surely, and not just on average. Therefore, it is evident that $\sigma^2_{FS}(v) \leq \sigma^2_Z(v)$.

5 Conclusions

We propose an empirical strategy to disentangle different asset pricing models. The strategy uses sequentially two necessary conditions that must be satisfied by an asset pricing model. The first necessary condition says that the predictability of returns should disappear when returns are discounted by the model-implied Stochastic Discount Factors (SDF). The second condition requires the variance of the model-implied SDF to be larger than the variance dictated by the predictors-based variance bounds, i.e. bounds on the variance of those SDF that price a given set of returns conditional on the information contained in a vector of returns predictors.

For the three leading classes of consumption-based asset pricing models considered in the paper (namely the long-run risks model of Bansal et al. (2016) and its modified version with multiple volatility components suggested by Schorfheide et al. (2016), the habit-formation model of Campbell and Cochrane (1999), and the rare disaster model of Nakamura et al. (2013)) we show that the condition of no predictability for discounted returns cannot be rejected. Therefore we use the predictors-based bounds to assess the performance of the three models at short and long horizons. We show that the predictors-based bounds provide a useful tool to discriminate between models.

When looking across models, we document that the rare disaster model handily meets the predictor-based bounds across horizons and asset classes, even after accounting for estimation uncertainty. Interestingly, the asset pricing models under scrutiny reproduce reasonably well the annual *unconditional* equity premium and the real risk-free return. However our evidence shows that, upon accounting for the information contained in a set of returns predictors, the variance of both the SDFs implied by the Bansal et al. (2016) classical long-run risks model and by the Campbell and Cochrane (1999) model fail to meet the lower bound restriction at 1-year horizon.

At the 5-year horizon, the *unconditional* HJ bound is satisfied by all the three models under scrutiny, yielding support to the intuition that unconditional asset pricing puzzles are less pronounced at longer horizons. This conclusion, however, is not robust to the introduction of conditioning information: in particular, only the rare disaster models maintain the capability of addressing the equity premium at long horizons when the set of traded assets is augmented to include also the returns from holding Treasury Bonds.

When concentrating on the class of long-run risks models, in particular, we document that multiple volatility components a-la Schorfheide et al. (2016) are needed for the long-run risks to satisfy the 1-year and 5-year predictors-based bounds constructed from equities. As for the rare disasters models, the version of Nakamura et al. (2013) that allows for partial recoveries after disasters that unfold over multiple years performs better than a permanent, one-period disaster model.

The predictors-based bounds represent a convenient tool for researchers. In fact, the dynamic asset pricing models under consideration are constructed from a mixture of assumptions about preferences (such as recursive utility or habit persistence) and exposure to fundamental shocks (such as jumps or persistent shocks). The bounds yield a graphical and intuitive comparison of the performance of asset pricing models. Consistent with the idea that all models are approximations of reality and as such likely to be mis-specified along some dimensions, the predictors-based bounds use the investment horizon and conditioning information as the fundamental ingredients to set apart models identical in terms of preferences (as in the case of Bansal et al. (2016) and Nakamura et al. (2013) at long horizons), or to identify the common behavior among apparently different models (as, in the case of the Campbell and Cochrane (1999) and Bansal et al. (2016) at the 1-year horizon). Importantly, the predictors-based bounds emerge as a criteria more effective than model-based predictability to discriminate among asset pricing models.

We conclude by observing that whereas we take the parameters of preferences and state variables dynamics as given, and we investigate the model performance at the light of the predictors-based bounds, one could also use the information contained in the predictors-based bounds as an additional constraint in the estimation/calibration of a model. For instance one could reverse engineer the bounds, to back out hard-to-detect parameters such as the probability of disaster risks. We leave this investigation to future research.

Appendices

A Data

We consider a set of quarterly equity and bond returns over the period 1952Q2 to 2012Q4. Our choice of the start date is dictated by the availability of data for our predictors. Real returns are computed by deflating nominal returns by the Consumer Price Index inflation. We obtain the time series of bond and stock returns using monthly daily returns on stocks and bonds.

1. Stock returns: Return data on the value-market index are obtained from the Center for Research in Security Prices (CRSP) at the University of Chicago. We use the NYSE/Amex value-weighted index with dividends as our market proxy, R_{t+1}. Quarterly returns are constructed by compounding their monthly counterparts. The h-horizon return is calculated as $R_{t+h} = \exp(r_{t,t+h}) = \exp(r_{t+1} + \ldots + r_{t+h})$ where $r_{t+j} = \ln(R_{t+j})$ is the 1-period log stock return between dates $t+j-1$ and $t+j$ and R_{t+j} is the simple gross return.
2. Bond returns: Returns on bonds are extracted from the US Treasuries and Inflation Indices File and the Stock Indices File of the Center of Research in Security Prices (CRSP) at the University of Chicago. The CRSP US Treasuries and Inflation Indices File provides returns on constant maturity coupon bonds, with maturities ranging from 1 year to 30 years, starting on January, 1942. Quarterly returns are constructed by compounding their monthly counterparts. The h-horizon return is calculated as (R_{t+j}) is the 1-period log stock return between dates $t+j-1$ and $t+j$ and R_{t+j} is the simple gross return. The nominal short-term rate $(R_{f,t+1})$ is the annualized yield on the 3-month Treasury bill taken from the CRSP treasury files.
3. Stock market predictors: consumption-wealth ratio, cay_t, see Lettau and Ludvigson (2001); the dividend-price ratio, dp_t, see Campbell and Shiller (1988a) and Campbell and Shiller (1988b).
4. Bond return predictors: the lagged yield spread, spr_t; the one month maturity US treasury-bill rates, y_t (see, e.g., Campbell (1987) and Fama and French (1989)). The term spread, spr_t, is the difference between the long term 5-year yield on government bonds and the Treasury-bill (see, e.g., Campbell (1987) and Fama and French (1989)).

5. Inflation: we use the seasonally unadjusted CPI from the Bureau of Labor Statistics. Quarterly inflation is the log growth rate in the CPI.

B Bootstrap algorithm of F-test

This appendix describes the steps to bootstrap the empirical distribution of the F-test for the predictive regressions of stock and bond returns at different horizons. To obtain the empirical distribution of the F-statistics for the long-horizon regressions, we follow the lead of Maio (2008).

First, recall that our predictive regression for asset return is

$$R^i_{t+h} = \beta^i_{0,h} + \beta^i_{1,h} Z^i_t + u^i_{t+h}$$

where $i = S, B$ stands for stocks and bonds, respectively, $h = 1,\ 4,\ 20$ quarters, and $Z_t = \left(Z^S_t, Z^B_t\right)$ denotes the vector of returns' predictors, potentially different for stock and bonds. The bootstrap algorithm consists of the following steps:

1. We first estimate the following predictive system, which imposes the null hypothesis of no predictability for asset returns, R^i_{t+h}, and assumes that each one of the persistent predictors follows an AR(1) process:

$$\begin{aligned} R^i_{t+h} &= \alpha^i_{0,h} + \varepsilon^i_{t+h} \\ z^{i,k}_{t+1} &= \psi^{i,k} + \phi^{i,k} z^{i,k}_t + v^{i,k}_{t+1} \end{aligned}$$

   where $k = 1,\ 2,\ ...,\ K$, $\left[z^{S,1}_t\ z^{S,2}_t\ \cdots\ z^{S,K}_t\right]' = Z^S_t$ and $\left[z^{B,1}_t\ z^{B,2}_t\ \cdots\ z^{B,K}_t\right]' = Z^B_t$. We save the estimated, $\widehat{\alpha}^i_{0,h}$, $\widehat{\psi}^{i,k}$, $\widehat{\phi}^{i,k}$, and residuals, $\widehat{\varepsilon}^i_{t+h}$ and $\widehat{v}^{i,k}_{t+1}$

2. Then in each replication $m = 1, ..., 10,000$, we generate pseudo-samples for the innovations in asset returns and predictors by drawing with replacement from the residuals. We first create the time indices, $s^m_1, s^m_2, ..., s^m_T$, from the orginal time sequence, $1, 2, ..., T$.

$$\begin{aligned} \left\{\widehat{\varepsilon}^{i,m}_{t+h}\right\}, t &= s^m_1, s^m_2, ..., s^m_T \\ \left\{\widehat{v}^{i,k,m}_{t+1}\right\}, t &= s^m_1, s^m_2, ..., s^m_T \end{aligned}$$

	Stock returns			Bond returns		
	1-quarter	1-year	5-year	1-quarter	1-year	5-year
10%	1.9867	7.5294	13.2696	2.2093	4.4909	14.1618
5%	2.4128	9.1719	19.0007	2.6438	5.3209	16.8890
1%	3.4357	12.9734	35.6098	3.5172	7.1635	23.7077

Table B1 Empirical distribution of F-statistics for predictive regressions. This table reports the critical value at 10%, 5% and 1% significance level of the bootstrapped distribution of the F-test for the predictive regressions of raw stock and bond returns at different horizons.

