PhD THESIS DECLARATION

The undersigned

SURNAME *YANG*
FIRST NAME *HAOXI*
PhD Registration Number *1465715*

Thesis title: *Essays in Asset Pricing*

PhD in *Finance*
Cycle *25*
Candidate's tutor *Professor Fulvio Ortu*
Year of thesis defence *2015*

DECLARES

Under *her* responsibility:

1) that, according to Italian Republic Presidential Decree no. 445, 28th December 2000, mendacious declarations, falsifying records and the use of false records are punishable under the Italian penal code and related special laws. Should any of the above prove true, all benefits included in this declaration and those of the temporary embargo are automatically forfeited from the beginning;

2) that the University has the obligation, according to art. 6, par. 11, Ministerial Decree no. 224, 30th April 1999, to keep a copy of the thesis on deposit at the "Bilioteche Nazionali Centrali" (Italian National Libraries) in Rome and Florence, where consultation will be permitted, unless there is a temporary embargo protecting the rights of external bodies and the industrial/commercial exploitation of the thesis;

3) that the Bocconi Library will file the thesis in its “Archivio istituzionale ad accesso aperto” (institutional registry) which permits online consultation of the complete text (except in cases of a temporary embargo);

4) that, in order to file the thesis at the Bocconi Library, the University requires that the thesis be submitted online by the candidate in unalterable format to Società NORMADEC (acting on behalf of the University), and that NORMADEC will indicate in each footnote the following information:

   - thesis *Essays in Asset Pricing*;
   - by *YANG HAOXI*;
   - defended at Università Commerciale “Luigi Bocconi” – Milano in *2015*;
   - the thesis is protected by the regulations governing copyright (Italian law no. 633, 22$^{\text{th}}$ April 1941 and subsequent modifications). The exception is the right of Università Commerciale “Luigi Bocconi” to reproduce the same for research and teaching purposes, quoting the source;

5) that the copy of the thesis submitted online to NORMADEC is identical to the copies handed in/sent to the members of the Thesis Board and to any other paper or digital copy deposited at the University offices, and, as a consequence, the University is absolved from any responsibility regarding errors, inaccuracy or omissions in the contents of the thesis;

6) that the contents and organization of the thesis is an original work carried out by the undersigned and does not in any way compromise the rights of third parties (Italian law, no. 633, 22nd April 1941 and subsequent integrations and modifications), including those regarding security of personal details; therefore the University is in any case absolved from any responsibility whatsoever, civil, administrative or penal, and shall be exempt from any requests or claims from third parties;

7) that the PhD thesis is not the result of work included in the regulations governing industrial property, was not produced as part of projects financed by public or private bodies with restrictions on the diffusion of the results, and is not subject to patent or protection registrations, and therefore not subject to an embargo;

Date *January 30, 2015*

SURNAME *YANG*
FIRST NAME *HAOXI*

Contents

List of Figures

List of Tables

Acknowledgements

During these six years, there are many people who supported and encouraged me throughout the journey. I would like to take this opportunity to thank all of them who helped me gain this invaluable experience.

I would like to express my special appreciation and thanks to my supervisors Fulvio Ortu and Carlo Favero. They make a great team when supervising, each of them taught me very different, but equally important aspects of doing research. I would like to thank them for encouraging my research and for allowing me to grow as a research scientist. Meanwhile, I want to thank the other member of my committee, Lorenzo Garlappi. He was my host when I visited Sauder Business School at UBC for a semester. That was a fantastic experience. Lorenzo gave me feedback on my papers, presentations, and introduced me to a lot of people in the field.

I would also like to express my gratitude toward all the other faculty members at Bocconi Uniersity. They are always patient with all my questions and give me their brilliant comments and suggestions.

Furthermore, I want to thank my co-author Arie Guzluklu and Andrea Tamoni. As senior Bocconi graduates, they are my role model. From them, I have learnt how to become a independent research.

I also want to devote some extra attention to all my lovely friends and colleges. They are always there to support and help me unconditionally. Without them, it would be impossible for me to reach this destination.

Last but not least, I would like to give special thanks to my family. Words can not express how grateful I am to my parents for all of the sacrifices that you have made on my behalf. Thank you for supporting me for everything, and especially I can not thank you enough for encouraging me throughout this experience.

Haoxi Yang
Milan, January 2015

Abstract

The main topic of this thesis is about implications of return predictability in asset pricing. This question is specified into three distinct chapters. The first chapter "Implications of Returns Predictability across Horizons for Asset Pricing Models", which is coauthored with Carlo Favero, Fulvio Ortu and Andrea Tamoni, analyzes 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. For an asset pricing model identified by its state variables, information structure and model SDF, we supply a sufficient condition under which our predictors-based bounds constitute legitimate lower bounds on the variance of the SDF of the model. Using predictors-based bounds, we analyze discount factors produced by the long-run risk, the habit and the rare disasters models. We document that consumption-based asset pricing models such as long-run risk and habit models do not produce SDFs volatile enough at the one-year horizon. The rare disasters model satisfies our predictors-based bounds at each horizon. The second chapter "A Robust Variance Bound of Pricing Kernels" proposes a data-based measure of model performance to discriminate among competing asset pricing models of return predictability. I form a set of variance bounds on pricing kernels based on different systems for predicting asset returns. For a given asset pricing model, I define the robust variance bound to be the tightest variance bound that this model-implied pricing kernel is able to satisfy. Using the diagnostic results of the robust variance bounds, I then construct a model performance index. This index quantifies the degree of return predictability which a given asset pricing model is able to obtain. I apply this method to examine the performance of three leading classes of asset pricing models: long run risk, external habit and rare disasters. The long run risk type of rare disaster model of Nakamura et al. (2013) performs best. The final chapter "Demographics and the Behavior of Interest Rates", which is coauthored with Carlo Favero and Arie Guzluklu, relates the common persistent component of the US term structure of interest rates to the age composition of population. Interest rates are very persistent. Modelling the persistent component of interest rates has important consequences for forecasting. Demographics determines the equilibrium rate in the monetary policy rule and therefore the persistent component in one-period yields. Fluctuations in demographics are then transmitted to the whole term structure via the expected policy rate components. We build an affine term structure model (ATSM) which exploits demographic information to capture the dynamics of yields and produce useful forecasts of bond yields and excess returns that provides economic value for long-term investors.

Chapter 1

Implications of Returns Predictability across Horizons for Asset Pricing Models

1.1 Introduction

If there is valuable information for predicting stock and bond prices over time, and the more so the longer the horizon, when and how can we use this information to discriminate among competing asset pricing models?[1] The answer we give in this chapter is both methodological and empirical. From a methodological point of view, we offer a simple condition under which a variance bound that incorporates conditioning information from a given a set of stock and bond predictors constitutes a legitimate lower bound on the variance of the Stochastic Discount Factor (SDF) of a given asset pricing model. From an empirical point of view, we examine three leading classes of asset pricing models: external habit formation (Abel [1990], Campbell and Cochrane [1999]), rare disaster (Rietz [1988], Barro [2006a], Nakamura et al. [2013]), and long-run risk (Bansal and Yaron [2004], Bansal et al. [2012a]). We show that all these three models satisfy our condition under which the predictors-based bounds to be a legitimate lower bounds on the variance of their SDFs, and we use these predictors-based bounds to explore the role that short and long horizon predictability plays in the econometric evaluation of these models.

The importance of our methodological contribution is based on the fact that, in practice, the most successful predictors for stock and bond returns follow from accounting identities rather than from specific models of asset pricing (see Campbell [2003]). The predictive relation between the (log) price-dividend ratio and stock market returns, for

[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

instance, follows from solving forward a linearized expression for returns (Campbell and Shiller [1988b]). Similarly, the predictive relation between the consumption-wealth ratio, the *cay* variable of Lettau and Ludvigson [2001], and stock market returns follows from a linearized version of the consumer's intertemporal budget constraint. Likewise, a successful bond market predictors such as the term spread follows from a linearization argument similar to the one that generates the dynamic dividend growth model for the stock market (see Campbell and Ammer [1993]). If the question is to see whether a given asset pricing model is able to generate a sufficiently volatile discount factor, one needs to bridge the information contained in a given set of predictors with the informational content of the state variables of the model under scrutiny. In fact, the variability of the discount factor of a given model is conditional on some model specific state variables that are in general different from the variables used in the predictive regressions. As a consequence, the possible discrepancy between the informational content of the predictors and that of the state variables may render the empirical evidence on predictability on one side not informative to reject one models, and on the other side decisive to draw inference on another model.

We illustrate in Proposition 1 our condition for the variance of a model SDF to satisfy a set of Hansen and Jagannathan [1991] variance bounds extended to incorporate conditioning information from a set of predictors. Our condition is on the returns discounted by the model SDF, and it requires the predictability of the discounted returns not to increase when the information in the predictors is added to the information in the state variables of the model. Equivalently, if the model SDF fails to achieve the variance threshold dictated by our predictors-based bound, then the discounted returns on some assets must become predictable when the information in the state variables is augmented with the information in the predictors.

In the empirical part of the paper we first provide robust evidence that predictability translates into tighter bounds on the variance of the SDFs. To this end we employ a linear predictive model to compute the first two conditional moments of asset returns. We then use these moments to compute the variance bounds based on conditioning information, see Gallant et al. [1990]. 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 conditional moments are known.

Our next step is to analyze three models representative of the classes of asset pricing models discussed above: the habit-formation model of Campbell and Cochrane [1999] as estimated by Aldrich and Gallant [2011a], the rare disaster model of Nakamura et al. [2013], and the long-run risk model of Bansal et al. [2012a]. There are two main reasons behind our choice. First, the fact that these three are all estimated models allows us to evaluate the effect of estimation uncertainty on the moments of a model-implied SDF. Second, and most importantly, we want to explicitly abstain from considering models that are specifically tailored to produce predictability. One of our results, in fact, is that although the Nakamura et al. [2013] model does not deliver as much predictability as there is in the data, yet it performs much better than the other two models at the light of our predictors-based bounds. This fact shows that, while predictability in the data is key to have sharp predictors-based bounds, the amount of predictability produced by a model is not a necessary condition for that model to satisfy a variance bound that incorporates information from a set of predictors.

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 a low frequency component that should make asset pricing puzzles less pronounced at longer horizons.[2] Consistent with these statements, the conclusion drawn from the standard unconditional Hansen and Jagannathan [1991] bounds is not surprising: all models satisfy these unconditional cups at medium and long horizons. To see if these conclusions are robust to the inclusion of information contained in our standard predictors we proceed 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 condition under which our predictors-based bounds constitute legitimate lower bounds on the variance of the SDF of the model. We show that, for all horizons and sets of assets, the condition cannot be rejected for either one of the models under scrutiny. Therefore our predictors-based bounds are legitimate lower bounds for the volatilities of the SDFs of the three models. Three main conclusions emerge from our horse race. First, the rare disaster model is by far the best performer, since it satisfies our predictors-based bounds across all horizons and for both sets of assets. Second, the long-run risk model is the one most challenged by the introduction of conditioning information: it basically struggles to meet the bound at all horizon and for both sets of assets, with the only possible exception the 5-years horizon and when the set of traded assets includes only equities and T-Bills. Finally, the habit model lies somehow in the middle: it performs quite well at all horizons when the set of assets includes only equities and T-Bills, but struggles to meet the predictors based bound at the quarterly and yearly horizon when the Treasury Bonds are added.

These results show the importance of understanding when and 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

[2]The Nakamura et al. [2013], differently from the Rietz-Barro model of permanent and instantaneous disasters, accounts for the partially transitory nature of disasters, and the fact that they unfold over multiple years: these two effects generate variation in expected consumption growth surrounding disasters.

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, our bounds are simultaneously able to (a) tell the habit model apart from the long-run risk model at the 1-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 when Treasury Bonds are added; (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.

Does the performance of an asset pricing model measured with the yardstick of our predictors-based bound translate one-to-one into the ability of the model to reproduce the return predictability observed in the data? To investigate this issue we run predictive regressions of future gross real returns on the price-dividend ratio over the 1, 3, 5 and 8-years horizons, with both actual and simulated data. Two main results emerge. First, since the R^2s in the data are within the 95% model-based confidence intervals generated by simulated regressions using alternatively all the three models under scrutiny, the evidence suggests that all these three models can, in principle, account for the returns predictability observed in the data. Second, and most importantly, while the median of the simulated R^2s for the habit model is able to basically match the data R^2, the median of the simulated R^2s for the rare disaster model is well below the data R^2. Recalling that the rare disaster model satisfies our predictors-based bounds across all horizons and for both sets of assets, while the habit model exhibits several violations of the bounds, we can conclude that the ability of a model to produce predictable returns is neither necessary nor sufficient to stand up to variance bounds that incorporate the actual predictability in the data. From this standpoint, our predictors-based bounds emerge as a more effective tool than model-generated predictability to discriminate among asset pricing models when returns are predictable.

Our work is related to Cochrane and Hansen [1992], who have been the first to look systematically at HJ 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*.

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 et al. [2011]). In particular we view our predictors-based bounds as a useful tool for understanding the high- and low-frequency components of such models. Finally, our work is also related to the recent information-theoretic literature that uses entropy bounds to restrict the admissible regions for the SDF and its components (see Bakshi and Chabi-Yo [2012] and Ghosh et al. [2011]). In particular our conclusions are in line with Backus et al. [2011b] who show that the entropy of a model should be sufficiently large to account for observed excess returns.

The rest of this chapter is organized as follows. Sections 2 introduces our predictors-based variance bounds and provides the condition under which these bounds are indeed legitimate lower bounds on the variance of the SDF of a given asset pricing model. 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 our bounds at different horizons. We then assess whether various SDF specifications are consistent with our predictors-based bounds. We conclude by discussing how the comparison of model-implied return predictability versus the historical one is connected to our predictors-based variance 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.

1.2 Variance Bounds, Predictability and Asset Pricing

In this section we first define our predictors-based variance bounds, which are bounds on the variance of the SDFs that price a given set of returns conditional on the information contained in a vector Z_t of return's predictors. Given then any asset pricing model with

SDF m^X_{t+h}, we ask: under what conditions does a predictors-based bound constitute a legitimate lower bound on the variance of m^X_{t+h}? We answer this question by identifying in Proposition 1 a simple condition under which the variance of m^X_{t+h} must indeed satisfy the bound obtained by conditioning on the predictors Z_t. In Proposition 2, moreover, we rephrase our sufficient condition in terms of an upper bound on the $\mathcal{R}^2$ from predictive regressions of future returns on the current values of the predictors.

1.2.1 Variance bounds when returns are predictable

We consider a set of N assets traded at a given time t, we denote the return on each asset by $R_{j,t+h}$, with $h = 1, 2...$ the investment horizon, and we let R_{t+h} denote the vector collecting these N returns. Alongside the returns we consider a vector Z_t of return's predictors, and we denote with $\mathcal{F}^Z_t$ the informational content of these predictors. By saying that Z_t predicts the return $R_{j,t+h}$ on some asset j we mean $Var\left[E\left(R_{j,t+h}\left|\mathcal{F}^Z_t\right.\right)\right] > 0$ over some holding period h.[3]

We denote with $\mathcal{M}^Z$ the set of SDFs that price returns conditionally on the realizations of the predictors Z_t, that is

$$\mathcal{M}^Z = \left\{m_{t+h} \;\middle|\; E(m^2_{t+h}) < \infty,\ E\left(m_{t+h}R_{t+h}\left|\mathcal{F}^Z_t\right.\right) = e\right\} \tag{1.1}$$

where e denotes the unit vector. We assume $\mathcal{M}^Z$ non-empty, which means that the Law of One Price holds in the linear space of payoffs obtained by managed portfolios that condition on the predictors' realization.

Given the SDFs in $\mathcal{M}^Z$, we call predictors-based variance bound, denoted with $\sigma^2_Z(v)$, the lower envelope of the set of all variances of SDFs in $\mathcal{M}^Z$, that is the map

$$\sigma^2_Z(v) = \inf\left\{Var\left(m_{t+h}\right) \;\middle|\; m_{t+h} \in \mathcal{M}^Z,\ E\left(m\right) = v,\ v \in \Re\right\} \tag{1.2}$$

where the variable v is the shadow price of a unit risk-free zero-coupon bond with maturity $t+h$. The parabolic function $\sigma^2_Z(v)$ represents an unconditional frontier for SDFs in the sense of Gallant et al. [1990]: since it considers all SDFs that price returns conditionally on Z_t it takes full advantage of the predictive power of the vector Z_t while maintaining the simplicity of concentrating on the unconditional moments of such SDFs.

[3]We assume that all the random variables are defined over a probability space $(\Omega, \mathcal{F}, \mathcal{P})$, so that $\mathcal{F}^Z_t$ is formally the $\sigma-$algebra generated by the vector of predictors Z_t, and hence $\mathcal{F}^Z_t \subset \mathcal{F}$. We also assume that returns have finite unconditional first and second moments, μ and $E^2 \triangleq E\left[R_{t+h}R'_{t+h}\right]$, with unconditional variance-covariance matrix $\Sigma = E^2 - \mu\mu'$. Moreover, we assume that the matrix $E^2_t \triangleq E\left[R_{t+h}R'_{t+h}\left|\mathcal{F}^Z_t\right.\right]$ of second moments conditional on the predictors is (almost surely) positive-definite and hence invertible, i.e. returns are linearly independent conditionally on information in the predictors. Therefore, denoting with $\mu_t \triangleq E\left[R_{t+h}\left|\mathcal{F}^Z_t\right.\right]$ the vector of conditional expected returns, both the conditional variance-covariance matrix $\Sigma_t = E^2_t - \mu_t\mu'_t$ and its unconditional counterpart Σ are positive-definite (and hence invertible) as well.

As observed by Bekaert and Liu [2004], when the conditional moments of returns are not correctly specified the predictors-based bound $\sigma_Z^2(v)$ may fail to be a valid lower bound for the volatility of SDFs in $\mathcal{M}^Z$. To obviate to this problem, one can extend to this conditional setting the duality between mean-variance frontiers for SDFs and maximum Sharpe ratios first illustrated for the unconditional case by Hansen and Jagannathan [1991] in their seminal work. More specifically, define the following set of returns from managed portfolios

$$\mathcal{R}^Z = \left\{ R_{t+h}^w \;\middle|\; R_{t+h}^w = w_t' R_{t+h},\ w_t\ \mathcal{F}_t^Z - measurable \text{ s.t. } E\left(w_t' e\right) = 1 \right\} \tag{1.3}$$

This set collects all the payoffs that are generated by trading strategies that exploit the information contained in the predictors at time t. As long as $\nu \neq 0$ one can show (see Bekaert and Liu [2004], Abhyankar et al. [2007] and Peñaranda and Sentana [2013]) that[4]

$$\sigma_Z^2(v) = \nu^2 \sup_{R_{t+h}^w \in \mathcal{R}^Z} \left(\frac{E(R_{t+h}^w) - \nu^{-1}}{Var\left(R_{t+h}^w\right)} \right)^2 \tag{1.4}$$

In words, for any given level of the risk-free rate the predictors-based bound is proportional to the square of the maximum Sharpe ratio that can be generated by managed portfolios that exploit the information contained in the predictors Z_t. Observing that a mis-specification of the conditional expected returns and variances introduces a duality gap in (2.4), Bekaert and Liu [2004] suggest to always use the right-hand side to actually compute a variance bound that incorporates conditioning information since, by the very own definition of sup, this right hand side will always constitute a valid lower bound on the variance of the SDFs in $\mathcal{M}^Z$ (albeit, not necessarily the highest lower bound if mis-specification of the first two conditional moments of returns actually occurs). This is why in the empirical part of this chapter we estimate our predictors-based bounds using the solution to the dual problem defined in Eq.(2.4).

Before proceeding, we remark that from the horizon perspective our predictors-based bounds are conservative since they are based on those SDFs that price trading strategies which, although they take full advantage of the information in the predictors Z_t, still are required to be "buy and hold" over the horizon h. The predictors-based bounds that would be obtained from those SDFs that price the (truly) dynamic trading strategies that are allowed to be rebalanced at the intermediate dates $t+1$, $t+2, \ldots . t+h$ would clearly impose a much harder yardstick. In fact, allowing for intertemporal rebalancing would expand the set $\mathcal{R}^Z$ of managed returns and, via the duality in (2.4), that would yield a

[4]When $\nu = 0$ the sup coincides with the reciprocal of the global minimum portfolio variance over the set $\mathcal{R}^Z$. Moreover, the sup is always attained with the only exception of the case in which ν is set equal to the expected return on the global minimum variance portfolio, case in which the sup is attained by a return whose expected price is zero.

much higher bound $\sigma^2_Z(v)$. In this chapter, however, we concentrate on "buy and hold" strategies and leave the more general framework to future research.

1.2.2 Predictors-based bounds and asset pricing modelling

Let's consider now the asset pricing modelling side of our argument. Our main interest is to understand when and how we can employ information contained in the set of predictors Z_t, and synthesized in the predictors-based bound $\sigma^2_Z(v)$, to evaluate a given asset pricing model. To formalize our discussion, we identify any given asset pricing model with the triple $\left(X_t, \mathcal{F}^X_t, m^X_{t+h}\right)$ where X_t denotes the set of state variables of the model, $\mathcal{F}^X_t$ denotes the informational content in state variables X_t,[5] and m^X_{t+h} denotes the SDF of the given asset pricing model. Since from the standpoint of a given asset pricing model agents maximize their utility based on the information contained in the state variables X_t, the SDF m^X_{t+h} together with the returns R_{t+h} must satisfy the first order condition

$$E\left(m^X_{t+h} R_{t+h} \middle| \mathcal{F}^X_t\right) = e \tag{1.5}$$

More generally, the SDF m^X_{t+h} must price all the managed portfolio that condition on the state variables X_t of the given asset pricing model.

To help the intuition it is useful to exemplify this general framework with the three asset pricing models that we analyze in the empirical part. The first example is the Bansal et al. [2012a] model of long-run risk, where the state variables are the first two conditional moments of log consumption growth g_t, that is $X_t = (x_t, \sigma^2_t)$, the information $\mathcal{F}^X_t$ is generated by the innovations in these first two conditional moments, and the SDF takes the form

$$\ln\left(m^X_{t+h}\right) = A + Bg_{t+h} + Cr_{a,t+h}$$

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. A second example 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 $\mathcal{F}^X_t$ is generated by the innovations in surplus consumption ratio, and the SDFs takes the form

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

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

[5]Formally, $\mathcal{F}^X_t$ is the $\sigma-$algebra generated by the vector of state variables X_t, and hence $\mathcal{F}^X_t \subset \mathcal{F}$ where $\mathcal{F}$ is the $\sigma-$algebra of the probability space $(\Omega, \mathcal{F}, \mathcal{P})$ over which all the random variables are defined.

The last example 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 $X_t = (I_t, z_t)$, the information F_t^X is generated by state variables, and the SDF takes the same functional form as in the long-run risk model.

The question we want to address is: under what conditions does the predictors-based bound $\sigma_Z^2(v)$ constitutes a *legitimate* lower bound on the variance of the SDF of a given asset pricing model? To address this question, given an asset pricing model $\left(X_t, \mathcal{F}_t^X, m_{t+h}^X\right)$ we denote with $\mathcal{F}_t^{X,Z}$ the information set obtained by adjoining to the vector of state variables X_t the vector Z_t of predictors.[6] Moreover, we let $\nu^X \equiv E\left(m_{t+h}^X\right)$ denote the price assigned by the SDF m_{t+h}^X to a unit zero-coupon bond with maturity $t+h$. With this notation at hand we are now able to state the following sufficient condition for the predictors-based bound σ_Z^2 to constitute a legitimate volatility bound for m_{t+h}^X.

Proposition 1.2.1 *Suppose that the asset pricing model* $\left(X_t, \mathcal{F}_t^X, m_{t+h}^X\right)$ *satisfies*

$$E\left(m_{t+h}^X R_{t+h} \middle| \mathcal{F}_t^X\right) = E\left(m_{t+h}^X R_{t+h} \middle| \mathcal{F}_t^{X,Z}\right) \qquad (*)$$

Then the predictors-based frontier for SDFs $\sigma_Z^2(v)$ *constitutes a legitimate lower bound for the volatility of* m_{t+h}^X*, in the sense that*

$$Var\left(m_{t+h}^X\right) \geqslant \sigma_Z^2\left(\nu^X\right)$$

is a necessary condition for (1.5) to hold.

Proof. The iterative property of conditional expectation implies that

$$E\left(m_{t+h}^X R_{t+h} \middle| \mathcal{F}_t^Z\right) = E\left[E\left(m_{t+h}^X R_{t+h} \middle| \mathcal{F}_t^{X,Z}\right) \middle| \mathcal{F}_t^Z\right]$$

This, together with the orthogonality condition $E\left[m_{t+h}^X R_{t+h} \middle| \mathcal{F}_t^X\right] = e$ and (*) implies that

$$E\left(m_{t+h}^X R_{t+h} \middle| \mathcal{F}_t^Z\right) = e$$

that is $m_{t+h}^X \in \mathcal{M}^Z$, from which

$$Var\left(m_{t+h}^X\right) \geqslant \sigma_Z^2\left(\nu^X\right)$$

follows readily■

To better place our result in the vast literature on predictability and asset pricing observe that, from Kirby [1998] on, it is standard in that literature to assume that the

[6] Formally, $\mathcal{F}_t^{X,Z} = \sigma\left(\mathcal{F}_t^X \cup \mathcal{F}_t^Z\right)$, i.e. it is the smallest $\sigma-$algebra that contains all the information in X_t and Z_t, therefore $\mathcal{F}_t^{X,Z} \subset \mathcal{F}$.

predictors belong to a general information set $\mathcal{F}_t^I$ which investors condition upon when pricing assets. More formally, in the literature it is customary to concentrate on those SDFs m_{t+h} which satisfy

$$E\left(m_{t+h}R_{t+h}|\,\mathcal{F}_t^I\right) = e \tag{1.6}$$

for some information set $\mathcal{F}_t^I$ such that $\mathcal{F}_t^I \supset \mathcal{F}_t^Z$. This perspective is clearly useful to investigate if the ability of Z_t to predict a given set of return is correctly priced by *some* rational asset pricing model, since whenever $\mathcal{F}_t^I \supset \mathcal{F}_t^Z$ then any SDF that satisfies (1.6) must also price those dynamic strategies that condition on the predictors Z_t. If, however, one wants 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*, then the information sets $\mathcal{F}_t^Z$ and $\mathcal{F}_t^X$ must be taken as given, there is no guarantee that $\mathcal{F}_t^Z \subset \mathcal{F}_t^X$, and this is where our condition (*) finds its bite.[7]

Whenever condition (*) holds, therefore, $\sigma_Z^2(v)$ is a legitimate lower bound on the volatility of the SDF of the given asset pricing model. If $Var\left(m_{t+h}^X\right) < \sigma_Z^2\left(\nu^X\right)$ but condition (*) fails, however, we cannot reject out of hand the asset pricing model $\left(X_t, \mathcal{F}_t^X, m_{t+h}^X\right)$, since in that case the orthogonality condition $E\left(m_{t+h}^X R_{t+h}|\,\mathcal{F}_t^X\right) = e$ is in principle compatible with a volatility level lower that the one dictated by conditioning on the predictors Z_t. An alternative, but logically equivalent, way to express the implication in Proposition 1 is contained in the next result.

Corollary 1 *If, given the predictors-based bound $\sigma_Z^2(\nu)$, an asset pricing model $\left(X_t, \mathcal{F}_t^X, m_{t+h}^X\right)$ satisfies $E\left(m_{t+h}^X R_{t+h}|\,\mathcal{F}_t^X\right) = e$ and $Var\left(m_{t+h}^X\right) < \sigma_Z^2\left(\nu^X\right)$, then for some return $R_{j,t+h}$*

$$Var\left[E\left(m_{t+h}^X R_{j,t+h}|\,\mathcal{F}_t^{X,Z}\right)\right] > 0$$

Proof. By Proposition 1, if $E\left(m_{t+h}^X R_{t+h}|\,\mathcal{F}_t^X\right) = e$ and $Var\left(m_{t+h}^X\right) < \sigma_Z^2\left(\nu^X\right)$ then $E\left[m_{t+h}^X R_{j,t+h}|\,\mathcal{F}_t^{X,Z}\right] \neq 1$ for some return $R_{j,t+h}$, hence

$$Var\left(E\left[m_{t+h}^X R_{j,t+h}|\,\mathcal{F}_t^{X,Z}\right]\right) = Var\left(E\left[m_{t+h}^X R_{t+h}|\,\mathcal{F}_t^{X,Z}\right] - 1\right) > 0$$

■

Corollary 1 supplies a dual interpretation of Proposition 1 in terms of predictability of discounted returns. Given an asset pricing model $\left(X_t, \mathcal{F}_t^X, m_{t+h}^X\right)$, the discounted returns $m_{t+h}^X R_{t+h}$ cannot be predicted by the state variables X_t alone if the model satisfies the Euler equation. If the SDF m_{t+h}^X satisfies the conditional Euler equation and yet it fails to achieve the variance threshold dictated by the predictors-based bound $\sigma_Z^2\left(\nu^X\right)$, however,

[7]Observe however that the relation $\mathcal{F}_t^Z \subset \mathcal{F}_t^X$ would need to hold in equilibrium if the vector of predictors Z were constituted only by equilibrium quantities, since in that case in equilibrium Z ought to depend functionally on the state variables of the asset pricing model employed. From this standpoint, condition (*) represents a direct equilibrium restriction whenever Z is made of endogenous variables.

then the discounted return of some asset becomes predictable upon augmenting the state variables X_t with the predictors Z_t.

1.2.3 Predictors-based bounds and predictive $\mathcal{R}^2s$

We show now that the predictors-based bound $\sigma_Z^2(\nu)$ generates also an upper bound for the $\mathcal{R}^2s$ from regressions of the returns R_{t+h} on the predictors Z_t. When taken together with Proposition 1, this implies that the variance of the SDF of any asset pricing model $\left(X_t, \mathcal{F}_t^X, m_{t+h}^X\right)$ that satisfies condition (*) bounds from above these predictive $\mathcal{R}^2s$ as well. We establish these facts in the next Proposition, under the assumption that a risk-free return $R_{f,t+h}$ is available to the investors.

Proposition 1.2.2 *Given an asset pricing model* $\left(X_t, \mathcal{F}_t^X, m_{t+h}^X\right)$ *and the return* $R_{j,t+h}$ *on a traded asset, suppose that* $Var\left(R_{j,t+h}\middle| \mathcal{F}_t^Z\right)$ *is constant and condition (*) holds. Then*

$$\mathcal{R}_j^2 \equiv \frac{Var\left[E\left(R_{j,t+h}\middle| \mathcal{F}_t^Z\right)\right]}{Var\left(R_{j,t+h}\right)} \leqslant R_{f,t+h}^2 \sigma_Z^2\left(\nu_{\min}\right) \leqslant R_{f,t+h}^2 Var\left(m_{t+h}^X\right) \tag{1.7}$$

where $\sigma_Z^2\left(\nu_{\min}\right)$ *is the global minimum variance over all SDFs in* $\mathcal{M}^Z$*.*

Proof. Since under condition (*) the second inequality follows readily from Proposition 1, we only need to establish the first inequality. To this end, denoting with $R_{j,t+h}^e = R_{j,t+h} - R_{f,t+h}$ the excess return on asset j, for any $m_{t+h} \in \mathcal{M}^Z$ we have $E\left(m_{t+h}R_{j,t+h}^e\middle| \mathcal{F}_t^Z\right) = 0$, that is

$$E\left(R_{j,t+h}^e\middle| \mathcal{F}_t^Z\right) = -R_{f,t+h}cov\left(m_{t+h}, R_{j,t+h}^e\middle| \mathcal{F}_t^Z\right)$$

Squaring both sides up, exploiting the conditional Cauchy-Schwarz inequality, taking expectations and exploiting the fact that the variance cannot exceed the second moment, we have

$$\begin{aligned} Var\left[E\left(R_{j,t+h}\middle| \mathcal{F}_t^Z\right)\right] &= Var\left[E\left(R_{j,t+h}^e\middle| \mathcal{F}_t^Z\right)\right] \\ &\leq R_{f,t+h}^2 E\left[Var\left(R_{j,t+h}\middle| \mathcal{F}_t^Z\right) Var\left(m_{t+h}\middle| \mathcal{F}_t^Z\right)\right] \\ &\leq R_{f,t+h}^2 Var\left(R_{j,t+h}\right) E\left[Var\left(m_{t+h}\middle| \mathcal{F}_t^Z\right)\right] \\ &\leq R_{f,t+h}^2 Var\left(R_{j,t+h}\right) Var\left(m_{t+h}\right) \end{aligned}$$

where the second inequality follows from the assumption of constant conditional variance and the last inequality follows from decomposing the total variance of m_{t+h} into the sum of average conditional variance plus variance of the conditional expectation. Therefore

$$\mathcal{R}_j^2 \equiv \frac{Var\left(E\left[R_{j,t+h}\middle| \mathcal{F}_t^Z\right]\right)}{Var\left(R_{j,t+h}\right)} \leqslant R_{f,t+h}^2 Var\left(m_{t+h}\right), \qquad \forall\ m_{t+h} \in \mathcal{M}^Z$$

from which the first inequality in (1.7) follows from the definition of $\sigma_Z^2(\nu)$ in (2.2)■

It is useful to compare this result with the literature, in particular with Proposition 5 in Ross [2005] (see also Zhou [2010]). In line with our general approach of allowing the information in the predictors to be not necessarily included in the information in the state variables, a first contribution of our Proposition 2 is to show that, if condition (*) is violated, the $\mathcal{R}^2$ from a predictive regression is not constrained to be below the volatility of the SDF of a given pricing model, that is, $R^2_{f,t+h}Var\left(m^X_{t+h}\right) < \mathcal{R}^2_j$ is potentially compatible with the model Euler equation $E\left(m^X_{t+h}R_{t+h}\middle|\,\mathcal{F}^X_t\right) = e$. This fact can not emerge from Proposition 5 in Ross [2005], since there agents are assumed to price conditionally on a generic information set $\mathcal{F}^I_t$ which is implicitly assumed to satisfy $\mathcal{F}^I_t \supset \mathcal{F}^Z_t$, i.e. to contain all the information in the predictors. Our proof, moreover, highlights the importance of assuming returns to have constant conditional variance, which implies

$$E\left[Var\left(R_{j,t+h}\middle|\,\mathcal{F}^Z_t\right)Var\left(m_{t+h}\middle|\,\mathcal{F}^Z_t\right)\right] \leq Var\left(R_{j,t+h}\right)Var\left(m_{t+h}\right)$$

from which the bound on $\mathcal{R}^2_j$ obtained in Proposition 2 follows. Without constant conditional second moments, that is, in the case of stochastic volatility, it is not obvious that a bound as the one in (1.7) can be established at all.

1.3 Empirical Results

In this section we put the theoretical framework introduced in the previous section to work. We first introduce the linear predictive model for returns and use it throughout the empirical part to compute the predictors-based bound $\sigma_Z^2(\nu)$ for different sets of assets and horizons. We then test condition (*) for the long run risk model, the external habit model and the rare disaster model. Since this condition holds, we go ahead and compare the volatilities of the model implied SDFs with our predictors-based bounds.

1.3.1 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;[8]

[8] 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

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).[9]

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]).[10]

Table 2.2 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 2.2]

In contrast to the simple random walk view, stock and bond returns do seem predictable, and markedly more so the longer the return 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:

$$R^i_{t+h} = \beta^i_{0,h} + \beta^i_{1,h} z^i_t + u^i_{t+h} \tag{1.8}$$

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

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 and particularly so at the 1-year horizon.

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

[10] 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.

above, the holding period ranges from one quarter to five years, i.e. $h = 1, 4, 20$ quarters. Table 1.2-Panel A presents regressions of the real stock returns R^s_{t+h} on the price-dividend ratio pd_t and the consumption-wealth ratio cay_t. The choice of these two stock market predictors is motivated by the present value logic, see Campbell and Shiller [1988b], and 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. Although the $R^2 = 4\%$ at quarterly horizon does not look that impressive, it then rises with the horizon, reaching a value of about 50%, at the 5 years horizon.[11] 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 1.2-Panel B presents regressions of the real returns from holding a 5-year maturity bond onto the lagged short-term nominal interest rate y_t and the lagged yield spread spr_t. The results show that our predictive system is able to capture fluctuations in bond returns at all horizons.[12] 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) and Viceira [2012] among others).

[Insert Table 1.2]

We conclude this section with three observations. First, our predictive model (2.6) is based on predictors that are not a direct consequence of any specific asset pricing model. In our notation, this means that $\mathcal{F}^Z_t$ is not restricted to be a subset of $\mathcal{F}^X_t$. We apply our methodological framework to assess whether the evidence of predictability at different horizons based on $\mathcal{F}^Z_t$ and reviewed above can in fact be used to construct legitimate bounds for an asset pricing model in which expectations are taken conditioning on the information $\mathcal{F}^X_t$ generated by the model's state variables. 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,

[11]Using excess returns yields similar results.

[12]The results at the quarterly and 5-year horizons are robust to the inclusion of the Cochrane and Piazzesi [2005] factor, CP_t, in the set of predictors used to forecast bond returns. Had we used the CP_t factor as an additional predictor, the main difference with the system incorporating only the spread and 3-month T-bill would be at the 1-year horizon: the Newey-West corrected t-statistic on the yield spread decreases from 3.66 to 1.66 and the R^2 increases from 19.1% to 29.1%. In this sense our predictors-based bounds are conservative.

2009) or the long-run past market variance (see Bandi and Perron, 2008), 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, 2008b) 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, i.e. for models where $Var\left[E\left(R_{j,t+h}\left|\mathcal{F}_t^Z\right.\right)\right] > 0$ over some holding period h.[13]

1.3.2 Predictors-based bounds across horizons

It seems apparent from Table 1.2 that expected returns vary over time. To evaluate the ability of the predictors to increase the utility of the variance bounds as a diagnostic tool, in this section we compare the predictors-based bounds to the classical, unconditional HJ bounds. Along with the predictive versus unconditional dimension, we also analyze the effect of altering the investment horizon. Whereas Cochrane and Hansen [1992] were the first to carry out this exercise on the unconditional variance bounds, our analysis extends their results and highlights the interaction of conditioning information with the horizon dimension.

To compute the predictors-based bounds we use the solution to the left-hand side of Eq. (2.4). Bekaert and Liu [2004] show that the optimal trading strategy w_t that incorporates information is:

$$w_t = \left(\mu_t \mu_t^\intercal + \Sigma_t\right)^{-1} \left(1 - \kappa \mu_t\right)$$

where μ_t and Σ_t are the vector of conditional expected returns and the conditional variance-covariance matrix, respectively, and $\kappa = (\nu - b)/(1 - d)$, $\nu = E[M_{t+1}]$, $b = E\left[e' \left(\mu_t \mu_t^\intercal + \Sigma_t\right)^{-1} \mu_t\right]$ and $d = E\left[\mu_t' \left(\mu_t \mu_t^\intercal + \Sigma_t\right)^{-1} \mu_t\right]$. To compute the first and second conditional moments of asset returns, μ_t and Σ_t, we use the linear predictive model in equation (2.6).[14] For simplicity, we assume the conditional covariance matrix for returns to be constant, and estimate it has the residual in the forecasting regressions.

Figure 1.3 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

[13] 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.

[14] We investigate the effect of potential model misspecification in our regressions on the construction of the bounds in Section 1.4.1.

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.[15] Similar to Cochrane and Hansen [1992], Figure 1.3 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 1.2 translates into a tight lower bound on the variance of the SDF. In particular Figure 1.3 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 1.73 (1.64, 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 considerable predictability documented in Table 1.2.

Figure 1.5 presents the same analysis for SET B. In this case, the use of conditioning information yields a bound that is about 1.3 (1.4, respectively) times the unconditional HJ bound at the 1-year horizon (at the 5-year horizon, respectively). Upon comparing Figures 1.3 with 1.5, 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 A.

Taken together Figures 1.3 and 1.5 impart two conclusions. First, these figures highlight the three effects that are at work simultaneously: the conditioning information embedded in the conditional 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, our predictors-based bounds reveal the fundamental role played by conditioning information already at short horizons.

1.3.3 Predictors-based bounds and asset pricing models

Since we want to employ our predictors-based bounds to assess features of a SDF that are necessary to price returns, we focus on models that embed different preferences and specify the long run and short run risk in distinct ways (see also Hansen [2009]). In particular

[15] 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.

we investigate three leading classes of asset pricing models: the habit model, the long-run risk model, and the rare disaster model. We take as representatives of these classes the Campbell and Cochrane [1999] model as estimated by Aldrich and Gallant [2011a], the model proposed and estimated by Bansal et al. [2012c], and the model proposed and estimated by Nakamura et al. [2013].[16,17, 18]

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 SDF. It is important to note that all three models are estimated using a long span of the data sampled on an annual frequency to better capture the overall low frequency variations in asset and macroeconomic data and to reduce the measurement errors that arise from seasonalities and other measurement problems (see e.g. Wilcox, 1992).[19] For our chosen specification of the long-run risk and habit models, calibrated parameters are available from Bansal et al. [2012a] and Campbell and Cochrane [1999], respectively. Therefore for these two models we also report for comparison the moments of the model-implied SDF m_{t+h}^{X} obtained using calibrated parameters instead of estimated values. Finally, the three models have solution methods that are well established, and that ease the computation of the first and second unconditional moments of their SDF.[20]

Tables 1.7, 1.8 and 1.9 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.

[Insert Tables 1.7, 1.8 and 1.9 about here]

[16] Aldrich and Gallant [2011a] 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.

[17] 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.

[18] 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. For each parameter, the posterior mean and posterior standard deviation are presented in Table 1.9. We refer to the posterior mean of each parameter as our point estimate for that parameter.

