Origins of Stock Market Fluctuations*

Daniel L. Greenwald
NYU

Martin Lettau
UC Berkeley, CEPR and NBER

Sydney C. Ludvigson
NYU and NBER

First draft: June 6, 2012
This draft: June 16, 2014

*Greenwald: Department of Economics, New York University, 19 W. 4th Street, 6th Floor, New York, NY 10012; Email: dlg340@nyu.edu. Lettau: Haas School of Business, University of California at Berkeley, 545 Student Services Bldg. #1900, Berkeley, CA 94720-1900; E-mail: lettau@haas.berkeley.edu; Tel: (510) 643-6349, http://faculty.haas.berkeley.edu/lettau. Ludvigson: Department of Economics, New York University, 19 W. 4th Street, 6th Floor, New York, NY 10012; Email: sydney.ludvigson@nyu.edu; Tel: (212) 998-8927; http://www.econ.nyu.edu/user/ludvigsons/. The authors are grateful to Jarda Borovicka and Eric Swanson and to seminar participants at the University of Chicago Booth School of Business and NYU for helpful comments.

Origins of Stock Market Fluctuations

Abstract

Three mutually uncorrelated economic shocks that we measure empirically explain 85% of the quarterly variation in real stock market wealth since 1952. We use a model to show that they are the observable empirical counterparts to three latent primitive shocks: a total factor *productivity* shock, a *risk aversion* shock that is unrelated to aggregate consumption and labor income, and a *factors share* shock that shifts the rewards of production between workers and shareholders. On a quarterly basis, risk aversion shocks explain roughly 75% of variation in the log difference of stock market wealth, but the near-permanent factors share shocks plays an increasingly important role as the time horizon extends. We find that more than 100% of the increase since 1980 in the deterministically detrended log real value of the stock market, or a rise of 65%, is attributable to the cumulative effects of the factors share shock, which persistently redistributed rewards away from workers and toward shareholders over this period. Indeed, without these shocks, today's stock market would be about 10% lower than it was in 1980. By contrast, technological progress that rewards both workers and shareholders plays a smaller role in historical stock market fluctuations at all horizons. Finally, the risk aversion shocks we identify, which are uncorrelated with consumption or its second moments, largely explain the long-horizon predictability of excess stock market returns found in data. These findings are hard to reconcile with models in which time-varying risk premia arise from habits or stochastic consumption volatility.

1 Introduction

Asset pricing theorists have long been concerned with explaining stock market *expected returns*, typically measured over monthly, quarterly or annual horizons. This is important because empirical evidence suggests that variation in the stock market price-dividend ratio is driven almost entirely by expected excess return variation (i.e., forecastable movements in equity premia).[1] Far less attention has been given to understanding the real (adjusted for inflation) *level* of the stock market, i.e., stock price variation, or the cumulation of returns over many decades. The profession spends a lot of time debating which risk factors drive expected excess returns, but little time investigating why real stock market wealth has evolved to its current level compared to 30 years ago. To understand the latter, it is necessary to probe beyond the role of stationary risk factors and short-run expected returns, to study the primitive economic shocks from which all stock market (*and* risk factor) fluctuations originate.

To see why, consider that some economic shocks may have tiny innovations but permanent or near-permanent effects on cash flows. Under rational expectations, permanent cash flow shocks have no influence on the price-dividend *ratio* (they are incorporated immediately into both prices in the numerator and dividends in the denominator), but they can have a dramatic influence on real stock market wealth as the decades accumulate. Such shocks are the sources of stochastic trends in stock prices that are by definition impossible to predict and not reflected in expected returns. On the other hand, fluctuations in expected returns may be associated with movements in risk premia and can persistently shift the real value of the stock market around its long-term trend. But because these fluctuations are transitory, their impact eventually dies out. Stock market wealth evolves over time in response to the cumulation of both transitory expected return and permanent cash flow shocks. The crucial unanswered questions are, what are the economic sources of these shocks? And what have been their relative roles in evolution of the stock market over time?

The objective of this paper is to address these questions. We begin by identifying three mutually orthogonal observable economic shocks that explain the vast majority (over 85%) of quarterly fluctuations in real stock market wealth since the early 1950s. Econometrically, these shocks are measured as specific orthogonal movements in consumption, labor income, and asset wealth (net worth), identified from a cointegrated vector autoregression (VAR) and extracted using a standard recursive identification procedure. Roughly speaking, this

[1] Expected dividend growth and expected short-term interest rates play little role empirically in price-dividend ratio variation (Campbell (1991); Cochrane (1991); Cochrane (2005); Cochrane (2008)).

methodology decomposes movements in log stock market wealth into those associated with (1) shocks to log aggregate consumption, (2) shocks to the log labor income-consumption ratio, holding *fixed* consumption, and (3) shocks to log stock wealth holding fixed *both* log consumption and log labor income. We investigate how these shocks have affected stock market wealth over time, with special attention paid to their relative importance over short versus long time horizons.

We then address the question of what these observable VAR shocks represent economically. Providing an economic interpretation of the shocks requires a theoretical framework. A common theoretical approach for explaining stock market behavior is to assume the existence of a representative agent who consumes the aggregate consumption stream.[2] But we find that understanding the observable sources of variation we find in stock prices in post-war data drives us to consider a framework with heterogeneity. So this paper develops a general equilibrium model of two types of agents: shareholders and workers. The representative shareholder in the model is akin to a large institutional investor or wealthy individual who derives all income from investments. The representative worker consumes a stream of labor income every period. Economic fluctuations in the model originate from three mutually orthogonal primitive shocks that differentially affect each type: a permanent total factor *productivity* (TFP) shock that governs the state of factor neutral technological progress and propels aggregate (shareholder plus worker) consumption, a near-permanent *factor shares* shock that reallocates the rewards from production between shareholders and workers without affecting the size of those rewards, and an independent shock to shareholder *risk aversion* that moves the stochastic discount factor pricing assets independently of stock market fundamentals or real variables such as consumption and labor income. The modifier "independent" in reference to shareholder risk aversion refers to the independence in the model of such shocks from macroeconomic fundamentals (dividends, earnings, consumption, labor income, output). Our findings indicate that this independence is essential for explaining the observed time-variation in the reward for bearing stock market risk, implying that quantitatively large component of risk premia fluctuations is acyclical. One interpretation of this is that risk premia are driven by intangible information that is largely unrelated to the current economic state. In this sense, our findings contrast with classic earlier studies that emphasized the countercyclicality of risk premia (e.g., Fama and French (1989)). We discuss this further below.

[2] This approach is widely taken, including among the leading asset pricing theories of the day. Influential examples include Campbell and Cochrane (1999) and Bansal and Yaron (2004).

The model is then employed to show that the mutually orthogonal VAR innovations that explain almost all stock market variation are the observable empirical counterparts to the latent primitive shocks in the theoretical framework. Specifically, we show that, if the model generated the data, the total factor productivity shock would be revealed as a *consumption shock* in the VAR, namely a movement in log consumption, c_t, that contemporaneously affects both log labor income, y_t, and log asset wealth, a_t, where all three move in the same direction. The factors share shock would be revealed as a *labor income shock* that moves a_t in the opposite direction of y_t but is restricted to have no contemporaneous impact on c_t. And the risk aversion shock would be revealed as a *wealth shock* that moves a_t but is restricted to have no contemporaneous impact on either c_t or y_t. We show that the dynamic responses to these mutually orthogonal VAR innovations produced from model generated data are remarkably similar to those produced from historical data. We refer to the three mutually orthogonal VAR innovations (consumption, labor income and wealth) interchangeably as productivity or TFP, factors share, and risk aversion innovations, respectively.

How have these shocks affected stock market wealth over time? We find that the vast majority of short- and medium-term stock market fluctuations in historical data are driven by risk aversion shocks, revealed as movements in wealth that are orthogonal to consumption and labor income, both contemporaneously (an identifying assumption), and at all subsequent horizons (a result). Although transitory, these shocks are quite persistent with a half life of over four years. On a quarterly basis, they explain approximately 75% of variation in the log difference of stock market wealth, but their contribution declines as the horizon extends. These facts are well explained by the model, in which the orthogonal wealth shocks originate from independent shifts in investors' willingness to bear risk.

At longer horizons, the relative importance of the shocks changes. The factors share shock explains a negligible fraction of variation in the stock market over shorter time horizons, but because its innovations are nearly permanent, it plays an increasingly important role as the time horizon extends. A spectral decomposition of variance by frequency shows that this shock explains virtually none of the variation in the real level of the stock market over cycles of a quarter or two, but it explains roughly 40% over cycles two to three decades long. These facts are well explained by the model economy, which is subject to small but highly persistent innovations that shift the allocation of rewards between shareholders and workers independently from the magnitude of those rewards.

By contrast, factor neutral productivity shocks, revealed as consumption shocks empirically, play a small role in the stochastic fluctuations of the stock market at all horizons, once a deterministic trend is removed. The model is consistent with this finding. The crucial aspect of the model that makes it consistent with this finding is its heterogeneous agent specification: given capital's smaller role in the production process, most gains and losses from total factor productivity shocks accrue to workers rather than shareholders, so their effect on stock market wealth is smaller than that of the other two shocks. This finding is a direct contradiction of representative agent asset pricing models where shocks that drive aggregate consumption plays a central role in stock market fluctuations.

As an example of the magnitude of these forces for the long-run evolution of the stock market, we decompose the percent change since 1980 in the deterministically detrended real value of stock market wealth that is attributable to each shock. The period since 1980 is an interesting one to consider, in which the cumulative effect of the factor shares shock persistently redistributed rewards away from workers and toward shareholders. (The opposite occurs from the mid 1960s to mid 1980s.) After removing a deterministic trend, we find that the cumulative effects of the factors share shocks have resulted in a 65% increase in real stock market wealth since 1980, an amount equal to 110% of the total increase in stock market wealth over this period. Indeed, without these shocks, today's stock market would be roughly 10% lower than it was in 1980. This finding underscores the extent to which the long-term value of the stock market has been profoundly altered by forces that reallocate the rewards of production, rather than raise or lower all of them.

Our calculations imply that an additional 38% of the increase in the detrended real value of the stock market since 1980, or a rise of 22%, is attributable to the cumulative effects of risk aversion shocks, which were on average lower in the last 30 years than earlier in the post-war period. By contrast, the cumulative effects of TFP shocks have made a negative contribution to change in stock market wealth since 1980, once a deterministic trend is removed. The importance of the TFP shock is uncharacteristically large over this period, a direct consequence of the string of unusually large negative draws for consumption in the Great Recession years from 2007-2009. These shocks account for -38% of the increase since 1980. Together, the three mutually orthogonal economic shocks we identify explain almost all of the increase in deterministically detrended real stock market wealth since 1980. (Specifically, they account for 110% of the increase, with the remaining -10% accounted for by a residual.)

Finally, our findings speak to the question of why the stock market is predictable. The process for shareholder risk aversion in the model is highly non-linear, with shareholders close to risk-neutral most of the time but subject to rare "crises" in their willingness to bear risk, captured in the model by infrequent, large spikes in risk aversion. These states lead to rare "flights to safety" in which the market crashes. Below we explain why this strong non-linearity is important for simultaneously matching the range of empirical evidence we study here. But even in normal times, the time-varying expectation that risk tolerance could crash in the future generates continuous fluctuations in the price-dividend ratio that are associated with forecastable variation in excess stock market returns, or a time-varying equity risk premium. Risk premia variation is observable as the orthogonal wealth shocks from the empirical VAR, which in turn generate excess return predictability in a linear forecasting regression. This model implication is closely matched in the data. In particular, the predictive content for long horizon excess stock market returns of common stock market forecasting variables (such as the price-dividend ratio or cay_t) is found to be largely subsumed by the information in lags of the independent risk aversion/wealth shock. These variables forecast excess returns only because they are correlated with the VAR wealth shocks.

An important aspect of these results is that the time-variation in the reward for bearing risk, both in the model and the data, is divorced from fluctuations in traditional macroeconomic fundamentals. The wealth shocks we identify are by construction orthogonal to movements in consumption and labor income contemporaneously, and, as we show, at all future horizons (a result).We also find that these innovations bear little relation to other traditional macroeconomic fundamentals such as dividends, earnings, consumption volatility, or broad-based macroeconomic uncertainty, and none of these variables forecast equity premia. These findings are hard to reconcile with models in which time-varying risk premia arise from habits (which vary with innovations in consumption), stochastic consumption volatility, or consumption uncertainty.

The rest of this paper is organized as follows. The next section discusses related literature. Section 3 describes the econometric procedure and data used to identify the three mutually uncorrelated empirical shocks from a VAR on consumption, labor income, and asset wealth. Section 4 describes the theoretical model that we use to interpret these shocks. Section 5 presents our findings, which are of two forms. The first are results on the performance of the model, including summary statistics for standard asset pricing implications. This section also demonstrates that the mutually orthogonal VAR innovations described in Section 4 are the observable empirical counterparts to the latent primitive shocks in the model. The second set of results studies the relative role of the observable shocks in historical stock market fluctuations, with special attention paid to how these roles depend on the time horizon over which one measures a change in stock market wealth. Here we compare the role of each shock in the dynamic responses and variance decompositions of stock market wealth when estimated from historical data with the same statistics when estimated from model-generated data and show that they are remarkably similar along a number of dimensions discussed below. The final subsection investigates the question of why stock returns are predictable. We show that the independent risk aversion shocks we identify largely explain why long-horizon excess stock market returns are predictable by common forecasting variables such as the log price-dividend ratio or the consumption-wealth variable cay_t. Section 6 briefly discusses the link between the factors share shock and economic inequality. Section 7 concludes.

2 Related Literature

The empirical part of this paper builds on Lettau and Ludvigson (2013). That paper provided empirical evidence in a purely statistical model, studying (a rotation of) the three VAR innovations described here and their relationship to different components of household wealth, consumption and labor income. The contribution of this paper is to provide an economic interpretation of these innovations and a detailed investigation of their implications for the stock market. Our model is also related to the work of several recent papers that have emphasized the weak empirical correlation between stock market behavior and innovations to consumption growth or its second moments (Duffee (2005), Albuquerque, Eichenbaum, and Rebelo (2012), Lettau and Ludvigson (2013)). And there is an important earlier literature in asset pricing that identified and distinguished cash-flow from discount rate "shocks" (e.g., Campbell (1991); Cochrane (1991)). This work was central to our understanding of how innovations in stock returns are related to forecastable movements in returns as compared to dividend growth, but it is silent on the underlying *economic* mechanisms that drive these forecastable changes. It is precisely these primitive economic shocks that are the subject of this paper.

We build on an earlier literature that emphasizes the importance of limited stock market participation for explaining stock return data (Mankiw and Zeldes (1991); Vissing-Jorgensen (2002); Guvenen (2009); Lettau and Ludvigson (2009); Malloy, Moskowitz, and Vissing-Jorgensen (2009)). Much of this literature uses data on household surveys to measure stockholder's consumption, such as the Consumer Expenditure Survey (CEX). While such survey's certainly contain useful information, they also have limitations. A drawback with the CEX is that answers to asset questions are missing for more than half the sample. Moreover, institutional ownership has reached almost half of total equity ownership in recent data

(Kacperczyk, Nosal, and Stevens (2014)), suggesting that representative household surveys may be missing important elements of shareholder risk.