Notice that the innovations in both asset returns and predictors have the same time index to account for their contemporaneous cross-correlation.

3. For each replication $m = 1, ..., 10,000$, we construct the pseudo-sample of asset return and predictors, by imposing the null:

$$\begin{aligned} R_{t+h}^{i,m} &= \widehat{\alpha}_{0,h}^{i} + \widehat{\varepsilon}_{t+h}^{i,m} \\ z_{t+1}^{i,k,m} &= \widehat{\psi}^{i,k} + \widehat{\phi}^{i,k} z_{t}^{i,k,m} + \widehat{v}_{t+1}^{i,k,m} \end{aligned}$$

4. Finally, we estimate the predictive regression by using the artificial data rather than the original data:

$$R_{t+h}^{i,m} = \beta_{0,h}^{i,m} + \beta_{1,h}^{i,m} Z_{t}^{i,m} + u_{t+h}^{i,m}$$

Within each replication we record the F-statistics for the regression. We then use these recorded F-stat to construct the F-statistics empirical distribution. The following table reports the critical values of the bootstrapped F-statistics at standard significance level.

C Quantifying uncertainty

This appendix explains how we account for uncertainty in the evaluation of the difference between the estimated model-implied standard deviation of the SDF, $\sigma\left(m_t^X\right)$, and the estimated predictors-based bound, $\sigma_{Z(v)}$. First, to compute the mean and variance of a given model-implied SDF, we take into account the uncertainty of the parameters of the exogenous state dynamics. Second, since the predictors-based bound are estimated from the data, we also account for the uncertainty surrounding the linear predictive model, which is used to compute the conditional moments of asset returns, see Eq. (12). Finally, we obtain the finite sample distribution of the difference, $\Delta = \sigma\left(m_t^X\right) - \sigma_{Z(v)}$, based on a related approach in Cecchetti et al. (1994) and Burnside (1994).

To make explicit the dependence of the moments of the SDF from the parameters of the model, we denote the model-implied mean and standard deviation of the SDF as

$$\mu_m(\phi,\psi)$$

$$\sigma_m(\phi,\psi)$$

where ϕ denotes the vector of parameters that characterize the preferences, and ψ contains all the parameters associated with the state dynamics. For instance, in the LRR model, $\phi = (\delta,\gamma,\psi)$ and $\psi = (\mu,\mu_d,\phi,\varphi_d,\rho_{dc},\rho,\varphi_e,\overline{\sigma},\upsilon,\sigma_\omega)$, (see Table D1). Then, for the Aldrich and Gallant (2011), Bansal et al. (2016), and Nakamura et al. (2013) models, we draw the parameters related to state dynamics, i.e. ψ, from normal distributions with mean and standard deviation given in Table D1, D3 and D4, respectively. Similarly, for the Schorfheide et al. (2016) model we use the distributions detailed in D2. In this latter case, we also verify that our results are robust to drawing parameters from normal distributions with 90% intervals covering the maximum among the 5% and 95% values reported in Table D2. Given these parameters, for each model we simulate an SDF of length equal to the size of our data, i.e. 742 months, and compute the model-implied mean, $\mu_m(\phi,\psi)$, and variance, $\sigma_m(\phi,\psi)$.

Moving to the predictors-based bounds, we draw the coefficients in Eq. (12), $\beta_0^S, \beta_1^{S,cay}, \beta_1^{S,dp}$, for stocks, and $\beta_0^B, \beta_1^{B,spr}, \beta_1^{B,y}$, for bonds, from normal distributions. For each parameter, the mean of the normal distribution is set to the sample estimates from the predictive regressions, and the standard deviation is provided by the Newey-West t-stats corrected value (see the regression results in Table 2). Given these parameters, we simulate a series of returns of length equal to 742 months, re-estimate the predictive regressions, and compute the predictors-based bounds using Eq. (6).

Finally, we compute the difference between $\sigma_m(\phi,\psi)$ and $\sigma_{Z(\mu_m)}$ and repeat this exercise

10000 times. We consider three different cases. In the first case (reported in the leftmost block in Table 4) the mean of the SDF, $\mu_m(\phi, \psi)$, is fixed to the long-run mean obtained from a simulations run of 600,000 months, v^{lr}. Given this (population) value of the mean SDF, the only source of uncertainty stems from the location of the bound, $\sigma_Z(v^{lr})$, and the standard deviation implied by the model, $\sigma_m(\phi, \psi)$. The second case (reported in the middle block in Table 4) keeps the parameters of the predictive regressions fixed at their point estimates so that (for a given value of the average SDF, $\mu_m(\phi, \psi)$) there is no uncertainty in the location of the bound, $\sigma_{Z(\mu_m)}$. However, there is uncertainty induced by the fact that the mean of the SDF for the model, $\mu_m(\phi, \psi)$, must be estimated, as well as uncertainty in the standard deviation implied by the model, $\sigma_m(\phi, \psi)$. The last case (reported in the the rightmost block) is the case when there is: (i) uncertainty induced by the fact that the mean of the SDF for the model, $\mu_m(\phi, \psi)$, must be estimated (ii) uncertainty in the location of the bound, $\sigma_{Z(\mu_m)}$ (for a given expected value of the SDF), (iii) and uncertainty in the standard deviation implied by the model, $\sigma_m(\phi, \psi)$ (for a given expected value of the SDF).

D Model parameters

	Parameter	Bansal-Kiku-Yaron (BKY)
Preferences		
Time preference	δ	0.9989 (0.0010)
Risk aversion	γ	7.42 (1.55)
EIS	ψ	2.05 (0.84)
Consumption growth dynamics, g_t		
Mean	μ	0.0012 (0.0007)
Dividends growth dynamics, $g_{d,t}$		
Mean	μ_d	0.0020 (0.0017)
Persistence	ϕ	4.45 (1.63)
Volatility parameter	φ_d	5.00 (1.39)
Comsumption exposure	π	0
Correlation between innovations	ρ_{dc}	0.49 (0.33)
Long-run risk, x_t		
Persistence	ρ	0.9812 (0.0086)
Volatility parameter	φ_e	0.0306 (0.0160)
Consumption growth volatility, σ_t		
Mean	$\overline{\sigma}$	0.0073 (0.0015)
Persistence	v	0.9983 (0.0021)
Volatility parameter	σ_w	2.62×10^{-5} (3.10×10^{-6})

Table D1 Parametrization of long-run risk asset pricing model by Bansal et al. (2016). The estimated values and corresponding standard errors (in parentheses) are taken from (Bansal et al., 2016). The model is simulated at the monthly frequency.

	Distr.	Schorfheide-Song-Yaron (SSY) Posterior		
		5%	50%	95%
Household Preferece				
δ	B	–	0.999	–
ψ	G	1.13	1.93	3.42
γ	G	5.44	8.60	12.97
Preference Risk				
ρ_λ	U	0.916	0.956	0.982
σ_λ	IG	0.0003	0.0005	0.007
Consumption Growth Process				
ρ	U	0.949	0.987	0.999
φ_x	U	0.139	0.232	0.506
σ	IG	0.0020	0.0032	0.0044
ρ_{h_c}	N^T	0.973	0.991	0.996
$\sigma^2_{h_c}$	IG	0.0074	0.0088	0.0100
ρ_{h_x}	N^T	0.987	0.994	0.999
$\sigma^2_{h_x}$	IG	0.0027	0.0039	0.0061
Dividend Growth Process				
ϕ	N	2.82	4.15	5.44
π	N	0.204	1.54	4.31
φ_d	U	3.56	5.02	7.83
ρ_{h_d}	N^T	0.948	0.967	0.984
$\sigma^2_{h_d}$	IG	0.0174	0.0393	0.0833
Consumption Measurement Error				
σ_ϵ	IG	0.0006	0.0010	0.0016
σ^a_ϵ	IG	0.0061	0.0231	0.0423

Table D2 Parametrization of long-run risk asset pricing model by Schorfheide et al. (2016). The estimated values are taken from Schorfheide et al. (2016). The model is simulated at the monthly frequency.

