Essays in Financial Economic Modeling

by

Tyler Jorge Beason

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved April 2021 by the
Graduate Supervisory Committee:

Rajnish Mehra, Co-Chair
Sunil Wahal, Co-Chair
Seth Pruitt
David Schreindorfer

ARIZONA STATE UNIVERSITY

May 2021

ABSTRACT

This dissertation consists of three essays studying topics in financial economics through the lens of quantitative models. In particular, I provide three examples of the effective use of data in the disciplining of financial economics models. In the first essay, I provide evidence of a significant transitory component of aggregate equity payout. Leading asset pricing models assume exogenous dividend growth processes which are inconsistent with this fact. I find that imposing market clearing for consumption and income in these models induces the relevant behaviors in dividend growth, even when dividend growth is obtained indirectly. In the second essay, I provide a novel decomposition of the unconditional equity risk premium. In the data, the majority of the equity premium is attributable to moderate left tail risks, not those associated with disaster states. In stark contrast to the data, leading asset pricing models do not predict that this intermediate left tail region meaningfully contributes to the equity premium. The shortcomings of the models can be pinned on unreasonably low prices of risk for tail events relative to the data. In the third essay, I document a large dispersion in household allocations to risky assets conditional on age. I show that while standard household portfolio choice models can be made to match the average risky share over the lifecycle, the models fall short of generating sufficient heterogeneity in the cross-section of household portfolios.

To Ashley, whose unwavering love, support, and encouragement made this possible.

ACKNOWLEDGMENTS

I am grateful to Rajnish Mehra, Sunil Wahal, Seth Pruitt, and David Schreindorfer for their continuous encouragement, inspiration, and thoughtful feedback. The mentorship and guidance that I received has been world-class and invaluable, and I will remain eternally indebted to them for their service. I also owe many thanks to the finance faculty and my fellow PhD colleagues for many enlightening discussions, for a good deal of laughter, and for contributing to what ultimately became a greatly rewarding experience for me.

Additionally, I would like to acknowledge the generous financial and non-financial support provided by the ASU Finance Department throughout my period of study and the W. P. Carey School of Business for what was, in retrospect, an unhealthy amount of free coffee.

Lastly, I am thankful for my family and friends, especially my parents, my wife Ashley, and my son Emmett, for always reminding me to look at the bigger picture. The pursuit of happiness is just as important, and demanding, as the pursuit of this degree.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Figure Page

Chapter 1

INTRODUCTION

Economic models allow economists to de-clutter an untidy world. Quantitative models are those that are further disciplined by data. In financial economics, quantitative models abound. The use of quantitative models in financial economics is a two-way street. While models deliver new insights about the world, data informs about correct model mechanisms. This dissertation comprises three examples of such a process at work. The first two studies examine aggregate asset pricing models, which are used in the study of the relative pricing of different asset classes, such as stocks and bonds. The third examines models used in the study of household portfolio choice decisions. In this introductory chapter, I highlight the main findings and contributions of each essay.

In the first essay (Chapter 2), I provide evidence that long term reversals are a robust property of dividend growth rates at the aggregate and portfolio levels, suggesting the robust presence of a transitory component in equity payout. Leading asset pricing models assume exogenous dividend growth rate processes which are incompatible with this stylized fact. I show that by imposing oft-ignored equilibrium market clearing conditions in these models, which require aggregate consumption to equal the sum of dividends and labor income, and asking the models to match stylized moments of labor income growth rates, the models inherit dividend growth rates which display the behaviors observed in the data. This occurs even though dividend growth is obtained only *indirectly* via the market clearing condition.

One characteristic of series with long term reversals is that they display downward sloping term structures of annualized volatility. An interpretation of this is that they are less risky in the long run, in the sense that there is less uncertainty surrounding long horizon forecasts of the series than there otherwise would be if they did not have long term reversals. Changing the dynamic behaviors of dividend growth in a model naturally has the potential to alter the prices it generates for claims to that stream of dividends. The extent of the changes is not obvious ex ante because it will depend on the pricing mechanism in each model and which asset prices are in question. I explore the asset pricing implications of my modified dividend process to show that these endowment features are quantitatively important.

In Chapter 3 (co-authored with David Schreindorfer), I provide a novel decomposition of the equity risk premium. Our decomposition presents a stronger and more direct test for theories seeking to explain the equity premium puzzle. While many theories can match the level of the unconditional equity premium, our decomposition informs about the average distribution of the equity premium across different states of the world.

We use option prices and realized returns to decompose risk premia into different parts of the return state space. In the data, 8/10 of the average equity premium is attributable to monthly returns below -10%, but returns below -30% matter very little. In contrast, leading asset pricing models based on habits, long-run risks, and rare disasters attribute the premium almost exclusively to returns above -10%, or to the extreme left tail.

A further decomposition of the importance of this region into what we call quantity of risk and price of risk is particularly informative. The intuition for the split is simple: a return state may have a very large price of risk, yet if it occurs once every one million years it will likely not contribute meaningfully to the equity premium.

When we conduct an exercise that controls for differences in the quantity of risk across models, all of the models we consider have virtually identical implications for our equity premium decomposition. We conclude that the discrepancy between model and data arises primarily from an unrealistically small price of risk for stock market tail events.

In the third study (Chapter 4), I explore the intersection of household heterogeneity and household portfolio choice. Using the Survey of Consumer Finances, I document a large dispersion in household allocations to risky assets conditional on age. I show that while standard household portfolio choice models can be made to match the average risky share over the lifecycle, the models fall short of generating sufficient heterogeneity in the cross-section of household portfolios. This can occur despite substantial heterogeneity in household characteristics, such as labor income risk. The evidence suggests standard household portfolio choice models are still missing important determinants of household portfolio choice.

The continual benchmarking of models to data will always be a worthwhile exercise. When model and data differ, there is room to learn. This dissertation applies thoughtful introspection to two important classes of problems in financial economics: the pricing of financial assets and the determinants of portfolio choice. It is my hope that my essays will act as stepping stones to additional interesting research in these areas.

Chapter 2

CASH FLOWS IN EQUILIBRIUM ASSET PRICING MODELS

2.1 Introduction

Finance theory suggests an asset's price is determined by the joint properties of its cash flows and the pricing kernel. Naturally, an asset pricing model requires a correct description of both to be relevant. Modern consumption-based asset pricing models have had success at matching high and time-varying equity risk premia with low and stable real risk free rates.[1] However, in the process of doing so, they assume away the largest component of variation in dividend growth and ignore conditions that ensure the existence of a long run economic equilibrium in the model. I show that the successes of these models are potentially sensitive to these modeling decisions.

I start by providing a breadth of evidence pointing to a strong transitory component in equity payout. Many aggregate equity payout growth rate series display features consistent with long run reversals. That is, if dividend growth was high this year, we might expect it to be lower in the future. Because periods of high growth tend to be followed by low growth (and vice versa), these effects offset as we look at longer and longer horizons. As a result, the annualized variance of these series tends to decline over time. In most cases, the twenty year annualized variance is less than half of its one year counterpart. Cochrane (1988), Cochrane and Sbordone (1988), and Lo and

[1] Examples include the external habit model (Campbell and Cochrane 1999; Bekaert and Engstrom 2017), the long run risks model (Bansal and Yaron 2004; Bollerslev, Tauchen, and Zhou 2009; Drechsler and Yaron 2011), and the rare disasters model (Rietz 1988; Barro 2006, 2009; Wachter 2013).

MacKinlay (1988) show that these variance ratios can be used to identify the relative importance of transitory shocks in a series.[2]

As further evidence, I examine aggregate equity payout growth during and following recessions. If there was no transitory component to the series, growth immediately following recessions would be no different than growth in the midst of a mature economic expansion. Instead, I find evidence suggesting that growth following recessions is higher than average, again consistent with reversals.[3]

To investigate the pervasiveness of this time series property, I turn to portfolio-level evidence. To the extent that the transitory dynamics feature prominently at the aggregate level, they would be expected to exist at somewhat dis-aggregated levels as well. I form portfolios based on size, book-to-market, momentum, market beta, and industry. In each case, portfolio-level payout behaviors align with that of the aggregate market. In combination with the aggregate level evidence, these findings suggest that long run reversals are a robust feature of equity payout growth.

In contrast, leading endowment economy asset pricing models assume dividends are only subject to permanent shocks. Without the transitory component that gives rise to long run reversals, longer horizon growth rates do not enjoy the offsetting effects of periods of high and low growth, and the term structure of annualized variance is not downward sloping. A further consequence of assuming away transitory components of dividends is that it assumes away any long run connection between dividends and consumption. Put differently, consumption and dividends cease to be cointegrated. In economic terms, the economy lacks balanced growth.

[2]Beeler and Campbell (2012), Belo, Collin-Dufresne, and Goldstein (2015), and Marfè (2017), among others, also use variance ratios to study the dynamic properties of dividend growth.

[3]Beaudry and Koop (1993), Sichel (1994), Kim and Nelson (1999), Kim, Morley, and Piger (2005), and Morley and Piger (2012) document and model bounce back growth in US output.

I find that making the models jointly consistent with empirical evidence for the growth rates of consumption, dividends, *and labor income* in such a way that is theoretically consistent mostly addresses the shortcomings. Implicit in endowment economy models is an equilibrium market clearing condition, which states that aggregate consumption is the sum of aggregate dividends and aggregate labor income.[4] I show that modeling consumption and labor income growth rates directly, when calibrated to match the relevant moments of those series, generate *implied* dividend growth rates, obtained *indirectly* through the market clearing condition, that are able to match these stylized facts about dividend growth. The approach is largely agnostic to the assumed consumption growth process and is independent of preferences, making it applicable to a wide range of models.

My paper adds to the literature modeling dividend dynamics in new and alternative ways. Belo, Collin-Dufresne, and Goldstein (2015) introduce financial leverage into the endowment economy and model dividends as the residual cash flows after payments to bondholders have been met. Hasler and Marfe (2016) model consumption and dividends jointly as a system having permanent and transitory components. Marfè (2017) incorporates incomplete markets and models dividends as resulting from equilibrium risk-sharing between workers and firm owners. As I do, these papers target both the short and long horizon properties of dividend growth. I show that no additional economic or statistical restrictions are needed to achieve this behavior beyond what was already present in the equilibrium conditions of the original endowment economy used by existing leading asset pricing models.

[4]For example, Bansal and Yaron (2004) state that "...the agent is implicitly assumed to have access to labor income."

My paper relates to the literature using new empirical moments to discipline economic models. A prominent example of this is using moments computed across horizons, commonly referred to as a term structure.[5] Binsbergen, Brandt, and Koijen (2012) use data on dividend strips to document that the term structure of equity risk premia is downward sloping.[6] They show that leading asset pricing models based on habits, long run risks, and disasters predict the opposite. Backus, Chernov, and Zin (2014) and Dew-Becker et al. (2017) use term structures of interest rates and returns on variance swaps, respectively, to further inform about the pricing of risk across horizons. Backus, Boyarchenko, and Chernov (2018) connect the term structure of excess returns on an asset to the term structure of its cash flow risk and its cash flow co-movement with the pricing kernel.

My paper also contributes to the growing literature at the intersection of labor and finance. There are a plethora of empirical asset pricing papers showing the cross-sectional pricing implications of labor frictions among publicly-traded firms (Chen, Kacperczyk, and Ortiz-Molina 2011; Belo, Lin, and Bazdresch 2014; Favilukis and Lin 2016; Kuehn, Simutin, and Wang 2017; Donangelo 2018; Donangelo et al. 2019). I take this as additional evidence that these frictions matter for investors and are important to consider at the aggregate level.

A separate strand of the labor and finance literature focuses on heterogeneity and incomplete markets, such as Constantinides and Duffie (1996) and Schmidt (2016).

[5]Examples from other areas of finance include Giglio and Kelly (2018) who use term structures to examine violations of the law of iterated expectations and Conrad and Wahal (2020) who study the term structure of liquidity provision at short horizons.

[6]This finding is not without controversy. Bansal et al. (2019) suggest the data display a downward sloping term structure due to a short sample bias, while Schulz (2016) suggests that differential tax treatment of dividends and capital gains gives rise to an unconditionally flat term structure, and Boguth et al. (2012) suggest the original results are plagued by microstructure noise. Van Binsbergen et al. (2013) and Gormsen (2018) provide additional evidence consistent with the original findings.

These papers propose *idiosyncratic* labor market risk as key drivers of asset prices, while I focus on the impacts of aggregate labor market frictions, in particular the quantity dynamics. A relatively smooth labor income series in aggregate and highly variable labor income at the individual or household level, as documented by Guvenen, Ozkan, and Song (2014), Guvenen et al. (2019), and Pruitt and Turner (2020), are not inconsistent observations. In fact, incomplete markets models often take aggregate labor income as given and specify individual risk as redistribution shocks, which indicates that my cash flow model could also be used in these contexts. I leave this connection for future work.

Finally, my paper relates to the literature studying the permanent and transitory features of the macroeconomy (Campbell and Mankiw 1987; Cochrane 1994; Aguiar and Gopinath 2007; Kim, Piger, and Startz 2007). While that literature typically focuses on output, consumption, and investment, my primary focus is on equity payout.

2.2 Long Run Reversals in Payout Growth

Figure 1 displays the term structure of volatility for different measures of aggregate equity payout growth alongside that of real per-capita non-durables and services consumption growth, for the period 1930-2019. The first equity payout series is a measure of the gross payout, defined as the sum of dividends and share repurchases, emanating from firms listed in the CRSP database. The second is just the dividend portion of the first. I compute these series following Bansal, Dittmar, and Lundblad (2005). The next 2 series are those constructed by Larrain and Yogo (2008), which I've extended using data through 2019. The first of the Larrain and Yogo (2008) series is net dividends of the US nonfinancial corporate business sector. The second of the Larrain and Yogo (2008) series captures net equity payout of the US nonfinancial corporate business sector. The final equity payout measure is per capita personal dividend income from the BEA National Income and Product Accounts (NIPA), which is a measure of dividends *received* by households rather than *paid* by firms. Each series is time aggregated to the annual frequency to avoid well-known seasonalities in equity payout is converted to real using the Consumer Price Index (CPI). See the Appendix for more detail on the construction of the sample. Annualized volatility is defined here as $\sqrt{\frac{1}{k}\text{Var}(x_{t+k} - x_t)}$ where k is the horizon and x_t is the log payout or consumption series.[7]

Looking at the left edge of this figure, it clear that the different measures of equity payout have very different annual volatilities. These differences have caused confusion as to what the appropriate measures of aggregate equity payout might be. However, only considering the one year volatility misses two important dimensions along which there is much more agreement amongst the series. First, the right edge of the plot (the annualized 20 year volatility) displays much less disagreement among payout series than the one year volatility. Second, there is near uniform agreement in the slope and shape of the term structure as one moves from left to right. That is, the lines are downward sloping, and more steeply in the beginning than in the end. This finding suggests that the differences between various aggregate payout series might be smaller than previously thought.

[7] I compute the variances in the same manner as Cochrane and Sbordone (1988),

$$\frac{1}{k}\widehat{\text{Var}}(x_{t+k} - x_t) = \frac{T}{k(T-k)(T-k+1)} \sum_{t=k}^{T} \left(x_t - x_{t-k} - \frac{k}{T}(x_T - x_0) \right)^2,$$

where T is the length of the growth rate time series. The scalar out front implements a small sample bias correction.

A convenient normalization allows for separate consideration of the level and slope of the volatility term structure. Let $VR_k(x_t) = \frac{\frac{1}{k}\mathrm{Var}(x_{t+k}-x_t)}{\mathrm{Var}(x_{t+1}-x_t)}$ be the variance ratio at horizon k for the series x_t. The behavior of the variance ratio as k varies captures all of the same information about the dynamics of the time series that are embedded in the volatility term structure. However, the variance ratio is also now a unitless quantity which is comparable across time series because it abstracts from differences in the level of volatility. Figure 2 plots the variance ratios which are analogous to the volatility term structures in Figure 1. The similarity between the variance ratios of the payout series more clearly suggests that the differences among these series relate almost exclusively to differences in the level of volatility. The dynamic behaviors of each series are strikingly similar.

A corollary of this finding is that variances beyond the one year horizon matter if one wants to match the observed dynamics of payout growth. Again here, variance ratios provide convenient insight. Lo and MacKinlay (1988) and Cochrane (1988) show that the variance ratio is a weighted sum of autocorrelations through order k,

$$VR_k = 1 + 2\sum_{j=1}^{k-1} \frac{k-j}{k}\rho_j, \tag{2.1}$$

where ρ_j is the j-th autocorrelation coefficient. Therefore, variance ratios, and, by extension, term structures of volatility, are directly related to the autocorrelation function. Table 2 displays the results of regression the CRSP gross payout growth measure on its lags. When using only one lag, the slope coefficient is small and positive, marginally significant, and the adjusted R^2 is 1.7%. The picture changes dramatically if more lags are included in the regression. When including six lags, all coefficients become negative, all but one are significant, and the adjusted R^2 rises to 17%. These differences explain the hump shaped variance ratios seen in Figure 2; dividend growth has positive low order and negative higher order autocorrelations.

Importantly, the regressions suggest that the explanatory power of the higher order terms cannot simply be ignored if one is concerned with appropriately modeling the dynamics of payout.

And finance researchers *should* be concerned with appropriately modeling the dynamics of payout. Two examples and a theory underscore this point. First, the variance ratios of *returns*, and answers to the question of whether returns exhibit long run reversals, are intimately tied to the variance ratios of payout growth. This is clear from looking at the annualized variance of (log) returns at horizon k:

$$\frac{1}{k}\text{Var}(r_{t,t+k}) = \frac{1}{k}\text{Var}\left(\sum_{j=0}^{k-1}\ln\frac{PD_{t+j+1}+1}{PD_{t+j}}\right) + \frac{1}{k}\text{Var}(d_{t+k}-d_t) + \frac{2}{k}\text{Cov}\left(\sum_{j=0}^{k-1}\ln\frac{PD_{t+j+1}+1}{PD_{t+j}}, d_{t+k}-d_t\right). \quad (2.2)$$

The connection to variance ratios can be made more explicit by substituting in using the definition of variance ratio:

$$VR_k(r_t)\text{Var}(r_{t,t+1}) = \frac{1}{k}\text{Var}\left(\sum_{j=0}^{k-1}\ln\frac{PD_{t+j+1}+1}{PD_{t+j}}\right) + VR_k(d_t)\text{Var}(d_{t+1}-d_t) + \frac{2}{k}\text{Cov}\left(\sum_{j=0}^{k-1}\ln\frac{PD_{t+j+1}+1}{PD_{t+j}}, d_{t+k}-d_t\right). \quad (2.3)$$

Clearly, the properties of long horizon payout growth are instrumental in determining the properties of long horizon returns. Matching variance ratios for returns while mismatching variance ratios for dividend growth necessarily implies a mischaracterization of one or both of the terms on the right hand side of the equation involving the evolution of price dividend ratios. The second example comes from the nascent literature on the term structure of the equity risk premium. Backus, Boyarchenko, and Chernov (2018) and Bansal et al. (2019) show that the log excess holding period return on zero coupon equity claims is $rx_{t+k} = d_{t+k} - d_t + k\,e_{k,t} - y^r_{k,t}$ where $e_{k,t}$ is the equity yield for horizon k and $y^r_{k,t}$ is the real rate of the same maturity. The future dividend is the only stochastic piece remaining in the equation, so the conditional variance and Sharpe ratio of zero coupon equity depend only on the conditional term structures of the expectation and volatility of dividend growth. Belo, Collin-Dufresne, and Goldstein (2015) were among the first to provide evidence that producing downward sloping term structures of dividend growth volatility in the external habit and long run risks models allows the models to generate downward sloping term structures of equity risk premia. Finally, asset prices are forward looking. The price of an equity claim is the expected discounted present value of its cash flows into the infinite future. By definition, long horizon dividend dynamics matter for asset prices.

An issue with using annual macroeconomic data is that the number of observations is not large (90, in this case). To gain a sense for the statistical precision of these estimates, Figure 3 plots $VR_k(d_t)$ and its standard error bands for the CRSP gross payout series. The standard errors are computed using the asymptotic Bartlett formula and are scaled by the estimated 30 year variance (or variance ratio, in this case) of the series.[8] The solid horizontal line, equal to one, displays what the variance profile would look like if the series was a random walk. While the error bands do widen with horizon, we can confidently state that for horizons 5 years or longer, the variance ratio is significantly lower than what would be predicted under random walk dynamics.[9] At

[8]Concretely, the standard errors are computed as $\sqrt{\frac{4k}{3T}} \times VR_{30}(d_t)$. See Cochrane (1988) and Cochrane and Sbordone (1988) for more detail.

[9]However, as noted by Cochrane (1988), the standard errors *under the null* of a random walk are much larger. So, while we may reject the hypothesis $\text{Var}(x_{t+1} - x_t) = \frac{1}{k}\text{Var}(x_{t+k} - x_t)$ based on the standard errors of the long horizon variances, we often cannot reject the possibility that the series contains a unit root component. This is not the focus of the paper, so I do not report the unit root tests.

the 20 year horizon, the standard error bands imply a variance ratio of 0.14 to 0.44, consistent with the range of estimates given by the other payout series in Figure 2.

Of course, one can always request a model to match more features of the data. The marginal benefit of matching additional empirical moments declines quickly and is not always clearly positive ex ante. Additionally, the term structure of volatility of dividend growth (or the variance ratio term structure) is not a single empirical moment, but a possibly infinite set of moments. In practice, requiring a model to match the entire term structure may be too cumbersome or restrictive. In the next subsection, I will cast the problem in a slightly different framework. Within this framework, I will show that long horizon variance ratios take on additional economic meaning, which reinforces their importance for describing the relevant dynamics of payout growth. In particular, the important statistic will be a *single* long horizon variance (or variance ratio) rather than a sequence of them, which greatly reduces the burden of matching model to data.

2.2.1 The Transitory Component in Equity Payout

Under quite general conditions, non-stationary time series can be decomposed into permanent and transitory components (Beveridge and Nelson 1981). Let x_t be a general non-stationary time series expressed in logs. We can find a decomposition of the form

$$x_t = \tau_t + z_t \tag{2.4}$$

where τ_t is the trend component, which has a unit root, and z_t is the transitory, or cyclical, component. The trend component also carries the title of permanent component because innovations in it never revert – they permanently raise or lower expectations of future x. Innovations in the transitory component do not impact long run expectations of x. Put differently, fluctuations in the transitory component by definition die out as one considers longer and longer horizons.

Naturally, the relative importance of the permanent and transitory components dictates the dynamic behavior of x_t. If the permanent component is much more volatile than the transitory component, x_t will inherit relatively more properties from its stochastic trend. Because the permanent component has a unit root and the transitory component does not, which series dominates is key to understanding the risk profile and dynamics of x_t.

One simple way to estimate the relative importance of these two components is by examining variance ratios (Cochrane 1988; Cochrane and Sbordone 1988; Lo and MacKinlay 1988; Cochrane 1994). Intuitively, the contribution of the transitory component to the variance of long horizon growth rates shrinks to 0 as the horizon becomes infinite. Thus, the variance of long horizon growth rates of x_t serves as an estimate of the variance of its permanent component, τ_t.

Recall that Figure 2 displays the univariate variance ratios for horizons of 1 year to 20 years for different measures of aggregate equity payout growth alongside the variance ratios of real per-capita non-durables and services consumption growth, for the period 1930-2019. The annual variance and the annualized 20 and 30 year variances of each series, the inputs to the variance ratios, are tabulated in Table 3. For all measures of equity payout, the variance ratios decrease with horizon to levels well below 1. In fact, all are below 1/2. To the extent 20 or 30 years is sufficiently long enough to nearly eliminate the effects of the transitory component, the variance ratio at 20 or 30 years is an estimate of the relative importance of the permanent component. Thus, for these series, the permanent component appears to be the minority component, with the bulk of the short term variation coming from the transitory component.

In contrast, the variance ratios of consumption growth are slightly above one. At the 20 year horizon, a variance ratio of 1.2 is about what we'd expect to find from a pure random walk.[10] Whereas the variance ratios for equity payout growth suggest a large transitory component in the series, the variance ratios for consumption growth predict a very small role for any transitory component. Indeed, previous papers have used this as a basis for employing consumption as a proxy for the stochastic trend component in output (Cochrane and Sbordone 1988; Fama 1992; Cochrane 1994), although Kim, Piger, and Startz (2007) find evidence of economically relevant transitory fluctuations in consumption during recessions.

One may wonder if these results are driven by the pre-WWII years, which include the Great Depression, a noticeably volatile period. Figure 2 shows the variance ratios for the same series for the years 1948-2019. Barring one exception, the conclusions are, if anything, stronger. The 20 year variance ratios for equity payout growth rates fall to approximately 1/4. The lone exception is the cash dividend series from CRSP, which has variance ratios above 1 at all horizons during this post-war period. This is likely the result of the well-known tendency for public firms to smooth cash dividends during this period, which introduces positive autocorrelation in the series. Starting in the 1980s, firms began shifting away from dividend payouts towards share repurchases, and continued to do so at a rapidly increasing rate. Therefore, the measures that exclude repurchases ignore the bulk of equity payout.

The term structure of volatility for consumption growth becomes steeply upward

[10] Working (1960) shows that time aggregation induces positive autocorrelation in the series. Thus, a random walk which is time aggregated would display a variance ratio slightly above one.

sloping in the post-war period, with the variance ratios rising to nearly 3.5 at the 20 year horizon. Interestingly, while the long sample variance ratios suggest consumption growth is essentially a random walk, the variance ratios in the post-war sample suggest a positively autocorrelated component embedded in the stochastic trend (time-varying drift), similar to what can be found in the long run risks literature (Bansal and Yaron 2004).[11] As a result, the contrast between the behavior of consumption growth and the behavior of equity payout growth becomes even more severe in this sample.

Cochrane and Sbordone (1988) show that a second viable approach to estimating the variance of the permanent component of a series is to use the variance of long horizon growth rates of a second time series which is cointegrated with the series of interest but is closer to a random walk. Building on the fact the consumption is nearly a random walk in the long sample, and borrowing from the literature that documents cointegration between consumption and dividends (Bansal, Dittmar, and Lundblad 2005; Hansen, Heaton, and Li 2008; Bansal, Dittmar, and Kiku 2009), I also implement this second approach. Precisely, this alternative variance ratio is $1/k$ times the variance of k differences of log consumption divided by the one period variance of the equity payout growth series. The outermost columns of Table 3 show that this alternative variance ratio is always lower than the univariate variance ratios. Therefore, it is possible that the transitory component of equity payout is even larger than univariate variance ratios suggest.[12]

[11] It is worth pointing out that it is also possible for the permanent and transitory shocks to be correlated in such a way as to produce hump shaped volatility term structures. Therefore, we still cannot rule out the possibility of a transitory component in consumption growth.

[12] Cochrane and Sbordone (1988) also suggest a third approach to estimating the variance of the permanent component, which is to use the covariance between long horizon growth rates of the two series. For the preferred equity payout series, CRSP dividends plus repurchases, this approach produces an estimate of 0.000734 at 30 years, which is similar to the estimate provided by consumption. They suggest that this covariance approach suffers from a downward bias that

2.2.2 Bounce-back Growth

Variance ratios are not the only way to identify the presence of a strong transitory component in a series. Kim, Morley, and Piger (2005) show that the behavior of output growth rates following recessions also indicates the presence of a transitory component in postwar output. They find that, on average, output growth in the first quarter following the end of a recession is more than 2 times above its mean. The growth rate slowly decays back to its long run mean over a period of 8 quarters. The top panel of Figure 5 replicates their evidence on the 1948-2019 sample. Rather than quarterly growth rates, I use a 4-quarter moving average of quarterly growth rates in order to be consistent with the payout series in the lower panel, which require the use of a moving average to remove some of the strong seasonality effects. The main effect is a simple phase shift to the right 4 quarters. Thus, while time 0 marks the first quarter of the expansion, time 3 marks the first quarter that is free from any recession periods in the trailing average, and is denoted by the solid black vertical line. Clearly, GDP growth peaks immediately following recessions and slowly declines back to its average value. This consistent pattern of high growth following recessions (times of low growth), what they refer to as "bounce back" growth, is compelling evidence in favor of a transitory component in output. If GDP was a random walk, growth following recessions would simply be expected to be average. The top panel of Figure 3 also displays the result of this exercise for consumption growth. Consistent with consumption growth being nearly a random walk, there is no apparent bounce back growth following recessions.

---

cannot be corrected without knowing the true correlation structure of the data, so I do not explore it further.

The bottom panel of Figure 5 displays the same estimates for the preferred equity payout growth measure, CRSP dividends plus repurchases. For the purpose of constructing this figure, I do not time aggregate the data to the annual horizon, leaving the series as a higher frequency smoothed quarterly growth rate series.[13] Bounce back growth emerges after a short delay. Whereas the bounce back growth in output starts immediately following the end of a recession (time 3 because of the phase shift), the bounce back growth in equity payout does not begin for an additional two quarters (time 5 due to the phase shift) following a recession. However, the effect is large, with average growth rates near 20% on an annualized basis – more than 4 times the average. These findings are consistent with the evidence presented in Hasler and Marfe (2016) of significant economic recoveries, especially in equity payout, following disasters. This bounce back growth is consistent with a strong transitory component in equity payout.

2.2.3 Cross-sectional Evidence

To further assess the pervasiveness of a strong transitory component in equity payout, and to better understand it, I construct equity payout for portfolios of firms sorted on commonly used characteristics: size, book-to-market, momentum, beta, and industry. I use 5 portfolios for each sort, for a total of 25 portfolios. For each portfolio, I construct both the cash dividend and cash dividend plus repurchases payout series following Bansal, Dittmar, and Lundblad (2005), and I time aggregate to the annual

[13] As in Bansal, Dittmar, and Lundblad (2005), a four quarter moving average is applied before computing the quarterly log growth rates. The difference between this approach and just taking a moving average of standard growth rates is generally negligible, so I use a standard moving average for the GDP and consumption growth series in the top panel. The spikiness in the lower panel is the result of the imperfect removal of the seasonality in the series.

frequency. Because the construction of the sorting variables is standard, I relegate those details to the appendix. The portfolios are rebalanced at the end of June in each year, I use NYSE breakpoints, and the data runs from 1930-2019.

Figures 4 through 8 plot the variance ratios for both payout series, for each portfolio, for horizons ranging from 1 to 20 years. The patterns seen in the aggregate data continue to shine through. For the total payout measure, found in the top panel in each figure, all 25 portfolios have 20 year variance ratios below 1. Most of the portfolios have 20 year variance ratios that agree with the values seen in Figure 1 (below 1/2), implying a large role for the transitory component in payout. The variance ratios of the strictly dividend growth series, shown in the bottom panel in each figure, largely corroborate the evidence. Just two portfolios have 20 year variance ratios exceeding 1.

The lack of much variation across portfolios in their variance ratio plots suggests that the permanent-transitory split is ubiquitous. In combination with the aggregate level evidence, this cross-sectional evidence suggests that transitory variation is the primary driver of short run fluctuations in payout.

2.2.4 Summary of Evidence, Robustness, & Discussion

The body of evidence favors an economically large and meaningful transitory component in equity payout. While I cannot conclusively prove that it exists, I find plenty of evidence in support of it. Equity payout displays many of the same features that have been used to argue in favor of a transitory component in output. It is marked by long run reversals, characterized by downward sloping term structures of volatility. Five different measures of aggregate equity payout provide coherent perspectives. Twenty-five dis-aggregated measures of equity payout at the portfolio level are in agreement with the aggregate findings. The permanent component of equity payout is estimated to have an annualized volatility in the 6 to 10 percent range when measured using payout alone or less than 3 percent when measured by consumption, both being less than the 10 to 20 percent annual volatility of the payout series.

Additional robustness checks tell a largely consistent story of a strong transitory component. The majority of the results presented above used annual data from the 1930-2019 or 1948-2019 periods. The results are similar if I use overlapping annual observations constructed from quarterly data for the period 1948-2019. There are two potential structural breaks in the early 1980s: the widely-documented break in macroeconomic volatility and the rapid increase in share repurchase activity. If I consider only the 1984-2019 period, the results are consistent, although this is an admittedly short sample to be thinking about long run dynamics. If I stop the sample in 2007 in order to exclude the Great Recession, the results are unchanged. All of the results presented thus far focus only on equity payout. Larrain and Yogo (2008) construct a series of payout that includes payments to debtholders; my results continue to hold for this broader measure of payout.