The factors share element of our paper is related to a separate macroeconomic literature that examines the long-run variation in labor's share, but has not studied the asset pricing implications of these movements (e.g., Karabarbounis and Neiman (2013)). An exception is the theoretical study of Lansing (2014), which we became aware of after circulating the first draft of this paper. Lansing presents a model similar in spirit to the one contained here, with a shareholder-worker redistribution shock that affects asset values. The factors share findings in this paper echo those from previous studies that use very different methodologies but like this paper find that returns to human capital are negatively correlated with those to non-human (stock market) capital (Lustig and Van Nieuwerburgh (2008); Lettau and Ludvigson (2009); Chen, Favilukis, and Ludvigson (2014))).

Our findings on risk premia variation are potentially consistent with other theories in which time-variation in risk premia is largely divorced from fluctuations in macroeconomic fundamentals and instead associated with independent market crises. Examples include the ambiguity aversion framework of Bianchi, Ilut, and Schneider (2013), in which a component of investor confidence varies in a manner like our risk aversion that is independent of shocks to economic fundamentals, or in models of rare events in which the probability of consumption disaster is a random variable independent of normal-times consumption shocks (e.g., Wachter (2013)), or intermediary-based models in which intermediaries' risk-bearing capacity varies independently of macroeconomic fundamentals (e.g., Gabaix and Maggiori (2013)), or models where restrictions on intermediaries risk-bearing capacity result in very strong nonlinearities between risk premia and macroeconomic state variables (e.g., Brunnermeier and Sannikov (2012); He and Krishnamurthy (2013); Muir (2014)).[3] Of course, the risk-bearing capacity of intermediaries is determined by the households who supply them with capital, and so must ultimately be traced to investors' willingness to bear risk. Each of theories, like the one presented here, describe mechanisms for generating risk premia variation that is not closely linked to macroeconomic fundamentals. We view the econometric evidence presented here as contributing to a growing body that forms the basis of an empirical rationale for such mechanisms. We discuss this further in the conclusion.

[3] In general these non-linearities must be quite extreme. In the model studied here, the extreme non-linearity of risk aversion is enough to explain why fluctuations in risk premia appear largely divorced from macroeconomic fundamentals even when we force the risk-aversion shock to be perfectly negatively correlated with the TFP shock. But, as we explain below, it is very difficult to match the full range of evidence presented here without an independent shock to risk aversion.

3 Econometric Analysis: Three Mutually Orthogonal Shocks

This section describes the empirical analysis used to identify three mutually orthogonal observable shocks from data on aggregate consumption, labor income and household wealth.

3.1 Data

Many empirical details of the estimation that follows are covered in Lettau and Ludvigson (2013). Here we outline the main elements and refer the reader to that paper for more information.[4]

We consider a cointegrated vector of variables in the data, denoted $\mathbf{x}_t = (c_t, y_t, a_t)'$, where c_t is log of real, per capita aggregate consumption, y_t is log of real, per-capita labor income, and a_t is log of real, per-capita asset wealth. Throughout this paper we use lower case letters to denote log variables, e.g., $\ln(A_t) \equiv a_t$. Lettau and Ludvigson (2013) provide updated evidence of a cointegrating relation among these variables, which can be motivated by considering the long-run implications of a standard household budget constraint (see Lettau and Ludvigson (2001)).

The Appendix contains a detailed description of the data used in this study. The log of asset wealth, a_t, is a measure of real, per capita household net worth, which includes all financial wealth, housing wealth, and consumer durables. It is compiled from the flow of funds accounts by the Board of Governors of the Federal Reserve.

We study the implications of the empirical shocks identified from the system $(c_t, y_t, a_t)'$ and subsequently relate these shocks to stock market wealth. We denote the log of stock market wealth s_t. Stock market wealth is a component of total asset wealth. Corporate equity was 23% of total asset wealth in 2010, and 29% of net worth. For comparison, we will often also study the implications of these same shocks for the Center for Research in Securities Prices (CRSP) value-weighted stock price index. We denote the log of the CRSP value-weighted stock price index p_t.

We use the log of real, per capita, expenditures on nondurables and services (excluding shoes and clothing), as a measure of c_t. From the household's budget constraint, an internally consistent cointegrating relation among c_t, y_t, and a_t, may then be obtained if we assume that the log of (unobservable) real total flow consumption is cointegrated with the log of real

[4]That paper showed how the innovations we study here can be econometrically identified as disturbances distinguished by whether their effects are permanent or transitory. An additional rotation of innovations was required for this interpretation, but the shocks obtained there are perfectly correlated with those obtained here and discussed below.

nondurables and services expenditures (Lettau and Ludvigson (2010)). The log of after-tax labor income, y_t, is also measured in real, per capita terms. Lettau and Ludvigson (2013) presents empirical evidence supportive of a single cointegrating relationships between c_t, a_t, and y_t in the post-war data used in the study.[5] Our data are quarterly and span the first quarter of 1952 to the third quarter of 2012.

3.2 Empirical Implementation

Before explaining the details of the estimation, we discuss a practical distinction between the data and the model. The model developed below is intended to focus on the implications of the empirical shocks for stock market wealth. As such, it has one form of risky capital (equity) and a risk-free bond in zero net supply. It follows that, in the model, all wealth is stock market wealth, which is the same as total wealth, which is the same as net worth. In historical data, total wealth contains non-equity forms of wealth, so total wealth and stock market wealth are two different variables. As noted above, stock market wealth accounts for about 23% of total wealth in 2010. In the data we distinguish the two by denoting log of asset wealth (net worth) by a_t and log of stock market wealth by s_t. In the model, $a_t = s_t$. In constructing the empirical VAR innovations discussed next, we use system of variables that contains consumption, labor income and total asset wealth a_t, and then subsequently relate these innovations to stock market wealth. We do *not* construct these innovations by restricting analysis to how consumption and labor income move only with stock market wealth. We do this for two reasons. First, a (factor neutral) TFP shock should affect the value of all productive capital, so a system that identifies such a shock from the data should include total wealth. If TFP shocks affect non-stock wealth but these components are omitted from the system, this could lead to spurious estimates of productivity and its dynamics, which would also contaminate estimates of the factors share and risk aversion shocks that are presumed orthogonal to the TFP shock. Second, consumption and labor income are cointegrated with total wealth, as expected from theory (Lettau and Ludvigson (2001)), but

[5] The data provide no evidence of a second linearly independent cointegrating relation (there can be at most two). In particular, bivariate log ratios of these variables appear to contain trends in our sample. Economic models often imply that bivariate log ratios (e.g., $y_t - a_t$) are stationary, and the model below assumes as much. In finite samples, it is impossible to distinguish a highly persistent but stationary series from a non-stationary one. For the empirical system, we follow the advice of Campbell and Perron (1991) and empirically model only the single, trivariate cointegrating relation for which we find direct statistical evidence of in our sample. For the model, we set parameter values that imply the log ratio $y_t - a_t$ is stationary but deviations from the common trend in y_t and a_t are so persistent one could not reject a unit root in samples of the size we currently face, consistent with the data.

there is no implication that these variables should be (or are) cointegrated with stock market wealth by itself, a component of total wealth. It is important to control empirically for these long-run relationships by imposing the restrictions implied by cointegration in a VAR for which wealth shares a common trend with consumption and labor income. When this is done, it will then be an empirical matter how closely the identified VAR shocks are related to stock market wealth, which is the subject of an extensive investigation below. Notice that there is no implication that the shocks identified from the c_t. y_t, a_t empirical system should explain all or even most of the variation in stock market wealth s_t. Although s_t can be related to these shocks, there will be an unexplained residual that in principal could be quite important, as we explain below.

Identification of the three mutually orthogonal empirical disturbances is achieved in several steps. First, we assume all of the series contained in $\mathbf{x}_t$ are first order integrated, or I(1), an assumption confirmed by unit root tests, available upon request. The cointegrating coefficient on consumption is normalized to one, and we denote the single cointegrating vector for $\mathbf{x}_t = [c_t, a_t, y_t]'$ as $\boldsymbol{\alpha} = (1, -\alpha_a, -\alpha_y)'$. The cointegrating parameters α_a and α_y are estimated using dynamic least squares, which generates "superconsistent" estimates of α_a and α_y (Stock and Watson, 1993).[6] We estimate $\widehat{\boldsymbol{\alpha}} = (1, -0.18, -0.70)'$. The Newey and West (1987) corrected t-statistics for these estimates are 20 and 56, respectively.

Second, a cointegrated VAR (or vector error correction mechanism–VECM) representation of $\mathbf{x}_t$ is estimated taking the form

$$\Delta\mathbf{x}_t = \boldsymbol{\upsilon} + \boldsymbol{\gamma}\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1} + \boldsymbol{\Gamma}(L)\boldsymbol{\Delta}\mathbf{x}_{t-1} + \mathbf{u}_t, \tag{1}$$

where $\Delta\mathbf{x}_t$ is the vector of log first differences, $(\Delta c_t, \Delta a_t, \Delta y_t)'$, $\boldsymbol{\upsilon}$, and $\boldsymbol{\gamma} \equiv (\gamma_c, \gamma_a, \gamma_y)'$ are (3×1) vectors, $\boldsymbol{\Gamma}(L)$ is a finite order distributed lag operator, and $\widehat{\boldsymbol{\alpha}} \equiv (1, -\widehat{\alpha}_a, -\widehat{\alpha}_y)'$ is the (3×1) vector of previously estimated cointegrating coefficients.[7] The term $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$ gives last period's equilibrium error, or cointegrating residual, a variable we denote with $cay_t \equiv \widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$. The coefficients $\boldsymbol{\gamma}$ are the vector of "adjustment" coefficients that tells us which variables subsequently adjust to restore the common trend when a deviation occurs. Throughout this paper, we use "hats" to denote the estimated values of parameters.

[6] We use eight leads and lags of the first differences of Δy_t and Δa_t in the dynamic least squares regression. Monte Carlo simulation evidence in both Ng and Perron (1997) and our own suggested that the DLS procedure can be made more precise with larger lag lengths.

[7] Standard errors do not need to be adjusted to account for the use of the "generated regressor," $\boldsymbol{\alpha}'\mathbf{x}_t$ in (1) because estimates of the cointegrating parameters converge to their true values at rate T, rather than at the usual rate $\sqrt{T}$ (Stock (1987)) .

The results of estimating a first-order specification of (1) are presented in Lettau and Ludvigson (2013). An important result is that, although consumption and labor income are somewhat predictable by lagged consumption and wealth growth, they are not predictable by the cointegrating residual $\widehat{\boldsymbol{\alpha}}'\mathbf{x}_{t-1}$. Estimates of γ_c and γ_y are economically small and insignificantly different from zero. By contrast, the cointegrating error cay_t is an economically large and statistically significant determinant of next quarter's wealth growth: γ_a is estimated to be 0.20, with a t-statistic equal to 2.3.[8] Thus, only wealth exhibits error-correction behavior. Wealth is mean reverting and adapts over long-horizons to match the smoothness in consumption and labor income.

Third, the individual series involved in the cointegrating relation can be represented as a reduced-form multivariate Wold representation:

$$\Delta\mathbf{x}_t = \boldsymbol{\delta} + \boldsymbol{\Omega}(L)\mathbf{u}_t, \tag{2}$$

where $\mathbf{u}_t$ is an $n \times 1$ vector of innovations, and where $\boldsymbol{\Omega}(L) \equiv \mathbf{I} + \boldsymbol{\Omega}_1 L + \boldsymbol{\Omega}_2 L + \boldsymbol{\Omega}_3 L + \cdots$. The parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}$, both of rank r, satisfy $\boldsymbol{\alpha}'\boldsymbol{\Omega}(1) = 0$ and $\boldsymbol{\Omega}(1)\boldsymbol{\gamma} = 0$ (Engle and Granger, 1987). The "reduced form" disturbances $\mathbf{u}_t$ are not necessarily mutually uncorrelated. To identify shocks that are mutually uncorrelated, we employ a recursive orthogonalization. Specifically, let $\mathbf{H}$ be a lower triangular matrix that accomplishes the Cholesky decomposition of $\mathrm{Cov}(\mathbf{u}_t)$, and define a set of orthogonal structural disturbances $\mathbf{e}$ such that

$$\mathbf{e} \equiv \mathbf{H}^{-1}\mathbf{u}_t.$$

Also define

$$\mathbf{C}(L) \equiv \boldsymbol{\Omega}(L)\mathbf{H}.$$

Then we may re-write the decomposition of $\Delta\mathbf{x}_t = (\Delta c_t, \Delta y_t, \Delta a_t)'$ as

$$\Delta\mathbf{x}_t = \boldsymbol{\delta} + \mathbf{C}(L)\mathbf{e}_t, \tag{3}$$

which now yields a vector of mutually uncorrelated innovations $\mathbf{e}_t$ in consumption, labor income, and asset wealth. We refer to the mutually orthogonal $\mathbf{e}_t$ shocks as "structural" disturbances, but the specific economic interpretation of these shocks must wait for the model presented below. Denote the individual consumption, labor income and wealth disturbances as $e_{c,t}$, $e_{y,t}$, and $e_{a,t}$. Note that these shocks are i.i.d.

[8] We also find that the *four*-quarter lagged value of the cointegrating error strongly predicts asset growth. This shows that the forecasting power of the cointegrating residual for future asset growth cannot be attributable to interpolation procedures used to convert annual survey data to a quarterly housing service flow estimate, part of the services component of c_t.

The ordering of variables will matter for this orthogonalization. We choose a particular ordering, and thus a particular orthogonalization that can be related to the primitive shocks in the model described next. Specifically, we restrict Δc to be first, Δy to be second, and Δa to be last in $\Delta \mathbf{x}_t$, as written above. Roughly speaking, these disturbances amount to (i) consumption shocks $e_{c,t}$: unforecastable movements in Δc that may contemporaneously affect Δy and Δa (ii) labor income shocks $e_{y,t}$: unforecastable movements in Δy and Δa holding fixed Δc contemporaneously, (iii) wealth shocks $e_{a,t}$: unforecastable movements in Δa holding fixed both Δc and Δy contemporaneously.[9]

3.3 Relating Structural Disturbances to Stock Market Fluctuations

The econometric procedure just described is applied to the system of Δc_t, Δy_t, and Δa_t. Ultimately our goal is to study the role each shock has played in the dynamic behavior of stock market wealth over time. To relate stock market wealth to the structural disturbances $e_{c,t}$, $e_{y,t}$, and $e_{a,t}$, we estimate empirical relationships taking the form

$$\Delta s_t = \kappa_0 + \kappa_c (L) e_{c,t} + \kappa_y (L) e_{y,t} + \kappa_a (L) e_{a,t} + \eta_t, \tag{4}$$

where s_t represents the log level of the stock market wealth, $\kappa_c (L)$, $\kappa_y (L)$, and $\kappa_a (L)$ are polynomial lag operators, and η_t is a residual that represents the component of stock wealth that is unexplained by the structural empirical disturbances $e_{c,t}$, $e_{y,t}$, and $e_{a,t}$. For comparison, we compute the same type of relationship for the CRSP value-weighted stock price index, replacing stock market wealth s_t with p_t on the left-hand-side:

$$\Delta p_t = \lambda_0 + \lambda_c (L) e_{c,t} + \lambda_y (L) e_{y,t} + \lambda_a (L) e_{a,t} + \xi_t. \tag{5}$$

Since $e_{c,t}$, $e_{y,t}$ and $e_{a,t}$ are mutually uncorrelated and i.i.d., we estimate these equations separately by OLS with $L = 16$ quarters.