[19] Aldrich and Gallant [2011a] use annual data from 1930 to 2008, Bansal et al. [2012c] from 1930 till 2009, and Nakamura et al. [2013] from 1890 to 2006. Thus 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.

[20] While the conditional variances are amenable to closed-form characterization, the unconditional variances are in general tractable only via simulations.

A number of recent papers (see e.g. Gourio [2008], 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. There are two main reasons behind our choice of concentrating on the Nakamura et al. [2013] instead of a framework with time-varying probability and severity of disasters. First, 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. Second, and most importantly, we want to explicitly abstain from models that are tailored to produce predictability. We will see that although the Nakamura et al. [2013] model does not deliver as much predictability as there is in the data, yet it stands out at the light of our predictors-based bounds. This fact highlights that while predictability in the data is key to have sharp predictors-based bounds, the amount of predictability implied by a model is not a necessary condition for that model to satisfy our bounds.

Testing condition (*)

As discussed in Section 2, a violation of condition (*) could in principle prevent us from using the predictors-based bounds as a legitimate diagnostic tool for the SDF of a given model. Therefore, to make sure that it is sensible to apply our predictors-based bound to the external habit, the long-run risk and the rare disaster models, we first check condition (*) for each model, asset class and horizon.

To test condition (*) we simulate from a given model a large number of paths for the SDF m_t^X and the state variables X_t. The length of the simulated path matches with the post-war sample period used to construct the predictors-based bounds. For each path we then run the following two regressions:

$$\begin{aligned} m^X_{t,t+h} R^i_{t,t+h} &= \alpha_{1,h} + \beta_{1,h} X_t + \varepsilon_{1,t+h} \\ m^X_{t,t+h} R^i_{t,t+h} &= \alpha_{2,h} + \beta_{2,h} X_t + \gamma_h Z_t + \varepsilon_{2,t+h} \end{aligned}$$

with $i = S, B$ denotes the asset class, $m^X_{t,t+h} = \prod_{i=1}^h m^X_{t+i}$ 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, X_t is a vector of state variables of a given model, and $Z_t = (pd_t, cay_t, spr_t, y_t)$ is a vector of all predictors employed. Our criterion is not to reject condition (*) as long as zero is included within the 90% confidence interval of the difference between the fitted values

$$\widehat{m^X_{t,t+h} R^i_{t,t+h}}^{X,Z} - \widehat{m^X_{t,t+h} R^i_{t,t+h}}^{X} \quad .$$

With the test assets, the set of candidate predictors and the model-implied stochastic discount factors at hand, we are now ready to test condition (*). Figure 1.1 displays the results for the long-run risk (Panel A), the external habit (Panel B) and rare disaster models (Panel C) when the test asset is the Equity Market. Empirically, there is no horizon at which we reject condition (*), and this is 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 constant maturity bonds.[21] Therefore our predictors-based bounds are legitimate lower bounds for the volatility of the SDF of the models under scrutiny and we can conclude that the predictors that we employ in our linear forecasting model are indeed useful in sharpening the diagnostic efficacy of variance bounds with respect to the unconditional case.

[Insert Figure 1.1 about here]

To conclude this section we remark that, although an analysis of the size and power properties of our test procedure is not an objective of this paper, still we are able to show that the condition 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. Figure 1.2 shows that in this case our procedure does reject condition (*) soundly, as one definitely expects.

[Insert Figure 1.2 about here]

Model-implied SDFs and predictors-based bounds

Since condition (*) cannot be rejected, we now assess the three asset pricing models under scrutiny using our predictors-based bounds as a diagnostic tool. 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 Bansal et al. [2012a] and Campbell and Cochrane [1999] models, and 50,000 annual observations for the rare disaster model of Nakamura et al. [2013].[22] From this long time series we then calculate the unconditional moments of the corresponding SDF.[23, 24]

To properly compare the moments of the simulated model-implied SDF with the predictor-based bounds we need to account for two sources of uncertainty. First, the

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

[22] Our parametrization of estimated values of rare disaster model are taken from Nakamura et al. [2013], in which all parameters are estimated at the annual frequency.

[23] 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).

[24] The triangles are obtained by averaging the mean and the standard deviation of the model-implied SDF obtained from 10 long (50,000 years) simulations.

volatility bound is estimated from the data, and hence it reflects the uncertainty surrounding the linear predictive model (2.6) used to compute the conditional moments of returns. Second, the computation of the mean and standard deviation of a model-implied SDF relies on the estimates of the exogenous state dynamics, and hence it reflects the uncertainty of the parameters describing these dynamics. To account for the first source of uncertainty, we construct confidence intervals for the predictors-based bounds using a block bootstrap scheme to draw 50,000 random samples from the data.[25] To account for the second source of uncertainty, we follow an approach similar to Cecchetti et al. [1994] and Burnside [1994] and further described in the Appendix.

Figures 1.3, 1.4, 1.5, and 1.6 display the predictors-based bounds and the SDFs generated by the three competing models for different horizons and for different sets of test assets. The triangles represent population values of the three SDFs obtained using estimated parameters. The uncertainty arising from the estimation of the parameters that govern the law of motion of the state variables is captured by the ellipses centered around the triangles, ellipses that are meant to represent confidence areas. For the long-run risk and the habit model we also report, for comparison, the population values of the SDFs obtained using calibrated parameters (the stars in the figure). Figures 1.3 and 1.5 display our predictors-based bound along with the classical HJ bound, for SET A and SET B respectively. Figure 1.4 and 1.6 report instead the 90% confidence interval, based on block bootstrap, for the predictors based-bounds obtained from SET A and SET B respectively.

[**Insert Figures 1.3, 1.4, 1.5, and 1.6 about here**]

A first conclusion that emerges clearly from Figures 1.4-Panel C and 1.6-Panel C is that independently from the horizon and test assets considered, the SDF of the rare disaster model falls by large into the predictors-based bound even after accounting for estimation uncertainty. When we consider the other two models, instead, we observe that at the one-year horizon, and using SET A, it is the long-run risk model that seems most challenged by our predictors-based bounds. As shown in Figure 1.4-Panel A, in fact, even after accounting for parameter uncertainty the SDF of this model touches only marginally the predictors-based bound. The SDF of the external habit fares much better and falls within the lowest 90% confidence interval of our predictors-based bound - see Figure 1.4-Panel B. However as soon as we expand the set of asset, we see from Figure 1.6-Panel B that, at the 1-year horizon, also the habit model fails to satisfy the predictor-based bound. At long horizons, finally, while the habit model satisfies the predictors-based bound independently from the test assets considered, the long-run risk model generates instead an SDF not volatile enough, even after accounting for uncertainty (see Figure 1.6-Panel A).

[25] We select the optimal block length for the bootstrap according to Politis and White [2004]. We thank Andrew Patton for making the code available on his website.

The above results point to an interesting fact about the role played by preferences and state dynamics. Although the rare disasters and long-run risk models share the same preferences for early resolution of uncertainty, the model-implied SDFs have very different behavior. This behavior can be explained by the different ways the long-run risk and the rare disaster 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 our predictors-based bounds. To reinforce this point we consider in Figure 1.7 two specifications of the rare disaster model. Both specification have the same recursive preferences, but allow for different disaster dynamics. In particular we compare the model-implied SDF of Nakamura et al. [2013] (triangles), with the SDF implied by a model with permanent, one-period disasters of the type analyzed in Barro [2006a] (stars). The SDF from the permanent disaster 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, but 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 very close to the upper 90% confidence predictors-based bound across all horizons and for all test assets.

[Insert Figures 1.7 about here]

Figures 1.3 and 1.5 provide a visual representation of the importance of jointly considering conditioning information and horizons for the equity premium puzzle. Figure 1.3 shows that the SDFs (triangles) of all three models satisfy the unconditional HJ bounds at the 1-year horizon. This is not surprising for the first two models, since they are calibrated closely to offer conformity with the historical unconditional annual real return on the risk-free bond and the equity market and the estimated parameters are close to the calibrated ones, a fact reflected visually by the vicinity of the stars to the triangles. The conclusions are different when we incorporate conditioning information. In this case the long-run risk model struggles to meet the bounds. The habit model meets the restriction, its SDF being exactly on top of the predictors-based bound, but the parameter uncertainty does not allows us to draw a robust conclusion. As expected, using SET B and hence expanding the set of assets exacerbates these results and both models fall largely below the bounds (see Figure 1.5). 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 our predictors-based bounds highlight that what really matters is the interaction between state dynamics and preferences.

Recent theoretical and empirical research in macro-finance 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 1.3 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 our 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 find it onerous to satisfy bounds at such a long horizon (see Panel A in Figures 1.3 and 1.5).

Finally, although for both the long-run risk and the external habit models the SDFs computed using estimated parameters fare a touch worse than the SDFs computed with calibrated parameters in terms of satisfying the bounds (see for instance the SDF implied by long-run risk at the 1 and 5 years horizons, and the habit SDF at the 5 years horizon), Figures 1.4 and 1.6 show that our conclusions are largely unaffected by the choice of calibrated versus estimated parameters. The evidence reveals that the variance of the SDF from the long-run risk and habit models fails to meet the predictors-based bound restriction at the 1-year horizon for both SET A and SET B, and of these two only the Campbell and Cochrane [1999] model has the potential to resolve the long horizon equity premium.

The results discussed so far are summarized in Table 1.3. The last two columns report the minimum value of the predictors-based bound for SET A and SET B, respectively. The remaining columns report the unconditional variance of the model-implied SDFs, using both estimated and calibrated parameters.[26] The table reveals that at the yearly horizon the standard deviation implied by the Bansal et al. [2012a] model using estimated parameters is 0.47 and therefore the long-run risk value lies below the minimum volatility from the predictors-based bound constructed from SET A, which is equal to 0.64. When we look at long horizons the habit and rare disaster models, with the sound values 4.29 and 6.18 using estimated parameters, lie comfortably below the bound. The table shows also that the population values for the volatility of the SDF implied by the long-run risk model and external habit model would be slightly greater when using calibrated

[26]For the rare disaster model we consider only estimated values, since there are no calibrated values of model parameters.

values. For the long-run risk model this is mainly driven by the lower persistence in the consumption growth volatility: compared to the benchmark calibration, where the half-life is essentially infinite (58-year), the estimated value implies a half-life slightly over 33-year.

[Insert Table 1.3 about here]

In sum, this section shows that by incorporating conditioning information from a well-established set of stock and bond predictors our 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 other two models, at a long 5-years investment horizon, and using SET B, it is the habit model that turns out to be able to generate enough volatility in the SDF, and not the long-run risk model. Finally, at the 1-year horizon, and using SET B, both the long-run risk and habit models fail to meet the restrictions imposed by the predictor-based bounds, with the standard deviation of their SDFs never approaching the bounds.

As a final remark, we note that the conclusions in this section are indeed conservative. As discussed at the end of Section 2.1, our predictors-based bounds could be tightened even more by enlarging the set of assets, the set of predictors, or in general by considering truly dynamic investment strategies with rebalancing. We leave the analysis of these extensions to future research.

1.3.4 Historical versus model-implied predictability

In this section we investigate whether the ability of an asset pricing model to satisfy the predictors-based bound is the counterpart of 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.

Table 1.4 displays results for the predictive regressions of future gross real returns over the 1, 3, 5 and 8 years horizons, with both the actual data and the model simulated data. Model-implied predictive regression results are listed in Column 2 (Bansal et al. [2012a]), Column 3 (Campbell and Cochrane [1999]) and Column 4 (Nakamura et al. [2013]). For these cases the table reports the median coefficient and its standard error, along with the median R^2 with associated 95% confidence interval, obtained from 1000 regressions run on simulated data.

[Insert Table 1.4 about here]

In the data, the predictability of gross returns increases as the horizon goes up, the R^2 rising from 7.9% at the 1-year horizon to about 32.8% at the 8-year horizon. When we look at the model-implied results, we observe that the long run risk model by Bansal et al. [2012a] and the rare disaster model by Nakamura et al. [2013] feature modest predictability, with the median R^2 in the range of 0.0% - 7.5% and 0.0% - 2.4%, respectively. The estimated slope coefficients of the long run risk model are close to those obtained by using real data across all the horizons, while for the rare disaster model they are too low at the short horizon, -0.05 at the one-year, and too high at the longer horizons, -1.28 at the eight-year. On the other hand the Campbell and Cochrane [1999] model exhibits a median R^2 that starts at 9% at the 1-year horizon, and then rises all the way to 37.5% at the 8-year horizon, a figure even higher than the one, 32.8%, generated from real data. As shown in Table 1.4, however, the very high R^2s in the habit model are paired with excessively large slope coefficients across all horizons. For example, the model implied slope coefficients are -0.29 at one-year horizon and -1.31 at eight-year horizon while the corresponding point estimates for the historical data are -0.13 and -0.63, respectively. The predictive power for the price-dividend ratio is amplified by a factor of almost two. Table 1.4 also presents the 95% confidence interval of model-implied R^2s. Since the R^2s in the data are within the 95% model-based confidence intervals, the evidence suggests that all these three models can, in principle, match the return predictability observed in the data. Figure 1.9 provides a visual representation of these results.

[Insert Figure 1.9 about here]

Our results show that the ability of a model to replicate the return predictability observed in real data is not a crucial criteria for model selection. In particular the explanatory power of the price-dividend ratio in predicting future returns at medium and short horizons is higher than in the data within the Campbell and Cochrane [1999] model, while it is close to zero in the Nakamura et al. [2013] model. Despite this difference, the two models perform comparably well at long horizons. Moreover at the 1-year horizon our predictors-based bound favors the rare disaster, not the habit model.[27]

In sum this section reinforces the fact that return predictability can be used to sharpen the variance bounds of SDFs (provided condition (*) is satisfied) and that our predictors-based bounds emerge as a criteria more effective than model-based predictability to discriminate among asset pricing models.

[27] A natural directions for future research would be to compare the performance of rare disaster models engineered to reproduce predictability, e.g. Gabaix [2012] and Wachter [2013a], to the Nakamura et al. [2013] model.

1.4 Extensions

In Section 1.3.1 we have shown that incorporating predictability of asset returns does make the variance bounds tighter and hence it imposes – when condition (*) is satisfied – a harder yardstick on asset pricing models. 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 then reinterpret our findings in terms of upper bounds on the R^2s from predictive regressions, and finally we analyze the robustness of our results to the possibility of misspecification of the model for the conditional moments of returns.

1.4.1 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 (2.6). In the first case (restriction I) stock returns are assumed to be unpredictable. In the second case (restriction II), it is instead the returns from the strategy rolling over Treasury bills that is assumed to be unpredictable.

Figure 1.8 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 1.8 about here**]

From Figure 1.8 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.[28] One might then wonder if it is the predictability of bond returns the key to

[28] 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.

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.[29]

Summing up, the shape of our 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.

Bounds, models and R^2

We now evaluate the asset pricing models scrutinized so far through the lenses of the R^2s of the predictive model (2.6) that underlies our predictive-based bounds. In Table 1.5, for each model, we compare the variance of the SDF scaled by the squared gross risk-free rate, with the R^2s of predictive regressions for stock and 5-year constant maturity government bond real gross returns. The table, moreover, reports the minimum variance of both the unconditional HJ bounds and of our predictors-based bounds, both scaled by the squared gross risk-free rate as well. As a proxy for the risk-free rate we use the returns from a strategy that rolls 3-month Treasury Bills over the horizon of interest.

[Insert Table 1.5 about here]

The implication of Proposition 2 in Section 2.3 is that, as long as Condition (*) holds, for any asset class the minimum scaled variance of the predictors-based bounds $R^2_{f,t+h}\sigma^2_Z(\nu_{\min})$ should be intermediate between the predictive R^2 and the scaled variance of the models SDFs. After the analysis carried on in the previous section, it is not surprising that this implication is challenged by the data at different horizons for the long-run risk and the habit models, while it is not challenged at all for the rare disaster model. What is more interesting here is to observe that at shorter horizons the scaled variance of the SDFs implied by the long-run risk and the habit models are very close to, and in certain cases outright below, the predictive R^2s. This happens, in particular, for the case of the long-run risk model and the R^2 of the 5-year constant maturity government bond real gross returns, where the variance of the scaled SDF is well below the predictive R^2 at the 1-quarter horizon (although, to be fair, these R^2 still belong to the confidence interval around the scaled variance of the SDF). When compared to the scaled minimum variance of the predictors-based bounds, however, the predictive R^2s line up nicely below the minimum values at all horizons.

The findings in this section are interesting from two points of view. First, the fact that the scaled variance of the long-run risk model falls below the R^2 of the 5-year constant

[29] 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).

maturity government bond real gross returns at the 1-quarter implies a short horizon challenge to the model that complements the long horizon challenge discussed above. Second, the fact that the inversion between the scaled variance of a model-implied SDF and a predictive R^2 may occur even when the R^2 falls nicely below the scaled minimum variance of the predictors-based bound, reinforces the point made in Proposition 2: the difference between the information sets associated with the predictors and that posited by a given model must be duly accounted for when employing predictive $\mathcal{R}^2$ as model diagnostics.

Predictability, model mis-specification and variance bounds

We conclude this section by investigating the performance of our bounds along two further dimensions: robustness and efficiency. Recall that the results presented so far are obtained 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 our bound we plot in Figure 1.10 alternative implementations of the variance bounds. Specifically, along with our 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 1.10 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 (2.4) in Section 2.1), must be a parabola under the null of correct moments specification. Figure 1.10 shows that in our case we obtain a smooth parabola indeed. The figure, moreover, shows that the GHT bound, obtained via the inf in (2.4), 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, see also Table 1.6: 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.[30] Although the FS approach yields the most

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

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

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.

1.5 Conclusions

We analyze predictors-based variance bounds, i.e. bounds on the variance of those SDFs that price a given set of returns conditional on the information contained in a vector of returns predictors. We identify a simple sufficient condition under which the predictors-based bounds constitute legitimate lower bounds on the variance of the SDF of a given asset pricing model. We use our predictors-based bounds to assess the performance of three leading consumption-based asset pricing models: the long run risk model of Bansal et al. [2012a], the habit-formation model of Campbell and Cochrane [1999] and the rare disaster model of Nakamura et al. [2013].

Our results point to the importance of jointly considering conditioning information and horizons. 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 the SDFs implied by Bansal et al. [2012a] fails to meet the lower bound restriction at 1-year horizon, while the SDF implied by Campbell and Cochrane [1999] displays violations when the set of traded assets is augmented to include also the returns from holding Treasury Bonds. Of the three models under scrutiny, in fact, only the rare disaster model of Nakamura et al. [2013] meets comfortably the 1-year predictors-based bound for all sets of assets considered in the paper.

Interestingly, the 5-year *unconditional* HJ bounds are 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, while the habit and rare disaster models maintain the capability of addressing the equity premium at long horizons, the long run risk model falls short since it produces an SDF that is not volatile enough.

Our 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). Our bounds yield a graphical

---

where

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

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 $\sigma^2_{FS}(v) \leq \sigma^2_{Z}(v)$.

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, our 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. [2012a] 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. [2012a] at the 1-year horizon). Importantly, our 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 our 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.

Appendix I: 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. 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-year log stock return between dates $t+j-1$ and $t+j$ and R_{t+j} is the simple gross return.

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}.

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. 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: price-dividend ratio, see Campbell and Shiller [1988b] and Campbell and Shiller [1988a]; cay, see Lettau and Ludvigson [2001]. Dividends are 12-month moving sums of dividends paid on the index.

4. Bond return predictors: Treasury-bill rates are the 3-Month Treasury Bill: SecondaryMarket Rate from the economic research data base at the Federal Reserve Bank at St. Louis (FRED) (see, e.g., Campbell [1987]). The Term Spread 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.

Appendix II: Confidence interval for SDFs

This appendix explains how we construct the confidence area of a model-implied SDF. These area account for the uncertainty of the estimated parameters used to compute the mean and standard deviation of the SDF. For a given consumption-based asset pricing model (LRR, external habit and rare disasters model in our paper), we obtain the model-implied SDF by simulating a long series (50,000 years) of consumption growth and of the state variables. Let us denote the model-implied mean and standard deviation of the SDF as

$$\mu_m(\phi,\theta)$$
$$\sigma_m(\phi,\theta)$$

where ϕ is denoted as the vector of parameters that characterize the preference, and θ contains all the parameters associated with the dynamics. For instance, in the LRR model, $\phi=(\delta,\gamma,\psi)$ and $\theta=(\mu,\mu_d,\phi,\varphi_d,\rho_{dc},\rho,\varphi_e,\overline{\sigma},\upsilon,\sigma_\omega)$, (see Table 1.9).

To construct the confidence area of (μ_m,σ_m), we calculate the standard deviation of $\mu_m\left(\phi,\widehat{\theta}\right)$, denoted as $\sigma^2_{\mu_m}$, and the standard deviation of $\sigma_m\left(\phi,\widehat{\theta}\right)$, denoted as $\sigma^2_{\sigma_m}$. These standard deviations exhibit the sensitivity of (μ_m,σ_m) with respect to estimated $\widehat{\theta}$. To obtain $\sigma^2_{\mu_m}$ and $\sigma^2_{\sigma_m}$ we follow the lead of Cecchetti, Lam and Mark (1994).

Based on equation (19) in Cecchetti, Lam and Mark (1994), the estimate for the variance of μ_m and of σ_m can be written as

$$\widehat{\sigma}^2_{\mu_m}=\left(\frac{\partial\mu_m}{\partial\theta'}\right)\bigg|_{\widehat{\theta}}\Sigma_{\widehat{\theta}}\left(\frac{\partial\mu_m}{\partial\theta}\right)\bigg|_{\widehat{\theta}}$$

$$\widehat{\sigma}^2_{\sigma_m}=\left(\frac{\partial\sigma_m}{\partial\theta'}\right)\bigg|_{\widehat{\theta}}\Sigma_{\widehat{\theta}}\left(\frac{\partial\sigma_m}{\partial\theta}\right)\bigg|_{\widehat{\theta}}$$

where $\Sigma_{\widehat{\theta}}$ is a diagonal variance matrix of $\widehat{\theta}$, and $\left(\frac{\partial\mu_m}{\partial\theta'}\right)\Big|_{\widehat{\theta}}\left(\left(\frac{\partial\sigma_m}{\partial\theta'}\right)\Big|_{\widehat{\theta}}\right)$ is a vector of the first derivative of $\mu_m\,(\sigma_m)$ with respect to θ, and evaluated at $\theta=\widehat{\theta}$. For instance, in the LRR model,

$$\Sigma_{\widehat{\theta}}=diag\left(\sigma^2_{\widehat{\mu}},\sigma^2_{\widehat{\mu}_d},\sigma^2_{\widehat{\phi}},...,\sigma^2_{\widehat{\omega}}\right)$$

where the values of $\sigma^2_{\widehat{\mu}},\sigma^2_{\widehat{\mu}_d},\sigma^2_{\widehat{\phi}},...,\sigma^2_{\widehat{\omega}}$ are listed in Table 1.9 and we numerically approximate the derivative by

$$\left(\frac{\partial\sigma_m}{\partial\theta}\right)\bigg|_{\widehat{\theta}}=\begin{pmatrix}\frac{\sigma_m\left(\phi,\widehat{\theta}_{-\widehat{\mu}},\widehat{\mu}+\varepsilon^{\mu}\right)-\sigma_m\left(\phi,\widehat{\theta}_{-\widehat{\mu}},\widehat{\mu}-\varepsilon^{\mu}\right)}{2\varepsilon^{\mu}}\\ \frac{\sigma_m\left(\phi,\widehat{\theta}_{-\widehat{\mu}_d},\widehat{\mu}_d+\varepsilon^{\mu_d}\right)-\sigma_m\left(\phi,\widehat{\theta}_{-\widehat{\mu}_d},\widehat{\mu}_d-\varepsilon^{\mu_d}\right)}{2\varepsilon^{\mu_d}}\\ \vdots\\ \frac{\sigma_m\left(\phi,\widehat{\theta}_{-\widehat{\omega}},\widehat{\omega}+\varepsilon^{\omega}\right)-\sigma_m\left(\phi,\widehat{\theta}_{-\widehat{\omega}},\widehat{\omega}-\varepsilon^{\omega}\right)}{2\varepsilon^{\omega}}\end{pmatrix}$$

in which $\varepsilon^{\mu}=\widehat{\mu}\times 10^{-7}$ and $\widehat{\theta}_{-\widehat{\mu}}$ is the vector of all estimated parameters excluded $\widehat{\mu}$.

With the standard deviations of μ_m and σ_m at hand, we describe the sensitivity property, and report our results graphically as the confidence ellipses in Figures 1.3 –

1.6: in particular the confidence area around (μ_m, σ_m) is covered by an ellipse centered $\left(\mu_m\left(\phi,\widehat{\theta}\right),\sigma_m\left(\phi,\widehat{\theta}\right)\right)$ with radius $\left(\frac{\widehat{\sigma}_{\mu_m}}{2},\frac{\widehat{\sigma}_{\sigma_m}}{2}\right)$.

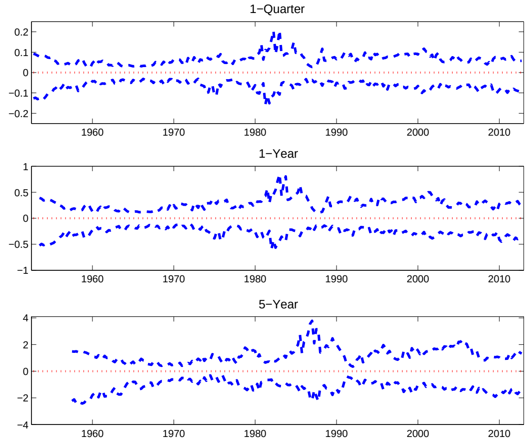

(a) Long Run Risk Model

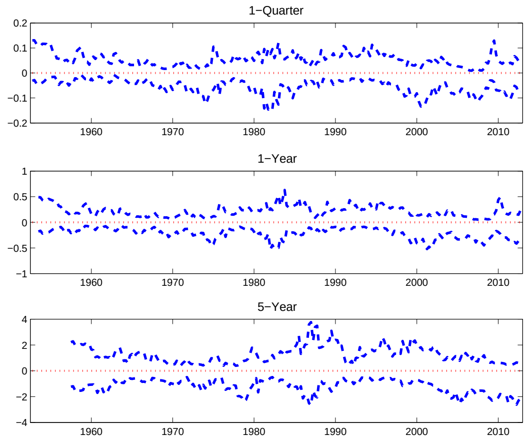

(b) External Habit Model

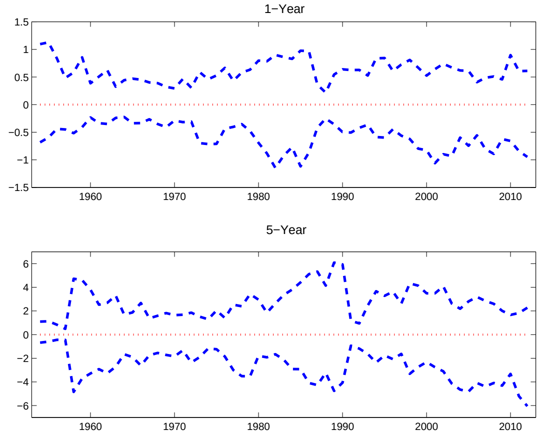

(c) Rare Disaster

Figure 1.1: Empirical verification of Condition(*) - Market returns.
Dashed blue lines give the 90% confidence interval of the differences between the estimated values of discounted returns with and without using predictors, at horizon 1-Quarter, 1-Year, 5-Year. The discounted returns are the product of model generated SDFs and real market index returns. Dotted red lines locate the benchmark of zero value. We do 1000 times simulations over 241 quarters for the long-run risk model by Bansal et al. [2012a] (Panel A), the external habit model by Campbell and Cochrane [1999] (Panel B) and the rare disaster model by Nakamura et al. [2013] (Panel C). We use *estimated* parameters as given in Tables 1.7), 1.8) and 1.9). We regress the discounted returns both on model generated state variables, and on model generated state variables plus predictors. Finally we plot the 90% confidence interval of the difference between the fitted discount returns from the two regressions. Sample: 1952Q2 - 2012Q3.

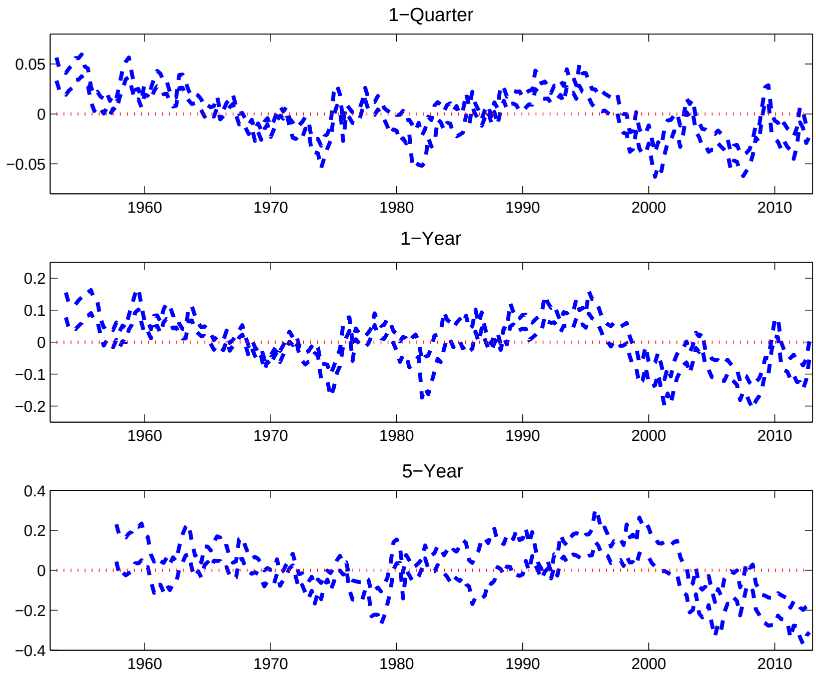

Figure 1.2: An analytical example in which Condition(*) fails - Market returns. Dashed blue lines give the 90% confidence interval of the differences between the estimated values of discounted returns with and without using predictors, at horizon 1-Quarter, 1-Year, 5-Year. The discounted returns are the product of model generated SDFs and real market index returns. Dotted red lines locate the benchmark of zero value. We do 1000 times simulations over 241 quarters of a consumption-based model with i.i.d consumption growth and CRRA utility. In this example, the risk aversion coefficient equals 7.42 and the subjective discounted factor equals 0.9989. We then regress the discounted returns both on model generated state variables, and on model generated state variables plus predictors. Finally we plot the 90% confidence interval of the difference between the fitted discount returns from the two regressions. Sample: 1952Q2 - 2012Q3.

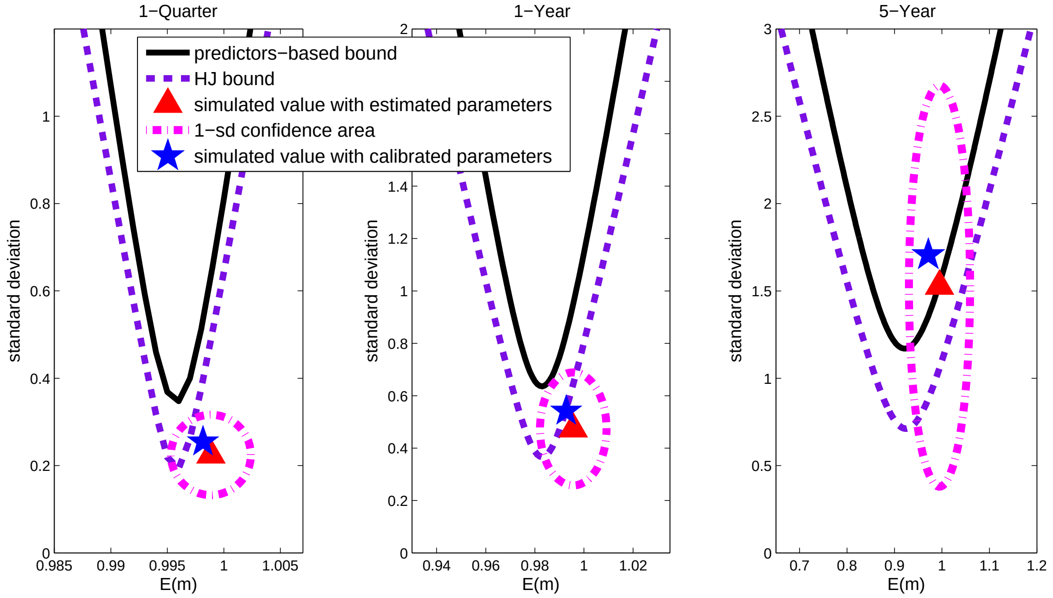

(a) Long Run Risk Model

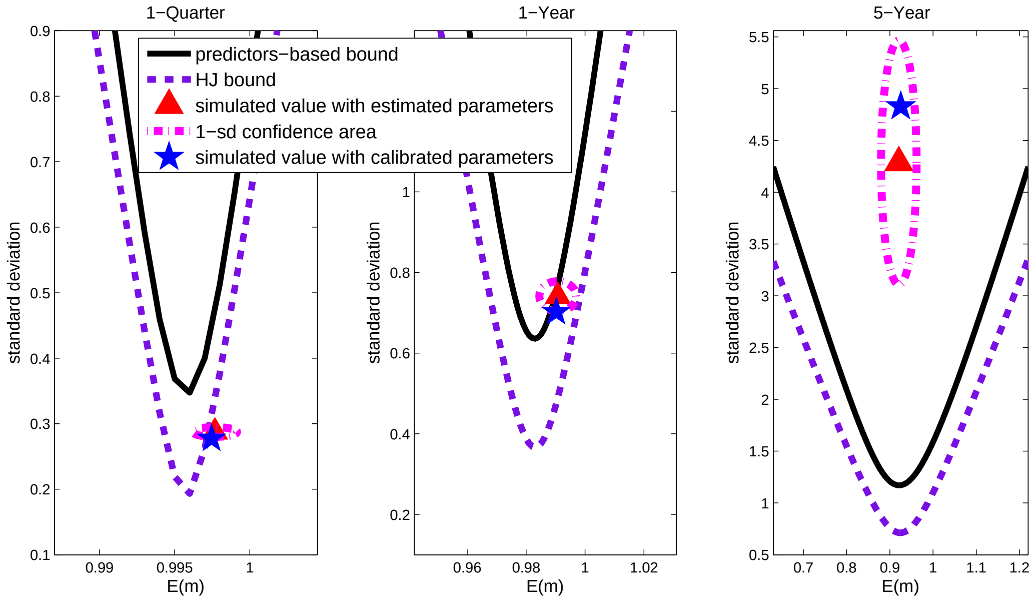

(b) External Habit Model

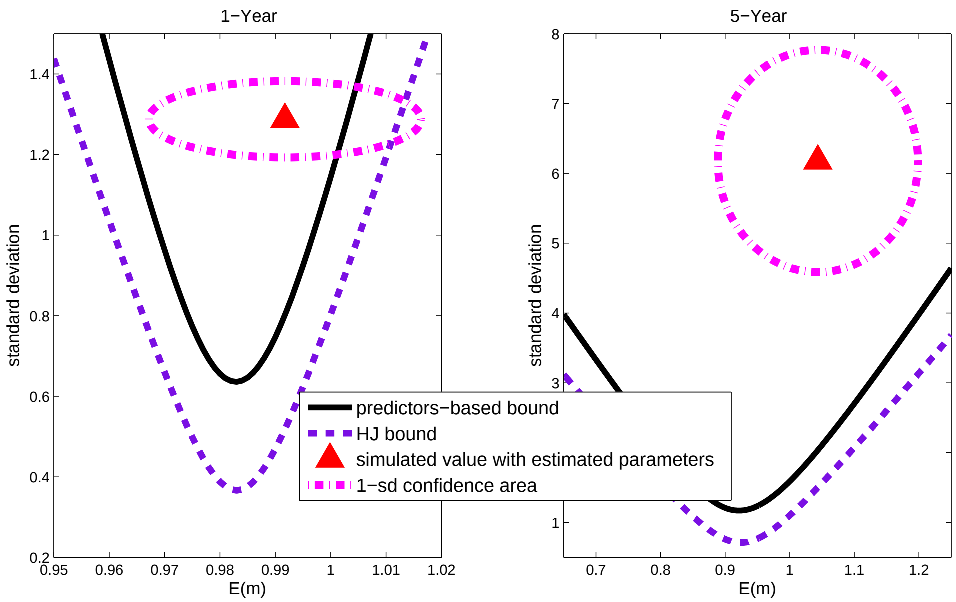

(c) Rare Disaster

Figure 1.3: Predictors-based bound $\sigma_Z^2(\nu)$, 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's (2004) specification. The red triangle reports average mean and standard deviation values from 10 simulations run of 600,000 months with *estimated* parameters. The blue star reports the same objects computed with calibrated parameters. The ellipse (dashed dotted area) shows the uncertainty (one standard deviation) in the calculation of the mean and standard deviation of the model implied SDFs. We only take account uncertainty in the parameter values for the state dynamics (i.e. we take as given parameters that characterize the preferences). Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3.

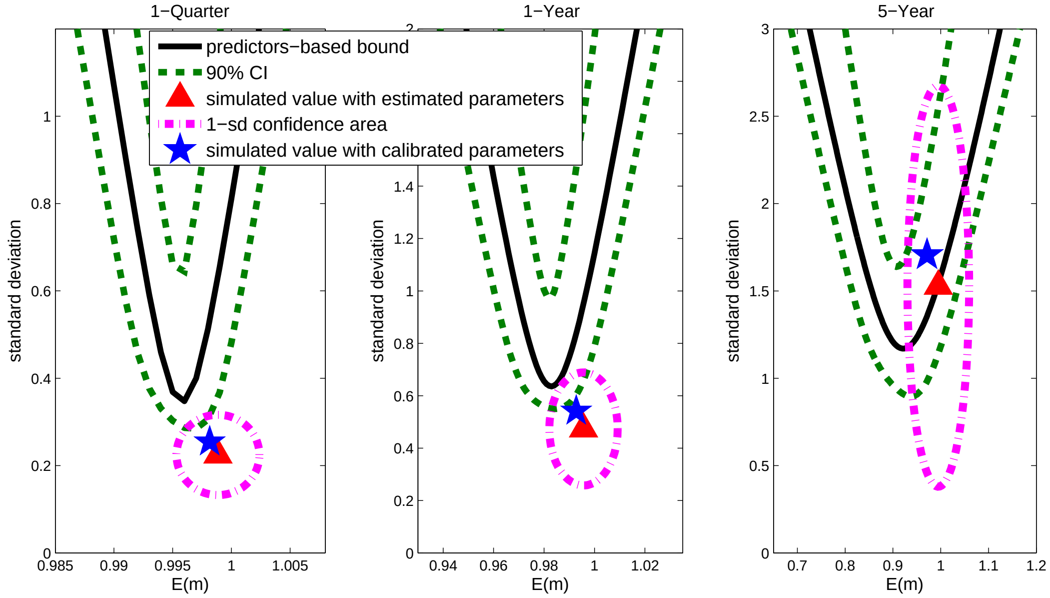

(a) Long Run Risk Model

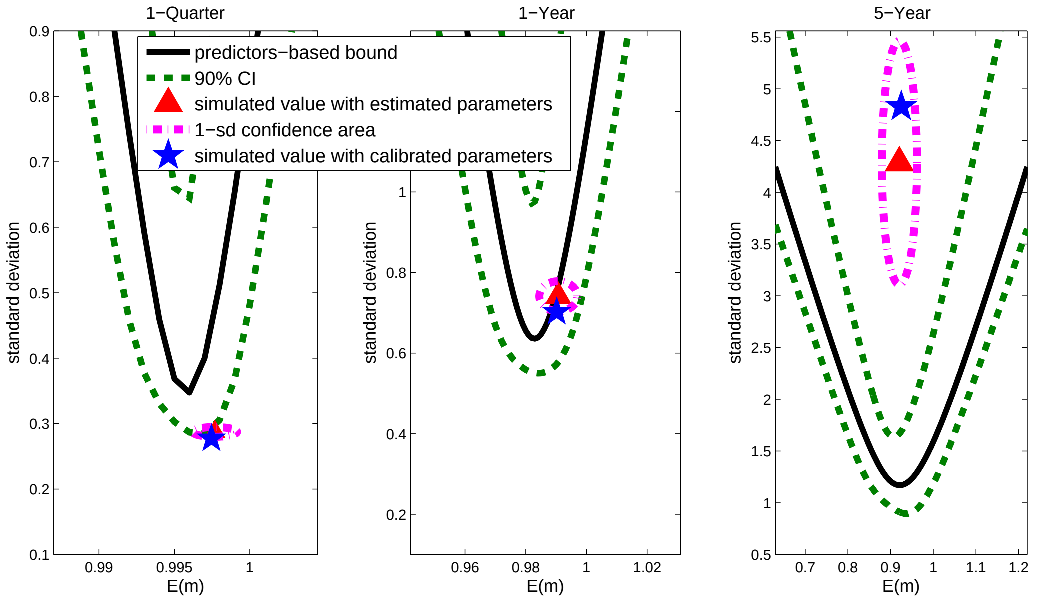

(b) External Habit Model

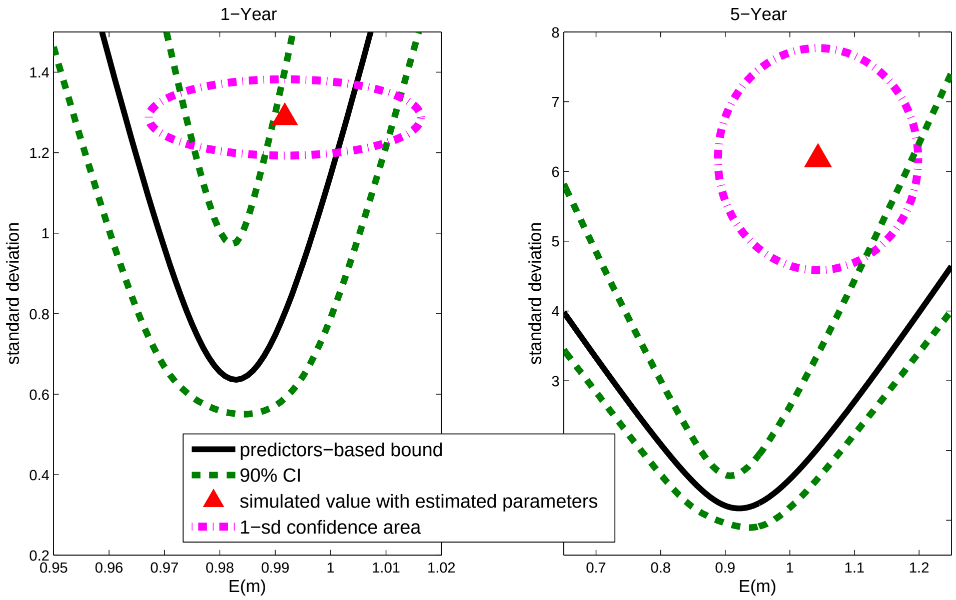

(c) Rare Disaster

Figure 1.4: Model-implied SDFs, predictors-based bounds and parameters uncertainty – SET A.