	Parameter	Campbell-Cochrane
Preferences		
Time preference	δ	0.9903 (0.0004)
Risk aversion	γ	1.9756 (0.0772)
Consumption growth dynamics, g_t		
Mean	$\overline{g}$	0.0017 (0.00007)
Volatility parameter	σ	0.0050 (0.00019)
Dividend growth dynamics, Δd_t		
Volatility parameter	σ_w	0.0319 (0.0014)
Corr between innovitions	ρ_{dc}	0.1945 (0.0093)
Steady state surplus consumption ratio	$\overline{S}$	0.0637
Persistence in consumption surplus ratio	ϕ	0.9877 (0.0003)
Log of risk-free rate	$r^f \times 10^2$	0.0854 (0.0365)

Table D3 Parametrization of external habit asset pricing model. The estimated values and corresponding standard errors (in parentheses) are taken from (Aldrich and Gallant, 2011). The model is simulated at the monthly frequency.

	Parameter	Annual
Preferences		
Time preference	δ	0.967
Risk aversion	γ	6.4
Elasticity of intertemporal substitution	ψ	2
Potential consumption dynamics, g_t, only for US		
Mean of potential consumption growth, x_t	μ	0.022 (0.003)
Volatility parameter	σ_ϵ	0.003 (0.002)
Volatility parameter	σ_η	0.018 (0.002)
Disaster parameters		
Probabilities of		
a world-wide disaster	p_W	0.037 (0.016)
a country will enter a disaster when a world disaster begins	p_{CbW}	0.623 (0.076)
a country will enter a disaster "on its own."	p_{CbI}	0.006 (0.003)
a country will stay at the disaster state	$1 - p_{Ce}$	0.835 (0.027)
Disaster gap process, z_t		
Persistence	ρ_z	0.500 (0.034)
a temporary drop in consumption caused by shock, ϕ_t		
Mean	ϕ	−0.111 (0.008)
Volatility parameter	σ_θ	0.121 (0.015)
a permanent shift in consumption caused by shock, θ_t		
Mean	θ	−0.025 (0.007)
Volatility parameter	σ_ϕ	0.083 (0.006)

Table D4 Parametrization of asset pricing model incorporating rare disasters. The value of the parameter estimates and corresponding standard errors (in parentheses) are taken from Nakamura et al. (2013). The model is simulated at the annual frequency.

References

Abel, A. B. 1990. Asset Prices under Habit Formation and Catching Up with the Joneses. *American Economic Review* 80:38–42.

Aldrich, E. M., and A. R. Gallant. 2011. Habit, Long-Run Risks, Prospect? A Statistical Inquiry. *Journal of Financial Econometrics* 9:589–618.

Backus, D., M. Chernov, and S. Zin. 2014. Sources of Entropy in Representative Agent Models. *Journal of Finance* 69:51–99.

Bakshi, G., and F. Chabi-Yo. 2012. Variance bounds on the permanent and transitory components of stochastic discount factors. *Journal of Financial Economics* 105:191–208.

Bandi, F. M., B. Perron, A. Tamoni, and C. Tebaldi. 2016. The scale of predictability. CIRANO Working Papers 2015s-21, CIRANO. URL `https://ideas.repec.org/p/cir/cirwor/2015s-21.html`.

Bansal, R., D. Kiku, and A. Yaron. 2012. An empirical evaluation of the long-run risks model for asset prices. *Critical Finance Review* 1:183–221.

Bansal, R., D. Kiku, and A. Yaron. 2016. Risks for the long run: Estimation with time aggregation. *Journal of Monetary Economics* 82:52 – 69.

Bansal, R., and I. Shaliastovich. 2013. A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets. *Review of Financial Studies* 26:1–33.

Bansal, R., and A. Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59:1481–1509.

Barro, R. J. 2006. Rare Disasters and Asset Markets in the Twentieth Century. *Quarterly Journal of Economics* pp. 823–866.

Beeler, J., and J. Y. Campbell. 2009. The long-run risks model and aggregate asset prices: an empirical assessment. Tech. rep., National Bureau of Economic Research.

Bekaert, G., and J. Liu. 2004. Conditioning Information and Variance Bounds on Pricing Kernels. *Review of Financial Studies* 17:339–378.

Bollerslev, T., G. Tauchen, and H. Zhou. 2009. Expected Stock Returns and Variance Risk Premia. *Review of Financial Studies* 22:4463–4492.

Borovicka, J., L. P. Hansen, M. Hendricks, and J. A. Scheinkman. 2011. Risk-Price Dynamics. *Journal of Financial Econometrics* 9:3–65.

Burnside, C. 1994. Hansen-Jagannathan Bounds as Classical Tests of Asset-Pricing Models. *Journal of Business and Economic Statistics* 12:pp. 57–79.

Campbell, J. Y. 1987. Stock returns and the term structure. *Journal of Financial Economics* 18:373 – 399.

Campbell, J. Y., and J. Ammer. 1993. What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns. *The Journal of Finance* 48:pp. 3–37.

Campbell, J. Y., and J. Cochrane. 1999. Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* 107:205–251.

Campbell, J. Y., and N. G. Mankiw. 1989. Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence. NBER Working Papers 2924, National Bureau of Economic Research, Inc. URL `http://ideas.repec.org/p/nbr/nberwo/2924.html`.

Campbell, J. Y., and R. J. Shiller. 1988a. The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *Review of Financial Studies* 1:195–228.

Campbell, J. Y., and R. J. Shiller. 1988b. Stock Prices, Earnings and Expected Dividends. *Journal of Finance* 43:661–676.

Campbell, J. Y., and R. J. Shiller. 1991. Yield Spreads and Interest Rate Movements: A Bird's Eye View. *Review of Economic Studies* 58:495–514.

Cecchetti, S., P.-s. Lam, and N. C. Mark. 1994. Testing Volatility Restrictions on Intertemporal Marginal Rates of Substitution Implied by Euler Equations and Asset Returns. *Journal of Finance* 49:123–52.

Cochrane, J. H. 2001. *Asset Pricing*. Princeton University Press.

Cochrane, J. H. 2008. The Dog That Did Not Bark: A Defense of Return Predictability. *Review of Financial Studies* 21:1533–1575.

Cochrane, J. H., and L. P. Hansen. 1992. Asset Pricing Explorations for Macroeconomics. NBER Working Papers 4088, National Bureau of Economic Research, Inc.

Cochrane, J. H., and M. Piazzesi. 2005. Bond Risk Premia. *American Economic Review* 95:138–160.

De Roon, F., and M. Szymanowska. 2012. Asset pricing restrictions on predictability: frictions matter. *Management Science* 58:1916–1932.

Fama, E. F., and R. R. Bliss. 1987. The information in long-maturity forward rates. *The American Economic Review* pp. 680–692.

Fama, E. F., and K. R. French. 1989. Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics* 25:23–49.

Ferson, W. E., and A. F. Siegel. 2003. Stochastic Discount Factor Bounds with Conditioning Information. *Review of Financial Studies* 16:567–595.

Ferson, W. E., and A. F. Siegel. 2009. Testing Portfolio Efficiency with Conditioning Information. *Review of Financial Studies* 22:2535–2558.

Gabaix, X. 2012. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics* 127:645–700.

Gallant, A. R., L. P. Hansen, and G. Tauchen. 1990. Using conditional moments of asset payoffs to infer the volatility of intertemporal marginal rates of substitution. *Journal of Econometrics* 45:141–179.

Ghosh, A., C. Julliard, and A. P. Taylor. 2011. What is the Consumption-CAPM missing? An informative-Theoretic Framework for the Analysis of Asset Pricing Models. Fmg discussion papers, Financial Markets Group.

Gourio, F. 2008a. Disasters and recoveries. *The American Economic Review* 98:68–73.

Gourio, F. 2008b. Time-series predictability in the disaster model. *Finance Research Letters* 5:191–203.

Goyal, A., and I. Welch. 2003. Predicting the Equity Premium with Dividend Ratios. *Management Science* 49:639–654.

Goyal, A., and I. Welch. 2008. A Comprehensive Look at The Empirical Performance of Equity Premium Prediction. *Review of Financial Studies* 21:1455–1508.