3.4 Decomposition of Levels

We will want to study the role that each shock has played in the evolution of the *levels* of stock market wealth over time. To do so, we decompose the log levels into components

[9] Lettau and Ludvigson (2013) showed how the innovations we study here can be given an additional interpretation, namely as disturbances distinguished by whether their effects have permanent or transitory effects on the variables in the system. An additional rotation of innovations was required for this interpretation, but the shocks obtained there are perfectly correlated with the $\mathbf{e}_t$ obtained here. What this shows is that the permanent and transitory shocks in the $\Delta \mathbf{x}_t = (\Delta c_t, \Delta y_t, \Delta a_t)'$ system are associated with specific orthogonal movements in consumption, labor income and wealth that readily be related to the latent model primitive shocks, as below.

driven by each structural disturbance. Let us rewrite the decomposition of growth rates for stock market wealth (4) as

$$\begin{aligned} \Delta s_t &= \kappa_0 + \kappa_c(L)\, e_{c,t} + \kappa_y(L)\, e_{y,t} + \kappa_a(L)\, e_{a,t} + \eta_t \\ &\equiv \kappa_0 + \Delta s_t^c + \Delta s_t^y + \Delta s_t^a + \eta_t, \end{aligned} \tag{6}$$

where $\Delta s_t^c \equiv \kappa_c(L)\, e_{c,t}$, and analogously for the other terms. The effect on the log levels of stock wealth of each disturbance is obtained by summing up the effects on the log differences, so that the log level of stock wealth may be decomposed into the following components:

$$\begin{aligned} s_t &= s_0 + \kappa_0 t + \sum_{k=1}^{t} \Delta s_k \\ &= s_0 + \kappa_0 t + \sum_{k=1}^{t} \Delta s_k^c + \sum_{k=1}^{t} \Delta s_k^y + \sum_{k=1}^{t} \Delta s_k^a + \sum_{k=1}^{t} \eta_k \\ &\equiv s_0 + \kappa_0 t + s_t^c + s_t^y + s_t^a + \sum_{k=1}^{t} \eta_k, \end{aligned} \tag{7}$$

where s_0 is the initial level of stock market wealth, s_t^c, s_t^y, and s_t^a, are the components of the level attributable to the (cumulation of) the consumption shock, the labor income shock, and the wealth shock, respectively. The final sum $\sum_{k=1}^{t} \eta_k$ is the component of s_t attributable to the unexplained residual. Note that $\gamma_0 t$ is the deterministic trend in stock market wealth, which in the model below is attributable to steady state technological progress. For the level decomposition reported below, we remove this trend and normalize the initial observation s_0 to zero in the quarter before the start of our sample. Expressions analogous to (6) and (7) are also computed for the log stock price index p_t.

This completes our description of the empirical shocks. The next section presents the theoretical model.

4 The Model

4.1 Production and Technology

Suppose that aggregate output $\mathcal{Y}_t$ is given by a constant returns to scale process:

$$\mathcal{Y}_t = \mathcal{A}_t N_t^{\alpha} K_t^{1-\alpha}, \tag{8}$$

where $\mathcal{A}_t$ is a factor neutral total factor *productivity* (TFP) shock, and N_t and K_t are inputs of labor and capital, respectively. We further assume that labor supply is fixed and that there is no capital accumulation, so that both N_t and K_t are constant over time and normalized to unity. Thus $\mathcal{Y}_t = \mathcal{A}_t$ is driven entirely by technological change.

The economy is populated by two types of representative households, each of whom consume an endowment stream. The first type, "shareholders," own a claim to shares of the dividend income stream (equity) generated from aggregate output $\mathcal{Y}_t$. There is no saving and no new shares are issued. Shareholders consume the dividend stream. The second type, "workers," own no assets, inelastically supply labor to produce $\mathcal{Y}_t$, and consume their labor income every period. Dividends, D_t, are equal to output minus a wage bill

$$D_t = \mathcal{Y}_t - W_t N_t, \tag{9}$$

where W_t is the wage rate paid to workers. With labor supply fixed at $N_t = 1$, log labor income, which equals $\ln(W_t N_t) = \ln(W_t)$, is alternatively denoted y_t, to be consistent with the notation above. The total number of shares is normalized to unity.

The wage rate W_t is given by marginal product of labor, multiplied by a time-varying function $f(Z_t)$:

$$\begin{aligned} W_t &= \left[\alpha \mathcal{A}_t N_t^{\alpha-1} K_t^{\alpha-1}\right] f(Z_t) \\ &= \alpha \mathcal{A}_t f(Z_t). \end{aligned} \tag{10}$$

The random variable Z_t over which the function is defined is referred to as a *factors share* shock. We specify $f(Z_t)$ to be a logistic function

$$f(Z_t) = \frac{1}{1+\exp(-Z_t)} + \psi,$$

where ψ is a constant parameter. The calibration we choose insures that the real wage is equal to its competitive value on average, but can be shifted away from this value by a multiplicative scale factor $f(Z_t)$ with $f(Z_t) = 1$ in the non-stochastic steady state. The logistic function insures that the level labor income is never negative and bounded above and below.

Although not modeled explicitly as such, $f(Z_t)$ could be interpreted as the time-varying bargaining parameter resulting from some underlying wage-bargaining problem that creates deviations from competitive equilibrium. Possible sources for such a shift could include changes in reliance on offshoring, outsourcing, part-time or temporary workers, or unionization. At a more basic level, it is a reduced-form way of capturing shifts in the allocation of rewards between shareholders and workers holding fixed the size of those rewards and could occur for a number of reasons (see the conclusion for more discussion). For example, an alternative interpretation of the model is that labor markets are competitive and $f(Z_t)$ represents exogenous technological change that alters how labor-intensive production is without changing the total amount produced. This could be modeled by specifying

$$\mathcal{Y}_t = \mathcal{A}_t N_t^{\alpha f(Z_t)} K_t^{1-\alpha f(Z_t)},$$

and $N = K = 1$. Higher values of $f(Z_t)$ imply a more labor intensive process, and lower values indicating more labor-saving but in this specification fluctuations in $f(Z_t)$ have no influence of aggregate output, which remains $\mathcal{Y}_t = \mathcal{A}_t$. As above, $f(Z)$ only affects how the rewards of production are divided, with $W_t = \alpha\mathcal{A}_t f(Z_t)$ and $D_t = \mathcal{Y}_t - W_t$. This specification has equilibrium allocations that are identical to the previous one.

With this specification for wages, dividends are given by

$$\begin{aligned} D_t &= \mathcal{Y}_t - \alpha\mathcal{A}_t f(Z_t) \\ &= \mathcal{A}_t (1 - \alpha f(Z_t)), \end{aligned} \tag{11}$$

and log dividends, d_t take the form

$$d_t = \mathfrak{a}_t + \ln(1 - \alpha f(Z_t)),$$

where $\mathfrak{a}_t \equiv \ln \mathcal{A}_t$. Notice that log dividends are a non-linear function of the factors share shock.

Aggregate consumption, C_t, is the sum of shareholder consumption (total dividends) and worker consumption (total labor income):

$$C_t = D_t + W_t = \mathcal{Y}_t - W_t + W_t = \mathcal{A}_t.$$

Aggregate (shareholder plus worker) consumption is determined solely by the TFP shock.

The log difference in the TFP shock, $\Delta\mathfrak{a}_t$, is assumed to follow a first-order autoregressive (AR(1)) stochastic process given by

$$\Delta\mathfrak{a}_t - \mu_{\mathfrak{a}} = \phi_{\mathfrak{a}} (\Delta\mathfrak{a}_{t-1} - \mu_a) + \sigma_{\mathfrak{a}}\varepsilon_{\mathfrak{a},t}, \quad \varepsilon_{\mathfrak{a},t} \sim i.i.d.\,(0,1)\,. \tag{12}$$

The factors share shock Z_t is assumed to follow a mean-zero AR(1) processes:

$$Z_t = \phi_z Z_{t-1} + \sigma_z\varepsilon_{z,t}, \quad \varepsilon_{z,t} \sim i.i.d.\,(0,1)\,. \tag{13}$$

Note that in a non-stochastic steady state, Z_t is identically zero, $f(Z_t) = 1$, and dividends are proportional to productivity: $D_t = \mathcal{A}_t (1 - \alpha)$. The above specification implies that the economy grows non-stochastically in steady state at the gross rate of $\mathcal{A}_t$, given by $1 + \mu_{\mathfrak{a}}$, the deterministic rate of technological progress.

4.2 Shareholder Preferences

Worker preferences play no role in asset pricing since they hold no assets. Shareholder preferences determine the pricing of equity. We assume that the economy is populated by a large number of identical shareholders. This leads to a representative *shareholder* model. This should be distinguished from the more common approach of modeling a representative *household* in which aggregate consumption is the source of systematic risk. In this model, there is no single representative household but instead two different types of households, where the distribution of resources between them play a role in asset pricing. The representative shareholder in this model is akin to a large institutional investor or wealthy individual who earns income only from investments. This is empirically relevant because the distribution of stock market wealth across the households is heavily skewed to the right, so that the average dollar invested in stock market wealth is held by such an investor. For this representative shareholder, dividends are the appropriate source of systematic risk.

Let C_{it}^s denote the consumption of an individual stockholder indexed by i at time t. Let β_t be a time-varying subjective discount factor, discussed below. Identical shareholders maximize the function

$$U = E \sum_{t=0}^{\infty} \prod_{k=0}^{t} \beta_k u\left(C_{it}^s\right) \tag{14}$$

with

$$u\left(C_{it}^s\right) = \frac{\left(C_{it}^s\right)^{1-x_{t-1}}}{1-x_{t-1}}, \tag{15}$$

and where $\beta_0 = 1$.

An important aspect of these preferences is that the parameter x_t is not constant but instead varies stochastically over time. It is interpreted as a risk aversion shifter, as discussed further below. It is important that this shifter have low (or zero) correlation with consumption and labor income fluctuations, in order to match evidence presented below that movements in risk premia are largely divorced from these and other traditional economic fundamentals. Thus, habit models (e.g., Campbell and Cochrane (1999)) where risk aversion varies endogenously with consumption shocks will not work. The model specified here also avoids the Brunnermeier and Nagel (2008) critique that portfolio proportions invested in risky assets appear to be cross-sectionally unrelated to liquid wealth, as they should be in many models with habits or consumption commitments.

Shareholder preferences are also subject to an externality in the subjective discount factor β_t, which is assumed to vary over time in a manner dependent on aggregate shareholder consumption (which in equilibrium equal dividends) as follows:

$$\beta_t \equiv \frac{\exp\left(-r_f\right)}{E_t\left[\frac{D_{t+1}^{-x_t}}{D_t^{-x_{t-1}}}\right]}, \tag{16}$$

where r_f is a parameter. We discuss this parametrization for β_t below. In equilibrium, identical individuals choose the same level of consumption, equal to per capita aggregate dividends D_t. We therefore drop the i subscript and simply denote the consumption of a representative shareholder $C_t^s = D_t$ from now on. Notice that aggregate shareholder consumption, given by D_t, is taken as given by individual shareholders and is therefore not internalized in the individual optimization problem.

With these preferences for shareholders, the intertemporal marginal rate of substitution in stockholder consumption is the stochastic discount factor (SDF) given by:

$$M_{t+1} = \frac{\exp\left(-r_f\right)\left(\frac{D_{t+1}}{D_t}\right)^{-x_t}}{E_t\left[\left(\frac{D_{t+1}}{D_t}\right)^{-x_t}\right]}. \tag{17}$$

This can be written

$$M_{t+1} = \beta_t \frac{\left(D_{t+1}\right)^{-x_t}}{\left(D_t\right)^{x_{t-1}}} = \frac{\exp\left(-r_f\right)\frac{\left(D_{t+1}\right)^{-x_t}}{\left(D_t\right)^{-x_{t-1}}}}{E_t\left[\frac{\left(D_{t+1}\right)^{-x_t}}{\left(D_t\right)^{-x_{t-1}}}\right]}. \tag{18}$$

Multiplying the numerator and denominator of (18) by $\left(D_t\right)^{x_t - x_{t-1}}$ gives

$$M_{t+1} = \exp\left[-r_f - \ln E_t \exp\left(-x_t \Delta d_{t+1}\right) - x_t \Delta d_{t+1}\right]. \tag{19}$$

We specify the stochastic risk aversion variable x_t so that it is always non-negative and bounded from above. Specifically, let x_t be specified as a logistic function of a stochastic variable $\widetilde{x}$ that itself can take unbounded values:

$$x_t = \frac{\theta}{1+\exp\left(-\widetilde{x}_t\right)},$$

$$\widetilde{x}_t - \mu_{\widetilde{x}} = \phi_{\widetilde{x}}\left(\widetilde{x}_{t-1} - \mu_{\widetilde{x}}\right) + \sigma_{\widetilde{x}} \varepsilon_{\widetilde{x},t}, \quad \varepsilon_{\widetilde{x},t} \sim i.i.d.\left(0,1\right). \tag{20}$$

In the above, θ is a parameter that controls the maximum value of x_t, which is a nonlinear function of shocks to a stationary first-order autoregressive stochastic process with mean $\mu_{\widetilde{x}}$, innovation variance $\sigma_{\widetilde{x}}^2$, and autoregressive parameter $0 < \phi_{\widetilde{x}} < 1$.

As a benchmark, we specify the three shocks in the model $\varepsilon_{\mathfrak{a},t}$, $\varepsilon_{z,t}$, and $\varepsilon_{\widetilde{x},t}$ are uncorrelated. We do this to demonstrate that the model corresponds closely with the empirical evidence presented here when risk premia are fluctuations are entirely divorced from fluctuations in traditional macroeconomic fundamentals. The orthogonal structure of the three shocks also provides a straightforward interpretation of the three VAR empirical shocks, which are by construction mutually orthogonal. We could allow for correlations among shocks, e.g., a small component of risk aversion could be countercyclical by allowing $\varepsilon_{\mathfrak{a},t}$ and $\varepsilon_{\widetilde{x},t}$ to be negatively correlated, and TFP could be negatively correlated with labor's share. We discuss these possibility below.

4.3 Pricing the Dividend Claim

We use the dividend claim to model the stock market claim. Let P_t denote the ex-dividend price of a claim to the dividend stream measured at the end of time t. The gross return from the end of period t to the end of $t+1$ is defined

$$R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t}. \tag{21}$$

The return on a risk-free asset whose value is known with certainty at time t is given by

$$R_{f,t+1} \equiv \left(E_t\left[M_{t+1}\right]\right)^{-1}.$$

We denote the log return on equity as $\ln\left(R_{t+1}\right) \equiv r_{t+1}$, and the log excess return $\ln\left(R_{t+1}/R_{f,t+1}\right) \equiv r^{ex}_{t+1}$.

From the shareholder's first-order condition for optimal consumption choice, the price-dividend ratio satisfies

$$\begin{aligned} \frac{P_t}{D_t}\left(\mathfrak{s}_t\right) &= E_t\left[M_{t+1}\left(\frac{P_{t+1}}{D_{t+1}}\left(\mathfrak{s}_{t+1}\right)+1\right)\frac{D_{t+1}}{D_t}\right] && (22) \\ &= E_t \exp\left(m_{t+1} + \Delta d_{t+1} + \ln\left(\frac{P_{t+1}}{D_{t+1}}\left(\mathfrak{s}_{t+1}\right)+1\right)\right), && (23) \end{aligned}$$

where $\mathfrak{s}_t$ is a vector of state variables, $\mathfrak{s}_t \equiv \left(\Delta\mathfrak{a}_t, Z_t, x_t\right)'$. Even with all shocks normally distributed, there is no closed-form solution to the functional equation (23), due to two sources of nonlinearities on the right-hand-side of (23). First, dividend growth is non-linear in the Z_t shock. Second, the term $\ln\left(\frac{P_{t+1}}{D_{t+1}}\left(\mathfrak{s}_{t+1}\right)+1\right)$ is itself a non-linear function of the shocks. We therefore solve the $\frac{P_t}{D_t}\left(\mathfrak{s}_t\right)$ function numerically on an $n \times n \times n$ dimensional grid of values for the state variables, replacing the continuous time processes with a discrete Markov approximation following the approach in Rouwenhorst (1995). Further details are given in the Appendix.