Solid black line gives the volatility bound using conditional information based on Bekaert and Liu's(2004) specification. Dashed dark green lines give the bootstrapped 90% confidence interval for corresponding variance bound. The red triangle reports average mean and standard deviation values from 10 simulations run of 600,000 months with estimated parameters. The blue star reports the same objects computed with calibrated parameters. The ellipse (dashed dotted area) shows the uncertainty (one standard deviation) in the calculation of the mean and standard deviation of the model implied SDFs. We only take account uncertainty in the parameter values for the state dynamics (i.e. we take as given parameters that characterize the preferences). Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3.

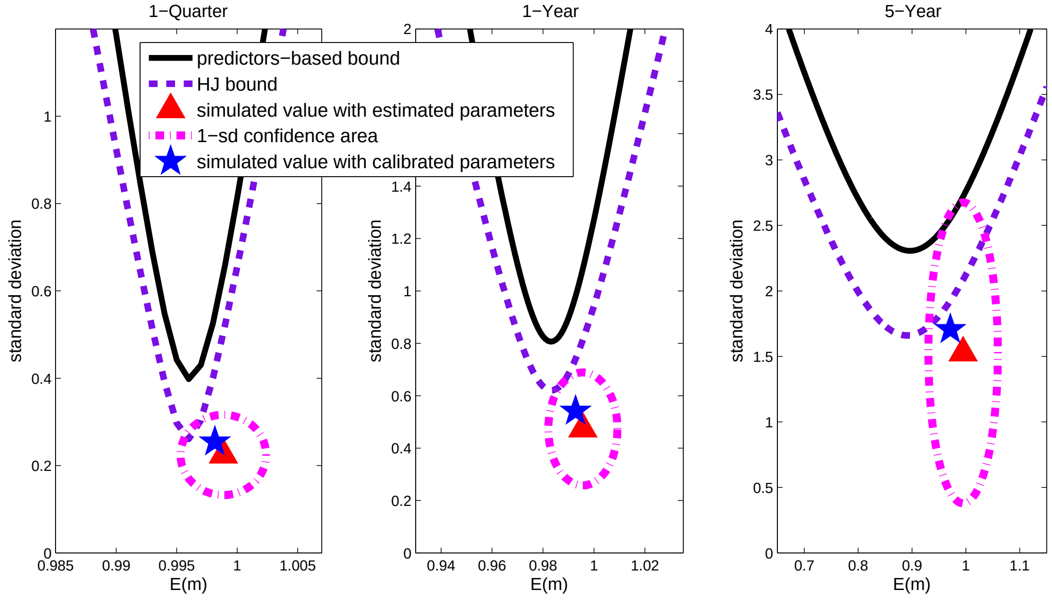

(a) Long Run Risk Model

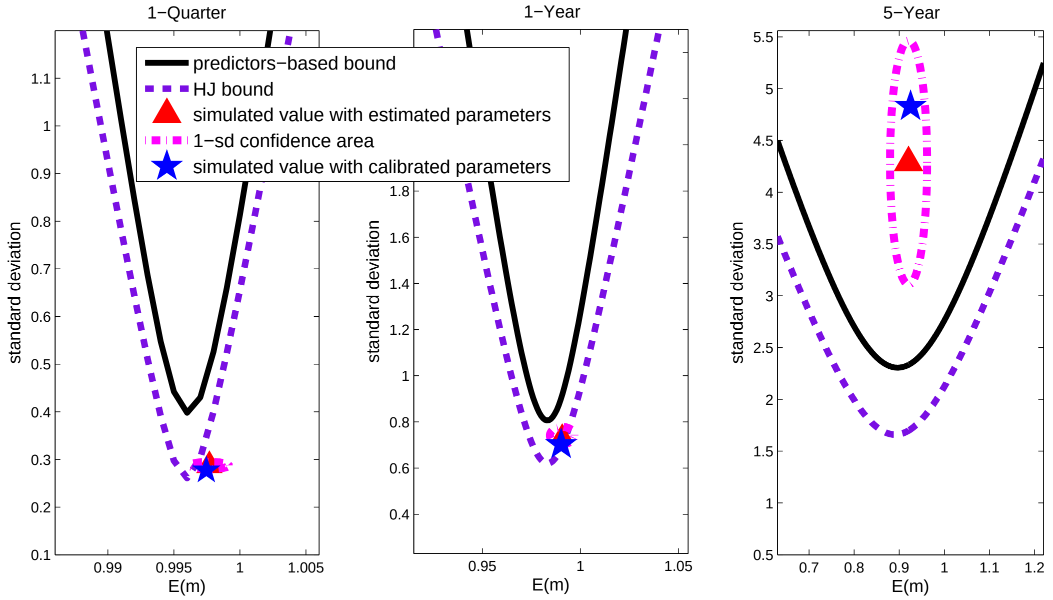

(b) External Habit Model

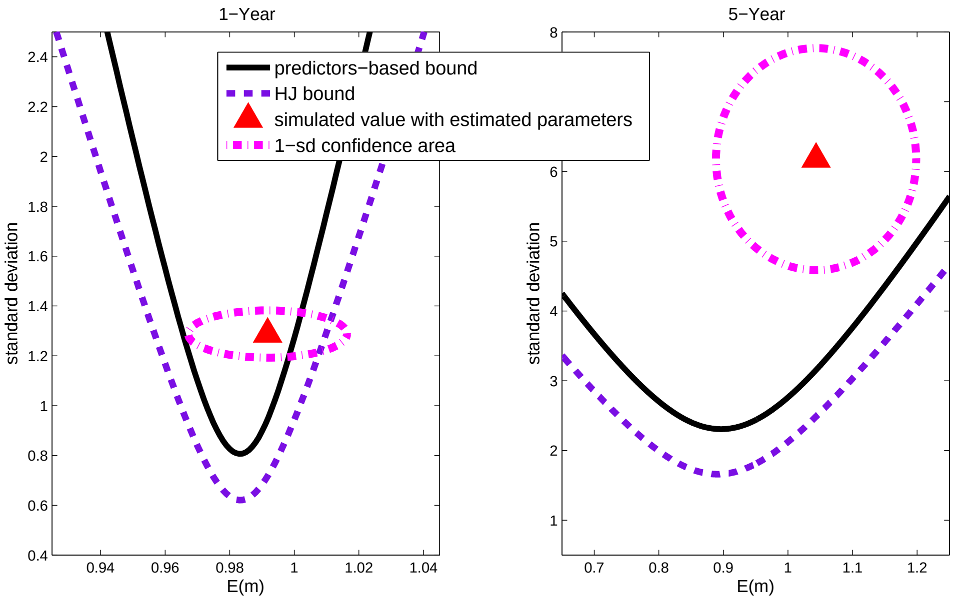

(c) Rare Disaster

Figure 1.5: Predictors-based bound $\sigma^2_Z(\nu)$, 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's (2004) specification. The red triangle reports average mean and standard deviation values from 10 simulations run of 600,000 months with *estimated* parameters. The blue star reports the same objects computed with calibrated parameters. The ellipse (dashed dotted area) shows the uncertainty (one standard deviation) in the calculation of the mean and standard deviation of the model implied SDFs. We only take account uncertainty in the parameter values for the state dynamics (i.e. we take as given parameters that characterize the preferences). Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3.

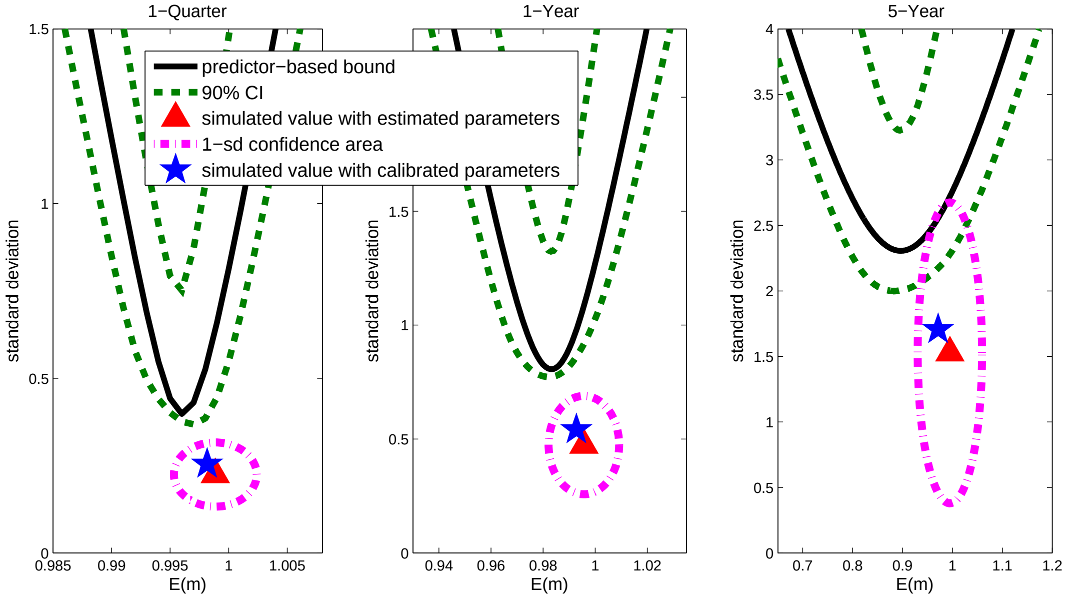

(a) Long Run Risk Model

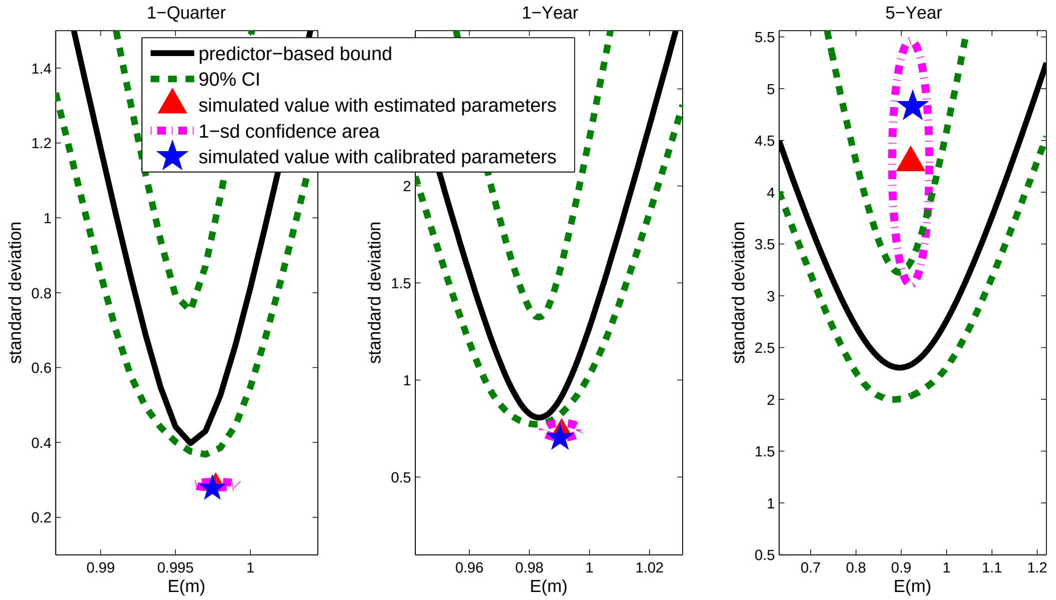

(b) External Habit Model

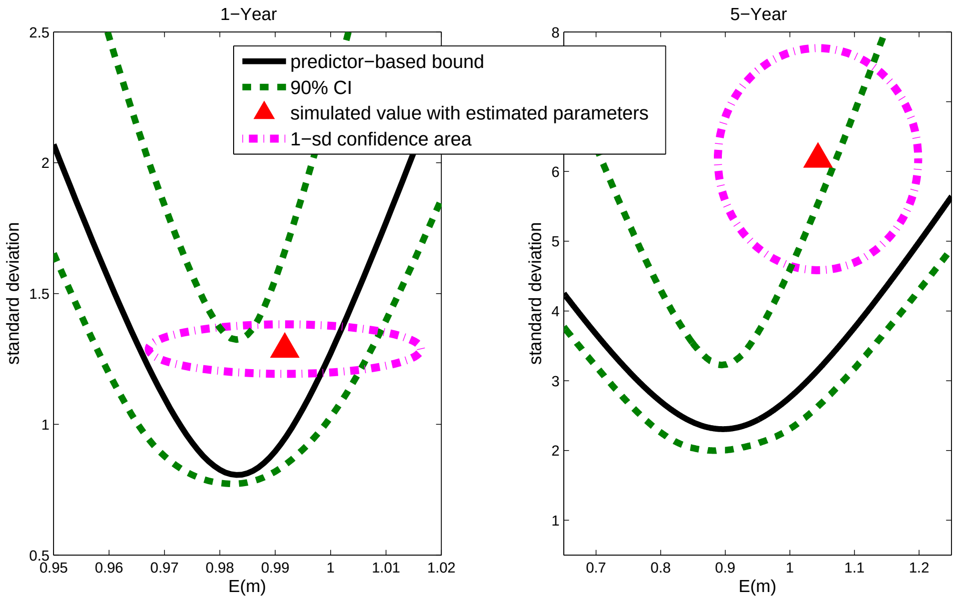

(c) Rare Disaster

Figure 1.6: Model-implied SDFs, predictors-based bounds and parameters uncertainty – SET B.
Solid black line gives the volatility bound using conditional information based on Bekaert and Liu's(2004) specification. Dashed dark green lines give the bootstrapped 90% confidence interval for corresponding variance bound. The red triangle reports average mean and standard deviation values from 10 simulations run of 600,000 months with estimated parameters. The blue star reports the same objects computed with calibrated parameters. The ellipse (dashed dotted area) shows the uncertainty (one standard deviation) in the calculation of the mean and standard deviation of the model implied SDFs. We only take account uncertainty in the parameter values for the state dynamics (i.e. we take as given parameters that characterize the preferences). Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3.

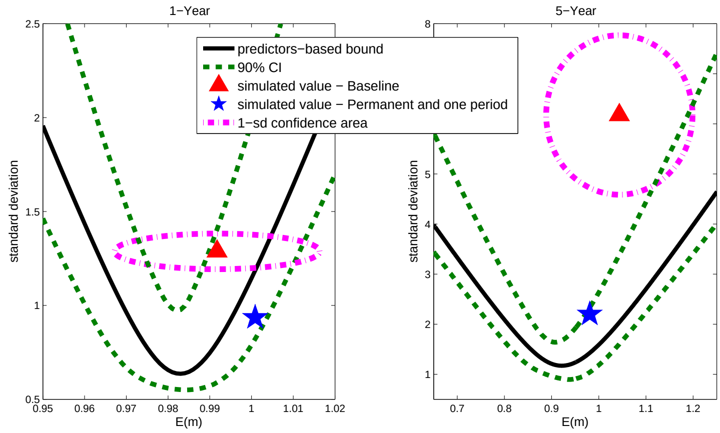

(a) SET A

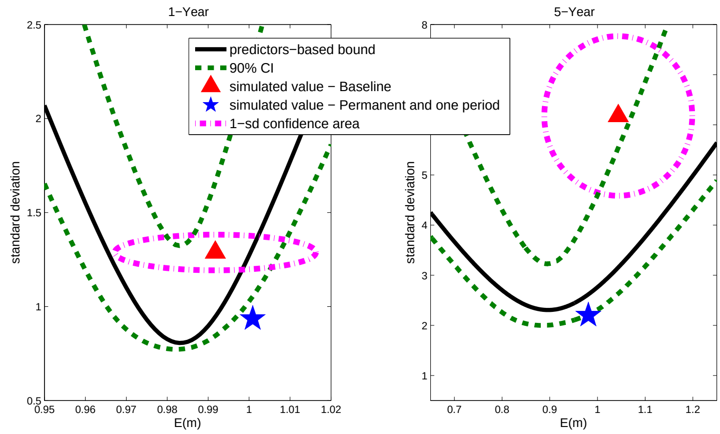

(b) SET B

Figure 1.7: 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's(2004) specification. Dashed dark green lines give the bootstrapped 90% confidence interval for corresponding variance bound. 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 star reports the same objects of the model with permanent, one-period disasters. The ellipse (dashed dotted area) shows the uncertainty (one standard deviation) in the calculation of the mean and standard deviation of the model implied SDFs. We only take account uncertainty in the parameter values for the state dynamics (i.e. we take as given parameters that characterize the preferences). Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3.

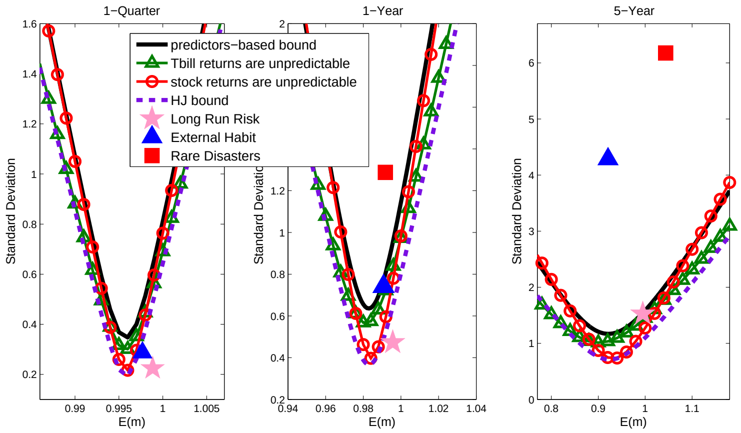

(a) SET A

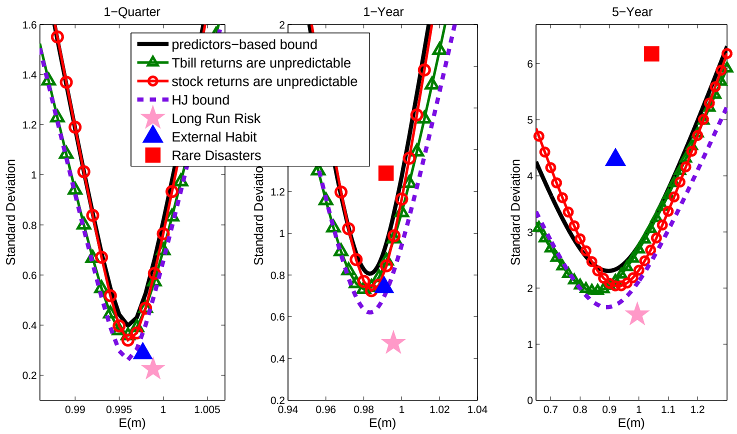

(b) SET B

Figure 1.8: 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 reports average mean and standard deviation values from 10 simulations run of 600,000 months with estimated parameters of the long run riks model. 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 SET A (see Panel A) and SET B (see Panel B). Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3.

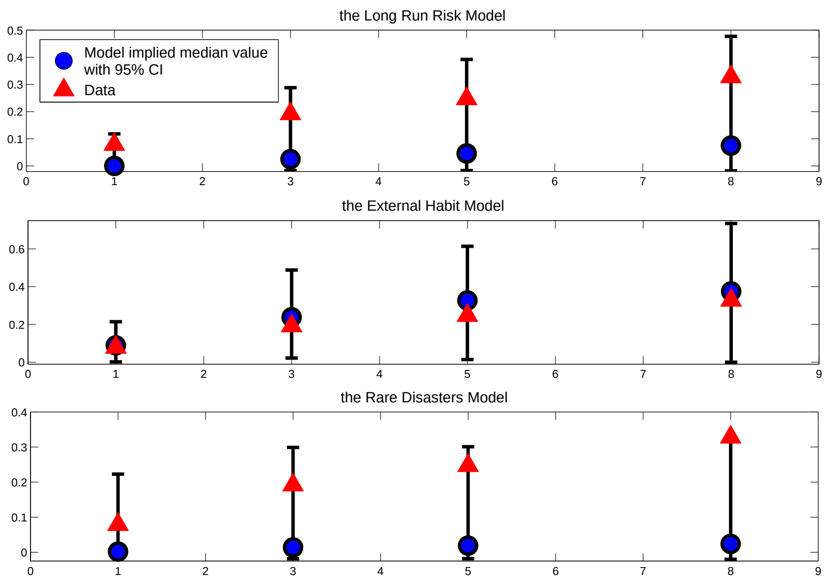

Figure 1.9: Historical versus model-implied predictability.
Predictability of gross returns, and Price-Dividend Ratios. Red triangle presents the adjusted $\mathcal{R}^2s$ from projecting one-, three-, five- and eight-year real gross return of the aggregate stock market portfolio onto lagged price-dividend ratio. The blue spot and black line provide the median values of adjusted $\mathcal{R}^2s$ 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. Sample: 1952Q2 - 2012Q3.

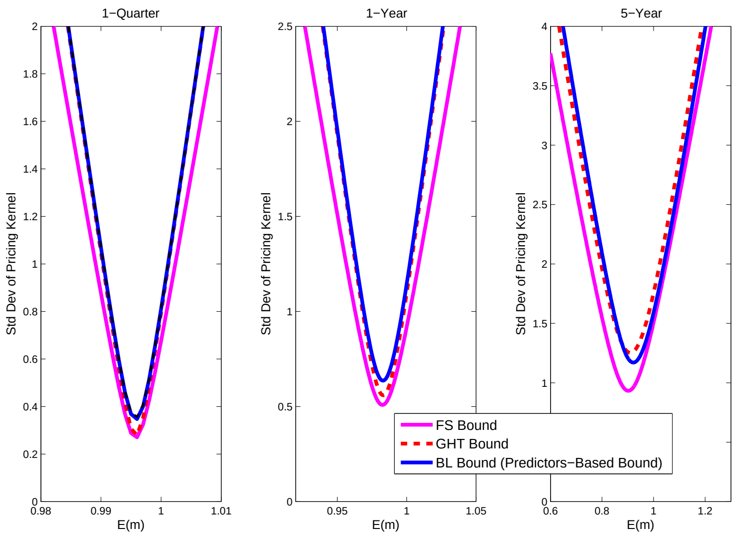

Figure 1.10: Alternative implementation of the HJ bounds – SET A.
This figure presents the volatility bounds using conditional information based on Ferson and Siegel (2001, 2003) (FS Bound), Bekaert and Liu (2004) (BL-predictors-based bound) and Gallant, Hansen, and Tauchen (1990) (GHT Bound) specifications, respectively. Sample: 1952Q2 - 2012Q3.

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

Table 1.1: Statistics of the Data.
This table reports sample statistics of quarterly nominal stock and 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: 1947Q2: 2012Q3.

Panel A: Predictive regressions for stock returns

Horizon h (in quarters)	pd_t [t-stat]	cay_t [t-stat]	$R^2(\%)$
1	-0.03	0.85	4.9
	(-1.99)	(2.93)	
4	-0.13	3.24	17.4
	(-2.37)	(2.75)	
20	-0.60	16.23	50.0
	(-4.70)	(4.17)	

Panel B: Predictive regressions for bond returns

Horizon h (in quarters)	spr_t [t-stat]	y_t [t-stat]	$R^2(\%)$
1	7.77	0.75	7.5
	(3.58)	(1.94)	
4	22.49	3.28	19.1
	(3.66)	(2.58)	
20	34.90	19.44	42.5
	(2.39)	(4.19)	

Table 1.2: Predictability of stock and bond returns.

Panel A reports quarterly overlapping regressions of multiple horizon real gross stock returns onto a constant, the log price-dividend ratio and cay_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 log short rate $y(t)$ and the yield spread $spr(t)$. The short rate is the log yield on the 30-day Treasury Bill from CRSP, and 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 Treasury Bill (T-bill). The table reports coefficient estimates, the R^2 of the regression, and, in brackets, the Newey-West corrected t-statistics. Sample: 1947Q2: 2012Q3.

	Long-Run Risk Bansal-Kiku-Yaron (calibration)	Long-Run Risk Bansal-Kiku-Yaron (estimation)	External Habit Campbell-Cochrane (calibration)	External Habit Campbell-Cochrane (estimation)	Rare Disaster Nakamura et al.	Lower Bound $\sigma_{\min}$ Data SET A	Lower Bound $\sigma_{\min}$ Data SET B
1-quarter	0.2546	0.2245	0.2773	0.2874		0.3475	0.3980
						[0.287 0.645]	[0.376 0.752]
1-year	0.5415	0.4734	0.7023	0.7412	1.2875	0.6358	0.8064
						[0.551 0.977]	[0.773 1.324]
5-year	1.7062	1.5274	4.8292	4.2860	6.1765	1.1701	2.3062
						[0.911 1.670]	[2.003 3.228]

Table 1.3: Model-implied SDFs and predictors-based bounds.

We compute the standard deviation of the SDFs across horizons, via simulations, respectively, for the models that incorporate long-run risk, external habit persistence and rare disasters, with both calibrated and estimated parameters (see Table 1.7), 1.8) and 1.9)). The reported values in the first five columns are the average values of standard deviation of model-implied SDFs from 10 single simulations run of 600,000 months. We then, in the last two columns, present the minimum standard deviation of the volatility bound using conditional information based on Bekaert and Liu's (2004) specification, together with associated 90% bootstrapped confidence intervals. To compute the confidence intervals of the volatility bound of SDFs, we create 50,000 random samples of sample size from the data, where the sampling in the block bootstrap is based on the optimal block length we calculated for each asset return regression residuals. Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3

Horizon (year)	Data β	Data R^2	Bansal-Kiku-Yaron β	Bansal-Kiku-Yaron R^2	Campbell-Cochrane β	Campbell-Cochrane R^2	Rare Disaster Nakamura et al. β	Rare Disaster Nakamura et al. R^2
1	−0.127 (0.053)	0.079	−0.080 (0.092)	0.000 [−0.016 0.118]	−0.291 (0.118)	0.089 [0.002 0.214]	−0.051 (0.294)	0.002 [−0.017 0.223]
3	−0.312 (0.084)	0.192	−0.274 (0.230)	0.025 [−0.017 0.288]	−0.732 (0.206)	0.238 [0.022 0.488]	−0.353 (0.486)	0.014 [−0.018 0.299]
5	−0.433 (0.075)	0.247	−0.523 (0.365)	0.046 [−0.017 0.392]	−1.020 (0.245)	0.327 [0.014 0.614]	−0.722 (0.582)	0.019 [−0.018 0.301]
8	−0.631 (0.101)	0.328	−0.929 (0.525)	0.075 [−0.018 0.477]	−1.3101 (0.264)	0.375 [0.000 0.734]	−1.275 (0.714)	0.024 [−0.020 0.323]

Table 1.4: Predictability of Gross Returns, and Price–Dividend Ratios.

This table presents adjusted $\mathcal{R}^2$s and slope coefficients from projecting model simulated one-, three-, five- and eight-year gross return of the aggregate stock market portfolio onto lagged model simulated price-dividend ratio. The results correspond to regressing $r_{t+1} + r_{t+2} \cdots + r_{t+h} = \alpha(h) + \beta(h) \log(P_t/D_t) + \varepsilon_{t+h}$, where r^e_{t+1} is the excess return, and h denotes the forecast horizon in year. Data statistics are reported in the "Data" column. The "Bansal-Kiku-Yaron" column presents predictability evidence implied by the Bansal Kiku and Yaron's (2012) model. The "Campbell-Cochrane" column presents predictability evidence implied by Campbell and Cochrane's (1999) model with estimated parameters reported in 1.8. "Nakamura et al." column presents predictability evidence implied by Nakamura et al.(2013) model. The entries for the models are based on 1,000 simulations each with 724 monthly observations that are time-aggregated to an annual frequency. The reported value for each model is the median of simulated slope and mean of simulated adjusted $\mathcal{R}^2$, along with 95% confidence interval in the square brackets. Standard errors which reported in parentheses are Newey-West corrected. Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3.

Panel A. Variance bound-based upper bound on R^2						
	HJ bounds		predictor-based bounds		regression R^2	
Horizon quarter	SET A	SET B	SET A	SET B	Stock Returns	Bond Returns
1	0.038	0.069	0.122	0.160	0.049	0.075
			[0.083 0.419]	[0.143 0.570]		
4	0.139	0.399	0.418	0.673	0.174	0.191
			[0.314 0.988]	[0.618 1.814]		
20	0.606	3.298	1.638	6.365	0.500	0.425
			[0.993 3.337]	[4.803 12.467]		

Panel B. Molde-based upper bound on R^2			
Horizon quarter	Bansal-Kiku-Yaron	Campbell-Cochrane	Rare Disaster Nakamura et al.
1	0.051	0.083	
	[0.018 0.101]	[0.079 0.088]	
4	0.232	0.569	1.716
	[0.069 0.491]	[0.513 0.627]	[1.473 1.978]
20	2.792	21.982	45.650
	[0.172 8.566]	[11.562 35.721]	[25.151 72.215]

Table 1.5: Upper bound on the $\mathcal{R}^2$ of return predictive regressions.
This table presents the upper bound on $\mathcal{R}^2$ of returns' predictive regressions. We compute the upper bound $R^2_{f,t+h}\sigma^2_Z(v_{\min})$ of $\mathcal{R}^2$, together with the corresponding 90% confidence intervals for both SET A and SET B. Panel A contains the results based on volatility bounds of SDFs. For the unconditional HJ bounds, we report in the first two columns the minimum value for each horizon, 1-quater, 1-year and 5-year. For the predictor-based bounds, we report, for each horizon, the minimum value of the predictor-based bounds, $\sigma^2_Z(v_{\min})$, along with the values of the 90% bootstraped confidence intervals (square brackets). For the scaled variance of the models, $R^2_{f,t+h}Var(m^X_{t+h})$, we report the values for the Bansal et al. [2012a] model, the Campbell and Cochrane [1999] model and the Nakamura et al. [2013] model (see Panel B). The 90% confidence interval of these scaled variances are reported in the brackets. The last two colomns report, from Table 2, the predictive regression $\mathcal{R}^2$ of stock returns and 5-year constant maturity government bond returns. For the risk-free rate, we adopt the mean of the 3-month T–bill returns. Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3.

	Lower Bound $\sigma_{\min}$	
	SET A	SET B
1-quarter	0.2705	0.4240
	[0.232 0.517]	[0.414 0.695]
1-year	0.5087	0.8395
	[0.440 0.793]	[0.790 1.296]
5-year	0.9326	1.8116
	[0.777 1.347]	[1.741 2.851]

Table 1.6: FS Bounds with conditioning information.
This table compute the volatiltiy bounds of SDFs using conditioning information, based on Ferson and Siegel's (2003) specification for both SET A and SET B. The reported values are the values of the minimum point of the volatility bound at each horizon, with associated bootstraped 90% confidence intervals. To compute the confidence intervals, we create 50,000 random samples of sample size from the data, where the sampling in the block bootstrap is based on the optimal block length we calculated for each asset return regression residuals. Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q3.

	Parameter	BKY2012a *calibration*	BKY2012b *estimation*
Preferences			
Time preference	δ	0.9989	0.9989 (0.0010)
Risk aversion	γ	7.5	7.42 (1.55)
EIS	ψ	1.5	2.05 (0.84)
Consumption growth dynamics, g_t			
Mean	μ	0.0015	0.0012 (0.0007)
Dividends growth dynamics, $g_{d,t}$			
Mean	μ_d	0.0015	0.0020 (0.0017)
Persistence	ϕ	2.5	4.45 (1.63)
Volatility parameter	φ_d	5.96	5.00 (1.39)
Comsumption exposure	π	2.6	0
Correlation between innovations	ρ_{dc}		0.49 (0.33)
Long-run risk, x_t			
Persistence	ρ	0.975	0.9812 (0.0086)
Volatility parameter	φ_e	0.038	0.0306 (0.0160)
Consumption growth volatility, σ_t			
Mean	$\overline{\sigma}$	0.0072	0.0073 (0.0015)
Persistence	v	0.999	0.9983 (0.0021)
Volatility parameter	σ_w	0.28×10^{-5}	2.62×10^{-5} (3.10×10^{-6})

Table 1.7: Appendix-I: Parametrization of long-run risk asset pricing model. Our parameterization of calibrated values is taken from Bansal et al. [2012a]. The value of the parameter estimates and corresponding standard errors (in parentheses) are taken from Bansal et al. [2012c]. The model is simulated at the monthly frequency.

	Parameter	Campbell-Cochrane *calibration*	*estimation*
Preference			
Time preference	δ	0.9903	0.9903 (0.0004)
Risk aversion	γ	2	1.9756 (0.0772)
Consumption growth dynamics, g_t			
Mean	$\overline{g}$	0.0016	0.0017 (0.00007)
Volatility parameter	σ	0.0043	0.0050 (0.00019)
Dividend growth dynamics, Δd_t			
Volatility parameter	σ_w	0.0323	0.0319 (0.0014)
Corr between innovitions	ρ_{dc}	0.2	0.1945 (0.0093)
Steady state surplus consumption ratio	$\overline{S}$	0.0570	0.0637
Persistence in consumption surplus ratio	ϕ	0.9884	0.9877 (0.0003)
Log of risk-free rate	$r^f \times 10^2$	0.0783	0.0854 (0.0365)

Table 1.8: Appendix-II: Parametrization of external habit asset pricing model. Our parameterization of calibrated values is taken from Campbell and Cochrane [1999]. The value of the parameter estimates and corresponding standard errors (in parentheses) are taken from Aldrich and Gallant [2011a]. 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 1.9: Appendix-III: 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.

Chapter 2

A Robust Variance Bound on Pricing Kernels

2.1 Introduction

In a seminal paper, Hansen and Jagannathan [1991] (HJ) proposed a method to assess whether an asset pricing model is consistent with observed returns. This is the so called HJ bound: a lower bound on the variance of the model-implied pricing kernel as a quadratic function of its mean. In its simplest version, the HJ bound defines a feasible region on the mean–variance plane of pricing kernels by using only the unconditional first and second moments of available asset returns. A series of subsequent studies (Gallant et al. [1990], Ferson and Siegel [2003], Bekaert and Liu [2004]) show that asset return predictability translates into tighter bounds. The tighter the variance bound, the stronger the rejection of the asset pricing models.

In this chapter, I investigate how the information contained in a set of return predictors can be used to efficiently discriminate among asset pricing models by proposing a data-based performance measure. Such a measure is based on pricing kernels and reveals how robustly a model generated pricing kernel can explain returns on financial assets. More specifically, I develop a robust variance bound on pricing kernels by taking into account asset return predictability and the predictive model uncertainty.

To understand the intuition behind the properties of predictor-based variance bound, it is useful to look at Figure 3.3(a). The predictor-based variance bounds are computed based on two assets, the US stock market index and the roll over three-month nominal risk-free rate.[1] I employ three different linear predictive systems to compute the first two conditional moments of asset returns. Then, I follow the duality-based approach of Bekaert and Liu [2004] to construct the predictor-based variance bound by using the conditional moments of asset returns. In the left panel of Figure 2.1, the tightest variance

[1]For a detailed description of the data, see Appendix I.

bound (red dashed curve) is associated with the highest system-wide R^2 (0.51), while the loosest bound (black dotted curve) is connected with the lowest system-wide R^2 (around 0). The blue solid parabola is linked with the case in which the system-wide R^2 equals 0.31. These empirical results are especially interesting for two reasons. First, a high degree of predictability of the asset returns (indicated high R^2 value in the predictive regression) would result a stringent variance bound. Second and more important, the diagnostic results of a given asset pricing model would be different compared with different predictor-based variance bounds. A question, here, naturally arises: Without knowledge of the correct return predictive system, how can the information contained in a set of predictors be efficiently used to discriminate among competitive asset pricing models?

I address this question by proposing a data-based robust variance bound for asset pricing models. Among all the candidate variance bounds computed based on different return predictive systems, the *robust variance bound* is the tightest one that the pricing kernel of *a given asset pricing model* can satisfy. To quantify the relation between a given asset pricing model and the degree of return predictability explained by such a model–generated pricing kernel , I then develop a *model performance index* relying on the diagnostic results of the robust variance bound. The value of the index is the proportion of the number of bounds above the robust variance bound for a given asset pricing model. The smaller the value of the index, the tighter the robust variance bound and the higher the degree of return predictability explained by such a model–generated pricing kernel.

In practice, the most successful predictors for stock and bond returns follow from accounting identities rather than from specific asset pricing models (Campbell [1987]). For instance, the predictive relation between the (log) price–dividend ratio and stock market returns follows from solving forward a linearized expression for returns (Campbell and Shiller [1988a]). The predictive relation between the consumption–wealth ratio, the cay variable of Lettau and Ludvigson [2001], and stock market returns follows from a linearized version of the consumer's intertemporal budget constraint. Similarly, a successful bond market predictor such as the term spread follows from a linearization argument similar to that which generates the dynamic dividend growth model for the stock market (Campbell and Ammer [1993]). Although these studies identify some potential predictors for both stock and bond returns, the 'correct' specification remains an open issue (Xia [2001], Avramov [2002], Barberis [2002], Welch and Goyal [2008a], Diris [2011], Ghysels et al. [2012], Pástor and Stambaugh [2012]). To reduce the selection bias of the return predictive system, I employ a set of relevant return predictors for both stocks and bonds and study all possible predictive systems based on them. Therefore, the robust variance bound is robust to return predictability and predictive model uncertainty.

The robust variance bound has several interesting properties. First, each given asset pricing model has its own robust variance bound on the pricing kernel, unless it is unable to reach any available candidate bounds. Second, the robust variance bound is data based.

Different choices of asset returns and return predictors lead to different robust variance bounds for a given asset pricing model. Third, a stringent robust variance bound of an asset pricing model is linked with a high degree of return predictability. This means that this model-implied pricing kernel is capable of pricing the given set of assets, even considering return predictability.

To provide a wider foundation for my empirical examination, I apply my method to a broad class of asset pricing models, namely, (i) long-run risk (Bansal and Yaron [2004], Bansal et al. [2012d], Bansal et al. [2012b]) , (ii) external habit formation (Campbell and Cochrane [1999], Bekaert and Engstrom [2010a]), and (iii) rare disasters (as parameterized by Backus et al. [2011a] and Nakamura et al. [2013]). Hansen [2008] and Bakshi and Chabi-Yo [2012] point out the importance of investigating long-run risk, external habit formation, and rare disasters under a common platform. I am then led to ask two questions of economic interest: (i) What can one learn about these asset pricing models that consistently price the risk-free bond, the equity market, and other assets by incorporating return predictability? (ii) How can one improve the ability of asset pricing models to explain asset returns with a higher degree of predictability?

My analysis yields a number of new insights about the performance of asset pricing models when equity and bond data are used. First, long-run risk models struggle to satisfy the variance bound at each horizon when asset returns are predictable, and high risk aversion does not improve their performance. The reason is that an increase in the variance of the pricing kernel cause an increase in its mean. Second, the external habit type of asset pricing model performs well enough to be consistent with return predictability at long horizons, but fails to reach the variance bound related to high return predictability at short horizons. In addition, asymmetric reactions to good and bad news modifications of consumption growth dynamics cannot increase the volatility of the model-implied pricing kernel. Finally, the long–run risk type of the rare disaster model of Nakamura et al. [2013] exhibits outstanding performance at each horizon. This occurs because this model allows disasters to unfold over multiple years and to be systematically followed by recoveries. The large drops in consumption raise the volatility of the model-implied pricing kernel at a short horizon, and the recoveries after disaster in consumption drive the volatility of model-implied pricing kernel up at a long horizon. However, the mean of the pricing kernel is not raised by the increase in its variance. Therefore, this model obtains the most stringent robust variance bound at each horizon among all asset pricing models. The model-implied pricing kernel can explain asset returns with high return predictability.

The paper makes four main contributions. First, it contributes to the literature of asset pricing model diagnostics by quantifying the relation between a given asset pricing model and the degree of return predictability explained by such a model–implied pricing kernel. It is the first paper to evaluate asset pricing model performance based on asset return predictability and predictive model uncertainty. Second, apart from contributing to theoretical study related to variance bounds on pricing kernels, it also lists many interesting empirical properties of the bound for predictable returns. For instance, besides return predictability, an increase in the number of assets also leads to a tight variance bound. Meanwhile, stock return predictability increases the variance bound, while bond return predictability narrows it. Third, this chapter also contributes to the literature of consumption-based asset pricing models. This performance measure provides new insight into the behavior of popular asset pricing models. It empirically shows for each asset pricing model the kind of modification (such as a higher risk aversion coefficient, higher correlation between different sources of shock, and the strong persistency of a certain state variable) that would lead to a higher variance of the pricing kernel to match the higher degree of return predictability. Last but not least, this chapter contributes to the literature of return predictive model seeking. It provides a new methodology for searching of return predictive systems for each given asset pricing model with certain preferences and dynamics.