Hansen, L. P. 2012. Dynamic Valuation Decomposition Within Stochastic Economies. *Econometrica* 80:911–967.

Hansen, L. P., and R. Jagannathan. 1991. Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy* 99:225–62.

Hansen, L. P., and J. A. Scheinkman. 2009. Long-term Risk: An Operator Approach. *Econometrica* 77:177–234.

Kirby, C. 1998. The restrictions on predictability implied by rational asset pricing models. *Review of Financial Studies* 11:343–382.

Lazarus, E., D. J. Lewis, J. H. Stock, and M. W. Watson. 2016. HAR Inference: Kernel Choice, Size Distortions, and Power Losses. Tech. rep., NBER Discussion Papers.

Lettau, M., and S. Ludvigson. 2001. Consumption, Aggregate Wealth, and Expected Stock Returns. *Journal of Finance* 56:815–849.

Lewellen, J. 2004. Predicting returns with financial ratios. *Journal of Financial Economics* 74:209–235.

Lustig, H., S. Van Nieuwerburgh, and A. Verdelhan. 2013. The wealth-consumption ratio. *Review of Asset Pricing Studies* 3:38–94.

Maio, P. F. 2008. The FED model and expected asset returns. *Available at SSRN 889931* .

Mehra, R., and E. C. Prescott. 1985. The equity premium: A puzzle. *Journal of monetary Economics* 15:145–161.

Nakamura, E., J. Steinsson, R. Barro, and J. Ursúa. 2013. Crises and recoveries in an empirical model of consumption disasters. *American Economic Journal: Macroeconomics* 5:35–74.

Rietz, T. A. 1988. The equity risk premium a solution. *Journal of Monetary Economics* 22:117–131.

Schorfheide, F., D. Song, and A. Yaron. 2016. Identifying long-run risks: A bayesian mixed-frequency approach. Tech. rep., National Bureau of Economic Research.

Tsai, J. 2015. Rare disasters and the term structure of interest rates. Tech. rep.

Tsai, J., and J. A. Wachter. 2015. Disaster risk and its implications for asset pricing. Tech. rep.

Viceira, L. M. 2012. Bond Risk, Bond Return Volatility, and the Term Structure of Interest Rates. *International Journal of Forecasting* 28:97–117.

Wachter, J. A. 2006. A consumption-based model of the term structure of interest rates. *Journal of Financial Economics* 79:365–399.

Wachter, J. A. 2013a. Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility? *The Journal of Finance* 68:987–1035.

Wachter, J. A. 2013b. Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility? *The Journal of Finance* 68:987–1035.

Wilcox, D. W. 1992. The construction of US consumption data: some facts and their implications for empirical work. *The American economic review* pp. 922–941.

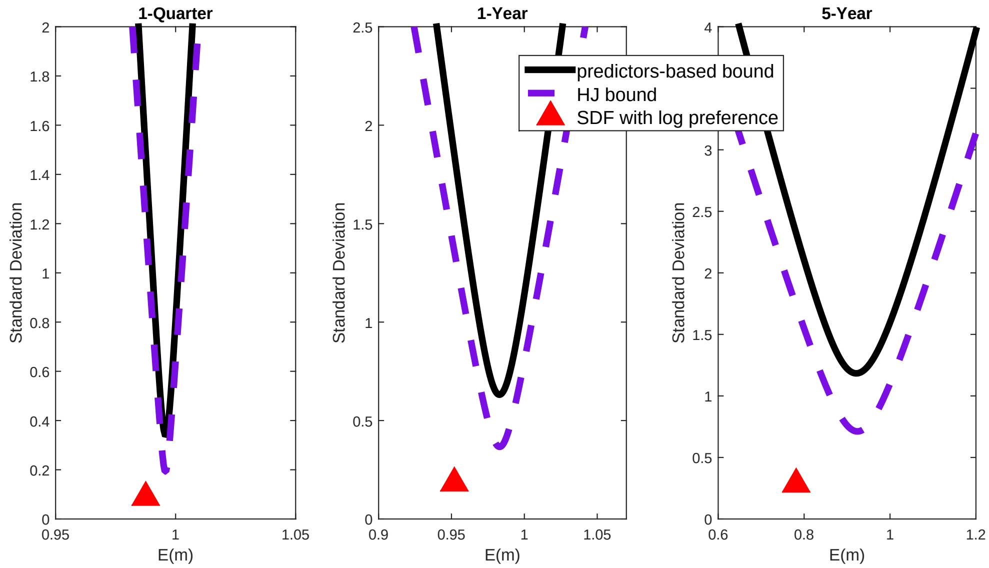

(a) SET A

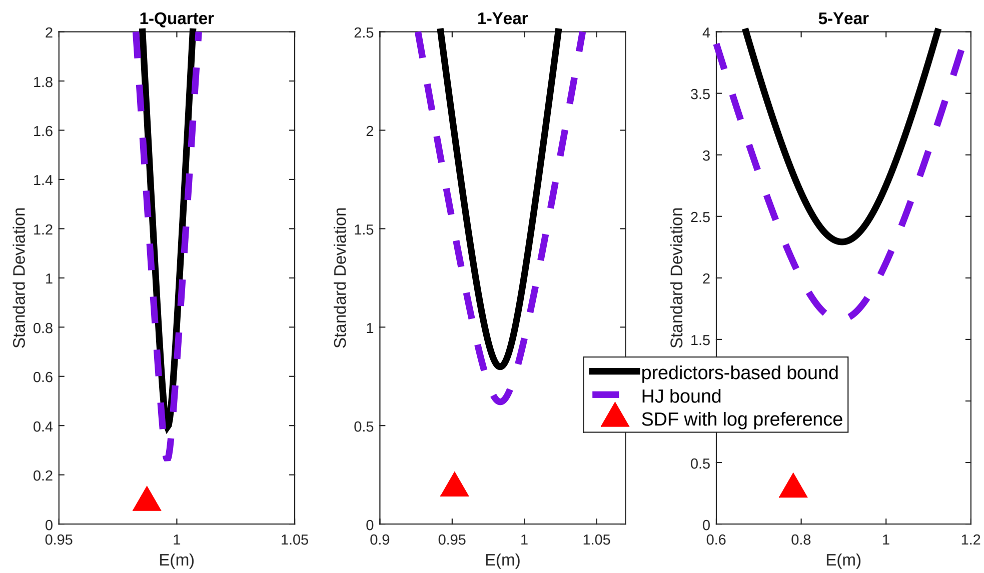

(b) SET B

Figure 1 Predictors-based bound $\sigma_Z^2(\nu)$, Hansen–Jagannathan (1991) bound, and model-implied SDFs across horizons. Dashed violet line gives the volatility bound when no conditional information is used. Solid black line gives the volatility bound using conditional information based on Bekaert and Liu (2004) specification. The red triangle reports the mean and standard deviation values of the SDF from a model in which the representative agent has log preference. In this example,the model-implied SDF, m_{t+h}^X, equals to $1/R_{t+h}^{MKT}$, where $1/R_{t+h}^{MKT}$ is the equity market return. The bounds are generated using asset returns included in SET A (see Panel A) and SET B (see Panel B). Sample: 1952Q2 - 2012Q3.