4.4 The Risk-free Rate, the Sharpe Ratio and Risk Aversion

It is well known that empirical measures of the risk-free interest rate are extremely stable, with a standard deviation one-tenth (or less) the size of that for an aggregate stock market index return or its premium over a risk-free interest rate. The specification for the subjective time discount factor in (16) is essential for obtaining a stable risk-free rate alone with a volatile equity premium. If instead the subjective time discount factor were itself a constant (as is common), shocks to x_t and dividend growth would generate counterfactual volatility in the risk-free rate. For simplicity, we calibrate the subjective time-discount factor β_t so that the risk-free rate is constant. Specifically, the externality involving the term $1/E_t\left[\frac{D_{t+1}^{-x_t}}{D_t^{-x_{t-1}}}\right]$ is a compensating factor in the stochastic discount factor that renders the risk-free rate constant even in the face of large shocks to the price of risk. While this compensating factor may seem unusual, it is in fact a generalization of a familiar compensating Jensen's term that appears in many lognormal models of the stochastic discount factor (e.g., Campbell and Cochrane (1999)) and (Lettau and Wachter (2007)).[10] The generalization given here holds for any distribution of shocks, not just lognormal. This specification can be understood intuitively by noting that shocks to x_t generate changes in risk aversion and a desire for precautionary savings that would (absent offsetting movements in the subjective rate of time-preference for saving) produce counterfactually large movements in the risk-free rate. Specifying the variation in β_t as a function of aggregate shareholder consumption, rather than individual shareholder consumption, greatly simplifies the implications for risk aversion in the model, as discussed below.

From (19) we have

$$\begin{aligned} E_t M_{t+1} &= E_t\left[\exp\left(-r_f\right)\exp\left(-\ln E_t \exp\left(-x_t \Delta d_{t+1}\right)\right)\exp\left(-x_t \Delta \ln d_{t+1}\right)\right] \\ &= \frac{\exp\left(-r_f\right) E_t \exp\left(-x_t \Delta \ln d_{t+1}\right)}{\exp\left(\ln E_t \exp\left(-x_t \Delta d_{t+1}\right)\right)} = \exp\left(-r_f\right). \end{aligned}$$

[10] Specifically, in the lognormal model of Lettau and Wachter (2007), the stochastic discount factor takes the form

$$M_{t+1} = \exp\left\{-r^f - \frac{1}{2}x_t^2 - x_t \epsilon_{d,t+1}\right\}$$

and the compensating factor is the Jensen's term $\frac{1}{2}x_t^2$. The pricing kernel in this paper is, however, quite different from those of Campbell and Cochrane (1999) and Lettau and Wachter (2007), since it exhibits sharp departures from lognormality generated by the non-linearity of risk aversion and factors share functions.

Denoting the gross risk-free rate $R_f \equiv \exp(r_f)$, it follows that

$$\begin{aligned} E_t [M_{t+1}]^{-1} &= \exp(r_f) = R_f => \\ -\ln E_t [M_{t+1}] &= r_f. \end{aligned}$$

Thus the log risk-free rate is equal to the constant parameter r_f.

Define the price of risk as $\sqrt{\text{Var}_t (M_{t+1})}/E_t (M_{t+1})$. This variable is equal to the largest possible Sharpe ratio $E_t \left(R^{ex}_{t+1}\right) / \sqrt{\text{Var}_t \left(R^{ex}_{t+1}\right)}$ on an asset with excess return R^{ex}_{t+1} (Hansen and Jagannathan (1991)). Given (19), the maximal Sharpe ratio can be shown to equal

$$\frac{\sqrt{\text{Var}_t [M_{t+1}]}}{E_t [M_{t+1}]} = \frac{\sqrt{\text{Var}_t [\exp(-x_t \Delta d_{t+1})]}}{E_t \exp(-x_t \Delta d_{t+1})}. \tag{24}$$

Because of the nonlinearities in this model, there is no simple analytical expression for (24) as a function of the underlying state variables $\mathfrak{s}_t$. Nevertheless, (24) shows that the price of risk will vary with innovations in risk aversion x_t that multiply shocks to dividend growth.

The computation of risk aversion in the full stochastic model is quite complicated numerically. However, it is straightforward to calculate risk aversion along a non-stochastic balanced growth path. Define the coefficient of relative risk aversion RRA_t as

$$RRA_t = \frac{-A_t E_t V''(A_{t+1})}{E_t V'(A_{t+1})}, \tag{25}$$

where $V(A_{t+1})$ is the representative shareholder's value function associated with optimal consumption choice, and A_t is this shareholder's asset wealth. Following the derivation in Swanson (2012), we show in the Appendix that risk aversion along the non-stochastic balanced growth path, RRA, is equal to

$$RRA = \frac{-C^{s\cdot}_t u''\left(C^{s\cdot}_{t+1}\right)}{u'\left(C^{s}_{t+1}\right)} = E(x_t) / (1 + \mu_a),$$

where $E(x_t)$ is the unconditional mean of x_t and $1 + \mu_a$ is the non-stochastic gross growth rate of the economy driven by steady state technological progress $\mathcal{A}_t$.[11] This shows that risk aversion in the model is a function of the process x_t, with its steady state value given by the mean of this process. We refer to x_t as a *risk aversion* shock.

[11] The $1 + \mu_a$ term in the expression for risk aversion is a nuisance term that arises because of balanced growth and the beginning of period timing assumption for assets.

4.5 Calibration

This section discusses the calibration of the model. We begin with a discussion of the calibration of the primitive shocks.

We choose parameters of the model so as to match key empirical moments of asset pricing data, inclusive of new evidence we present below. Table 1 presents a list of all parameters and their calibrated values. The parameter α, the exponent on labor in the production function, is set to 0.667, a value that is standard in real business cycle modeling. The constant value for the quarterly log risk-free rate is set to match the mean of the quarterly log 3-month Treasury bill rate. We set $\phi_{\mathfrak{a}} = 0$ so that the log level of TFP, $\mathfrak{a}_t$, follows a unit root stochastic process with drift. The mean and standard deviation of productivity $\Delta\mathfrak{a}_t$ is set to roughly match the mean and standard deviation of the quarterly log difference of consumption in the data. The factor share shock Z_t is set to be very persistent, with $\phi_z = 0.995$, to match the extreme persistence of the empirical labor income shock found in the data. The parameter ψ in $f(Z_t)$ is set to $\psi = 0.5$ so $f(Z_t)$ lies in the interval $[0.5, 1.5]$ and equals unity in the non-stochastic steady state. The symmetry of the (normal) distribution for Z_t insures that the mean of the factor's share shifter is also unity $E(f(Z_t)) = 1$. This calibration, along with the calibration of the volatility of Z_t, allow the model to roughly match the standard deviation of dividend growth, which is over ten times that of aggregate consumption growth. Matching evidence for a volatile dividend growth process also has important implications for the model's ability to match the frequency decomposition of stock price changes.

Finally, the parameters of the risk aversion process $\sigma_{\widetilde{x}}$, $\phi_{\widetilde{x}}$, and θ, are set so as to come as close as possible to simultaneously matching (i) the mean equity premium, (ii) the forecastability of the equity premium and the average level of the price-dividend ratio. Matching both features of the data simultaneously requires a risk-aversion process that takes values close to zero most of the time but is highly skewed to the right, exhibiting infrequent spikes upward.

To understand why, observe that shareholders who consume out of dividends are exposed to much greater systematic risk than would be the case for a representative household who consumes the stable aggregate consumption stream. With dividend growth this volatile, shareholder risk aversion must be close to zero in most states or the model generates a counterfactually *high* equity premium. But matching evidence for a *time-varying* equity premium requires risk aversion to fluctuate. With risk aversion bounded below at zero, fluctuations must be skewed upward. If the upper bound on risk aversion is too restrictive, however, not only does the model generate too little variation in risk premia, there are also important regions of the state space that imply (almost) unbounded values for the price-dividend ratio, resulting in equilibria that would drastically overshoot the mean price-dividend ratio. Taken together, these factors necessitate extreme non-linearities in risk aversion: a very low value in normal times, but infrequent spikes upward to extraordinary values. This is an economic result of the model. In a world where investors are willing to tolerate to a high degree of systematic risk most of the time, and in which the process for $f(Z)$ generates states in which investors can expect to persistently transfer rewards from labor, stocks would simply be too attractive to explain the observed mean price-dividend ratio unless these factors are offset by the expectation of rare states in which the market's risk tolerance implodes, leading to a "flight to safety" and a market crash.[12]

Figure 2 displays the density of our risk aversion process, given the parameters chosen. Almost all of the mass is close to zero, with the median and mode equal to unity to close approximation. The mean of 30 is reached far more infrequently and there is a small amount of mass near the maximum value for risk-aversion, set to 450.[13] This could be compared to the risk aversion variation in the Campbell and Cochrane (1999) habit model, which in their benchmark calibration has a minimum value of 60 reaches toward infinity in states where consumption is very close the habit level. The Campbell and Cochrane (1999) risk aversion process is less non-linear than the process here, with much higher median risk aversion.

The risk aversion process in the model should be thought of as an externality–the market's willingness to bear risk. One interpretation of such independent variation is that it is driven by intangible information. The extreme non-linearities of this process are akin to a "threshold" model in which risk aversion fluctuates modestly around relatively low values most of the time but, once a certain threshold is crossed, spikes up to extreme values. Because the risk aversion shifter is persistent, small fluctuations, even in normal times, change the conditional expectation that the threshold will be crossed. These changing expectations generate fluctuations in the price-dividend ratio that are far less non-linear in the state are than the risk aversion dynamics itself (though fluctuations in $p_t - d_t$ are naturally largest in crisis times). These fluctuations in the price-dividend ratio are driven by time-variation in the equity risk premium.

It is far more difficult to match these facts if risk aversion does not have a large inde-

[12] Investors can expect to persistently transfer rewards from labor in high Z states (recall Z is mean reverting but highly persistent). These states also endogenously create low *volatility* of dividend growth because of the boundedness of $f(Z)$. Both factors push up the price-dividend ratio.

[13] This highly skewed distribution for risk aversion is not an artifact of the logistic function chosen for $f(Z)$. A truncated Normal distribution of Z that generates similar equilibrium allocations also requires low risk aversion most of the time with infrequent extreme values.

pendent component statistically. For example, introducing a negative correlation between $\varepsilon_{\mathfrak{a},t}$ and $\varepsilon_{\widetilde{x},t}$ reduces correlation between the stochastic discount factor and dividend growth, causing the model to substantially overshoot the unconditional equity premium over a range of parameter values. In principle this could be addressed lowering the upper bound for aversion, which reduces its volatility. But, as explained above, these changes make difficult if not impossible to match evidence for a sufficiently *time-varying* equity premium along with a plausible mean price-dividend ratio.

5 Results

5.1 Model Summary Statistics

Table 2 presents summary asset pricing statistics of the model and compares them to those in post-war data. The average annual excess return on equity in the model closely matches, which is 4.9% in the data and 5.1% in the model, and the model standard deviation closely matches the annualized standard deviation of this return, which is 16.5% in the data and 16.9% in the model. The model also matches the mean and standard deviation of the log price-dividend ratio. By construction, the model exactly matches the mean risk-free rate. The model also does a good job of matching the evidence the volatility of dividend growth is much higher than that of consumption growth; the annual standard deviation of log dividend growth is 12.0% in the data and 12.3% in the model. Because dividends are subject to the factors share shock, they are more volatile than aggregate consumption.

We investigate the model's implications for the dynamic relationship between the log price dividend ratio, $p_t - d_t$, and future long horizon excess equity returns, $\sum_{j=0}^{h} r_{t+j+1}^{ex}$, consumption growth, $\sum_{j=0}^{h} \Delta c_{t+j+1}$, and dividend growth, $\sum_{j=0}^{h} \Delta d_{t+j+1}$. Table 3 reports regression results of one through five year log excess equity returns on the lagged price-dividend ratio, as well as one through five year log differences in consumption growth or dividend growth on the lagged price-dividend ratio. The columns on the right present results from our model; the columns on the left present the corresponding results from historical data. The model results are averages across 1000 simulations of length 238 quarters, the same length as our quarterly historical data.

Table 3 shows that the log price-dividend ratio predicts future excess returns with statistically significant negative coefficients in the model, while the coefficients for consumption growth are statistically indistinguishable from zero at all horizons and those for dividend growth are indistinguishable from zero for all but the 5 year horizon. Persistent but stationary variation in risk aversion in the model generates forecastable variation in equity premia but there is no forecastability of consumption or dividend growth, with the exception of the 20 quarter horizon for dividend growth where the model implies a modest amount of predictability. For the most part, these implications are consistent with the data: the price-dividend ratio exhibits long-horizon predictive power for equity premia, but not consumption growth or dividend growth. The adjusted R-squared statistics for forecasting excess returns are comparable between model and data: they range from 0.139 to 0.244 in the model for one to five year horizons, and from 0.068 to 0.188 in the data. Thus the model is consistent with the well known "excess volatility" property of stock market returns, namely that fluctuations in stock market valuation ratios are informative about future equity risk premia, but not about future fundamentals on the stock market (i.e., dividend or earnings growth, LeRoy and Porter (1981), Shiller (1981)), or future consumption growth (Lettau and Ludvigson (2001); Lettau and Ludvigson (2004)).

While the model is broadly consistent with these benchmark asset pricing moments, it is limited in matching the data in some ways. One is, although the model correctly implies that labor income growth is more volatile than consumption growth, the standard deviation is too high: 5% annually compared to 2% in the data. In the simplified model environment here, it not possible to simultaneously match evidence for both a volatile dividend growth process and a stable labor income growth process, since the two are tied together by the volatility of the factor shares shock. Future work could explore extensions of the model that exhibit a wage smoothing or stickiness in the presence of a factors share shock. Another is that, at very long horizons, the model implies some predictability of Δd_{t+h} by $p_t - d_t$. Such predictability is not present in estimates using historical data but is an unavoidable feature of the model here because of the requirement that Z be bounded in order to prevent the level of dividends and labor income from going negative. If Z is very highly persistent, however, the predictability in the model can be pushed out to very long horizons, making it more difficult to detect in finite samples.

5.2 Comparing the Model Primitive Shocks to the Model VAR shocks

We next investigate the connection in the model between the observable VAR shocks and the latent primitive shocks. To do so, we take model simulated data, compute the VAR shocks and compare them to the primitive shocks.