The rest of this chapter is organized as follows. Section 2 presents a simple example to motive my study and bring forth the research question. Section 3 defines a robust variance bound on the pricing kernel and proposes an asset pricing model performance index related to return predictability. Specifically, it first clarifies the relation between the variance bound on the pricing kernel and return predictability. Then, it provides a formal definition of a robust variance bound and presents the features of asset pricing models studied in this chapter. Section 4 summarizes the main empirical findings by applying my methodology to three leading class of asset pricing models. Section 5 provides insights into how economic fundamentals are linked to pricing kernels and how the performance of an asset pricing model could be improved by altering the properties of return predictability. Section 6 concludes the paper.

2.2 Illustration with a Simple Example

I start by presenting a simple example to motivate my study and bring up the research question. This simple example involves two assets, the stock market index and the roll over three-month nominal risk-free rates. I use two well-known return predictors: the log price–dividend ratio and the term spread between long-term and short-term government bonds for both stocks and bonds. The choice of these two predictors is motivated by present value logic (for stock markets and bonds, respectively, see Campbell and Shiller [1988b] and Campbell and Ammer [1993]).

Let $R_{t+h} = \left[\begin{array}{cc} R^S_{t+h} & R^B_{t+h} \end{array} \right]'$ be a vector of realized asset return and $Z_t = \{pd_t,\, spr_t\}$ be a set of suspected return predictors. I use quarterly data from 1952Q1 to 2012Q4. All returns are gross and deflated using the consumer price index (CPI).[2]

For simplification, I use the linear return predictive system of asset returns

$$R^S_{t+h} = \mu^S_{t+h} + \Phi^S z^S_t + \varepsilon^S_{t+h}$$
$$R^B_{t+h} = \mu^B_{t+h} + \Phi^B z^B_t + \varepsilon^B_{t+h}$$

where $z^i_t \in \left\{0, pd_t, spr_t, \left[\begin{array}{cc} pd_t & spr_t \end{array}\right]\right\}$, for any $i = S, B$ and $\varepsilon_{t+h} = \left[\begin{array}{cc} \varepsilon^S_{t+h} & \varepsilon^B_{t+h} \end{array}\right]'$ is normally distributed with conditional mean zero and variance–covariance matrix Σ_h. There are $2 \times 2^2 = 16$ regression specifications. In this example, I only study the case in which the holding horizon is one year , that is, $h = 4$. The predictive system of asset returns is estimated by using seemingly uncorrelated regression (SUR). I report the system-wide R^2 values of each regression specifications and associated return predictors in Table 2.1. The system-wide R^2 values increase from about 0.0% when no predictor is included to about 38.9% when all predictors are adopted for both stock and bond returns. The value of R^2 increases when more return predictors are included. Meanwhile, using *pd* to predict the stock market and *spr* for bonds, the system-wide R^2 reaches 31.0%, while it is only 11.0% if *spr* is used to predict the stock market and *pd* is used for bonds. This reveals that misspecification of asset returns reduces the system-wide R^2.

[Insert Table 2.1 about Here]

To compute the predictor-based bounds, I solve the left-hand side of Eq. (2.4). The variance of the pricing kernel is a quadratic function of its unconditional mean. The coefficients in this quadratic function are determined by the conditional moments of asset returns based on each return predictive system. Figure 1 plots all 16 variance bounds on the pricing kernels based on different regression specifications. The shape of the variance bounds on the pricing kernels essentially depends on the predictive system chosen to predict the asset returns. The tightest variance bound (red curve) is associated with the highest system-wide R^2, while the lowest bound (blue cup) is connected with the lowest system-wide R^2. The higher the system-wise R^2, the tighter the variance bound on the pricing kernels.

Furthermore, I use these variance bounds to analyze three asset pricing models: the long-run risk model of Bansal et al. [2012b], the external habit model of Campbell and Cochrane [1999], and the rare disaster model of Nakamura et al. [2013]. The three triangles in Figure 2.2 indicate the average mean and variance of the model-implied pricing kernels.[3] These empirical results are interesting for two reasons. First, asset return predictability translates into a stringent variance bound on the pricing kernels. Second and

[2]For a detailed description of the data construction, see Appendix I.

[3]The triangles are obtained by averaging the mean and standard deviation of the model-implied pricing kernels obtained from 10 long (50,000–year) simulations. For further details on the asset pricing models, see Sections 3.3 and 4.4.

more important, the diagnostic results of a given asset pricing model would be different for different predictor-based variance bounds, as for the case of the external habit model in Figure 2.2. The question I address is which legitimate variance bound can be used to diagnose a given asset pricing model.[4]

[**Insert Figure 2.2 about Here**]

In this chapter, I propose a data-based *robust variance bound* on pricing kernels. The robust variance bound for a given asset pricing model is the tightest variance bound among all available bounds the model-implied pricing kernel can satisfy . In this example, the pricing kernel of the long-run risk model fails to lie within any of the variance bounds. The robust variance bound of the external habit model is the black dotted curve plotted in Figure 2, while the robust variance bound of the rare disaster model is the tightest red cup. Therefore, the rare disaster model satisfies all the variance bounds related to the different return predictive systems and has the tightest robust variance bound of the three asset pricing models. Furthermore, I quantify the asset pricing model performance related to return predictability by using the robust variance bound. The model performance index α is

$$\alpha = \frac{K}{2^{2\times 2}} \in [0,1]$$

where K is the number of variance bounds that lie above the robust variance bound of each given asset pricing model and α represents the proportion of variance bounds that reject this model-implied pricing kernel. When α equals zero, the model-implied pricing kernel satisfies all the variance bounds, such as the rare disaster model, in the example. On the other hand, when α equals zero , the model-implied pricing kernel is rejected by all variance bounds, for instance, the long-run risk model in the example. For the external habit model, α equals 0.375. Specifically, the external habit model no longer falls within the variance bounds once the log price–dividend ratio is included as a predictor of stock market return. Among these three asset pricing models, the rare disaster model is capable of satisfying the variance bound associated with the highest degree of return predictability, while the long-run risk model is rejected by financial data, even when no conditioning information is involved.

[**Insert Figure 2.3 about Here**]

In sum, the robust variance bound on the pricing kernels is a criterion for discriminating among asset pricing models based on the return predictability of financial data. The model performance index is an intuitive measure that relates asset pricing models with the degree of return predictability of financial data. This index shows how robustly a given asset pricing model is in pricing a given set of asset returns.

[4] The tightest variance bound is not always adopted to diagnose asset pricing models because of the issue of overfitting (seeWhite [2000]).

2.3 Robust Variance Bounds and Asset Pricing

This section presents theoretical bounds related to the unconditional variance of pricing kernels. First, I document the properties of predictor-based variance bounds, which are 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 Z_t of return predictors. Second, I define the robust variance bounds on the pricing kernels under a general framework and then construct the model performance index. Finally, I document the features of the three classes of asset pricing models evaluated by my method.

2.3.1 Predictor-Based Variance Bound on Pricing Kernels

In modern asset pricing theory, asset returns contain information about the pricing kernel that generate them. In this section, I present the theoretical variance bounds on the pricing kernels, which are bounds on the variance of the SDFs that price a given set of returns conditional on the information contained in a vector Z_t of the returns' predictors.

I adopt notation similar to that Alvarez and Jermann [2005] and let $\{M_t\}$ be the process of strictly positive pricing kernels. Similar to Hansen and Richard [1987], I use the absence of arbitrage opportunities to specify the current return of an asset that is R_{t+h} at time $t+h$ as

$$\begin{aligned} 1 &= E_t\left[\frac{M_{t+h}}{M_t}R_{t+h}\right] \\ &= E_t\left[m_{t+h}R_{t+h}\right] \end{aligned}$$

where $E_t[\cdot]$ represents the conditional expectation operator. The SDF from t to $t+h$ is represented by $m_{t+h} = \frac{M_{t+h}}{M_t}$.

Now, I consider a set of N assets traded at a given time t and denote the return on each asset by $R_{j,t+h}$, with investment horizon $h = 1, 2...$, and let R_{t+h} denote the vector collecting these N returns. Alongside the returns, I consider a vector Z_t of the returns' predictors and denote their informational content by $\mathcal{I}_t^Z$

I denote by $\mathcal{M}^Z$ the set of SDFs that price returns conditionally on the realizations of the predictors Z_t, that is,

$$\mathcal{M}^Z = \left\{m_{t+h} \;\middle|\; E(m_{t+h}^2) < \infty,\ E\left(m_{t+h}R_{t+h} \left|\mathcal{I}_t^Z\right.\right) = e\right\} \tag{2.1}$$

where e denotes the unit vector. I assume $\mathcal{M}^Z$ is non-empty, which means that the law of one price holds in the linear space of payoffs obtained by managed portfolios conditional on the predictors' realization.

Given the SDFs in $\mathcal{M}^Z$, the predictor-based variance bound, denoted $\sigma_Z^2(v)$, is the lower envelope of the set of all variances of SDFs in $\mathcal{M}^Z$,, that is, the map

$$\sigma_Z^2(v) = \inf \left\{ Var\left(m_{t+h}\right) \mid m_{t+h} \in \mathcal{M}^Z,\ E\left(m\right) = v,\ v \in \Re \right\} \tag{2.2}$$

where the variable v is the shadow price of a unit risk-free zero-coupon bond with maturity $t+h$. The parabolic function $\sigma_Z^2(v)$ represents an unconditional frontier for SDFs in the sense of Gallant et al. [1990]: Since it considers all SDFs that price returns conditionally on Z_t, it takes full advantage of the predictive power of the vector Z_t while maintaining the simplicity of concentrating on the unconditional moments of such SDFs. By saying that Z_t predicts the return $R_{j,t+h}$ on some asset j, I mean $Var\left[E\left(R_{j,t+h}\left|\mathcal{I}_t^Z\right.\right)\right] > 0$ over some holding period h.[5] The HJ variance bound is an extreme case under my setup, when $Var\left[E\left(R_{j,t+h}\left|\mathcal{I}_t^Z\right.\right)\right] = 0$. Furthermore, if $\mathcal{I}_t^{Z_1} \subset \mathcal{I}_t^{Z_2}$, then $\mathcal{M}^{Z_1} \subset \mathcal{M}^{Z_2}$. Hence, the lower variance bound generated based on Z_1, $\sigma_{Z^1}^2(v)$, is higher than that computed based on the predictor set Z_2. The predictability of asset returns translates into a tighter variance bound on the pricing kernels.

As observed by Bekaert and Liu [2004], when the conditional moments of returns are not correctly specified, the predictor-based bound $\sigma_Z^2(v)$ can fail to be a valid lower bound for the volatility of SDFs in $\mathcal{M}^Z$. To obviate this problem, one can extend this conditional setting with the duality between the mean–variance frontiers for SDFs and maximum Sharpe ratios first illustrated for the unconditional case in HJ's seminal work. More specifically, I define the following set of returns from the managed portfolios:

$$\mathcal{R}^Z = \left\{ R_{t+h}^w \mid R_{t+h}^w = w_t' R_{t+h},\ w_t \text{ is } \mathcal{I}_t^Z - \text{measurable s.t. } E\left(w_t' e\right) = 1 \right\} \tag{2.3}$$

This set collects all the payoffs that are generated by trading strategies that exploit the information contained in the predictors at time t. As long as $\nu \neq 0$, one can show (Bekaert

[5] I assume that all the random variables are defined over a probability space $(\Omega, \mathcal{I}, \mathcal{P})$, so that $\mathcal{I}_t^Z$ is formally the $\sigma-$ algebra generated by the vector of predictors Z_t and hence $\mathcal{I}_t^Z \subset \mathcal{I}$. I also assume that returns have finite unconditional first and second moments, μ and $E^2 \triangleq E\left[R_{t+h}R_{t+h}'\right]$, with an unconditional variance–covariance matrix $\Sigma = E^2 - \mu\mu'$. Moreover, I assume that the matrix $E_t^2 \triangleq E\left[R_{t+h}R_{t+h}'\left|\mathcal{I}_t^Z\right.\right]$ of second moments conditional on the predictors is (almost surely) positive- definite and hence invertible; that is, the returns are linearly independent conditional on information in the predictors. Therefore, denoting by $\mu_t \triangleq E\left[R_{t+h}\left|\mathcal{I}_t^Z\right.\right]$ the vector of conditional expected returns, both the conditional variance–covariance matrix $\Sigma_t = E_t^2 - \mu_t\mu_t'$ and its unconditional counterpart Σ are positive definite (and hence invertible) as well.

and Liu [2004], Abhyankar et al. [2007], Peñaranda and Sentana [2013]) that[6]

$$\sigma_Z^2(v) = \nu^2 \sup_{R_{t+h}^w \in \mathcal{R}^Z} \left(\frac{E(R_{t+h}^w) - \nu^{-1}}{Var\left(R_{t+h}^w\right)} \right)^2 \tag{2.4}$$

In other words, for any given level of the risk-free rate, the predictors-based bound is proportional to the square of the maximum Sharpe ratio that can be generated by the managed portfolios that exploit the information contained in the predictors Z_t.

Observing that a misspecification of the conditional expected returns and variances introduces a duality gap in Eq. (2.4), Bekaert and Liu [2004] suggest always using the right-hand side to compute a variance bound that incorporates conditioning information, since, by the very definition of sup, this right–hand side always constitutes a valid lower bound on the variance of the SDFs in $\mathcal{M}^Z$ (albeit not necessarily the highest lower bound if the first two conditional moments of the returns are misspecified). Therefore, in the empirical part of this chapter, I follow the lead of Bekaert and Liu [2004], and I compute the predictor-based bounds using the solution to the dual problem defined in Eq.(2.4).

The above predictor-based bound is conservative, since it is based on those SDFs that price trading strategies which are required to be "buy and hold" over the horizon h. The predictor-based bounds that would be obtained from those SDFs that price the (truly) dynamic trading strategies that are allowed to be rebalanced at the intermediate dates $t+1, t+2, ..., t+h$ would clearly impose much stricter criteria . Allowing for intertemporal rebalancing would expand the set $\mathcal{R}^Z$ of managed returns and, via the duality in Eq. (2.4), would yield a much higher bound $\sigma_Z^2(v)$. In this chapter, however, I concentrate on "buy and hold" strategies and leave the more general framework to future research.

2.3.2 Robust Variance Bound on Pricing Kernels

To compute the predictor-based bounds, I use the solution to the left-hand side of Eq. (2.4). Bekaert and Liu [2004] show that the optimal trading strategy w_{t+h} that incorporates information is:

$$w_{t+h} = \left(\mu_{t+h}\mu_{t+h}^\intercal + \Sigma_{t+h}\right)^{-1} (1 - \kappa\mu_{t+h})$$

[6]When $\nu = 0$, the supremum coincides with the reciprocal of the global minimum portfolio variance over the set $\mathcal{R}^Z$. Moreover, the sup is always attained, with the single exception of the case in which ν is set equal to the expected return on the global minimum variance portfolio, where the supremum is attained by a return whose expected price is zero.

where μ_{t+h} and Σ_{t+h} are the vector of conditional expected returns and the conditional variance–covariance matrix, respectively, and

$$\begin{aligned}\kappa &= (\nu - b)/(1-d), \nu = E[m_{t+h}] \\ b &= E\left[e'\left(\mu_{t+h}\mu_{t+h}^{\intercal} + \Sigma_{t+h}\right)^{-1}\mu_{t+h}\right] \\ d &= E\left[\mu_{t+h}'\left(\mu_{t+h}\mu_{t+h}^{\intercal} + \Sigma_{t+h}\right)^{-1}\mu_{t+h}\right]\end{aligned}$$

The predictor-based variance bound depends on the conditional distribution function only through the first and second conditional moments. It is decreasing in the conditional variance and is not monotonic in the conditional means. The first and second conditional moments of asset returns, μ_{t+h} and Σ_{t+h}, can be computed by using the selected predictive system of asset returns. A different choice of return predictive system leads to different conditional moments of the asset returns. For simplicity, I assume the conditional covariance matrix for the returns is constant, and estimate that it has residual in the forecasting regressions . In this chapter, I use a linear predictive system to predict future rates of returns on both stocks and bonds.

Consider the following representation of predictive regressions for arbitrary asset i at different horizons:

$$\begin{aligned}R_{t+h}^{i} &= \beta_{0,h}^{i} + \beta_{1,h}^{i} z_{t}^{i} + v_{t}^{i} \\ v_{t}^{i} &\sim N\left(0, \sigma_{i}^{2}\right) \\ \mu_{t+h}^{i} &= \beta_{0,h}^{i} + \beta_{1,h}^{i} z_{t}^{i}\end{aligned} \tag{2.5}$$

where $z_t^i \subseteq Z_t$ is a vector of predictors for the return R_{t+h}^i observable at time t. The conditional first and second moments of traded returns are (μ_{t+h}, Σ_h). The empirical evidence indicates that the performance of model (2.5) in predicting returns improves with the horizon. At a short horizon, there is very little difference between the conditional and unconditional moments, since the predictability is negligible; however as the forecasting horizon increases, substantial predictability materializes . As illustrated in the simple example, return predictability drives the difference between the variance bounds with and without the incorporation of conditioning information. Meanwhile, the choice of the predictor vector z_t^i of each asset return also drives the differences among predictor-based variance bounds.

Suppose Q explanatory variables are potential asset return predictors are contained in Z_t. There are $2^{N\times Q}$ competing regression specifications. Hence, $2^{N\times Q}$ variance bounds on pricing kernels can be computed by using the first and second conditional moments of asset returns obtained with each predictive system. Now, consider the side of the asset pricing model. To formalize this discussion, I identify any given asset pricing model with the triple $\left(X_t, \mathcal{F}_t^X, m_{t+h}^X\right)$, where X_t denotes the set of state variables of the model, $\mathcal{F}_t^X$ denotes the informational content in the state variables X_t,[7] and m_{t+h}^X denotes the SDF of the given asset pricing model. Moreover, I let $\nu^X \equiv E\left(m_{t+h}^X\right)$ denote the price assigned by the SDF m_{t+h}^X to a unit zero-coupon bond with maturity $t+h$. With this notation, I can now define the robust variance bound.

For a given asset pricing model, the diagnostic results vary according to the model–generated pricing kernel, m_{t+h}^X, with different variance bounds, $\sigma_{\widetilde{Z}}^2(v)$. There are a total of $2^{N\times Q}$ candidates . The *robust variance bound*, denoted $\sigma_{X,Z}^2(\nu)$, for this given asset pricing model is the tightest variance bound the model–implied pricing kernel, m_{t+h}^X, is able to satisfy.[8]

$$\sigma_{Z,X}^2(v) = \sup\left\{\sigma_{\widetilde{Z}}^2(v)\middle|\, \widetilde{Z} \subseteq Z, Var\left(m_{t+h}^X\right) \geqslant \sigma_{\widetilde{Z}}^2(\nu), v \in \Re\right\}$$

I now discuss several properties of the robust variance bound. First, each given asset pricing model has its own robust variance bound on the pricing kernel, unless it is unable to reach any available candidate bounds, that is $Var\left(m_{t+h}^X\right) < \sigma_{\widetilde{Z}}^2(\nu)$ for $\forall \widetilde{Z} \subseteq Z$. Second, the robust variance bound is data based. Different choices of R_{t+h} and Z_t lead to different robust variance bounds for a given asset pricing model. Third, a stringent robust variance bound of the asset pricing model is associated with a high degree of return predictability. As shown in the previous section, a higher degree of return predictability translates into a tighter variance bound. Let $\widetilde{Z}^X$ be the set of predictors linked to the robust variance bound of the given asset pricing model, that is, $\sigma_{Z,X}^2(v) = \sigma_{\widetilde{Z}^X}^2(\nu)$. If $\sigma_{Z,X^1}^2(v) > \sigma_{Z,X^2}^2(v)$, the degree of predictability of the predictive system based on $\widetilde{Z}^{X^1}$ is higher than that based on $\widetilde{Z}^{X^2}$. Therefore asset pricing model 1 is capable to pricing a given set of returns even though more effective conditioning information is incorporated. The robust variance bound for a given asset pricing model is robust to the existence of return predictability and predictive model uncertainty.

Then, based on the robust variance bound of the asset pricing model, I construct an model performance index, denoted α, which is the ratio of the number of variance bounds that lie above the robust bound to the total number of available candidate bounds:

$$\alpha = \frac{K}{2^{N\times Q}} \in [0, 1]$$

where K is also the number of the variance bounds in which such a model-generated pricing kernel is not located . There two extreme cases: When $\alpha = 0$, the model-implied pricing kernel is able to satisfy all the available variance bounds; when $\alpha = 1$, this model-

[7] Formally, $\mathcal{F}_t^X$ is the σalgebra generated by the vector of state variables X_t and hence $\mathcal{F}_t^X \subset \mathcal{F}$, where $\mathcal{F}$ is the σ algebra of the probability space $(\Omega, \mathcal{F}, \mathcal{P})$ over which all the random variables are defined.

[8] As a remark, if $\widetilde{Z} \subseteq \widetilde{Z}' \subseteq Z$, then $\sigma_{\widetilde{Z}}^2(\nu) < \sigma_{\widetilde{Z}'}^2(\nu)$. However, if there is no guarantee that $\widetilde{Z} \subseteq \widetilde{Z}'$ or $\widetilde{Z}' \subseteq Z$, the two parabolas, $\sigma_{\widetilde{Z}}^2(\nu)$ and $\sigma_{\widetilde{Z}'}^2(\nu)$, may intersect.

implied pricing kernel fails to obtain any available variance bounds. Therefore, a small (large) value of α is connected with a high (low) degree of return predictability. This model performance measure shows how robustly a given asset pricing model can satisfy the variance bounds when asset returns are predictable. Therefore, I can discriminate among asset pricing models by comparing the α of each model.

2.3.3 Asset Pricing Models

I apply this methodology to three popular classes of asset pricing models, namely, (i) the long-run risk model (Bansal and Yaron [2004], Bansal et al. [2012d], Bansal et al. [2012b]), (ii) the external habit formation model (Campbell and Cochrane [1999], Bekaert and Engstrom [2010a]), and (iii) the rare disaster model (as parameterized by Backus et al. [2011a] and Nakamura et al. [2013]). All the asset pricing models examined in this chapter have recursive preference. The first class consists of the long-run risk models, proposed by Bansal and Yaron [2004] and extended by Bansal et al. [2012d] and Bansal et al. [2012b]. The dynamics of this model are described as the following:

$$\begin{aligned}
x_{t+1} &= \rho x_t + \varphi_e \sigma_t e_{t+1} \\
g_{t+1} &= \mu + x_t + \sigma_t \eta_{t+1} \\
g_{d,t+1} &= \mu_d + \phi x_t + \varphi_d \sigma_t u_{t+1} + \pi \sigma_t \eta_{t+1} \\
\sigma_{t+1} &= \overline{\sigma}^2 + v_1 \left(\sigma_t^2 - \overline{\sigma}^2\right) + \sigma_w w_{t+1}
\end{aligned}$$

where g_{t+1} and $g_{d,t+1}$ are the logarithms of consumption growth and dividend growth , respectively, and $e_{t+1}, u_{t+1}, \eta_{t+1}, w_{t+1} \sim i.i.d.\ N(0,1)$. Following Bansal et al. [2012d], x_t is a small persistent component and σ_t^2 is the stochastic volatility. According to Bansal and Yaron [2004], dividend growth dynamics are not exposed to the consumption growth shock, that is, $\pi = 0$. The SDF is

$$\ln\left(\frac{M_{t+1}}{M_t}\right) = \theta \ln \delta - \frac{\theta}{\psi} g_{t+1} + (\theta - 1)\, r_{a,t+1}$$

where $r_{a,t+1}$ is the return on the aggregate consumption claim, which is unobservable, δ is the time discount factor, γ is the risk aversion parameter, and $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$.

A second class of asset pricing models consists of the external habit formation type. Campbell and Cochrane [1999] were the first to propose a model with external habit formation. The driving processes for this model is

$$\begin{aligned}
g_{t+1} &= \overline{g} + v_t \\
g_{d,t+1} &= \overline{g} + w_t \\
s_{t+1} - \overline{s} &= \phi \left(s_{t++2} - \overline{s}\right) + \lambda \left(s_t\right) v_t
\end{aligned}$$

where g_{t+1} and $g_{d,t+1}$ are the logarithms of consumption growth and dividend growth , respectively, and $\begin{bmatrix} v_t \\ w_t \end{bmatrix} \sim NID\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma^2 & \rho\sigma\sigma_w \\ \rho\sigma\sigma_w & \sigma_w^2 \end{bmatrix}\right)$. The term $s_t \equiv \ln S_t$, where S_t is the surplus consumption ratio, is defined as

$$S_t \equiv \frac{C_t^a - X_t}{C_t^a}$$

where C_t^a denotes the average consumption of all individuals in the economy, X_t reflects the individual habit response to average consumption, and $\lambda(s)$ is the sensitivity function of s. The SDF is

$$\ln\left(\frac{M_{t+1}}{M_t}\right) = \ln\delta - \gamma\left(s_{t+1} - s_t + g_{t+1}\right)$$

Bekaert and Engstrom [2010a] propose a variant of the model of Campbell and Cochrane [1999] in which the dynamics of consumption growth g_t consist of two fat-tailed skewed distributions. Uncertainty in this economy is described by

$$\begin{array}{ll} g_{t+1} = \overline{g} + x_t + \sigma_{gp}\omega_{p,t+1} - \sigma_{gn}\omega_{n,t+1} & x_t = \rho_x x_{t-1} + \sigma_{xp}\omega_{p,t} + \sigma_{xn}\omega_{n,t} \\ q_{t+1} = \mu_q + \rho_q q_t + \sigma_{qp}\omega_{p,t+1} + \sigma_{qn}\omega_{n,t+1} & \omega_{p,t+1} = ge_{t+1} - p_t \text{ and } \omega_{n,t+1} = be_{t+1} - n_t \\ p_t = \overline{p} + \rho_p\left(p_{t-1} - \overline{p}\right) + \sigma_{pp}\omega_{p,t} & n_t = \overline{n} + \rho_n\left(n_{t-1} - \overline{n}\right) + \sigma_{nn}\omega_{n,t} \\ ge_{t+1} \sim Gamma(p_t, 1) & be_{t+1} \sim Gamma\left(n_t, 1\right) \end{array}$$

where p_t and n_t are the conditional means of good and bad environment shocks, denoted ge_{t+1} and be_{t+1}, respectively. The SDF is

$$\log\left(\frac{M_{t+1}}{M_t}\right) = \ln\delta + \gamma\left(q_{t+1} - q_t\right) - \gamma g_{t+1}$$

where δ is the time preference parameter, γ is the curvature parameter, and $q_t \equiv \ln\left(\frac{C_t}{C_t - X_t}\right) = \ln\left(\frac{1}{S_t}\right)$; hence, q_t represents the log of the inverse surplus consumption ratio.

The last asset pricing model class consists of the rare disaster type. In this chapter, I study two versions of the asset pricing model of the rare disaster type. One is the version of Rietz [1988], Barro [2006b], and Gabaix [2012]. I use the one proposed by Backus et al. [2011a]. The second version of the rare disaster model is that proposed by Nakamura et al. [2013]. Instead of a one-period permanent disaster, it allows for a slow recovery process from a big jump of consumption growth. Backus et al. [2011a] model the dynamics of consumption growth follows:

$$g_{t+1} = w_{t+1} + z_{t+1},\ w_{t+1} \sim i.i.d.N(0,1)$$
$$z_{t+1} |\ J_{t+1} \sim i.i.d.\ N(\theta J_{t+1}, \delta^2 J_{t+1})$$

where J_{t+1} is an i.i.d. Poisson random variable with density $\frac{\omega^j}{j!}e^{-\omega}$ for $j \in \{0, 1, 2, ...\}$ and mean ω. In the model, w_{t+1} and z_{t+1} are mutually independent over time and the SDF is

$$\ln\left(\frac{M_{t+1}}{M_t}\right) = \ln\delta - \gamma g_{t+1}$$

where δ is the time discount rate and γ is the coefficient of relative risk aversion. Eventually, the model is proposed by Nakamura et al. [2013] . The dynamics of this model are described as the following:

$$\begin{aligned} c_{t+1} &= x_{t+1} + z_{t+1} + \epsilon_{t+1} \\ \triangle x_{t+1} &= \mu_{t+1} + \eta_{t+1} + I_{t+1}\theta_{t+1} \\ z_{t+1} &= \rho z_t - I_{t+1}\theta_{t+1} + I_{t+1}\phi_{t+1} + v_{t+1} \end{aligned}$$

where c_{t+1} denotes the log consumption at time $t+1$, x_t denotes potential consumption, and z_t denotes the disaster gap, that is, the amount by which consumption differs from its potential due to current and past disasters, and ϵ_{t+1} is $i.i.d.N(0, \sigma_\epsilon^2)$. In the potential consumption and disaster gap processes, μ_{t+1} is the average growth rate of the consumption trend and η_{t+1} is an $i.i.d.N(0, \sigma_\eta^2)$ shock. The occurrence of disasters is governed by a Markov process I_{t+1} that contains two types of disaster components: a world component and an idiosyncratic component. Disasters affect consumption in two ways, as a short-run drop and as a long-run drop. The term θ_{t+1} denotes a one-off permanent shift in the level of potential consumption due to a country–specific disaster at time $t+1$ and $\theta_{t+1} \sim N(\theta, \sigma_\theta^2)$. The term ϕ_{t+1} denotes a shock that causes a temporary drop in consumption due to the disaster at time $t+1$ and $\phi_{t+1} \sim tN\left(\phi, \sigma_\phi^2, -\infty, 0\right)$. The SDF is

$$\ln\left(\frac{M_{t+1}}{M_t}\right) = \xi\ln\delta - \frac{\xi}{\psi}g_{t+1} + (\theta - 1)\, r_{a,t+1}$$

where g_{t+1} is consumption growth, $r_{a,t+1}$ is the return on the aggregate consumption claim, which is unobservable, and $\xi = \frac{1-\gamma}{1-\frac{1}{\psi}}$. This model is an extension of the long-run risk type of asset pricing model. In this case, the persistent component of consumption growth is the linear combination of potential consumption and the disaster gap, $x_{t+1} + z_{t+1}$. Intuitively, this model-implied pricing kernel has higher volatility compared with that of the long-run risk models.

2.4 Empirical Application

To lay the groundwork for the empirical examination, this section's analysis starts by documenting the properties of the linear predictive systems and asset returns. Then I compute the variance bounds on pricing kernels for each asset set across horizons based on the available return predictive systems. Finally, I elaborate on the performance of asset pricing models given my evaluation metrics.

2.4.1 Predictive Systems

To conduct my empirical analysis, I consider two horizons and two alternative sets of assets for each horizon. Specifically, I concentrate on horizons of $h = 4$ and $h = 20$ quarters and for each horizon I alternatively consider 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 and[9]

2. SET B: the returns from SET A plus the returns from growth and value portfolios formed by the book-to-market ratio and returns holding over the horizon h and Treasury bonds with a constant maturity of five years.

In particular, the market return is the (gross) return on the value- weighted portfolio of all stocks traded in the New York Stock Exchange, the American Stock Exchange, and NASDAQ. All returns are gross and deflated using the CPI.[10] 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 one to examine whether the equity premium puzzle (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 stock portfolios and government bond.[11]

Table **2.2** presents full-sample statistics for the quarterly stock and bond returns for the common sample period, 1952Q2 to 2012Q4. Over this sample period, the mean nominal return on stock market was 11.41% per annum and the mean nominal returns on growth and value portfolios were 10.93% and 15.16%, respectively. Meantime, the mean nominal return on five-year maturity government bonds was 6.28% per annum and the mean short-term interest rate was about 5.01% per annum. The standard deviation of

[9] As an alternative to the returns from rolling over the three-month Treasury- bill, I also considered the yield to maturity on a zero-coupon bond with maturity matching the horizon h. The results were basically unchanged and particularly so at the one-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 I 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, the results were unaffected and particularly so at the one-year horizon.

[10] For a detailed description of the data construction, see Appendix I.

[11] The empirical term structure literature has typically employed unsmoothed Fama–Bliss zero-coupon yields from the Center for Research in Security Prices (CRSP). When the holding period is one year, the correlation between the returns on the five-year constant-maturity Treasury bond and the Fama–Bliss five-year zero coupon is 0.989. I use constant maturity coupon bonds from the CRSP instead of the Fama–Bliss data set.

stock market nominal returns was 16.65% per annum and the standard deviations of the growth and value stock portfolios were 18.27% and 19.13%. The standard deviation of the five-year maturity bond returns was 5.76% per annum and the standard deviation of the short-term interest rate was about 1.62%.

[**Insert Table 2.2 about Here**]

There are six potential relevant returns predictors: 1) the log price–dividend ratio, pd_t (Campbell and Shiller (1988a, 1988b); 2) the stock variance of the S&P500 , $svar_t$ (Goyal and Welch [2003], Bandi and Perron [2008]); 3) the consumption–wealth ratio of Lettau and Ludvigson [2001], cay_t; 4) the short term yield on a one-month government bill, tbl_t; 5) the term spread between a five-year and a one-month government bond yield, spr_t (Campbell [1987], Fama and French [1989]); and 6) a linear combination of the forward rate factor of Cochrane and Piazzesi [2005], CP_t. [12].

In contrast to the simple random walk view , stock and bond returns do seem predictable and markedly more so the longer the return horizon. As in the simple example in Section 2, I consider the following linear predictive system:

$$R^i_{t+h} = \beta^i_{0,h} + \beta^i_{1,h} z^i_t + u^i_{t+h} \tag{2.6}$$

where $i = 1, 2, ..., N$ stands for different assets and $Z_t = \left(z^1_t, z^2_t, ..., z^N_t\right)$ denotes the vector of return predictors. Hence, one finds $2^{6\times 2} = 4096$ competitive predictive systems by using asset returns contained in SET A and $2^{6\times 5} = 1073\,741\,824$ competitive predictive systems by using asset returns contained in SET A. As mentioned above, the holding period ranges from one year to five years, that is, $h = 4$ and $h = 20$ quarters, respectively . Each return predictive system is estimated by using SUR at each horizon.

[**Insert Table 2.3 about Here**]

Table 2.3 reports the quantile values of the system-wise R^2 of predictive regressions across the horizons of both asset sets. The median system-wide R^2 ranges from 31.48% at the one-year horizon to 38.74% at the five-year horizon for SET A, while for SET B it ranges from 19.25% at the one-year horizon to 29.54% at the five-year horizon. Similarly, comparison of the quantile values of the system-wide R^2 at the five-year horizon with the associated values at the one-year horizon empirically shows that the performance of return predictive systems in predicting returns improves with the horizon. Moving on from a comparison of different holding horizons to a comparison of different asset sets, I find that

[12]Initially, there were 21 return predictors selected based on previous return predictability studies (Lettau and Ludvigson [2001], Welch and Goyal [2008a], Ludvigson and Ng [2009a]). To simplify the estimation task, I keep only six of them, based on both in-sample fit and out-of-sample forecasting performance (for more details, see Appendix I).

the median of system R^2 of the predictive regression at the 1one-year horizon for SET A is 31.48%, while the median is only 19.25% for SET B after employing the additional three more assets, two stock portfolios, and a long term government bond. Similarly, the 95% quantile value changes from 49.08% to 31.50% at the one-year horizon from 54.02% to 42.53% at the five-year horizon. These results are consistent with much of the recent empirical research on the predictability of stocks and bonds.

In conclusion, this section notes three observations. First, the return predictive model 2.6 is independent any specific asset pricing model. Second, since I am not seeking the best performance return predictive system across horizons for each asset, I consider all possible predictive systems using a given set of assets and returns predictors. I then compute the variance bounds on pricing kernels based on the results from each predictive system. Meanwhile, the empirical results reported in Table 2.3 provide some intuition about the greater influence of return predictability on longer–horizon variance bounds . Third, my conclusions may appear weakened by the results of Goyal and Welch (2003, 2008a) , where return forecasts based on dividend yields and a number of other variables do not work out of sample. However, my analysis requires only *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.

2.4.2 Variance Bounds across Horizons

It is apparent from the previous section that expected returns vary over time. To evaluate the ability of return predictability to improve the utility of the variance bounds as a diagnostic tool, in this section I compare the variance bounds with different choices of return predictors across horizons. Whereas previous studies e.g., Cochrane and Hansen [1992]) carry out the exercise on the variance bounds across horizons, my analysis extends all these results and highlights the interaction of return predictability with predictive model uncertainty.

As shown in Section 2, the predictor-based bounds are computed by using the left-hand side of Eq. (2.4). The variance of the pricing kernel is a quadratic function of its unconditional mean and its coefficients are determined by the conditional moments of asset returns based on each predictive system. For simplicity, I assume the conditional covariance matrix for returns to be constant and estimate that it has a residual in the predictive regressions.

[**Insert Figure 2.1 about Here**]

Section 3.1 theoretically shows that return predictability translates into a tight variance bound on pricing kernels. Figure 2.1 presents the empirical results for both SET

A and SET B. I consider investment horizons of one year and five years. In each panel I report three variance bounds that are generated with different return predictive systems sorted by level of system-wide R^2. The lowest parabola (with the black dash–dot line)presents the variance bound connected with the minimum system-wide R^2 among all competing return predictive systems, the most stringent parabola (red dashed line) is linked with the maximum system-wide R^2, and the blue solid lines located in the middle is associated with the median system-wide R^2. The figure shows that the predictability across horizons documented in Table 2.3 translates into a stringent lower bound on the variance of the pricing kernel. In particular, Figure 2.1, Panel A, reveals that the variance bounds become sharper once the system-wide R^2 values increase. For instance, at the one-year (five-year) horizon) for SET A, the minimum point of the frontier associated with a 51.3% system-wide R^2 is about 1.69 (1.80) times sharper than the frontier associated with the 62.1% system-wide R^2. Panel B reports the same analysis for SET B. The minimum of the frontier at the one-year (five-year) horizon linked to the 34.1% system-wide R^2 is about 1.68 (1.70) times sharper than the frontier linked to the 48.9% system-wide R^2. The difference between the bounds associated with the lower and higher system-wide R^2 values across horizons reflects the considerable predictability documented in Table 2.3.

[Insert Figure 2.4 about Here]

To display the feasible region of the pricing kernels formed with different variance bounds, for each given value of the unconditional mean v, I sort all available bounds by variance and present the minimum (dark green dash–dotted line), median (blue solid line), and maximum (red dashed line) of the variance in Figure 2.4 for both SET A and SET B across horizons. Compared with the results in Figure 2.1, the lowest boundaries are wider than the variance bound without return predictability and not all the curves are parabolas . One question may arise: whether this evidence contradicts the statement in the previous section, that is, return predictability translates into a tight variance bound. The answer is no, because the variance bounds intersect each other for different return predictive systems. For instance, in Figure 2.2, the variance bound (black solid parabola) computed by using the log price–dividend ratio to predict both stocks and bonds and the variance bound (blue solid parabola) computed by using the log price–dividend ratio and the term spread to predict bond returns and no predictor for stocks, are crossed . The main reason is that stock predictability is the main force increasing the variance bound, while bond predictability is the main force narrowing it (Favero et al. [2013]).

To conclude this section, I note that the tightness of the volatility bounds is a combination of three simultaneous forces: the predictors embedded in the conditional moments of the returns, the horizon at which these predictors becomes more relevant, and the set of assets available for investment. Specifically, first, the bound becomes tighter as the number of assets increases; second, stock return predictability increases the variance bound; and, third, bond return predictability narrows the variance bound.

2.4.3 Performance of Asset-Pricing Models

To employ my method to assess features of a pricing kernel that are necessary to price returns, I focus on three leading classes of asset pricing models (see Section 3.3): the long-run risk class, the external habit class and the rare disasters class. To simulate the series of model-implied pricing kernels I use either calibrated or estimated values for the parameters. The complete specification of the parameter values for preferences and exogenous dynamics, are reported in Table 2.7 - 2.11.

[Insert Tables 2.7 - 2.11 about here]

To compute the mean and the volatility of the SDFs of each model, I simulate 600, 000 monthly observations (50, 000 years) of the model-implied SDF for all the asset pricing models, except 50, 000 annual observations for the model of Nakamura et al. [2013]. From this long time series, I then calculate the unconditional moments of the corresponding pricing kernel.[13] The reported values in Table 2.4 are obtained by averaging the mean and the standard deviation of the model-implied SDF obtained from 10 long (50, 000–year) simulations. Within the long-run risk class, Table 2.4 reveals that, at the one- year (five- year) horizon, the standard deviations of the SDF implied by Bansal and Yaron [2004] and Bansal et al. (2012d, 2012b) are 0.4267 (1.1118), 0.8885 (4.0487), and 0.4734 (1.5274). The yearly pricing kernel implied by Bansal et al. [2012d] has a high standard deviation with a mean above one, which yields a negative risk–free rate. For the external habit class, I study the model of Campbell and Cochrane [1999] with both calibrated and estimated parameters,[14] and a variant proposed by Bekaert and Engstrom [2010a]. The calibrated version and the estimated version of Campbell and Cochrane [1999] yield similar results for both the mean and standard deviation, where the mean of the model-implied pricing kernels of Bekaert and Engstrom [2010a] is smaller. Finally, the volatility implied by Nakamura et al. [2013][15] is 1.2921 at the one-year horizon and 6.0472 at the five–year horizon, while the volatility implied by Backus et al. [2011a] is 0.8089 at one year and 1.2921 at five years. The volatility of the model-implied SDF is significantly

[13] 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].

[14] Aldrich and Gallant [2011b] estimate the habit model of Campbell and Cochrane [1999] and present for each parameter the posterior mean and posterior standard deviation. I refer to the posterior mean of each parameter as my point estimate for that parameter.