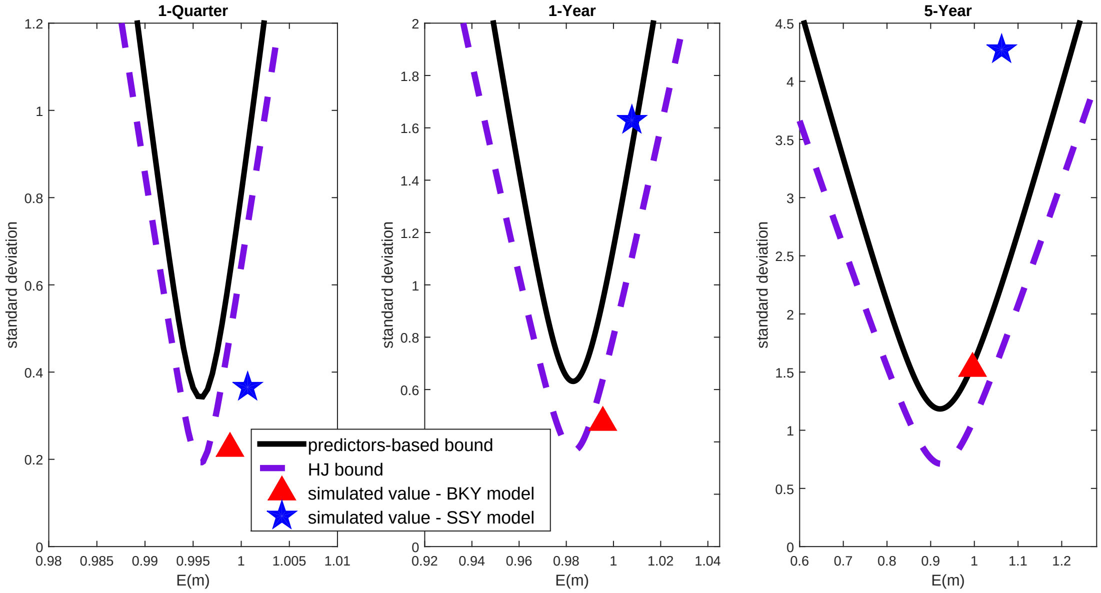

(a) Long-Run Risks Models

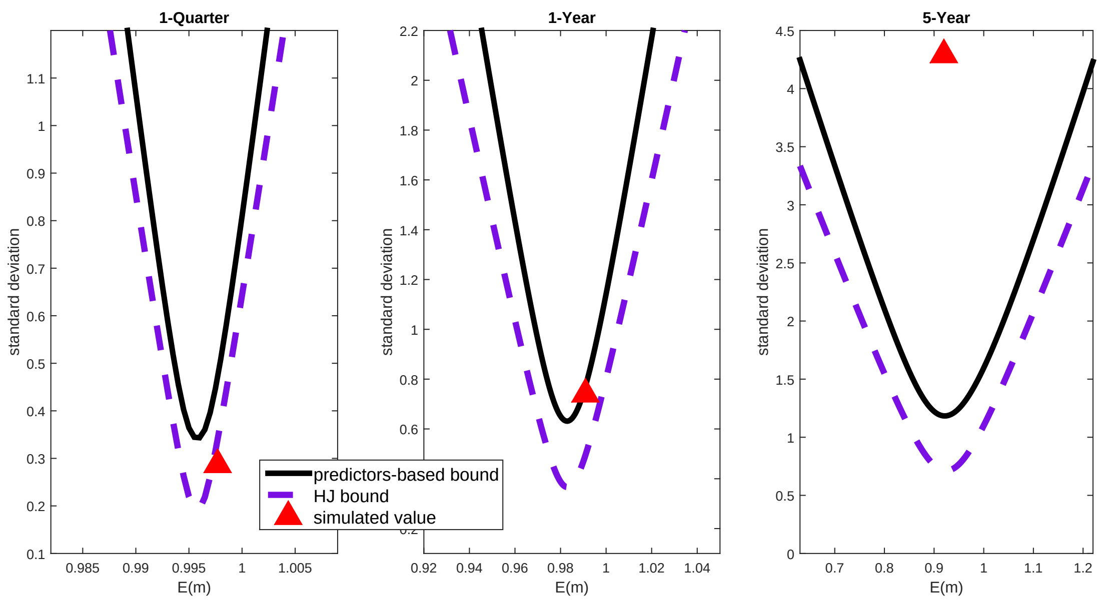

(b) External Habit Model

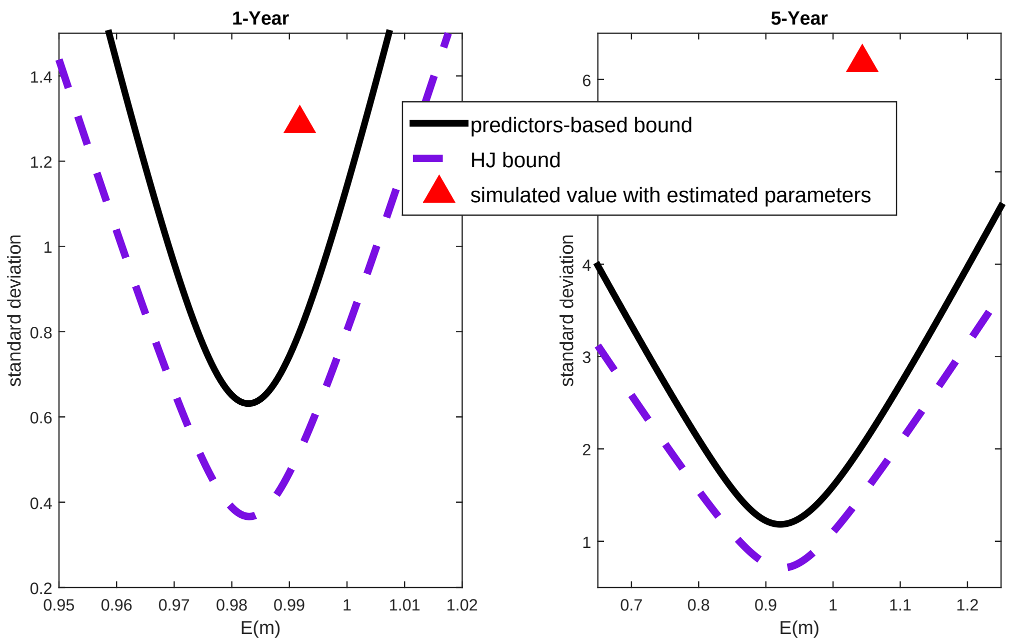

(c) Rare Disasters Model

Figure 2 Predictors-based bound $\sigma^2_Z(\nu)$, Hansen–Jagannathan (1991) bound, and model-implied SDFs across horizons– SET A. Dashed violet line gives the volatility bound when no conditional information is used. Solid black line gives the volatility bound using conditional information based on Bekaert and Liu (2004) specification. The red triangle and blue star (in Panel A) report average mean and standard deviation values from 10 simulations run of 600,000 months of the long-run risks, the external habit and the rare disasters model. In Panel A, we report the simulated values of the long-run risks model of Bansal et al. (2016) (BKY model, red triangle) and Schorfheide et al. (2016) (SSY model, blue star), respectively. The bounds are generated using asset returns included in SET A. Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3.