Figure 2 shows two sets of cumulative dynamic responses of Δc_t, Δa_t, and Δy_t. The left column shows the cumulative responses of these variables to the three primitive shocks in the model. The top panel shows the responses of these variables to the TFP shock, the middle panel shows the responses to the factors share shock, and the bottom panel shows the responses to the risk aversion shock. These responses are calculated by applying, for each shock one at a time, a one standard deviation change in the direction that increases Δa_t at time $t = 0$,and a zero value at all subsequent periods, and then simulating forward using the solved policy functions with the other two shocks set to zero in every period. This implies we plot the responses to a one standard deviation increase in $\varepsilon_{\mathfrak{a},t}$, and a one standard deviation decrease in $\varepsilon_{z,t}$ and $\varepsilon_{\widetilde{x},t}$. The right column uses model simulated data to calculate the mutually orthogonal VAR innovations $\mathbf{e}_t$ (3) and plots dynamic responses to one standard deviation change in each $\mathbf{e}_t$ shock, again in the direction that increases Δa_t. The top right panel shows the responses of consumption, labor income and wealth to a consumption shock, $e_{c,t}$, the middle right panel shows the responses to a labor income shock, $e_{y,t}$, and the bottom right panel shows the responses to a wealth shock, $e_{a,t}$.[14]

The key result shown in Figure 2 is that the dynamic responses of aggregate consumption, labor earnings, and asset wealth to the VAR innovations in the right column are almost identical to the theoretical responses of the same variables to the productivity, factors share, and risk aversion shocks, respectively, in the left column. The small deviations that do exist from perfect correlation for some responses are attributable to small numerical errors and to nonlinearities in the model not captured by the linear VAR, as discussed above. But these deviations are small. The responses of Δc_t, Δy_t and Δa_t to the consumption shock, $e_{c,t}$, are all perfectly correlated with the responses of these variables to the TFP shock $\varepsilon_{\mathfrak{a},t}$; the response of Δc_t to the labor income shock $e_{y,t}$ is perfectly correlated with the response of Δc_t to the factors share shock $\varepsilon_{z,t}$, and the responses of Δc_t, Δy_t, and Δa_t to the wealth shock $e_{a,t}$ are all perfectly correlated with the responses of Δc_t, Δy_t, and Δa_t to the risk aversion shock $\varepsilon_{\widetilde{x},t}$. We verify, from a long simulation of the model, that the correlation between the consumption shock $e_{c,t}$ and the productivity shock $\varepsilon_{\mathfrak{a},t}$ is unity, the correlation between labor income shock $e_{y,t}$ and first difference of the factors share shifter $\Delta \ln f(Z_t)$ is unity, and the correlation between the wealth shock $e_{a,t}$ and the innovation in Δa_t attributable only to risk aversion shocks $\varepsilon_{\widetilde{x},t}$ is 0.97.[15]

In presenting the above, we do not claim that the mutually uncorrelated VAR shocks

[14] Since we compare the model-based VAR responses to the in-population theoretical responses to primitive shocks, we rid the VAR responses of small sample estimation biases by computing them from a single simulation of the model with very long length (238,000 quarters). The size of the primitive shocks are normalized so that they are the same as the empirical shocks in the right column.

[15] The innovation in Δa_t attributable to risk aversion shocks is computed as $\Delta a_t - E\left[\Delta a_t | \mathfrak{s}_{t-1}, Z_t, \Delta \mathfrak{a}_t\right]$.

$(e_{c,t}, e_{y,t}, e_{a,t})$ exactly equal the primitive shocks $(\varepsilon_{\mathfrak{a},t}, \varepsilon_{z,t}, \varepsilon_{\widetilde{x},t})$, respectively. Exact equality is impossible because the endogenous variables in the model are nonlinear functions of the primitive shocks, while the VAR imposes a linear relation between these variables and the VAR shocks. Moreover the comparable innovations are in different units so a rescaling is necessary. What the above does show is that, if the model generated the data, the VAR disturbances would, to a very close approximation, serve as the observable empirical counterparts to the innovations originated from the latent primitive shocks.

5.3 The Role of the Empirical Shocks in Quarterly Stock Market Fluctuations

With this theoretical interpretation of the VAR disturbances in mind, we now study the role of the empirical shocks for historical stock market data. In each investigation, we compare the outcomes for stock market wealth using historical data with those using model-simulated data. We begin by studying the dynamic responses to shocks.

5.3.1 Impulse Responses

Figure 3 again shows two sets of cumulative dynamic responses of Δc_t, Δa_t, and Δy_t. The left column uses model simulated data to calculate model-based responses to the mutually orthogonal VAR innovations (3). These are the same responses that are shown in the right panel of Figure 2, except that the responses in Figure 3 are averages across 1000 samples of size 238 quarters, rather than over one very long sample. The right column of Figure 3 shows the cumulative dynamic responses of Δc_t, Δa_t, and Δy_t in historical data to the VAR innovations (3) estimated from historical data. In each column, the top panel shows the responses of consumption, labor income and wealth to an e_c shock, the middle right panel shows the responses to an e_y shock, and the bottom right panel shows the responses to an e_a shock.

Figure 3 shows that the cumulative dynamic responses Δc_t, Δa_t, and Δy_t to the observable VAR shocks are remarkably similar across model and data. These in turn are reasonably comparable to the responses of Δc_t, Δa_t, and Δy_t to the model primitive shocks in Figure 2. In each case, a positive innovation in the consumption e_c shock leads to an immediate increase in c_t, a_t, and y_t, both in the data and the model. From Figure 2, it is clear that this shock in the model reveals the effects of the total factor productivity shock. The model responses of c, y, and a, to the consumption shock lie on top of each other because the levels of these variables are all proportional to TFP, so the log responses are the same. Moreover, in the model the TFP shock is the innovation to a random walk, so all three variables move immediately to a new, permanently higher level. In the data, this adjustment is not exactly immediate but it is relatively quick: full adjustment of consumption occurs within 3 quarters or less, very close to what one would expect from a random walk shock. Cochrane (1994) makes the same observation when studying a bivariate cointegrated VAR for consumption and GNP and argues that consumption is sufficiently close to a random walk so as to effectively define the stochastic trend in GNP.

The second rows of Figure 3 displays the dynamic responses of c_t, a_t, and y_t to the labor income shock $e_{y,t}$, which Figure 2 shows reveals the effects of the factors share shock in the model. In the data, the response of consumption to this shock is economically negligible. This is true by construction on impact as a result of our identifying assumption. But it is also true in all subsequent periods, an important result that is not part of our identifying assumption. Instead, this shock affects asset wealth and labor income, driving a_t and y_t in opposite directions. A positive value for this shock raises asset wealth a_t and lowers labor income y_t, both of which are moved to new long-run levels. The effect on labor earnings is large and immediate: labor income jumps to a new lower level within the quarter. Below we present evidence that changes in a_t resulting from this shock are driven by the stock market. The model dynamic responses in the left column show the same basic patterns: the $e_{y,t}$ shock is one that has no affect on consumption at any horizon, but drives a near-permanent wedge between asset values and labor income.[16] Lettau and Ludvigson (2013) show that this factor shares shock has almost no effect on housing wealth or non-stock market financial wealth, so that it is effectively a shock to *shareholder* wealth.

One puzzling aspect of the data impulse responses exhibited by Figure 3 is that asset wealth responds sluggishly to the factors share shock, suggesting that the information revealed in the innovation is incorporated only slowly into stock prices. This pattern is attributable to the behavior of stock market wealth, as shown in Figure 4, discussed below. This could reflect a composition effect in the stock market, if for example an increasing fraction of firms going public over the sample employ labor-saving technologies. It could

[16] The model responses in Figure 2 were obtained in the same way as the data responses, namely from a cointegrated VAR that imposes cointegration among c_t, a_t, and y_t but does not impose a second linearly independent cointegrating relation between y_t and a_t. The model shocks to $y_t - a_t$ are very persistent (inheriting the persistence of the Z_t shock), but ultimately stationary. This degree of persistence implies that we cannot reject a unit root for $y_t - a_t$ in either the model or the data in samples of the size we currently have. Thus the model is consistent with the data in this respect. Given the extreme persistence of $y_t - a_t$, we follow the advice of Campbell and Perron (1991) to *not* impose stationarity for $y_t - a_t$ in the estimated cointegrated VAR for either the model or the data.

also reflect, in part, imperfect observability of factors share shocks by shareholders who own shares in many independently managed firms and have to learn over time about the pervasiveness and persistence across the broader economy of the ultimate sources of such shocks, which could include labor-saving technological changes, shifts toward or away form offshoring and outsourcing, changing reliance on temporary or part-time workers, and the bargaining power of unions. All of these factors could play a role in factors share shocks, but it could take time to fully observe them and grasp their long-run implications for cash-flows. Future research is needed to formally investigate these and other possibilities.[17]

The third row of Figure 3 shows the effects of a positive wealth shock, $e_{a,t}$, which is driven by a decline in risk aversion in the model. In both the data and the model, this shock leads to a sharp increase in asset wealth, but has no impact on consumption and labor earnings at any future horizon. The zero responses of c_t and y_t on impact are the result of our identifying assumptions, but the finding that this shock has no subsequent influence on consumption or labor income at any future horizon is a result. The Appendix describes a bootstrap procedure for computing error bands, and shows that these zero responses are within 90% error bands. By contrast, the effect of this risk aversion shock on a_t is strongly significant over periods from a quarter to many years. Although transitory, the degree of persistence of this shock is quite high. Its influence on a_t has a half-life of over 4 years in historical data.

This result in central to understanding why the risk aversion shock must be modeled as (largely) independent of consumption and labor earnings shocks: the large transitory movements in wealth captured by the $e_{a,t}$ shock (and which we show below drive forecastable variation in equity risk premia) are unrelated to consumption and labor income at all horizons. The response of a_t to this shock in the model is of the same pattern, but less persistent than in the data. As in the data, the model-based responses in the left column to the $e_{a,t}$ shock are driven by the risk aversion shock and have big affects on a_t but negligible affects on c_t or y_t.

Figure 3 showed dynamic responses of asset wealth a_t. Figure 4 shows the cumulative

[17] The left column of Figure 3 shows that, even in the model, there is slight sluggish response of stock market wealth to the factors share shock, although it is far less pronounced than in the data. Interestingly, this is a finite sample effect: the responses in Figure 3 for the model are averages over 1000 simulations of size 238 quarters (the same size as our historical dataset). In many samples of this size, the response appears sluggish, even though the population response displays no sluggishness. This can be seen via a comparison with Figure 2, which, unlike the response in Figure 3, is computed from one very long simulation of the model, rather than from averages over many short ones. The response of a_t to a y shock in Figure 2 displays no sluggishness (middle row, right column).

dynamic responses of stock market wealth Δs_t in historical data to a one-standard deviation innovation in each VAR disturbance, along with 90% error bands computed from the bootstrap procedure described in the Appendix. The responses are constructed using the OLS estimates of (4) for stock wealth or (5) for CRSP stock price Δp_t. It is clear that the responses of stock wealth or stock price to the wealth and labor income shock mimic those of asset wealth to these same shocks, indicating that they are primarily shocks to *shareholder* wealth, not other forms of wealth. This is consistent with the evidence in Lettau and Ludvigson (2013) which finds that other forms of wealth are not closely related to these wealth shocks. To avoid clutter, the responses for the CRSP stock price Δp_t using (5) are not reported, but we confirm that they are very similar to those for stock wealth.

5.3.2 Variance Decompositions

The impulse responses tell us about the dynamic effects of each shock on stock market wealth, but are less informative about the quantitative importance of each shock. Using (4), we characterize the relative quantitative importance of each empirical disturbance, as well as that of the residual, for quarterly stock market wealth using a variance decomposition. Table 4 displays the fraction of the unconditional variance in the log difference of stock market wealth that is attributable to each empirical shock. The variance decompositions in the data are also computed for the log difference in the CRSP value-weighted stock price index using (5). The model-based variance decompositions are computed by estimating (4) on simulated data and reporting average statistics over 1000 simulations of size 238 quarters.

Table 4 shows that the wealth shock e_a explains the largest fraction of quarterly stock wealth growth Δs_t and accounts for 76% of its quarterly volatility (75% of the quarterly variation in Δp_t). The two other shocks account for very small amounts, 4% and 6% for the labor income and consumption shocks, respectively. Note that the residual explains only 13% of quarterly stock market fluctuations, showing that the vast majority of quarterly stock market fluctuations are explained by these three shocks.

The model is broadly comparable with the data along these lines: the vast bulk of quarterly fluctuations in stock wealth in the model are attributable to the wealth/risk aversion shock, with much smaller roles for the consumption/TFP and labor income/factors share shocks. Under this calibration model gets the role of the labor income shock almost exactly right, but somewhat overstates the role of the wealth shock and slightly understates the role of the consumption shock. It is important to bear in mind that, by construction, there is no "residual" in the model version of equation (4), since the productivity, factors share, and risk aversion shocks explain 100% of the variability in stock market wealth in the model. Since there is no residual in the model but there is in the data, the variance decomposition in the model must differ from the data somewhere. It is encouraging that the residual accounts for less than 15% of total variation in stock market wealth in the data, implying that the model implications for the sources of stock market fluctuations are reasonably correct.

An important result in Table 4 is that, both in the model and in the data, aggregate consumption shocks play a very small role in quarterly stock market fluctuations. We show below that this is also true at longer horizons (where, by contrast, the factors share shocks play a large role.) The model above implies that factor neutral technological progress should reward labor and all productive capital, thus it should be a driver of aggregate (shareholder plus worker) consumption. Given labor's greater role (on average) in the production process (steady state labor share is two-thirds), most gains and losses from total factor productivity shocks accrue to workers rather than shareholders, so these shocks are less important. This finding is difficult if not impossible to reconcile with representative agent models where consumption shocks play the key role in asset price fluctuations. As an example, the last row of Table 4 gives the corresponding numbers for the Campbell and Cochrane (1999) habit model, in which 84% of quarterly stock price growth is driven by consumption shocks.[18]

5.4 Short- Versus Long-Term Determinants of Stock Market Wealth

Above we examined the role of each shock for quarterly stock market growth. A key question this paper seeks to address is how the sources of stock market fluctuation might vary by the time frame over which a change in stock market wealth is measured. This section examines the short- versus long-term determinants of stock market wealth in two ways.

5.4.1 Frequency Decomposition

One way to examine how the stock wealth is affected at different time horizons is to decompose the variance of the stock wealth by frequency, using a spectral decomposition. This decomposition tells us what proportion of sample variance in Δs_t is attributable to cycles of different lengths. We estimate the population spectrum for the deterministically detrended log difference in stock wealth $\Delta s_t - \kappa_0$ or stock price $\Delta p_t - \lambda_0$ using (4) or (5). Noting that $\Delta s_t - \kappa_0$ in (6) is a function of three components, Δs_t^c, Δs_t^y, Δs_t^a, plus an i.i.d. residual η_t,

[18]The fraction of variance explained by consumption shocks is less than 100% only because the Campbell Cochrane model is non-linear, while the variance decomposition is computed from a linear VAR.

and using the fact that the spectrum of the sum is the sum of the spectra, we estimate the fraction of the total variance in stock market wealth that is attributable to each component at cycles of different lengths, in quarters. This is computed as the estimated spectrum for the component, e.g., Δs_t^c, divided by the estimated spectrum for $\Delta s_t - \kappa_0$. The Appendix provides additional details. This decomposition is potentially quite informative about relative importance of different sources of variation for the long-run evolution of stock market wealth.

Figure 5 exhibits these decompositions for the model (top panel), and for the data using stock market wealth (middle panel) or stock price (bottom panel). The horizontal axis shows the length of the cycle in quarters. The vertical axis gives the frequency decomposition of variance. Consider the line marked "a" in the middle panel for historical stock market wealth. This line shows that, for short cycles (i.e., periods of a few quarters), the fraction of variance in stock wealth that is attributable to the independent risk aversion/wealth shock is very high, close to 80%. The high frequency, short-horizon variability of the stock market in post-war data is dominated by independent shocks to investors' willingness to bear risk. As these cycles become longer, the fraction of variance in stock wealth explained by this shock declines and asymptotes to roughly 40%. Thus the lower frequency, longer horizon variability of the stock market is less dominated by such shocks.