[15] Nakamura et al. [2013] estimates their model using data for 24 countries over more than 100 years. To be consistent with other studies, my paper only uses the estimations of the United States for country–specific parameters. For each parameter, the posterior mean and posterior standard deviation are presented. I refer to the posterior mean of each parameter as my point estimate for that parameter.

improved by incorporating a disaster-related persistent component into the consumption growth process. To get a feeling for the differences between all the models, I also plot the average mean and standard deviation of the model–implied SDFs in Figure 2.4. This figure shows that the asset pricing model performance depends on both the unconditional mean and the standard deviation of the model–implied SDFs . For instance, compared to the external habit model of Campbell and Cochrane [1999], it is more difficult for Bansal et al. [2012d] to reach the variance bound related with a high degrees of return predictability, even though the standard deviation of this model-implied SDF is higher.

[Insert Table 2.4 about Here]

I then compute the value of the model performance index, α, by searching for the robust variance bound on the pricing kernel for each asset pricing model and report the results in Table 2.5. At the one–year horizon for both SET A and SET B, the model performance index α of all long-run risk types of models is about one. This implies that, even without considering return predictability, the long-run risk type of asset pricing model fails to price the given sets of assets. However, the models of Bansal and Yaron [2004] and Bansal et al. [2012b] are able to satisfy 82.35% and 56.32% variance bounds among all available candidates at the five–year horizon for SET A. For the external habit class, all the asset pricing models exhibit outstanding performance at the five-year horizon for both sets of assets. These model-implied pricing kernels can be found within all available variance bounds based on the different return predictive systems. Moreover, at the one–year horizon, the values of α of Campbell and Cochrane [1999] with a calibrated parameter and an estimated parameter from SET A (SET B) are 0.4880 (0.9875) and 0.3506 (0.9692). The equity premium puzzle is more pronounced at the yearly horizon than at the longer horizon. Generally, external habit type asset pricing models are capable of being associated with a high degree of return predictability at a long horizon and at a short horizon they still perform better than the long-run risk types. Finally, within the class of rare disaster models, both the asset pricing models satisfy all the available variance bounds for SET A at each horizon, while for SET B the model of Backus et al. [2011a] is rejected by 64.79% of the variance bounds at one year and by 29% of them at five years.

[Insert Table 2.5 about Here]

The above results highlight an interesting fact: To obtain a tighter robust variance bound, that is, a lower model performance index, it is not enough to only improve the volatility of the model-implied SDF. The mean of the model-based SDFs is crucial. In particular, return predictability leads to a sharper variance bound on the pricing kernels. A tiny shift of the mean would lead to a big increase in volatility. The variance bound becomes tougher for asset pricing models to obtain . For instance, comparing the models of Backus et al. [2011a] and Bansal et al. [2012d] at a one–year horizon, even the standard deviation is lower, the model of Backus et al. [2011a] lies within all the available variance bounds for SET A, and the model of Bansal et al. [2012d] is rejected by almost all the candidate bounds.

In sum, this section shows that the robust variance bound and the model performance index are a simple but useful tool for assessing the performance of candidate asset pricing models at multiple horizons. Each selected asset pricing model parameterization approximates quite reasonably the (annual) unconditional equity premium and the real risk-free return while simultaneously closely calibrating the first two moments of consumption growth. However, my evidence reveals that the variance of the SDFs from the long-run risk and external habit models fail to meet the restrictions imposed by the variance bounds at the one- year horizon once the asset returns are predictable. Among all the asset pricing models, that of Nakamura et al. [2013] performs the best. The persistent disaster process significantly increases the volatility of model-implied SDFs without increaseing the mean.

2.5 Implications for Asset Pricing Models

The thrust of this section is to examine whether the model-implied pricing kernels are sensitive to changes in the preference or dynamics of the state variables. Such analysis can enable insights into how economic fundamentals are linked to pricing kernels and how the performance of an asset pricing model could be improved by altering return predictability.

Within each class of asset pricing models, one has estimated parameters. By using the standard deviation of each estimation, I quantify the impact of parameter uncertainty on the mean and volatility of their pricing kernels. It is important to note that all three models are estimated using a long span of the data sampled at an annual frequency to better capture the overall low–frequency variations in asset and macroeconomic data and to reduce the measurement errors that arise from seasonalities and other measurement problems.[16]

2.5.1 Long Run Risk Models

At the outset, I simulated the model-implied pricing kernel of Cecchetti et al. [1994] by using different combinations of risk aversion and the intertemporal elasticity of substitution . I report the variance of the model-implied SDF and the model performance index

[16] Aldrich and Gallant [2011b] use annual data from 1930 to 2008, Bansal et al. [2012b] data from 1930 till 2009, and Nakamura et al. [2013] data from 1890 to 2006. Thus the rare disaster model is at an advantage, since it can rely also on periods encompassing both World War I and the great influenza epidemic of 1918-1920 to make inferences on the dynamics of consumption.

α in Table 2.6 for both asset sets at each horizon. The aim of this exercise is to see how this model generates SDFs sensitive to preference parameters.

The results show that a high degree of risk aversion will significantly increase the variance of model-implied SDFs , while there is no conclusive statement of how the elasticity of substitution will affect the variance of the pricing kernels. For instance, in the comparison of columns (3) and (7) in Table 2.6 , the standard deviation increases from 0.5367 (1.6375) to 0.8885 (4.0487) at the one–year (five–year) horizon. However, the increase of the variance accompanies the increase of the mean of the SDFs . Therefore, when the risk aversion coefficient increases from 7.5 through 8.5 to 10, the model performance index varies slightly, from 0.8984 through 0.8801 to 0.7185 at the one-year horizon for SET A. As mentioned in the previous section, the mean of the model-implied SDFs is crucial, particularly when the return is predictable. To conclude, increasing the risk aversion coefficient of the long-run risk model can not produce a tighter robust variance bound for this model, that is, a higher degree of return predictability.

Going further, I follow an approach similar to that of Cecchetti et al. [1994] and the further description in the Appendix II to analyze the source of parameter uncertainty of the long-run risk model. Figure 2.5(a) displays the model–generated SDF of Bansal et al. [2012b] with a one standard deviation confidence ellipse that only takes into account preference parameter uncertainty, while in Figure 2.5(b) the ellipse represents the confidence area based on the uncertainty of the parameters describing the dynamics. The results reveal that the long-run risk model-implied pricing kernels are sensitive to both preference- and dynamics- relevant parameters . In comparison, the parameter uncertainty relative to the dynamics of the state variables has a more significant effect on the model–generated pricing kernel. This brings about the insight that, by change the dynamic process of the state variable of the long-run risk model, this model's performance can be improved to obtain a tighter robust variance bound associated with a high degree of return predictability.

[Insert Table 2.6 about here]
[Insert Figure 2.5(a), 2.5(b) about here]

2.5.2 External Habit Models

Judging by my results, the approach of Bekaert and Engstrom [2010a] does not appear to substantially improve upon the variance of the SDF in to the approach of Campbell and Cochrane [1999]. This finding is somewhat unexpected, but consistent with the empirical evidence of Figures 2.6(a) and 2.6(b), where I report the average mean and variance of the model–simulated SDFs of Campbell and Cochrane [1999] with an estimated parameter with a one standard deviation confidence ellipse. Unlike in the long-run risk model, external habit model-implied SDF is sensitive to neither changes in preference parameters nor changes in dynamics- related parameters. As Figure 6 shows, particularly at the one-year horizon, the confidence ellipses in both cases almost overlap with the triangles. Hence, changing either the preference parameters or the dynamics of the state variables is an ineffective way to improve the external habit model's performance related to return predictability .

[**Insert Figure 2.6(a), 2.6(b) about here**]

2.5.3 Rare Disaster Models

In the previous section, I show that the rare disaster model of Backus et al. [2011a] performs well at the one–year horizon, since jumps contained in the consumption growth process increase the volatility of model–generated SDFs without changing the mean. However, this type of rare disaster model still struggles at the long horizon. Moreover, among all the asset pricing models examined in this chapter, the rare disaster model of Nakamura et al. [2013] exhibits outstanding performance regarding predictability across horizons.

As for the long-run risk and external habit models, I first study the parameter uncertainty effect of the model-implied pricing kernels. Figure 2.7 presents the average mean and standard deviation of the model-implied SDFs together with a one standard deviation confidence ellipse.[17] At the one-year horizon, the mean of these model-implied SDFs is more sensitive to parameter uncertainty than the variance is . At the five–year horizon, both the mean and variance are sensitive to changes in dynamics–related parameters.

Going further, it is important to determine what the economic fundamentals are to improve the asset pricing model performance. Section 3.3 documents that the model of Nakamura et al. [2013] is an extension of the long-run risk model. The persistent component of consumption growth in this case is the linear combination of potential consumption and the disaster gap. First, the disaster process increases the volatility of short-term model-implied SDFs . Then, the persistent components add to the volatility of the long-term model-implied SDFs . Therefore, the model of Nakamura et al. [2013] is capable of being within all the available candidate bounds for the different return predictive systems .

[**Insert Figure 2.7 about here**]

In summary, for the long-run risk type of rare disaster model, Nakamura et al. [2013] shows that unexpected shocks and the persistency of state variables within asset pricing models are the most important economic fundamentals for improving the models' performance .

[17] In the model of Nakamura et al. (2013), the preference related parameters are chosen based on the previous literature. Hence, in this case, I only analyze the parameter uncertainty of model dynamics.

2.6 Conclusions

The HJ variance bound is a conventional diagnostic tool for asset pricing models. It indicates whether the pricing kernel of an asset pricing model can explain a given set of asset returns. On the other hand, the predictability of asset returns translates into a tight variance bound that is more difficult for asset pricing models to satisfy. Although financial economists have identified variables that predict stock and bond returns through time, the true predictive regression specification has remained an open issue. this chapter attempts to answer the question of how to efficiently use the information contained in a set of return predictors to discriminate between competitive asset pricing models without the knowledge of the correct return predictive system.

To address this question, I propose a robust variance bound on pricing kernels to discriminate among asset pricing models according to the property of the return predictability of the financial data. By taking into account return predictability and predictive model uncertainty, I form a set of variance bounds on pricing kernels based on different asset return predictive systems. For a given asset pricing model, the robust variance bound is the tightest among all the candidates that this model-implied pricing kernel can satisfy. The tightness of the robust variance bound reveals the degree of return predictability a given model pricing kernel is capable of explaining. Based on the results of the robust variance bounds, I form a model performance index. This index quantifies the degree of return predictability a given asset pricing model is able to obtain. It provides a new way to evaluate competing asset pricing models in terns of return predictability.

I then apply my method to a broad class of asset pricing models: (i) the long-run risk model (Bansal and Yaron [2004], Bansal et al. (2012d, 2012b)) , (ii) the external habit formation model (Campbell and Cochrane [1999], Bekaert and Engstrom [2010a]), and (iii) the rare disaster model (as parameterized Backus et al. [2011a] and Nakamura et al. [2013]). The long-run risk type of model is struggling to be associated with any degree of return predictability at each horizon. A high level of risk aversion raises the volatility of these model-implied pricing kernels, but it also increases their mean. Under dual effects, the model-implied pricing kernels are unable to explain the given set of assets once the returns are predictable. However, the long–run risk type of the rare disaster model of Nakamura et al. [2013]exhibits outstanding performance at each horizon.

This is the first study of the relation between asset pricing model performance and asset return predictability. Apart from its methodological contribution, it also brings great deal of interesting empirical evidence that points out directions for improving the performance of current asset pricing models. For instance, at a short horizon, large transitory shocks on consumption increase the ability of asset pricing models to match a high degree of return predictability; at a long horizon, the persistency of state variables improves the performance of the asset pricing models.

There are several extensions for future research. In my current exercise, SUR is the only estimation method. Given a predictive model and parameter uncertainty, a Bayesian approach would be intuitive . A Bayesian approach assigns posterior probabilities to a wide set of competing predictive systems, providing the associated distribution of the variance bound for each given value of the mean. Furthermore, a growing body of literature reveals that ambiguity in the preferences of representative agents helps to explain the preceding asset pricing puzzles . It would be interesting to investigate whether ambiguity preference would significantly improve the performance of asset pricing models in terms of consistent with return predictability. Ultimately, the empirical evidence reveals that the return predictability using US data is more significant than that using Japanese or UK data. Therefore, it is more objective to evaluate asset pricing models by applying my method with international data.

Appendix I: Data

I consider a set of quarterly equity and bond returns over the period 1952Q2 to 2012Q4. My choice of the start date is dictated by the availability of data for my 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. 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-year log stock return between dates $t+j-1$ and $t+j$ and R_{t+j} is the simple gross return.

1. Stock returns: Return data on the value-market index are obtained from the Center for Research in Security Prices (CRSP). I use the NYSE/Amex value-weighted index with dividends as our market proxy, R_{t+1}. Meanwhile, return data on growth and value portflios are obtained from the website of

   Kenneth R. Frenchformed. The portfolios are formed on Book to Market ratio. Growth portfolio is with low Book to Market ratio, while value portfolio is with high Book to Market ratio.

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. 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. Returns predictors: price-dividend ratio, see Campbell and Shiller [1988b] and Campbell and Shiller [1988a]; svar, see Goyal and Welch (2008); cay, see Lettau and Ludvigson [2001]. Dividends are 12-month moving sums of dividends paid on the index. TB1M, Treasury-bill rates are the 1-Month Treasury Bill: Secondary-Market Rate from the economic research data base at the Federal Reserve Bank at St. Louis (FRED) (see, e.g., Campbell [1987]); tms, the Term Spread 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]); CP, the single CP factor is constructed as in Cochrane and Piazzesi (2005) from contemporaneous forward rates, see Cochrane and Piazzesi (2005).

Series Number	Short Name	Descriptions
1	pd	(log) Price Divident Ratio
2	svar	Stock Variance on S&P 500
3	cay	Consumption-Wealth Ratio
4	TB1M	Short Term Yield on 1-Month Treasury Bill
5	tms	Term Spread
6	CP	Cochrane-Piazzesi Factor

4. Inflation: I use the all items seasonally unadjusted CPI from the Bureau of Labor Statistics. Quarterly inflation is the log growth rate in the CPI.

Appendix II: Confidence interval for SDFs

This appendix explains how we construct the confidence area of a model-implied SDF. These area account for the uncertainty of the estimated parameters used to compute the mean and standard deviation of the SDF. For a given consumption-based asset pricing model (LRR, external habit and rare disasters model in our paper), we obtain the model-implied SDF by simulating a long series (50,000 years) of consumption growth and of the state variables. Let us denote the model-implied mean and standard deviation of the SDF as

$$\mu_m(\phi, \theta)$$
$$\sigma_m(\phi, \theta)$$

where ϕ is denoted as the vector of parameters that characterize the preference, and θ contains all the parameters associated with the dynamics. For instance, in the LRR model, $\phi = (\delta, \gamma, \psi)$ and $\theta = (\mu, \mu_d, \phi, \varphi_d, \rho_{dc}, \rho, \varphi_e, \overline{\sigma}, \upsilon, \sigma_\omega)$, (see Table 1.9).

To construct the confidence area of (μ_m, σ_m), we calculate the standard deviation of $\mu_m\left(\phi, \widehat{\theta}\right)$, denoted as $\sigma^2_{\mu_m}$, and the standard deviation of $\sigma_m\left(\phi, \widehat{\theta}\right)$, denoted as $\sigma^2_{\sigma_m}$. These standard deviations exhibit the sensitivity of (μ_m, σ_m) with respect to estimated $\widehat{\theta}$. To obtain $\sigma^2_{\mu_m}$ and $\sigma^2_{\sigma_m}$ we follow the lead of Cecchetti, Lam and Mark (1994).

Based on equation (19) in Cecchetti, Lam and Mark (1994), the estimate for the variance of μ_m and of σ_m can be written as

$$\widehat{\sigma}^2_{\mu_m} = \left(\frac{\partial \mu_m}{\partial \theta'}\right)\bigg|_{\widehat{\theta}} \Sigma_{\widehat{\theta}} \left(\frac{\partial \mu_m}{\partial \theta}\right)\bigg|_{\widehat{\theta}}$$

$$\widehat{\sigma}^2_{\sigma_m} = \left(\frac{\partial \sigma_m}{\partial \theta'}\right)\Big|_{\widehat{\theta}} \Sigma_{\widehat{\theta}} \left(\frac{\partial \sigma_m}{\partial \theta}\right)\Big|_{\widehat{\theta}}$$

where $\Sigma_{\widehat{\theta}}$ is a diagonal variance matrix of $\widehat{\theta}$, and $\left(\frac{\partial \mu_m}{\partial \theta'}\right)\Big|_{\widehat{\theta}}$ $\left(\left(\frac{\partial \sigma_m}{\partial \theta'}\right)\Big|_{\widehat{\theta}}\right)$ is a vector of the first derivative of $\mu_m\,(\sigma_m)$ with respect to θ, and evaluated at $\theta = \widehat{\theta}$. For instance, in the LRR model,

$$\Sigma_{\widehat{\theta}} = diag\left(\sigma^2_{\widehat{\mu}}, \sigma^2_{\widehat{\mu}_d}, \sigma^2_{\widehat{\phi}}, ..., \sigma^2_{\widehat{\omega}}\right)$$

where the values of $\sigma^2_{\widehat{\mu}}, \sigma^2_{\widehat{\mu}_d}, \sigma^2_{\widehat{\phi}}, ..., \sigma^2_{\widehat{\omega}}$ are listed in Table 1.9 and we numerically approximate the derivative by

$$\left(\frac{\partial \sigma_m}{\partial \theta}\right)\Big|_{\widehat{\theta}} = \begin{pmatrix} \frac{\sigma_m\left(\phi, \widehat{\theta}_{-\widehat{\mu}}, \widehat{\mu}+\varepsilon^{\mu}\right) - \sigma_m\left(\phi, \widehat{\theta}_{-\widehat{\mu}}, \widehat{\mu}-\varepsilon^{\mu}\right)}{2\varepsilon^{\mu}} \\ \frac{\sigma_m\left(\phi, \widehat{\theta}_{-\widehat{\mu}_d}, \widehat{\mu}_d+\varepsilon^{\mu_d}\right) - \sigma_m\left(\phi, \widehat{\theta}_{-\widehat{\mu}_d}, \widehat{\mu}_d-\varepsilon^{\mu_d}\right)}{2\varepsilon^{\mu_d}} \\ \vdots \\ \frac{\sigma_m\left(\phi, \widehat{\theta}_{-\widehat{\omega}}, \widehat{\omega}+\varepsilon^{\omega}\right) - \sigma_m\left(\phi, \widehat{\theta}_{-\widehat{\omega}}, \widehat{\omega}-\varepsilon^{\omega}\right)}{2\varepsilon^{\omega}} \end{pmatrix}$$

in which $\varepsilon^{\mu} = \widehat{\mu} \times 10^{-7}$ and $\widehat{\theta}_{-\widehat{\mu}}$ is the vector of all estimated parameters excluded $\widehat{\mu}$.

With the standard deviations of μ_m and σ_m at hand, we describe the sensitivity property, and report our results graphically as the confidence ellipses in Figures 1.3 – 1.6: in particular the confidence area around (μ_m, σ_m) is covered by an ellipse centered $\left(\mu_m\left(\phi, \widehat{\theta}\right), \sigma_m\left(\phi, \widehat{\theta}\right)\right)$ with radius $\left(\frac{\widehat{\sigma}_{\mu_m}}{2}, \frac{\widehat{\sigma}_{\sigma_m}}{2}\right)$.

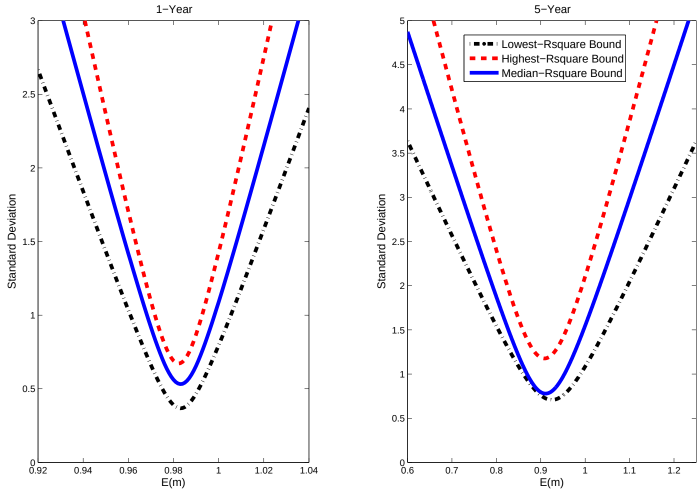

(a) Set A

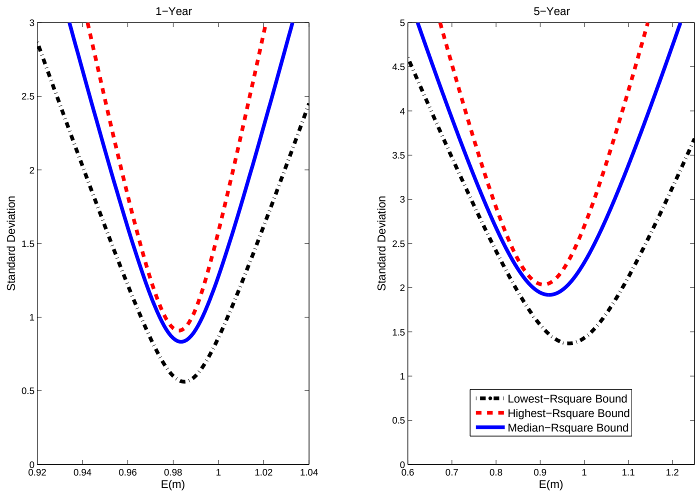

(b) Set B

Figure 2.1: Variance bounds on pricing kernel ranked with system-wide R-square for different investment horizons.
I compute the variance bound on pricing kernels wiht all availabe return predictive system for both asset sets across horizons. Red dashed line presents the vairance bound, which is computed based on the return predictive system with the highest system-wide R-square, blue solid line is variance bound with the median system-wide R-square, and the black dashed dotted line indicates the bound with the lowest system-wide R-square. The bounds are generated using SET A (see Panel A) and SET B (see Panel B). Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2-2012Q4

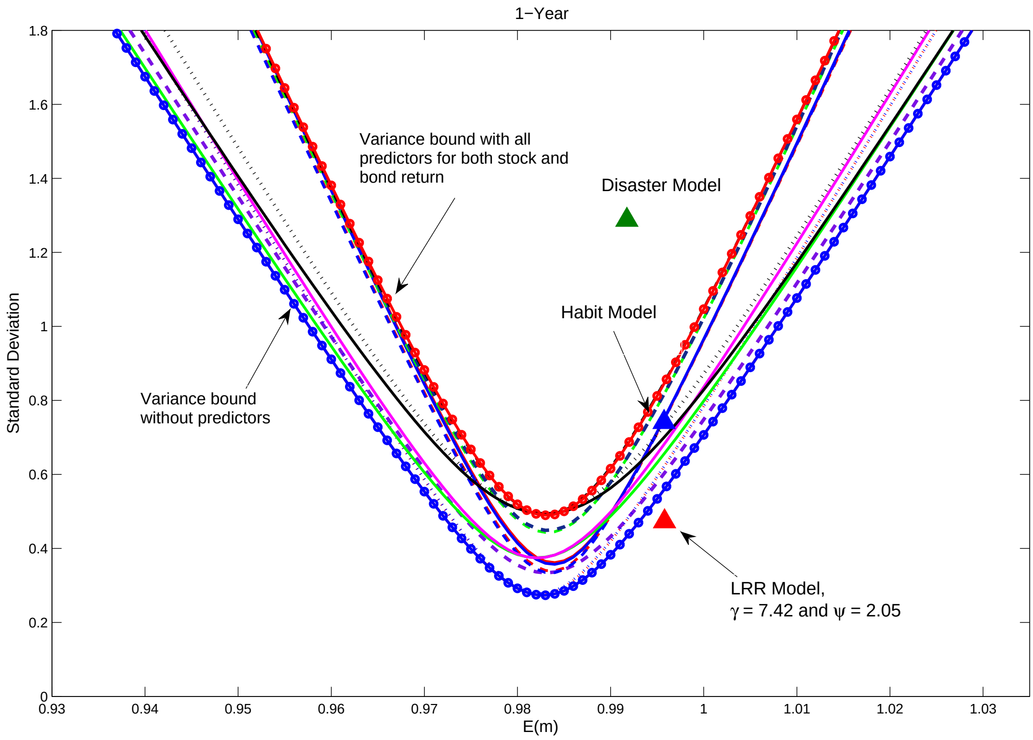

Figure 2.2: Variance Bounds and Model-Implied SDFs.

I report all the variance bounds on pricing kernels which are computed based on different predictive systems. Triangles presents the average mean and standard deviation from 10 single simulations run of 600,000 month, for the models incorporating long-run risk, external habit persistence, and rare disasters.Sample: 1952Q2-2012Q4.

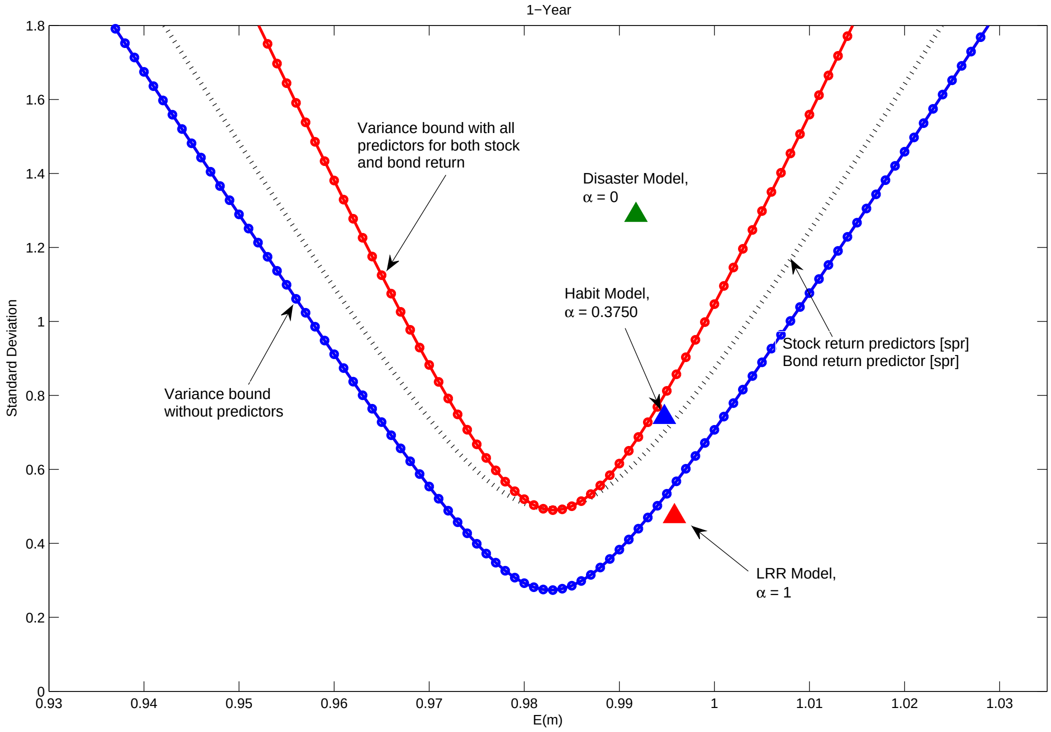

Figure 2.3: Robust Variance Bounds and Model-Implied SDFs.
This figure presents the robust variance bounds on pricing kernels for each asset pricing models. Triangles presents the average mean and standard deviation from 10 single simulations run of 600,000 month, for the models incorporating long-run risk, external habit persistence, and rare disasters. In this figure, I also point out the model performance index value of each asset pricing models. Sample: 1952Q2-2012Q4.

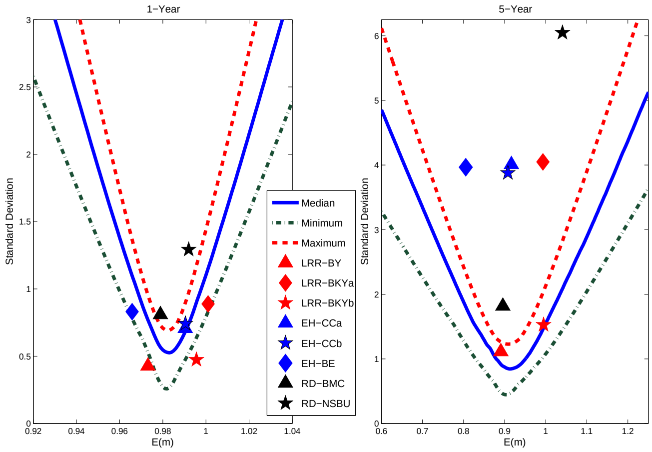

(a) Set A

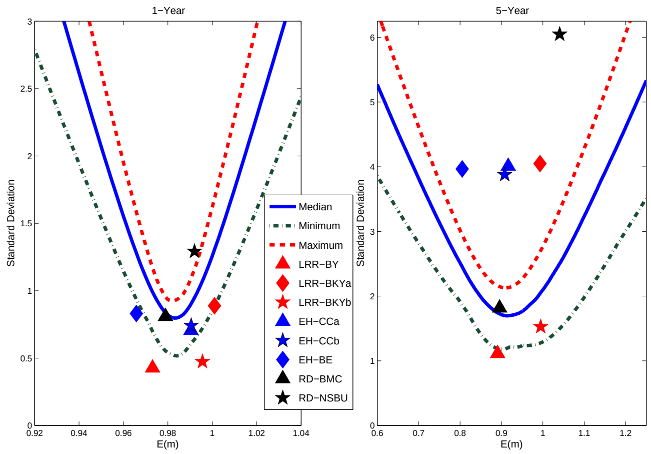

(b) Set B

Figure 2.4: Feasible region of pricing kernels.
I compute all the variance bounds on pricing kernels based on different return predictive system for each set across horizons. Then to form the feasible region of pricing kernels, I plot the minimum, median and maximum variance of SDF for each given mean of it in this figure. I also report average mean and standard deviation values from 10 simulations run of 600,000 months with calibrated (triangles and diamonds) and estimated parameters (pentagrams), reapectively. The long-run risk models are presented by red symbols, the external habit models are presented by blue symbols, and the rare disasters models are presented by black symbols. Sample: 1952Q2-2012Q4.

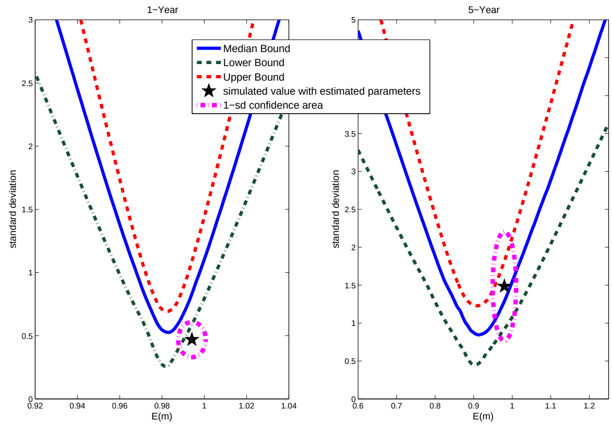

(a) Preference

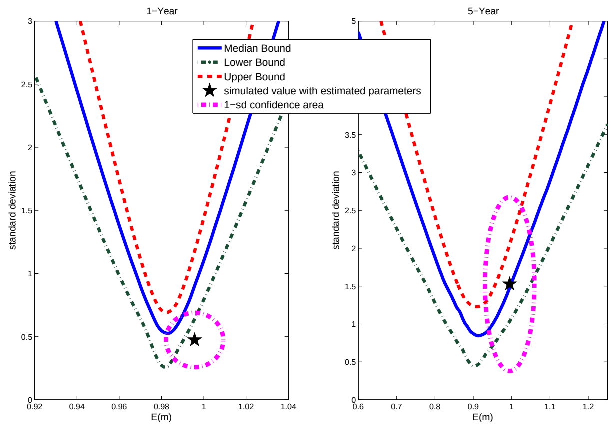

(b) Dynamics

Figure 2.5: Parameter sensitivity of the long-run risk model.
The black pentagrams report average mean and standard deviation values from 10 simulations run of 600,000 months with estimated parameters of Bansal, Kiku and Yaron(2012a). The ellipse (dashed dotted area) shows the uncertainty (one standard deviation) in the calculation of the mean and standard deviation of the model implied SDFs. Panel A take account uncertainty in the parameter values for preferences (i.e. we take as given parameters that characterize the state dynamics). Panel B take account uncertainty in the parameter values for the state dynamics (i.e. we take as given parameters that characterize the preferences). I also report the feasible region provided by all candidate variance bound in this figure. Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q4.

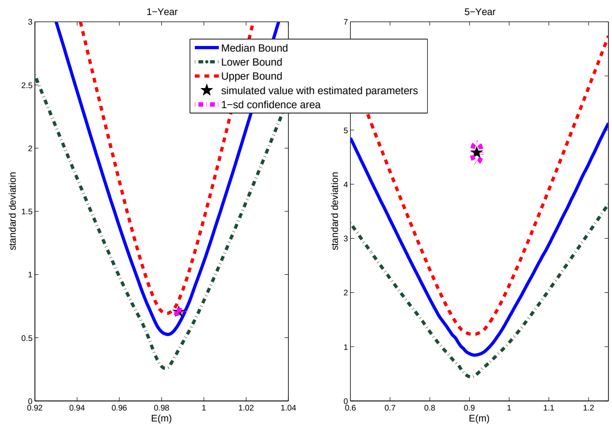

(a) Preference

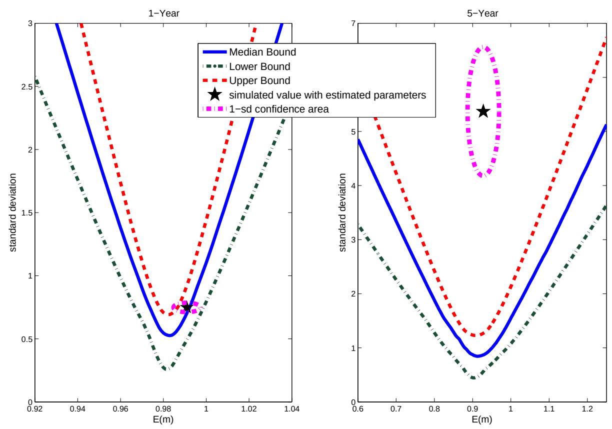

(b) Dynamics

Figure 2.6: Parameter sensitivity of the external habit model.
The black pentagrams report average mean and standard deviation values from 10 simulations run of 600,000 months with estimated parameters of Campbell and Cochrane(1999). The ellipse (dashed dotted area) shows the uncertainty (one standard deviation) in the calculation of the mean and standard deviation of the model implied SDFs. Panel A take account uncertainty in the parameter values for preferences (i.e. we take as given parameters that characterize the state dynamics). Panel B take account uncertainty in the parameter values for the state dynamics (i.e. we take as given parameters that characterize the preferences). I also report the feasible region provided by all candidate variance bound in this figure. Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q4.

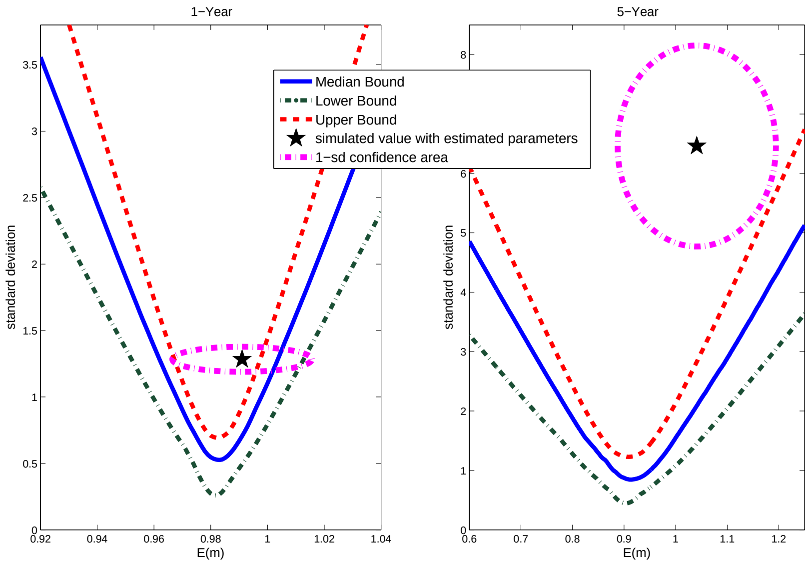

(a) Dynamics

Figure 2.7: Parameter sensitivity of the rare disasters model.
The black pentagrams report average mean and standard deviation values from 10 simulations run of 600,000 months with estimated parameters of Bansal, Kiku and Yaron(2012a). The ellipse (dashed dotted area) shows the uncertainty (one standard deviation) in the calculation of the mean and standard deviation of the model implied SDFs. For this model I only take account uncertainty in the parameter values for the state dynamics (i.e. we take as given parameters that characterize the preferences). I also report the feasible region provided by all candidate variance bound in this figure. Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q4.

	Stock Return Predictors	Bond Return Predictors	System wide R^2
1	$\varnothing$	$\varnothing$	0.000
2	*spr*	$\varnothing$	0.065
3	*pd*	$\varnothing$	0.067
4	$\varnothing$	*pd*	0.107
5	$\left[\ pd\quad spr\ \right]$	$\varnothing$	0.109
6	*spr*	*pd*	0.110
7	*pd*	*pd*	0.212
8	$\left[\ pd\quad spr\ \right]$	*pd*	0.254
9	$\varnothing$	*spr*	0.295
10	*spr*	*spr*	0.297
11	$\left[\ pd\quad spr\ \right]$	*spr*	0.309
12	*pd*	*spr*	0.310
13	$\varnothing$	$\left[\ pd\quad spr\ \right]$	0.322
14	*spr*	$\left[\ pd\quad spr\ \right]$	0.324
15	*pd*	$\left[\ pd\quad spr\ \right]$	0.385
16	$\left[\ pd\quad spr\ \right]$	$\left[\ pd\quad spr\ \right]$	0.389

Table 2.1: Predictive System Estimation - A Simple Example.
This table report the system-wide R^2 of sixteen predictive system regressions in the simple example. In this example, there are two asset returns, real gross stock returns and real gross roll-over three month Treasury Bill returns. Two well-known return predictors are considered, log price-dividend ratio (*pd*) and term spread between long term and short term government bond (*spr*). Column (1) and (2) list the associated return predictors of each return predictive system. All contents are sorted by the system-wide R^2. Quarterly inflation is the log growth rate in the CPI. Sample: 1952Q2: 2012Q4.

	Stock			Bond	
	Market Ind	Value Ptf.	Growth Ptf.	3-month	5-year
Mean return (% p.a.)	11.41%	15.16%	10.93%	5.01%	6.28%
Standard deviation (% p.a.)	16.65%	19.31%	18.27%	1.62%	5.76%

Table 2.2: Statistics of Data.
This table reports sample statistics of quarterly nominal stock and bond total returns. Market Index are returns on the stock total returns on the value weighted portfolio of all stocks traded in the NYSE, the AMEX, and NASDAQ from CRSP. Value and Groth Portfolio are returns on the stock portfolio formed with Book-to-Market ratio. 3-month and 5-year bond returns are real returns on the 3-month and 5-year constant maturity bond from the CRSP Fixed Term Indices File. Real returns are computed by deflating the nominal returns by the seasonally unadjusted Consumer Price Index inflaion from the Bureau of Labor Statistics. Quarterly inflation is the log growth rate in the CPI. Sample: 1952Q2: 2012Q4.

Quantiles		5%	10%	25%	50%	75%	90%	95%
SET A	1-year	0.0672	0.1123	0.2009	0.3148	0.4158	0.4820	0.4908
	5-year	0.0773	0.1248	0.2860	0.3874	0.4623	0.5337	0.5402
SET B	1-year	0.0683	0.0925	0.1392	0.1925	0.2512	0.3051	0.3150
	5-year	0.1011	0.1245	0.2109	0.2954	0.3572	0.4092	0.4253

Table 2.3: System-wide R^2 of Predictative Regressions.
This table reports the 5%, 10%, 25%, 50%, 75%, 90% and 95% quantile values of the system-wide R^2 of predictive regressions across horizon for both SET A and SET B. I estimated $2^{N \times Q}$ predictive models by using seemingly unrelated regression (SUR) at each horizon of each asset set. Then based on the system-wide R^2, I search for the associated values with 5%, 25%, 50%, 75% and 90% quantile. Long horizon returns are computed by compounding quarterly returns. Sample: 1952Q2 - 2012Q4.

		Long-Run Risk			External Habit			Rare Disasters	
		Bansal-Yaron(2004)	Bansal, Kiku-Yaron(2012a)	Bansal, Kiku-Yaron(2012b)	Campbell-Cochrane calibration	Campbell-Cochrane estimation	Bekaert-Engstrom	Backus et al. (2011)	Nakamura et al. (2013)
1-year	mean	0.9731	1.0010	0.9956	0.9904	0.9905	0.9658	0.9789	0.9920
	std. dev.	0.4267	0.8885	0.4734	0.7058	0.7420	0.8309	0.8089	1.2921
5-year	mean	0.8908	0.9932	0.9949	0.9164	0.9080	0.8053	0.8957	1.0408
	std. dev.	1.1118	4.0487	1.5274	4.0077	3.8780	3.9655	1.8167	6.0472

Table 2.4: Model-Implied Pricing Kernels.

I compute the mean and standard deviation of the model implied pricing kernels across horizons, via simulations, respectively, for the models that incorporate long-run risk, external habit persistence, and rare disasters. All calculations are based on model parameters tabulated in Tables Appendix-I through Appendix-V, and reported values are the averages from 10 single simulations run of 600,000 month.