My results reveal that the chasm between the dynamics of consumption and equity payout is large. Consumption is quite possibly nearly a random walk, and any transitory dynamics it has are drowned out completely by fluctuations in the permanent component. Perhaps somewhat surprisingly, the dichotomy between consumption and equity payout is mostly absent from consumption-based asset pricing. While the literature has long recognized the contrasting levels of annual volatility between the two series, the differential treatment mostly stops there. As I discussed in the preceding sections, the term structure of volatility contains a lot of information, beyond simply its level, regarding the dynamics of a series. As was the case for equity payout, differing levels of one period (or short-term) volatility need not imply different dynamics. However, this appears not to be the case for consumption and equity payout. Equity payout growth is an order of magnitude more volatile than consumption growth *and* has vastly different dynamics – the latter point having been essentially cast aside.

There are a number of ways to model dividends in a Lucas (1978) endowment economy, with the most common being directly specifying an exogenous stochastic process for log dividend growth, Δd. The approach often involves specifying the same "type" of stochastic process for dividend growth as that of consumption growth, but with potentially different parameters.[14] For example, if consumption is modeled as a random walk with drift, then dividends are modeled as a random walk with drift, but the drift and the volatility parameters are allowed to differ from those for consumption. Additionally, the shocks may be allowed to have imperfect correlation. By allowing the volatility parameters to differ, one can simultaneously match the relatively low volatility of consumption growth and the relatively high volatility of dividend growth. Indeed, this is the endowment specification in the Campbell and Cochrane (1999) external habit model, one of the leading examples of modern consumption-based asset pricing. Unfortunately, this dividend growth specification is wholly inconsistent the evidence presented in the preceding sections. Campbell and Cochrane are not alone;

[14]Abel (1990) suggested modeling dividend growth as a levered claim to consumption, $d_t = \lambda c_t$ with $\lambda > 1$. Cecchetti, Lam, and Mark (1993) noted that the Abel (1990) was quite restrictive in that it always predicted perfect correlation and that it scaled both the mean and volatility of consumption growth by λ. They relaxed these assumptions, allowing for imperfect correlation and differential scaling on the mean and volatility.

the great majority of the consumption-based asset pricing literature has coalesced around this approach to modeling dividend growth.

In the setup just described, there are two shortcomings. The first is a near impossible tension between matching the short and long run dynamics of payout simultaneously. This shortcoming is not obvious in the existing literature because the focus has been entirely on matching only the short run dynamics of payout, such as the annual volatility and first order autocorrelation. In truth, virtually any stochastic process can be calibrated to replicate these moments within a reasonable tolerance. Therefore, they are not very informative about the actual dynamics of payout. Cochrane (1988) makes a similar observation regarding models of GNP. The outcome is that models have matched the short run dynamics, but virtually all of the leading consumption-based asset pricing models predict upward sloping term structures of payout volatility, which means they overstate the degree of long run variation. Ultimately, all else equal, this implies that they overstate the riskiness of the infinite cash flow stream. To concretely examine this practice, in Figure 11 I present the variance ratios for two leading asset pricing models: the Campbell and Cochrane (1999) external habit model and the Bansal and Yaron (2004) long run risks model. As is evident, both models do a poor job matching the variability of payout as seen in the data beyond short horizons. By ignoring the transitory component in payout, these models counterfactually assign all variation to that of the permanent component.

The second shortcoming is a lack of cointegration between consumption and payout. Modeled in the manner described above, both consumption and payout contain a unit root. However, the restriction that they have *identical* unit roots is not imposed. Therefore, $\Delta d - \Delta c$ also contains a unit root and is non-stationary. The implication is that equity payout as a share of consumption will tend to zero or infinity. That divergence is a limiting result, the thing that occurs as the horizon becomes infinite. However, that does not mean it can safely be ignored. Because payout is much more volatile than consumption, this divergence is relevant even in short samples. To document this, I conduct a simple Monte Carlo simulation exercise. I assume that monthly log consumption and payout growth rates are drawn in an IID fashion from a bivariate Normal distribution with $\mu_c = \mu_d = 0.02/12$, $\sigma_c = 0.02/\sqrt{12}$, $\sigma_c = 0.15/\sqrt{12}$, and $\rho_{cd} = 0.2$. I simulate 100,000 samples of 240 months, and I initialize the level of payout to 10% of consumption ($\frac{D}{C} = 0.1$). After only 10 years, the payout-consumption ratio ranges from 0.01 to 0.66. After only 20 years, the payout-consumption ratio ranges from 0 to 1.5 – every dollar of consumption is financed by 1.5 dollars of payout. This range will only continue to widen as the sample length increases. Thus, lack of cointegration in this scenario is economically meaningful even in a sample of only 70 or 90 years.

These two shortcomings are mostly driven a single source – model misspecification error. Matching the level of volatility while mismatching the term structure of volatility implies misspecification of the transitory component in payout, which necessarily implies misspecification of the permanent component as well. For consumption and payout to be cointegrated, they must share a common permanent component.

It deserves to be noted that a second less-common approach to modeling dividends, in which the log dividend-consumption ratio is modeled with a stationary process, does not suffer from lack of cointegration.[15] This approach has the advantage of breaking the strong tension between simultaneously matching the dynamics of consumption

[15] Longstaff and Piazzesi (2004), Bekaert, Engstrom, and Xing (2009), and Marfè (2016) are some examples of papers that utilize this approach.

and dividends jointly. However, it is still common practice to focus almost entirely on matching the short run dynamics of dividend growth. Hasler and Marfe (2016) and Marfè (2017) are two exceptions; both papers introduce consumption and dividends as having permanent and transitory components, and they target the variance ratios for both series. My results suggest that their approaches are promising avenues for additional exploration.

Appropriately capturing the permanent-transitory decomposition of payout appears absolutely critical to delivering the relevant dynamics of payout growth in the short run and the long run. With few exceptions, the literature has focused only on the former. The practical guidance I can deliver is to include a long horizon annualized volatility or variance ratio, on the order of 20 or 30 years or more, in the set of targeted moments, in addition to the one year volatility and autocorrelation. While matching those three moments does not pin down intermediate dynamics at business cycle frequencies, eg. it might miss the hump shape seen in most series' variance ratios in Figure 2, it will allow the model to capture the bulk of the information embedded in the volatility term structure. In the next section, I develop an economically-motivated way to model dividend growth in endowment economy asset pricing models designed to reproduce these stylized facts.

2.3 A New Cash Flow Model

The economic environment underlying most consumption-based asset pricing models is a simple representative agent endowment economy, such as the one described in detail in Campbell (1996) and Lettau and Ludvigson (2001). The same setup will be employed here.

Time is discrete. There is a representative agent whose consumption bundle in each period, C_t, is composed of period labor income, L_t, and dividends from financial asset holdings, D_t.[16] Labor income should be thought of as representing all payments made for provision of human capital, H_t.[17] All wealth, including human capital, is tradable, so that aggregate wealth, W_t, is human capital plus financial asset holdings, A_t. The share of aggregate wealth composed of financial assets is A_t/W_t. The gross return on aggregate wealth is defined

$$R^w_{t+1} = \frac{A_t}{W_t} R^a_{t+1} + (1 - \frac{A_t}{W_t}) R^h_{t+1} \tag{2.5}$$

where R^a_{t+1} and R^h_{t+1} are the gross returns on financial assets and human capital, respectively.

The representative agent is endowed with preferences over his consumption stream, $U(c)$. Consumption growth, Δc_{t+1}, is specified exogenously, assumed to be the outcome of optimizing behavior on the behalf of the representative agent.[18] With these two components, the equilibrium stochastic discount factor (M_{t+1}), consumption-wealth ratio ($\frac{C_t}{W_t}$), and the return on aggregate wealth (R^W_{t+1}) become computable objects. These are the model primitives upon which asset pricing models should be assessed, as they determine the pricing of all assets in the economy.

Market clearing requires that $C_t = L_t + D_t$. As demonstrated in the previous section, when consumption and dividends are not cointegrated, the two series can drift

[16] Aggregate savings (net borrowing or lending) is assumed to be zero.

[17] As Lettau and Ludvigson (2001) note, there are multiple ways to model the connection between labor income and returns to human capital. Campbell (1996) and Jagannathan and Wang (1996) model labor income as the dividend paid by human wealth, while Lettau and Ludvigson (2001) model labor income as the dividend and the capital gains from human wealth. I do not believe this distinction is important for my purposes.

[18] Lowercase letters will denote variables under the natural log transformation, eg. $c_t = \ln C_t$. Log growth rates are denoted with Δ, eg. $\Delta c_{t+1} = c_{t+1} - c_t$.

apart. Thus, the share of consumption financed by dividends can become arbitrarily large or small. This implies that the stock market will either consume the aggregate economy or cease to exist, which of course has undesirable mirror image effects on the stock of human capital. To quote Lettau and Ludvigson (2014):

> Despite these statistical concerns, there are good economic reasons to believe that there is a common long-run relationship between c_t (consumption), a_t (aggregate wealth), and y_t (labor income): cointegration is implied by an aggregate budget constraint identity (Campbell and Mankiw 1989; Lettau and Ludvigson 2001). Just as no reasonable economic model would imply that the log price-dividend ratio is nonstationary (where stationarity follows from a Taylor approximation to the equation defining the log stock return), no reasonable model would imply that c, a, and y are not cointegrated, or equivalently that the system is characterized by three independent random walks.

Because it is not economically reasonable to have a model that predicts a large *negative* labor income stream, consumption, dividends, and labor income need to be modeled in such a way as to always respect the market clearing condition that ties them together. More succinctly, they need to be cointegrated.

With a model for Δc_{t+1} already in hand, the only models that can exist for labor income or dividends which continue to satisfy the constraint and produce economically meaningful representations of both series are ones which model the share of those variables in consumption. Clearly, because the shares must always sum to one, only one of the shares can be modeled explicitly. I propose to explicitly model the consumption share of labor income and leave the dividend share as the residual piece.

Let $S_t = \frac{L_t}{C_t}$ represent the share of consumption financed by labor income. By the aggregate resource constraint, $1 - S_t$ is the share of consumption financed by dividends. Log labor income and dividend growth rates, $\Delta \ell_{t+1}$ and Δd_{t+1} respectively, are given by

$$\Delta \ell_{t+1} = \ln \frac{C_{t+1} S_{t+1}}{C_t S_t} = \Delta c_{t+1} + \Delta s_{t+1}, \tag{2.6}$$

$$\Delta d_{t+1} = \ln \frac{C_{t+1}(1 - S_{t+1})}{C_t(1 - S_t)} = \Delta c_{t+1} + \Delta dc_{t+1} \tag{2.7}$$

where Δs_{t+1} is the change in log labor share and Δdc_{t+1} is the change in log dividend share.[19] Equations (2.6) and (2.7) are simply accounting identities which are constructed to respect the market clearing condition.

This approach is non-standard and requires explanation. First, labor income is the largest component of consumption. To the extent that we care about properly matching moments of macroeconomic quantities, labor income would appear to be a first order concern based on magnitude alone. Second, obtaining dividends from market clearing is consistent with theory in which dividends are the residual cash flow after all other financial commitments, such as labor payments, have been met. Third, I will show that proceeding in this manner generates heteroskedasticity and predictability in dividend growth rates whenever the share of labor income varies over time in a persistent fashion, properties that Ang and Liu (2007) list as critical for any serious model of dividend growth to display.

Continuing with the assumption, discussed in the previous section, that consumption is lacking a transitory component, I specify a stationary process for the labor share. I assume that the log of the labor income share of consumption follows an autoregressive process of the form

$$s_{t+1} = \bar{s}(1 - \beta_s) + \beta_s s_t + \mathcal{U}^s_{t+1} \tag{2.8}$$

[19] $\Delta dc_{t+1} = \ln(1 - S_{t+1}) - \ln(1 - S_t)$.

where $\mathcal{U}^s_{t+1}$ is the innovation to the share.[20] The parameter $\bar{s}$ represents the unconditional mean of the log labor share, while the parameter β_s controls the speed of mean reversion in the share. To link this cash flow model back to the model for the pricing kernel I place some additional structure on the innovation, specifying

$$\mathcal{U}^s_{t+1} = \beta_c \mathcal{U}^c_{t+1} + \sigma_\eta \eta_{t+1} \tag{2.9}$$

where $\mathcal{U}^c_{t+1} = \Delta c_{t+1} - \mathbb{E}_t[\Delta c_{t+1}]$ is unexpected consumption growth and η_{t+1} is an independent standard Normal innovation. The parameter β_c controls the direction and magnitude of the cyclicality of the share. When β_c is negative, the share becomes countercyclical – it rises when consumption growth unexpectedly falls. Given the assumption of only permanent shocks in consumption, β_c also controls correlation between the permanent shocks to labor income ($\mathcal{U}^c_{t+1}$) and the transitory shocks (η_{t+1}).

Because the labor share is modeled as a stationary series, labor income and dividends are cointegrated with consumption. Thus, all three series grow at the same rate on average, never drifting apart. In the long run, all three series are equally risky. As I did above for consumption and dividends, I now turn to the data on aggregate labor income to see if there exists support for these assumptions. Using the NIPA Table 2.1, I construct four plausible measures of aggregate labor income. The first, and preferred, is the after-tax labor income series used by Lettau and Ludvigson (2001). The second is the line item Compensation of Employees. It is well known that the splitting of proprietors income into labor and capital components is ambiguous. For the third measure, I allocate two-thirds of proprietors income to labor, which is in line with the basic adjustments used in the existing literature. Finally, I examine the after-tax labor income series excluding personal current transfer

[20] I have also explored using this model for the level of the labor share, which does not alter my results.

receipts, so that the effects of automatic-stabilization and welfare policies are absent. All series are per-capita and converted to real by the implicit personal consumption expenditures deflator. Descriptive statistics are shown in Table 1, variance ratios of these labor income series are on display in Figure 12, and the long horizon variance estimates given in Table 4. Consistent with the assumption of stationarity of s_t, the long horizon variances of labor income growth are essentially the same as those of consumption growth. Furthermore, the difference in the level of volatility between these two series suggests that the transitory component of labor income, modeled here as s_t, does not account for a large portion of the volatility of labor income growth at any horizon. This is consistent with Ludvigson and Lettau (2003), who also find that both consumption and labor income are mainly driven by permanent shocks. Additionally, the correlation between labor income growth and consumption growth is 0.7 or higher, consistent with fluctuations in the labor share not driving a large wedge between these series.[21]

Dividends become more exposed to consumption shocks when the labor share of consumption is countercyclical. The intuition for why that is is simple. A countercyclical labor share of consumption implies that labor income moves less than one-for-one with consumption. Because the market clearing condition states all movements in consumption must be attributed to dividends or labor income, dividends must move more than one-for-one with consumption. This has some mild empirical support. The consumption beta of dividend growth is 3.75 over the long sample, which is 3 times higher than that of labor income growth (1.28). Another source of indirect evidence for a countercyclical labor share comes from the aggregate income statement. Gross value added is a measure of output and is, as its name implies, not net of any

[21]The correlation would be 1 if the labor share was constant.

costs. Net operating surplus is a profits-like measure that takes gross value added and subtracts away consumption of fixed capital (depreciation, essentially), taxes on production, and labor payments – this last being the largest operating cost by far. Thus, comparing the term structures of volatility of gross value added and net operating surplus provides indirect evidence on the effects of labor payments. Figure 13 does just that using annual gross value added and net operating surplus for the nonfinancial corporate business sector from NIPA Table 1.14. Clearly, there is a substantial transitory component to net operating surplus that is not in gross value added. Because labor costs are the bulk of the difference in these two series, this is suggestive of a transitory component in labor income. Because the volatility of net operating surplus is much larger than gross value added, this suggests that the labor share is countercyclical. Marfè (2017) develops an incomplete markets asset pricing model founded upon the countercyclical nature of the labor share. His model is able to reproduce many asset pricing facts as well as the main features of the volatility term structures presented here for consumption, dividend, and labor income.

As a concrete example of this new approach, consider the case when log consumption growth is IID Normal as in the Campbell and Cochrane (1999) model: $\Delta c_{t+1} = \mu + \sigma \varepsilon^c_{t+1}$. In that case, Equation (2.8) is isomorphic to

$$s_{t+1} = \bar{s}(1 - \beta_s) + \beta_s s_t + \tilde{\sigma}_\eta \tilde{\eta}_{t+1}, \tag{2.10}$$

where $\tilde{\eta}_{t+1} \sim N(0,1)$, $\tilde{\sigma}_\eta = \sqrt{\beta_c^2 \sigma^2 + \sigma_\eta^2}$, and $\text{Corr}(\varepsilon^c_{t+1}, \tilde{\eta}_{t+1}) = \beta_c \sigma / \tilde{\sigma}_\eta$. This is simply a standard AR(1) process with Gaussian innovations correlated with consumption growth innovations whenever β_c is nonzero. Despite the simplicity of the model, I show in the next that it is quite capable of generating the observed properties of the shares and growth rates of labor income and, by implication, dividends.

Even though the process for the (log) labor income share of consumption is relatively simple, the process for its residual complement, the dividend share, is not.[22] Luckily, we can qualitatively describe a number of its properties without too much effort. The first is that dc_{t+1} must be mean reverting whenever s_{t+1} is. The flip side to the labor share being above its long run mean and expected to decline is the dividend share being below its long run mean and expected to grow. This embeds some predictability into Δdc_{t+1} and, by extension, Δd_{t+1}. This is the insight delivered by Santos and Veronesi (2006).

Second, a leverage effect takes place. The resource constraint implies that changes in these two shares must net to zero (in levels). Due to the relative size of these components, a given change in the labor share translates to a much larger change in the dividend share. Therefore, the volatility of Δdc_{t+1} is much higher than that of Δs_{t+1}. One way to see this is to examine a first order approximation to the definition of the log change in the dividend share,

$$\Delta dc_{t+1} \approx -\frac{\mathrm{e}^{s_t}}{1-\mathrm{e}^{s_t}}\Delta s_{t+1}. \tag{2.11}$$

The term $-\frac{\mathrm{e}^{s_t}}{1-\mathrm{e}^{s_t}}$ represents the elasticity of changes in the dividend share to changes in the labor share. It is negative, large in magnitude, and time-varying because it depends on the current level of the labor share, ranging from about -4 when the labor income share of consumption is near its lowest observed value in the data to beyond -30 when the labor income share is near its highest observed value.[23] Marfè (2017) calls this the cyclicality effect of income insurance. The time variation in this elasticity imparts heteroskedasticity into dividend growth rates.

[22] I derive relevant properties of the labor share and log changes in the labor share in Appendix A.

[23] If a loglinear approximation was used instead, this elasticity would be constant, with s_t replaced by $\bar{s}$.

This leverage effect is asymmetric. Figure 14 shows the true nonlinear elasticity between share changes, ie. not the one from the first order approximation in Eq. 2.11. Not only is the elasticity, which is always negative, increasing in magnitude with the current level of the labor share (s_t), but it is also increasing in the realized change in the labor share (Δs_{t+1}). The implication is that positive changes to the log labor share lead to additional amplification and negative changes to the log labor share lead to reduced amplification. Because the labor share is countercyclical, positive changes are associated with consumption declines on average. As a result, the volatility of dividend growth and its correlation with consumption growth should both be larger in these states of the world. Indeed, Schreindorfer (2020) finds evidence of increasing downside correlations between consumption and dividend growth rates.

2.4 Results

I now embed the cash flow model of the previous section into two leading asset pricing models to examine its quantitative performance. The habit model of Campbell and Cochrane (1999) and the long run risks model of Bansal and Yaron (2004) deliver important insights about the driving forces behind risk premia. They do so with very different mechanisms and with very different assumptions about the behavior of consumption growth, making them ideal testbeds for this exercise. In the habit model, consumption growth is IID Lognormal. In the long run risks model, consumption growth exhibits extremely persistent swings in its conditional mean and variance. Both models counterfactually assume that dividend growth behaves quite similarly to consumption growth, and, as a result, they ignore the important transitory component in payout.

For each model, I simulate the original specification and calibration from the cited paper and I simulate the modified model that replaces the original dividend growth process with the indirect model of the previous section. The consumption growth process and the preference parameters remain unchanged. The parameters for the labor share are calibrated to target moments of labor income growth rates; the chosen values are $\bar{s} = \ln(0.9)$, $\beta_s = 0.965$, $\beta_c = -0.3$, and $\sigma_\eta = 0.0037$. With β_c negative, the labor share becomes countercyclical. With $\beta_s = 0.965$, the labor share is persistent. The average labor share of consumption is set to 90%, which is about its average value over the postwar period. It implies an average dividend share of 10%, which is a few percentage points higher than its average value. This occurs because the market clearing condition does not perfectly hold in the data. To generate model simulated data, I simulate 100,000 small samples of 1,092 months (91 years) from the model before time-aggregating to the annual frequency.

Summary statistics are shown in Table 5, where I compute the statistic of interest in each small sample and take the average across samples. Both models are able to match the basic features of labor income growth rates – slightly more volatile than consumption growth, about as persistent as consumption growth, and highly correlation with consumption growth. There are three main features on display in Table 5. First, for both models, the annual volatility of dividend growth under the modified model is quite close to the original specification. This is surprising given the circumstances and it needn't have been the case. Although dividend growth is obtained indirectly via the market clearing condition, and the consumption growth process is the same, it is not mechanical that we would arrive at the same level of volatility with our new approach. Second, the modified models produce smaller first order autocorrelation coefficients for dividend growth, more in line with the data.

Third, the modified models produce slightly higher correlation between consumption and dividend growth rates. This happens because a large portion of the variation in the dividend share is ultimately driven by unexpected consumption growth. Different sources of risk could be added to enhance the model and allow it to reproduce a lower correlation.

The variance ratios for these models were given in Figure 11. The long horizon variances and variance ratios are given in Table 6. As was discussed previously, both model originally imply upward sloping term structures of volatility. The result is that both models significantly overstate the volatility of dividend growth at long horizons. For example, the original models imply 11% and 19% annualized volatility at the 30 year horizon, whereas the empirical value is likely in the range of 6% to 8%. When my cash flow process is embedded into the existing models, the resulting variance ratios exhibit an impressive fit to the data. Notably, the patterns match well even though both two models have very different assumptions about the nature of consumption risk. The modified models do a much better job matching the term structure, producing long horizon volatilities of 5% and 6%.

In addition to matching the univariate variance ratios well, the Campbell and Cochrane (1999) does a good job at matching the alternative, bivariate variance ratios, too. In the data, the bivariate consumption-dividend variance ratio was 0.023 and 0.034 at the 20 and 30 year horizons, in the model those values are 0.024 and 0.025. Beyond recognizing a good model fit, these values tell us something more. In this setting, by construction, consumption is the permanent component in dividends. The fact that these bivariate variance ratios are far below their univariate counterparts is indicative of the effects of the transitory component still being large. Therefore, the 30 year variance of dividend growth is not equal to the variance of the permanent component. Nevertheless, it still provides a wealth of information about the term structure, and approaches the variance of the permanent component as the as the importance of the transitory component shrinks. This is evident because the univariate and bivariate variance ratios for labor income growth show much more aggreeance.

The data is characterized by a strongly downward sloping term structure – something the modified models capture well. The empirically observed behavior of consumption and dividend growth rates differs markedly. This suggests that modeling the two series with the same stochastic process, as is done in the original specifications of these models, is unlikely sufficient to simultaneously match the stylized facts about the series' univariate and joint behaviors. The new cash flow model breaks this tight link, and is able to capture the dynamics of payout well. The fact that realistic dividend growth rate dynamics can be obtained via the market clearing condition in a setting with realistic behaviors of consumption and labor income growth rates is suggestive evidence that the market clearing condition, and the endowment economy that it represents, might be a better approximation to the world than commonly thought.

Having explored the implications for fundamentals, I now ask in what way my approach alters the asset pricing predictions. I solve the habit model and leave the long run risks model, as well as others, for future research. As discussed in Section 2.2, there are a multitude of reasons why we might expect the modified model to display different pricing behaviors than that of its original speicification. In particular, the original models overstate the degree of long run variability in dividend growth – they make the cash flow stream seem riskier than it is.

Note that because the preferences and consumption growth process from the original model specification remains unchanged, the pricing kernel does as well. Therefore, the behavior of interest rates in the model is the same in the existing and modified specifications. I solve the model numerically on a discretized grid for the state space (surplus consumption and the labor share), using Gaussian quadrature to numerically compute expectations.

For the habit model, the level of the equity premium is virtually unchanged, shown in Table 7. Annual log returns have a mean of 5% in the original model and 5.4% in the modified version. The modified model produces too little return volatility relative to the original specification, 10.3% compared to 15%. Because the volatility of annual dividend growth rates were nearly equal under the original and modified specifications, the discrepancy in return volatilities is linked to differences in the behavior of price-dividend ratios under the two specifications. The variance of changes in the price-dividend ratio is actually larger under the modified model, suggesting that the correlation between changes in the price dividend ratio and changes in dividend growth as much more correlated in this model. Table 7 also shows the variance ratio for log returns at the 1, 2, 5, and 10 years. Clearly, the modified model produces a much too steeply downward sloping variance ratio profile for returns. This is to be somewhat expected. In order to match the small negative slope of the variance ratio profile seen in the data, the original habit model put all of the action into the pricing kernel. Now, the modified model introduces long run reversals into dividend growth. The combination of the two effects is much too large relative to the data. This also suggests that return predictability is too strong in the modified model.

A more stringent test is to examine the term structure of risk premia. As documented by Van Binsbergen, Brandt, and Koijen (2012), this is steeply upward sloping in the original model and downward sloping in the data. Figure 15 showcases the term structure of equity yields in the habit model combined with my proposed model for cash flows as well as in the original speicification of the model, from 1 to 120 months.

The overall effect of the modified cash flow process is to shift more of the risk premium, relatively speaking, to shorter maturities, which is evident because the equity yields in the modified model line are always above the original model. Furthermore, over this horizon, the term structure becomes essentially flat, whereas it is always upward sloping in the original model. Both cases remain far from the data, however – not steeply downward sloping enough and an order of magnitude too low.

2.5 Conclusion

The evidence points towards a strong transitory component in equity payout that generates long run reversals. A majority of consumption-based asset pricing models specify a dividend growth process that is fundamentally at odds with this empirical evidence.

Equilibrium models contain market clearing conditions that must hold in equilibrium. I propose a model for cash flows that respects the equilibrium market clearing condition that states aggregate consumption is the sum of capital and labor income. Instead of modeling dividends directly, I model the log labor income-consumption ratio with a simple autoregressive process. I use this model and the market clearing condition to back out, or imply, dividend growth rates period-by-period. These implied dividend growth rates exhibit predictability and heteroskedasticity, as in the data. The cash flow model can be embedded into existing discrete time endowment economies to capture the effects that labor market frictions have on labor income and dividends. The approach is tractable, delivers cointegration between macroeconomic quantities, and allows models to simultaneously match the stylized facts about the univariate and joint behavior of consumption, dividend, and labor income growth rates.

When embedded into the habit and long run risks models, the cash flow model generates downward sloping term structures of dividend volatility and upward sloping term structures of labor income volatility, as in the data. Beyond the properties of cash flows, the method also has positive implications for asset prices, bringing the term structure of dividend strip expected returns and volatility closer to the data, and continuing to deliver a high and volatile equity risk premium.

Figure 1. Term Structure of Volatility of Consumption and Payout Growth Rates

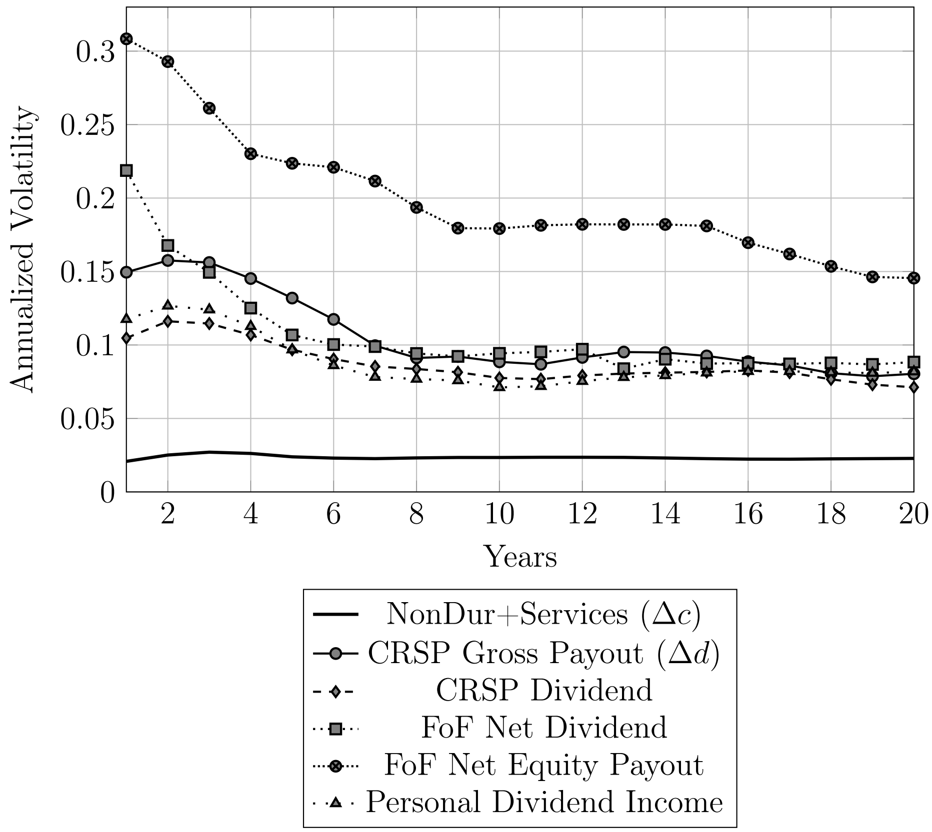

Note. Annualized volatility is plotted against horizon. The data is annual and spans 1930-2019. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction, after which I take the square root.

Figure 2. Variance Ratios of Consumption and Payout Growth Rates: 1930-2019

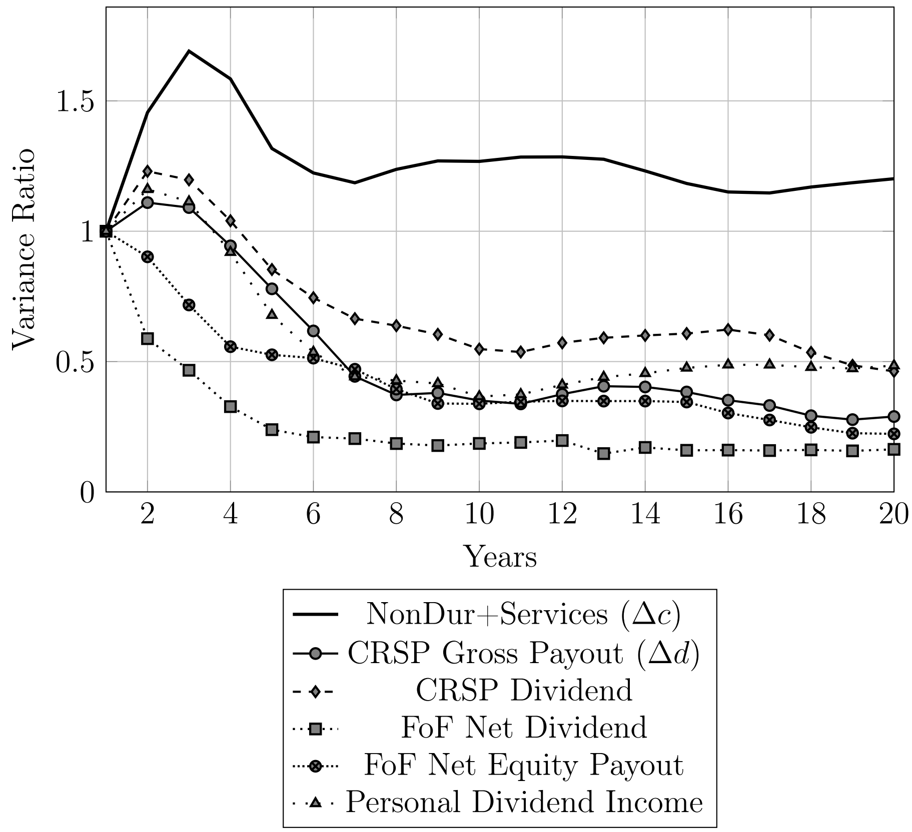

Note. Variance ratios plotted against horizon. The data is annual and spans 1930-2019. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. The variance ratio is defined $\frac{1}{k}\text{Var}(x_{t+k} - x_t)/\text{Var}(x_{t+1} - x_t)$.