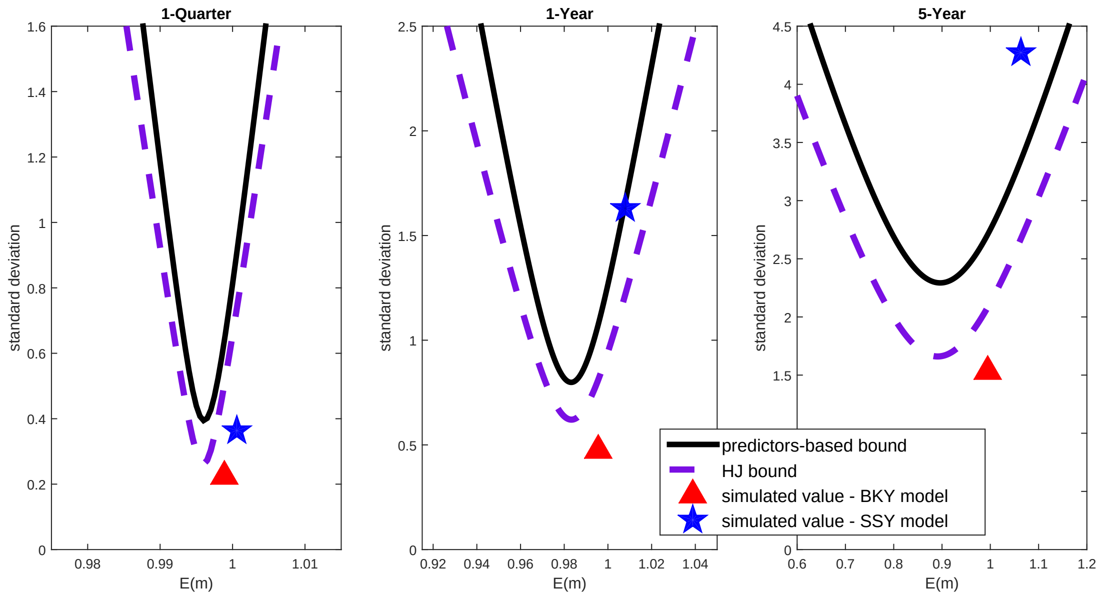

(a) Long-Run Risks Models

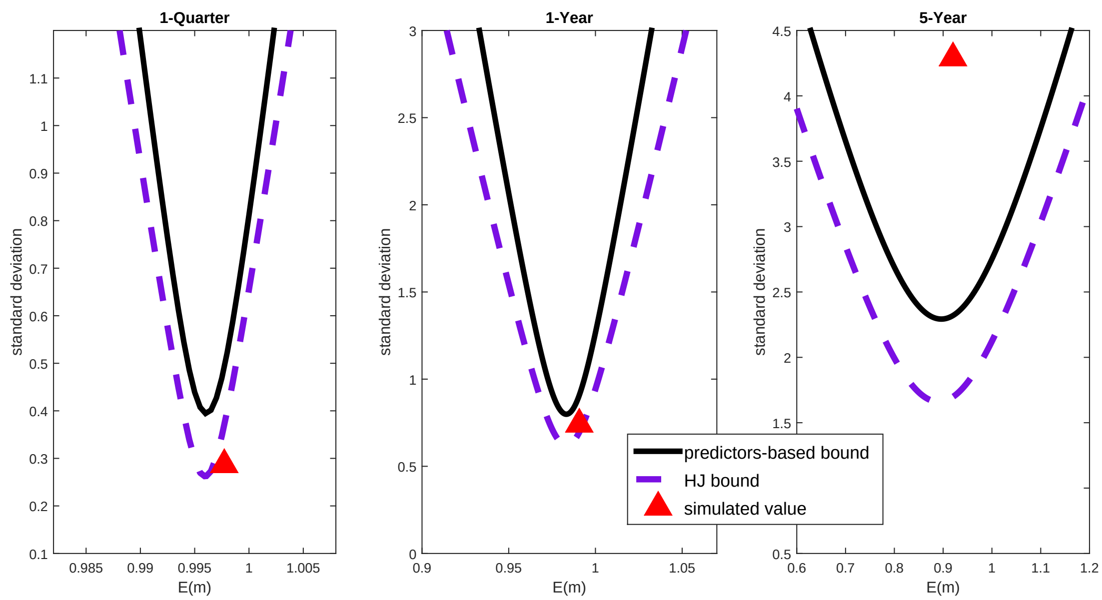

(b) External Habit Model

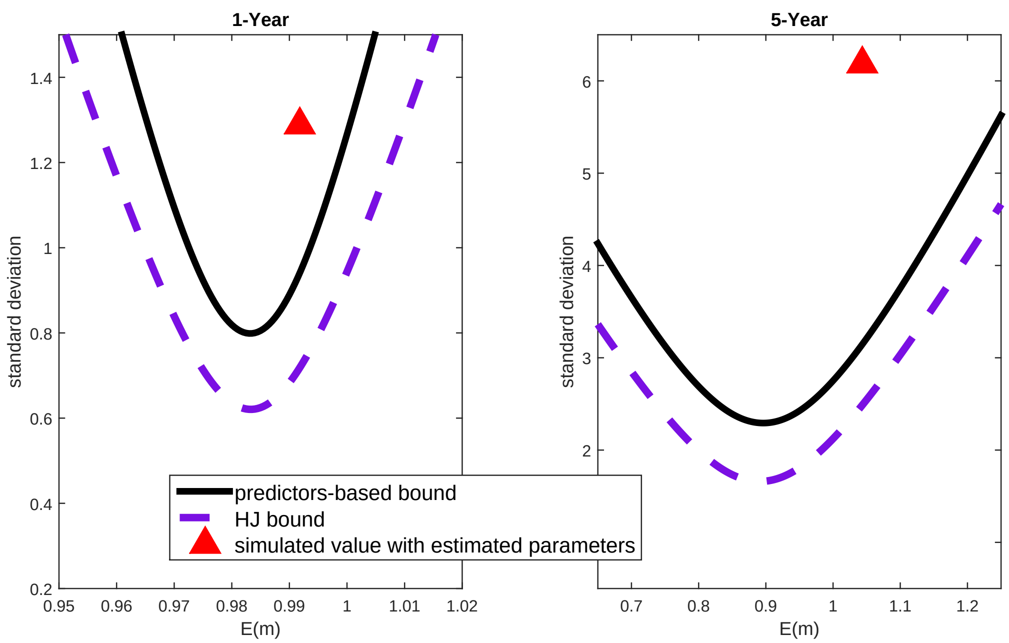

(c) Rare Disasters Model

Figure 3 Predictors-based bound $\sigma_Z^2(\nu)$, Hansen–Jagannathan (1991) bound, and model-implied SDFs across horizons – SET B. Dashed violet line gives the volatility bound when no conditional information is used. Solid black line gives the volatility bound using conditional information based on Bekaert and Liu (2004) specification. The red triangle and blue star (in Panel A) report average mean and standard deviation values from 10 simulations run of 600,000 months of the long-run risks, the external habit and the rare disasters model. In Panel A, we report the simulated values of the long-run risks model of Bansal et al. (2016) (BKY model, red triangle) and Schorfheide et al. (2016) (SSY model, blue star), respectively. The bounds are generated using asset returns included in SET B. Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3.

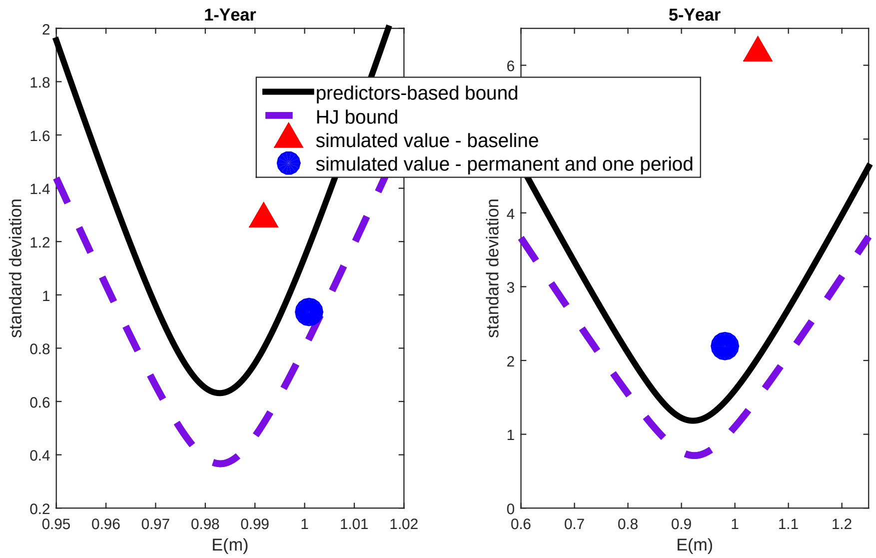

(a) SET A

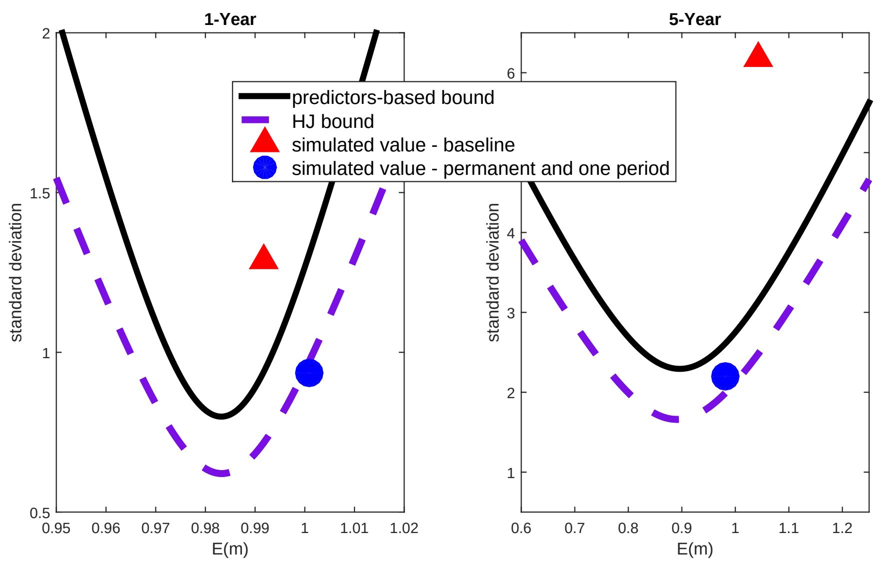

(b) SET B

Figure 4 Rare disaster Model-implied SDFs, predictors-based bounds and parameters uncertainty – SET A and SET B. Solid black line gives the volatility bound using conditional information based on Bekaert and Liu (2004) specification. The red triangle reports average mean and standard deviation values from 10 simulations run of 600,000 months with estimated parameters of the baseline model in Nakamura et al. (2013). The blue circle reports the same objects of the model with permanent, one-period disasters. Long horizon returns are computed by compounding quarterly returns. The bounds are generated using asset returns included in SET A (see Panel A) and SET B (see Panel B). Sample: 1952Q2 - 2012Q3.