		Long-Run Risk			External Habit			Rare Disasters	
		Bansal-Yaron(2004)	Bansal, Kiku-Yaron(2012a)	Bansal, Kiku-Yaron(2012b)	Campbell-Cochrane calibration	Campbell-Cochrane estimation	Bekaert-Engstrom	Backus et al. (2011)	Nakamura et al. (2013)
SET A	1-year	1.0000	0.9919	1.0000	0.4880	0.3506	0.9565	0	0
	5-year	0.1865	0	0.4368	0	0	0	0	0
SET B	1-year	1.0000	1.0000	1.0000	0.9875	0.9692	1.0000	0.6479	0
	5-year	1.0000	0	0.9585	0	0	0	0.2900	0

Table 2.5: Model Performance Index α.

I compute the mean and standard deviation of the model implied SDFs across horizons, via simulations, respectively, for the models that incorporate long-run risk, external habit persistence, and rare disasters. All calculations are based on model parameters tabulated in Tables Appendix-I through Appendix-VI. I construct variance bounds of SDFs at different horizon, using Bekaert and Liu (2004) approach, based on each return predictive system of each asset set. I then compare a given model based SDFs with all $2^{N\times Q}$ variance bounds of SDFs at different horizon and count the number, K, of variance bounds, which this model based SDFs is able to exceed. The reported value is the model performance measure α, which represent the proportion of $K/2^{N\times Q}$. $N = 2$ and $Q = 6$ in the case of SET A, while in SET B $N = 5$ and $Q = 6$. Sample: 1952Q2 - 2012Q3.

	RA = 7		RA = 7.5		RA = 8.5		RA = 10		RA = 10	
	IES = 1.5		IES = 1.5		IES = 1.5		IES = 1.5		IES = 2.5	
	SET A	SET B	SET A	SET B	SET A	SET B	SET A	SET B	SET A	SET B
1-year	0.4798	0.4798	0.5367	0.5367	0.6625	0.6625	0.8885	0.8885	0.8977	0.8977
	{0.9705}	{1.0000}	{0.8984}	{1.0000}	{0.8801}	{0.9995}	{0.7185}	{0.9819}	{1.0000}	{1.0000}
5-year	1.3950	1.3950	1.6375	1.6375	2.2299	2.2299	4.0487	4.0487	3.5489	3.5489
	{0}	{0.9219}	{0}	{0.6345}	{0}	{0.0073}	{0}	{0}	{0}	{0}

Table 2.6: Parameter Sensitivity Analysis.

I study the sensitivity of parameters, mainly about risk aversion and intertemporal elasticity of substitution coefficients, for Bansal, Kiku and Yaron's (2012a) model. I compute the mean and standard deviation of the model implied SDFs across horizons, via simulations, for the long-run risk model with different combination of risk aversion and intertemporal elasticity substitution coefficients. All calculations are based on model parameters tabulated in Tables Appendix-I, and reported values are the averages from single simulations run of 600,000 month. The value in the brackets are the model performance measure α, which represent the proportion of $K/2^{N\times Q}$. $N=2$ and $Q=6$ in the case of SET A, while in SET B $N=5$ and $Q=6$. Sample: 1952Q2 -2012Q3.

	Parameter	Bansal-Yaron (2004)	Bansal-Kiku-Yaron(2012a) *calibration*	Bansal-Kiku-Yaron(2012b) *estimation*
Preferences				
Time preference	δ	0.998	0.9989	0.9989 (0.0010)
Risk aversion	γ	7.5	10	7.42 (1.55)
EIS	ψ	1.5	1.5	2.05 (0.84)
Consumption growth dynamics,	g_t			
Mean	μ	0.0015	0.0015	0.0012 (0.0007)
Dividends growth dynamics,	$g_{d,t}$			
Mean	μ_d	0.0015	0.0015	0.0020 (0.0017)
Persistence	ϕ	3	2.5	4.45 (1.63)
Volatility parameter	φ_d	4.5	5.96	5.00 (1.39)
Comsumption exposure	π	0	2.6	0
Correlation between innovations	ρ_{dc}			0.49 (0.33)
Long-run risk, x_t				
Persistence	ρ	0.979	0.975	0.9812 (0.0086)
Volatility parameter	φ_e	0.044	0.038	0.0306 (0.0160)
Consumption growth volatility, σ_t				
Mean	$\overline{\sigma}$	0.0078	0.0072	0.0073 (0.0015)
Persistence	v	0.987	0.999	0.9983 (0.0021)
Volatility parameter	σ_w	0.23×10^{-5}	0.28×10^{-5}	2.62×10^{-5} (3.10×10^{-6})

Table 2.7: Appendix-I: Parametrization of long-run risk asset pricing model. Our parameterization of calibrated values is taken from Bansal and Yaron [2004] and [Bansal et al., 2012d]. The value of the parameter estimates and corresponding standard errors (in parentheses) are taken from Bansal et al. [2012b]. The model is simulated at the monthly frequency.

	Parameter	Campbell-Cochrane	
Preferences		*calibration*	*estimation*
Time preference	δ	0.9909	0.9903 (0.0004)
Risk aversion	γ	2	0.9756 (0.0772)
Consumption growth dynamics, g_t			
Mean	$\overline{g}$	0.0016	0.0017 (0.00007)
Volatility parameter	σ	0.0043	0.0050 (0.00019)
Divdend growth dynamics, $\triangle d_t$			
Volatility parameter	σ_w	0.0323	0.0319 (0.0014)
Corr between innovitions	ρ_{dc}	0.2	0.1945 (0.0093)
Steady state surplus consumption ratio	$\overline{S}$	0.0570	0.0637
Persistence in consumption surplus ratio	ϕ	0.9884	0.9877 (0.0003)
Log of risk-free rate	$r^f \times 10^2$	0.0783	0.0854 (0.0365)

Table 2.8: Appendix-II: Parametrization of external habit asset pricing model. Our parameterization of calibrated values is taken from Campbell and Cochrane [1999]. The value of the parameter estimates and corresponding standard errors (in parentheses) are taken from Aldrich and Gallant [2011b]. The model is simulated at the monthly frequency.

	Parameter	Bekaert-Engstrom
Preferences		
Time preference	δ	0.9970
Risk aversion	γ	3.0690
Consumption growth dynamics, g_t		
Mean	$\overline{g}$	0.0031
Long-run risk, x_t		
Persistence	ρ_x	0.9467
Volatility parameter	$\sigma_{xp} \times 10^3$	-0.0016
Volatility parameter	$\sigma_{xn} \times 10^3$	-0.0140
Volatility parameter	σ_{gp}	0.0012
Volatility parameter	σ_{gn}	0.0052
p_t *dynamics*		
Mean	$\overline{p}$	12.1244
Persistence	ρ_p	0.9728
Volatility parameter	σ_{pp}	0.4832
n_t *dynamics*		
Mean	$\overline{n}$	0.6342
Persistence	ρ_n	0.9810
Volatility parameter	σ_{nn}	0.1818
Log of inverse consumption surplus, q_t		
Mean	$\overline{q}$	1.0000
Persistence	ρ_q	0.9841
Volatility parameter	σ_{qp}	0.0005
Volatility parameter	σ_{qn}	0.0479

Table 2.9: Appendix-III: Parametrization of external habit asset pricing model. Our parameterization of calibrated values is taken from Bekaert and Engstrom [2010a]. The model is simulated at the monthly frequency.

	Parameter	Annual	Monthly
Preferences			
Time preference	δ	1.0448	$1.0448^{\frac{1}{12}}$
Risk aversion	γ	5.190	5.190
Gaussian component of consumption growth, w_t			
Mean	μ	0.023	0.023/12
Volatility parameter	σ	0.01	$0.01/\sqrt{12}$
Non-gaussian component of consumption growth, z_t			
Mean of Poisson density	ω	0.010	0.010/12
Mean of z_{t+1} conditional on J_{t+1}	θ	-0.30	-0.30
Variance of z_{t+1} conditional on J_{t+1}	δ^2	0.15^2	0.15^2
Mapping between dividend and consumption, $D_t = C_t^{\phi}$	ϕ	2.60	2.60

Table 2.10: Appendix-IV: Parametrization of asset pricing models incorporating rare disasters.
Our parameterizations are based on Backus et al. [2011a], Table II, annualized parameters. Then the monthly time preference parameter is the annual time preference parameter raised to the power 1/12, see Campbell and Cochrane [1999]. The leverage parameter λ which follows Wachter [2013b], Table 1.

	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_ϕ	0.121 (0.015)
a permanent shift in consumption caused by shock, θ_t		
Mean	θ	−0.025 (0.007)
Volatility parameter	σ_θ	0.083 (0.006)

Table 2.11: Appendix-V: 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.

Chapter 3

Demographics and The Behavior of Interest Rates

3.1 Introduction

Recent evidence shows that the behavior of interest rates is consistent with the decomposition of spot rates in the sum of two processes, (i) a very persistent long term expected value and (ii) a mean-reverting component (Fama [2006]; Cieslak and Povala [2013]). Traditionally, models of the term structure concentrate mainly on the mean-reverting component as only stationary variables are used to determine yields. Partial adjustments to equilibrium yields are then used to rationalize the persistence in observed data (see Figure 1). This chapter offers a novel interpretation for the persistent long-term component of interest rates by relating it to the age structure of the US population.

Modelling the persistent component of interest rates has important consequences for forecasting. Consider Affine Term Structure Models (ATSM). In this framework, given the dynamics of the short term rate, a stationary VAR representation for the factors is used to project the entire term structure. The risk premia are identified by posing a linear (affine) relation between the price of risk and the factors. In this case the no-arbitrage assumption allows to pin down the dynamics of the entire term structure by imposing a cross-equation restrictions structure between the coefficients of the state model (the VAR for the factors) and the measurement equations that maps the factors in the yields at different maturities (Ang et al. [2007]; Dewachter and Lyrio [2006]). The potential problem with this general structure is that while yields contain a persistent component, the state evolves as a stationary VAR which is designed to model a mean-reverting process and cannot capture the time series behavior of persistent variables. This discrepancy might therefore explain the, somewhat disappointingly, mixed results from the forecasting performance of affine term structure models (Duffee [2002]; Favero et al. [2012]; Sarno et al. [2014]). Enlarging the information set by explicitly considering a large number of macroeconomic variables as factors (Mönch [2008]) has generated some clear improvement without addressing the discrepancy between the stationarity of the factors and the high persistence of interests rates. In fact, forecasting interest rates in the presence of a highly persistent component in rates requires the existence of a factor capable of modeling the persistence.

To explicitly address this problem we argue that when the monetary policy authorities set the policy rate, they do not react only to cyclical swings reflected in the transitory (expected) variations of output from its potential level and of (expected) inflation from its target, but also consider the slowly evolving changes in the economy, i.e., trends, which take place at lower frequency (see, for example, Bernanke [2006]).[1] In particular, the target for the policy rate is set by implicitly taking into account the life-cycle savings behavior of the population to determine the equilibrium policy rate. Linking the target policy rates to demographics makes Taylor-type rule of monetary policy capable of generating observed persistence in interest rates (Diebold and Li [2006]; Diebold and Rudebusch [2013]).[2]

Yields at different maturities depend on the sum of short rate expectations and the risk premium. While it is less plausible to consider the risk premium as a non-mean reverting component (e.g., Dai and Singleton [2002]), the presence of a persistent component related to demographics can be rationalized in terms of smooth adjustments in short-rate expectations that take decades to unfold. In particular, we consider a demographic variable MY, a proxy for the age structure of the US population originally proposed by Geanakoplos et al. [2004](GMQ from now onwards) and defined as the ratio of middle-aged (40-49) to young (20-29) population in the US as the relevant demographic variable to determine the persistent component of interest rates.[3]

First, we illustrate our permanent-transitory decomposition using demographic information. Then we propose an affine term structure model (ATSM) which parsimoniously incorporates demographic channel in one-period yield via the central bank reaction function and all yields at longer maturities as the sum of future expected policy rates and the term premium. The advantage of an ATSM is that the term premia are explicitly modeled using both observable and unobservable factors.[4] This framework provides a the natural

[1]"... *adequate preparation for the coming demographic transition may well involve significant adjustments in our patterns of consumption, work effort, and saving ...*" Chairman Ben S. Bernanke, Before The Washington Economic Club, Washington, D.C.,October 4, 2006.

[2]When young adults, who are net borrowers, and the retired, who are dissavers, dominate the economy, savings decline and interest rates rise. The idea, is certainly not a new one as it can be traced in the work of Wicksell [1936], JOHN MAYNARD [1936], Modigliani and Brumberg [1954], but it has received relatively little attention in the recent literature.

[3]In principle there are many alternative choices for the demographic variable, MY. However, using a proxy derived from a model is consistent with economic theory (Altavilla et al. [2014]) and reduces the risk of a choice driven by data-mining. Importantly, MY is meant to capture the relative weights of active savers investing in financial markets.

[4]The literature is vast, few related examples are Ang and Piazzesi [2003], Diebold et al. [2006], **?**,

complement to the Taylor rule. In this specification, given the dynamics of the short term rate, a stationary VAR representation for the factors is used to project the entire term structure. We show that the demographic ATSM not only provides improved yield forecasts with respect to traditional benchmarks considering statistical accuracy (Carriero and Giacomini [2011]), but it also provides economic gains for long term investors in the context of portfolio allocation (Sarno et al. [2014]; Gargano et al. [2014]).

To our knowledge, the potential relation between demographics and the target policy rate in a reaction function has never been explored in the literature. This analysis is relevant for two reasons. First, the persistence of policy rates cannot be modeled by the mainstream approach to central bank reaction functions that relate monetary policy exclusively to cyclical variables. Second, putting term structure model at work to relate the policy rate to all other yields requires very long term projections for policy rates. For example, in a monthly model, 120 step ahead predictions of the one-month rate are needed to generate the ten-year yield. However, long-term projections are feasible in a specification where the persistent component of the policy rates is modelled via demographics while macroeconomic factors capture the cyclical fluctuations. For instance, a standard VAR could be used to project the stationary component, while the permanent component is projected by exploiting the exogeneity of the demographic variable and its high predictability even for a very long-horizon.[5]

The trend-cycle decomposition of interest rates has been also recently investigated by Fama [2006] and Cieslak and Povala [2013], who argue that the predictive power of the forward rates for yields at different maturities could be related to the capability of appropriate transformations of the forward rates to capture deviations of yields from their permanent component. These authors propose time-series based on backward looking empirical measures of the persistent component; in particular Fama [2006] considers a five-year backward looking moving average of past interest rates and Cieslak and Povala [2013] consider a ten-year discounted backward-looking moving average of annual core CPI inflation. We propose instead a forward looking measure for which reliable forecasts are available for all the relevant horizons. Figure 2 illustrates the existence of a persistent component in interest rates by relating it to different measures of slowly evolving trends. The Figure reports the yield to maturity of one-Year US Treasury bond, along with the persistent components as identified by Fama [2006] and Cieslak and Povala [2013], and the demographic variable, MY.

The Figure shows that MY not only strongly co-moves with the alternative estimates of the persistent component, but it is also capable of matching exactly the observed peak in yields at the beginning of the eighties. The very persistent component of yields is

Hördahl et al. [2006], Rudebusch and Wu [2008], Bekaert et al. [2010].

[5]The Bureau of Census currently publishes on its website projections for the age structure of the population with a forecasting horizon up to fifty years ahead.

common to the entire term structure of interest rates: Figure 1 illustrates this point by reporting the US nominal interest rates at different maturities. The visual evidence reported in Figures 1-2 motivates the formal investigation of the relative properties of the different observable counterparts for the unobservable persistent component of the term structure.

Our framework brings together four different strands of the literature: i) the one analyzing the implications of a persistent component on spot rates predictability, ii) the one linking demographic fluctuations with asset prices, iii) the empirical literature modeling central bank reaction functions using the rule originally proposed by Taylor [1993] and iv) the term structure models with observable macro factors and latent variables.

The literature on spot rates predictability has emerged from a view in which forecastability is determined by the slowly mean-reverting nature of the relevant process. Recently, it moved to a consensus that modeling a persistent component is a necessary requirement for a good predictive performance (Bali et al. [2009]; Duffee [2012]). Early literature attributes this predictability to the mean reversion of the spot rate toward a constant expected value. This view has been recently challenged; the predictability of the spot rate captured by forward rates is either attributed to a slowly moving, yet still stationary, mean (Balduzzi et al. [1998]) or to the reversion of spot rates towards a time-varying very persistent long-term expected value (Fama [2006]; Cieslak and Povala [2013]).[6]

Our choice of the variable determining the persistent component in short term rate is funded in the literature linking demographic fluctuations with asset prices and in the empirical approach to central bank reaction functions based on Taylor's rule. Taylor rule models policy rates as depending on a long term equilibrium rate and cyclical fluctuations in (expected) output and inflation. The long term equilibrium rate is the sum of two components: the equilibrium real rate and equilibrium inflation, which is the (implicit) inflation target of the central bank. Evans [2003] shows that over longer horizons, expectation of the nominal and real yields rather than the inflation expectations dominate in the term structure. The long-term equilibrium is traditionally modeled as a constant. However, Woodford [2001] highlights the importance of a time-varying constant in the feedback rule to avoid excess interest rate volatility while stabilizing inflation and output gap. This chapter allows for a time-varying target for the equilibrium policy rate by relating it to the age structure of population. The use of a demographic variable allows us to explicitly model the change of regime in the spot rate proving a natural alternative to regime-switching specifications (for example, Gray [1996]; Ang and Bekaert [2002]).

The idea of using demographics to determine the persistent component of the whole

[6]There are other alternative views in the literature which argue for a unit root in the spot rates (Dewachter and Lyrio [2006]; Christensen et al. [2011] or suggest a near unit root process to model the persistent component (Cochrane and Piazzesi [2005]; Jardet et al. [2013]; Osterrieder and Schotman [2012]

term structure complements the existing literature that uses demography as an important variable to determine the long-run behavior of financial markets (Abel [2001]). While the literature agrees on the life-cycle hypothesis[7] as a valid starting point, there is disagreement on the correct empirical specification and thus the magnitude of demographic effects (Poterba [2001]; Goyal [2004]). Substantial evidence is available on the impact of the demographic structure of the population on long-run stock-market returns (Ang and Maddaloni [2003]; Bakshi and Chen [1994]; Goyal [2004]; DellaVigna and Pollet [2007]). However, the study of the empirical relation between demographics and the bond market is much more limited, despite the strong interest for comovements between the stock and the bond markets (Lander et al. [1997]; Vuolteenaho and Campbell [2004]; Bekaert and Engstrom [2010b]).

Geanakoplos et al. [2004] consider an overlapping generation model in which the demographic structure mimics the pattern of live births in the U.S., that have featured alternating twenty-year periods of boom and busts. They conjecture that the life-cycle portfolio behavior (Bakshi and Chen [1994]) plays an important role in determining equilibrium asset prices. Consumption smoothing by the agents, given the assumed demographic structure requires that when the MY ratio is small (large), there will be excess demand for consumption (saving) by a large cohort of retirees (middle-aged) and for the market to clear, equilibrium prices of financial assets should adjust, i.e., decrease (increase), so that saving (consumption) is encouraged for the middle-aged (young). The model predicts that the price of all financial assets should be positively related to MY and it therefore also predicts the negative correlation between yields and MY. Note that we use the results of the GMQ model to rationalize the target for policy rate at generational frequency, in this framework there is no particular reason why the ratio of middle-aged to young population should be directly linked to aggregate risk aversion.[8] Following this intuition, we take a different approach from the available literature that studies the relationship between real bond prices and demographics through the impact on time-varying risk (Brooks [1998]; Bergantino [1998]; Davis and Li [2003]).

We concentrate on the relation between equilibrium real interest rate and the demographic structure of population as we consider the target inflation rate as set by an independent central bank who is not influenced by the preferences of the population. However, a possible relation between the preference of the population and inflation has been investigated in other studies (Lindh and Malmberg [2000], Gozluklu [2014]) which show evidence on the existence of an age pattern of inflation effects. Our approach is consistent with McMillan and Baesel [1988] who analyze the forecasting ability of a slightly

[7]Life cycle investment hypothesis suggests that agents should borrow when young, invest for retirement when middle-aged, and live off their investment once they are retired.

[8]Recent literature also shows that consumption smoothing across time rather than the risk management across states is the primary concern of the households (Rampini and Viswanathan [2014]).

different demographic variable, prime savers over the rest of the population. Our work is also related to Malmendier and Nagel [2013], who show that an aggregate measure that summarizes the average life-time inflation experiences of individuals at a given point in time is useful in predicting excess returns on long-term bonds.

Our approach to monetary policy rule has an important difference from the one adopted in the monetary policy literature. In this literature monetary policy has been described by empirical rules in which the policy rate fluctuates around a constant long-run equilibrium rate as the central bank reacts to deviations of inflation from a target and to a measure of economic activity usually represented by the output gap. The informational and operational lags that affect monetary policy (Svensson [1997]) and the objective of relying upon a robust mechanism to achieve macroeconomic stability (Evans and Honkapohja [2003]), justify a reaction of current monetary policy to future expected values of macroeconomic targets. As the output-gap and the inflation-gap are stationary variables, this framework per se is not capable of accommodating the presence of the persistent component in policy rates. One outstanding empirical feature of estimated policy rules is the high degree of monetary policy gradualism, as measured by the persistence of policy rates and their slow adjustment to the equilibrium values determined by the monetary policy targets (Clarida et al. [2000]; Woodford [2001]). Rudebusch [2002]Rudebusch [2002] and Söderlind et al. [2005] have argued that the degree of policy inertia delivered by the estimation of Taylor-type rules is heavily upward biased. In fact, the estimated degree of persistence would imply a large amount of forecastable variation in monetary policy rates at horizons of more than a quarter, a prediction that is clearly contradicted by the empirical evidence from the term structure of interest rates.[9] Rudebusch [2002] relates the "illusion" of monetary policy inertia to the possibility that estimated policy rules reflect some persistent shocks that central banks face. The introduction of demographics allows to model this persistent component of the policy rate as the time-varying equilibrium interest rate is determined by the age-structure of the population.

We shall implement the formal investigation in four stages. First, we illustrate the potential of the temporary-permanent decomposition to explain fluctuations of the term structure using the demographic information. Second, we introduce a formal representation of our simple framework, by estimating a full affine term structure model with time varying risk premium. Third, we run a horse-race analysis between a random walk benchmark, standard Macro ATSM and proposed ATSM with demographic information. We consider several measures of statistical accuracy and economic value for different investment horizon. Fourth, we investigate the relative performance of MY and other backward

[9]In a nutshell, high policy inertia should determine high predictability of the short-term interest rates, even after controlling for macroeconomic uncertainty related to the determinants of the central bank reaction function. This is not in line with the empirical evidence based on forward rates, future rates (in particular federal funds futures) and VAR models.

looking measures proposed in the literature to model the persistent component of interest rates. Finally, after assessing the robustness of our empirical findings, the last section concludes.

3.2 Demographics and the Structure of Yield Curve

We motivate our analysis with a simple framework, in which the yield to maturity of the 1-period bond, $y_t^{(1)}$, is determined by the action of the monetary policy maker and all the other yields on n-period (zero-coupon) bonds can be expressed as the sum of average expected future short rates and the term premium, $rpy_t^{(n)}$:

$$\begin{aligned} y_t^{(n)} &= \frac{1}{n}\sum_{i=0}^{n-1} E_t[y_{t+i}^{(1)} \mid I_t] + rpy_t^{(n)} \\ y_t^{(1)} &= y_t^* + \beta(E_t\pi_{t,k} - \pi^*) + \gamma E_t x_{t,q} + u_{1,t+1} \end{aligned} \tag{3.1}$$

In setting the policy rates, the Fed reacts to variables at different frequencies. At the high frequency the policy maker reacts to cyclical swings reflected in the output gap, $x_{t,q}$, i.e., transitory discrepancies of output from its potential level, and in deviation of inflation, $\pi_{t,k}$, from the implicit target of the monetary authority. Monetary policy shocks, $u_{1,t+1}$,also happen. As monetary policy impacts on macroeconomic variable with lags, the relevant variables to determine the current policy rate are k-period ahead expected inflation and q-period ahead expected output gap. However, cyclical swings are not all that matter to set policy rates. We posit that the monetary policy maker determines the equilibrium level of interest rates y_t^* (which is determined by the sum of a time varying real interest rate target and the inflation target π^*) accordingly to the slowly evolving changes in the economy that take place at a generational frequency, i.e., those spanning several decades. We relate this to the age structure of population, MY_t as it determines savings behavior of middle-aged and young population.

The relation between the age structure of population and the equilibrium real interest rate is derived by GMQ in a three-period overlapping generation model in which the demographic structure mimics the pattern of live births in the US. Live births in the US have featured alternating twenty-year periods of boom and busts. Let q_o (q_e) be the bond price and $\{\mathrm{c}_y^o,\mathrm{c}_m^o,\mathrm{c}_r^o\}$ ($\{\mathrm{c}_y^e,\mathrm{c}_m^e,\mathrm{c}_r^e\}$) the consumption stream (young, middle, old) in two consecutive periods, namely odd and even. In the simplest deterministic setup, following the utility function over consumption

$$\begin{aligned} U(c) &= E(u(c^y) + \delta u(c^m) + \delta^2 u(c^r)) \\ u(x) &= \frac{x^{1-\alpha}}{1-\alpha} \qquad \alpha > 0 \end{aligned}$$

The agent born in an odd period then faces the following budget constraint

$$c_y^o + q_e c_m^o + q_o q_e c_r^o = w^y + q_e w^m \tag{3.2}$$

and in an even period

$$c_y^e + q_o c_m^e + q_o q_e c_r^e = w^y + q_o w^m \tag{3.3}$$

Moreover, in equilibrium the following resource constraint must be satisfied

$$Nc_y^o + nc_m^o + Nc_r^o = Nw^y + nw^m + D \tag{3.4}$$

$$nc_y^e + Nc_m^e + nc_r^e = nw^y + Nw^m + D \tag{3.5}$$

where D is the aggregate dividend for the investment in financial markets.

In this economy an equilibrium with constant real rates is not feasible as it would lead to excess demand either for consumption and saving. When the MY ratio is small (large), i.e., an odd (even) period, there will be excess demand for consumption (saving) by a large cohort of retirees (middle-aged) and for the market to clear, equilibrium prices of financial assets should adjust, i.e., decrease (increase), so that saving (consumption) is encouraged for the middle-aged. Thus, letting q_t^b be the price of the bond at time t, in a stationary equilibrium, the following holds

$$\begin{aligned} q_t^b &= q_o \text{ when period odd} \\ q_t^b &= q_e \text{ when period even} \end{aligned}$$

together with the condition $q_o < q_e$. In the absence of risk, the substitutability of bond and equity together with the no-arbitrage condition implies that

$$\frac{1}{q_t^b} = 1 + y_t = \frac{D + q_{t+1}^e}{q_t^e}$$

where q_t^e is the real price of equity and $t \in \{odd, even\}$.

So, since the bond prices alternate between q_o^b and q_e^b, then the price of equity must also alternate between q_t^e and q_t^e. Hence the model predicts a positive correlation between real asset prices and MY, and a negative correlation between MY and (expected) bond yields; in other words the model implies that a bond issued in odd (even) period and maturing in even (odd) period offers a high (low) yield, since the demographic structure is characterized by a small (large) cohort of middle-aged individuals, hence low MY ratio in odd (even) periods.

Therefore, the main prediction of the model is that real interest rates and the dividend price ratio should fluctuate with the age structure of population. Unfortunately real interest rates are not observable for most of our sample. Inflation-indexed bonds (TIPS, the Treasury Income Protected Securities) have traded only since 1997 and the market of these instruments faced considerable liquidity problem in its early days. Ang et al. [2008] have solved the identification problem of estimating two unobservables, real rates and inflation risk premia, from only nominal yields by using a no-arbitrage term structure model that imposes restrictions on the nominal yields. These pricing restrictions identify the dynamics of real rates (and the inflation risk premia). We report in first panel of Figure 3 the time series behavior of the 5-year real rate identified by Ang et al. [2008] together with MY. In the second panel we consider instead MY and the dividend price ratio which is the readily observable stock market variable predicted to comove with demographics by the GMQ model. Both panels in Figure 3 illustrate that the co-movement between the (log of) dividend price ratio and 5-year real rates with MY cannot falsify the predictions of the GMQ model.[10]

Consistently with the GMQ model we consider the following permanent-transitory decomposition for the 1-period policy rates:

$$\begin{aligned}
&y_t^{(1)} = P_t^{(1)} + C_t^{(1)} = \rho_0 + \rho_1 MY_t + \rho_2 X_t \\
&P_t^{(1)} \equiv \rho_0 + \rho_1 MY_t = y_t^* \\
&C_t^{(1)} \equiv \beta(E_t \pi_{t,k} - \pi^*) + \gamma E_t x_{t,q} + u_{1,t+1} = \rho_2 X_t
\end{aligned}$$

and, assuming that the inflation gap and the output gap can be represented as a stationary VAR process, yields at longer maturity can be written as follows

$$\begin{aligned}
y_t^{(n)} &= \rho_0 + \frac{1}{n}\sum_{i=0}^{n-1} \rho_1 MY_{t+i} + b^{(n)} X_t + rpy_t^{(n)} \qquad (3.6) \\
y_t^{(n)} &= P_t^{(n)} + C_t^{(n)} \\
P_t^{(n)} &= \rho_0 + \frac{1}{n}\sum_{i=0}^{n-1} \rho_1 MY_{t+i} \\
C_t^{(n)} &= b^{(n)} X_t + rpy_t^{(n)}
\end{aligned}$$

The decomposition of yields to maturity in a persistent component, reflecting demographics, and a cyclical components reflecting macroeconomic fluctuations and the risk premia, is consistent with the all the stylized facts reported so far documenting the presence of a slow moving component common to the entire term structure. Moreover, the relation between the permanent component and the demographic variable is especially appealing for forecasting purposes as the demographic variable is exogenous and highly predictable even for very long-horizons. No additional statistical model for MY_{t+i} is needed to make the simple model operational for forecasting, as the bureau of Census projections can be readily used for this variable, as it can be safely considered strongly

[10] The implications of this evidence for stock market predictability are further investigated in Favero et al. [2011].

exogenous for the estimation and the simulation of the model to our interest.

3.3 An ATSM with Demographics

We now propose an ATSM which parsimoniously incorporates demographic channel in one-period yield via the central bank reaction function and models all yields at longer maturities as the sum of future expected policy rates and the term premium. Hence we consider the role of demographics within a more structured specification that explicitly incorporates term premia. In particular, we estimate the following Demographic ATSM:

$$\begin{aligned} y_t^{(n)} &= -\frac{1}{n}\left(A_n + B_n' X_t + \Gamma_n \mathrm{MY}_t^n\right) + \varepsilon_{t,t+1} \qquad \varepsilon_{t,t+n} \sim N(0, \sigma_n^2) \qquad (3.7) \\ y_t^{(1)} &= \delta_0 + \delta_1' X_t + \delta_2 \mathrm{MY}_t \\ X_t &= \mu + \Phi X_{t-1} + \nu_t \quad \nu_t \sim i.i.d.N(0, \Omega) \end{aligned}$$

where $\Gamma_n = \left[\gamma_0^n, \gamma_1^n \cdots, \gamma_{n-1}^n\right]$, and $\mathrm{MY}_t^n = \left[\mathrm{MY}_t, \mathrm{MY}_{t+1} \cdots, \mathrm{MY}_{t+n-1}\right]'$, $y_t^{(n)}$ denotes the yield at time t of a zero-coupon government bond maturing at time $t+n$, the vector of the states $X_t = \left[f_t^o, \ f_t^u\right]$, where $f_t^o = \left[f_t^\pi, f_t^x\right]$ are two observable factors extracted from large-data sets to project the inflation and output gap using all relevant output and inflation information which the Fed uses to set the monetary policy rate in a data-rich environment (Bernanke and Boivin (2003), Ang et al. [2007], while $f_t^u = \left[f_t^{u,1}, f_t^{u,2}, f_t^{u,3}\right]$ contain unobservable factor(s) capturing fluctuations in the unobservable interest rate target of the Fed orthogonal to the demographics fluctuations, or interest rate-smoothing in the monetary policy maker behavior. Consistently with the previous section and recent literature (e.g., Ang and Piazzesi [2003]; Huang and Shi [2011]; Barillas [2011]), we extract the two observable stationary factors from a large macroeconomic dataset following Ludvigson and Ng [2009b] to capture output and inflation information (see Appendix B).

Our specification for the one period-yield is a generalized Taylor rule in which the long-term equilibrium rate is related to the demographic structure of the population, while the cyclical fluctuations are mainly driven by the output gap and fluctuations of inflation around the implicit central bank target. Note that in our specification the permanent component of the 1-period rate is modelled via the demographic variable and the vector of the states X_t is used to capture only cyclical fluctuations in interest rates. Hence, it is very natural to use a stationary VAR representation for the states that allows to generate long-term forecasts for the factors and to map them into yields forecasts. MY_t is not included in the VAR as reliable forecasts for this exogenous variable up to very long-horizon are promptly available from the Bureau of Census. The model is completed by assuming a linear (affine) relation between the price of risk, Λ_t,and the states X_t by specifying the pricing kernel, m_{t+1}, consistently and by imposing no-arbitrage restrictions (see, for example, Duffie and Kan [1996], Ang and Piazzesi [2003]). We solve the coefficients A_{n+1}, B'_{n+1} and Γ_{n+1} recursively (see Appendix A). We study the modified affine term structure model in assuming the more general case of time varying risk premium, i.e. the market prices of risk are affine in five state variables $\lambda_0 = \begin{bmatrix} \lambda_0^\pi & \lambda_0^x & \lambda_0^{u,1} & \lambda_0^{u,2} & \lambda_0^{u,3} \end{bmatrix}$ where λ_0 is a non-zero vector and λ_1 is a diagonal matrix;

$$\Lambda_t = \lambda_0 + \lambda_1 X_t$$

$$m_{t+1} = \exp(-y_{t,t+1} - \frac{1}{2}\Lambda_t' \Omega \Lambda_t - \Lambda_t \varepsilon_{t+1})$$

$$A_{n+1} = A_n + B_n' (\mu - \Omega\lambda_0) + \tfrac{1}{2} B_n' \Omega B_n + A_1$$
$$B_{n+1}' = B_n' (\Phi - \Omega\lambda_1) + B_1'$$
$$\Gamma_{n+1} = [-\delta_2, \Gamma_n]$$

Note that the imposition of no-arbitrage restrictions allows to model the impact of current and future demographic variables on the term structure in a very parsimonious way, as all the effects on the term structure of demographics depend exclusively on one parameter: δ_2. Our structure encompasses traditional ATSM with macroeconomic factors, and no demographic variable, labelled as Macro ATSM, as this specification is obtained by setting $\delta_2 = 0$. In other words, the traditional Macro ATSM, which omits the demographic variable, is a restricted version of the more general Demographic ATSM. The no arbitrage restrictions guarantees that when $\delta_2 = 0$ also $\Gamma_n = 0$: as demographics enter the specification of yields at longer maturities only via the expected one-period yield, the dynamics of yields at all maturities become independent from demographics if MY_t does not affect the one-period policy rate. However, when the restriction $\delta_2 = 0$ is imposed, the structure faces the problem highlighted in the previous section of having no structural framework for capturing the persistence in policy rates. In fact, to match persistence in the policy rates, some of the unobservable factors must be persistent as the observable factors are, by construction, stationary. Then, the VAR for the state will include a persistent component which will make the long-term forecasts of policy rates, necessary to model the long-end of the yield curve, highly uncertain and unreliable. In the limit case of a non-stationary VAR, long-term forecast become useless as the model is non-mean reverting and the asymptotic variance diverges to infinite.

3.3.1 Model Specification and Estimation

We estimate the model on quarterly data by considering the 3-month rate as the policy rate. The properties of the data are summarized in Table 1. The descriptive statistics reported in Table 1 highlights the persistence of all yields which is not matched by the persistence of the macroeconomic factors extracted from the large data-set and it is instead matched by the persistence of the demographic variable MY.

We evaluate the performance of our specification with MY against that of a benchmark discrete-time ATSM obtained by imposing the restriction $\delta_2 = 0$ on our specification. Following the specification analysis of Pericoli and Taboga [2008], we focus on a parsimonious model including three latent factors and only contemporaneous values of the macro variables. We use the Chen and Scott [1993] methodology; given the set of parameters and observed yields latent variables are extracted by assuming that number of bonds which are priced exactly is equal to the number of unobserved variables. Hence we assume that 3-month, 2-year and 5-year bond prices are measured without error and estimate the model with maximum likelihood. We assume the state dynamics to follow a VAR(1). We impose the following restrictions on our estimation (Favero et al. [2012]):

i) the covariance matrix Ω is block diagonal with the block corresponding to the unobservable yield factor being identity, and the block corresponding to the observable factors being unrestricted, i.e.

$$\Omega = \begin{bmatrix} \Omega^o & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$$

ii) the loadings on the factors in the short rate equation are positive, $0 \leq$ -A_1

iii) $f_0^u = 0$.

We first estimate the model for the full sample 1964Q1-2013Q4, the estimated results are reported in Table 2. The results show significant evidence of demographics in the reaction function. The additional parameter δ_2 in the Demographic ATSM is highly significant with the expected negative sign. Moreover, we notice that while the unobservable level factor picks up the persistence in the Macro ATSM specification, the demographic variable dominates the level factor which becomes negligible in the Demographic ATSM. This observation is especially relevant in the context of out-of-sample forecasting. The omission of the demographic variable results in overfitting of the restricted model. Such a restricted model may be useful in explaining the in-sample patterns of the data, but does not reflect the true data generating process of bond yields (Duffee [2011]). We also notice that the estimated dynamics of the unobservable factors, especially the level factor, is very different when the benchmark model is augmented with MY. In fact, in the Macro ATSM model the third factors is very persistent and the matrix $(\Phi - I)$ describing the long-run properties of the system is very close to be singular, while this near singularity disappears when the persistent component of yields at all maturities is captured by the appropriate sum of current and future age structure of the population. In this case the VAR model for the states becomes clearly stationary and long-term predictions are more precise and reliable.

3.3.2 Out-of-Sample Forecasts

We complement the results of full sample estimation by analyzing the properties of out-of-sample forecasts of our model at different horizons. The key challenge facing ATSM models is that they are good at describing the in-sample yield data and explain bond excess returns, but often fail to beat even the simplest random walk benchmark, especially in long horizon forecasts (Duffee [2002]; Guidolin and Thornton [2010]; Sarno et al. [2014]). In our multi-period ahead forecast, we choose iterated forecast procedure, where multiple step ahead forecasts are obtained by iterating the one-step model forward

$$\widehat{y}^{(n)}_{t+h|t} = \widehat{a}_n + \widehat{b}_n \hat{X}_{t+h|t} + \widehat{\Gamma}_n \mathrm{MY}^n_{t+h}$$
$$\hat{X}_{t+h|t} = \sum_{i=0}^{h} \hat{\Phi}^i \widehat{\mu} + \hat{\Phi}^h \hat{X}_t$$

where $\widehat{a}_n = -\frac{1}{n}\widehat{A}_n, \widehat{b}_n = -\frac{1}{n}\widehat{B}_n$ are obtained by no-arbitrage restrictions. Forecasts are produced on the basis of rolling estimation with a rolling window of eighty observations. The first sample used for estimation is 1961Q3-1981Q2. We consider 5 forecasting horizons (denoted by h): one to five years. For example, for the one year forecasting horizon, we provide a total of 126 forecasts for the period 1982Q2 - 2013Q4, while the number of forecasts reduces to 111 for 5-year ahead forecasts.

Forecasting performance is measured by the ratio of the root mean squared forecast error (RMSFE) of the Demographic ATSM to the RMSFE of a random walk forecast and to the RMSFE of the benchmark Macro ATSM without the demographic variable. In parentheses, we report the p-values of the forecasting test due to Giacomini and White [2006] which is a two-sided test of the equal predictive ability of two competing forecasts. In addition, we compute the Clark and West [2006, 2007] test statistics and associated p-values testing the forecast accuracy of nested models. The additional Clark and West statistics are useful in evaluation the forecasting performance, because it corrects for finite sample bias in RMSFE comparison between nested models. Without the correction, the more parsimonious model might erroneously seem to be a better forecasting model if we only consider the ratio of RMSFE. Forecasting results from different models are reported in Table 3. Panel A compares the forecasts of Demographic ATSM against the random walk benchmark, while Panel B uses the restricted Macro ATSM ($\delta_2 = 0$) as the benchmark.

The evidence on statistical accuracy using different tests shows that the forecasting performance of the Demographic ATSM dominates the traditional Macro ATSM, especially in longer horizon starting from 2 years. Including demographic information in term structure models seems decisive to generate a better forecasting performance. By using an affine structure to model time-varying risk one can impose more structure on the yield dynamics and still improve on the forecasting performance of a simpler model once demographics is incorporated into the model to project future bond yields. The finding is striking in light of earlier evidence from the above cited literature which highlights the difficulty of forecasting future yields using ATSM specification.

In order to demonstrate the importance of a common demographics related component to explain the common persistent component in the term structure, we conduct the following dynamic simulation exercise: using the full-sample estimation results, both Macro ATSM and Demographic ATSM are simulated dynamically from the first observation onward to generate yields at all maturities. The simulated time series in Figure 4 show that, while the model without demographics converges to the sample mean, the model with demographics feature projections that have fluctuations consistent with those of the observed yields, except the recent period of quantitative easing whose start is indicated by the vertical line in 2008Q3. These simulations confirm that fitting a persistent level factor does not necessarily result in accurate out-of-sample forecasts.