Figure 3. Representative Standard Error Illustration

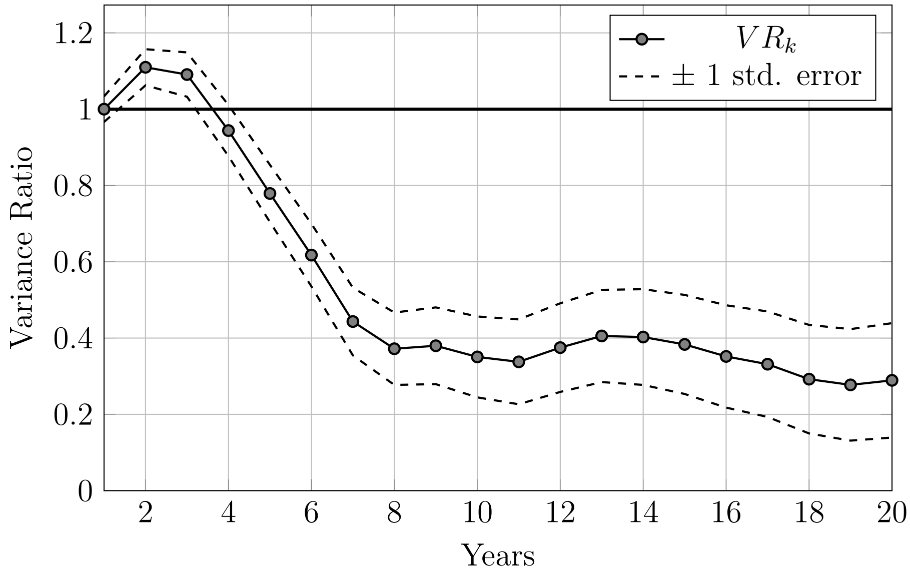

Note. Variance ratio plotted against horizon for the CRSP gross payout series, along with 1 standard error bands. The data is annual and spans 1930-2019. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. Asymptotic Bartlett standard errors scaled by the 30 year variance of the series. The horizontal solid line is the IID benchmark, equal to one.

Figure 4. Variance Ratios of Consumption and Payout Growth Rates: 1948-2019

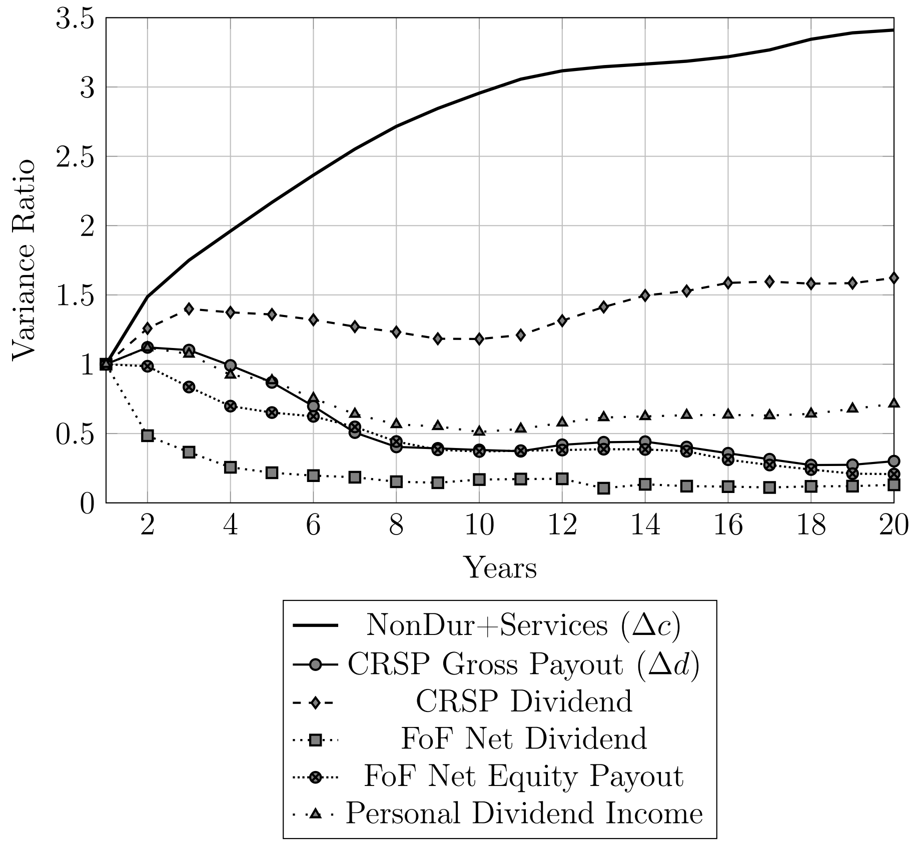

Note. Variance ratios plotted against horizon. The data is annual and spans 1948-2019. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. The variance ratio is defined $\frac{1}{k}\text{Var}(x_{t+k} - x_t)/\text{Var}(x_{t+1} - x_t)$.

Figure 5. Economic Recoveries

(A) Output and Consumption

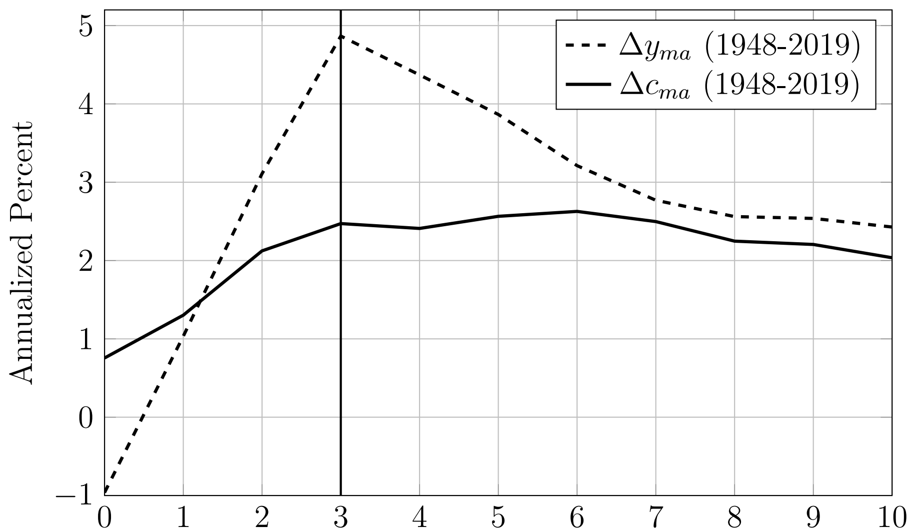

(B) Payout

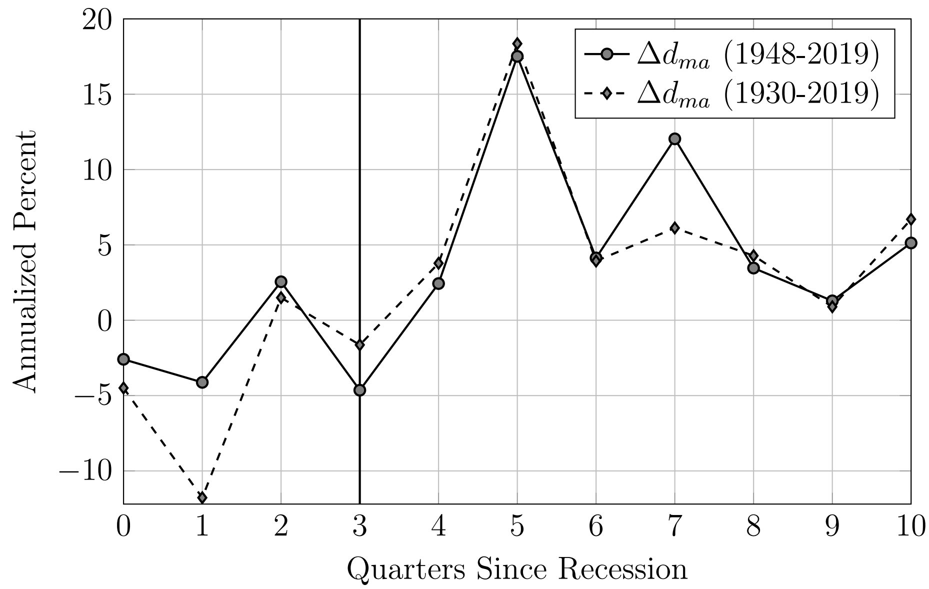

Note. Average annualized growth following the end of a recession. Recessions are dated using the midpoint method. In panel A, a MA(4) filter is applied to quarterly output (Δy) and consumption (Δc) growth rates. In panel B, the quarterly growth rate in a trailing 4 quarter average of CRSP gross payout is used as in Bansal, Dittmar, and Lundblad (2005). Due to the smoothing, the observation for the first quarter following a recession is an average of 1 expansion quarter and 3 recession quarters. The vertical solid line on the plot at time 3 indicates the first observation without any recession growth rate in its average.

Figure 6. Variance Ratios of Size Portfolio Payout Growth Rates

(A) Dividends+Repurchases

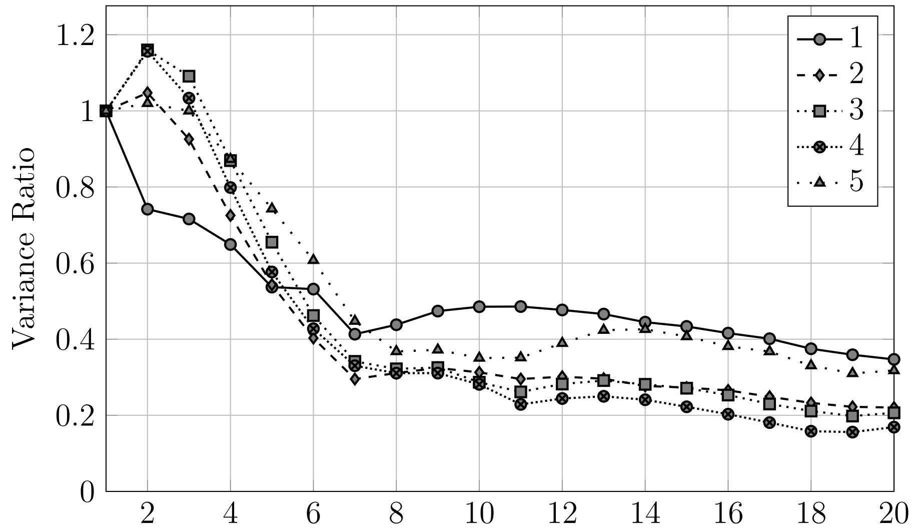

(B) Dividends

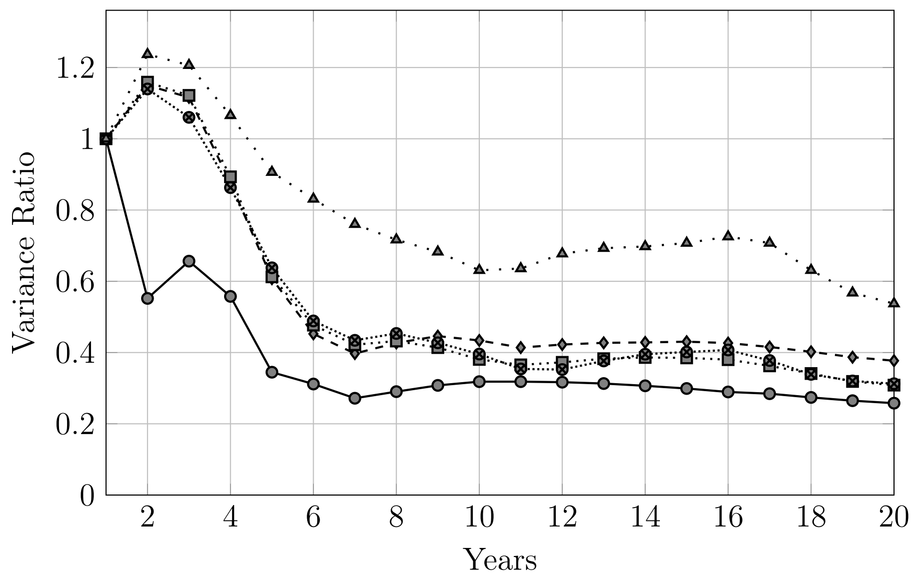

Note. Variance ratios plotted against horizon for five portfolios formed on firm size. The data is annual and spans 1930-2019. End of June market capitalization is used as the sorting variable with NYSE breakpoints, portfolios are value-weighted and rebalanced at the end of June, and 5 is the portfolio with largest firms. The top panel shows gross payout, while the bottom shows cash dividends. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. The variance ratio is defined $\frac{1}{k}\text{Var}(x_{t+k} - x_t)/\text{Var}(x_{t+1} - x_t)$.

Figure 7. Variance Ratios of Book-to-Market Portfolio Payout Growth Rates

(A) Dividends+Repurchases

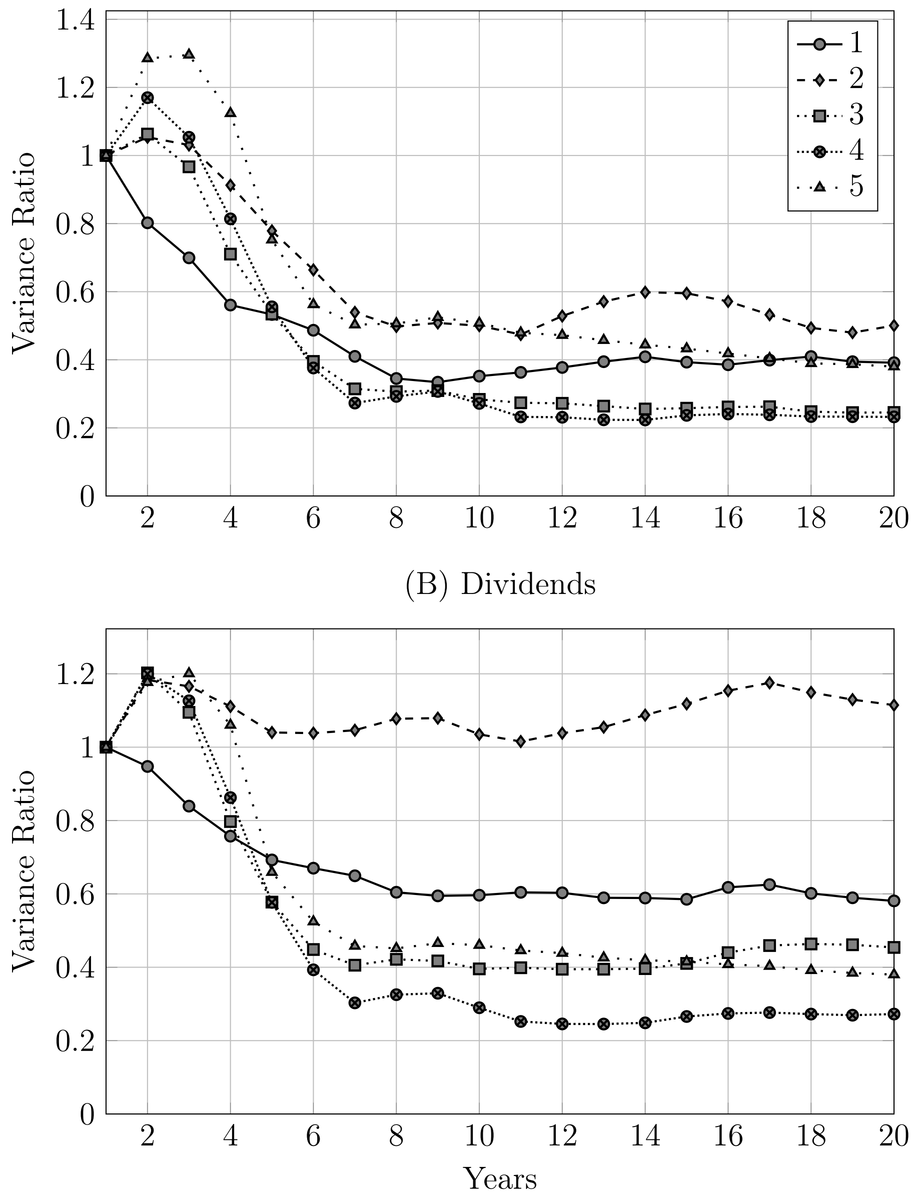

Note. Variance ratios plotted against horizon for five portfolios formed on firm book-to-market. The data is annual and spans 1930-2019. Book-to-market . . . market capitalization is used as the sorting variable with NYSE breakpoints, portfolios are value-weighted and rebalanced at the end of June, and 5 is the portfolio with largest firms. The top panel shows gross payout, while the bottom shows cash dividends. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. The variance ratio is defined $\frac{1}{k}\text{Var}(x_{t+k} - x_t)/\text{Var}(x_{t+1} - x_t)$.

Figure 8. Variance Ratios of Momentum Portfolio Payout Growth Rates

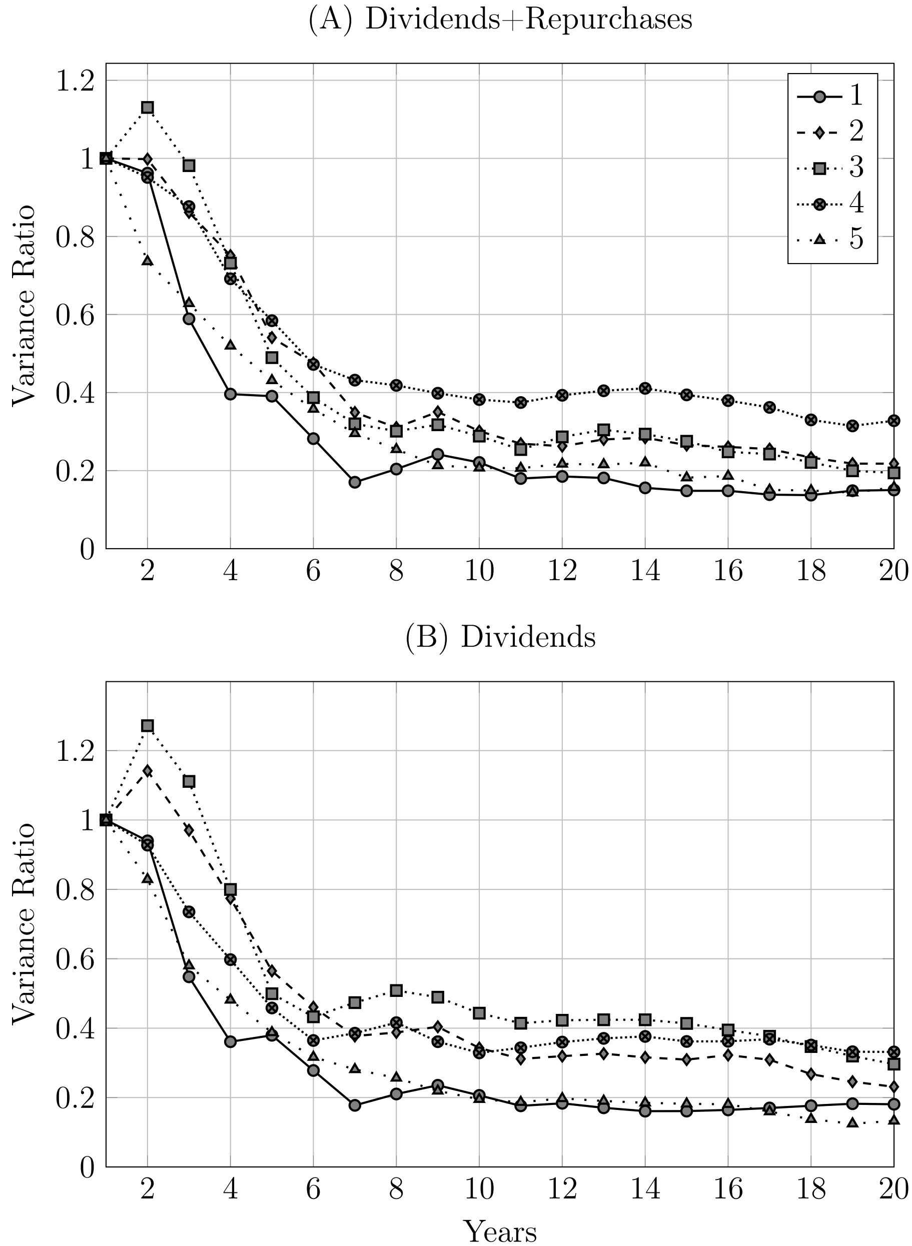

Note. Variance ratios plotted against horizon for five portfolios formed on return momentum. The data is annual and spans 1930-2019. Past [-11,-2] return is used as the sorting variable with NYSE breakpoints, portfolios are value-weighted and rebalanced at the end of June, and 5 is the portfolio with highest returns. The top panel shows gross payout, while the bottom shows cash dividends. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. The variance ratio is defined $\frac{1}{k}\text{Var}(x_{t+k} - x_t)/\text{Var}(x_{t+1} - x_t)$.

Figure 9. Variance Ratios of Beta Portfolio Payout Growth Rates

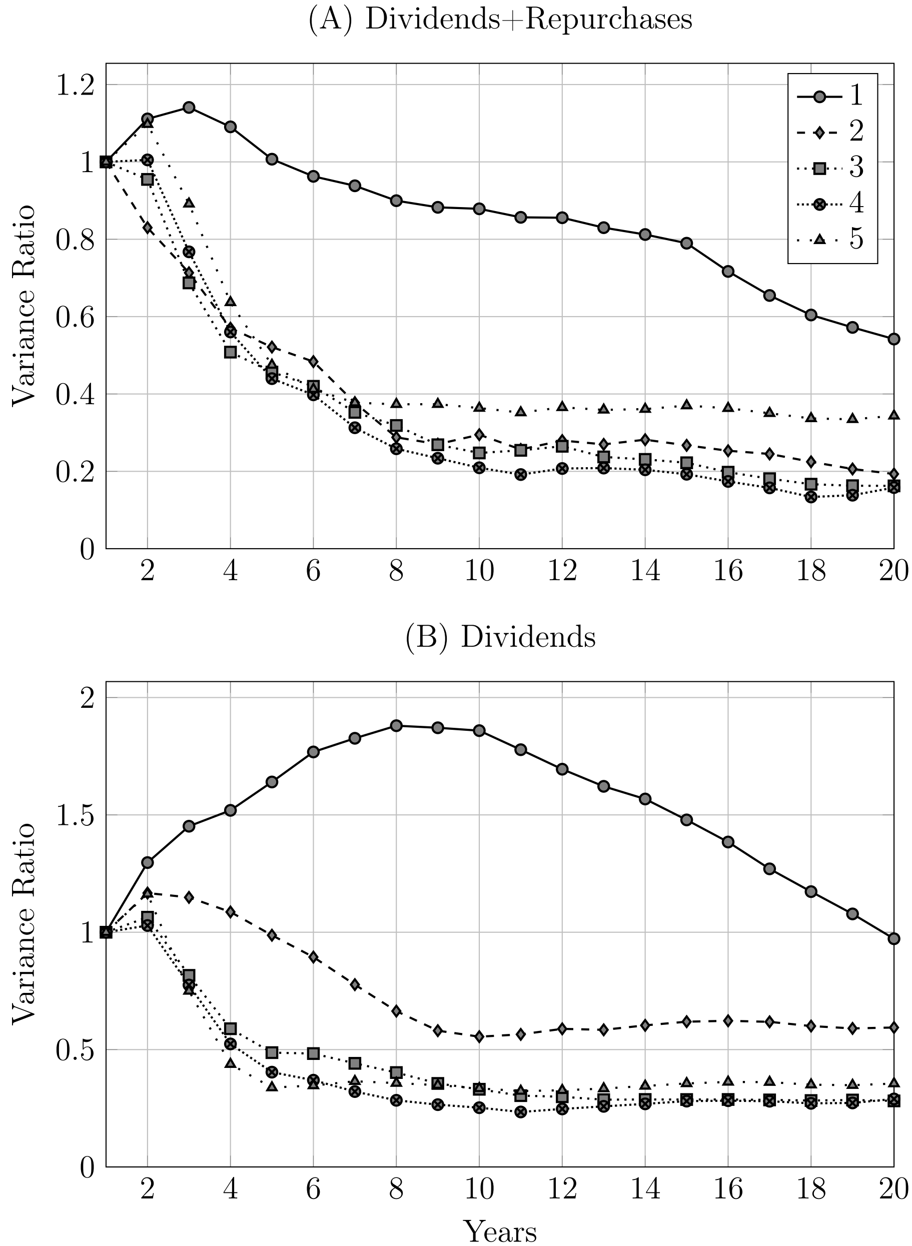

Note. Variance ratios plotted against horizon for five portfolios formed on market beta. The data is annual and spans 1930-2019. Beta computed using 60 months of past returns is used as the sorting variable with NYSE breakpoints, portfolios are value-weighted and rebalanced at the end of June, and 5 is the portfolio with highest beta. The top panel shows gross payout, while the bottom shows cash dividends. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. The variance ratio is defined $\frac{1}{k}\text{Var}(x_{t+k} - x_t)/\text{Var}(x_{t+1} - x_t)$.

Figure 10. Variance Ratios of Industry Portfolio Payout Growth Rates

(A) Dividends+Repurchases

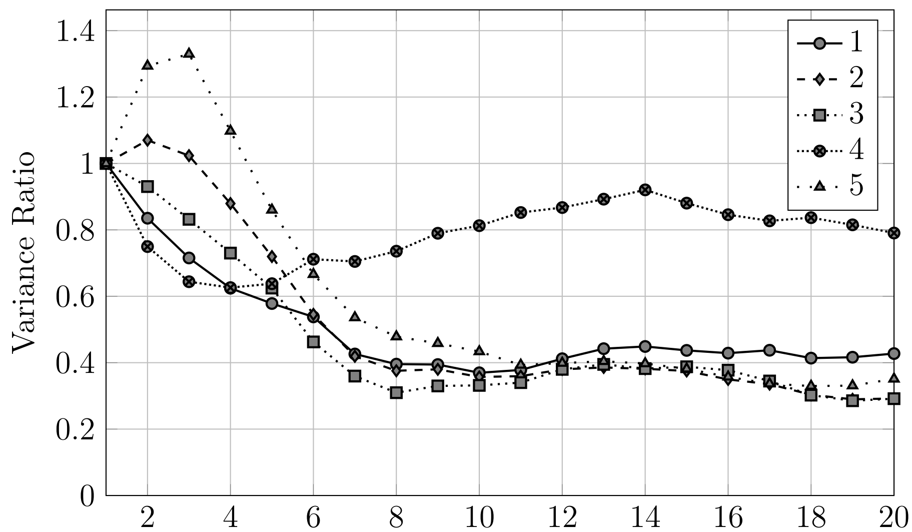

(B) Dividends

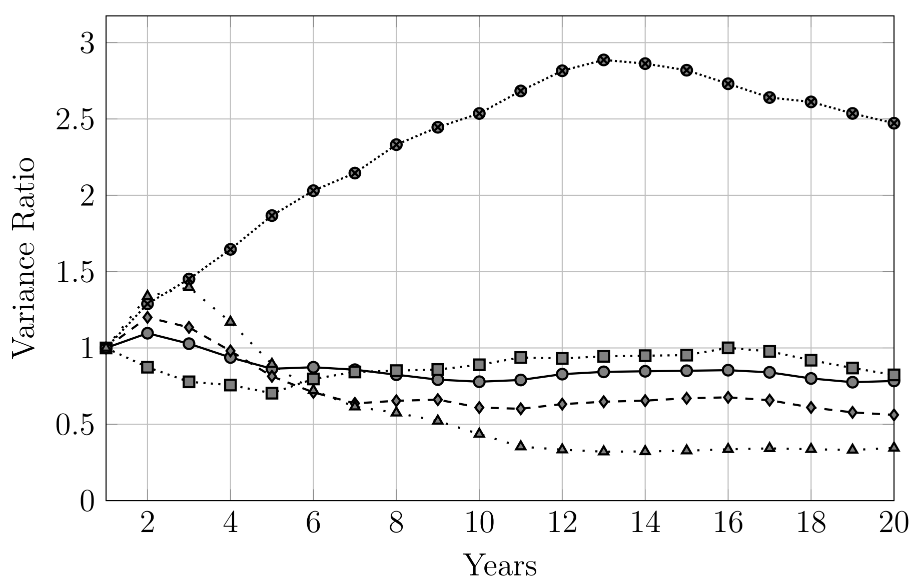

Note. Variance ratios plotted against horizon for the five Fama-French industry classifications. The data is annual and spans 1930-2019. Portfolios are value-weighted and rebalanced at the end of June. The top panel shows gross payout, while the bottom shows cash dividends. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. The variance ratio is defined $\frac{1}{k}\text{Var}(x_{t+k} - x_t)/\text{Var}(x_{t+1} - x_t)$.

Figure 11. Variance Ratios of Payout Growth Rates in Models

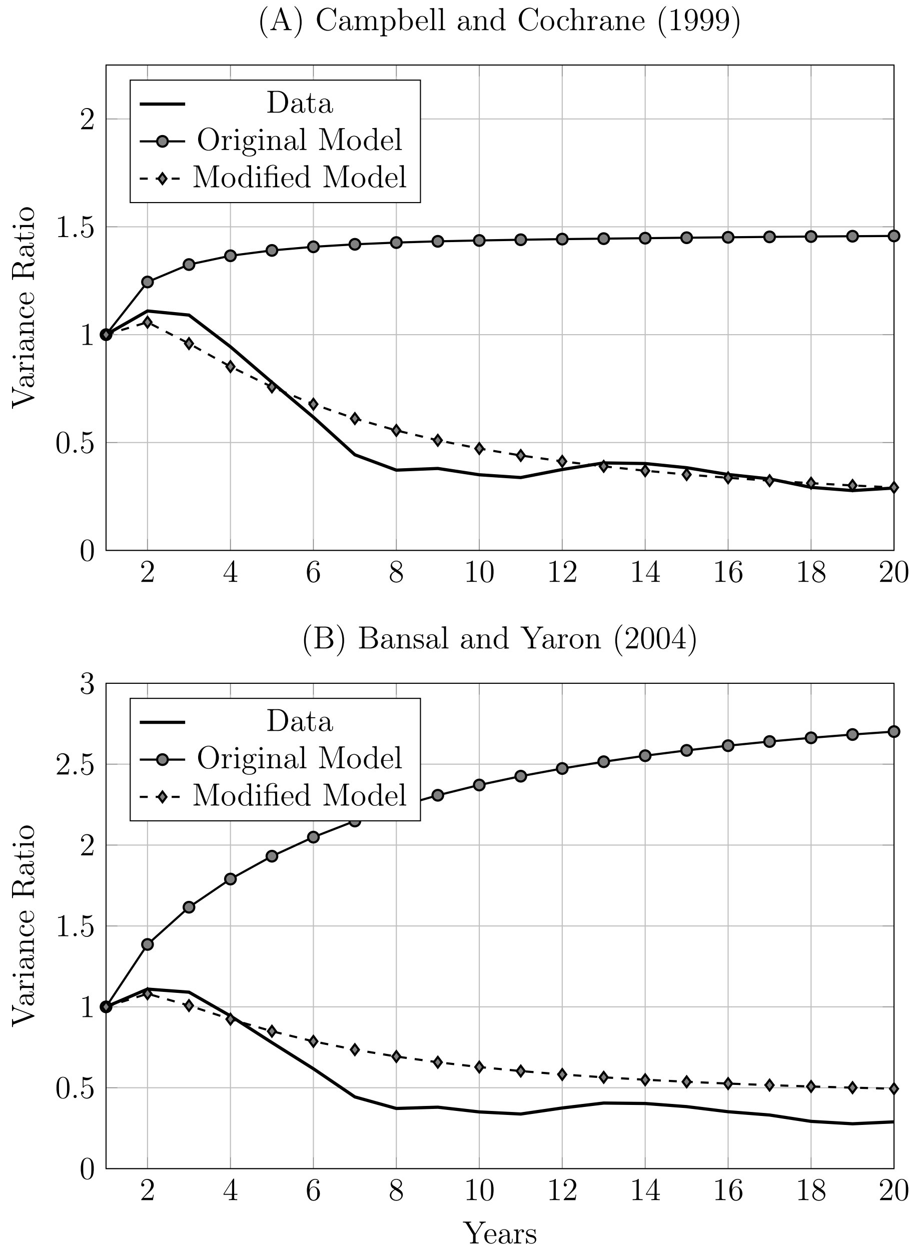

Note. Variance ratios plotted against horizon. The data is annual growth in CRSP gross payout and spans 1930-2019. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. The variance ratio is defined $\frac{1}{k}\text{Var}(x_{t+k} - x_t)/\text{Var}(x_{t+1} - x_t)$.