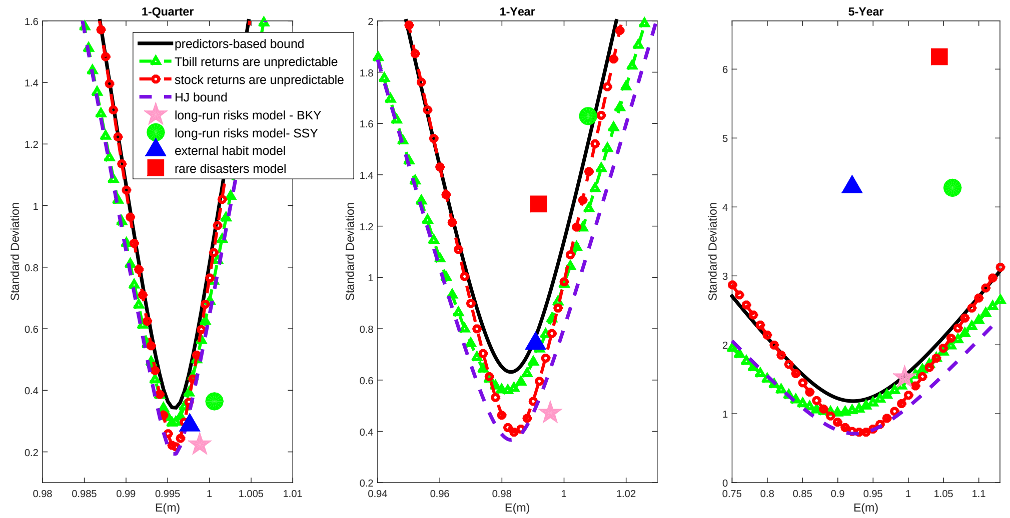

(a) SET A

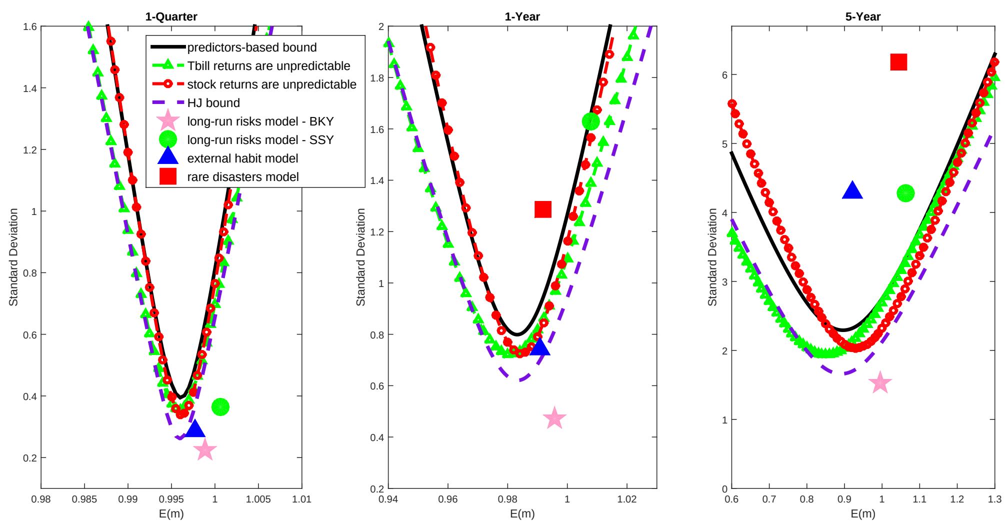

(b) SET B

Figure 5 Stock-based versus bonds-based variance bounds. Volatility bounds on stochastic discount factors for different investment horizons. Violet dashed line gives the volatility bound when no conditional information is used. Solid black line gives the Bekaert and Liu (2004) volatility bound using conditional information. Dashed red line with circles gives the volatility bound when imposing that stock returns are unpredictable. Dashed green line with triangles gives the volatility bound when imposing that roll over 3-month Treasury bill returns are unpredictable. The pink star and green circle report average mean and standard deviation values from 10 simulations run of 600,000 months with estimated parameters of the long run riks model by Bansal et al. (2016) and Schorfheide et al. (2016), respectively. The blue triangle and red square report the same objects computed of the external habit, and the rare disaster model, respectively. The bounds are generated using asset returns included in SET A (see Panel A) and SET B (see Panel B). Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3.

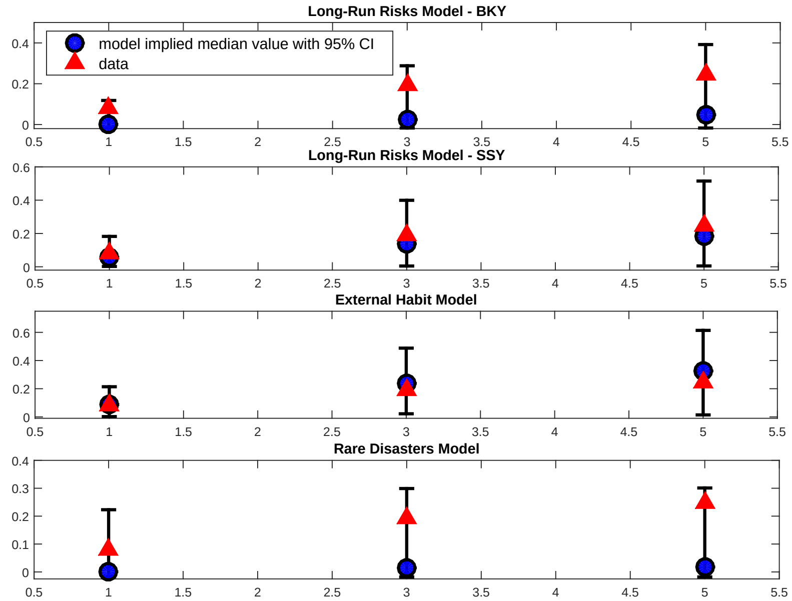

Figure 6 Historical versus model-implied predictability. Predictability of gross returns, and Price-Dividend Ratios. Red triangle presents the adjusted R^2s from projecting one-, three-, and five-year real gross return of the aggregate stock market portfolio onto lagged price-dividend ratio. The blue spot and black line provide the median values of adjusted R^2s from 1,000 simulations run of 724 monthly observations and associated 95% confidence interval. In each simulation, we project model simulated aggregate gross returns of stock market portfolio onto laggated model simulated price-dividend ratio. For the long-run risks type of model, we report the results of the model by Bansal et al. (2016)(BKY) and Schorfheide et al. (2016)(SSY), respectively. Sample: 1952Q2 - 2012Q3.

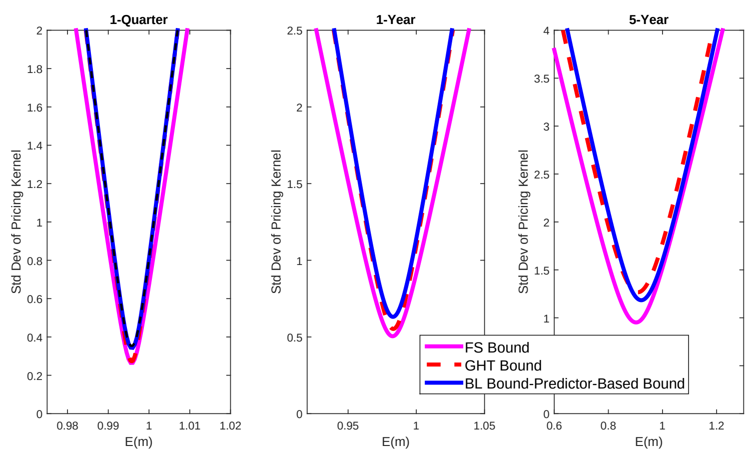

Figure 7 Alternative implementation of the HJ bounds – SET A. We present the volatility bounds using conditional information based on Ferson and Siegel (2003, 2009) (FS Bound), Bekaert and Liu (2004) (BL-predictors-based bound) and Gallant et al. (1990) (GHT Bound) specifications, respectively. The bounds are generated using asset returns included in SET A. Sample: 1952Q2 - 2012Q3.