Next consider the line marked "y" in the middle panel of Figure 5. This line shows that, for short cycles of a few quarters, the fraction of variance in stock wealth that is attributable to the factors share/labor income shock is very low, close to zero. Thus the high frequency, short horizon, variability of the stock market in post-war data is virtually unrelated to these factors share shocks. But because these innovations are nearly permanent, they play an increasingly important role as the time horizon extends. Figure 5 shows that as the cycle become longer, the fraction of variance in stock wealth explained by the factors share shock steadily rises and asymptotes to roughly 40%, equal in importance to the risk aversion shock at long horizons. Thus the lower frequency, long horizon, variability of the stock market is quite significantly effected by the factors share shock, even though these shocks play virtually not role in quarterly or even annual stock market fluctuations.

Finally in this middle panel, consider the line marked "c." This line shows that, no matter what the length of the cycle, the fraction of variance in stock wealth that is attributable to the TFP/consumption shock is very low, close to zero. The line marked "residual" shows the contribution of component of stock market fluctuations that is unexplained by these three mutually orthogonal innovations. This residual explains less than 20% of the variability in the stock market at all frequencies, and asymptotes to around 10% as the horizon extends. There are no important differences in these results if we instead use stock price in place of stock market wealth (bottom panel).

The model captures this frequency decomposition well. The top panel of Figure 5 shows model-based orthogonal wealth shocks that originate from independent shifts in investors' willingness to bear risk explain almost all of the variation in stock market fluctuations over short horizons and less as the horizon extends, consistent with the data. The model-based labor income/factors share shock explains a negligible fraction of variation in the stock market over shorter time horizons, but is equal in importance to the risk aversion shock as the time horizon extends. Indeed these two shocks each explain about 45% of the variation in the stock market over cycles three decades long, a fraction comparable to the 40% found in the data for these shocks. And the model provides an explanation for why the productivity shocks plays such a small role in the stochastic fluctuations of the stock market at all horizons: Given capital's smaller role in the production process, most gains and losses from total factor productivity shocks accrue to workers, so these shocks are less important for stock market wealth than are the other two sources of variation.

5.4.2 The Level of the Stock Market in the Post-War Period

A second way we investigate the short- versus long-run determinants of stock market wealth is to study the role each shock has played in driving stock market wealth at specific points in out sample. This can be accomplished by examining the levels decomposition of stock market wealth (7) over time. Figure 6 plots the levels decomposition for stock wealth (left column) and stock price (right column) over our sample. The top panels of each column shows the sum of all components, which equals the log level of the variable (stock wealth or stock price) after removing the deterministic trend. The panels below the top panel show to the component attributable to the cumulation of the consumption/TFP shocks, the labor/factors share shock, the wealth/risk aversion shock, and the residual. Each component (and the sum) is normalized so that its value in 1952:Q1 (one period prior to the beginning of our sample) is zero.

It is immediately clear from Figure 6 that the TFP component contributes relatively little to the variation in stock market wealth consistently throughout the sample. This component does take a noticeable drop at the end of the sample during and after the recession of 2007-2009, but it is still quite modest compared to the variation in other components. The bottom panel shows that the variation attributable to the unexplained residual is also reasonably small. Instead, the big movers of stock market wealth are the factors share shock and the risk aversion shock. Figure 6 shows that the low frequency movements in the level of stock market wealth are dominated by the cumulative swings in the factor shares component, while shorter-lived peaks and troughs in the stock market have coincided with spikes up or down in the risk aversion component.

The large role of the factors share shock in driving stock market wealth over horizons of many years is particularly striking in this decomposition. Let us consider the last 25 years as an example. The third panel of Figure 6 shows that the cumulative effect of the factors share shock has persistently boosted stock market wealth over this time. From the dynamic responses in Figure 3, it's clear that these very shocks also persistently lowered the level of labor earnings over this period. By contrast, from the mid 1960s to the mid 1980s, the cumulative effect of this shock persistently boosted labor earnings and lowered stock market wealth. The large low frequency swings in the stock market over the sample are exactly mirrored in the low frequency swings of the factors share component. Figure 7 plots the factors share components of both log stock market wealth and log labor income over time. The figure shows that there is a stark inverse relationship over time between labor earnings and the stock market that is the result of the cumulative reallocative outcomes of the factors share shock. In the last 15 years, the factors share component of labor income rose sharply, explaining the decline in the factors share component of stock market wealth over the recent period.

As an example of the quantitative importance of such shocks over long-horizons, we use this levels decomposition to calculate the percentage change since 1980 in the deterministically detrended real value of stock market wealth that is attributable to each shock. This is trivial to compute once we have (7), because the change in the total value of the stock market on the left-hand-side of is the sum of three structural components and a residual on the right-hand-side, all of which are orthogonal to one another. As mentioned, the period since 1980 is an interesting one to study, given the large cumulative effects of the factor shares shock on stock market wealth and labor earnings over this time. The decomposition (7) implies that the cumulative affects of the factors share shock, which persistently boosted stock market wealth at the expense of labor earnings over this period, have resulted in a 65% increase in the deterministically detrended real value of the stock market since 1980, an amount that exceeds 100% of the total increase. (Precisely, these shocks account for 110.5% of the increase.) The result provides a timely example of how small but near-permanent shocks to cash-flows can dramatically affect real stock market wealth over long time horizons. It also underscores the extent to which the long-term value of the stock market has been far more influenced by forces that redistribute the rewards of production, rather than by forces that raise or lower the total amount of rewards.

The levels decomposition implies that an additional 38% of the increase in stock market wealth since 1980, or a rise of 22%, is attributable to the cumulative effects of risk aversion shocks, which tended to be lower in the last 30 years than in earlier in the post-war period. By contrast, the cumulative affects of TFP shocks have made a negative contribution to change in stock market wealth since 1980, once a deterministic trend is removed. Although the role of this shock is in general quite modest, over this particular period it is non-trivial because its contribution was close to an all-time high (in our sample) in 1980 and close to an all-time low at the end of the sample. This latter observation is a direct consequence of the string of large negative draws for the consumption/productivity shock in the Great Recession years from 2007-2009. Our calculations imply that the stock market would be 22% higher today compared to 1980 had these productivity shocks been zero. These shocks accounted for -38% of the total increase since 1980. The residual accounts for the remaining -10.5% of the increase.

The findings are very similar if we instead study the real CRSP value-weighted stock price index plotted in the right-hand column of Figure 6. The decomposition implies that 103% of the increase in the deterministically detrended CRSP price index since 1980, or a rise of 59%, is attributable to the cumulative effects of the factors share shock; 34%, or a rise of 19%, is attributable to the cumulative effects of the risk aversion shock, and -22%, or a rise of -12.6%, is attributable to the cumulative effects of the productivity shock. The residual accounts for the remaining -14.5% of the increase.

We also investigated the relationship between the mutually orthogonal VAR innovations and earnings on the Standard & Poor (S&P) 500 stock market index. To do so, we estimate an empirical relationship analogous to (4), but instead of putting the log difference in stock wealth on the left-hand-side, we put log earnings changes on the left-hand-side. Quarterly S&P earnings are known to be noisy.[19] We therefore use long-term earnings growth (four year log changes in real per capita earnings) rather than quarterly log changes.

Figure 8 shows the dynamic response (not cumulated) of the four-year (16 quarter) log difference in real per capita earnings to the $\mathbf{e}_t$ shocks. The response of earnings to a con-

[19] S&P earnings fluctuate with time-variation in the application of accounting principles that affect the definition of some receipts and expenses and the timing with which some receipts and expenses are recorded (e.g., Petrick (2001)). These fluctuations may be strategic and therefore correlated with the economy or stock prices.

sumption/TFP shock is positive after a couple of quarters but much smaller than that the response to a factors share shock. A positive factors share shock (one that shifts rewards *toward* shareholders) leads to a much larger jump in long-run earnings that persists for many quarters. One aspect of these results that is not explained by the present version of the model with uncorrelated shocks is that there is a temporary response of earnings to the wealth/risk aversion shock: it is small but positive and grows over time before falling back again. This result is reminiscent of empirical findings showing that expected stock market returns and expected dividend growth rates are correlated (Lettau and Ludvigson (2005)). The model proposed above could be made more consistent with this last result by allowing shocks to be modestly correlated.

We close this section with a brief discussion of data on TFP and labor's share. It is tempting to consider comparing our VAR disturbances, which we take as proxies for the model primitive shocks, to external estimates of similar variables, where they exist. Figure A.1 in the Appendix shows that the low frequency variation in the model-implied productivity (consumption) shock computed from historical data exhibits a reasonably strong coherence with the estimate of TFP provided by Fernald (2009). There are some notable discrepancies at higher frequencies, but this is not surprising since the stylized model studied here has no variable labor supply or capital stock, as does the model and data underlying the Fernald TFP estimate. We find it encouraging that the low frequency variation in our consumption innovations corresponds so closely to the same variation in external estimates of TFP.

Similarly, it is tempting to compare our factor share component of labor income with external measures of labor's share of output, such as the Bureau of Labor Statistics (BLS) labor share measure. In this case, however, the two are not directly comparable. The BLS labor share is a measure of labor income *divided by GDP*, while our factor share component of labor income is (roughly) a measure of labor income *divided by consumption*, *holding fixed consumption.* These two measures can behave quite differently over time and have done so, especially in the last 15 years. The BLS labor share measure has fallen sharply in during this period, while the factor share component of labor income actually rose significantly over this period (Figure 7). Notice that the BLS labor share can fall without any change in labor income if profits not distributed to labor income are spent on investment or consumption. These changes in the BLS labor share would result in no change in the factor share component of labor income because they would not move the $y - c$ ratio holding fixed c. Part of the reason for the discrepancy between the two measures is that the BLS measure is negatively correlated with changes in the size of the consumption pie and in particular with TFP growth (Figure A.2 in the Appendix shows a plot). By contrast, our factor share component of labor income is by construction orthogonal to the size of the pie and productivity shocks, as required to identify the independent effects of these two innovations. Thus the factors share shock can be thought of as a component of labor's share that serves as an instrument for distinguishing the effects of labor's share shocks from productivity shocks. But since much of the movement in the BLS labor share is correlated with productivity changes, the two series can behave quite differently in some time periods.[20]

5.5 Why is the Stock Market Predictable?

We now turn to the subject of stock market predictability. A large and well known body of evidence finds that excess stock returns are forecastable over longer horizons, suggesting that the reward for bearing risk changes over time.[21] Several theories have been put forth to explain this forecastability, including habit formation (Campbell and Cochrane (1999)), or stochastic consumption volatility (Bansal and Yaron (2004)), but there is far less evidence that these mechanisms generate the observed predictability. This section provides evidence on this question by investigating sources of variation in common predictor variables such as the price-dividend ratio or the consumption-wealth variable cay_t (Lettau and Ludvigson (2001)). The results are presented in Table 5, with the top half showing results from historical data, and the bottom half showing results from the model.

Table 5 has several panels. The left panel reports regression results of one through five year log excess equity returns on the lagged price-dividend ratio alone. Moving rightward, the next panel reports regression results of one through five year log excess equity returns on lagged cay_t alone. We will also discuss the predictability of equity premia by measures of stochastic consumption volatility and uncertainty–the next panel reports regression results of log excess returns on one measure of stochastic consumption volatility. The panel to the right of this one, headed "$\mathbf{e_a}$ only," reports regression results of one through five year log excess equity returns on multiple lags of the i.i.d. wealth/risk aversion innovations $e_{a,t}$ and its lags. (The table reports the sum of the coefficients on all lags.) The next two columns show results when returns are predicted either by the component of the price-dividend ratio

[20] The model here could be extended to capture the stark decline in the BLS labor share in recent years by allowing a component of the factors share shock to be negatively correlated with the TFP shock. Of course, as explained, we would have no way of identifying the independent effects on endogenous variables of TFP and labor share shocks arising from this component.

[21] For extensive reviews of this evidence see Campbell, Lo, and MacKinlay (1997), Cochrane (2005), Lettau and Van Nieuwerburgh (2008), and Lettau and Ludvigson (2010).

that is *un*related to the wealth shocks, pd^{orth}, (movements in pd that are orthogonal to $e_{a,t}$ and its lags), or by the component of cay_t that is driven *only* by the wealth shocks, denoted cay_a. We motivate these regressions further below.

Several points bear noting. First, as previously shown in Table 3, in both the model and the data, the log price-dividend ratio and cay_t predict future excess returns with statistically significant coefficients and sizable adjusted R-squared statistics. By contrast, time-varying stochastic consumption volatility has no predictive power for equity premia at any horizon.[22] These results are robust to using additional lags of the stochastic volatility measure, to using first differences of the stochastic volatility measure, to using GARCH measures of consumption growth volatility, and to using measures of stochastic consumption growth volatility looking out over horizons greater than one quarter. In addition, broad-based measures of *macro*economic uncertainty developed in Jurado, Ludvigson, and Ng (2013) also exhibit no forecasting power for equity premia at any horizon. (These results are omitted to conserve space but are available upon request.) And none of these measures display any significant relationship with the empirical wealth shocks $e_{a,t}$ or their lags. These results provide no evidence that stock return predictability is driven by time-varying second moments of consumption growth or broad-based macroeconomic uncertainty.

Second, by comparison with pd or cay, lags of the $e_{a,t}$ wealth shocks exhibit greater forecasting power than either of these variables at every horizon except the five-year horizon. The adjusted R-squared statistics range from 14% at a one year horizon to 31% at a 3 year horizon before falling back to 18% at a 5 year horizon. A Wald test strongly rejects the hypothesis that the sum of squared coefficients on the lags of these shocks is zero.[23] A positive innovation for the wealth shock *in*creases asset wealth, so the negative coefficients in this forecasting regression imply that increases in wealth holding fixed consumption and labor income are transitory and forecast lower future returns.

Third, the predictive content for long horizon excess stock market returns contained in the pd_t and cay_t is subsumed by the information in lags of the wealth shocks at all but very long horizons, while at the longest horizons both variables have independent forecasting power. These results are in the two far right panels of Table 5. To show this, we construct

[22] Consumption volatility is measured as the conditional (on time t information) expectation of the squared innovation in consumption growth one quarter ahead. A stochastic volatility model is used to estimate $E_t\left(\left[\Delta \ln C_{t+h} - E_t\left[\Delta \ln C_{t+h}\right]\right]^2\right)$, for different horizons h. This estimate is taken from Jurado, Ludvigson, and Ng (2013).

[23] Wald tests similarly reject the hypothesis that the coefficients are jointly zero, and that the sum of coefficients is zero.

fitted residuals from regressions of the price-dividend ratio on the current and lagged values of the $e_{a,t}$ shock, denoted pd^{orth}. These residuals give the component of pd that is *not* attributable to the wealth shocks. We also construct the component of cay_t that is driven *only* by the wealth shocks as $cay_{a,t}$.[24] At all horizons except the five year, the pd residual components pd^{orth} have no statistically significant forecasting power for equity premia. And though the fitted residual components are statistically significant at the 5 year horizon, the R-square statistics are much smaller than those regressions for which the wealth shocks are used as predictor variables. There is no horizon at which the wealth shocks are not strongly statistically marginally significant. This suggests that the wealth shocks are still explaining the bulk of predictive power at the 5 year horizon even though the residual components have some predictive power at the this horizon. The results from the model are similar along this dimension.