Figure 12. Variance Ratios of Labor Income Growth Rates

(A) 1930-2019

(B) 1948-2019

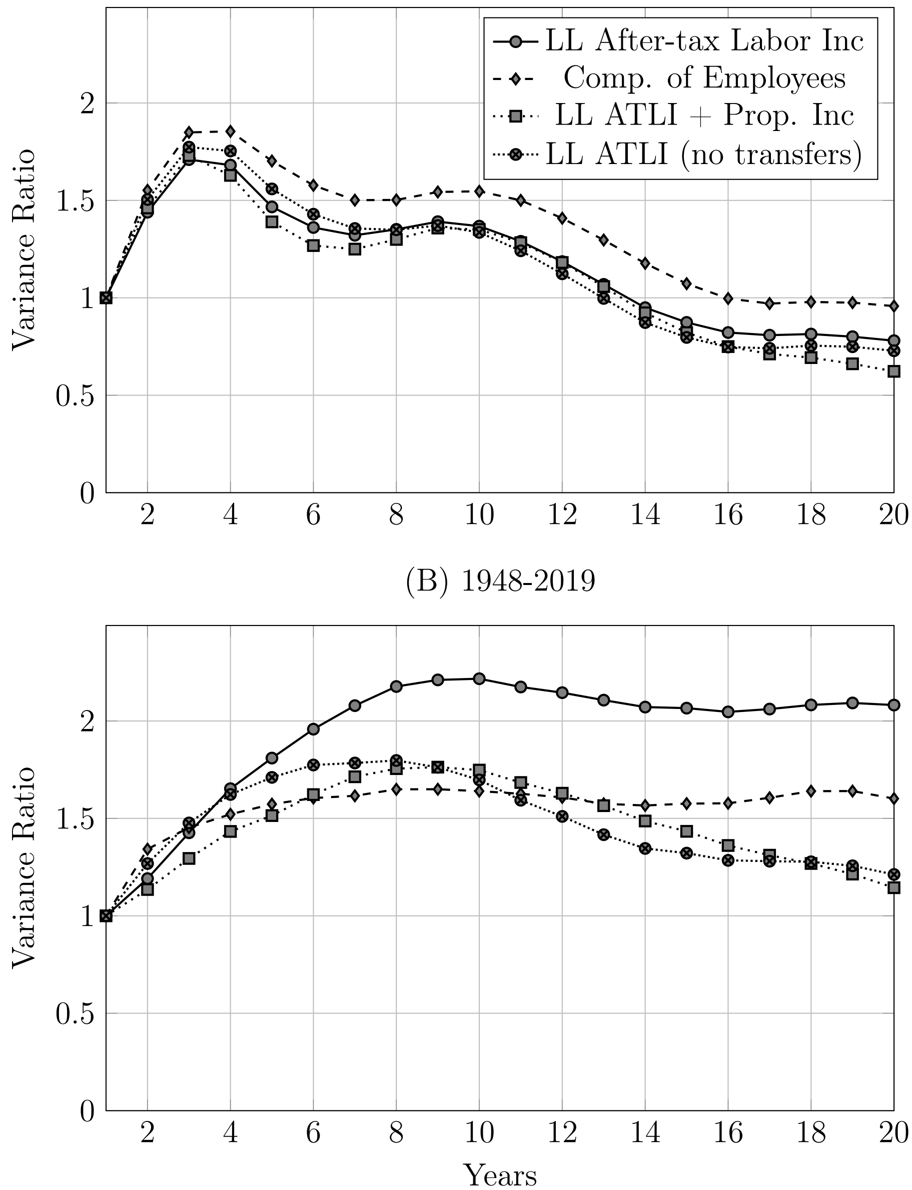

Note. Variance ratios plotted against horizon. The data is annual and spans either 1930-2019 (panel A) or 1948-2019 (panel B). Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. The variance ratio is defined $\frac{1}{k}\text{Var}(x_{t+k} - x_t)/\text{Var}(x_{t+1} - x_t)$.

Figure 13. Variance Ratios of Earnings Growth Rates

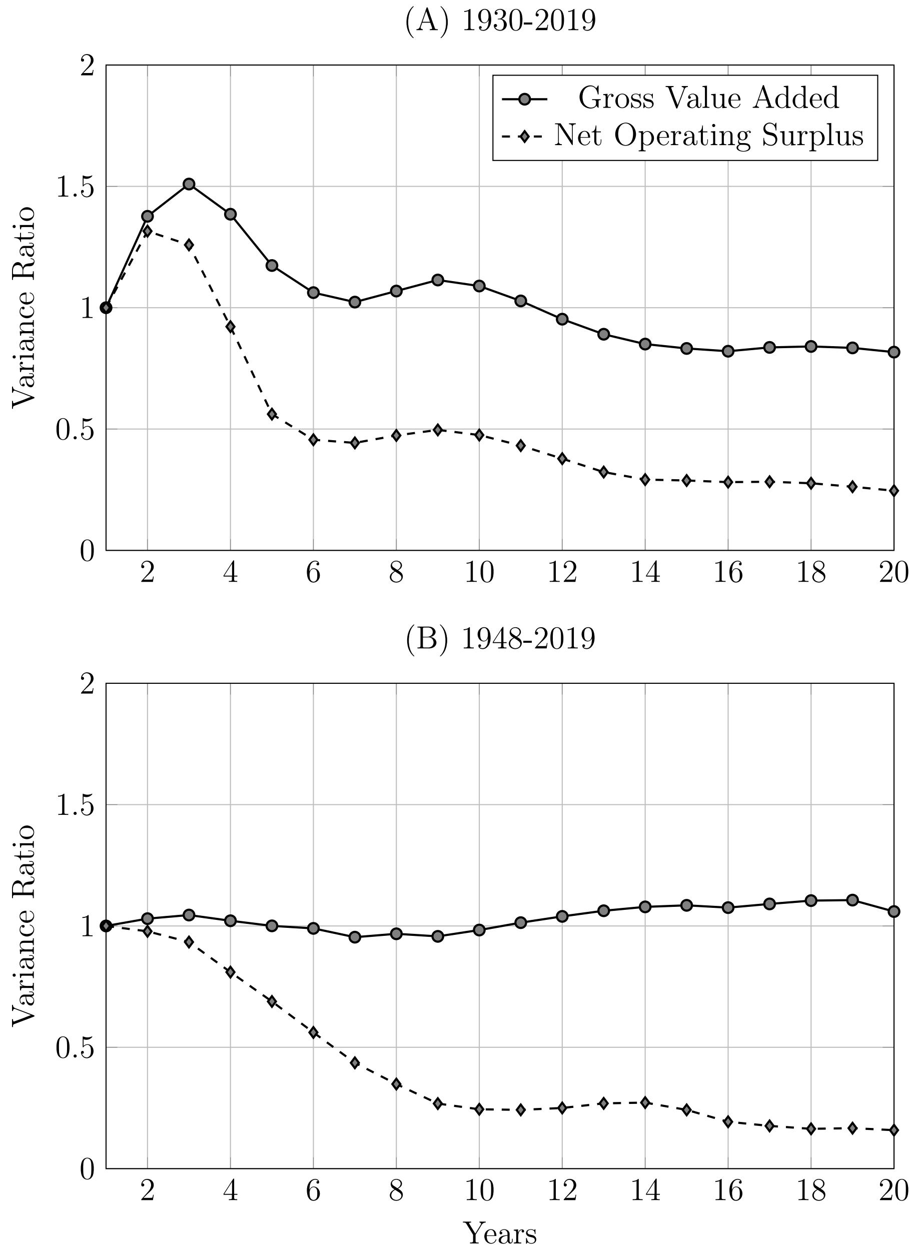

Note. Variance ratios plotted against horizon. The data is annual and spans either 1930-2019 (panel A) or 1948-2019 (panel B). Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. The variance ratio is defined $\frac{1}{k}\text{Var}(x_{t+k} - x_t)/\text{Var}(x_{t+1} - x_t)$.

Figure 14. Leverage Effect

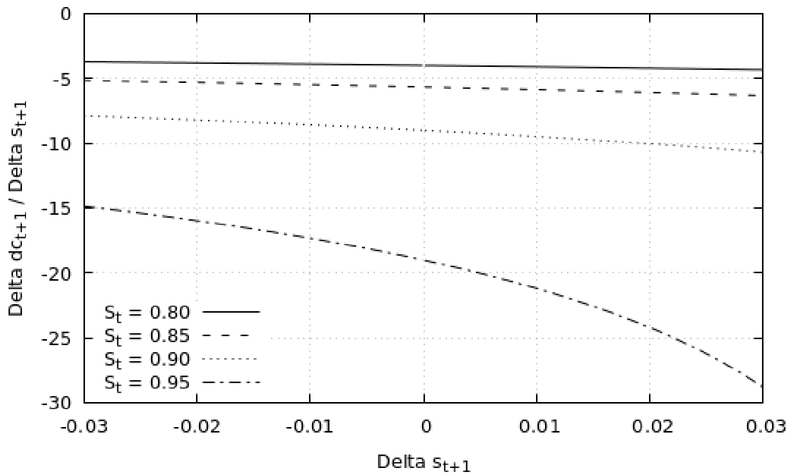

Note. The ratio of the log dividend share growth rate (Δdc_{t+1}) to the log labor share growth rate (Δs_{t+1}) conditional on a value for Δs_{t+1} (x-axis), given four different values for S_t.

Figure 15. Term Structure of Equity Yields: Habit Model

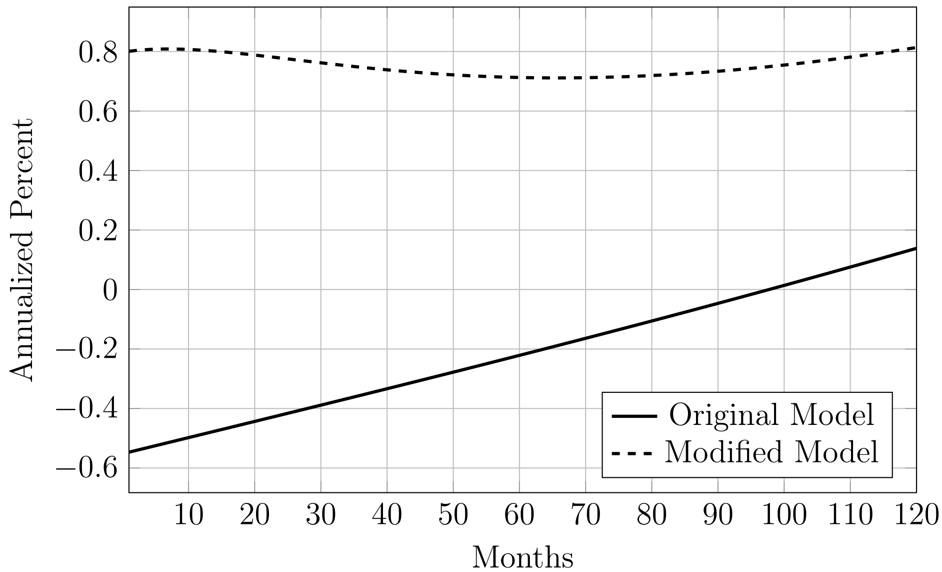

Note. Equity yields are computed as $-\frac{12}{k}\ln\left(PD_{t,t+k}\right)\times 100$, where $PD_{t,t+k}$ is the price-dividend ratio at t for the zero coupon equity claim with maturity k. Displayed is the unconditional average of the conditional term structures, averaged by interpolation over a long simulation of the model states.

Table 1. Summary statistics

x	$\mathbb{E}[x]$	$\sigma(x)$	$\gamma(x)$	$\rho_1(x)$	$\text{Corr}(x, \Delta c)$
			(A) 1930-2019		
Consumption					
NDS (Δc)	1.82	2.08	-1.50	0.47	
Equity Payout					
$\text{CRSP}_{\text{total}}$ (Δd)	1.12	14.95	-1.09	0.15	0.53
CRSP_{div}	1.65	10.48	-0.38	0.20	0.45
FoF_{div}	2.36	21.87	-0.12	-0.34	0.27
$\text{FoF}_{\text{total}}$	4.72	30.84	0.98	-0.02	0.12
Dividend Inc.	1.90	11.76	-0.78	0.14	0.58
Labor Income					
ATLI ($\Delta \ell$)	2.19	3.72	0.03	0.43	0.72
COE	2.16	4.53	0.10	0.53	0.68
$\text{ATLI}_{\text{withPI}}$	2.11	4.10	-0.26	0.46	0.76
$\text{ATLI}_{\text{notransfers}}$	1.87	4.16	-0.08	0.49	0.68
			(B) 1948-2019		
Consumption					
NDS (Δc)	1.89	1.20	-0.34	0.46	
Equity Payout					
$\text{CRSP}_{\text{total}}$ (Δd)	2.59	12.89	-1.20	0.09	0.26
CRSP_{div}	2.56	7.04	0.71	0.23	0.04
FoF_{div}	3.32	22.38	-0.08	-0.43	0.09
$\text{FoF}_{\text{total}}$	4.39	29.46	0.44	-0.04	0.30
Dividend Inc.	2.92	8.45	-1.59	0.09	0.29
Labor Income					
ATLI ($\Delta \ell$)	2.16	1.70	0.07	0.16	0.72
COE	2.03	2.47	-0.31	0.31	0.80
$\text{ATLI}_{\text{withPI}}$	2.05	1.79	-0.17	0.11	0.74
$\text{ATLI}_{\text{notransfers}}$	1.86	2.15	-0.32	0.24	0.76

Note. Annual summary statistics for the series used in the analysis. The mean and volatility are expressed in percent. γ = skewness, ρ_1 = autocorrelation. The data is annual and spans either 1930-2019 (panel A) or 1948-2019 (panel B).

Table 2. Long horizon predictability in payout

Dependent Variable: Δd_{t+1}	(1)	(2)
Constant	0.015	0.046***
	(0.014)	(0.017)
Δd_t	0.155*	-0.085
	(0.096)	(0.090)
Δd_{t-1}		-0.279^{*}
		(0.177)
Δd_{t-2}		-0.280^{***}
		(0.102)
Δd_{t-3}		-0.144^{*}
		(0.106)
Δd_{t-4}		-0.205^{**}
		(0.100)
Δd_{t-5}		-0.266^{***}
		(0.071)
N	89	84
Adj. R^2	0.017	0.170

Note. Regressions of equity payout growth on its lags. Equity payout growth is defined as CRSP dividends plus repurchases. The data is annual and spans 1930-2019. Newey-West standard errors in parentheses. Stars indicate significance at the 10 (*), 5 (**), and 1 (***) percent confidence levels.

Table 3. Long horizon variances and variance ratios: Consumption and Payout

	1 year	20 years	30 years	VR_{20}	VR_{30}	VR^b_{20}	VR^b_{30}
			(A) 1930-2019				
NDS (Δc)	0.000433	0.000520	0.000756	1.201	1.746		
	0.0208	0.0228	0.0275				
$\text{CRSP}_{\text{total}}$ (Δd)	0.022352	0.006463	0.006157	0.289	0.275	0.023	0.034
	0.1495	0.0804	0.0785				
CRSP_{div}	0.010977	0.005075	0.003929	0.462	0.358	0.047	0.069
	0.1048	0.0712	0.0627				
FoF_{div}	0.047843	0.007812	0.007431	0.163	0.155	0.011	0.016
	0.2187	0.0884	0.0862				
$\text{FoF}_{\text{total}}$	0.095089	0.021175	0.024838	0.223	0.261	0.005	0.008
	0.3084	0.1455	0.1576				
Dividend Inc.	0.013824	0.006699	0.007763	0.485	0.562	0.038	0.055
	0.1176	0.0818	0.0881				
			(B) 1948-2019				
NDS (Δc)	0.000143	0.000488	0.000599	3.411	4.189		
	0.0120	0.0221	0.0245				
$\text{CRSP}_{\text{total}}$ (Δd)	0.016616	0.004997	0.005517	0.301	0.332	0.029	0.036
	0.1289	0.0707	0.0743				
CRSP_{div}	0.004953	0.008036	0.011339	1.623	2.289	0.099	0.121
	0.0704	0.0896	0.1065				
FoF_{div}	0.055010	0.006521	0.004906	0.130	0.098	0.010	0.012
	0.2238	0.0808	0.0700				
$\text{FoF}_{\text{total}}$	0.086798	0.018010	0.018030	0.207	0.208	0.006	0.007
	0.2946	0.1342	0.1343				
Dividend Inc.	0.007140	0.005100	0.005294	0.714	0.741	0.068	0.084
	0.0845	0.0714	0.0728				

Note. The first 3 columns of this table display the annualized variance (first row) and corresponding volatility (second row) for horizons of 1, 20, and 30 years. The middle columns provide an estimate of the variance ratios for these series at 20 and 30 year horizons. The final 2 columns are computed as the 20 and 30 year annualized variances for consumption growth relative to the variance of dividend growth at the first horizon. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. The data is annual and spans either 1930-2019 (panel A) or 1948-2019 (panel B).

Table 4. Long horizon variances and variance ratios: Labor income

	1 year	20 years	30 years	VR_{20}	VR_{30}	VR^b_{20}	VR^b_{30}
			(A) 1930-2019				
NDS (Δc)	0.000433	0.000520	0.000756	1.201	1.746		
	0.0208	0.0228	0.0275				
ATLI ($\Delta \ell$)	0.001383	0.001079	0.001078	0.780	0.780	0.376	0.547
	0.0372	0.0328	0.0328				
COE	0.002056	0.001968	0.002071	0.958	1.007	0.253	0.368
	0.0453	0.0444	0.0455				
$\text{ATLI}_{\text{withPI}}$	0.001677	0.001044	0.000848	0.623	0.506	0.310	0.451
	0.0410	0.0323	0.0291				
$\text{ATLI}_{\text{notransfers}}$	0.001729	0.001260	0.001245	0.728	0.720	0.301	0.437
	0.0416	0.0355	0.0353				
			(B) 1948-2019				
NDS (Δc)	0.000143	0.000488	0.000599	3.411	4.189		
	0.0120	0.0221	0.0245				
ATLI ($\Delta \ell$)	0.000291	0.000605	0.000551	2.082	1.898	1.679	2.062
	0.0170	0.0246	0.0235				
COE	0.000611	0.000979	0.000977	1.602	1.599	0.798	0.980
	0.0247	0.0313	0.0313				
$\text{ATLI}_{\text{withPI}}$	0.000321	0.000367	0.000247	1.144	0.770	1.521	1.868
	0.0179	0.0192	0.0157				
$\text{ATLI}_{\text{notransfers}}$	0.000463	0.000561	0.000508	1.212	1.098	1.054	1.294
	0.0215	0.0237	0.0225				

Note. The first 3 columns of this table display the annualized variance (first row) and corresponding volatility (second row) for horizons of 1, 20, and 30 years. The middle columns provide an estimate of the variance ratios for these series at 20 and 30 year horizons. The final 2 columns are computed as the 20 and 30 year annualized variances for consumption growth relative to the variance of income growth at the first horizon. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. The data is annual and spans either 1930-2019 (panel A) or 1948-2019 (panel B).

Table 5. Select Model Moments

x	$\mathbb{E}[x]$	$\sigma(x)$	$\gamma(x)$	$\rho_1(x)$	$\text{Corr}(x, \Delta c)$
		(A) Campbell and Cochrane (1999)			
Consumption					
Δc	1.89	1.22	0.00	0.23	
Equity Payout					
Original Model	1.89	9.11	0.00	0.23	0.20
Modified Model	1.89	9.43	−0.01	0.05	0.40
Labor Income					
Modified Model	1.89	1.32	0.00	0.20	0.71
		(B) Bansal and Yaron (2004)			
Consumption					
Δc	1.81	2.85	0.00	0.47	
Equity Payout					
Original Model	1.82	11.26	0.00	0.37	0.30
Modified Model	1.80	11.43	−0.04	0.07	0.56
Labor Income					
Modified Model	1.81	2.66	0.00	0.53	0.92

Note. Annual summary statistics for the model series. Original Model moments are from the original specification and calibrations of the model, while Modified Model moments are computed from the model combined with my proposed cash flow process. The mean and volatility are expressed in percent. γ = skewness, ρ_1 = autocorrelation. Values are the mean taken across 10,000 small sample simulations of length 1,092 months (91 years), time aggregated to the annual frequency (90 years).

Table 6. Long horizon variances and variance ratios: Models

	1 year	20 years	30 years	VR_{20}	VR_{30}	VR^b_{20}	VR^b_{30}
		(A) Campbell and Cochrane (1999)					
Δc	0.000149	0.000216	0.000218	1.455	1.465		
	0.0122	0.0147	0.0148				
Original Δd	0.008302	0.012101	0.012135	1.458	1.462		
	0.0911	0.1100	0.1101				
Modified Δd	0.008896	0.002595	0.002150	0.292	0.242	0.024	0.025
	0.0943	0.0509	0.0464				
Modified $\Delta \ell$	0.000176	0.000225	0.000225	1.281	1.282	1.233	1.242
	0.0132	0.0150	0.0150				
		(B) Bansal and Yaron (2004)					
Δc	0.000810	0.002919	0.003055	3.605	3.773		
	0.0285	0.0540	0.0553				
Original Δd	0.012688	0.034275	0.035595	2.701	2.805		
	0.1126	0.1851	0.1887				
Modified Δd	0.013056	0.006449	0.006001	0.494	0.460	0.224	0.234
	0.1048	0.0712	0.0627				
Modified $\Delta \ell$	0.000708	0.002875	0.003013	4.060	4.256	4.123	4.315
	0.0266	0.0536	0.0549				

Note. The first 3 columns of this table display the annualized variance (first row) and corresponding volatility (second row) for horizons of 1, 20, and 30 years. The middle columns provide an estimate of the variance ratios for these series at 20 and 30 year horizons. The final 2 columns are computed as the 20 and 30 year annualized variances for consumption growth relative to the variance of dividend growth at the first horizon. Annualized variance computed following Cochrane and Sbordone (1988) with small sample bias correction. Values are the mean taken across 10,000 small sample simulations of length 1,092 months (91 years), time aggregated to the annual frequency (90 years).

Table 7. Select Return Moments: Habit Model

	Original	Modified
$\mathbb{E}[r]$	5.0	5.4
$\sigma(r)$	15.0	10.3
$VR_1(r)$	0.97	0.93
$VR_2(r)$	0.94	0.87
$VR_5(r)$	0.89	0.79
$VR_{10}(r)$	0.79	0.56

Note. Select moments of annual log returns in the Habit model. Model statistics computed from 100,000 small samples of 90 years.

Chapter 3

DISSECTING THE EQUITY PREMIUM

3.1 Introduction

Most economists would agree that macroeconomic risk is an important source of stock market risk premia. Indeed, equilibrium asset pricing theories typically assume that fundamental shocks represent *the only source* of priced risk. There is far less agreement, however, about the types of macroeconomic shocks reflected in asset prices or investors' risk attitudes towards them. An important reason is that theories based on vastly different fundamental risks, such as Campbell and Cochrane (1999), Bansal and Yaron (2004), Rietz (1988), and Barro (2009), have similar implications for many key data features, including the average equity premium.

This paper proposes a data-based metric for discriminating between asset pricing theories based on *how* they generate the equity premium. Our metric builds on the concept of Arrow-Debreu securities – claims to a unit payoff in a particular state of nature. The absence of arbitrage opportunities allows us to write the average equity premium as

$$\mathbb{E}[R_{t+1} - R_t^f] = \int_{-1}^{\infty} R\,[f(R) - f^*(R)]\,dR, \tag{3.1}$$

where $f(R)$ and $f^*(R) = \mathbb{E}[f_t^*(R)]$ are, respectively, the average payoff and average forward price of an Arrow-Debreu security that pays \$1 if the realized stock market return equals R. The average payoff can equivalently be interpreted as an unconditional return distribution. The forward price, $f_t^*(R)$, can equivalently be interpreted as a preference-adjusted, or "risk-neutralized", return distribution under which the expected stock market return equals the risk-free rate. Equation (3.1) represents a powerful tool for understanding how stock prices are formed. Specifically, it allows us to assess how much individual states contribute to the overall equity premium by evaluating the integral separately for different regions.

We implement the decomposition empirically for the S&P 500, a broad US stock market index, by estimating f from returns and f^* from option prices. Using these estimates, the solid line in Figure 16 shows the fraction of the equity premium associated with returns below different thresholds, a function we coin "$EP(x)$". By construction, increasing segments of the $EP(x)$ curve reflect states that contribute positively to the equity premium, whereas decreasing segments reflect states that contribute negatively. Strikingly, the figure shows that the 80/100 of the equity premium is associated with monthly returns below -10%, while 67/100 of it stems from the gray-shaded region between -30% and -10%. Such returns occur roughly at a business cycle frequency, and they have historically coincided with events like Iraq's invasion of Kuwait in 1990, the collapse of Long Term Capital Management in 1998, the September 11 attacks in 2001, the bankruptcy of Lehman Brothers in 2008, and the Covid-19 pandemic in 2020. Investors are therefore predominantly compensated for shocks that coincide with infrequent, large, but not extremely large negative stock market returns.

We show that the shape of $EP(x)$ to the right of the gray-shaded region, including the fact that the curve rises above one, is consistent with a pricing kernel that is slightly U-shaped when projected onto returns. Because this non-monotonicity has been extensively documented in prior work and mechanisms to explain it exist, we do not make it a focal point of our paper.[24]

[24] Cuesdeanu and Jackwerth (2018) provide a recent survey of this literature.

The remaining lines in Figure 16 show $EP(x)$ for different asset pricing theories. The habit model of Campbell and Cochrane (1999) and the long-run risks model of Bansal and Yaron (2004) build on very different economic mechanisms, but nevertheless attribute the equity premium to nearly identical return states. In both models, positive returns account for about half of the equity premium, while returns below -10% account for very little. The rare disaster model of Barro (2009) attributes 92/100 of the equity premium to states on the left of the gray-shaded region, and more than half to returns below -89%. By extending the Barro model with time variation in the probability of disasters and a recursive utility agent, Wachter (2013) combines the long-run risks and disaster mechanisms. We find that her model attributes about 2/3 of the equity premium to the far left tail and 1/3 to states on the right of the gray-shaded region. This evidence suggests that sources of stock market risk premia in popular asset pricing models differ substantially from those in the data.

To understand the origin of these discrepancies, we separate the equity premium contribution of different return states into their probability and price of risk. This analysis shows that the aforementioned models undershoot risk prices for stock market tail events by a factor of nearly 2. Fundamental shocks that coincide with large market corrections are therefore considerably more painful to investors than commonly assumed. We examine the robustness of this result by looking at a model with disappointment aversion risk preferences, an extension of expected utility theory that overweights lower tail outcomes and rationalizes the Allais paradox (Gul (1991), Routledge and Zin (2010)). The model produces high risk prices and attributes 81/100 of the equity premium to returns between -30% and -10%, close to the value in the data. Interestingly, the same utility function that rationalizes puzzling implications of expected utility theory in experimental settings can account for puzzling implications of standard utility functions in a market-based setting. One interpretation of our empirical evidence is therefore that common utility functions are misspecified. As such, our finding has relevance for other areas of economics where decision making under uncertainty plays a role.

Related Literature.– Our finding complements that of Binsbergen, Brandt, and Koijen (2012), who use options data to show that claims to near term dividends earn larger risk premia than claims to far term dividends. Their evidence speaks primarily to the pricing of shocks with different persistence levels, whereas ours speaks primarily to the pricing of shocks with different magnitudes. It may appear intuitive that a common mechanism underlies both findings. We find, however, that Belo, Collin-Dufresne, and Goldstein (2015)'s modification of the long-run risks model, which replicates the downward-sloping term structure of expected returns on dividend strips, does not materially alter the model's implications for $EP(x)$. Similarly, the Generalized Disappointment Aversion model of Schreindorfer (2020) comes close to replicating our empirical evidence, but makes no headway in matching the term structure evidence. An important step for future work is therefore the development of a theory of the equity premium that matches the pricing of all shocks to aggregate dividends: transient, persistent, small, and large.

We also relate closely to Bollerslev and Todorov (2011), who quantify the importance of jumps for the *conditional* equity premium in the data, and conclude that "*any satisfactory equilibrium-based asset pricing model must be able to generate large and time-varying compensations for the possibility of rare disasters*". Methodologically, our approach differs from theirs because we study the *unconditional* equity premium. As a result, we do not require an extreme value theory to approximate the probability of large low-frequency (e.g. monthly) returns based on the probability of large high-frequency (e.g. 5 minute) returns. Economically, we reach the opposite conclusion regarding the importance of rare disasters. In particular, our evidence shows that the disaster models of Barro (2009) and Wachter (2013) attribute the equity premium primarily to the extreme left tail, while the data attributes it predominantly to an intermediate left tail region.

Welch (2018) studies protected equity to evaluate the rare disaster mechanism. He finds that investments in the market, combined with a 1-month put option struck 15% below the money, have historically earned average excess returns of 5% per year, compared with 7% on unprotected equity. Based on this evidence, he argues that disasters account for at most 2% (in absolute terms) of the equity premium. This reasoning overlooks, however, that put options only protect against losses in financial assets, not against losses in human capital or consumption. If a protected equity position experiences a return of -15% during a macroeconomic disaster, for example, that return will still coincide with high levels of marginal utility. Hence, even protected equity carries a premium for exposure to extreme tail events. Seo and Wachter (2019) formalize this argument by showing that the disaster model of Wachter (2013) replicates the large returns on protected equity, despite attributing the equity premium predominantly to disaster risk. Returns on protected equity are therefore unsuitable for learning about sources of stock market risk premia.

More broadly, Backus, Chernov, and Martin (2011) and Martin (2017) show that many representative agent models are inconsistent with option prices, in that they either imply too little or too much skewness in the option-implied return distribution. Our evidence is conceptually distinct from this finding because risk premia reflect *differences* between the option-implied and historical return distributions. Indeed, our analysis of Backus et al.'s preferred calibration, which matches option prices by construction, shows that consistency with option prices does not imply consistency with $EP(x)$.

Lastly, we build on a large literature in empirical option pricing, including Rubinstein (1994) and Broadie, Chernov, and Johannes (2009), which shows that out-of-the-money put options earn large negative returns on average. This prior evidence suggests that stock market crashes coincide with high marginal utility for investors, but it does not allow us to conclude that crash states are central to the equity premium. For example, put options that are struck 50% below the money may generate average returns of nearly -100%, while the associated payoff states nevertheless contribute little to the equity premium due to their rarity. We contribute to this literature by quantifying the importance of different return states for the equity premium.

3.2 A State-space Decomposition of the Equity Premium

Let $R_{t+1}^{cum} = \frac{S_{t+1}+D_{t+1}}{S_t} - 1$ be the cum-dividend return on the market and $f_t(R)$ its conditional probability density function (PDF). We assume the absence of arbitrage opportunities, which guarantees the existence of a risk-neutral probability measure, f_t^*, under which the expected return equals the risk-free rate,

$$\int_{-1}^{\infty} R f_t^*(R) dR = R_t^f. \tag{3.2}$$

The conditional PDF can be interpreted as the expected payoff of an Arrow-Debreu security that pays \$1 if $R_{t+1}^{cum} = R$ and \$0 otherwise, and $f_t^*(R)$ as the security's forward price. Using (3.2), we can write the equity premium as

$$\mathbb{E}_t[R_{t+1}^{cum}] - R_t^f = \int_{-1}^{\infty} R \left[f_t(R) - f_t^*(R) \right] dR \tag{3.3}$$

and interpret it as the difference between the expected payoff and the forward price of a portfolio of Arrow-Debreu securities.

In principle, (3.3) allows us to assess how much individual states contribute to the overall equity premium by evaluating the integral separately for different regions. Two issues hinder our ability to estimate the integrant empirically, however. First, option prices only identify state prices for ex-dividend, rather than cum-dividend, returns. Second, estimates of the conditional return distribution require strong statistical assumptions and necessarily suffer from a mismatch between investors' and the econometrician's information set.

We address issue one by assuming that $(t+1)$-dividends are in investors' time-t information set. This is reasonable for short horizons such as one month, which are the focus of our empirical analysis, because S&P 500 companies announce dividends on average three weeks in advance of the ex-dividend date (Schulz (2016), Table I).[25] The assumption implies that dividends drop out of (3.3), allowing us to interpret f_t and f_t^* as the expected payoff and forward price of Arrow-Debreu securities that are written on the *ex-dividend* return. To address issue two, we take the unconditional expectation of (3.3) and note that $\mathbb{E}[f_t(R)]$ equals the unconditional PDF, $f(R)$. Unlike the conditional PDF, f is straightforward to estimate from historical data. We also define $f^* \equiv \mathbb{E}[f_t^*]$, which yields the average equity premium as (3.1) in the introduction.

To decompose the equity premium, we define

$$EP(x) \equiv \frac{\int_{-1}^{x} R\,[f(R) - f^*(R)]\,dR}{\int_{-1}^{\infty} R\,[f(R) - f^*(R)]\,dR}, \tag{3.4}$$

[25] Binsbergen, Brandt, and Koijen (2012) show that claims to "near-term" dividends earn larger risk premia than the market itself. Because this finding is based on claims to market dividends over the coming 1.6 years (on average), however, it is not inconsistent with the absence of risk premia on very short-term dividends.

which measures the fraction of the average equity premium that is associated with returns below x. Note that the total equity premium in the denominator of (3.4) is a normalization constant that does not depend on x. The normalization ensures that, like a CDF, $EP(x)$ approaches zero for return thresholds in the far left tail and approaches one for return thresholds in the far right tail. Unlike a CDF, however, theory does not restrict $EP(x)$ to be monotonically increasing. Instead, the function is increasing for states that contribute positively to the equity premium and decreasing for states that contribute negatively.