3.3.3 Forecast Usefulness and Economic Value

Out-of-sample forecasting results reported in Table 5 suggest that the random walk model which does not impose any structure on yield dynamics and risk premium is still a valid benchmark, especially for short horizon forecasts up to one year. So, in the context of out-of-sample forecasting, the question is whether to choose a completely parsimonious model with no economic structure or a full fledged ATSM specification which models risk dynamics while capturing the persistence in interest rates via common demographic component. In this section, we follow the framework proposed by Carriero and Giacomini [2011] which is flexible enough to allow for forecast combination and assess the usefulness of two competing models, by both using a statistical and an economic measure of forecast accuracy. In particular, in the former case given a particular type of loss function, e.g., quadratic loss, the forecaster finds the optimal weight λ^* which minimizes the expected out-of-sample loss of the following combined forecast

$$y^{(n),*}_{t+h|t} = \widehat{y}^{(n),RW}_{t+h|t} + (1-\lambda)(\widehat{y}^{(n),DATSM}_{t+h|t} - \widehat{y}^{(n),RW}_{t+h|t})$$

where$\widehat{y}^{(n),RW}_{t+h|t}$ ($\widehat{y}^{(n),DATSM}_{t+h|t}$) is the h-period ahead yield forecast at time t of the random walk model (Demographic ATSM) of a bond maturing in n periods.

If estimated λ^* is close to one, then it suggests that only the random walk models is useful in forecasting bond yields. If on the other hand estimated λ^* is close to zero, than Demographic ATSM model dominates the random walk benchmark in out-of-sample forecasting. Estimated λ^* close to 0.5 implies that both models are equally useful in forecasting. In Table 4 Panel A, we provide estimated λ^*, and t-statistics $t^{\lambda=0}$ and $t^{\lambda=1}$ to test the null hypotheses $\lambda = 0$ and $\lambda = 1$, respectively. Results are broadly in line with the evidence reported in Table 5; while the parsimonious random walk model is useful for 1-year ahead forecasts, more structured Demographic ATSM provides more useful long horizon yield forecasts.

So far the evidence is limited to statistical forecast accuracy, but recent literature finds that statistical accuracy in forecasting does not necessarily imply economic value in portfolio choice, especially for bond excess returns (Thornton and Valente [2012]; Sarno et al. [2014]; Gargano et al. [2014]). Carriero and Giacomini [2011] framework can be extended to find the optimal portfolio weight w^* as a function λ^* by minimizing the utility loss of an investor with quadratic utility who has to choose among m risky bonds. We implement this test for 1-year and 2-year holding periods. In the first case, m=4, namely the investor chooses among 2-year to 5-year bonds. In the second case, the investment opportunity set consists of 3 bonds given the data we use in our forecasting exercise. Let the bond excess returns (net of 3-month spot rate) be a 4x1 vector, $\mathrm{rx}_{t+1} = [rx^{(2)}, rx^{(3)}, rx^{(4)}, rx^{(5)}]$ in case of 1-year holding period and a 3x1 vector $\mathrm{rx}_{t+2} = [rx^{(3)}, rx^{(4)}, rx^{(5)}]$ for 2-year holding period. Given our yield forecasts we can compute the bond excess returns

$$
\begin{aligned}
rx_{t+1} &= -n\ y_{t+1}^{(n)} + (n+1)y_t^{(n+1)} - y_t^{(n/4)} \\
rx_{t+2} &= -n\ y_{t+2}^{(n)} + (n+2)y_t^{(n+2)} - y_t^{(n/4)}
\end{aligned}
$$

and using our forecasting models we obtain excess return forecasts

$$
\begin{aligned}
\widehat{rx}_{t+1} &= -n\ \widehat{y}_{t+1|t}^{(n)} + (n+1)y_t^{(n+1)} - y_t^{(n/4)} \\
\widehat{rx}_{t+2} &= -n\ \widehat{y}_{t+2|t}^{(n)} + (n+2)y_t^{(n+2)} - y_t^{(n/4)}
\end{aligned}
$$

Panel B in Table 4 reports the estimated forecast combination weight λ^*, and associated t-statistics $t^{\lambda=0}$ and $t^{\lambda=1}$ to test the null hypotheses $\lambda = 0$ and $\lambda = 1$, respectively. As before, we consider the random walk specification as the benchmark model and compare the forecast combination weight λ^* of either the Demographic ATSM or Macro ATSM models against the random walk benchmark. For 1-year holding period, the random walk model clearly dominates Macro ATSM model in line with earlier evidence. However, the optimal weight is not statistically different from 0.5 if we combine the random walk model with Demographic ATSM, suggesting that both models are equally relevant for an investor with 1-year horizon. On the other hand, for long term investors it is evident that the Demographic ATSM is the only model that is useful for forecasting bond excess returns.

3.3.4 Long Term Projections

One of the appealing features of the demographic ATSM specification is that the availability of long-term projections for the age-structure of the population which can be exploited to produce long-term projections for the yield curve. In our specification, yields at time $t+j$ with maturities $t+j+n$ are functions of all realization of MY between $t+j$ and $t+j+n$. The exogeneity of the demographic variable and the availability of long term projections is combined in the affine model with a parsimonious parameterization generated by the no-arbitrage restrictions that allow to weight properly all future values of MY with the estimation of few coefficients. As a result future paths up to 2045 can be generated for the entire term structure, given the availability of demographic projections up to 2050.[11] In Figure 5, we compare the in-sample estimation and out-of-sample forecasts for both the 3-month spot rate and 5-year bond yield. While the in-sample estimation results are very similar, the long term projections reveal that the Macro ATSM is not able to capture the persistence in true data generating process. In particular, spot rate forecasts of the Macro ATSM model immediately converge to the unconditional mean, while it takes approx. 15 years (around 2030) for the Demographic ATSM forecasts to reach the unconditional mean.

3.4 Alternative Specifications of Permanent Component

The existence of a permanent component in spot rates has been identified in the empirical literature by showing that predictors for return based on forward rates capture the risk premium and the business cycle variations in short rate expectations. Fama [2006] explains the evidence that forward rates forecast future spot rates in terms of a mean reversion of spot rates towards a non-stationary long-term mean, measured by a backward moving average of spot rates. Cieslak and Povala [2013] explain the standard return predictor based on the tent-shape function of forward rates proposed by Cochrane and Piazzesi [2005] as a special case of a forecasting factor constructed from the deviation of yields from their persistent component. The latter is measured by a discounted moving-average of past realized core inflation.

In this section we use the standard framework to assess the capability of MY to capture the permanent component of spot rates against that of the different proxies proposed by Fama [2006] and Cieslak and Povala [2013]. This framework is designed to compare the forecasting ability of the spot rates deviations from their long term expected value and forward spot spreads. We implement it by taking three different measures of the

[11] The Bureau of Census websites provides projections for demographics variable up to 2050 and the current 5-year yield depends on the values of MY over the next five years.

permanent component: our proposed measure based on the age composition of population, the measure adopted by Fama [2006] based on a moving average of spot rates, and the measure proposed by Cieslak-Povala based on a discounted moving average of past realized core inflation.

Given the decomposition of the spot interest rates, $y_t^{(n)}$ [12] in two processes: a long term expected value $P_t^(n)$, that is subject to permanent shocks, and a mean reverting component $C_t^(n)$:

$$y_t^{(n)} = C_t^{(n)} + P_t^{(n)}$$

The following models are estimated

$$y_{t+4x}^{(1)} \text{-} y_t^{(1)} = a^x + b^x D_t + c^x [f_{t,t+4x}^{(1)} \text{-} y_t^{(1)}] + d^x [y_t^{(1)} \text{-} P_t^{(1),i}] + \varepsilon_{t+4x}$$

$$P_t^{(1),1} = \frac{1}{20}\sum_{i=1}^{20} y_{t-i-1}^{(1)} \tag{3.8}$$

$$P_t^{(1),2} = \frac{\sum_{i=1}^{40} \upsilon^{i-1}\pi_{t-i-1}}{\sum_{i=1}^{40} \upsilon^{i-1}} \tag{3.9}$$

$$P_t^{(1),3} = e^x \frac{1}{4}\sum_{i=1}^{4} MY_{t+i-1} \tag{3.10}$$

where $f_{t,t+4x}^{(1)}$ is the one-year forward rate observed at time t of an investment with settlement after $3x$ years and maturity in $4x$ years, $y_t^{(1)}$is the one-year spot interest rate, π_t is annual core CPI inflation from time $t-4$ to time t, υ is a gain parameter calibrated at 0.96 as in Cieslak and Povala [2013], and MY_t is the ratio of middle-aged (40-49) to young (20-29) population in the US, D_t is a step dummy, introduced by Fama [2006] in his original study, taking a value of one for the first part of the sample up to August 1981 and zero otherwise. This variable captures the turning point in the behavior of interest rates from a positive upward trend to a negative upward trend occurred in mid-1981.

The specification is constructed to evaluate the predictor based on the cyclical component of rates against the forward spot spread. In his original study, Fama found that, conditional on the inclusion of the dummy in the specification, this was indeed the case. This evidence is consistent with the fact the dominant feature in the spot rates of an upward movement from the fifties to mid-1981 and a downward movement from 1981 onwards is not matched by any similar movement in the forward-spot spread which looks like a mean reverting process over the sample 1952-2004. We extend the original results

[12] We adopt Cochrane and Piazzesi [2005] notation for log bond prices: $p_t^{(n)}$ = log price of n-year discount bond at time t. The continuously compounded spot rate is then $y_t^{(n)} \equiv -\frac{1}{n}p_t^{(n)}$

by considering alternative measures of the permanent component over a sample up to the end of 2013[13]. The results from estimation on quarterly data are reported in Table 5.

We consider forecasts at the 2, 3, 4 and 5-year horizon. For each horizon we estimate first a model with no cyclical component of interest rates but only the forward spot spread, then we include the three different proxies for the cyclical components of interest rates. The estimation of the model with the restriction $d^x = 0$ delivers a positive and significant estimate of c^x with a significance increasing with the horizon x. However, when the restriction $d^x = 0$ is relaxed, then the statistical evidence on the significance of c^x becomes much weaker. In fact, this coefficient is much less significant when the cycle is specified using the demographic variable to measure the permanent component and when any measure of the cycle in interest rates is introduced in the specification. The inclusion of the dummy is necessary only in the case of the Fama-cycle, while in the cases of the inflation based cycle and the demographic cycle the inclusion of the dummy variable is not necessary anymore to capture the turning points in the underlying trend. This confirms the capability of demographics and smoothed inflation of capturing the change in the underlying trend affecting spot rates. The performance the demographic cycle, however, dominates the inflation cycle at each horizon. The estimated coefficient on the demographic variable is very stable at all horizons, while the one on the discounted moving average of past inflation is more volatile.

3.5 Robustness

This section examines the robustness of our results along three dimensions. First, we extend our results to international data, since all the empirical results reported are based on US data. Second, in all forward projections we have implemented so far we have treated MY_{t+i} at all relevant future horizons as a known variable. Predicting MY requires projecting population in the age brackets 20-29, and 40-49. Although these are certainly not the age ranges of population more difficult to predict[14] the question on the uncertainty surrounding projections for MY is certainly legitimate. Therefore, we consider projections under different fertility rates and consider foreign holdings of US debt securities. Third, one might object that our statistical evidence on MY_t and the permanent component of interest rates is generated by the observation of a couple of similar paths of nonstationary random variables. Although the spurious regression problem is typical in static regression and all the evidence reported so far is based on estimation of dynamic time-series model,

[13]1-year Treasury bond yields are taken from Gurkaynak et al. dataset. Middle-young ratio data is available at annual frequencies from Bureau of Census (BoC) and it has been interpolated to obtain quarterly series.

[14]Improvement in mortality rates that have generated over the last forty years difference between actual population and projected population are mostly concentrated in older ages, after 65.

some simulation based evidence might be helpful to strengthen our empirical evidence.

3.5.1 International Evidence

We provide international evidence to evaluate the evidence so far on a larger and different dataset. In particular, the demographic variable MY_t is constructed for a large panel of 35 countries over the period 1960-2011 (unbalanced panel)[15]. We consider the performance of augmenting autoregressive models for nominal bond yields [16] against the benchmark where the effects of demographics is restricted to zero.

The results from the estimation are reported in Table 6.

The evidence on the importance of MY in capturing the persistent component of nominal yields is confirmed by the panel estimation. Note that the coefficient on MY is significant with the expected sign even if once we control for the autoregressive component.

3.5.2 The Uncertainty on Future MY

To analyze the uncertainty on projections on MY we use the evidence produced by the Bureau of Census 1975 population report, which publishes projections of future population by age in the United States from 1975 to 2050.[17] The report contains projections based on three different scenarios for fertility, which is kept constant and set to 1.7, 2.1 and 2.7, respectively. All three scenarios are based on the estimated July 1, 1974 population and assume a slight reduction in future mortality and an annual net immigration of 400,000 per year. They differ only in their assumptions about "future fertility". Since there is only 5-year forecasts from 2000-2050, we interpolate 5-year results to obtain the annual series. Then we construct MY_t ratio by using this annually projection results of different fertility rates from 1975 to 2050.

To evaluate the uncertainty surrounding projections for our relevant demographic variable, Panel A in Figure 6 reports plot actual MY_t and projected MY_t in 1975.

The actual annual series of MY_t is constructed based on information released by BoC until December, 2010, while, for the period 2011 to 2050 we use projections contained in the 2008 population report. The figure illustrates that the projections based on the central value of the fertility rate virtually overlaps with the observed data up to 2010 and with the later projections for the period 2011-2050 (Davis and Li [2003]). Different assumptions on fertility have a rather modest impact on MY.

[15] The results are robust when we contruct a smaller panel with balanced data. The demographic data is collected from Worldbank database.

[16] Bond yield are collected from Global Financial data. Long term bond yields are 10-year yields for most of the countries, except Japan (7-year), Finland, South Korea, Singapore (5-year), Mexico(3-year), Hong Kong(2-year).

[17] The report provides annual forecasts from 1975 to 2000 and five-year forecasts from 2000 to 2050.

Another concern about the uncertainty on future MY is regarding the foreign holdings of US debt securities. The theoretical justification of the demographic effect comes from a closed economy model, i.e., it assumes segmented markets. As long as the foreign demographic fluctuations do not counteract the US demographic effect, this assumption should be innocuous. Therefore we compute a demographic variable which takes into account the foreign holdings of US securities, in particular total debt and US Treasury holdings. Following the last report by FED New York published in April 2013, we identify the countries with most US security holdings and compute the middle age-young ratio for those countries, namely Japan, China, UK, Canada, Switzerland, Belgium, Ireland, Luxembourg, Hong Kong.[18] We compute the MY ratio adjusted for foreign holdings; the MY ratio is a weighted average of the MY ratios of those countries with most US security holdings. The weights are computed based on the relative US security holdings reported in Table 7 of the report. In our estimation, we keep the weights fixed at 2012 holdings.

As we see from Panel B in Figure 6, the shape of the demographic variable does not change substantially once we take into account either total debt or treasury holdings. We observe that during the early 2000s, for a short period, the predictions of the original MY variable, and the MY variable adjusted for foreign treasury holdings differ. However, the discrepancy between the two series is temporary and the variables start to co-move again in the out-of-sample period. While foreign holdings of US Treasuries have been increasing during the last decade, there is no reason to think that the trend will continue forever (e.g., Feldstein [2011]).

3.5.3 A Simulation Experiment

To assess the robustness of our results we started from the estimation of a simple autoregressive model for 3-month rates over the full sample. By bootstrapping the estimated residuals we have then constructed one thousand artificial time series for the short-rates. These series are very persistent (based on an estimated AR coefficient of 0.948) and generated under the null of no-significance of MY in explaining the 3-month rates. We have then run one thousand regression by augmenting an autoregressive model for the artificial series with MY_t.

Figure 7 shows that the probability of observing a t-stat of -2.91 on the coefficient on MY_t is 0.039 (the t-stat on MY_t in the actual regression of the 3-month rate, its own lags and the demographic variable). This small fraction of simulated t-stat capable of replicating the observed results provides clear evidence against the hypothesis that our statistical results on demographics and the permanent component of interest rates are

[18]We do not have age structure data for Cayman Islands, Middle East countries and rest of the world. So we account for 60% of foreign bond holdings as of June 2012. Source: Demographic data 1960-2000 from World Bank Population Statistics, Data 2011-2050 from US Census International Database.

spurious.

3.6 Conclusion

The entire term structure of interest rates features a common persistent component. Our evidence has shown that such a persistent component is related to a demographic variable, to ratio of middle-aged to young population, MY_t. The relation between the age structure of population and the equilibrium real returns of bonds is derived in an overlapping generation model in which the demographic structure mimics the pattern of live births in the US. The age composition of the population defines the persistent component in one-period yields as it determines the equilibrium rate in the central bank reaction function. The presence of demographics in short-term rates allows more precise forecast of future policy rates, especially at very long-horizon, and helps modeling the entire term structure. Term structure macro-finance models with demographics clearly dominate traditional term-structure macro-finance models and random walk benchmarks. When demographics are entered among the determinants of short-term rates, a simple model based on a Taylor rule specification for yields at longer maturities outperforms in forecasting traditional term structure models. Better performance is not limited to statistical accuracy, but also confirmed by utility gains using the demographic information. There is a simple intuitive explanation for these results: traditional Taylor-rules and macro finance model do not include an observed determinant of yields capable of capturing their persistence. Linking the long-term central bank target for interest rates to demographics allows for the presence of a slowly moving target for policy rates that fits successfully the permanent component observed in the data. Rudebusch [2002] relates the "illusion" of monetary policy inertia to the possibility that estimated policy rules reflect some persistent shocks that central banks face. Our evidence illustrates that such persistent component is effectively modeled by the age structure of the population. The successful fit is then associated to successful out-of-sample predictions because the main driver of the permanent component in spot rates is exogenous and predictable. Overall, our results show the importance of including the age-structure of population in macro-finance models. As pointed out by Bloom et al. [2003] one of the remarkable features of the economic literature is that demographic factors have so far entered in economic models almost exclusively through the size of population while the age composition of population has also important, and probably neglected, consequences for fluctuations in financial and macroeconomic variables. This chapter has taken a first step in the direction of linking fluctuations in the term structure of interest rates to the age structure of population.

Appendix I: Derivation of Demographic ATSM

We consider the following model specification for pricing bonds with macro and demographic factors:

$$
\begin{aligned}
y_{t,t+n} &= -\frac{1}{n}\left(A_n + B_n' X_t + \Gamma_n \mathrm{MY}_t^n\right) + \varepsilon_{t,t+1} \qquad \varepsilon_{t,t+n} \sim N(0,\sigma_n^2) \\
X_t &= \mu + \Phi X_{t-1} + \nu_t \qquad \nu_t \sim i.i.d.N(0,\Omega) \\
y_{t,t+1} &= \delta_0 + \delta_1' X_t + \delta_2 \mathrm{MY}_t \\
\Lambda_t &= \lambda_0 + \lambda_1 X_t \\
m_{t+1} &= \exp(-y_{t,t+1} - \frac{1}{2}\Lambda_t' \Omega \Lambda_t - \Lambda_t' v_{t+1}) \\
P_t^{(n)} &\equiv \left[\frac{1}{1+Y_{t,t+n}}\right]^n, \; y_{t,t+n} \equiv \ln\left(1+Y_{t,t+n}\right) \\
\Gamma_n \mathrm{MY}_t^n &\equiv \left[\gamma_0^n, \, \gamma_1^n \cdots, \gamma_{n-1}^n\right] \begin{bmatrix} \mathrm{MY}_t \\ \mathrm{MY}_{t+1} \\ \vdots \\ \mathrm{MY}_{t+n-1} \end{bmatrix} \quad X_t = \begin{bmatrix} f_t^\pi \\ f_t^x \\ f_t^{u,1} \\ f_t^{u,2} \\ f_t^{u,3} \end{bmatrix}
\end{aligned}
$$

Bond prices can be recursively computed as:

$$
\begin{aligned}
P_t^{(n)} &= E_t[m_{t+1} P_{t+1}^{(n-1)}] = E_t[m_{t+1} m_{t+2} P_{t+2}^{(n-2)}] \\
&= E_t[m_{t+1} m_{t+2} \cdots m_{t+n} P_{t+n}^{(0)}] \\
&= E_t[m_{t+1} m_{t+2} \cdots m_{t+n} 1] \\
&= E_t[\exp(\sum_{i=0}^{n-1}(-y_{t+i,t+i+1} - \frac{1}{2}\Lambda_{t+i}' \Omega \Lambda_{t+i} - \Lambda_{t+i}' \nu_{t+i+1}))] \\
&= E_t[\exp(A_n + B_n' X_t + \Gamma_n' \mathrm{MY}_t^n)] \\
&= E_t[\exp(-n y_{t,t+n})] \\
&= E_t^Q[\exp(-\sum_{i=0}^{n-1} y_{t+i,t+i+1})]
\end{aligned}
$$

where E_t^Q denotes the expectation under the risk-neutral probability measure, under which the dynamics of the state vector X_t are characterized by the risk neutral vector of constants μ^Q and by the autoregressive matrix Φ^Q

$$
\mu^Q = \mu - \Omega \lambda_0 \text{ and } \Phi^Q = \Phi - \Omega \lambda_1
$$

To derive the coefficients of the model, let us start with $n = 1$:

$$P_t^{(1)} = \exp(-y_{t,t+1}) = \exp(-\delta_0 - \delta_1' X_t - \delta_2 \mathrm{MY}_t)$$

$A_1 = -\delta_0$, $B_1 = -\delta_1$ and $\Gamma_1 = \gamma_0^1 = -\delta_2$, Then for $n+1$,we have $P_t^{(n+1)} = E_t[m_{t+1} P_{t+1}^{(n)}]$

$$= E_t[\exp(-y_{t,t+1} - \frac{1}{2}\Lambda_t'\Omega\Lambda_t - \Lambda_t'\nu_{t+1})\exp(A_n + B_n' X_{t+1} + \Gamma_n \mathrm{MY}_{t+1}^n)]$$
$$= \exp(-y_{t,t+1} - \frac{1}{2}\Lambda_t'\Omega\Lambda_t + A_n) E_t[\exp(-\Lambda_t'\nu_{t+1} + B_n' X_{t+1} + \Gamma_n \mathrm{MY}_{t+1}^n)]$$
$$= \exp(-y_{t,t+1} - \frac{1}{2}\Lambda_t'\Omega\Lambda_t + A_n + \Gamma_n \mathrm{MY}_{t+1}^n) E_t[\exp(-\Lambda_t'\nu_{t+1} + B_n'(\mu + \Phi X_t + \nu_{t+1}))]$$
$$= \exp[-\delta_0 - \delta_1' X_t - \delta_2 MY_t - \frac{1}{2}\Lambda_t'\Omega\Lambda_t + A_n + \Gamma_n \mathrm{MY}_{t+1}^n + B_n'(\mu + \Phi X_t)] E_t[\exp(-\Lambda_t'\nu_{t+1} + B_n'\nu_{t+1})]$$
$$= \exp[-\delta_0 - \delta_1' X_t - \frac{1}{2}\Lambda_t'\Omega\Lambda_t + A_n - \delta_2 \mathrm{MY}_t + B_n'(\mu + \Phi X_t)$$
$$+ \Gamma_n \mathrm{MY}_{t+1}^n] \exp\{E_t[(-\Lambda_t' + B_n')\nu_{t+1}] + \frac{1}{2} var[(-\Lambda_t' + B_n')\nu_{t+1}]\}$$
$$= \exp[-\delta_0 - \delta_1' X_t - \frac{1}{2}\Lambda_t'\Omega\Lambda_t + A_n + B_n'(\mu + \Phi X_t)$$
$$+ \left[-\delta_2, \gamma_0^n, \gamma_1^n \cdots, \gamma_{n-1}^n\right] \mathrm{MY}_t^{n+1}] \exp\{\frac{1}{2} var[(-\Lambda_t' + B_n')\nu_{t+1}]\}$$

To simplify the notation we define $[-\delta_2, \Gamma_n] \equiv \left[-\delta_2, \gamma_0^n, \gamma_1^n \cdots, \gamma_{n-1}^n\right]$

$$= \exp\left\{-\delta_0 - \delta_1' X_t - \frac{1}{2}\Lambda_t'\Omega\Lambda_t + A_n + B_n'(\mu + \Phi X_t) + [-\delta_2, \Gamma_n]\mathrm{MY}_t^{n+1}\right\}$$
$$\exp\left\{\frac{1}{2} E_t[(-\Lambda_t' + B_n')\nu_{t+1}\nu_{t+1}^{'}(-\Lambda_t + B_n)]\right\}$$
$$= \exp\left\{-\delta_0 - \delta_1' X_t - \frac{1}{2}\Lambda_t'\Omega\Lambda_t + A_n + B_n'(\mu + \Phi X_t) + [-\delta_2, \Gamma_n]\mathrm{MY}_t^{n+1}\right\}$$
$$\exp\left\{\frac{1}{2}[\Lambda_t'\Omega\Lambda_t - 2B_n'\Omega\Lambda_t + B_n'\Omega B_n)]\right\}$$
$$= \exp\left\{-\delta_0 + A_n + B_n'\mu + (B_n'\Phi - \delta_1')X_t - B_n'\Omega\Lambda_t + \frac{1}{2}B_n'\Omega B_n + [-\delta_2, \Gamma_n] MY_t^{n+1}\right\}$$
$$= \exp\left\{-\delta_0 + A_n + B_n'\mu + (B_n'\Phi - \delta_1')X_t - B_n'\Omega(\lambda_0 + \lambda_1 X_t) + \frac{1}{2}B_n'\Omega B_n + [-\delta_2, \Gamma_n]\mathrm{MY}_t^{n+1}\right\}$$
$$= \exp\left\{A_1 + A_n + B_n'(\mu - \Omega\lambda_0) + \frac{1}{2}B_n'\Omega B_n + (B_n'\Phi - B_n'\Omega\lambda_1 + B_1')X_t + [-\delta_2, \Gamma_n]\mathrm{MY}_t^{n+1}\right\}$$

Then we can find the coefficients following the difference equations

$$A_{n+1} = A_1 + A_n + B_n'(\mu - \Omega\lambda_0) + \frac{1}{2}B_n'\Omega B_n$$
$$B_{n+1}' = B_n'\Phi - B_n'\Omega\lambda_1 + B_1'$$
$$\Gamma_{n+1} = [-\delta_2, \Gamma_n]$$

Appendix II: Data Description

Demographic Variables: The U.S. annual population estimates series are collected from U.S. Census Bureau and the sample covers estimates from 1900-2050. Middle-aged to young ratio, MY_t is calculated as the ratio of the age group 40-49 to age group 20-29. Past MY_t projections for the period 1950-2013 are hand-collected from various past Census reports available at http://www.census.gov/prod/www/abs/p25.html. MY projections under different fertility rates are based on BoC's 1975 population estimation and projections report.

Spot rate: 3-Month Treasury Bill rate is taken from Goyal and Welch (2008) extended collecting data from St. Louis FRED database.

Bond yields: Bond yields are collected from Gurkaynak, Wright and Sack (2007) dataset, end of month data.

Core Inflation: Time-series of core inflation are collected from St. Louis FRED database.

International data: International bond yields are collected from Global Financial Data up to 2011. Benchmark bond yield is the 10-year constant maturity government bond yields. For Finland and Japan, shorter maturity bonds, 5-year and 7-year, respectively, are used, since a longer time-series is available. International MY_t estimates for the period 1960-2008 are from World Bank Population estimates and projections from 2009-2050 are collected from International database (US Census Bureau).

Macro factors: Stationary output and inflation factors are constructed following the data appendix of Ludvigson and Ng [2009b]. Data series of Group 1 (output) and Group 7 (prices) are extended up to 2013Q4 using data from Bureau of Economic Analysis (BEA) and St. Louis FRED databases.

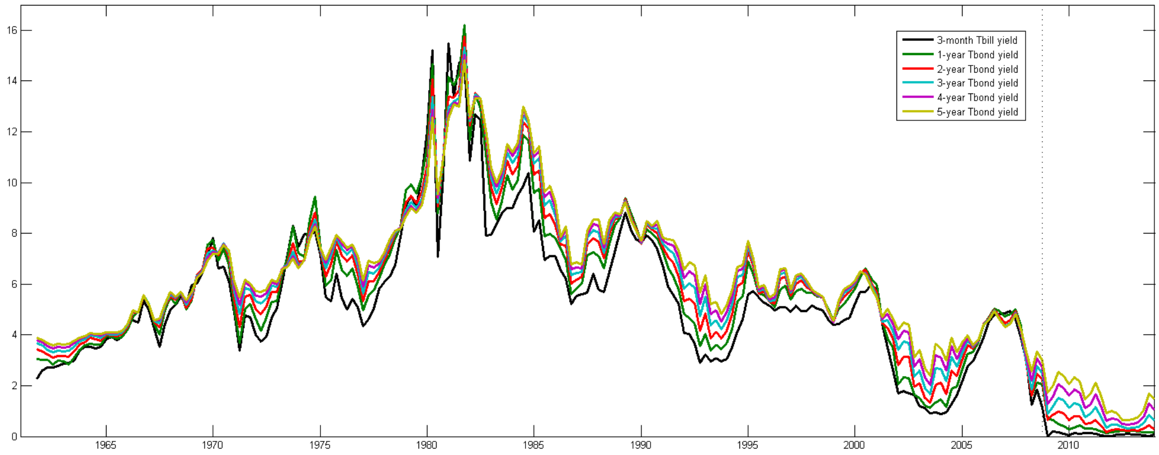

Figure 3.1: Nominal Bond Yields.
This figure shows the US post-war nominal yields. The dotted black vertical line indicates 2008Q3, the beginning of the first round of quantitative easing. Sample 1961Q3-2013Q4.

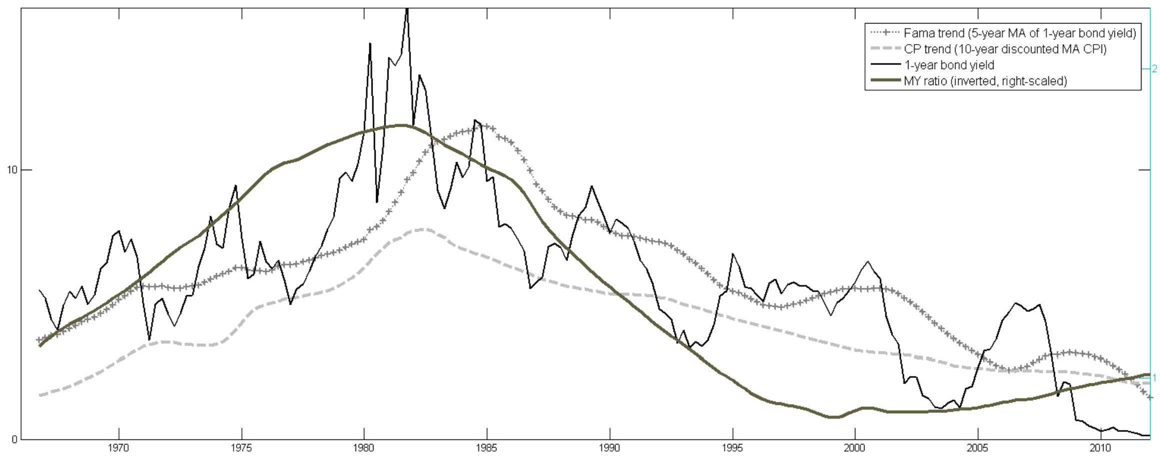

Figure 3.2: 1-Year US Treasury bond yields and the permanent component.
This figure compares the middle-aged to young ratio, MY (inverted, right-scaled, solid dark grey line), FAMA trend (dashed grey line with plus), i.e., 5-year moving average of 1-year Treasury bond yield, CP trend, i.e.,10 year moving average of core inflation (dashed light grey line) with 1-year Treasury bond yield (solid black line). Quarterly sample 1966Q3-2013Q4.

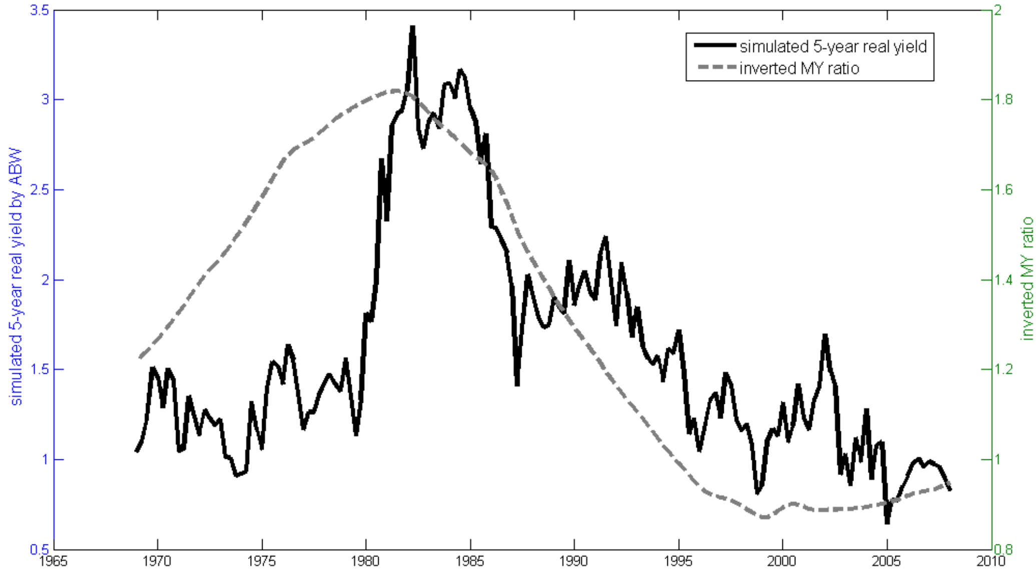

(a) US 5-year real bond yield

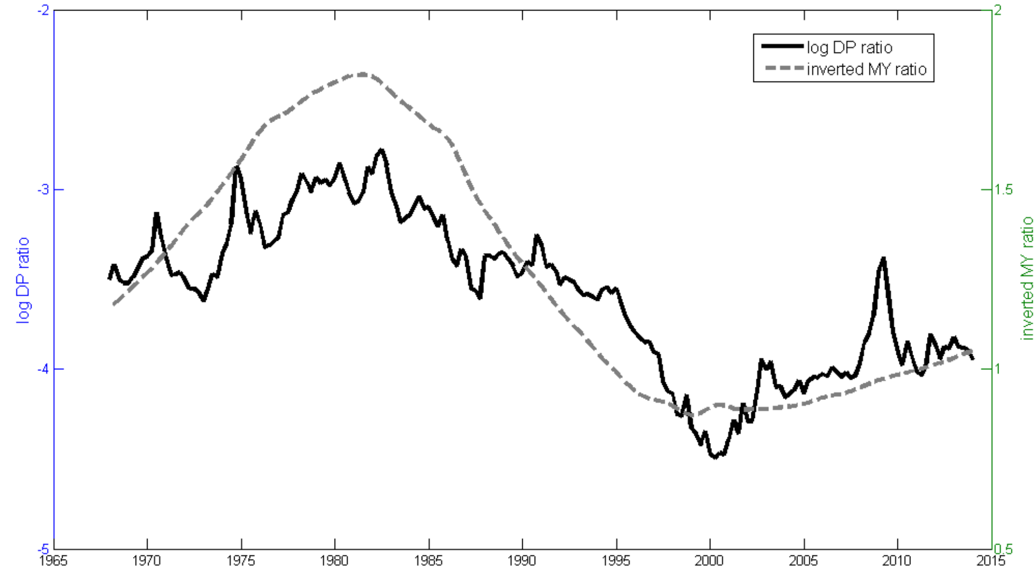

(b) US (log) dividend-price ratio

Figure 3.3: US real bond yield, dividend-price ratio and demgraphics. Panel A plots the US 5-year real bond yield (Ang, Bekaert and Wei (2008), left-scale) and MY (inverted, right-scale). Sample: 1967Q4-2007Q4. Panel B plots the US (log) dividend-price ratio (left-scale) and middle-aged young ratio, MY (inverted, right-scale). Sample: 1967Q4-2013Q4.

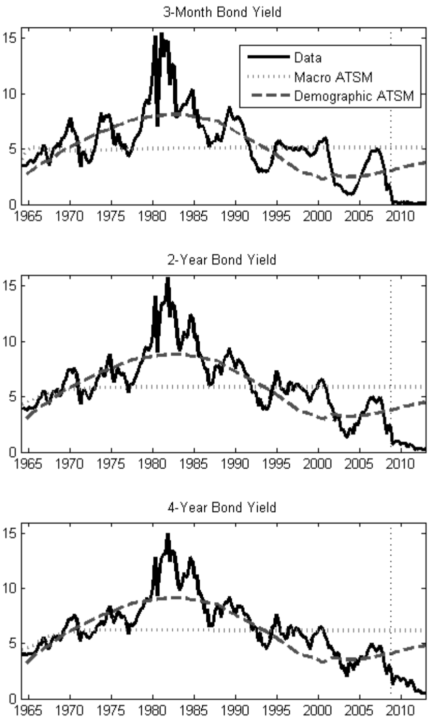

Figure 3.4: Dynamic Simulations.
This figure plots the time series of bond yields (maturity: 3m, 1y, 2y, 3y, 4y, 5y) along with those dynamically simulated series from the benchmark Macro ATSM (dashed light grey line) and Demographic ATSM (solid dark grey line). The affine models with time-varying risk premia are estimated over the full sample and dynamically solved from the first observation onward. The dotted black vertical line indicates 2008Q3, the beginning of the first round of quantitative easing. Sample: 1964Q1-2013Q4.

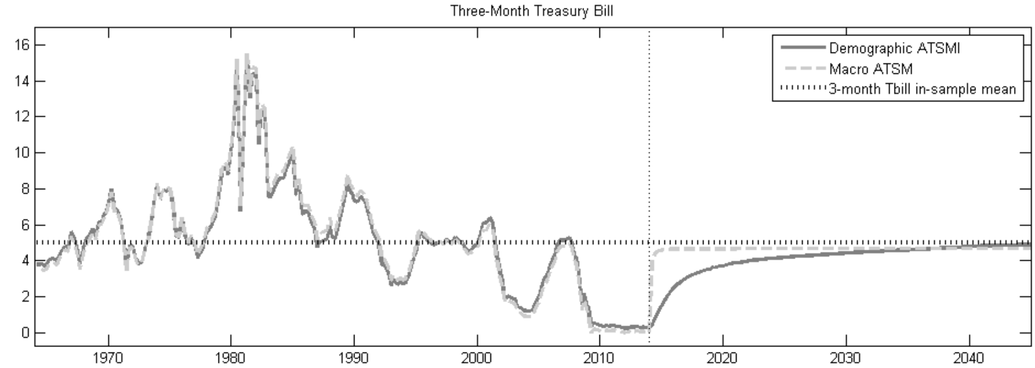

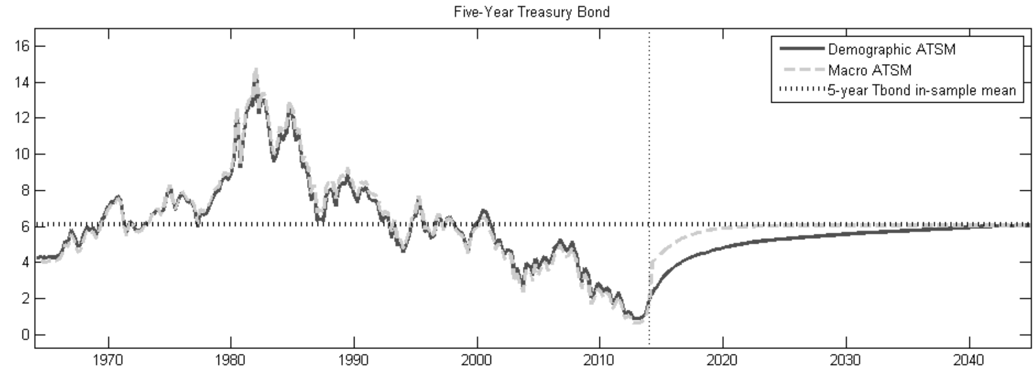

Figure 3.5: In-sample fitted values and dynamically simulated out-of-sample predictions. This figure plots the in-sample estimated values (1964Q1-2013Q4) and out-of-sample predictions (2014Q1-2045Q4) of: 3-month (reported in the upper panel) and 5-year (reported in the lower panel) yields. The Demographic ATSM (solid dark grey lines) and Macro ATSM (dashed light grey lines) are estimated over the whole sample 1964Q1-2013Q4. Using the estimated model parameters, models are solved dynamically forward starting from 1964Q1. The black dash lines are in-sample mean of associated yields, and the vertical dash line shows the end of in-sample estimation period. Sample 1964Q1-2013Q4.

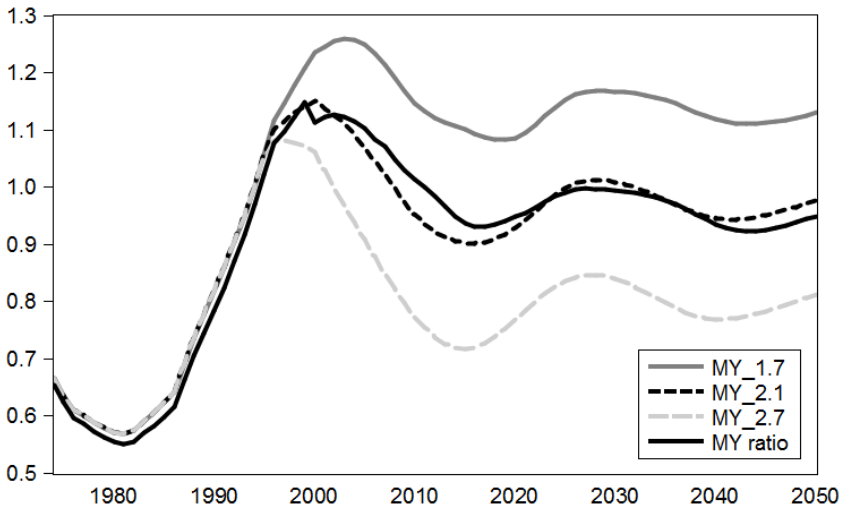

(a) MY projections and fertility rates.