The component of cay_t that is driven solely by the orthogonal wealth shocks, $e_{a,t}$, is responsible for the forecasting power of cay_t over all horizons except the 5 year. The adjusted R^2 statistic is, if anything, higher when using $cay_{a,t}$ rather than cay_t as a predictor variable to forecast equity premia. This evidence is difficult to reconcile with models in which risk premia vary with consumption shocks (e.g., Campbell and Cochrane (1999)), since the wealth shocks that "explain" most of the forecastability of excess returns are orthogonal to movements in consumption.

We have also computed the correlation between the wealth/risk aversion shocks we identify and stock market dividend growth in our sample. We find that they are contemporaneously unrelated, with a close-to-zero correlation of 0.06. The contemporaneous correlation with earnings growth is only slightly higher, 0.127. These results provide little evidence that the wealth/risk aversion shocks we identify that are, by construction, uncorrelated with consumption instead originate from shocks to measures of fundamental stock market value such as dividends or earnings.

In summary, according to this evidence, changes in the reward for bearing stock market risk are attributable to sources that are unrelated to traditional macroeconomic fundamentals, including aggregate consumption, labor income, measures of uncertainty or stochastic consumption volatility, dividend growth, or earnings growth.

[24] This is computed by inserting $\boldsymbol{\epsilon}_t \equiv \mathbf{H}\mathbf{e}_t$ into (1), using the initial values of $\mathbf{x}_t$ and $\Delta\mathbf{x}_t$ in the sample along with the estimated parameter values, and rolling forward (1) to obtain components of c_t, a_t, and y_t, attributable to each element of $\mathbf{e}_t$. These may then be combined with the estimated cointegrating vector $\boldsymbol{\alpha}$, to obtain components of cay_t attributable to each shock. See appendix for further details.

6 Economic Inequality

The causes and consequences of the upward trend in economic inequality over the last 30 years are hotly debated (see for example, Heathcote, Perri, and Violante (2010)). The model above has strong implications for this debate and indicates that fluctuations in economic inequality should be closely related to movements in the factors share shock. Figure 9 provides suggestive evidence that the cumulative effects of the factors share shock since 1980 may be associated with the observed rise in consumption inequality over this period. Figure 9 plots the consumption Gini coefficient from Heathcote, Perri, and Violante (2010), which uses data from the Consumer Expenditure Survey (CEX) over the period 1980 to 2006. Along with this series, we plot the cumulated factors share shock from the empirical VAR, and the model-implied consumption Gini obtained by feeding the observed sequence of factors share shocks from 1980 to 2006 into the model. (The Appendix gives the mapping between the consumption Gini and the cumulated factors share shocks in the model.) In the model, the consumption Gini is almost perfectly correlated with the cumulated factors share shocks, both of which rise over the 1980-2006 period as rewards shifted away from workers and toward shareholders.[25] This is not surprising since, in the model, all inequality is between group inequality across shareholders and workers, which is driven by the factors share shock. But there is also a striking low frequency correlation shown in Figure 9 between the rise in consumption inequality in the CEX data and the observed cumulated factors share shock, suggesting that the shift in rewards away from workers and toward shareholders over the last thirty years could be a driving force behind the rise in consumption inequality.

7 Conclusion

No comprehension of stock market behavior can be complete without understanding the origins of its fluctuations. Surprisingly little research has been devoted to this question. As a consequence, we have only a dimly lit view of why the real value of stock market wealth has evolved to its current level compared to five, or ten, or thirty years ago.

The starting point of this paper is to decompose real stock market fluctuations into components attributable to three mutually orthogonal observable economic shocks that explain the vast majority of fluctuations since the early 1950s. We then propose a model to interpret

[25] This calculation makes the (empirically relevant) assumption that equity holders' share of aggregate consumption is greater than their share in the population so that a shift in rewards toward shareholders increases rather than decreases consumption inequality.

these shocks and show that they are the observable empirical counterparts to three latent primitive shocks: a total factor productivity shock that benefits both workers and shareholders, a factors share shock that shifts the rewards of production between workers and shareholders without affecting the size of those rewards, and an independent risk aversion shock that shifts the stochastic discount factor pricing equities but is unrelated to aggregate consumption, labor earnings, or measures of fundamental value in the stock market.

Once we remove deterministic growth attributable to technological progress, the model implies that there are two big drivers of stock market wealth over time. One is a discount rate shock driven by fluctuations in investors' willingness to bear risk that is unrelated to real economic activity, including consumption, labor income, stock market dividends and earnings.[26] The other is a cash-flow innovation that redistributes the rewards of production between shareholders and workers with no change aggregate consumption. The independent discount rate shock dominates stock market volatility over periods of several quarters and a few years, while the factors share shock plays an increasingly important role as the time horizon extend into decades. Technological progress that raises aggregate consumption and benefits both workers and shareholders plays a small role in historical stock market fluctuations at all horizons.

A particularly striking example of the long-run implications of these economic shocks is provided by examining the period since 1980, a period in which the cumulative effect of the factor shares shock persistently rewarded shareholders at the expense of workers. After removing a deterministic trend, we find that factors share shocks have resulted in a 65% increase in real stock market wealth since 1980, an amount that exceeds 100% of the increase in stock market wealth over this period. Indeed, without these shocks, today's stock market would be about 10% lower than it was in 1980. This finding underscores the potential for small but near-permanent cash-flow shocks to dramatically affect equity values over long periods of time. It also implies that the *type* of shocks responsible for big historical movements in stock market wealth are not those that raise or lower aggregate rewards, but are instead those that redistribute a given level of rewards between workers and shareholders. This can be understood by noting that a unit increase in total factor productivity raises shareholder income by the fraction one minus the labor share, while a dollar transferred from workers raises shareholder income by 100% of the amount transferred.

[26] One real variable that in the data is related to the wealth shock is investment (Lettau and Ludvigson (2013)). But this is theoretically consistent with a discount rate shock, which should affect the present discounted value of marginal profits and therefore the optimal rate of investment (e.g., Abel (1983); Cochrane (1996)).

Our final results pertain to the question of why equity premia are forecastable over time. We find that most of the predictive content contained in common forecasting variables such as the price-dividend ratio or cay_t for future equity premia is subsumed by the information in lags of the empirical wealth shocks we identify. The predictability of excess stock market returns must therefore be understood as originating from sources largely unrelated to aggregate consumption, labor income, stock market earnings or dividends, measures of stochastic consumption volatility, or broad-based macroeconomic uncertainty. These facts are well explained by the model presented above, in which shareholders exhibit changes in their willingness to bear risk that is largely unrelated to the underlying economic fundamentals represented by these variables.

These latter findings accord well with a growing body of evidence from other recent studies which conclude that forecastable variation in equity market excess returns cannot be well explained by shocks to traditional economic fundamentals. Bianchi, Ilut, and Schneider (2013) find that business cycle variation in the price-GDP ratio must be explained by shocks to ambiguity-aversion unrelated to economic fundamentals. Muir (2014) conducts a cross-country analysis and finds that risk premia vary a lot during financial crises when consumption and consumption volatility are comparatively flat, while consumption falls a lot during wars when risk premia are comparatively flat. And Andersen, Fusari, and Todorov (2013) use data on stock market index options to find that risk-price variation responsible for time-varying equity premia is found to be largely unrelated to, and unidentifiable from, the stochastic evolution of economic risks.

The model presented here is deliberately stylized on the quantity side of the economy, abstracting from capital accumulation and fluctuations in employment. We have taken this approach in order to embed our analysis into an empirically plausible stock market environment. But there is growing evidence of important changes in the labor market over the last 30 years that have coincided with sharp declines in labor's share of output. These include the permanent disappearance of "middle skill" occupations not accompanied by a commensurate rise in employment elsewhere (Acemoglu (1999); Autor, Katz, and Kearney (2006); Jaimovich and Siu (2013)). Notably, the disappearance of these jobs does is not associated with a decline in output but instead reflects progress in technologies that more cheaply substitute for labor (Autor, Levy, and Murnane (2003)). Although the model we have explored here does not have an explicit role for changes in the level of employment, we consider the factors share shifter a reduced-form way of modeling these phenomena for understanding stock market behavior. Our findings imply that these forces for redistribution between shareholders and workers–whatever their cause–have had a profound effect on stock market wealth over longer periods of time.

Appendix

Data Description

CONSUMPTION

Consumption is measured as either total personal consumption expenditure or expenditure on nondurables and services, excluding shoes and clothing. The quarterly data are seasonally adjusted at annual rates, in billions of chain-weighted 2005 dollars. The components are chain-weighted together, and this series is scaled up so that the sample mean matches the sample mean of total personal consumption expenditures. Our source is the U.S. Department of Commerce, Bureau of Economic Analysis.

LABOR INCOME

Labor income is defined as wages and salaries + transfer payments + employer contributions for employee pensions and insurance - employee contributions for social insurance - taxes. Taxes are defined as [wages and salaries/(wages and salaries + proprietors' income with IVA and CCADJ + rental income + personal dividends + personal interest income)] times personal current taxes, where IVA is inventory valuation and CCADJ is capital consumption adjustments. The quarterly data are in current dollars. Our source is the Bureau of Economic Analysis.

POPULATION

A measure of population is created by dividing real total disposable income by real per capita disposable income. Our source is the Bureau of Economic Analysis.

WEALTH

Total wealth is household net worth in billions of current dollars, measured at the end of the period. A break-down of net worth into its major components is given in the table below. Stock market wealth includes direct household holdings, mutual fund holdings, holdings of private and public pension plans, personal trusts, and insurance companies. Nonstock wealth includes tangible/real estate wealth, nonstock financial assets (all deposits, open market paper, U.S. Treasuries and Agency securities, municipal securities, corporate and foreign bonds and mortgages), and also includes ownership of privately traded companies in noncorporate equity, and other. Subtracted off are liabilities, including mortgage loans and loans made under home equity lines of credit and secured by junior liens, installment consumer debt and other. Wealth is measured at the end of the period. A timing convention for wealth is needed because the level of consumption is a flow during the quarter rather than a point-in-time estimate as is wealth (consumption data are time-averaged). If we think of a given quarter's consumption data as measuring spending at the beginning of the quarter, then wealth for the quarter should be measured at the beginning of the period. If we think of the consumption data as measuring spending at the end of the quarter, then wealth for the quarter should be measured at the end of the period. None of our main findings discussed below (estimates of the cointegrating parameters, error-correction specification, or permanent-transitory decomposition) are sensitive to this timing convention. Given our finding that most of the variation in wealth is not associated with consumption, this timing convention is conservative in that the use of end-of-period wealth produces a higher contemporaneous correlation between consumption growth and wealth growth. Our source is the Board of Governors of the Federal Reserve System. A complete description of these data may be found at http://www.federalreserve.gov/releases/Z1/Current/.

STOCK PRICE, RETURN, DIVIDENDS

The stock price is measured using the Center for Research on Securities Pricing (CRSP) value-weighted stock market index covering stocks on the NASDAQ, AMEX, and NYSE. The data are monthly. The stock market price is the price of a portfolio that does not reinvest dividends. The CRSP dataset consists of $vwretx(t) = (P_t/P_{t-1}) - 1$, the return on a portfolio that doesn't pay dividends, and $vwretd_t = (P_t + D_t)/P_t - 1$, the return on a portfolio that does pay dividends. The stock price index we use is the price P_t^x of a portfolio that does not reinvest dividends, which can be computed iteratively as

$$P_{t+1}^x = P_t^x \left(1 + vwretx_{t+1}\right),$$

where $P_0^x = 1$. Dividends on this portfolio that does not reinvest are computed as

$$D_t = P_{t-1}^x \left(vwretd_t - vwretx_t\right).$$

The above give monthly returns, dividends and prices. The annual log return is the sum of the 12 monthly log returns over the year. We create annual log dividend growth rates by summing the log differences over the 12 months in the year: $d_{t+12} - d_t = d_{t+12} - d_{t+11} + d_{t+11} - d_{t+10} + \cdots + d_{t+1} - d_t$. The annual log price-dividend ratio is created by summing dividends in levels over the year to obtain an annual dividend in levels, D_t^A, where t denotes a year hear. The annual observation on P_t^x is taken to be the last monthly price observation of the year, P_t^{Ax}. The annual log price-dividend ratio is $\ln\left(P_t^{Ax}/D_t^A\right)$.

PRICE DEFLATOR

The nominal after-tax labor income and wealth data are deflated by the personal consumption expenditure chain-type deflator (2005=100), seasonally adjusted. In principle, one would like a measure of the price deflator for total flow consumption here. Since this variable is unobservable, we use the total expenditure deflator as a proxy. Our source is the Bureau of Economic Analysis.

Risk Aversion Along a Balanced Growth Path

The budget constraint for the representative shareholder can be written

$$A_t = \theta_t \left(P_t + D_t\right) + B_t \tag{26}$$

$$C_t^s + B_{t+1}q_t + \theta_{t+1}P_t \leq A_t, \tag{27}$$

where A_t are period t assets, θ_t are shares held in equity, P_t is the ex-dividend price of these shares, B_t is the beginning of period value of bonds held, and $q_t = 1/\left(1+R_f\right)$ is the risk-free rate paid on bonds. Along the non-stochastic balanced growth path, the equity return is equal to the risk-free bond rate. Rewrite (26) as

$$\begin{aligned} A_{t+1} &= P_t\theta_{t+1}\left(\frac{P_{t+1}+D_{t+1}}{P_t}\right) + B_{t+1} \\ &= P_t\theta_{t+1}R_{t+1} + B_{t+1} => \\ P_t\theta_{t+1} &= \frac{A_{t+1}}{R_{t+1}} - \frac{B_{t+1}}{R_{t+1}}. \end{aligned} \tag{28}$$

Plugging (28) into (27) and evaluating (27) at the equilibrium value of equality, we obtain

$$\begin{aligned} A_t &= C_t^s + B_{t+1}q_t + \frac{A_{t+1}}{R_{t+1}} - \frac{B_{t+1}}{R_{t+1}} \\ &= C_t^s + \frac{A_{t+1}}{R_{t+1}} \\ &= C_t^s + \frac{A_{t+1}}{R_f} \end{aligned}$$

where the last equality follows because $q_t = 1/R_f = 1/R_{t+1}$ along the equilibrium balanced growth path. Thus we have a beginning of period assets:

$$A_{t+1} = R_f\left(A_t - C_t^s\right)$$

or solving forward

$$A_t = \sum_{i=0}^{\infty}\left(\frac{1}{R_f}\right)^i C_{t+i}^s. \tag{29}$$

The value function is defined

$$V(A_t) = \max_{C_t^s} \{u(C_t^s) + \beta_t E_t V(A_{t+1})\}$$

or using (27)

$$V(A_t) = \max_{B_{t+1},\theta_{t+1}} \{u(A_t - B_{t+1}q_t + \theta_{t+1}P_t) + \beta_t E_t V(A_{t+1})\}.$$

Following the derivation in Swanson (2012), the coefficient of relative risk aversion RRA_t is

$$RRA_t = \frac{-A_t E_t V''(A_{t+1})}{E_t V'(A_{t+1})}. \tag{30}$$

It can be shown (see below) that

$$\begin{aligned} V'(A_{t+1}) &= u'\left(C_{t+1}^s\right) \Rightarrow && (31) \\ V''(A_{t+1}) &= u''\left(C_{t+1}^s\right)\frac{\partial C_{t+1}^{s*}}{\partial A_{t+1}}, && (32) \end{aligned}$$

where the notation "*" denotes the shareholder's optimal choice of C_{t+1}^s.