3.3 Estimation

This section discusses data sources and our approach for estimating average payoffs and forward prices of Arrow-Debreu securities. We use the S&P 500 index as a proxy for the aggregate stock market and focus on a return horizon of 1 month. Because the availability of options data limits our analysis to 1990-2019, we sample daily to maximize the efficiency of our estimates. Specifically, we rely on a daily sample of $T = 7,556$ overlapping 30 calendar day returns, $R_{t:t+30}$, and estimate state prices for the corresponding return interval, $f_t^*(R_{t:t+30})$. The appendix details data filters and the estimation of f^*.

3.3.1 Payoffs

Return data comes from the Center for Research in Security Prices (CRSP). The payoff of an Arrow-Debreu security for state R at time $t+30$ is given by the indicator function $\mathbf{1}\{R_{t:t+30} = R\}$. We estimate average payoffs, $f(R) = \mathbb{E}[\mathbf{1}\{R_{t:t+30} = R\}]$, with the empirical PDF, which attaches a probability of $1/T$ to each return observation. In our sample, monthly returns have an annualized volatility of 14.9%, skewness of -0.70, kurtosis of 7.20, and a probability of 1.63% of falling below -10%. Returns below -10% have therefore occurred every $\frac{1}{12\times 0.0163} = 5.1$ years on average, whereas returns below -30% have not occurred in our sample.

3.3.2 Forward Prices

Each realized return $R_{t:t+30}$ is drawn from an unobserved conditional PDF, $f_t(R_{t:t+30})$. We rely on a classic result by Breeden and Litzenberger (1978) to recover the risk-neutral PDF for the corresponding return horizon from options prices. Specifically, the absence of arbitrage opportunities implies that the time-t price of a European put option with strike K and maturity $t+\tau$ equals $P_t(K,\tau) = \frac{1}{R_t^f}\int_0^K (K - S_{t+\tau}) f_t^*(S_{t+\tau}) dS_{t+\tau}$, the risk-neutral expectation of its payoff discounted at the risk-free rate. Differentiating this expression twice with respect to K yields the conditional risk-neutral PDF of the future price level as

$$f_t^*(S_{t+\tau}) = R_t^f \times \left.\frac{\partial^2 P_t(K,\tau)}{\partial K^2}\right|_{K=S_{t+\tau}}. \tag{3.5}$$

The risk-neutral PDF of ex-dividend returns follows from the change of variables $R_{t:t+\tau} = \frac{S_{t+\tau}}{S_t} - 1$ as $f_t^*(R_{t:t+\tau}) = S_t \times f_t^*(S_{t+\tau})$. Hence, (3.5) allows us to recover $f_t^*(R)$ from prices of options with different strikes and a common maturity.

Option price quotes are obtained from the Chicago Board Options Exchange (CBOE). The dataset contains end-of-day information on European exercise style options written on the S&P 500 index (underlying SPX). We apply standard filters to remove observations with low liquidity and obvious data errors and measure prices by the average of bid and ask quotes. Because quotes are only available for a discrete set of moneyness-maturity pairs, and contracts for extreme moneyness levels are often unavailable altogether, it is necessary to interpolate and extrapolate observed prices to obtain sensible estimates of f_t^*. We apply a standard approach from the options literature for doing so (see the appendix for details). f_t^* is then computed based on interpolated option prices and (3.5) via finite differences. Each of the resulting 7,556 daily conditional density estimates covers the 30 calendar day interval of one realized return.

The estimated $f^*(R) = \frac{1}{T}\sum_{t=1}^{T} f_t^*(R)$ is shown in Figure 17, together with a smoothed estimate of $f(R)$.[26] Under f^*, probability mass is shifted towards less favorable outcomes, and monthly returns have an annualized volatility of 19.4%, skewness of -1.42, kurtosis of 12.09, and a 4.29% probability of falling between -30% and -10% (one event every 1.9 years). This probability is 2.6 times higher than the corresponding historical probability. The dash-dotted line shows that the f^*-to-f probability ratio rises from 1.6 for returns of -10% to 9.1 for returns of -30%.

3.3.3 $EP(x)$

Using the estimates of f and f^*, we can evaluate the expression for $EP(x)$ in (3.4). We compute the integral $\int_{-1}^{x} Rf(R)dR$ based on the empirical PDF as $\frac{1}{T}\sum_{t=1}^{T} R_{t:t+30}\mathbf{1}\{R_{t:t+30} \leq x\}$, and evaluate the integral $\int_{-1}^{x} Rf^*(R)dR$ numerically using a global adaptive quadrature routine. Figure 16 in the introduction shows the resulting $EP(x)$ curve. The far left tail contributes $EP(-30\%) \approx 1/10$ to the equity

[26] The empirical PDF consists of point masses and is therefore challenging to plot. For Figure 17, we approximate f with a 10-th order polynomial between percentiles 10 and 90 and with a generalized extreme value (GEV) distribution in both extreme tails. The empirical PDF (without smoothing) is used for all other results in the paper.

premium, whereas the intermediate left tail contributes $EP(-10\%) - EP(-30\%) \approx 2/3$. The majority of the equity premium therefore represents compensation for shocks that coincide with large, but not extremely large negative returns.

The supplement repeats our analysis for a quarterly horizon and finds that most of the equity the equity premium continues to be attributable to an intermediate left tail region in the data, but to other states for the models in Figure 1. The choice of a monthly horizon is therefore not critical for our results.

3.3.4 Statistical Precision

To assess the amount of statistical uncertainty associated with estimates of $EP(x)$, we use a block bootstrap to construct the joint sampling distribution of $EP(-30\%)$ and $EP(-10\%) - EP(-30\%)$. The supplement provides details on the construction of bootstrap samples, and Figure 18 depicts the resulting sampling distribution. Clearly, it is likely that most of the equity premium is attributable to return states between -30% and -10% in population: $EP(-10\%) - EP(-30\%)$ exceeds one half in 90.1% of bootstrap samples. In contrast, returns below -30% are very likely not the main source of stock market risk premia, as $EP(-30\%)$ falls below one half in 99.5% of all samples. The same conclusion applies to returns above -10% (the omitted category of the partition), as $1 - EP(-10\%)$ falls below one half in 98.9% of all samples. These results suggest that our findings cannot be attributed to sampling error.

One issue the bootstrap cannot account for is the potential underrepresentation of disasters in our 1990-2019 sample. If more disasters had occurred, however, $f(R)$ would have had *more* probability mass in the far left tail, the difference between f^* and f in that region would have been *smaller*, and such states would have accounted for *less* of the equity premium. If anything, the low importance of states below -30% therefore reflects the opposite of a "Peso problem".

3.4 Economic Interpretation

We connect $EP(x)$ to economic primitives by examining it through the lens of asset pricing theories, focusing on models with representative agents as they provide an excellent setting to isolate contributing factors. Computational details are discussed in the supplement.

As we saw in Figure 16, the $EP(x)$ curve highlights significant discrepancies between sources of the equity premium in the data and in popular representative agent models. To understand which assumptions drive these discrepancies, it is useful to write the probability spread $f - f^*$ in (3.1) as $f(1 - f^*/f)$ in order to separate the quantity of risk for different return states (f) from their price of risk (f^*/f).[27] In the representative agent setting, state probabilities primarily reflect assumptions about the distribution of fundamental shocks, which are arguably secondary for the models' economic mechanisms. In contrast, risk prices reflect preferences and lie at the heart of their mechanisms. It is therefore important to ask whether the models fail to match $EP(x)$ due to state probabilities or risk prices.

Table 8 shows the equity premium contribution, probability, and price of risk of returns in the region $R \in [-30\%, -10\%]$. The $EP(R)$ column shows that these states contribute 2/3 of the equity premium in the data, but an order of magnitude less in

[27] Note that $f^*/f = \mathbb{E}[f_t^*]/\mathbb{E}[f_t]$ differs from $\mathbb{E}[f_t^*/f_t] = \mathbb{E}\left[\frac{\mathbb{E}_t[M_{t+1}|R_{t+1}]}{\mathbb{E}_t[M_{t+1}]}\right]$, the average of the pricing kernel's projection onto returns. Our metric has the advantage that it is straightforward to estimate empirically. However, we show in the supplement that it is nearly indistinguishable from the average pricing kernel for all models considered in Table 8. It is therefore appropriate to interpret f^*/f as measuring risk prices.

the models of Campbell and Cochrane (1999), Bansal and Yaron (2004), Barro (2009), and Wachter (2013). The reason is that the models understate both the quantity and price of risk for such returns, as shown in the last two columns. The probability of $R \in [-30\%, -10\%]$ equals 0.016 in the data but 0.013 or less in the four models, while the price of risk equals 2.627 in the data and 1.603 or less in the models.[28]

In representative agent models, the probability of stock market tail events is primarily driven by the distribution of fundamental shocks. Because this functional form assumption is secondary for the models' economic mechanism, however, it is important to ask whether it accounts for their inconsistency with $EP(x)$. We construct counterfactual $EP(x)$ curves by combining risk prices from the models with the empirical return distribution, i.e. we evaluate $EP(x)$ based on $f_{\text{data}}\,(1 - f^*_{\text{model}}/f_{\text{model}})$ instead of $f_{\text{model}} - f^*_{\text{model}}$. Because this experiment imposes that return probabilities are identical across models, and equal to the values from the data, it allows us to assess how much of the differences between $EP(x)$ curves is attributable to differences in the price of risk. The resulting curves in Figure 19 show that (i) risk prices are very similar across models but very different from the data and that (ii) even when combined with a realistic return distribution, the models attribute less than 1/4 of the equity premium to return states between -30% and -10%. The models' inconsistency with the data therefore stems from unrealistically small risk prices for stock market tail events. Importantly, this issue can not be resolved by simply increasing the coefficient of relative risk aversion because doing so would increase the price of risk globally and result in an unrealistically large equity premium.

[28]The Barro (2009) model assumes low risk aversion of 3 but nevertheless implies a relatively high price of risk of 1.603. The reason is that the model counterfactually implies that probability mass is concentrated around the lower end of the return interval $[-30\%, -10\%]$, where f^*/f is largest, whereas the data (and other models) imply that mass is concentrated around the upper end of $[-30\%, -10\%]$, where f^*/f is smallest.

One interpretation of this finding is that common utility functions are misspecified. This view finds support in literature on decision theory, which has documented numerous inconsistencies between the behavior implied by expected utility and observed decisions under uncertainty (see, e.g., Allais (1953), Kahneman and Tversky (1979), and Rabin (2000)), many of which concern attitudes towards extreme risks.[29] One response to these puzzles has been the development of disappointment aversion preferences, an axiomatic extension of expected utility theory that overweights left tail outcomes and rationalizes the Allais paradox (Gul (1991), Routledge and Zin (2010)). Schreindorfer (2020) shows that disappointment aversion allows a simple representative agent model to match many features of equity index options, in addition to traditional asset pricing moments. We find that returns between -30% and -10% contribute 81/100 of the equity premium and are associated with a price of risk of 3.6 in his model. Interestingly, disappointment aversion therefore not only rationalizes behavior that appears puzzling under expected utility at the micro level, but it also explains risk prices that otherwise appear puzzling at the macro level.

3.5 Last Thoughts

We use data from options markets to decompose risk premia into different parts of the return state space. In the data, the equity premium compensates investors primarily for shocks that coincide with monthly returns between -30% and -10%, which tend to occur at a business cycle frequency. This fact is quantitatively at odds with leading asset pricing theories, including models with external habits, long-run risks,

[29] Although many of these inconsistencies have been documented for power utility functions, Epstein-Zin and habit preference functions exhibit the same shortcomings as they relate to our context.

rare disasters, and incomplete markets, because they imply a counterfactually low price of risk for stock market tail events. We show that different assumptions about shocks to fundamentals do not alter this result. The most plausible explanation of our findings is therefore that standard utility functions imply too little aversion against tail events. As such, our market-based evidence agrees with well-known findings at the micro level. We believe that future generations of asset pricing models should pay close attention to the composition of risk premia, in addition to their level and predictability. $EP(x)$ provides a tool for doing so.

Figure 16. Sources of the equity premium.

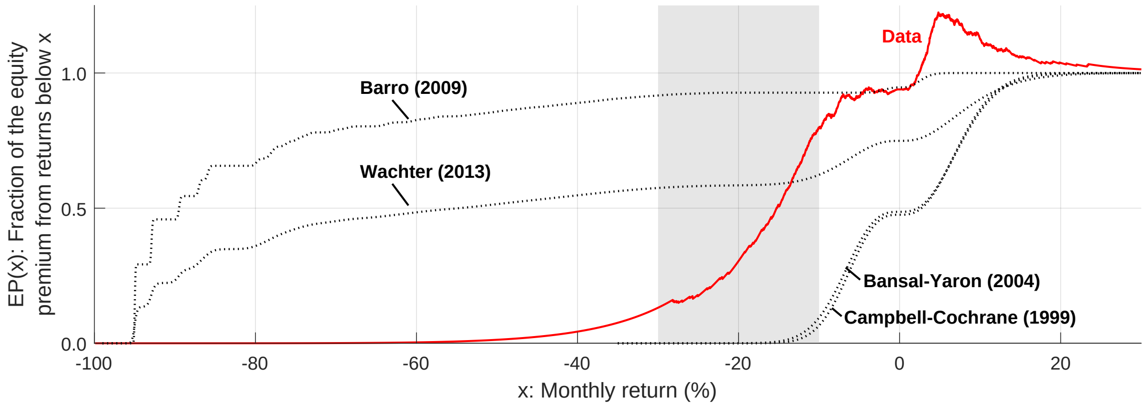

Note. Integral (3.1) is shown as a function of the upper limit of integration, x, and normalized by the total equity premium: $EP(x) \equiv \frac{\int_{-1}^{x} R[f(R)-f^*(R)]dR}{\mathbb{E}[R_{t+1}-R_t^f]}$. The shaded area marks monthly returns between -30% and -10%, a set of states that accounts for 2/3 of the equity premium in the 1990-2019 sample.

Figure 17. Average Arrow-Debreu payoffs and forward prices in 1990-2019 data.

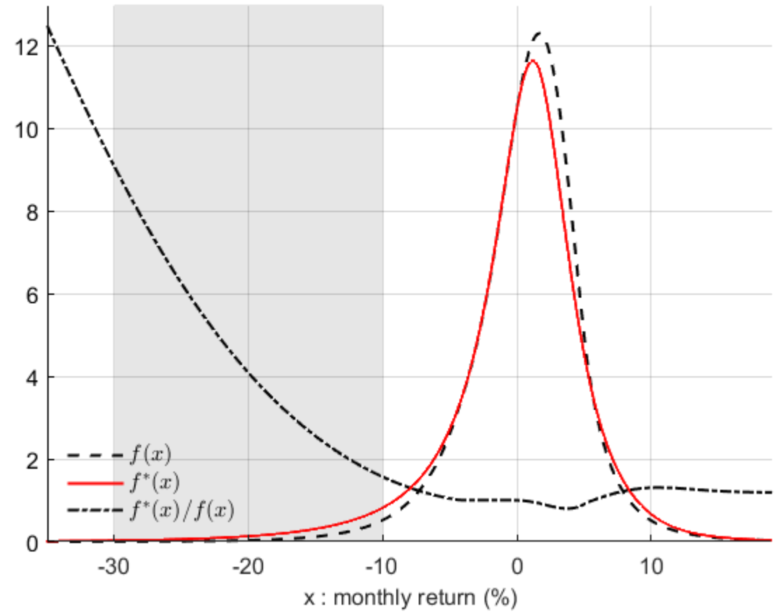

Figure 18. Sampling distribution of $EP(x)$.

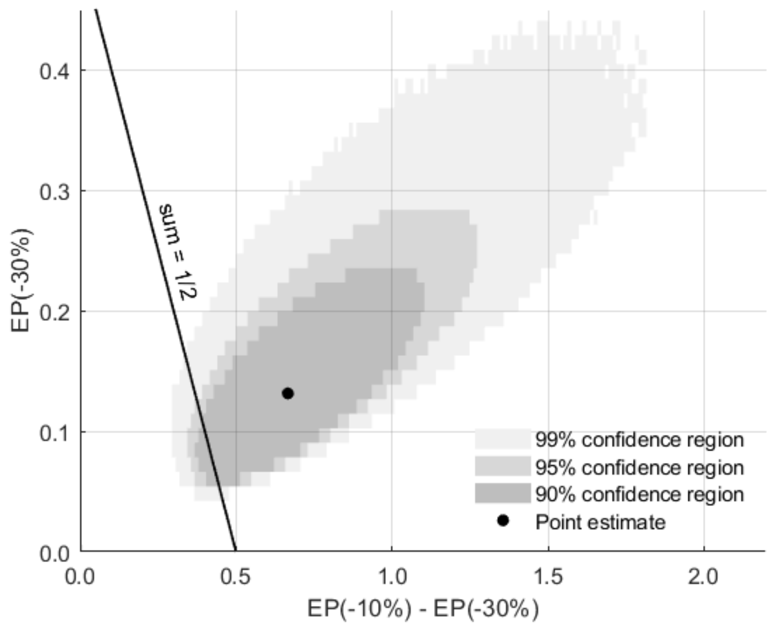

Note. For a partition of the return state space, the figure shows the joint sampling distribution of each region's contribution to the equity premium (as a fraction of the total equity premium). The distribution is constructed with a block bootstrap, using a block length of 21 trading days and 10 million bootstrap samples.

Figure 19. Counterfactual $EP(x)$.

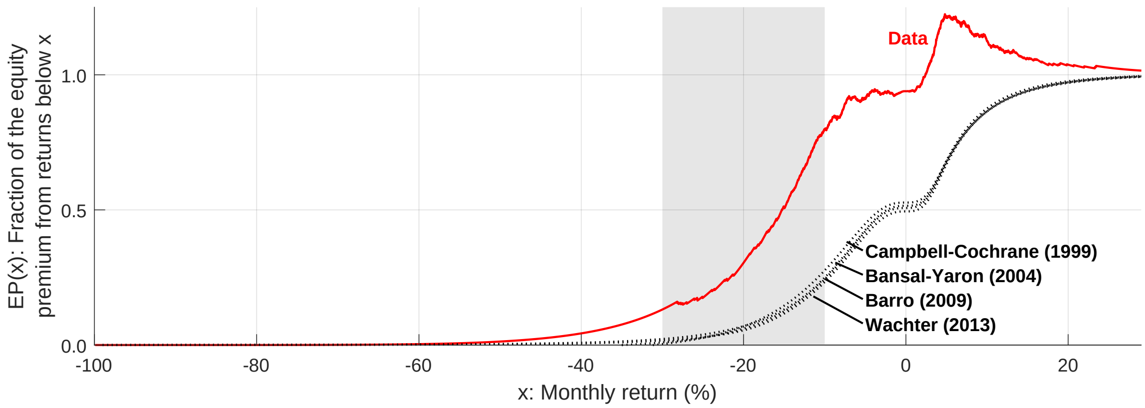

Note. We compute $EP(x)$ based on the empirical return distribution f_{data} and the price of risk implied by different models, $f^*_{\text{model}}/f_{\text{model}}$. Because the quantity of risk matches the data by construction, remaining gaps between the $EP(x)$ curves reflect differences in the price of risk.

Table 8. Characteristics of influential return states: $R \in [-30\%, -10\%]$.

Mechanism	Paper	$EP(R)$	$\int f(R)dR$	$\frac{\int f^*(R)dR}{\int f(R)dR}$
	Data, 1990-2019	0.666	0.016	2.627
Habits	Campbell-Cochrane (1999)	0.066	0.007	1.341
Long-run risks	Bansal-Yaron (2004)	0.095	0.013	1.280
Disasters	Barro (2009)	0.010	0.000	1.603
	Wachter (2013)	0.049	0.012	1.207
Disap. aversion	Schreindorfer (2020)	0.811	0.014	3.648

Note. $EP(R)$: fraction of the equity premium associated with monthly returns between -30% and -10%; $\int f(R)dR$: probability of such returns; $\frac{\int f^*(R)dR}{\int f(R)dR}$: ratio of their average risk-neutral to historical probability.

Chapter 4

HETEROGENEOUS LABOR INCOME RISK AND HOUSEHOLD PORTFOLIO CHOICE

4.1 Introduction

Modern portfolio choice theory predicts that human capital represents an important determinant of households' optimal portfolio allocations. In the canonical model of life cycle portfolio choice in the presence of uninsurable labor income risk, where labor income modeled as a Gaussian Markov process, young households optimally allocate most of their liquid wealth to risky assets because human capital behaves like a bond. This prediction is at odds with empirical observations on household portfolio allocations. Additionally, as shown in the left panel of Figure 20, the model produces very little heterogeneity in portfolio shares among households with similar age. For comparison, the right panel of Figure 20 shows that risky asset portfolio shares computed from the Survey of Consumer Finances data are nearly uniformly distributed from 0% to 100% at each stage of the life cycle.[30]

In this paper, I show that several of the counter-factual predictions of the standard life cycle portfolio choice model may simply be a result of mis-specified income dynamics. I document that there exists a tight link between the shortcomings of the standard model and the labor income process employed. After modifying the income process to one that more closely replicates observed income dynamics, some of the

[30]Details on the model specification and simulation procedure, along with data construction, are provided in later sections.

counter-factual portfolio choice predictions of the standard model disappear or are attenuated.

Using a large panel of U.S. administrative data, Guvenen et al. (2016) document new facts about income growth. They show that a Gaussian Markov process, the same type that is used in the standard model of household portfolio choice, does not properly capture the amount of dispersion or skewness in the conditional earnings growth distribution, the variation in those statistics over the life cycle, the way those statistics interact with income levels, or the way those statistics differ across households. They estimate a new flexible stochastic process that is able to closely reproduce these key properties, which are observable in the data. The portfolio choice implications of this new model of income dynamics have yet to be documented, and I find that they are potentially substantial.

Because the standard Gaussian Markov process for income misses important features of actual income dynamics, the risk profile of human capital in the standard model of life cycle portfolio choice is also mis-characterized. I embed the stochastic income process from Guvenen et al. (2016), which they call the Heterogeneous Income Profile (HIP) model, into an otherwise standard model of life cycle portfolio choice and demonstrate that my model, which I will also refer to as the HIP model, displays predictions that are much more consistent with observed household investment behavior. The HIP model of life cycle portfolio choice predicts that households aged 20 to 30 will optimally hold anything from 20% to 100% of their portfolio in risky assets, which is nearly the dispersion that we observe in the data. Furthermore, the HIP model predicts that the average young household invests only slightly above half of its financial wealth in risky assets.

The basic intuition behind these results is that different households have different income risk profiles. The human capital of a futures trader has different characteristics than the human capital of an elementary school teacher. The findings of Guvenen et al. (2016) imply that both the level and the risk of human capital vary a lot more over the life cycle and in the cross-section of households than previously thought, and the HIP model shows how this heterogeneity can produce quantitatively large and realistic differences in household risky asset shares.

The conclusions from this paper complement Danthine, Donaldson, and Mehra (2017) in providing support for personalized investment advice. While the current industry trend is for robo-advisors and one-size-fits-all portfolio advice, this paper shows the importance of customized solutions based on individual situations.

Section 2 discusses related papers, both theoretical and empirical, in the life cycle portfolio choice literature. Section 3 describes the HIP life cycle portfolio choice model and how it nests the standard model. Section 4 explores the predictions from the model and some welfare consequences of following suboptimal portfolio advice. Section 5 concludes.

4.2 Related Literature

The problem of a life cycle investor has been studied extensively. While early efforts, such as Merton (1969, 1971) and Bodie, Merton, and Samuelson (1992), focused on complete markets, more recent literature has focused on incomplete markets and highlighted the importance of uninsurable labor income risk and its portfolio choice implications. Much of the work, normative or positive, has focused on average outcomes over the life cycle. I demonstrate that predictions about average outcomes in these models are related to their predictions about heterogeneity.

The standard way that risky labor income is modeled in the household finance literature is as a first-order autoregressive Markov process with individual-specific innovations being drawn from a log-normal distribution common to all households.[31] The process is typically calibrated to match the average income path over the life cycle, along with some moments relating to the dispersion in income growth rates in the cross-section. The portfolio choice implications of modeling income in such a way are well-documented. Campbell et al. (2001) and Cocco, Gomes, and Maenhout (2005) both attempt to quantitatively explain observed portfolio shares by incorporating only income risk and uncertain lifespan. Gomes and Michaelides (2003) use the same labor income process to study the effects of participation costs and internal habit preferences with limited success in explaining the data. Gomes and Michaelides (2005) extend the standard model and are able to replicate the basic facts about stock market participation and portfolio allocations by appealing to preference heterogeneity in risk aversion and EIS. Cocco (2004) uses a similar labor income process and includes housing. Heaton and Lucas (2000) examines a more general notion of background risk, but continues to use a similar log-normal specification. The same style of income process is used to study precautionary savings motives in Hubbard, Skinner, and Zeldes (1995) and Gourinchas and Parker (2002).

My results add to the existing literature on household portfolio heterogeneity. Carroll (2000) documents that conditional risky asset shares are increasing in wealth. Wachter and Yogo (2010) introduce non-homothetic preferences over basic and luxury goods to generate an increasing relationship between wealth and portfolio shares, while Campanale, Fugazza, and Gomes (2015) show that concerns over asset liquidity

[31] An exception is Benzoni, Collin-Dufresne, and Goldstein (2007) where they model stock and labor markets as being co-integrated.

and transaction costs can produce the same increasing relationship. Both papers continue to use the standard Gaussian process for labor income. Thus, these papers produce portfolio heterogeneity in cross-section of households that is not due to the income risk heterogeneity mechanism used here. Importantly, the heterogeneous labor income risk that I exploit is directly observable in the data.

There is a large literature on the measurement and modeling of labor income. The strand of literature most relevant for this particular paper are the papers about income heterogeneity. Guvenen (2007) uses the Panel Survey of Income Dynamics to estimate a labor income process that incorporates more heterogeneity than the standard Gaussian process. Guvenen (2009) provides further support for heterogeneous income profiles. Using U.S. administrative data, Guvenen et al. (2016) show that labor income risk varies over the life cycle *and* with the level of recent earnings. The log-normal model is too restricted to reproduce these facts. In its place, the authors estimate a new stochastic process that produces the income heterogeneity and fat-tailed income growth distributions seen in the data.

There is a growing literature on household investment mistakes (Christelis, Jappelli, and Padula 2010; Goetzmann and Kumar 2008) and financial literacy (Van Rooij, Lusardi, and Alessie 2011, 2012). I view this literature as complementary to the analysis done here. In particular, it is possible for investors to find an optimal trade-off between risky and safe assets, while still being subject to behavioral biases when investing their equity portfolio.

Additionally, it need not be important that investors understand exactly the complexity of their future income stream, but only that they understand how it differs from certain reference income streams. Chang, Hong, and Karabarbounis (2018) show that learning about the labor market can also have significant impact on household portfolio allocations. An interesting avenue for future research would be to combine the mechanisms of their paper with the labor income process used in the model presented in this paper.

Empirically, labor income risk has been shown to be a determinant in household portfolio choice. Guiso, Jappelli, and Terlizzese (1996) find a small negative effect of income risk on risky asset portfolio shares in a sample of Italian households. Angerer and Lam (2009) also find a small negative effect of income risk on risky asset portfolio shares in the National Longitudinal Survey of Youth 1979 Cohort. Fagereng, Guiso, and Pistaferri (2017) find sizable negative marginal effects of income risk on portfolio choice, but also document that there exists considerable firm-worker insurance so that the overall size of the effect is much smaller in magnitude. While the empirical results lend support to the basic intuition that human capital is important for portfolio choice, the standard quantitative models of life cycle portfolio choice do not immediately appear to be consistent with this evidence. I show that this inconsistency arises because of the chosen labor income process. In the next section, I will develop a model that showcases a strong link between labor income risk and risky asset portfolio shares.

4.3 The HIP Model

To understand the effects of heterogeneous income risk on portfolio choice, I modify the workhorse life cycle portfolio choice model of Cocco, Gomes, and Maenhout (2005) to include a stochastic process estimated to match the high amount of heterogeneity in income growth from Guvenen et al. (2016). Apart from slight differences in calibration and notation, the model is otherwise equivalent to the benchmark model developed in Cocco, Gomes, and Maenhout (2005) with a new income process.

Let risky labor income Y_t^i for household i of normalized age t be the composition of (i) a deterministic function of age $g(t)$ (to capture a common age trend in income), (ii) an AR(1) stochastic process z_t^i with persistence parameter ρ_z, (iii) an unemployment shock ν_t^i, and (iv) household-specific income effects (α^i, β^i).[32] The innovations to the persistent component of income, z_t^i, are drawn from a mixture of two distributions with mixing probabilities that depend on normalized age and z_{t-1}^i, where one of the distributions is degenerate and the other is Normal with household specific standard deviation σ_z^i.

$$Y_t^i = (1-\nu_t^i)\exp(g(t) + \alpha^i + \beta^i t + z_t^i) \tag{4.1}$$

$$z_t^i = \rho_z z_{t-1}^i + \eta_t^i \tag{4.2}$$

$$\eta_t^i \sim \begin{cases} -p_{z,t}^i \mu_z & \text{with pr. } 1-p_{z,t}^i \\ N((1-p_{z,t}^i)\mu_z, \sigma_z^i) & \text{with pr. } p_{z,t}^i \end{cases} \tag{4.3}$$

The nonemployment shock ν_t^i occurs with probabilities that depend on normalized age and z_t^i. The nonemployment shock is drawn from a truncated Exponential distribution, so it is possible for households to have either full or partial unemployment in a year.

$$\nu_t^i \sim \begin{cases} 0 & \text{with pr. } 1-p_{\nu,t}^i \\ \min\{1, \text{Expon.}(\lambda)\} & \text{with pr. } p_{\nu,t}^i \end{cases} \tag{4.4}$$

[32]Normalized age is defined as $t = \frac{age-24}{10}$. The function $g(t)$ is a polynomial of order 4 in normalized age, estimated in Guvenen et al. (2016). $g(t) = 9.4895 + 1.222t - 0.6799t^2 + 0.2013t^3 - 0.0240t^4$.

The mixture probabilities for both innovations are logistic function transformations of linear equations with inputs of normalized age (t) and a value of the income shock (z).

$$p^i_{z,t} = \frac{\exp(\xi^i_{z,t-1})}{1+\exp(\xi^i_{z,t-1})}, \qquad p^i_{\nu,t} = \frac{\exp(\xi^i_{\nu,t})}{1+\exp(\xi^i_{\nu,t})} \tag{4.5}$$

$$\xi^i_{j,t} = a_j + b_j t + c_j z^i_t + d_j z^i_t t \tag{4.6}$$

The household-specific effects (α^i, β^i) are distributed in the cross-section of households according to a Normal distribution with zero mean and covariance matrix Σ. The household-specific income shock volatility parameter is distributed in the cross-section of households according to a log-Normal distribution with variance $\tilde{\sigma}^2_z$.

$$\begin{aligned}(\alpha^i, \beta^i) &\sim MVN(0, \Sigma), \qquad \Sigma = \left[\sigma^2_\alpha,\ \sigma_{\alpha\beta};\ \sigma_{\alpha\beta},\ \sigma^2_\beta\right] \\ \ln \sigma^i_z &\sim N(\ln \bar{\sigma}_z - \frac{\tilde{\sigma}^2_z}{2}, \tilde{\sigma}^2_z)\end{aligned}$$

Each household begins working at age 20 and its income is governed by the process above until retirement at age 65. During retirement, the agent receives a fraction Λ of his last working-year of potential income (that is, income before any unemployment) in each period for the remainder of his lifetime. Following Hubbard, Skinner, and Zeldes (1995), agents face some probability of death each period. I will use $\tilde{p}_t$ to denote the conditional probability that an agent alive at date t continues to be alive for time $t+1$.[33] It is not possible for households to live beyond age 100 and there are no bequests.

There is a riskfree asset paying R_f in each period and there is a risky asset paying $R_t = R_f + \pi_t$ in each period, where the excess return π_t is i.i.d. and Normally

[33] I will slightly abuse time notation, so that t at times refers to a households (normalized) age and at other times refers to a model period, such as year t. Because this is a partial equilibrium model, the confusion will hopefully remain contained, as I will never need to refer to a household aged t in year t – the concepts are one-and-the-same in this environment.

distributed with mean μ_π and variance σ_π^2. I do not model correlation between innovations to the income process (η_t^i, ν_t^i) and asset returns (π_t). The choice of no correlation follows in the footsteps of Cocco, Gomes, and Maenhout (2005) and is also motivated by the recent findings of Guvenen et al. (2017), who show that correlations between household earnings growth and stock market returns, although exhibiting heterogeneity, are generally negligible in magnitude.