	Asset	
	Stocks	Bonds
Mean return (% p.a.)	11.45	6.31
Standard deviation (% p.a.)	16.68	5.77

Table 1 Statistics of the Data. This table reports sample statistics of quarterly nominal stock and 5-year maturity US government bond total returns. Stock returns are nominal returns on the stock total returns on the value weighted portfolio of all stocks traded in the NYSE, the AMEX, and NASDAQ from CRSP. Bond returns are nominal returns on the 5-year constant maturity bond from the CRSP Fixed Term Indices File. Sample: 1952Q2: 2012Q3.

Panel A: Predictive regressions for stock returns

Horizon (quarters)	cay_t ($t-stat$)	dp_t ($t-stat$)	$R^2(\%)$
1	0.84	0.03	4.7
	(2.11)	(3.08)	
4	3.19	0.12	17.0
	(2.61)	(3.30)	
20	6.06	0.60	51.2
	(4.14)	(4.71)	

Panel B: Predictive regressions for bond returns

Horizon (quarters)	spr_t ($t-stat$)	y_t ($t-stat$)	$R^2(\%)$
1	7.77	0.75	8.3
	(8.82)	(3.67)	
4	22.49	3.28	19.8
	(6.73)	(3.45)	
20	34.90	19.44	43.0
	(2.39)	(4.23)	

Table 2 Predictability of stock and bond returns. Panel A reports quarterly overlapping regressions of multiple horizon real gross stock returns onto a constant, cay_t, and the log dividend-price ratio dp_t. Panel B reports monthly overlapping regressions of multiple horizon real gross return on a 5-year constant maturity coupon bond from CRSP onto a constant, the yield spread spr_t, and the short-term US treasury-bill rates y_t. The spread is the difference between the log yield on a 5-year artificial zero-coupon bond from the CRSP Fama-Bliss Discount Bond File, and the log yield on the one month maturity Treasury Bill (T-bill). The table reports coefficient estimates, the R^2 of the regression, and, Newey-West t-statistics in parentheses. We follow Lazarus et al. (2016) and we set the Newey-West truncation parameter to $S_T = \left(\frac{3}{2B}\right) T$, where $B = 8$. We evaluate the test statistic using critical values from the t_B. Critical values of Student's t-distribution with 8 degrees of freedom are 1.860, 2.306 and 3.355 at 10%, 5% and 1% significant level, respectively. Sample: 1952Q2: 2012Q3.

Panel A. Stock Market Index

Horizon (quarters)	Returns	Discounted Stock Returns				
		CRRA	Long-Run Risks BKY	Long-Run Risks SSY	External Habit	Rare Disasters
	F-stat	F-stat (p-value)	F-stat (p-value)	F-stat (p-value)	F-stat (p-value)	F-stat (p-value)
1	14.365*** (0.000)	6.576*** (0.004)	0.864 (0.460)	1.014 (0.401)	1.040 (0.392)	
4	13.223** (0.012)	18.909*** (0.002)	2.620 (0.401)	3.499 (0.294)	1.208 (0.656)	3.765 (0.269)
20	37.001** (0.033)	60.033*** (0.005)	0.737 (0.931)	7.623 (0.485)	4.981 (0.622)	3.031 (0.747)

Panel B. 5-Year Maturity Bond

Horizon (quarters)	Returns	Discounted Stock Returns				
		CRRA	Long-Run Risks BKY	Long-Run Risks SSY	External Habit	Rare Disasters
	F-stat	F-stat (p-value)	F-stat (p-value)	F-stat (p-value)	F-stat (p-value)	F-stat (p-value)
1	16.449*** (0.000)	7.559*** (0.001)	0.486 (0.607)	2.207* (0.069)	1.425 (0.241)	
4	16.210** (0.025)	17.518** (0.019)	3.247 (0.609)	3.249 (0.608)	2.399 (0.715)	0.257 (0.975)
20	19.440* (0.091)	35.506** (0.043)	0.395 (0.936)	1.533 (0.782)	2.461 (0.675)	0.189 (0.969)

Table 3 Predictability of raw and discounted stock and bond returns. Panel A reports the F-statistics from quarterly overlapping regressions of multiple horizon real raw (first column) and discounted (columns "Discounted Returns") stock returns onto a constant, the consumption-wealth ratio, cay_t, and the dividend-price ratio, dp_t. Panel B reports the F-statistics from quarterly overlapping regressions of multiple horizon real raw (first column) and discounted return (columns "Discounted Returns") on a 5-year constant maturity coupon bond from CRSP onto a constant, the lagged yield spread spr_t, and the short-term US treasury-bill rates y_t. To discount returns, we compute the $m^X_{t,t+h}$ from the consumption-based models with i.i.d consumption growth and CRRA utility, the long-run risks, the external habit, and the rare disasters. We use the realized consumption growth rate, formed by the real consumption of the nondurable goods and services, to compute $m^X_{t,t+h}$. For the long-run risks model of Bansal et al. (2016)(BKY) and the rare disasters model of Nakamura et al. (2013), we use the real stock returns on the value weighted portfolio of all stocks traded in the NYSE, AMEX, and NASDAQ from CRSP to approximate the consumption related asset returns. For the long run risk model of Schorfheide et al. (2016)(SSY), we use the median of the filtered series (kindly provided by the authors) to proxy for the latent states in the SDF. For the external habit model, we form the surplus consumption from quarterly real consumption of nondurable goods and services using the Campbell and Cochrane (1999) specification. We compute $m^X_{t,t+h}$ using estimated parameters. For any given set of computed $m^X_{t,t+h}$, we regress the discounted stock (bond) returns on the stock (bond) return predictors and compute the F-statistics of the regression. We report in parentheses the p-value of F-statistics by comparing with the bootstrapped empirical distribution. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. Sample: 1952Q2: 2012Q3.

$E[m]$	Fixed		Random		Random	
Bounds	Random		Fixed		Random	
	SET A	SET B	SET A	SET B	SET A	SET B
Horizon(quarters)	Long Run Risk Models – BKY					
1	0.041	0.063	−0.162	−0.216	−0.131	−0.162
4	0.555	0.461	−0.193	−0.346	−0.154	−0.278
20	1.907	1.058	−0.145	−1.193	−0.066	−0.696
	Long Run Risk Models – SSY					
1	−0.175	−0.150	−0.096	−0.150	−0.064	−0.093
4	0.664	0.579	−0.012	−0.166	0.028	−0.101
20	2.618	1.802	0.063	−0.981	0.142	−0.516
	External Habit Models					
1	−0.036	−0.019	−0.093	−0.148	−0.072	−0.113
4	0.457	0.349	−0.059	−0.219	−0.030	−0.158
20	3.028	2.178	0.174	−0.900	0.238	−0.496
	Rare Disasters Models					
4	0.821	0.691	0.327	0.168	0.389	0.168
20	4.250	3.655	2.344	1.170	2.480	1.170

Table 4 Distance between model-implied SDFs and the predictors-based bounds. The table displays the distance computed as the difference between the model implied standard deviation of the SDF and the volatility bound, $\Delta = \sigma\left(m_t^X\right) - \sigma_Z(v)$. A positive value means the model cannot be rejected at the five percent level. We consider three asset pricing models: the long-run risks, the external habit and the rare disasters models. For the long-risks type of model, we report the results of the model by Bansal et al. (2016)(BKY) and Schorfheide et al. (2016)(SSY), respectively. We compute the finite sample distribution of the distance as described in Appendix C. The distribution accounts for parameter and small-sample uncertainty. The parameter uncertainty reflects estimation uncertainty in both the asset pricing model and the return predictive regressions. For all cases, we present the 95% quantile of the distance. The table reports results for three different cases. The leftmost block is the case when the mean of the SDF, $E[m]$, is fixed but there is uncertainty surrounding the predictors-based bounds. In this case, the mean of the SDF is fixed to the long-run mean obtained from a simulations run of 600,000 months. In the second case reported in the middle block, when $E[m]$ is random but there is no parameter uncertainty in the predictors-based bounds. The rightmost block reports the results when both $E[m]$ and the predictors-based bounds are random. For further description of the algorithm used . Sample: 1952Q2: 2012Q3.