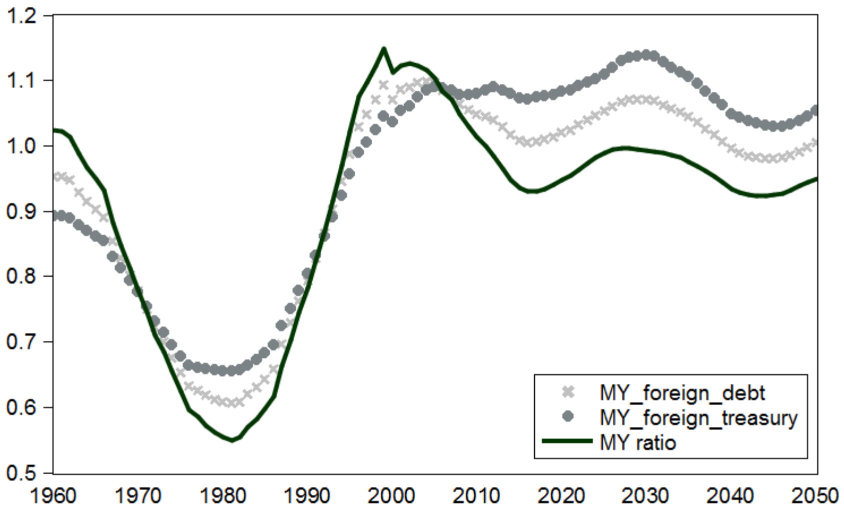

(b) MY projections and foreign holdings.

Figure 3.6: US real bond yield, dividend-price ratio and demgraphics.
This figure plots the middle-aged young (MY) ratio and its long run projections based on alternative scenarios for the fertility rate and foreign holdings. The MY ratio (solid black line) is based on annual reports of BoC while MY_1.7 (solid grey line), MY_2.1 (dashed black line) and MY_2.7 (dashed grey line) in Panel A are predicted in 1975 under 1.7, 2.1 and 2.7 fertility rates, respectively. All the projection information in Panel A is from BoC's 1975 population estimation and projections report. Panel B projections are based on authors' calculation from New York Fed's report on foreign portfolio holdings of U.S. Securities (April 2013).

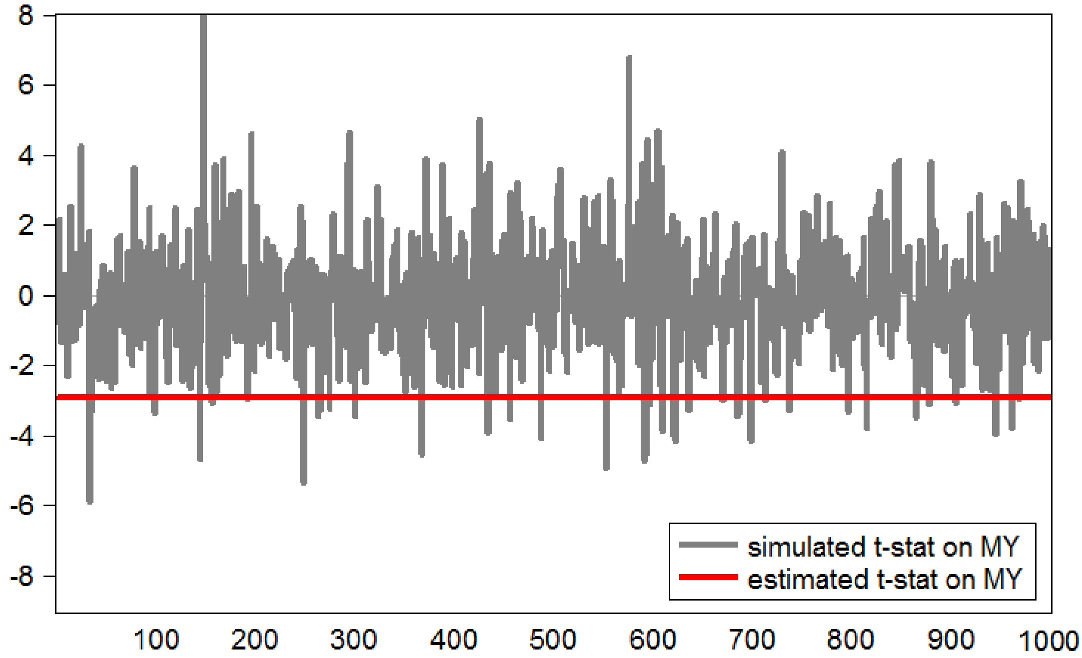

Figure 3.7: Simulated vs. estimated t-statistics.
This figure shows simulated t-statistics on MY ratio which is obtained from an autoregressive model where the dependent variable is an artificial series bootstrapped (1000 simulations) from an autoregressive model for the 3-month rate. The estimated t-statistics is the observed value of the t-statistics on MY ratio in an autoregressive model for the actual 3-month rate augmented with MY ratio.

	Central Moments				Autocorrelations		
	mean	Stdev	Skew	Kurt	Lag 1	Lag 2	Lag 3
3-month	5.0031	3.0855	0.6914	4.1167	0.9351	0.8874	0.8642
1-year	5.4697	3.1833	0.5313	3.5486	0.9499	0.9087	0.8784
2-year	5.6861	3.1087	0.4494	3.3606	0.9553	09208	0.8931
3-year	5.8577	3.0221	0.4371	3.2724	0.9597	0.9290	0.9024
4-year	6.0020	2.9374	0.4605	3.2360	0.9628	09343	0.9084
5-year	6.1273	2.8589	0.4999	3.2282	0.9650	0.9377	0.9123
LN output factor	0.0674	0.9899	−0.4323	5.7257	0.2506	0.0835	0.1342
LN inflation factor	0.0504	1.0065	−0.2071	8.7346	0.0638	0.0228	−0.0302
middle-young ratio	0.8620	0.2023	−0.2075	1.5614	0.9974	0.9936	0.9887

Table 3.1: Summary Statistics of the Data in Term Structure Models.
This table reports the summary statistics. 1, 4, 8, 12, 16, 20 quarter yields are annualized (in percentage) zero coupon bond yields from Fedral Reserve Board (Gurkaynak, Sack and Wright[2006]). LN Inflation and real activity refer to the price and output factors extracted from large dataset using extended time series according to Ludvigson and Ng [2009b]. Quarterly sample 1961Q3-2013Q4.

	Demographic ATSM					Macro ATSM				
Companion form Φ										
	-0.125 (0.082)	0.137 (0.123)	-0.153 (0.140)	-0.253 (0.135)	0.165 (0.111)	-0.133 (0.095)	0.134 (0.104)	0.067 (0.105)	-0.311 (0.132)	0.240 (0.192)
	-0.057 (0.073)	0.348 (0.087)	0.147 (0.090)	0.079 (0.125)	-0.220 (0.112)	-0.054 (0.072)	0.380 (0.104)	-0.092 (0.110)	0.066 (0.168)	-0.279 (0.118)
	-0.028 (0.040)	0.041 (0.026)	0.764 (0.142)	-0.251 (0.040)	0.101 (0.068)	-0.015 (0.009)	0.059 (0.041)	0.981 (0.023)	0.036 (0.120)	-0.087 (0.112)
	-0.017 (0.028)	0.057 (0.021)	-0.178 (0.040)	0.622 (0.174)	0.060 (0.032)	-0.015 (0.039)	0.075 (0.024)	-0.039 (0.060)	0.608 (0.141)	0.172 (0.043)
	-0.002 (0.018)	0.001 (0.021)	0.240 (0.075)	0.189 (0.076)	0.754 (0.095)	0.020 (0.040)	-0.044 (0.052)	0.018 (0.107)	0.305 (0.034)	0.681 (0.127)
Short rate parameters										
δ_1	-0.006 (0.039)	0.157 (0.121)	0.000 (0.000)	0.000 (0.000)	2.739 (0.372)	-0.007 (0.059)	0.263 (0.119)	2.321 (0.588)	0.000 (0.000)	1.544 (0.957)
δ_2	-0.010 (0.0037)					0				
Price of risk λ_0 and λ_1										
$(\lambda_0)^T$	-0.004 (0.014)	-0.004 (0.002)	0.003 (0.003)	0.004 (0.002)	-0.003 (0.001)	-0.108 (0.261)	-0.008 (0.013)	-0.002 (0.009)	-0.002 (0.010)	0.008 (0.014)
λ_1	-0.045 (0.325)		$\cdots$		$\mathbf{0}$	-0.000 (0.004)		$\cdots$		$\mathbf{0}$
		0.685 (0.297)					-0.012 (0.046)			
	$\vdots$		-0.017 (0.053)		$\vdots$	$\vdots$		-0.016 (0.067)		$\vdots$
				-1.162 (0.565)					-1.129 (0.619)	
	$\mathbf{0}$		$\cdots$		-0.972 (0.708)	$\mathbf{0}$		$\cdots$		0.000 (0.000)
Innovation covariance matrix $\Omega^o \times 10^5$										
	0.537	0.033				0.535	0.046			
		0.487					0.494			

Table 3.2: ATSM Full-Sample Estimates.

This table reports the maximum likelihood estimation results for the system (3.7) with time-varying risk premium. The left panel contains estimated results for the unrestricted model which includes the demographic variable MY. The right panel reports estimated results of the system with the restriction δ_2 equal to zero. Standard errors are provided within parentheses. Sample 1964Q1-2013Q4.

Panel A. Random-walk Benchmark										
h	4		8		12		16		20	
	FRMSE (GW)	CW (pvalue)	FRMSE (GW)	CW (pvalue)	FRMSE (GW)	CW (pvalue)	FRMSE (GW)	CW (pvalue)	FRMSE (GW)	CW (pvalue)
$\widehat{y}^{(1/4)}_{t+h\|t}$	1.224	0.814	0.941	5.624	0.813	8.118	0.832	7.057	0.932	5.803
	(0.016)	(0.208)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\widehat{y}^{(1)}_{t+h\|t}$	1.158	0.338	0.923	5.188	0.821	7.466	0.839	6.359	0.935	5.391
	(0.010)	(0.368)	(0.006)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)
$\widehat{y}^{(2)}_{t+h\|t}$	1.158	−0.145	0.951	4.317	0.874	6.088	0.897	5.083	0.991	4.281
	(0.034)	(0.558)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)	(0.013)	(0.000)
$\widehat{y}^{(3)}_{t+h\|t}$	1.158	−0.337	0.982	3.649	0.926	4.890	0.948	4.070	1.036	3.341
	(0.008)	(0.632)	(0.393)	(0.000)	(0.001)	(0.000)	(0.113)	(0.000)	(0.258)	(0.000)
$\widehat{y}^{(4)}_{t+h\|t}$	1.154	−0.397	1.008	3.126	0.969	3.892	0.990	3.286	1.070	2.651
	(0.000)	(0.654)	(0.065)	(0.001)	(0.390)	(0.000)	(0.002)	(0.001)	(0.090)	(0.004)
$\widehat{y}^{(5)}_{t+h\|t}$	1.147	−0.387	1.027	2.705	1.003	3.076	1.023	2.689	1.096	2.182
	(0.000)	(0.651)	(0.002)	(0.003)	(0.075)	(0.001)	(0.172)	(0.004)	(0.016)	(0.015)
Panel B. Macro ATSM Benchmark										
h	4		8		12		16		20	
	FRMSE (GW)	CW (pvalue)	FRMSE (GW)	CW (pvalue)	FRMSE (GW)	CW (pvalue)	FRMSE (GW)	CW (pvalue)	FRMSE (GW)	CW (pvalue)
$\widehat{y}^{(1/4)}_{t+h\|t}$	1.060	1.496	0.894	8.410	0.778	9.673	0.747	9.203	0.756	6.962
	(0.000)	(0.067)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)	(0.002)	(0.001)
$\widehat{y}^{(1)}_{t+h\|t}$	1.014	2.531	0.859	9.218	0.761	10.689	0.744	9.757	0.760	7.158
	(0.002)	(0.006)	(0.011)	(0.000)	(0.010)	(0.000)	(0.005)	(0.000)	(0.001)	(0.000)
$\widehat{y}^{(2)}_{t+h\|t}$	0.989	2.967	0.837	9.487	0.752	11.280	0.743	10.080	0.766	7.485
	(0.000)	(0.002)	(0.001)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\widehat{y}^{(3)}_{t+h\|t}$	0.975	3.181	0.825	9.596	0.749	11.476	0.745	10.122	0.773	7.687
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\widehat{y}^{(4)}_{t+h\|t}$	0.965	3.394	0.817	9.713	0.746	11.491	0.748	10.083	0.778	7.829
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\widehat{y}^{(5)}_{t+h\|t}$	0.959	3.598	0.811	9.801	0.745	11.387	0.751	9.980	0.784	7.906
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

Table 3.3: Affine Model Out-of-Sample Forecasts.

This table provides yield forecast comparison of Demographic ATSM against the Random Walk model (Panel A) and Macro ATSM (Panel B) benchmarks. We use the in-sample estimators, from 1961Q3 to 1981Q2, to generate out-of-sample forecasts until 2013Q4. h indicates 4, 8, 12, 16, 20 quarter out-of-sample forecasts. We measure forecasting performance as the ratio of the root mean squared forecast error (RMSFE) of our model against the benchmarks. We report in parentheses the p-values of the forecasting test due to Giacomini and White (2006) in the columns with FRMSE. A p-value below 0.01 (0.05, 0.10) indicates a significant difference in forecasting performance at the 1% (5%, 10%) level. We also measure forecasting performance using Clark and West (2006, 2007) test. We report the test statistics in the columns CW for each horizon together with p-values in parentheses below. Quarterly sample 1981Q3- 2013Q4.

Panel A. Bond Yields - Quadratic Loss					
h	4	8	12	16	20
	$\widehat{\lambda}$ $(t^{\lambda=0})$ $[t^{\lambda=1}]$	$\widehat{\lambda}$ $(t^{\lambda=0})$ $[t^{\lambda=1}]$	$\widehat{\lambda}$ $(t^{\lambda=0})$ $[t^{\lambda=1}]$	$\widehat{\lambda}$ $(t^{\lambda=0})$ $[t^{\lambda=1}]$	$\widehat{\lambda}$ $(t^{\lambda=0})$ $[t^{\lambda=1}]$
$\widehat{y}^{(1/4)}_{t+h\|t}$	0.816	0.098	−0.238	−0.035	0.307
	(4.60***)	(0.56)	(−1.19)	(−0.15)	(2.02**)
	[−1.04]	[−5.16***]	[−6.17***]	[−4.36***]	[−4.55***]
$\widehat{y}^{(1)}_{t+h\|t}$	0.708	−0.040	−0.232	−0.016	0.316
	(3.61***)	(−0.30)	(−1.17)	(−0.07)	(2.26**)
	[−1.49]	[−7.83***]	[−6.19***]	[−4.59***]	[−4.88***]
$\widehat{y}^{(2)}_{t+h\|t}$	0.726	−0.076	−0.134	0.112	0.413
	(3.17***)	(−0.59)	(−0.73)	(0.59)	(3.51***)
	[−1.20]	[−8.35***]	[−6.18***]	[−4.70***]	[−4.98***]
$\widehat{y}^{(3)}_{t+h\|t}$	0.744	−0.068	−0.019	0.226	0.490
	(2.97***)	(−0.50)	(−0.11)	(1.34)	(4.63***)
	[−1.02]	[−7.74***]	[−6.04***]	[−4.60***]	[−4.81***]
$\widehat{y}^{(4)}_{t+h\|t}$	0.754	−0.035	0.091	0.320	0.548
	(2.83***)	(−0.23)	(0.57)	(2.04**)	(5.48***)
	[−0.92]	[−6.80***]	[−5.68***]	[−4.34***]	[−4.53***]
$\widehat{y}^{(5)}_{t+h\|t}$	0.755	0.011	0.188	0.395	0.590
	(2.71***)	(0.07)	(1.20)	(2.62***)	(6.05***)
	[−0.88]	[−5.93***]	[−5.16***]	[−4.01***]	[−4.20***]

Panel B. Bond Excess Returns - Portfolio Utility Loss		
holding period	1-year	2-year
Demographic ATSM	0.595	0.316
	(1.57)	(1.85*)
	[−1.07]	[−4.00***]
Macro ATSM	0.611	0.707
	(2.61***)	(1.94*)
	[−1.67*]	[−0.80]

Table 3.4: Out-of-Sample Forecast Usefulness.
This table provides results on forecasting usefulness according to Carriero and Giacomini (2011) test. Panel A shows yield forecast comparison of Demographic ATSM against the Random Walk benchmark. Panel B shows bond excess return forecast comparison of Demographic and Macro ATSM against the Random Walk benchmark. We use the in-sample estimators, from 1961Q3 to 1981Q2, to generate out-of-sample forecasts until 2013Q4. h indicates 4, 8, 12, 16, 20 quarter out-of-sample forecasts. We report $\widehat{\lambda}$, the weight on the restricted (random walk) model, and the test statistics associated with $\lambda = 0$ and $\lambda = 1$ in the parentheses below. Quarterly sample 1981Q3- 2013Q4.

$$y^{(1)}_{t+4x}\text{-}y^{(1)}_{t} = a^x + b^x D_t + c^x[f^{(1)}_{t,t+4x}\text{-}y^{(1)}_{t}] + d^x[y^{(1)}_{t}\text{-}P^{(1),i}_{t}] + \varepsilon_{t+12x}$$

$$P^{(1),1}_{t} = \frac{1}{20}\sum_{i=1}^{20} y^{(1)}_{t-i-1}\ (FAMA),\quad P^{(1),2}_{t} = \frac{\sum_{i=1}^{40} 0.96^{i-1}\pi_{t-i-1}}{\sum_{i=1}^{40} 0.96^{i-1}}(CP),\ P^{(1),3}_{t} = e^x \frac{1}{4}\sum_{i=1}^{4} MY_{t+i-1}$$

	a^x (s.e.)	b^x (s.e.)	c^x (s.e.)	d^x (s.e.)	e^x (s.e.)	R^2	
no cycle	−1.99 (0.26)	2.36 (0.134)	1.29 (0.17)			0.28	
Fama cycle no dummy	−0.74 (0.25)		0.87 (0.28)	−0.01 (0.11)		0.11	
Fama cycle	−1.88 (0.25)	3.30 (0.38)	0.42 (0.24)	−0.54 (0.12)		0.35	$x = 2$
CP cycle	0.78 (0.38)		−0.17 (0.27)	−0.63 (0.13)		0.20	
MY cycle	6.83 (1.16)		0.11 (0.20)	−0.54 (0.08)	−0.093 (0.009)	0.27	
no cycle	−3.04 (0.26)	3.50 (0.33)	2.01 (0.16)			0.50	
Fama cycle no dummy	−1.42 (0.27)	1.79 (0.31)	0.20 (0.13)			0.22	
Fama cycle	−2.93 (0.25)	4.35 (0.37)	1.20 (0.24)	−0.50 (0.11)		0.54	$x = 3$
CP cycle	0.22 (0.44)		0.45 (0.32)	−0.58 (0.15)		0.26	
MY cycle	8.13 (1.26)		0.49 (0.21)	−0.65 (0.09)	−0.095 (0.010)	0.39	
no cycle	−3.56 (0.25)	4.18 (0.32)	2.23 (0.16)			0.59	
Fama cycle no dummy	−1.75 (0.29)		2.21 (0.32)	0.36 (0.13)		0.25	
Fama cycle	−3.46 (0.24)	4.90 (0.136)	1.55 (0.23)	−0.43 (0.11)		0.62	$x = 4$
CP cycle	−0.17 (0.49)		0.77 (0.34)	−0.47 (0.17)		0.25	
MY cycle	9.36 (01.30)		0.50 (0.22)	−0.75 (0.09)	−0.094 (0.010)	0.43	
no cycle	−3.57 (0.26)	4.37 (0.33)	2.00 (0.16)			0.56	
Fama cycle no dummy	−1.65 (0.30)		1.98 (0.32)	0.36 (0.14)		0.18	
Fama cycle	−3.46 (0.25)	5.15 (0.37)	1.27 (0.24)	−0.46 (0.11)		0.59	$x = 5$
CP cycle	−0.41 (0.53)		0.77 (0.36)	−0.33 (0.018)		0.17	
MY cycle	10.48 (1.35)		0.18 (0.23)	−0.83 (0.010)	−0.092 (0.010)	0.41	

Table 3.5: Predictive Regressions for the 1-year Spot Rate.

This table shows predictive regressions with alternative permanant components. $f^{(1)}_{t,t+4x}$ is one-year forward rate observed at time t of an investment with settlement after $3x$ years and maturity in $4x$ years, $y^{(1)}_{t}$is 1-year spot rate, π_t is annual core CPI inflation, MY_t is the middle aged to young ratio, D_t is a time dummy($D_t = 1$ from 1961Q3 to 1981Q2). Standard errors are Hansen-Hodrick (1980) adjusted. Sample: 1961Q3-2013Q4.

Benchmark model:	$R_{lt} = \alpha_0 + \alpha_1 R_{lt-1} + \varepsilon_t$		
Augmented model:	$R_{lt} = \beta_0 + \beta_1 R_{lt-1} + \beta_2 \mathrm{MY}_t + \varepsilon_t$		
Specification	R_{lt-1}	MY_t	$\bar{\mathrm{R}}^2$
(1)	0.729 (8.39***)		0.55
(2)	0.676 (7.29***)	−0.044 (−3.78***)	0.58

Table 3.6: International Panel.

This table reports international evidence. Pooled regression coefficients account for country fixed effects. R_{lt} is the nominal bond yield. Specification (1) is the benchmark model and specification (2) is the augmented model with MY_t. The reported t-statistics are based on Driscoll-Kraay (1998) standard errors robust to general forms of cross-sectional (spatial) and temporal dependence. Asterisks *, ** and *** indicate significance at the 10 percent, 5 percent and 1 percent levels, respectively. Last column report within group R^2. There are 35 countries, and 1530 observations in an (unbalanced) panel. Annual sample 1960-2011.

Bibliography

Andrew B Abel. Asset Prices under Habit Formation and Catching Up with the Joneses. *American Economic Review*, 80(2):38–42, May 1990.

Andrew B Abel. Will bequests attenuate the predicted meltdown in stock prices when baby boomers retire? *Review of Economics and Statistics*, 83(4):589–595, 2001.

Abhay Abhyankar, Devraj Basu, and Alexander Stremme. Portfolio efficiency and discount factor bounds with conditioning information: An empirical study. *Journal of Banking & Finance*, 31(2):419–437, 2007.

Eric M Aldrich and A Ronald Gallant. Habit, long-run risks, prospect? a statistical inquiry. *Journal of Financial Econometrics*, 9(4):589–618, 2011a.

Eric M Aldrich and A Ronald Gallant. Habit, long-run risks, prospect? a statistical inquiry. *Journal of Financial Econometrics*, 9(4):589–618, 2011b.

Carlo Altavilla, Raffaella Giacomini, and Giuseppe Ragusa. Anchoring the yield curve using survey expectations. 2014.

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

Andrew Ang and Geert Bekaert. Regime switches in interest rates. *Journal of Business & Economic Statistics*, 20(2):163–182, 2002.

Andrew Ang and Angela Maddaloni. Do demographic changes affect risk premiums? evidence from international data. Technical report, National Bureau of Economic Research, 2003.

Andrew Ang and Monika Piazzesi. A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary economics*, 50(4):745–787, 2003.

Andrew Ang, Sen Dong, and Monika Piazzesi. No-arbitrage taylor rules. Technical report, National Bureau of Economic Research, 2007.

Andrew Ang, Geert Bekaert, and Min Wei. The term structure of real rates and expected inflation. *The Journal of Finance*, 63(2):797–849, 2008.

D. Avramov. Stock return predictability and model uncertainty. *Journal of Financial Economics*, 64(3):423–458, 2002.

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

David Backus, Mikhail Chernov, and Stanley E. Zin. Sources of entropy in representative agent models. NBER Working Papers 17219, National Bureau of Economic Research, Inc, July 2011b.

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

Gurdip S Bakshi and Zhiwu Chen. Baby boom, population aging, and capital markets. *Journal of Business*, pages 165–202, 1994.

Pierluigi Balduzzi, Sanjiv Ranjan Das, and Silverio Foresi. The central tendency: A second factor in bond yields. *Review of Economics and Statistics*, 80(1):62–72, 1998.

Turan Bali, Massoud Heidari, and Liuren Wu. Predictability of interest rates and interest-rate portfolios. *Journal of Business & Economic Statistics*, 27(4):517–527, 2009.

Federico M Bandi and Benoit Perron. Long-run risk-return trade-offs. *Journal of Econometrics*, 143(2):349–374, 2008.

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

Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509, 08 2004.

Ravi Bansal, Dana Kiku, and Amir Yaron. An empirical evaluation of the long-run risks model for asset prices. *Critical Finance Review*, 1(1):183–221, January 2012a.

Ravi Bansal, Dana Kiku, and Amir Yaron. Risks for the long run: Estimation with time aggregation. Technical report, National Bureau of Economic Research, 2012b.

Ravi Bansal, Dana Kiku, and Amir Yaron. Risks for the long run: Estimation with time aggregation. Technical report, National Bureau of Economic Research, 2012c.

Ravi Bansal, Dana Kiku, and Amir Yaron. An empirical evaluation of the long-run risks model for asset prices. Technical Report 1, January 2012d.

N. Barberis. Investing for the long run when returns are predictable. *The Journal of Finance*, 55(1):225–264, 2002.

Francisco Barillas. Can we exploit predictability in bond markets? *Available at SSRN 1787567*, 2011.

Robert J. Barro. Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, pages 823–866, 2006a.

Robert J Barro. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, pages 823–866, 2006b.

Jason Beeler and John Y. Campbell. The long-run risks model and aggregate asset prices: An empirical assessment. Nber working papers, National Bureau of Economic Research, Inc, March 2009.

Geert Bekaert and Eric Engstrom. Asset return dynamics under bad environment good environment fundamentals. NBER Working Papers 15222, National Bureau of Economic Research, Inc, August 2010a.

Geert Bekaert and Eric Engstrom. Inflation and the stock market: Understanding the fed model. *Journal of Monetary Economics*, 57(3):278–294, 2010b.

Geert Bekaert and Jun Liu. Conditioning information and variance bounds on pricing kernels. *Review of Financial Studies*, 17(2):339–378, 2004.

Geert Bekaert, Seonghoon Cho, and Antonio Moreno. New keynesian macroeconomics and the term structure. *Journal of Money, Credit and Banking*, 42(1):33–62, 2010.

Steven Michael Bergantino. *Life cycle investment behavior, demographics and asset prices.* PhD thesis, Massachusetts Institute of Technology, 1998.

Ben S Bernanke. The coming demographic transition: Will we treat future generations fairly? *Speech before the Washington Economic Club, Washington DC, October*, 4, 2006.

Jaroslav Borovicka, Lars P. Hansen, Mark Hendricks, and José A. Scheinkman. Risk-price dynamics. *Journal of Financial Econometrics*, 9(1):3–65, 2011.

Robin Jermyn Brooks. *Asset market and savings effects of demographic transitions.* PhD thesis, Yale University, 1998.

Craig Burnside. Hansen-jagannathan bounds as classical tests of asset-pricing models. *Journal of Business and Economic Statistics*, 12(1):pp. 57–79, 1994.

John Y. Campbell. Stock returns and the term structure. *Journal of Financial Economics*, 18(2):373 – 399, 1987.

John Y Campbell. Consumption-based asset pricing. *Handbook of the Economics of Finance*, 1:803–887, 2003.

John Y Campbell and John Ammer. What moves the stock and bond markets? a variance decomposition for long-term asset returns. *The Journal of Finance*, 48(1):3–37, 1993.

John Y. Campbell and John Cochrane. Force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251, 1999.

John Y. Campbell and Robert J. Shiller. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, 1(3):195–228, 1988a.

John Y. Campbell and Robert J. Shiller. Stock prices, earnings and expected dividends. *Journal of Finance*, 43:661–676, July 1988b.

Andrea Carriero and Raffaella Giacomini. How useful are no-arbitrage restrictions for forecasting the term structure of interest rates? *Journal of Econometrics*, 164(1): 21–34, 2011.

Stephen Cecchetti, Pok-sang Lam, and Nelson C. Mark. Testing volatility restrictions on intertemporal marginal rates of substitution implied by euler equations and asset returns. *Journal of Finance*, 49(1):123–52, 1994.

Ren-Raw Chen and Louis Scott. Maximum likelihood estimation for a multifactor equilibrium model of the term structure of interest rates. *The Journal of Fixed Income*, 3 (3):14–31, 1993.

Jens HE Christensen, Francis X Diebold, and Glenn D Rudebusch. The affine arbitrage-free class of nelson–siegel term structure models. *Journal of Econometrics*, 164(1):4–20, 2011.

Anna Cieslak and Pavol Povala. Expected returns in treasury bonds. *Unpublished working paper, Northwestern University working paper. Evanston and Chicago, IL*, 2013.

Richard Clarida, Jordi Gali, and Mark Gertler. Monetary policy rules and macroeconomic stability: evidence and some theory. *Quarterly Journal of Economics*, 115(1):147–180, 2000.

Todd E Clark and Kenneth D West. Using out-of-sample mean squared prediction errors to test the martingale difference hypothesis. *Journal of Econometrics*, 135(1):155–186, 2006.

Todd E Clark and Kenneth D West. Approximately normal tests for equal predictive accuracy in nested models. *Journal of econometrics*, 138(1):291–311, 2007.

John H. Cochrane. *Asset Pricing*. Princeton University Press, December 2001.

John H. Cochrane. The dog that did not bark: A defense of return predictability. *Review of Financial Studies*, 21(4):1533–1575, 2008.

John H Cochrane and Lars Peter Hansen. Asset pricing explorations for macroeconomics. In *NBER Macroeconomics Annual 1992, Volume 7*, pages 115–182. MIT Press, 1992.

John H. Cochrane and Monika Piazzesi. Bond risk premia. *American Economic Review*, 95(1):138–160, March 2005.

Qiang Dai and Kenneth J Singleton. Expectation puzzles, time-varying risk premia, and affine models of the term structure. *Journal of financial Economics*, 63(3):415–441, 2002.

E Philip Davis and Christine Li. Demographics and financial asset prices in the major industrial economies. 2003.

Stefano DellaVigna and Joshua M Pollet. Demographics and industry returns. *The American Economic Review*, pages 1667–1702, 2007.

Hans Dewachter and Marco Lyrio. Macro factors and the term structure of interest rates. *Journal of Money, Credit and Banking*, pages 119–140, 2006.

Francis X Diebold and Canlin Li. Forecasting the term structure of government bond yields. *Journal of econometrics*, 130(2):337–364, 2006.

Francis X Diebold and Glenn D Rudebusch. *Yield Curve Modeling and Forecasting: The Dynamic Nelson-Siegel Approach*. Princeton University Press, 2013.

Francis X Diebold, Glenn D Rudebusch, and S Boragan Aruoba. The macroeconomy and the yield curve: a dynamic latent factor approach. *Journal of econometrics*, 131(1): 309–338, 2006.

Bart F Diris. Model uncertainty for long-term investors. *Available at SSRN 1786587*, 2011.

Gregory Duffee. Forecasting interest rates. *Handbook of Economic Forecasting*, 2:385–426, 2012.

Gregory R Duffee. Term premia and interest rate forecasts in affine models. *The Journal of Finance*, 57(1):405–443, 2002.

Gregory R Duffee. Sharpe ratios in term structure models. Technical report, Working papers//the Johns Hopkins University, Department of Economics, 2011.

Darrell Duffie and Rui Kan. A yield-factor model of interest rates. *Mathematical finance*, 6(4):379–406, 1996.

George W Evans and Seppo Honkapohja. Adaptive learning and monetary policy design. *Journal of Money, Credit and Banking*, pages 1045–1072, 2003.

Martin DD Evans. Real risk, inflation risk, and the term structure*. *The Economic Journal*, 113(487):345–389, 2003.

Eugene F Fama. The behavior of interest rates. *Review of Financial Studies*, 19(2): 359–379, 2006.

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

Eugene F. Fama and Kenneth R. French. Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics*, 25(1):23–49, November 1989.

Carlo A Favero, Arie E Gozluklu, and Andrea Tamoni. Demographic trends, the dividend-price ratio, and the predictability of long-run stock market returns. *Journal of Financial and Quantitative Analysis*, 46(05):1493–1520, 2011.

Carlo A Favero, Linlin Niu, and Luca Sala. Term structure forecasting: No-arbitrage restrictions versus large information set. *Journal of Forecasting*, 31(2):124–156, 2012.

Carlo A Favero, Fulvio Ortu, Andrea Tamoni, and Haoxi Yang. Implications of returns predictability across horizons for asset pricing models. *Available at SSRN 2184254*, 2013.

Martin S Feldstein. What's next for the dollar? Technical report, National Bureau of Economic Research, 2011.

Wayne E. Ferson and Andrew F. Siegel. Stochastic discount factor bounds with conditioning information. *Review of Financial Studies*, 16(2):567–595, 2003.

Wayne E. Ferson and Andrew F. Siegel. Testing portfolio efficiency with conditioning information. *Review of Financial Studies*, 22(7):2535–2558, July 2009.

Xavier Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics*, 127(2):645–700, 2012.

A. Ronald Gallant, Lars Peter Hansen, and George Tauchen. Using conditional moments of asset payoffs to infer the volatility of intertemporal marginal rates of substitution. *Journal of Econometrics*, 45(1-2):141–179, 1990.

Antonio Gargano, Davide Pettenuzzo, and Allan G Timmermann. Bond return predictability: Economic value and links to the macroeconomy. *Available at SSRN 2471273*, 2014.

John Geanakoplos, Michael Magill, and Martine Quinzii. Demography and the long-run predictability of the stock market. *Brookings Papers on Economic Activity*, 2004(1): 241–325, 2004.

Anisha Ghosh, Christian Julliard, and Alex P. Taylor. What is the consumption-capm missing? an informative-theoretic framework for the analysis of asset pricing models. Fmg discussion papers, Financial Markets Group, October 2011.

Eric Ghysels, Casidhe Horan, and Emanuel Moench. Forecasting throuth the rear-view mirror: Data revisions and bond return predictability. Technical report, Staff Report, Federal Reserve Bank of New York, 2012.

Raffaella Giacomini and Halbert White. Tests of conditional predictive ability. *Econometrica*, 74(6):1545–1578, 2006.

Francois Gourio. Time-series predictability in the disaster model. *Finance Research Letters*, 5(4):191–203, 2008.

Amit Goyal. Demographics, stock market flows, and stock returns. *Journal of Financial and Quantitative Analysis*, 39(01):115–142, 2004.

Amit Goyal and Ivo Welch. Predicting the equity premium with dividend ratios. *Management Science*, 49(5):639–654, 2003.

Morin A. Gozluklu, A. E. Inflation, stock market and demographic fluctuations. *Mimeo, University of Warwick*, 2014.

Stephen F Gray. Modeling the conditional distribution of interest rates as a regime-switching process. *Journal of Financial Economics*, 42(1):27–62, 1996.

Massimo Guidolin and Daniel L Thornton. Predictions of short-term rates and the expectations hypothesis. *Federal Reserve Bank of St. Louis Working Paper Series No*, 2010.

Lars Peter Hansen. Modeling the long run: Valuation in dynamic stochastic economies. Technical report, National Bureau of Economic Research, 2008.

Lars Peter Hansen. Modeling the long run: Valuation in dynamic stochastic economies. NBER Working Papers 14243, National Bureau of Economic Research, Inc, August 2009.

Lars Peter Hansen and Ravi Jagannathan. Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–62, April 1991.

Lars Peter Hansen and Scott F Richard. The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica: Journal of the Econometric Society*, pages 587–613, 1987.

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

Peter Hördahl, Oreste Tristani, and David Vestin. A joint econometric model of macroeconomic and term-structure dynamics. *Journal of Econometrics*, 131(1):405–444, 2006.

Jing-zhi Huang and Zhan Shi. Determinants of bond risk premia. Technical report, Technical report, Penn State University, 2011.

Caroline Jardet, Alain Monfort, and Fulvio Pegoraro. No-arbitrage near-cointegrated var (p) term structure models, term premia and gdp growth. *Journal of Banking & Finance*, 37(2):389–402, 2013.

KEYNES JOHN MAYNARD. The general theory of employment, interest, and money. 1936.

C Kirby. The restrictions on predictability implied by rational asset pricing models. *Review of Financial Studies*, 11(2):343–382, 1998.

Joel Lander, Athanasios Orphanides, and Martha Douvogiannis. Earnings forecasts and the predictability of stock returns: evidence from trading the s&p. *The Journal of Portfolio Management*, 23(4):24–35, 1997.

Martin Lettau and Sydney Ludvigson. Consumption, aggregate wealth, and expected stock returns. *Journal of Finance*, 56(3):815–849, 06 2001.

Thomas Lindh and Bo Malmberg. Can age structure forecast inflation trends? *Journal of Economics and Business*, 52(1):31–49, 2000.

Sydney C Ludvigson and Serena Ng. Macro factors in bond risk premia. *Review of Financial Studies*, page hhp081, 2009a.

Sydney C Ludvigson and Serena Ng. Macro factors in bond risk premia. *Review of Financial Studies*, page hhp081, 2009b.

Ulrike Malmendier and Stefan Nagel. Learning from inflation experiences. *NBER Working Paper*, 2013.

Henry McMillan and Jerome B Baesel. The role of demographic factors in interest rate forecasting. *Managerial and Decision Economics*, 9(3):187–195, 1988.

Rajnish Mehra and Edward C Prescott. The equity premium: A puzzle. *Journal of monetary Economics*, 15(2):145–161, 1985.

Franco Modigliani and Richard Brumberg. Utility analysis and the consumption function: An interpretation of cross-section data. *Franco Modigliani*, 1, 1954.

Emanuel Mönch. Forecasting the yield curve in a data-rich environment: A no-arbitrage factor-augmented var approach. *Journal of Econometrics*, 146(1):26–43, 2008.

Emi Nakamura, Jón Steinsson, Robert Barro, and José Ursúa. Crises and recoveries in an empirical model of consumption disasters. *American Economic Journal: Macroeconomics*, 5(3):35–74, 2013.

Daniela Osterrieder and Peter C Schotman. The volatility of long-term bond returns: Persistent interest shocks and time-varying risk premiums. *Available at SSRN 2123412*, 2012.

L'uboš Pástor and Robert F Stambaugh. Are stocks really less volatile in the long run? *The Journal of Finance*, 67(2):431–478, 2012.

Francisco Peñaranda and Enrique Sentana. A unifying approach to the empirical evaluation of asset pricing models. CEPR Discussion Papers 7943, C.E.P.R. Discussion Papers, 2013.

Marcello Pericoli and Marco Taboga. Canonical term-structure models with observable factors and the dynamics of bond risk premia. *Journal of Money, Credit and Banking*, 40(7):1471–1488, 2008.

Dimitris N Politis and Halbert White. Automatic block-length selection for the dependent bootstrap. *Econometric Reviews*, 23(1):53–70, 2004.

James M Poterba. Demographic structure and asset returns. *Review of Economics and Statistics*, 83(4):565–584, 2001.

Adriano A Rampini and S Viswanathan. Household risk management. *Available at SSRN 2431133*, 2014.

Thomas A Rietz. The equity risk premium a solution. *Journal of monetary Economics*, 22(1):117–131, 1988.

Stephen A Ross. *Neoclassical finance*. Princeton University Press, 2005.

Glenn D Rudebusch. Term structure evidence on interest rate smoothing and monetary policy inertia. *Journal of monetary economics*, 49(6):1161–1187, 2002.

Glenn D Rudebusch and Tao Wu. A macro-finance model of the term structure, monetary policy and the economy*. *The Economic Journal*, 118(530):906–926, 2008.

Lucio Sarno, Paul Schneider, and Christian Wagner. The economic value of predicting bond risk premia: Can anything beat the expectations hypothesis? 2014.

Paul Söderlind, Ulf Söderström, and Anders Vredin. Dynamic taylor rules and the predictability of interest rates. *Macroeconomic Dynamics*, 9(03):412–428, 2005.

Lars EO Svensson. Inflation forecast targeting: Implementing and monitoring inflation targets. *European Economic Review*, 41(6):1111–1146, 1997.

John B Taylor. Discretion versus policy rules in practice. In *Carnegie-Rochester conference series on public policy*, volume 39, pages 195–214. Elsevier, 1993.

Daniel L Thornton and Giorgio Valente. Out-of-sample predictions of bond excess returns and forward rates: An asset allocation perspective. *Review of Financial Studies*, 25(10): 3141–3168, 2012.

Luis M. Viceira. Bond risk, bond return volatility, and the term structure of interest rates. *International Journal of Forecasting*, 28(1):97–117, February 2012.

Tuomo Vuolteenaho and John Campbell. Inflation illusion and stock prices. 2004.

Jessica A. Wachter. A consumption-based model of the term structure of interest rates. *Journal of Financial Economics*, 79(2):365–399, February 2006.

Jessica A Wachter. Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, 68(3):987–1035, 2013a.

Jessica A Wachter. Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, 68(3):987–1035, 2013b.

I. Welch and A. Goyal. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508, 2008a.

Ivo Welch and Amit Goyal. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508, July 2008b.

Halbert White. A reality check for data snooping. *Econometrica*, 68(5):1097–1126, 2000.

Knut Wicksell. *Interest and prices.* Ludwig von Mises Institute, 1936.

Michael Woodford. The taylor rule and optimal monetary policy. *American Economic Review*, pages 232–237, 2001.

Yihong Xia. Learning about predictability: The effects of parameter uncertainty on dynamic asset allocation. *The Journal of Finance*, 56(1):205–246, 2001.

Guofu Zhou. How much stock return predictability can we expect from an asset pricing model? *Economics Letters*, 108(2):184–186, 2010.