Households have constant relative risk aversion utility functions, with parameter γ, for which they maximize expected lifetime utility by choosing in each period how much wealth to consume and how to allocate savings between the risky and riskless asset. Households discount future values using discount factor δ. The timing is as follows: (i) a household enters the period with some level of wealth W_t^i; (ii) income for the period, Y_t^i, is realized; (iii) the household makes the consumption and portfolio choice decisions. For notational ease, define cash on-hand $X_t^i = W_t^i + Y_t^i$ and the adjusted discount factor $\tilde{\delta}_t = \delta \tilde{p}_t$. Then, a bit more formally, a household having already received its income for the period solves the following problem:

$$V_t^i(X_t^i, z_t^i) = \max_{C_t^i, \theta_t^i} \; u(C_t^i) + \tilde{\delta}_t E_t[V_{t+1}^i(X_{t+1}^i, z_{t+1}^i)|z_t^i] \tag{4.7}$$

such that

$$X_{t+1}^i = Y_{t+1}^i + (X_t^i - C_t^i)(R_f + \theta_t^i \pi_{t+1}) \tag{4.8}$$

$$C_t^i \in [0, X_t^i], \quad \theta_t^i \in [0,1]$$

$$Y_{t+1}^i \text{ as above}$$

where θ_t^i is the fraction of the portfolio allocated to the risky asset. The constraint $\theta_t^i \in [0,1]$ amounts to a no short-sale constraint on both the risky and risk-free asset. The constraint $C_t^i \in [0, X_t^i]$ represents that households cannot borrow against their labor income stream.

The canonical model of household life cycle portfolio choice in the presence of uninsurable labor income risk, where z_t^i takes the form of a permanent shock Gaussian process, is nested in the above specification by setting $p_{\nu,t}^i \equiv 0$, $p_{z,t}^i \equiv 1$, $\rho_z \equiv 1$, $\tilde{\sigma}_z \equiv 0$, $\Sigma = \mathbf{0}_{2\times 2}$, and appropriately adjusting the other calibrated parameters.

4.3.1 Solution and Simulation

A solution to the model consists of policy functions $C_t^{i*}(X_t^i, z_t^i)$ and $\theta_t^{i*}(X_t^i, z_t^i)$. For notational convenience, I will temporarily drop household-specific indexing and functional dependence on state variables from these policy functions. The first order conditions for C_t^* and θ_t^* are

$$u'(C_t) = \tilde{\delta}_t E_t[V'(X_{t+1}, z_{t+1})(R_f + \theta_t \pi_{t+1})|z_t] \tag{4.9}$$

$$0 = \tilde{\delta}_t E_t[V'(X_{t+1}, z_{t+1})(W_t - C_t)\pi_{t+1}|z_t] \tag{4.10}$$

respectively. After recasting the problem in terms of end-of-period savings, I solve the model numerically using the endogenous grid method (Carroll 2006). Additional details on numerical computation can be found in Appendix C.

The model is calibrated to match observed life cycle wealth accumulation and portfolio choice at an annual frequency. The parameters of the model are visible in Tables 9 and 10. There are many parameters in the model, but most of them are embedded in the exogenously specified income process, which I take as given because it was estimated precisely in Guvenen et al. (2016).[34] Therefore, the parameters of interest are risk aversion, which I set to 9, and the asset returns. I set the risk-free

[34] I had to estimate the parameters $(a_\nu, b_\nu, c_\nu, d_\nu)$ using additionally reported statistics from their paper because the authors mistakenly do not report these parameters in their calibration. I am in correspondence with the authors to obtain the precise parameter values.

rate to a level of 2% and I specify an equity premium of 4% and risky asset standard deviation of 18%. It is typical in this branch of literature to use high values of risk aversion (Cocco, Gomes, and Maenhout (2005) use $\gamma = 10$) and a value for the equity premium below the historical average of 6% that Mehra and Prescott (2003) find (Cocco, Gomes, and Maenhout (2005) use 4%).

For the standard model, I use the Gaussian calibration from Guvenen et al. (2016), visible in table 11. This particular calibration is slightly different from that presented in Cocco, Gomes, and Maenhout (2005), but the portfolio choice implications are the same.

It is infeasible to solve the HIP model for all combinations of the household-specific parameters. Therefore, I discretize the joint distribution of $(\alpha^i, \beta^i, \sigma_z^i)$ into a finite set of household *types*. I use 4 distinct (α^k, β^k) types and 3 distinct σ_z^k types, for a total of 12 household types. I then solve and simulate the model for only these 12 types of households.

After solving the model, I simulate a large panel of households starting at age 20 until death. For the HIP model, the panel consists of households from a number of different income types. I use the quadrature weights from the discretization of household types to determine how many households of each type to simulate. From the simulated panel of households, I compute averages and cross-sectional percentiles of household risky asset shares of wealth at every age, conditional on survival.

4.4 Results

Figure 21 displays the 10^{th}, 25^{th}, 50^{th}, 75^{th}, and 90^{th} percentiles of the cross-sectional distribution of risky asset portfolio shares of financial wealth by age from the

HIP model. Contrary to the standard model with a Gaussian labor income process, the HIP model generates large heterogeneity in portfolio choice for households over most of the working life. A household aged between 20-30 at the 90^{th} percentile of the cross-sectional distribution of portfolio allocations holds of their portfolio in risky assets, which is more-or-less in line with the standard model's prediction. However, the median young household allocates between 60% and 70% of their portfolio to risky assets, and a young household at the 10^{th} percentile allocates only about a quarter of wealth to risky assets.

The spread between the 90^{th} and 10^{th} percentiles of the distribution in the HIP model is 75% for young households, which is comparable to the 70-80% that we see in the SCF data. For the Gaussian model, the spread is 0% for the youngest households and moves to 20-30%.

Another notable outcome of the HIP model is that because many young households choose to hold much safer portfolios, the average age profile of risky shares is not as sharply decreasing at it is in the standard model. The average age profile in the data is nearly flat, but in the standard life cycle portfolio choice model that average age profile is steeply decreasing over the working life. Figure 22 plots the average age profile of risky asset shares for the standard Gaussian model, the HIP model, and the data. While neither model matches the SCF data particularly well, it is clear that the HIP model deviates much less from the data. Thus, matching the cross-section of household portfolios by age lead to dramatic improvements in the more traditional target, the average household portfolio, as we would expect.

The differences between the standard model and the HIP model are largest precisely when theory tells us they should be – when human capital hedging demands are strongest. At the end of the working life and into retirement, the two models have nearly identical predictions for the average portfolio share and the heterogeneity in the cross-section. This is to be expected. The effect of labor income risk on portfolio choice during these years is zero or extremely minor. Because the two models only differ in the income process considered, they are identical following retirement, where all differences portfolio shares must come from wealth effects.

A second lens through which to view the results is as an answer to the question: "How much heterogeneity in household portfolio shares is due to differences in labor income risk?" For young households, according to the HIP model, heterogeneity in labor income risk could plausibly account for a significant amount of the observed heterogeneity in portfolio allocations. However, as households enter their prime working years, differences in income risk begin to contribute very little to differences in risky asset portfolio shares. Therefore, future work could focus on heterogeneity in major determinants of portfolio choice that might begin to outweigh labor income risk for middle-aged households, such as differences in bequest motives, retirement benefits, and health.

There are two reasons that using the HIP income model generates substantially different portfolio choice predictions than the Gaussian model. The first is that the Gaussian model understates the labor income risk that working households face. As shown in Guvenen et al. (2016), the conditional distribution of earnings growth is fat-tailed and negatively skewed. All else equal, a household facing a fat-tailed and negatively skewed distribution of income growth would optimally choose a lower risky asset share than a household facing a thin-tailed and symmetric distribution of income growth.

The second reason is that the Gaussian model does not generate the amount of heterogeneity in labor income risk that exists in the cross-section of households. Some households have stable careers, other households have volatile income streams. The Gaussian model does not capture these differences. In a sense, the Gaussian model only generates heterogeneity in levels of income, not income risk. That is because if you consider the variance of the log of income, it is the same across all households regardless of income level or age. This is not the case in the HIP model, where differences in income risk still exist after conditioning on income level and age, even within the different income types. Thus, the HIP model provides both ex ante and ex post heterogeneity in labor income risk, the Gaussian model essentially has neither.

The HIP model is able to generate these results without relying on other plausible mechanisms that have been put forth in the literature. Guvenen, Ozkan, and Song (2014) document the cyclicality of earnings risk. Additionally, Davis and Willen (2000) and Guvenen et al. (2017) study the covariance of income shocks and asset returns. In the HIP model presented here, all shocks are independent. Therefore, the systematic components common to many other papers are absent in this model. If those channels were to be included, it may be the case that for some households income risk is substantially higher than depicted here, which would have even stronger portfolio choice implications. These are questions I leave for future research.

Portfolio choice, although the main focus of this paper, is not the only important aspect of these models. Income, consumption, and wealth are also important to examine. Income dynamics in these models are exogenous, and in this case it is by design that the HIP model more closely replicates the important features of income dynamics. Consumption and wealth are also choice variables for the household that are outcomes of the model solution. Just because a model performs well along one dimension, does not mean it will fit the data in all dimensions.

Figure 23 shows the average age profile of wealth for the two models under consideration alongside the SCF data. As is well-known, observed wealth patterns do not tend to decline to low levels after retirement, but rather they stay relatively stable. Because the models in this paper do not include bequest motives, it is optimal for households to drawdown wealth and have no wealth at the end of life. During the working years, wealth accumulation for both models happens too quickly compared to the data. Additionally, households in both models accumulate too much wealth. In Appendix C, I introduce an income tax to show that wealth accumulation levels in each model can be made to be mostly consistent with the data.

Unconditionally, both models have trouble producing a wealth distribution that looks very much like the one in the SCF data. Table 12 shows the percentiles from the unconditional wealth distribution for both models and the data. There are not nearly enough low wealth households or extremely wealthy households in the simulated panels from each model. However, the HIP model generates more inequality than the Gaussian model. In the HIP economy, the $P90/P10$ ratio is 48, meaning the household at the 90^{th} percentile has 48 times more wealth than the household at the 10^{th} percentile. In the Gaussian model, the same ratio is only 33. In the SCF data, it is 95. Because wealth also plays a role in the portfolio choice decision, the fact that the HIP model generates more wealth inequality than the Gaussian model is another factor that leads it to generate more portfolio heterogeneity in the cross-section.

4.4.1 Sources of Heterogeneity

There are several sources of heterogeneity in the HIP model. There are ex ante differences in the slope of the income profile, ex ante differences in the risk associated with innovations to income, and ex post differences because households receive different income shocks. In the standard Gaussian model, only the ex post heterogeneity arises. While Guvenen et al. (2016) stress the importance of each source in matching observed income dynamics, it is useful to see which sources of heterogeneity are most important for matching the various aspects of portfolio choice that are the focus of this paper.

Figure 24 shows the cross-section of household portfolio shares computed from two different specifications of the main HIP portfolio choice model. The left panel shuts down heterogeneity in the slope of the income profile. Comparing the left panel of Figure 24 to Figure 21, it is clear that there are essentially no implications for portfolio choice in this scenario.

On the other hand, the right panel of Figure 24 shuts down ex ante heterogeneity in the volatility of the labor income process. After turning this feature of the income process off, the model produces very little cross-sectional dispersion in household portfolios over the life cycle. Thus, it is the case in the HIP model that heterogeneity in risky asset portfolio shares comes largely from heterogeneity in labor income risk.

An additional observation from the right panel of Figure 24 is that the age profile of risky asset shares remains mostly as it is in the full HIP model. This observation underscores the fact that realistic variation in household labor income risk over the life cycle is important in generating realistic portfolio allocation behavior.

4.5 Conclusion

Campbell (2006) emphasizes the importance of having both normative and positive research in household finance. It is widely believed that the standard model of life cycle portfolio choice in the presence of uninsurable labor income risk, a normative theory, is inconsistent with observed household investment behavior. In this paper, I show that correcting mis-specifications in one of the most important features of the existing models, risky labor income, leads to significantly different predictions about household portfolio allocation decisions. My results suggest that the gap between normative and positive household finance with respect to portfolio choice may not be as big as previously thought. Even if one takes the view that these theories should only be normative in nature, the HIP model is still informative for wealth managers providing financial planning advice.

The heterogeneous income profile (HIP) model of life cycle portfolio choice is consistent with observed income dynamics and has portfolio choice predictions that more closely align with those seen in the Survey of Consumer Finances data. The HIP model predicts a large amount of cross-sectional heterogeneity in risky asset portfolio shares for young households and an age profile of average risky asset shares that is only slightly downward sloping – both being dimensions along which the standard model fails to match the data.

My findings have implications for much of the previously published work in household portfolio choice. Since nearly all extensions of the standard model do not change the income process, it is likely that most existing household life cycle portfolio choice models misrepresent the characteristics of human capital risk in similar ways as the standard model. Thus, it may be important to revisit previous extensions and ask how, or if, heterogeneous labor income risk impacts the conclusions. I leave these tasks for future research.

Figure 20. Cross-section of household risky asset shares from the standard model (left) and the data (right).

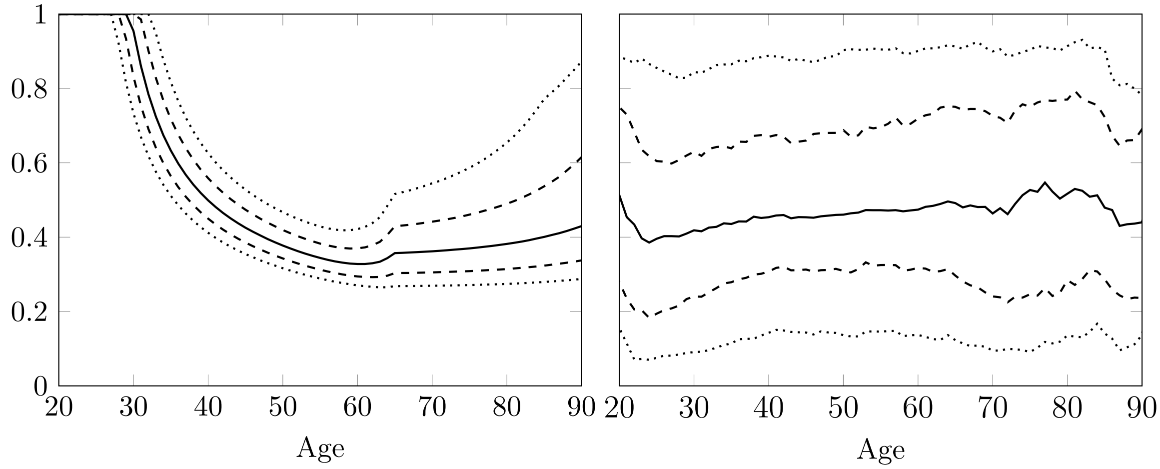

Note. Median (solid), 75th and 25th (dashed), 90th and 10th percentiles (dotted). For the left figure, I simulate a large panel of households from the standard life cycle portfolio choice model and compute the plotted statistics conditional on age. For the right figure, I compute each statistic using sampling weights within each of the 5 replicates for each survey year. I then average across the 5 replicates for a given survey year and then average across all surveys. For each age t, I group households aged $t-2$ to $t+2$ in order to increase sample size and smooth estimates. Data is the Survey of Consumer Finances from 1989 to 2016 and is conditional on participating in risky asset markets. Additional details available in Appendix C..

Figure 21. Cross-section of simulated household risky asset shares from the HIP model.

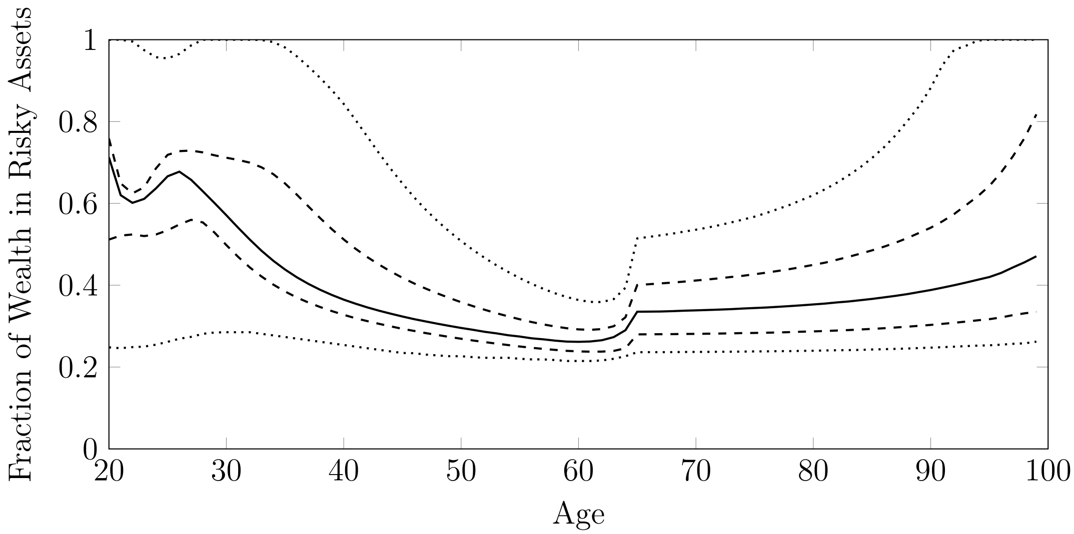

Note. Median (solid), 75th and 25th (dashed), 90th and 10th percentiles (dotted). I simulate a large panel of households from the HIP life cycle portfolio choice model and compute the plotted statistics conditional on age.

Figure 22. Mean risky asset share for the two models and the SCF data.

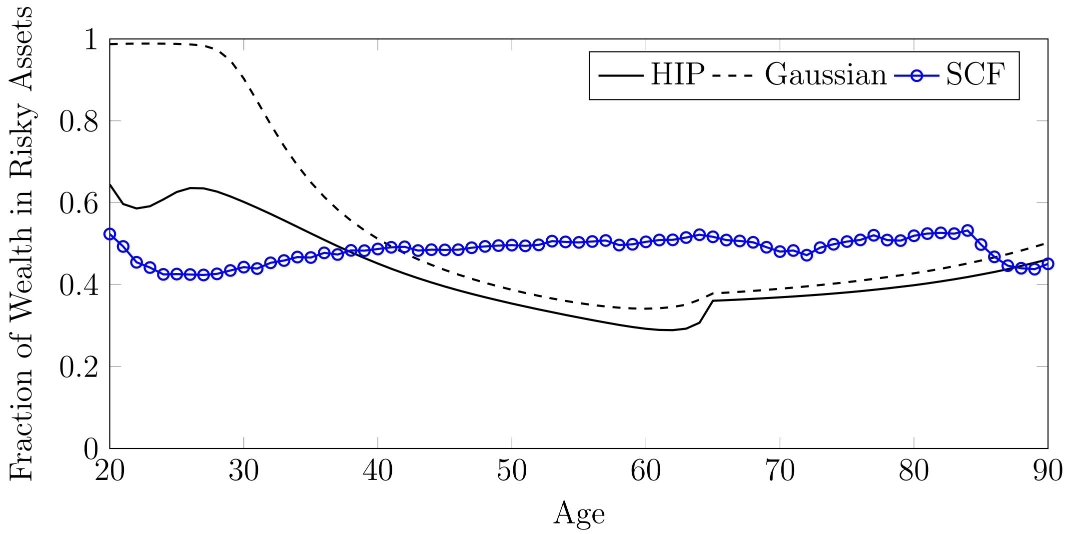

Note. Average risky asset portfolio share by age from the HIP model (solid line), standard model (dashed line), and the Survey of Consumer Finances data (blue marked line). For the SCF, I compute the mean by age using sampling weights within each of the 5 replicates for each survey year. I then average across the 5 replicates for a given survey year and then average across all surveys. For each age t, I group households aged $t-2$ to $t+2$ in order to increase sample size. Data is the Survey of Consumer Finances from 1989 to 2016. Additional details available in Appendix C.

Figure 23. Mean wealth for the two models and the SCF data.

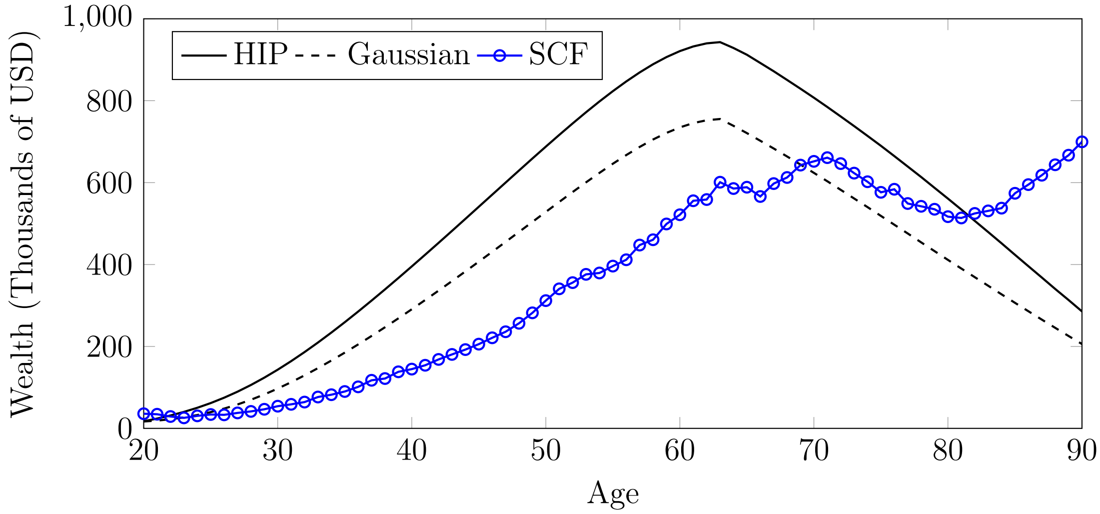

Note. Household wealth by age from the HIP model (solid line), standard model (dashed line), and the Survey of Consumer Finances data (blue marked line). For the SCF, I compute the mean by age using sampling weights within each of the 5 replicates for each survey year. I then average across the 5 replicates for a given survey year and then average across all surveys. For each age t, I group households aged $t-2$ to $t+2$ in order to increase sample size. Data is the Survey of Consumer Finances from 1989 to 2016. Additional details available in Appendix C.

Figure 24. Cross-section of household risky asset shares from different model specifications.

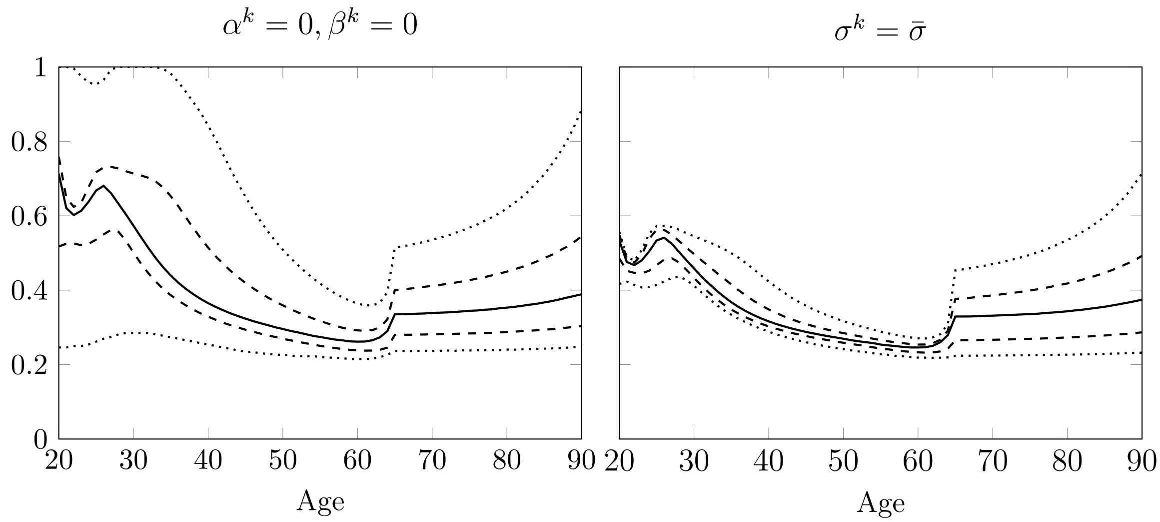

Note. Median (solid), 75th and 25th (dashed), 90th and 10th percentiles (dotted). I simulate a large panel of households from each model and compute the plotted statistics conditional on age. For the left figure, I turn off the ex ante heterogeneity in the slope of household income profiles. For the right figure, I turn off the ex ante heterogeneity in the volatility of the income shock process.

Table 9. Lifecycle Model parameters

Parameter	Symbol	Value
Risk aversion	γ	9
Discount factor	δ	0.96
Gross risk-free rate	R_f	1.02
Equity premium	μ_π	0.04
Standard deviation of excess stock returns	σ_π	0.18
Retirement income replacement rate	Λ	0.681

Note. These parameters are common to all households in both the standard model and the HIP model.

Table 10. HIP income process parameters.

This calibration is taken from Guvenen et al. (2016), specification #2.

Parameter	Symbol	Value
Std of α	σ_α	0.477
Std of β $\times 10$	$\sigma_\beta \times 10$	0.208
Correlation of α & β	$corr_{\alpha,\beta}$	0.307
Persistence of z	ρ_z	0.459
Mean of z	μ_z	-0.599
	$\bar{\sigma}_z$	0.417
	$\tilde{\sigma}_z$	0.764
	$\sigma_{z,0}$	0.426
Jump intensity	λ	0.075
Coefficients in shock probabilities		
	a_z	-0.810
	b_z	0.765
	c_z	-1.682
	d_z	-1.288
	a_ν	-5.353
	b_ν	0.384
	c_ν	-4.228
	d_ν	1.240

Table 11. Gaussian income process parameters.

This calibration is taken from Guvenen et al. (2016), specification #7.

Parameter	Symbol	Value
Std of α	σ_α	0.372
Persistence of z	ρ_z	1.0
Mean of z	μ_z	0.0
	$\bar{\sigma}_z$	0.156
	$\tilde{\sigma}_z$	0.0
	$\sigma_{z,0}$	0.354

Table 12. Percentiles of the unconditional wealth distribution.

Wealth in thousands of US Dollars. For the SCF, I compute each statistic using sampling weights within each of the 5 replicates for each survey year. I then average across the 5 replicates for a given survey year and then average across all surveys. Data is the Survey of Consumer Finances from 1989 to 2016. Additional details available in Appendix C.

	$P10$	$P25$	$P50$	$P75$	$P90$
SCF Data	5.7	17.6	59.4	189.9	539.3
Gaussian Model	28.0	87.5	233.3	503.0	918.5
HIP Model	27.3	89.6	302.4	680.4	1305.6

Chapter 5

CONCLUSION

Models are necessarily abstractions from reality. Nonetheless, quantitative models are most useful when they represent an accurate depiction of reality. Therefore, researchers ask models to be consistent with a broad array of empirical facts. Importantly, these facts represent a small subset of the possible features of reality they could have asked to have their model reproduce. It is therefore imperative to assess models within the context of their intended use.

Many asset pricing models seek to explain the equity premium puzzle. In the first two essays of this dissertation, I re-examine a number of leading asset pricing models in the wake of new and important empirical evidence with which these models should be consistent. Because the equity premium involves the return on a claim to the aggregate dividend stream, generating an equity premium based on dividend growth processes that mischaracterize its largest sources of variation could lead to fragile conclusions. Therefore, it is natural to request that models reproduce the robust fact that dividend risk is concentrated at short horizons. Furthermore, a complete explanation of the equity premium must be consistent with not only its average level, but also its average decomposition, such as the one developed in the second essay of this document.

To the extent that models are inconsistent with this new evidence, that provides economists with guidance on how to proceed with future research. For example, the first essay introduces an economically-motivated way to model the transitory component in aggregate dividends, but it is only one out of an infinite number of possible specifications. Are there better ways to capture that component? Does including additional information, such as output growth, add value? In a similar spirit, recalling the second essay, what is the relationship between the equity premium decomposition in the return space that we provide and the decomposition in the horizon dimension as documented by Van Binsbergen, Brandt, and Koijen (2012)? Similarly, the third essay documents levels of dispersion in household portfolios far beyond what can be achieved with simple lifecycle portfolio choice models. Is this a failure of the models or a peculiarity of the data? I intend to explore these questions, and more, in future work.

REFERENCES

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

———. 1999. "Risk premia and term premia in general equilibrium." *Journal of Monetary Economics* 43 (1): 3–33.

Aguiar, Mark, and Gita Gopinath. 2007. "Emerging market business cycles: The cycle is the trend." *Journal of political Economy* 115 (1): 69–102.

Allais, Maurice. 1953. "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine." *Econometrica* 21 (4): 503–546.

Ang, Andrew, and Jun Liu. 2007. "Risk, return, and dividends." *Journal of Financial Economics* 85 (1): 1–38.

Angerer, Xiaohong, and Pok-Sang Lam. 2009. "Income risk and portfolio choice: An empirical study." *The Journal of Finance* 64 (2): 1037–1055.

Asness, Clifford, and Andrea Frazzini. 2013. "The Devil in HML's Details." *Journal of Portfolio Management* 39 (4): 49–68.

Backus, David, Nina Boyarchenko, and Mikhail Chernov. 2018. "Term structures of asset prices and returns." *Journal of Financial Economics* 129 (1): 1–23.

Backus, David, Mikhail Chernov, and Ian Martin. 2011. "Disasters Implied by Equity Index Options." *Journal of Finance* 66 (6): 1969–2012.

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

Bansal, Ravi, Robert Dittmar, and Dana Kiku. 2009. "Cointegration and consumption risks in asset returns." *The Review of Financial Studies* 22 (3): 1343–1375.

Bansal, Ravi, Robert F Dittmar, and Christian T Lundblad. 2005. "Consumption, dividends, and the cross section of equity returns." *Journal of Finance* 60 (4): 1639–1672.

Bansal, Ravi, Shane Miller, Dongho Song, and Amir Yaron. 2019. *The term structure of equity risk premia.* Technical report. National Bureau of Economic Research.

Bansal, Ravi, and Amir Yaron. 2004. “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles.” *Journal of Finance* 57 (4): 1481–1509.

Barro, Robert J. 2006. “Rare disasters and asset markets in the twentieth century.” *Quarterly Journal of Economics* 121 (3): 823–866.

———. 2009. “Rare Disasters, Asset Prices, and Welfare Costs.” *American Economic Review* 99 (1): 243–264.

Beaudry, Paul, and Gary Koop. 1993. “Do recessions permanently change output?” *Journal of Monetary economics* 31 (2): 149–163.

Beeler, J., and John Y. Campbell. 2012. “The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment.” *Critical Finance Review* 1 (1): 141–182.

Bekaert, Geert, and Eric Engstrom. 2017. “Asset Return Dynamics under Habits and Bad Environment–Good Environment Fundamentals.” *Journal of Political Economy* 60 (3): 713–760.

Bekaert, Geert, Eric Engstrom, and Yuhang Xing. 2009. “Risk, uncertainty, and asset prices.” *Journal of Financial Economics* 91 (1): 59–82.

Belo, Frederico, Pierre Collin-Dufresne, and Robert S. Goldstein. 2015. “Dividend Dynamics and the Term Structure of Dividend Strips.” *Journal of Finance* 70 (3): 1115–1160.

Belo, Frederico, Xiaoji Lin, and Santiago Bazdresch. 2014. “Labor hiring, investment, and stock return predictability in the cross section.” *Journal of Political Economy* 122 (1): 129–177.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S Goldstein. 2007. “Portfolio choice over the life-cycle when the stock and labor markets are cointegrated.” *The Journal of Finance* 62 (5): 2123–2167.

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

Beveridge, Stephen, and Charles R Nelson. 1981. “A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the ‘business cycle’.” *Journal of Monetary economics* 7 (2): 151–174.

Binsbergen, Jules van, Michael Brandt, and Ralph Koijen. 2012. “On the Timing and Pricing of Dividends.” *American Economic Review* 102 (4): 1596–1618.

Black, Fischer, and Myron Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81 (3): 637–654.

Bodie, Zvi, Robert C Merton, and William F Samuelson. 1992. "Labor supply flexibility and portfolio choice in a life cycle model." *Journal of Economic Dynamics and Control* 16 (3-4): 427–449.

Boguth, Oliver, Murray Carlson, Adlai Fisher, and Mikhail Simutin. 2012. "Leverage and the Limits of Arbitrage Pricing: Implications for Dividend Strips and the Term Structure of Equity Risk Premia." Working Paper.

Bollerslev, Tim, George Tauchen, and Hao Zhou. 2009. "Expected stock returns and variance risk premia." *The Review of Financial Studies* 22 (11): 4463–4492.

Bollerslev, Tim, and Viktor Todorov. 2011. "Tails, fears, and risk premia." *Journal of Finance* 66 (6): 2165–2211.

Breeden, Douglas, and Robert H. Litzenberger. 1978. "State Contingent Prices Implicit in option Prices." *Journal of Business* 51 (4): 3–24.

Broadie, Mark, Mikhail Chernov, and Michael Johannes. 2009. "Understanding Index Option Returns." *Review of Financial Studies* 22 (11): 4493–4529.

Campanale, Claudio, Carolina Fugazza, and Francisco J. Gomes. 2015. "Life-cycle portfolio choice with liquid and illiquid financial assets." *Journal of Monetary Economics* 71:67–83.

Campbell, John Y, and N Gregory Mankiw. 1987. "Are output fluctuations transitory?" *Quarterly Journal of Economics* 102 (4): 857–880.

Campbell, John Y. 1996. "Understanding risk and return." *Journal of Political Economy* 104 (2): 298–345.

———. 2006. "Household finance." *Journal of Finance* 61 (4): 1553–1604.

Campbell, John Y., João F. Cocco, Francisco J. Gomes, and Pascal J. Maenhout. 2001. "Investing Retirement Wealth: A Life-Cycle Model." In *Risk Aspects of Investment-Based Social Security Reform,* 439–482. University of Chicago Press.

Campbell, John Y., and John H. Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107 (2): 205–251.

Carroll, Christopher D. 2000. *Portfolios of the Rich.* Technical report. National Bureau of Economic Research.

———. 2006. "The method of endogenous gridpoints for solving dynamic stochastic optimization problems." *Economics letters* 91 (3): 312–320.

Cecchetti, Stephen G, Pok-sang Lam, and Nelson C Mark. 1993. "The equity premium and the risk-free rate: Matching the moments." *Journal of Monetary Economics* 31 (1): 21–45.

Chang, Yongsung, Jay H. Hong, and Marios Karabarbounis. 2018. "Labor Market Uncertainty and Portfolio Choice Puzzles." *American Economic Journal: Macroeconomics* 10 (2): 222–62.

Chen, Huafeng Jason, Marcin Kacperczyk, and Hernan Ortiz-Molina. 2011. "Labor unions, operating flexibility, and the cost of equity." *Journal of Financial and Quantitative Analysis* 46 (1): 25–58.

Christelis, Dimitris, Tullio Jappelli, and Mario Padula. 2010. "Cognitive abilities and portfolio choice." *European Economic Review* 54 (1): 18–38.

Cocco, Joao F. 2004. "Portfolio choice in the presence of housing." *The Review of Financial Studies* 18 (2): 535–567.

Cocco, Joao F, Francisco J Gomes, and Pascal J Maenhout. 2005. "Consumption and portfolio choice over the life cycle." *The Review of Financial Studies* 18 (2): 491–533.

Cochrane, John H. 1988. "How big is the random walk in GNP?" *Journal of political economy* 96 (5): 893–920.

———. 1994. "Permanent and transitory components of GNP and stock prices." *The Quarterly Journal of Economics* 109 (1): 241–265.

Cochrane, John H, and Argia M Sbordone. 1988. "Multivariate estimates of the permanent components of GNP and stock prices." *Journal of economic dynamics and control* 12 (2-3): 255–296.

Conrad, Jennifer, and Sunil Wahal. 2020. "The term structure of liquidity provision." *Journal of Financial Economics* 136 (1): 239–259.

Constantinides, George M, and Darrell Duffie. 1996. "Asset pricing with heterogeneous consumers." *Journal of Political economy* 104 (2): 219–240.

Cuesdeanu, H., and J.C. Jackwerth. 2018. "The pricing kernel puzzle: survey and outlook." *Annals of Finance* 14:289–329.

Danthine, Jean-Pierre, John Donaldson, and Rajnish Mehra. 2017. "A 21st Century Perspective on Asset Management." *Finance in Society,* 55–68.

Davis, Steven J, and Paul Willen. 2000. *Occupation-level income shocks and asset returns: their covariance and implications for portfolio choice.* Technical report. National Bureau of Economic Research.

Dew-Becker, Ian, Stefano Giglio, Anh Le, and Marius Rodriguez. 2017. "The price of variance risk." *Journal of Financial Economics* 123 (2): 225–250.

Donangelo, Andres. 2018. "Untangling the Value Premium with Labor Shares." *The Review of Financial Studies.*

Donangelo, Andres, Francois Gourio, Matthias Kehrig, and Miguel Palacios. 2019. "The cross-section of labor leverage and equity returns." *Journal of Financial Economics* 132 (2): 497–518.

Drechsler, Itamar, and Amir Yaron. 2011. "What's Vol Got to Do with It." *Review of Financial Studies* 24 (1): 1–45.

Epstein, Larry G., and Stanley E. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57 (4): 937–969.

———. 1991. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis." *Journal of Political Economy* 99 (2): 263–286.

Fagereng, Andreas, Luigi Guiso, and Luigi Pistaferri. 2017. "Portfolio Choices, Firm Shocks, and Uninsurable Wage Risk." *The Review of Economic Studies* 85 (1): 437–474.

Fama, Eugene F. 1992. "Transitory variation in investment and output." *Journal of Monetary Economics* 30 (3): 467–480. https://doi.org/https://doi.org/10.1016/0304-3932(92)90007-O.

Favilukis, Jack, and Xiaoji Lin. 2016. "Does wage rigidity make firms riskier? Evidence from long-horizon return predictability." *Journal of Monetary Economics* 78:80–95.

Gatheral, Jim. 2004. "A parsimonious arbitrage-free implied volatility parametrization with application to the valuation of volatility derivatives." Presentation at Global Derivatives.

———. 2006. *The Volatility Surface: A Practitioner's Guide.* Wiley Finance.

Gatheral, Jim, and Antoine Jacquier. 2011. "Convergence of Heston to SVI." *Quantitative Finance* 11 (8): 1129–1132.

Giglio, Stefano, and Bryan Kelly. 2018. "Excess volatility: Beyond discount rates." *The Quarterly Journal of Economics* 133 (1): 71–127.

Goetzmann, William N, and Alok Kumar. 2008. "Equity portfolio diversification." *Review of Finance,* 433–463.

Gomes, Francisco, and Alexander Michaelides. 2003. "Portfolio choice with internal habit formation: A life-cycle model with uninsurable labor income risk." *Review of Economic Dynamics* 6 (4): 729–766.

———. 2005. "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence." *Journal of Finance* 60 (2): 869–904.

Gormsen, Niels Joachim. 2018. "Time variation of the equity term structure." *Available at SSRN 2989695.*

Gourinchas, Pierre-Olivier, and Jonathan A Parker. 2002. "Consumption over the life cycle." *Econometrica* 70 (1): 47–89.

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese. 1996. "Income risk, borrowing constraints, and portfolio choice." *The American Economic Review,* 158–172.

Gul, F. 1991. "A Theory of Disappointment Aversion." *Econometrica* 59 (3): 667–686.

Guvenen, Fatih. 2007. "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review* 97 (3): 687–712.

———. 2009. "An empirical investigation of labor income processes." *Review of Economic dynamics* 12 (1): 58–79.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2016. *What do data on millions of US workers reveal about life-cycle earnings risk?* Technical report. Working paper.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2019. *What Do Data on Millions of US Workers Reveal about Life-Cycle Earnings Dynamics?* Technical report. Federal Reserve Bank of New York.

Guvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy* 122 (3): 621–660. http://www.jstor.org/stable/10.1086/675535.

Guvenen, Fatih, Sam Schulhofer-Wohl, Jae Song, and Motohiro Yogo. 2017. "Worker Betas: Five Facts about Systematic Earnings Risk." *American Economic Review* 107 (5): 398–403.

Hansen, Lars Peter, John C Heaton, and Nan Li. 2008. "Consumption strikes back? Measuring long-run risk." *Journal of Political Economy* 116 (2): 260–302.

Hasler, Michael, and Roberto Marfe. 2016. "Disaster recovery and the term structure of dividend strips." *Journal of Financial Economics* 122 (1): 116–134.

Heaton, John, and Deborah Lucas. 2000. "Portfolio choice in the presence of background risk." *The Economic Journal* 110 (460): 1–26.

Hubbard, R Glenn, Jonathan Skinner, and Stephen P Zeldes. 1995. "Precautionary saving and social insurance." *Journal of Political Economy* 103 (2): 360–399.

Jagannathan, Ravi, and Zhenyu Wang. 1996. "The conditional CAPM and the cross-section of expected returns." *Journal of finance* 51 (1): 3–53.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decisions under Risk." *Econometrica* 47 (2): 263–292.

Kim, Chang-Jin, James Morley, and Jeremy Piger. 2005. "Nonlinearity and the permanent effects of recessions." *Journal of Applied Econometrics* 20 (2): 291–309.

Kim, Chang-Jin, and Charles R Nelson. 1999. "Friedman's plucking model of business fluctuations: tests and estimates of permanent and transitory components." *Journal of Money, Credit and Banking,* 317–334.

Kim, Chang-jin, Jeremy M Piger, and Richard Startz. 2007. "The dynamic relationship between permanent and transitory components of US business cycles." *Journal of Money, Credit and Banking* 39 (1): 187–204.

Kuehn, Lars-Alexander, Mikhail Simutin, and Jessie Jiaxu Wang. 2017. "A labor capital asset pricing model." *The Journal of Finance* 72 (5): 2131–2178.

Larrain, Borja, and Motohiro Yogo. 2008. "Does firm value move too much to be justified by subsequent changes in cash flow?" *Journal of Financial Economics* 87 (1): 200–226.

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

———. 2014. "Shocks and crashes." *NBER Macroeconomics Annual* 28 (1): 293–354.

Lo, Andrew W, and A Craig MacKinlay. 1988. "Stock market prices do not follow random walks: Evidence from a simple specification test." *The review of financial studies* 1 (1): 41–66.

Longstaff, Francis A, and Monika Piazzesi. 2004. "Corporate earnings and the equity premium." *Journal of financial Economics* 74 (3): 401–421.

Lucas, Robert E. 1978. "Asset Prices in an Exchange Economy." *Econometrica* 46 (6): 1429–1445.

Ludvigson, Sydney, and Martin Lettau. 2003. *Understanding trend and cycle in asset values: Reevaluating the wealth effect on consumption.* National Bureau of Economic Research.

Marfè, Roberto. 2016. "Corporate fraction and the equilibrium term structure of equity risk." *Review of Finance* 20 (2): 855–905.

———. 2017. "Income insurance and the equilibrium term structure of equity." *The Journal of Finance* 72 (5): 2073–2130.

Martin, Ian. 2017. "What is the expected return on the market." *Quarterly Journal of Economics* 132 (1): 367–433.

Mehra, Rajnish, and Edward C Prescott. 2003. "The equity premium in retrospect." *Handbook of the Economics of Finance* 1:889–938.

Merton, Robert C. 1969. "Lifetime portfolio selection under uncertainty: The continuous-time case." *The Review of Economics and Statistics,* 247–257.

———. 1971. "Optimal portfolio and consumption rules in a continuous-time model." *Journal of Economic Theory* 3 (4): 373–413.

Morley, James, and Jeremy Piger. 2012. "The asymmetric business cycle." *Review of Economics and Statistics* 94 (1): 208–221.

Pruitt, Seth, and Nicholas Turner. 2020. "Earnings Risk in the Household: Evidence from Millions of US Tax Returns." *American Economic Review: Insights* 2 (2): 237–54.

Rabin, Matthew. 2000. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica* 68 (5): 1281–1992.

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

Routledge, Bryan R., and Stanley E. Zin. 2010. "Generalized Disappointment Aversion and Asset Prices." *Journal of Finance* 65 (4): 1303–1332.

Rubinstein, Mark. 1994. "Implied Binomial Trees." *Journal of Finance* 49 (3): 771–818.

Santos, Tano, and Pietro Veronesi. 2006. "Labor income and predictable stock returns." *The Review of Financial Studies* 19 (1): 1–44.

Schmidt, Lawrence. 2016. "Climbing and falling off the ladder: Asset pricing implications of labor market event risk." *SSRN.*

Schreindorfer, David. 2020. "Macroeconomic Tail Risks and Asset Prices." *Review of Financial Studies* 33 (8): 3541–3582.

Schulz, Florian. 2016. "On the Timing and Pricing of Dividends: Comment." *American Economic Review* 106 (10): 3185–3223.

Seo, Sang Byung, and Jessica Wachter. 2019. "Option Prices in a Model with Stochastic Disaster Risk." *Management Science* 65 (8): 3449–3469.

Sichel, Daniel E. 1994. "Inventories and the three phases of the business cycle." *Journal of Business & Economic Statistics* 12 (3): 269–277.

Tauchen, George, and Robert Hussey. 1991. "Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models." *Econometrica,* 371–396.

Van Binsbergen, Jules, Michael Brandt, and Ralph Koijen. 2012. "On the timing and pricing of dividends." *American Economic Review* 102 (4): 1596–1618.

Van Binsbergen, Jules, Wouter Hueskes, Ralph Koijen, and Evert Vrugt. 2013. "Equity yields." *Journal of Financial Economics* 110 (3): 503–519.

Van Rooij, Maarten, Annamaria Lusardi, and Rob Alessie. 2011. "Financial literacy and stock market participation." *Journal of Financial Economics* 101 (2): 449–472.

———. 2012. "Financial literacy, retirement planning and household wealth." *The Economic Journal* 122 (560): 449–478.

Wachter, Jessica. 2005. "Solving models with external habit." *Finance Research Letters* 2 (4): 210–226.

———. 2013. "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" *Journal of Finance* 68:987–1035.

Wachter, Jessica, and Motohiro Yogo. 2010. "Why do household portfolio shares rise in wealth?" *The Review of Financial Studies* 23 (11): 3929–3965.

Welch, Ivo. 2018. "The (Time-Varying) Importance of Disaster Risk." *Financial Analysts Journal* 72 (5): 14–30.

Working, Holbrook. 1960. "Note on the Correlation of First Differences of Aggregates in a Random Chain." *Econometrica* 28:916–918.

APPENDIX A

APPENDIX TO CHAPTER 2

A.1 Labor Share and Labor Income Growth Properties

Begin with the following decomposition of log consumption growth,

$$\Delta c_{t+1} = \mu_c + \mathbb{E}_t[\Delta c_{t+1} - \mu_c] + \mathcal{U}^c_{t+1} \tag{A.1}$$

where unexpected consumption growth $\mathcal{U}^c_{t+1}$ has zero expectation and is unpredictable using past information. Then $\mathrm{Var}_t(\Delta c_{t+1}) = \mathrm{Var}_t(\mathcal{U}^c_{t+1})$, $\mathrm{Var}(\Delta c_{t+1}) = \mathrm{Var}(\mathbb{E}_t[\Delta c_{t+1} - \mu_c]) + \mathrm{Var}(\mathcal{U}^c_{t+1})$.

The log labor share is given by

$$s_{t+1} = \bar{s}(1-\beta_s) + \beta_s s_t + \beta_c \mathcal{U}^c_{t+1} + \sigma_\eta \eta_{t+1} \tag{A.2}$$

where η_{t+1} is an independent standard Normal innovation. The conditional mean and variance of the labor share are $\mathbb{E}_t[s_{t+1}] = \bar{s}(1-\beta_s) + \beta_s s_t$ and $\mathrm{Var}_t(s_{t+1}) = \sigma_\eta^2 + \beta_c^2 \mathrm{Var}_t(\mathcal{U}^c_{t+1})$, while the unconditional variance is $\mathrm{Var}(s_{t+1}) = \left(\sigma_\eta^2 + \beta_c^2 \mathrm{Var}(\mathcal{U}^c_{t+1})\right)/(1-\beta_s^2)$.

The conditional covariance between the log labor share and consumption growth is $\mathrm{Cov}_t(s_{t+1}, \Delta c_{t+1}) = \beta_c \mathrm{Var}_t(\mathcal{U}^c_{t+1})$. To compute the unconditional covariance, it is first necessary to compute $\mathrm{Cov}(s_t, \mathbb{E}_t[\Delta c_{t+1} - \mu_c])$. By recursive substitution, we find

$$\mathrm{Cov}(s_t, \mathbb{E}_t[\Delta c_{t+1} - \mu_c]) = \beta_c \sum_{k=0}^{\infty} \beta_s^k \mathrm{Cov}(\mathcal{U}^c_{t-k}, \mathbb{E}_t[\Delta c_{t+1} - \mu_c]), \tag{A.3}$$

which is capturing the relationship between past realized consumption growth innovations and current expected consumption growth. This would represent a correlation between the innovations to Δc_{t+1} and x_{t+1} in the long run risks model, for example. The unconditional covariance of the log labor share and consumption growth is then

$$\mathrm{Cov}(s_{t+1}, \Delta c_{t+1}) = \beta_c \mathrm{Var}(\mathcal{U}^c_{t+1}) + \beta_c \sum_{k=0}^{\infty} \beta_s^{k+1} \mathrm{Cov}(\mathcal{U}^c_{t-k}, \mathbb{E}_t[\Delta c_{t+1} - \mu_c]). \tag{A.4}$$

The conditional covariance is negative when β_c is negative, while the unconditional covariance is negative as long as $\sum_{k=0}^{\infty} \beta_s^{k+1} \mathrm{Cov}(\mathcal{U}^c_{t-k}, \mathbb{E}_t[\Delta c_{t+1} - \mu_c]) > -\mathrm{Var}(\mathcal{U}^c_{t+1})$. Changes in the log labor share exhibit many of the same (or similar) properties because

$$\Delta s_{t+1} = (\beta_s - 1)s_t + \beta_c \mathcal{U}^c_{t+1} + \sigma_\eta \eta_{t+1}. \tag{A.5}$$

Labor income growth is, after substituting Eqs. (A.1) and (A.5) into (2.6) and some simplification, given by

$$\Delta \ell_{t+1} = \mu_c + \mathbb{E}_t[\Delta c_{t+1} - \mu_c] + (\beta_s - 1)s_t + (1+\beta_c)\mathcal{U}^c_{t+1} + \sigma_\eta \eta_{t+1}. \tag{A.6}$$

It is quick to verify that the first two conditional moments are $\mathbb{E}_t[\Delta\ell_{t+1}] = \mu_c + \mathbb{E}_t[\Delta c_{t+1} - \mu_c] + (\beta_s - 1)s_t$ and $\text{Var}_t(\Delta\ell_{t+1}) = \sigma_\eta^2 + (1+\beta_c)^2\text{Var}_t(\mathcal{U}_{t+1}^c)$. The conditional covariance with consumption growth is $\text{Cov}_t(\Delta\ell_{t+1}, \Delta c_{t+1}) = (1+\beta_c)\text{Var}_t(\mathcal{U}_{t+1}^c)$.

The unconditional expectation of labor income grown is the same as that of consumption growth. The unconditional variance is

$$\begin{aligned}\text{Var}(\Delta\ell_{t+1}) = \text{Var}(\mathbb{E}_t[\Delta c_{t+1} - \mu_c]) + (\beta_s - 1)^2\text{Var}(s_t) + \sigma_\eta^2 \\ + (1+\beta_c)^2\text{Var}(\mathcal{U}_{t+1}^c) + 2(\beta_s - 1)\text{Cov}(s_t, \mathbb{E}_t[\Delta c_{t+1} - \mu_c])\end{aligned} \quad (\text{A.7})$$

and the unconditional covariance with consumption growth is

$$\begin{aligned}\text{Cov}(\Delta\ell_{t+1}, \Delta c_{t+1}) = \text{Var}(\mathbb{E}_t[\Delta c_{t+1} - \mu_c]) + (1+\beta_c)\text{Var}(\mathcal{U}_{t+1}^c) \\ + (\beta_s - 1)\text{Cov}(s_t, \mathbb{E}_t[\Delta c_{t+1} - \mu_c]).\end{aligned} \quad (\text{A.8})$$

In many cases, the formulas above simplify considerably. Many models do not include time-variation in expected consumption growth. For those that do, the term $\text{Cov}(s_t, \mathbb{E}_t[\Delta c_{t+1} - \mu_c])$ shown in Eq. (A.3) will be zero in most models, as they do not generally incorporate such correlation.

A.2 Data Construction

This appendix provides the necessary details to construct the data used in the analysis of this paper.

A.2.1 Consumption

The series used for aggregate consumption is the sum of real per-capita nondurables and services consumption from BEA NIPA Table 7.1. Consumption growth is computed as log changes in this series, as in Beeler and Campbell (2012).

A.2.2 Equity Payout

I use three sources for equity payout variables. For payout from public firms, I construct a CRSP dividend and CRSP dividend plus repurchase series following Bansal, Dittmar, and Lundblad (2005). For aggregate payout, I extend the net dividend and net equity payout series developed in Larrain and Yogo (2008), mainly sourcing from the Flow of Funds data. Net equity payout is computed from the Flow of Funds as the difference between net dividends paid and changes in corporate equity. Both series are taken from Table F.103 for Nonfinancial Corporate Business (lines 3 and 43, respectively). The final series I use is a measure of dividends received by the US household sector, called personal dividend income, which is from BEA NIPA Table 2.1 line 15.

For the portfolios, I sort firms by the sorting variable into 1 of 5 portfolios as of June 30 in each year. Portfolios are value-weighted and held through June 30 of the next year. Size portfolios sort on firm market capitalization Book-to-market portfolios sort on book-to-market, constructed following Asness and Frazzini (2013). I extend the series back to 1926 by using the historical Moody's book equity data available on the Ken French Data Library. Momentum portfolios are formed on the prior [-11,-2] return. Market beta portfolios sort on market beta, constructed using 60 months of data. For industry portfolios, I use the Fama-French 5 classifications.

The growth rates are converted to real using the growth in the CPI.

A.2.3 Labor Income

Aggregate labor income is computed following Lettau and Ludvigson (2001), using BEA NIPA Table 2.1. The series is after-tax labor income (ATLI) and its definition is

$$\begin{aligned}\text{After-tax labor income} = &\text{Wages and salaries} + \text{Transfer payments} + \\ &+ \text{Employer contrib. for employee pensions and insurance} \\ &- \text{Employee contrib. for social insurance} \\ &- \text{Taxes.}\end{aligned}$$

For robustness, I also consider several alternative series, all using data from BEA NIPA Table 2.1. The first is Compensation of Employees (CoE) series from line 2. The second is ATLI plus two-thirds of proprietors income with capital consumption and inventory adjustments (line 9). The third is ATLI minus personal current transfers (line 16). The variables are converted to real per-capita series using the implicit PCE deflator and population series from the BEA.

A.2.4 Earnings

I use Gross Value Added and Net Operating Surplus, both from the nonfinancial corporate business sector, lines 17 and 24 from BEA NIPA Table 1.14. The series is converted to real using the implicit GDP deflator.

A.2.5 Returns

My proxy for the return on the aggregate equity claim is return on the CRSP value-weighted index of all common stocks. Real returns are created by adjusting the nominal returns by the growth in the CPI.

APPENDIX B

APPENDIX TO CHAPTER 3

B.1 Estimating f^*

To recover risk-neutral densities from options based on (3.5), it is necessary to observe option prices for the desired maturity and a continuum of strikes. We generate these prices via interpolation and extrapolation of observed quotes as follows. For each day in the sample, we use Black's formula (a version of Black and Scholes (1973)) to convert observed option prices to implied volatility (IV) units, fit an interpolant to them, evaluate the interpolant at a maturity of 30 calendar days and a fine grid of strike prices, map interpolated IVs back to option prices, and finally compute f_t^* via finite differences based on (3.5). Importantly, this approach does not assume the validity of the Black-Scholes model because Black's formula is merely used to map back-and-forth between two spaces. The mapping relies on LIBOR rates that are linearly interpolated to the options' maturities and forward prices for the underlying. The remainder of this appendix details the interpolation of IVs.

B.1.1 The SVI Method

We interpolate IVs based on Jim Gatheral's SVI (stochastic volatility inspired) method.[35] SVI describes implied variance (the square of IV) for a given maturity τ with the function

$$\sigma^2_{BSM}(x) = a + b\left(\rho(x-m) + \sqrt{(x-m)^2 + \sigma^2}\right), \tag{B.1}$$

where $x = \log(K/F_{t,\tau})$ is the option's log-moneyness, $F_{t,\tau}$ the forward price for maturity τ, and a, b, ρ, m, σ are parameters. The method is widely used in financial institutions because it is parsimonious, yet known to provide a very good approximation to IVs, both in the data and in fully-specified option pricing models.

We make two modifications to the basic SVI method to allow for interpolation in the maturity, in addition to the moneyness dimension. First, we parameterize σ^2_{BSM} as a function of standardized moneyness, $\kappa \equiv \frac{\log(K/F_{t,\tau})}{VIX_t/100\times\sqrt{\tau}}$, rather than x. Doing so limits the extent to which the shape of the IV curve varies with maturity and makes it easier to fit the τ-dimension of the IV surface. Second, we specify linear functions of τ for the five coefficients, e.g.,

$$a = a_0 + a_1\tau, \tag{B.2}$$

[35] SVI was devised at Merrill Lynch and disseminated publicly by Gatheral (2004). See Gatheral and Jacquier (2011) for theoretical properties of the SVI method, Gatheral (2006) for a textbook treatment, and Berger, Dew-Becker, and Giglio (2020) for a recent application in economics.

and similarly for (b, ρ, m, σ). Jointly, (B.1) and (B.2) describe IVs as a bivariate function of standardized moneyness and maturity that depends on the parameter vector $\theta \equiv (a_0, a_1, b_0, b_1, \rho_0, \rho_1, m_0, m_1, \sigma_0, \sigma_1)$.

An important criterion for the successful interpolation and extrapolation of IVs is that the corresponding option prices respect theoretical no arbitrage restrictions, i.e. that they are (i) non-negative, (ii) monotonic in K, (iii) convex in K, and (iv) imply (via Equation 3.5) a density $f_t^*(R)$ that integrates to one. Because the function $\sigma_{BSM}(\kappa, \tau; \theta)$ is not guaranteed to satisfy these constraints for every parameter vector θ, we impose them in the estimation as further described below.

B.1.2 Data and Implementation

We apply standard filters to remove quotes with low liquidity and obvious data errors. Specifically, observations are dropped if they (i) violate the static no-arbitrage bounds $P \leq K/R^f$ or $C \leq S$, (ii) have a best bid quote of zero, (iii) have the CBOE's error code 999 for ask quotes or 998 for bid quotes, (iv) have non-positive bid-ask spreads, (v) have midquotes less than \$0.50, (vi) are singles (a call quote without a matching put quote or vice versa), (vii) are PM settled, or (viii) have annualized IVs less than 2% or more than 200%. To detect any additional outliers, we fit a linear function in maturity and standardized moneyness to IVs on each date, and remove observations that are highly influential based on their Cook's distance (a common statistical metric for detecting outliers). Finally, we restrict the sample to put options with a standardized moneyness below 0.5, call options with a standardized moneyness above -0.5, and maturities between 8 and 120 calendar days, i.e. we exclude long-term and in-the-money options.

For each day in the sample, we estimate the SVI parameter vector θ by minimizing the root mean squared error between observed IVs and the SVI interpolant,

$$\hat{\theta}_t = \underset{\theta}{\operatorname{argmin}} \sqrt{\frac{1}{N_t} \sum_{i=1}^{N_t} \left[\sigma_{BSM,t,i} - \sigma_{BSM}(\kappa_{t,i}, \tau_{t,i}; \theta)\right]^2}, \tag{B.3}$$

where N_t is the number of observations on day t. We use a particle swarm algorithm to minimize the objective function globally over the parameter space and discard parameter vectors for which SVI-implied option prices violate no arbitrage constraints. The positivity, monotonicity, and convexity of option prices are checked on a bivariate grid for κ and τ.[36] We extend the moneyness grid over -20 to +10 standard deviations

[36]The κ-grid includes the integers from -20 to -11, 61 equally-spaced points between -10 and 5, and the integers from 6 to 10, for a total of 76 points. The τ-grid is equally-spaced with 12 points between 10 and 120 days to maturity.

units to ensure that even extrapolated option prices are arbitrage-free. At every maturity in the τ-grid, we numerically integrate f_t^* over the κ-region and discard parameter vectors for which these integrals do not fall within (a numerical error tolerance of) 1 basis point of one.

The fit to IVs results in an average (median) R^2 of 98.8% (99.6%) across the 7,556 trading days in our sample. The online supplement provides examples of the SVI fitting procedure for a few days, and shows that our main results are robust to different functional forms of the IV interpolant.

B.2 Longer return horizons

Figure 25. $EP(x)$ for quarterly returns.

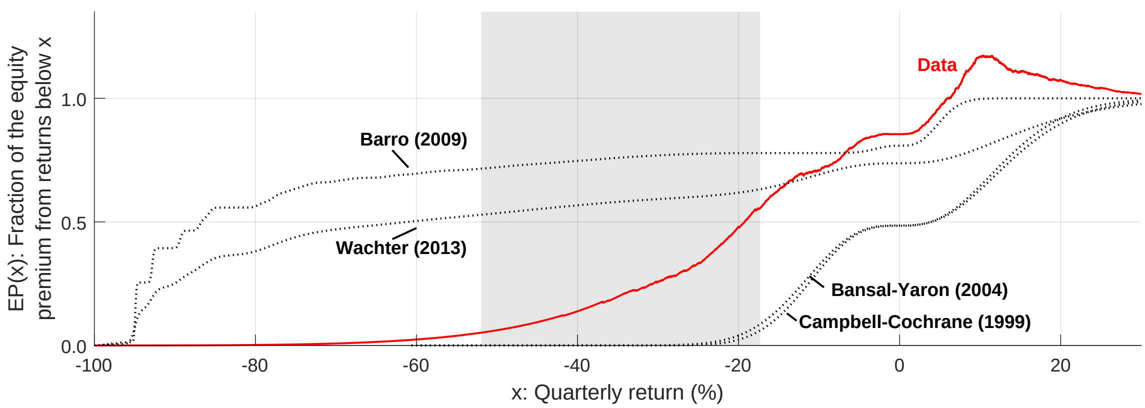

The paper focuses on a return horizon of one month. To understand the economic mechanism of existing models, and to relate our results to literature on the term structure of asset prices, it is similarly useful to decompose the equity premium for longer horizons. The fact that option prices and returns can only be observed jointly for 30 years, however, makes it difficult to compute the necessary estimate of the physical return distribution with reasonable precision at very long horizons.

With this caveat in mind, we examine the importance of different return states at horizons of a few months. Figure 25 replicates our main analysis at the quarterly (90 calendar day) horizon. To account for the larger volatility of quarterly returns, we scale the left tail region of interest by a factor of $\sqrt{3}$. In the data, quarterly returns in the interval $\sqrt{3} \times [-30\%, -10\%]$ account for 50.4/100 of the equity premium, while they account for 10.1/100 or less in the four depicted models. At the quarterly horizon, the models therefore perform similarly poorly in capturing sources of the equity premium.

The habit and long-run risks models continue to attribute risk premia to less extreme and the disaster model to more extreme left tail states than the data.

Figure 26 provides a closer look at the term structure dimension by plotting the equity premium contribution (left panel) and price of risk (right panel) for returns in $\sqrt{h} \times [-30\%, -10\%]$ as a function of the return horizon, h. The left panel shows that, in the data, left tail events contribute slightly less to the equity premium at longer horizons. In contrast, they contribute equally much (habits and long-run risks) or more (disasters) in the models. The figure suggests that the gap between models and data remains significant for horizons well beyond three months.

The right panel shows that the ratio of average risk-neutral to physical probabilities – our measure for the price of risk – increases at longer horizons in the data. While the models capture this increasing pattern, they continue to substantially undershoot the price of risk at longer horizons. The finding that common utility functions imply too little aversion against shocks that coincide with stock market tail events is therefore robust to the horizon over which such tail events are measured.

Figure 26. The importance of left tail events at different return horizons.

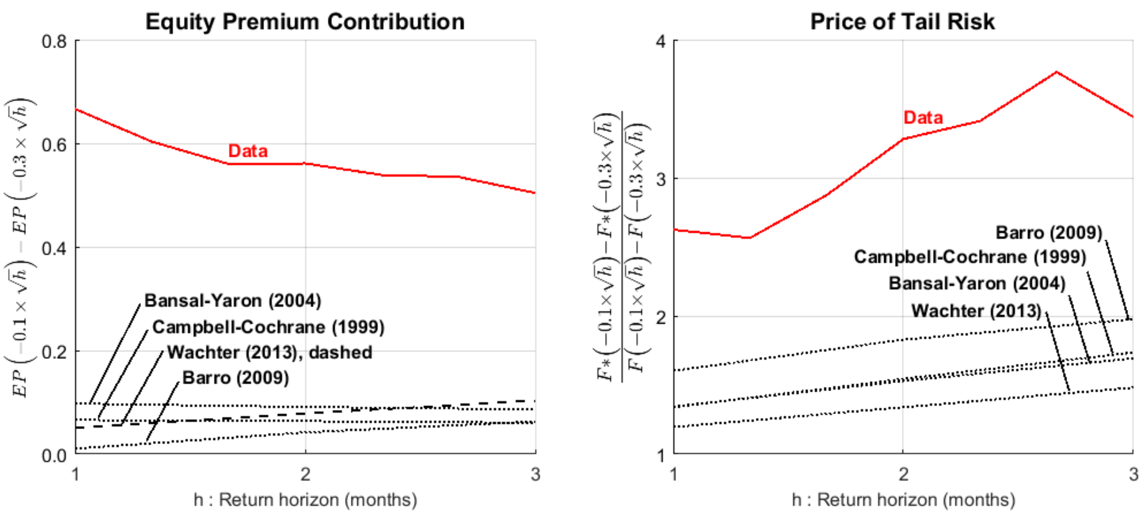

B.3 Bootstrap for $EP(x)$

We detail the block bootstrap used to compute the empirical sampling distribution for $EP(x)$ in Figure III of the main text. Our empirical sample consists of joint observations for $\{R_{t:t+30}, f^*_t(R_{t:t+30})\}$, $t = 1, ..., T$, where $f^*_t(R_{t:t+30})$ is the conditional risk-neutral PDF of a 30 calendar day return at time t, and $R_{t:t+30}$ is the realized return over the subsequent 30 calendar days. Using this sample, $EP(x)$ is computed as

$$EP(x) = \frac{\frac{1}{T}\sum_{t=1}^{T} R_{t:t+30}\mathbf{1}\{R_{t:t+30} \leq x\} - \int_{-1}^{x} R\left[\frac{1}{T}\sum_{t=1}^{T} f_t^*(R)\right] dR}{\frac{1}{T}\sum_{t=1}^{T} R_{t:t+30} - \int_{-1}^{\infty} R\left[\frac{1}{T}\sum_{t=1}^{T} f_t^*(R)\right] dR}.$$

To compute the joint sampling distribution of $EP(-30\%)$ and $EP(-10\%) - EP(-30\%)$, we use a block bootstrap with 10 million bootstrap samples and a block length of 21 trading days (the average number of trading days over periods of 30 calendar days). Specifically, the j^{th} pair $\{EP(-30\%), EP(-10\%)\}$ is created as follows.

1. Randomly draw an integer i between 1 and $T - 20$.
2. Add observations $i, ..., i + 20$ to the bootstrap sample.
3. Repeat steps 1 and 2 until the bootstrap sample contains at least T observations. In practice, there will be slightly more than T observations if $T/21$ does not equal an integer.
4. Use the bootstrap sample created in steps 1 through 3 to compute $EP(-30\%)$ and $EP(-10\%)$ based on the above formula for $EP(x)$.

The sampling distribution is created by repeating steps 1 through 4 ten million times.

B.4 Model Solutions

This section discusses our solution approach for each model and points out which results from the original studies were replicated for validation. The calculation of $EP(x)$ relies on option prices, which, among the papers we consider, only Backus, Chernov, and Martin 2011 and Schreindorfer 2020 solve for. We therefore derive our own solutions for the remaining models. Regardless of the model, we can write the put-to-spot price ratio for a strike price K and a 1-period maturity as

$$\begin{aligned}\mathcal{P}(X, \xi_t) =& \frac{1}{S_t}\mathbb{E}_t\left[M(\Delta c_{t+1}, \xi_t, \xi_{t+1}) \max\{0, K - S_{t+1}\}\right] \\ =& \mathbb{E}_t\left[M(\Delta c_{t+1}, \xi_t, \xi_{t+1})\left(X - \frac{PD(\xi_{t+1})}{PD(\xi_t)}e^{\Delta d_{t+1}}\right)\mathbf{1}_{\Delta d_{t+1} \leq \log\left(X\frac{PD(\xi_t)}{PD(\xi_{t+1})}\right)}\right],\end{aligned} \tag{B.4}$$

where $X = K/S_t$ equals moneyness, the pricing kernel M is a function of log consumption growth Δc_{t+1}, today's state ξ_t, and tomorrow's state ξ_{t+1}, and the price-dividend ratio PD is a function of the state as well. In what follows, we detail how the expectation on the RHS of (B.4) is evaluated in each model, and we defer details on the calculation of $EP(x)$ to Section B.5. Throughout, we rely on the same notation as the original studies and refer interested readers there for definitions and additional details.

B.4.1 Campbell and Cochrane (1999)

The external habits model does not admit analytical solutions. Wachter 2005 shows that the numerical solution in the original study is inaccurate and proposes the "series method" as a more precise alternative. We follow her solution approach. Specifically, the model is solved on a grid of 1001 unequally-spaced points for the model's state S_{t+1} (the surplus consumption ratio) – "Grid 3" in Wachter (2005). We employ Gauss-Chebyshev quadrature with 500 points covering ± 7 standard deviations when computing expectations. Cubic splines are used to interpolate the price-dividend ratio to off-grid values resulting from the quadrature. Our results closely match Table 2 in Wachter (2005).

The Campbell-Cochrane model features a single shock to each consumption and dividend growth, but no separate shock to the state. We evaluate the expectation in (B.4) in two steps. First, we condition on Δc_{t+1} and evaluate the expectation over Δd_{t+1} based on standard results for truncated normal random variables[37], which yields

$$\mathcal{P}(X,\xi_t) = \mathbb{E}_t\Bigg[M(\Delta c_{t+1},\xi_t)\Bigg(X\Phi\left(\nu(\Delta c_{t+1},\xi_t)\right) \\ - e^{\mathbb{E}[\Delta d_{t+1}|\Delta c_{t+1}]+\sigma[\Delta d_{t+1}|\Delta c_{t+1}]^2/2}\frac{PD(\Delta c_{t+1},\xi_t)}{PD(\xi_t)}\Phi\left(\nu(\Delta c_{t+1},\xi_t) - \sigma[\Delta d_{t+1}|\Delta c_{t+1}]\right)\Bigg)\Bigg]$$

where

- The functional form for the pricing kernel M is shown in Equation 5 of Campbell and Cochrane (1999)
- $\nu(\Delta c_{t+1},\xi_t) = \frac{\log\left(X\frac{PD(\xi_t)}{PD(\Delta c_{t+1},\xi_t)}\right) - E[\Delta d_{t+1}|\Delta c_{t+1}]}{\sigma[\Delta d_{t+1}|\Delta c_{t+1}]}$
- $E[\Delta d_{t+1}|\Delta c_{t+1}] = g + \rho\sigma_w v_{t+1}$
- $\sigma[\Delta d_{t+1}|\Delta c_{t+1}] = \sqrt{1-\rho^2}\sigma_w$

and $\Phi(\cdot)$ denotes the CDF of a standard normal. The remaining expectation over Δc_{t+1} is evaluated based on the aforementioned quadrature method.

B.4.2 Bansal and Yaron (2004)

We solve the long-run risks model numerically rather than relying on the log-linear approximation provided in the original paper. The model's bivariate state $\xi_t = [x_t, \sigma_t^2]$ is discretized with grids of size $N_x = 200$ and $N_\sigma = 100$, respectively. The grids are linearly-spaced, centered around the unconditional mean of each process, and

[37] For details, see, e.g. Lemma 1 in Schreindorfer (2020)

extend 5 standard deviations in each direction. For this purpose, we set the standard deviation of the conditional mean process to the one implied by the largest grid value for volatility. We use Gauss-Chebyshev quadrature with 30 nodes covering ± 7 standard deviations for shocks to the conditional mean and variance processes when computing expectations. We solve for the value-consumption and price-dividend ratios by iterating on the the Euler equations for both ratios until convergence, using bivariate cubic splines to interpolate to off-grid values when necessary. Our solution replicates Table 2 in Beeler and Campbell (2012), who provide an examination of long-run risk models and show more moments than the original paper.

The model features four IID normal shocks, one to consumption growth, one to dividend growth, and one to each of the model's two states. We apply standard results for truncated normal random variables to integrate over the shocks to consumption and dividend growth, which yields

$$\mathcal{P}(X,\xi_t) = \mathbb{E}_t\Bigg[M_1(\xi_t,\xi_{t+1})X\Phi\left(\nu(\xi_t,\xi_{t+1})\right) - M_2(\xi_t,\xi_{t+1})\frac{PD(\xi_{t+1})}{PD(\xi_t)}\Phi\left(\nu(\xi_t,\xi_{t+1})-\varphi_d\sigma_t\right)\Bigg]$$

where, based on the Hansen, Heaton, and Li (2008) formulation of the Epstein and Zin (1989) pricing kernel with VC being the utility-consumption ratio and $\mathcal{R}$ its certainty equivalent,

- $\nu(\xi_t,\xi_{t+1}) = \frac{\log\left(X\frac{PD(\xi_t)}{PD(\xi_{t+1})}\right)-\mu-\phi x_t}{\varphi_d\sigma_t}$
- $M_1(\xi_t,\xi_{t+1}) = \beta\left(\frac{VC(\xi_{t+1})}{\mathcal{R}(\xi_t)}\right)^{\frac{1}{\psi}-\gamma} e^{-\gamma\mu-\gamma x_t+\gamma^2\sigma^2/2}$
- $M_2(\xi_t,\xi_{t+1}) = \beta\left(\frac{VC(\xi_{t+1})}{\mathcal{R}(\xi_t)}\right)^{\frac{1}{\psi}-\gamma} e^{(1-\gamma)\mu+(\phi-\gamma)x_t+(\gamma^2+\varphi_d^2)\sigma^2/2}$

and $\Phi(\cdot)$ denotes the CDF of a standard normal. The remaining expectation over shocks to the model's future state is evaluated based on the aforementioned quadrature method.

B.4.3 Barro (2009)

The rare disaster model is cast in discrete time, but Barro (2009) solves it based on continuous time approximations. Because such approximations are not available for option prices, we instead rely on exact, analytical, discrete time solutions for the price-consumption ratio and risk-free rate. We find that these solutions differ only marginally from the values implied by the approximations in Equations 5, 7, and 12 of the original paper.

Barro (2009) provides only a back-of-the-envelope calculation for the (levered) equity premium, i.e. he does not formally define a dividend process. Because we require the price of an equity claim in order to price equity index options, we follow Abel (1999) in defining dividends as $D_t = C_t^{\phi}$, where ϕ captures leverage. We set leverage to a value of 2.6, the same number used by Wachter (2013). All of our results for the Barro model are based on this equity claim.

The model relies on a mixture of Gaussian shocks and disasters, where disaster sizes are drawn from the empirical distribution of international disaster occurrences. We follow Barro (2009) in assuming that the probability of multiple disasters in a short time period (in our case, a month) is negligible. To compute option prices, we condition on the disaster realization, evaluate expectations over Gaussian shocks analytically, and then average over disasters based on their empirical distribution. Hence, option prices are computed analytically as

$$\begin{aligned}\mathcal{P}(X,\xi_t) = (1-p)\left(Xe^{-\gamma g+\gamma^2\sigma^2/2}\Phi\left(\nu+\gamma\sigma\right) - e^{(\phi-\gamma)g+(\phi-\gamma)^2\sigma^2/2}\Phi\left(\nu-(\phi-\gamma)\sigma\right)\right)\\ + \frac{p}{N}\sum_{i=1}^{N} Xe^{-\gamma(g+Z_i)+\gamma^2\sigma^2/2}\Phi\left(\nu-\frac{Z_i}{\sigma}+\gamma\sigma\right)\\ - \frac{p}{N}\sum_{i=1}^{N} e^{(\phi-\gamma)(g+Z_i)+(\phi-\gamma)^2\sigma^2/2}\Phi\left(\nu-\frac{Z_i}{\sigma}-(\phi-\gamma)\sigma\right)\end{aligned}$$

where

- $\nu = \frac{\log(X)/\phi - g}{\sigma}$
- N denotes the number of empirical disaster realizations
- Z_i is the log consumption drop observed during disaster i

and $\Phi(\cdot)$ denotes the CDF of a standard normal.

B.4.4 Wachter (2013)

We solve a discrete time version of Wachter (2013) in order to simplify the calculation of option prices. Log consumption growth Δc and the disaster probability λ follow

$$\begin{aligned}\Delta c_{t+1} &= \mu dt + \sigma\sqrt{dt}\varepsilon^c_{t+1} + Z_{t+1}\nu_{t+1}\\ \lambda_{t+1} &= \lambda_t + \kappa(\bar{\lambda}-\lambda_t)dt + \sigma_l\sqrt{dt}\sqrt{\lambda_t}\varepsilon^l_{t+1},\end{aligned} \tag{B.5}$$

where $\varepsilon^c_t, \varepsilon^l_t \sim N(0,1)$, $N_t \sim$Berloulli$(\lambda_t dt)$, and the disaster size Z_t is drawn from the empirical disaster size distribution.[38] All shocks are IID. The discretization assumes that either zero or one disaster occurs every period, with a disaster probability of $\lambda_t dt$ per period. Dividend growth is given by $\Delta d_t = \phi \times \Delta c_t$, and the agent has Epstein and Zin 1991 utility.

Figure 27. Replication of Wachter (2013).

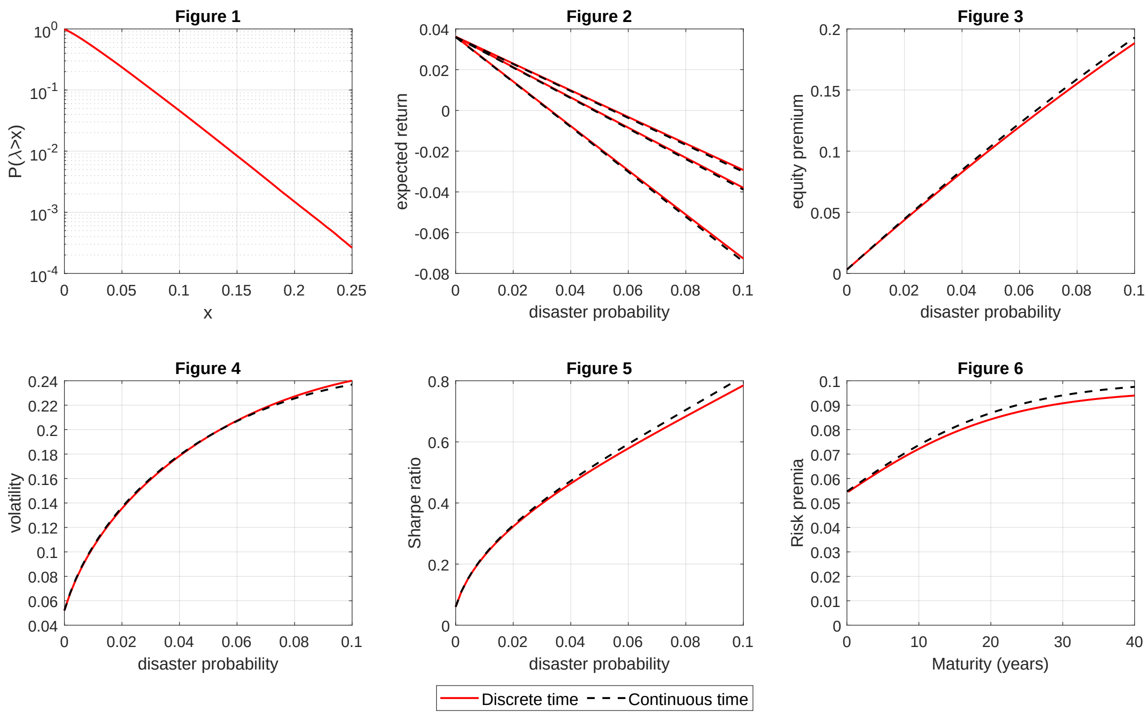

Note. We replicate figures 1 through 6 from Wachter (2013) for the original continuous time model, and compare it to our discrete time version of the model.

The calibration is identical to that in Wachter (2013). We solve the model at a monthly frequency ($dt = 1/12$), using an equally-spaced grid for λ_t with 500 points between 1e-16 and 1. Expectations over ε^c_{t+1}, Z_{t+1}, and N_{t+1} are evaluated analytically, while expectations over ε^l_{t+1} are evaluated via Gauss-Chebyshev quadrature with 500 points between ± 7 standard deviations. We solve the value-consumption ratio and the price-dividend ratio by iterating on the system until convergence. Both ratios are interpolated with cubic splines to any off-grid values of λ_{t+1} that result from

[38]The distribution of $1 - e^{Z_t}$ is shown in Panel A of Figure 7 in Wachter 2013. We rely on the same data.

the quadrature. Figure 27 replicates figures 1 through 6 from Wachter (2013), and compares them to the results from our discrete time version of the model. The plots show that the discretized model provides a good approximation to its continuous time counterpart.

The calculation of option prices is similar to the one for Barro (2009). We condition on the disaster size realization, evaluate integrals over ε^c_{t+1} and ν_{t+1} analytically, integrals over ε^l_{t+1} with the aforementioned quadrature method, and finally evaluate expectations over disaster realizations based on their empirical distribution.

B.4.5 Schreindorfer (2020)

The disappointment aversion model of Schreindorfer (2020) admits analytical solutions for all asset prices, including options. We solve the model based on the replication code on David Schreindorfer's website, www.davidschreindorfer.com.

B.5 Computing $EP(x)$ in models

Computing $EP(x)$ in models simply requires estimates of f and f^*. This section provides the necessary detail for how we compute these objects in each model, and how they are used to construct $EP(x)$.

To begin, we construct an equally-spaced grid $r_1, \ldots, r_N$ for returns. The edges of the grid differ by model, but are set wide enough to cover the entire model-implied return density. Table 13 provides the bounds used in each model. We set $N = 10,000$. Our estimates of f and f^* are defined on this 10,000 point grid.

Table 13. Return Grid Bounds

Model	Lower	Upper
Campbell-Cochrane (1999)	-0.35	0.35
Bansal-Yaron (2004)	-0.35	0.35
Barro (2009)	-0.99	$3\sqrt{1/12}$
Wachter (2013)	-0.99	$3\sqrt{1/12}$
Schreindorfer (2020)	-0.60	0.50

We obtain the unconditional return density $f(R)$ in two steps. We first obtain a highly accurate estimate of $f_t(R; \xi_t)$ via quadrature (where ξ_t is again being used to denote the model state vector). If the model is IID, this is equivalent to f, so we are done. If the model is not IID, we take the pointwise average of the conditional densities over the ergodic distribution of the model state space. Specifically, we simulate the model state variables L periods to obtain points from the ergodic distribution, and we average the densities as in Equation (B.6).

$$f(r_i) = \frac{1}{L}\sum_{s=1}^{L} f_t(r_i;\xi_s), \quad i = 1, ..., N \tag{B.6}$$

When needed, we interpolate $f_t(R;\xi_t)$ across state variables using cubic spline interpolation. We set $L = 10^7$.

We obtain f^* by first computing $f_t^*(R;\xi_t)$. Breeden and Litzenberger 1978 show that $f_t^*(R;\xi_t)$ can be computed as the second derivative of put option prices, $P(X,\xi_t)$, with respect to the strike price. Appendix Section B.4 contains detail regarding our computation of put prices in each model. We compute put option prices for each point in the return grid, $r_1, \ldots, r_N$, as well as at the points $r_i \pm \epsilon$.[39] We set $\epsilon =$ 1e-4. We compute $f_t^*(r_i;\xi_t)$ using second order central differences as in Equation (B.7).

$$f_t^*(r_i;\xi_t) = (1 + R_f(\xi_t))\frac{P(1 + r_i - \epsilon, \xi_t) + P(1 + r_i + \epsilon, \xi_t) - 2P(1 + r_i, \xi_t)}{\epsilon^2} \tag{B.7}$$

We compute f^* by averaging $f_t^*(R;\xi_t)$ in the same manner used in Equation (B.6).

Having estimates of f and f^* in hand, $EP(x)$ is given by

$$EP(x) = \frac{\sum_{i=1}^{N} \mathbf{1}\left\{r_i \leq x\right\} r_i(f(r_i) - f^*(r_i))\delta}{\sum_{j=1}^{N} r_j(f(r_j) - f^*(r_j))\delta}, \tag{B.8}$$

where $x \in [0, \infty)$.

[39] Recall that we compute put-to-spot ratios with moneyness $X = K/S_t$ where K is the strike price and S_t is the underlying spot price. Because our return grid represents net ex-dividend returns, $X_k = 1 + r_k$ for r_k in our return grid.

APPENDIX C

APPENDIX TO CHAPTER 4

C.1 Data

The data on household portfolios come from the Board of Governor's of the Federal Reserve Survey of Consumer Finances, using all regular surveys between 1989 and 2016.[40] The Survey of Consumer Finances (SCF) data contains approximately 4,000 households every three years from 1989 to 2016.[41] The survey contains information related to demographics, income, net worth, investments, private businesses, and many other financial categories. For the purposes of this paper, it is necessary to map the survey variables into a classification system consisting of only safe or risky assets. While there are many ways to do this, I will follow the classification system used in Chang, Hong, and Karabarbounis (2018).[42] For convenience, I will summarize the procedure here. The interested reader is invited to refer to the source for additional detail.

Safe assets are the sum of checking accounts, money-market accounts, savings accounts, certificates of deposits, life insurance, and the safe portions of mutual fund investments, pensions, bonds, trusts, annuities, and IRAs. Risky assets are the sum of brokerage accounts, direct stockholdings, non-actively managed businesses, and the risky portions of mutual fund investments, pensions, bonds, trusts, annuities, and IRAs. For mutual fund investments, pensions, bonds, trusts, annuities, and IRAs, the risky/safe split is imputed based on answers to auxiliary survey questions or assumed to be a conservative value if no additional information is available. Financial wealth is the sum of risky and safe assets, and the risky asset portfolio share is then the fraction of financial wealth comprised of risky assets. All of the analysis in my paper will be conditional on participation in the risky asset markets, meaning that I condition on risky asset shares being above 0%.

C.2 Additional Model Details

Because the utility function is time-separable, the first order condition for consumption, combined with the Envelope condition, can be inverted to become a function of tomorrow's consumption, today's portfolio allocation, and tomorrow's realization of

[40]The additional surveys from the 2007-09 Panel Survey of Consumer Finances were not included in the analysis.

[41]The data is available from the Board of Governors of the Federal Reserve System website: https://www.federalreserve.gov/econres/scfindex.htm.

[42]Other ways to classify assets of the household balance sheet can be seen in Gomes and Michaelides (2005) and Wachter and Yogo (2010), for example.

income and return shocks, as shown in (C.1).

$$C_t = u'^{-1}(\tilde{\delta}_t E_t[u'(C_{t+1})(R_f + \theta_t \pi_{t+1})|z_t]). \tag{C.1}$$

Carroll (2006) showed that by recasting the problem in terms of end-of-period savings it is possible to eliminate the dependence in (C.1) on any current period endogenous variables in some settings. This method, known as the endogenous grid method, when possible to implement, can provide significant improvements in model solution speed and accuracy. Because the endogenous grid method provides numerous advantages over traditional value function iteration, I will use it to numerically solve the model. Given a consumption and wealth policy for period $t+1$, and a portfolio allocation policy for period t, we can evaluate the integrals in (11-12) by interpolating between the discrete pairs of consumption and wealth.

Because the first order condition for θ_t^* does not directly depend on today's consumption, it is helpful to solve for θ_t^* first. The first order condition for θ_t^* must be solved using a numerical root-finding algorithm. Once we find the optimal risky asset share, we can simply plug it into equation (C.1) to get the optimal consumption level. Because we are using the endogenous grid method, from the optimal consumption rule we obtain the optimal cash-on-hand rule as well (X_{t+1}^*).

The model is solved by the Endogenous Grid Method of Carroll (2006). Begin by specifying a grid for end-of-period assets, a_t. I use $N_a = 350$ points with triple exponential spacing, which creates a dense grid of points for low values of cash-on-hand where the policy functions have more curvature.

It is also necessary to create a grid for the income shock process z_t. In previous papers, this process takes the form of a permanent shock and thus the problem can be detrended by the current shock value. In those cases, it is not necessary to specify a grid of current shock values because it is no longer a state variable. In the HIP model, the shock is a state variable. Unfortunately, existing techniques for discretizing autoregressive markov processes, like Tauchen and Hussey (1991), do not apply. Therefore, I follow a simulation-based approach to generate an appropriate grid of z_t values for which to solve the problem.

I begin by simulating the income shock process z_t for a large number of households of a single HIP type. Guvenen et al. (2016) then use a linearly spaced grid between the unconditional extrema of z from the simulated panel. However, the distribution for z is highly negatively skewed, and I found that this approach generated a large number of points in regions of the state space that were infrequently visited. Therefore, I compute the unconditional empirical distribution function for the simulated z values and define my grid in terms of quantiles of this distribution. Specifically I use the minimum observed value, maximum value, and a linearly spaced grid of 48 points (so that $N_z = 50$) between the 0.01% and the 99.99% of the distribution. This technique ensures that the bulk of the values in the z grid are frequently visited during the simulation, which reduces error due to discretization and interpolation.

After the state space is constructed the problem is solved by iterating backwards in time starting from the terminal conditions. For each (a_t, z_t) pair, I need to solve for (X_{t+1}, C_t, θ_t). I begin by solving for θ_t using an iterated (or zooming) grid search. That is, define a coarse grid of potential candidates for θ_t^*. Find the one that comes closest to satisfying the FOC determining θ_t^*. Define a new grid surrounding the previous optimal θ_t and repeat until the range of the grid is below a set tolerance (I use 0.01). With θ_t^* in hand, I can compute C_t^* directly from (C.1), which also gives the optimal cash-on-hand for next period X_{t+1}^*. I use linear interpolation in all cases to evaluate the optimal policies at points that are not on the state space grid.

Not surprisingly, the crux of the HIP model lies in the *heterogeneity*. The parameters $(\alpha_k, \beta_k, \sigma_{z,k})$ are individual-specific. We assumed that these parameters have a certain distribution throughout the cross-section of households.

$$
\begin{aligned}
(\alpha_k, \beta_k) &\sim MVN(0.0, \Sigma), \qquad \Sigma = \left[\sigma_\alpha^2, \sigma_\alpha\sigma_\beta\rho_{\alpha\beta}; \sigma_\alpha\sigma_\beta\rho_{\alpha\beta}, \sigma_\beta^2\right] \\
\ln \sigma_{z,k} &\sim N(\ln \bar{\sigma} - \frac{\sigma_i^2}{2}, \sigma_i^2)
\end{aligned}
$$

It is necessary to additionally discretize the household type distribution to reduce the dimensionality of the problem. To this end, I use 4 (α_k, β_k) types and 3 $\sigma_{z,k}$ types for a total of 12 ex ante heterogeneous households types. I solve the problem for each of the 12 types and use the Gaussian weights associated with each type to determine the appropriate number of households of each type to simulate. I combine all of the simulated panels together to compute economy-wide statistics, such as the average portfolio share and the cross-sectional quantiles of the portfolio shares at each age.

I use the programming language Julia and the supercomputing resources at Arizona State University to solve the model in parallel.

C.3 Robustness

Housing

One omission from the definition of household assets in the results presented in the paper is housing. As an alternative definition, I will include investment housing as an additional risky asset and primary, owner-occupied housing as an additional safe asset. Figure 28 plots the cross-section of household portfolio allocations using this alternative definition, with the statistics being computed as in Figure 20. There continues to be significant heterogeneity in the cross-section of portfolio shares at each stage in the life cycle. I do not include housing as an additional asset in the model, although it would be interesting it explore it following the examples of Cocco (2004) and Gomes and Michaelides (2005).

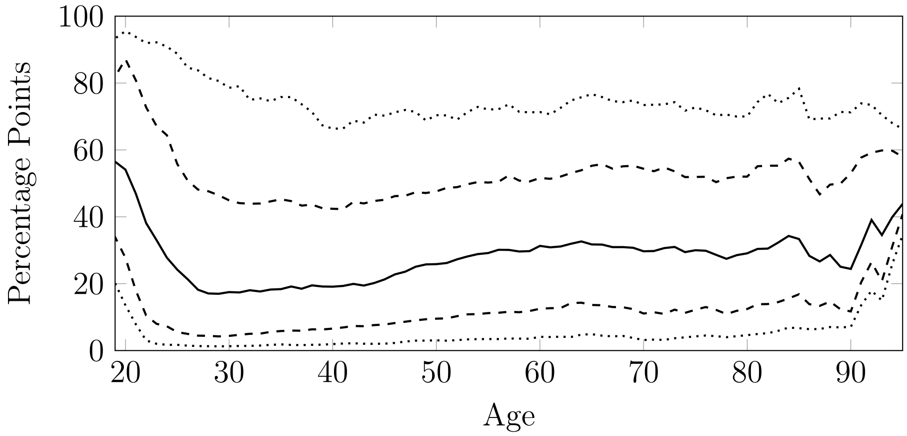

Figure 28. Cross-section of SCF household risky asset shares (alternative definition), conditional on participation.

Taxes

The income processes used in the paper were calibrated to pre-tax income data, whereas most previous papers used after-tax incomes. A result of this is that wealth accumulation is artificially high in the Gaussian and HIP models in this paper. Adding a flat income tax of 30% has limited effects on the portfolio allocation decisions, but brings the models more in-line with wealth levels seen in the SCF data. Table 14 displays percentiles of the unconditional wealth distribution from the model solved with a 30% income tax applied in every period and Figure 29 shows the mean wealth at each age.

Table 14. Percentiles of the unconditional wealth distribution. II

	$P10$	$P25$	$P50$	$P75$	$P90$
SCF Data	5.7	17.6	59.4	189.9	539.3
Gaussian Model	21.5	63.4	169.3	362.7	649.6
HIP Model	20.2	62.6	211.2	476.2	912.1

Note. Wealth in thousands of US Dollars. For the SCF, I compute each statistic using sampling weights within each of the 5 replicates for each survey year. I then average across the 5 replicates for a given survey year and then average across all surveys. Data is the Survey of Consumer Finances from 1989 to 2016.

Figure 29. Mean wealth for the two models and the SCF data.

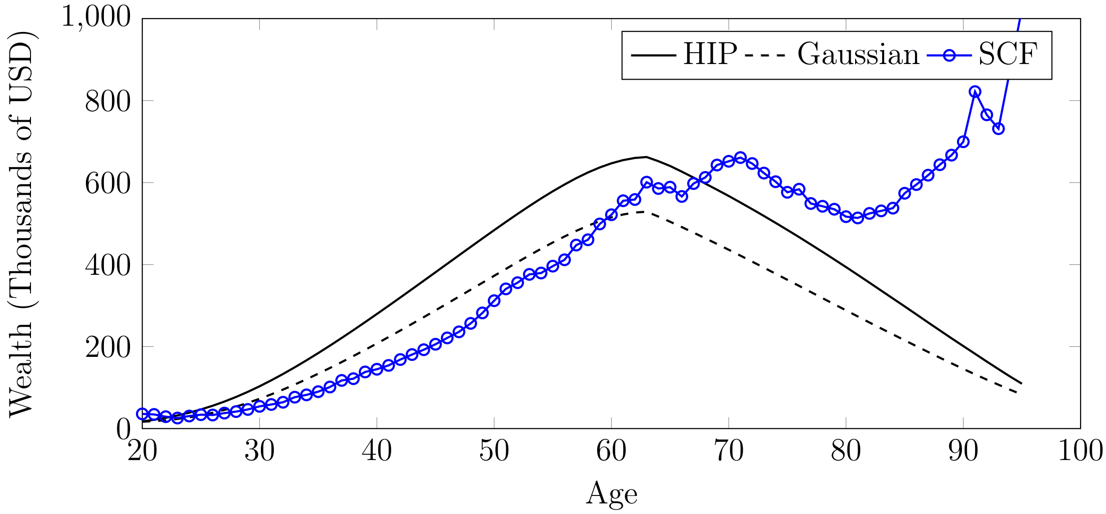

Note. Household wealth by age from the HIP model (solid line), standard model (dashed line), and the Survey of Consumer Finances data (blue marked line). For the SCF, I compute the mean by age using sampling weights within each of the 5 replicates for each survey year. I then average across the 5 replicates for a given survey year and then average across all surveys. For each age t, I group households aged $t-2$ to $t+2$ in order to increase sample size. Data is the Survey of Consumer Finances from 1989 to 2016.

APPENDIX D

CO-AUTHORSHIP STATEMENT

Chapter 3 of this dissertation, titled Dissecting the Equity Premium, forms the core of a paper of the same name, co-authored with Dr. David Schreindorfer and included in this document with his permission.