Swanson (2012) derives a relative risk aversion coefficient for dynamic models at a non-stochastic steady state or along a balanced growth path. We use $G = 1 + \mu_{\mathfrak{a}}$ to denote steady state growth. Along a balanced growth path where for $z \in \{A, C^s, D, C, Y\}$ we have $z_{t+k} = G^k z_t$,where $G \in (0, R_f)$ we have the following derivation in this setting. Note that the steady state value of $\beta_t \equiv \frac{\exp(-r_f)}{E_t\left[\frac{D_{t+1}^{-x_t}}{D_t^{-x_{t-1}}}\right]}$ is given by

$$\beta = \frac{\exp(-r_f)}{G^{\overline{x}}},$$

where the mean of x_t is denoted $\overline{x}$. Using the first order condition for optimal consumption choice in steady state:

$$\begin{aligned} u'\left(C_t^{s*}\right) &= \beta R_f u''\left(C_{t+1}^{s*}\right) => \\ u''\left(C_t^{s*}\right)\frac{\partial C_t^s}{\partial A_t} &= \beta R_f u''\left(C_{t+1}^{s*}\right)\frac{\partial C_{t+1}^{s*}}{\partial A_t} => \\ -\overline{x}\left(C_t^s\right)^{-\overline{x}-1}\frac{\partial C_t^{s*}}{\partial A_t} &= \frac{\exp(-r_f)}{G^{\overline{x}}} R_f\left[-\overline{x}\left(GC_t^{s*}\right)^{-\overline{x}-1}\right]\frac{\partial C_{t+1}^{s*}}{\partial A_t} => \\ \frac{\partial C_t^{s*}}{\partial A_t} &= G^{-1}\frac{\partial C_{t+1}^{s*}}{\partial A_t}. && (33) \end{aligned}$$

Applying the same transformation to the first order condition at time $t+1$ we have

$$\begin{aligned} u''\left(C_{t+1}^{s*}\right)\frac{\partial C_{t+1}^{s}}{\partial A_t} &= \beta R_f u''\left(C_{t+2}^{s*}\right)\frac{\partial C_{t+2}^{s*}}{\partial A_t} => \\ \frac{\partial C_{t+1}^{s*}}{\partial A_t} &= G^{-1}\frac{\partial C_{t+2}^{s*}}{\partial A_t} \end{aligned} \tag{34}$$

and combining (33) and (34) we obtain

$$\begin{aligned} \frac{\partial C_t^{s*}}{\partial A_t} &= G^{-2}\frac{\partial C_{t+2}^{s*}}{\partial A_t} => \\ \frac{\partial C_{t+2}^{s*}}{\partial A_t} &= G^{2}\frac{\partial C_t^{s*}}{\partial A_t}, \end{aligned}$$

and iterating obtain

$$\frac{\partial C_{t+i}^{s*}}{\partial A_t} = G^{i}\frac{\partial C_t^{s*}}{\partial A_t}.$$

Now differentiate (29) evaluated along the balanced growth path with respect to A_t:

$$\begin{aligned} 1 &= \sum_{i=0}^{\infty}\left(\frac{1}{R_f}\right)^{i}\frac{\partial C_{t+i}^{s*}}{\partial A_t} \\ 1 &= \frac{\partial C_t^{s*}}{\partial A_t}\left[1+\frac{G}{R_f}+\left(\frac{G}{R_f}\right)^2+\left(\frac{G}{R_f}\right)^3+\cdots\right] \\ &= \frac{\partial C_t^{s*}}{\partial A_t}\left(\frac{R_f}{R_f-G}\right), \end{aligned}$$

implying

$$\frac{\partial C_t^{s*}}{A_t} = \frac{R_f-G}{R_f}. \tag{35}$$

Assets A_t along a non-stochastic balanced growth path are

$$\begin{aligned} A_t &\equiv \sum_{i=0}^{\infty}\left(\frac{1}{R_f}\right)^{i} C_{t+i}^{s*} \\ &= \sum_{i=0}^{\infty}\left(\frac{G}{R_f}\right)^{i} C_t^{s*} \\ &= \frac{R_f C_t^{s*}}{(R_f-G)}. \end{aligned} \tag{36}$$

Plugging (31), (32), (35), and (36) into (30), we obtain a value for risk aversion along a balanced growth path equal to

$$RRA_t = \frac{-C_t^{s}u''\left(C_{t+1}^{s}\right)}{u'\left(C_{t+1}^{s}\right)} = \overline{x}/G.$$

Derivation of (31). First-order condition for B_{t+1}:

$$-u'\left(C_t^s\right) q_t + \beta_t V'\left(A_{t+1}\right) \underbrace{\frac{\partial A_{t+1}}{\partial B_{t+1}}}_{=1 \text{ from (26)}} = 0. \tag{37}$$

First-order condition for θ_{t+1}:

$$-u'\left(C_t^s\right) P_t + \beta_t V'\left(A_{t+1}\right)\left(P_{t+1} + D_{t+1}\right) = 0. \tag{38}$$

Differentiate the value function

$$V\left(A_t\right) = \max_{B_{t+1}, \theta_{t+1}} \left\{u\left(A_t - B_{t+1} q_t + \theta_{t+1} P_t\right) + \beta_t E_t V\left(A_{t+1}\right)\right\}$$

with respect to A_t, keeping in mind

$$C_t^s = A_t - B_{t+1} q_t + \theta_{t+1} P_t$$

and

$$A_{t+1} = \theta_{t+1}\left(P_{t+1} + D_{t+1}\right) + B_{t+1}.$$

We have

$$\begin{aligned} V'\left(A_t\right) &= u'\left(C_t^s\right) + \left[-u'\left(C_t^s\right) q_t + \beta_t V'\left(A_{t+1}\right) \underbrace{\frac{\partial A_{t+1}}{\partial B_{t+1}}}_{=1}\right] \frac{\partial B_{t+1}}{\partial A_t} \\ &\quad + \left[-u'\left(C_t^s\right) P_t + \beta_t V'\left(A_{t+1}\right)\left(P_{t+1} + D_{t+1}\right)\right] \frac{\partial \theta_{t+1}}{\partial A_t}. \end{aligned}$$

Evaluating at the optimum using (37) and (38), the terms in brackets are zero, leaving

$$V'\left(A_t\right) = u'\left(C_t^s\right).$$

Numerical Solution

The price-dividend ratio satisfies

$$\begin{aligned} \frac{P_t}{D_t}\left(s_t\right) &= E_t\left[M_{t+1}\left(\frac{P_{t+1}}{D_{t+1}}\left(\mathfrak{s}_{t+1}\right) + 1\right) \frac{D_{t+1}}{D_t}\right] \\ &= E_t \exp\left(m_{t+1} + \Delta d_{t+1} + \ln\left(\frac{P_{t+1}}{D_{t+1}}\left(s_{t+1}\right) + 1\right)\right), \end{aligned}$$

where $\mathfrak{s}_t$ is a vector of state variables, $\mathfrak{s}_t \equiv (\Delta \ln a_t, Z_t, x_t)'$. We therefore solve the function numerically on an $n \times n \times n$ dimensional grid of values for the state variables, replacing the continuous time processes with a discrete Markov approximation following the approach in Rouwenhorst (1995). The continuous function $\frac{P_t}{D_t}(\mathfrak{s}_t)$ is then replaced by the $n \times n \times n$ functions $\frac{P_t}{D_t}(i, j, k)$, $i, j, k = 1, ..., N$, each representing the price-dividend ratio in state $\Delta \ln a_i$, Z_j, and x_k, where the functions are defined recursively by

$$\frac{P}{D}(i, j, k) = \sum_{l=1}^{n} \sum_{m=1}^{n} \sum_{n=1}^{n} \pi_{i,l} \pi_{k,n} \pi_{j,m} \exp\left(m(l, m, n) + \Delta d(l, m, n) + \ln\left(\frac{P}{D}(l, m, n) + 1\right)\right),$$

where $m(l, m, n)$ refers to the values m_{t+1}can take on in each of the states, and analogously for the other terms. We set $N = 35$.

Estimating Population Spectrum for the Level of Stock Market Wealth

Here we discuss the level decomposition of variance based on a spectral decomposition. The reference for this procedure is Hamilton (1994), chapter 6. The procedure may be summarized as follows. First, we estimate (4) and plug the estimated parameters into formulas for the population spectrum for each component in (6) Δs_t^c, Δs_t^y, Δs_t^a, and i.i.d. residual η_t. Since $\Delta s_t - \kappa_0 = \Delta s_t^c + \Delta s_t^y + \Delta s_t^a + \eta_t$, the sum of the estimated spectra for each component gives the estimated spectrum for $\Delta s_t - \kappa_0$, denoted $S_{\Delta s}(\omega)$ as a function of cycles of frequency ω. Notice that we remove the deterministic trend from the log level of stock market wealth by subtracting κ_0 from Δs_t on the right-hand-side. Thus we have $S_{\Delta s}(\omega) = S_{\Delta s^c}(\omega) + S_{\Delta s^y}(\omega) + S_{\Delta s^a}(\omega) + S_\eta(\omega)$, where these right hand terms are the spectra for the individual components of $\Delta s_t - \kappa_0$. Roughly speaking, the proportion of sample variance in $\Delta s_t - \kappa_0$ attributable to cycles with frequency near ω is given by $S_{\Delta s}(\omega) 4\pi/T$, where T is the sample size. The fraction of the variance in the $\Delta s_t - \kappa_0$ at cycles with frequency near ω that is attributable to the consumption shock is

$$\frac{S_{\Delta s^c}(\omega)}{S_{\Delta s}(\omega)}, \tag{39}$$

and fraction of the variance in the $\Delta s_t - \kappa_0$ at cycles with frequency ω that is attributable to the other components are defined analogously. Recalling that, if the frequency of the cycle is ω, the period of the cycle is $2\pi/\omega$. Thus we plot (39), which is a function of frequencies $\omega_j = 2\pi j/T$, against periods $2\pi/\omega_j = T/j$ (here in units of quarters), where T is the sample size.

Bootstrap Procedure for Error Bands

Confidence intervals for parameters of interest are generated from a bootstrap following Gonzalo and Ng (2001). The procedure is as follows. First, the cointegrating vector is estimated, and conditional on this estimate, the remaining parameters of the VECM and subsequent regressions are estimated. The fitted residuals from the system

$$\begin{aligned} \Delta \mathbf{x}_t &= \widehat{\nu} + \widehat{\gamma}\widehat{\alpha}' \mathbf{x}_{t-1} + \widehat{\Gamma}(L)\Delta \mathbf{x}_{t-1} + \widehat{H} e_t \\ \Delta s_t &= \widehat{\kappa}_0 + \widehat{\kappa}_c(L) e_{c,t} + \widehat{\kappa}_y(L) e_{y,t} + \widehat{\kappa}_a(L) e_{a,t} + \eta_t, \end{aligned} \tag{40}$$

denoted $(\widehat{e}_{c,t}, \widehat{e}_{y,t}, \widehat{e}_{a,t}, \widehat{\eta}_t)$ are obtained and a new sample of data is constructed (conditional on our initial observations $\mathbf{x}_{-1}, \mathbf{x}_0$ and s_0,) using the initial VECM and stock wealth OLS parameter estimates by random sampling of $(\widehat{e}_{c,t}, \widehat{e}_{y,t}, \widehat{e}_{a,t}, \widehat{\eta}_t)$ with replacement. Denote the new randomly sampled (via block bootstrap) values for the residuals $(\widetilde{e}_{c,t}, \widetilde{e}_{y,t}, \widetilde{e}_{a,t}, \widetilde{\eta}_t)$ for $t = 1, \ldots, T$. The new bootstrapped sample of observable data, $(\widetilde{\mathbf{x}}_t, \widetilde{s}_t)$, is constructed from

$$\begin{aligned} \Delta \widetilde{\mathbf{x}}_t &= \widehat{\nu} + \widehat{\gamma}\widehat{\alpha}' \mathbf{x}_{t-1} + \widehat{\Gamma}(L)\Delta \mathbf{x}_{t-1} + \widehat{H} \widetilde{e}_t \\ \Delta \widetilde{s}_t &= \widehat{\kappa}_0 + \widehat{\kappa}_c(L) \widetilde{e}_{c,t} + \widehat{\kappa}_y(L) \widetilde{e}_{y,t} + \widehat{\kappa}_a(L) \widetilde{e}_{a,t} + \widetilde{\eta}_t. \end{aligned}$$

Given this new sample of data, all parameters in (40) (as well as the cointegrating coefficients) are re-estimated, and the impulse responses, variance decompositions, and other statistics of interest stored. This is repeated 5,000 times. The empirical 90% confidence intervals are evaluated from these 5,000 samples of the bootstrapped parameters. The bands for the impulse responses in Figure 3 are reported in Figure A.2. The bands are reasonably tight for most responses except for the response of a_t to a labor income shock $e_{y,t}$. However, Figure 4 shows that the response of stock wealth s_t (rather than net worth) to an $e_{y,t}$ shock is estimated much more precisely, reflecting the fact that the factors share shock affects the stock wealth component of net worth almost exclusively, but shows little relation to other forms of wealth included in net worth.

Consumption Gini

We explain how the model-implied consumption Gini coefficient is computed over the same sample period as in the empirical consumption Gini series of Heathcote, Perri, and Violante (2010). In the model, inequality is entirely attributable to the division of consumption between shareholders and workers (agents in each group are identical, so there is no within-group inequality). The level of inequality in the model is measured, as in the data, using the Gini coefficient for consumption. To calculate inequality in the model, we assume that the fraction of shareholders is smaller than their share of aggregate consumption, so that shareholders consume a disproportionately large fraction of aggregate consumption. This assumption ensures that a shift of income away from workers toward shareholders (i.e., a negative ε_z shock) has the effect of *increasing* consumption inequality. The share of aggregate consumption that accrues to workers is $\alpha f(Z_t)$. If we denote q to be the fraction of the population in the shareholder group, then we assume $q < 1 - \alpha f(z)$ for all z.

Under these assumptions, the Gini coefficient takes the simple form

$$G = 1 - q - \alpha f(z). \tag{41}$$

To see this, it is helpful to consider Figure A.4, which shows the consumption distribution in the model. The Gini coefficient is defined to be the ratio $A/(A+B)$, where A and B are the areas of the relevant labeled areas in Figure A.4. The area B is the sum of the areas of a triangle with base $1-q$ and height $\alpha f(z)$, a rectangle with base q and height $\alpha f(z)$, and a triangle with base q and height $1-\alpha f(z)$. Basic geometry then implies that

$$B = \frac{1}{2}(q + \alpha f(z)).$$

$A + B$ is a triangle with base 1 and height 1, so $A + B = 1/2$. Combining results, we obtain that

$$A = \frac{1}{2}(1 - q - \alpha f(z)).$$

Since $G = A/(A+B)$, this completes the derivation of (41).

Given this form for the Gini coefficient, it is clear that in the model, the Gini coefficient can be determined up to the constant q given values for $f(z)$. In the model, the e_y shock is nearly perfectly correlated with $\Delta \ln f(Z_t)$, so that

$$\Delta \ln f(Z_t) \simeq b_1 e_{y,t} \tag{42}$$

for some constant b_1. The constant is estimated by running the relevant regression using long time series simulated from the model.

Using our estimates $\widehat{e}_{y,t}$ from the empirical VAR, we can now construct an implied series for the Gini coefficient in the model. Note from (42), we have

$$\ln f(Z_t) = \ln f(Z_0) + b_1 \sum_{i=1}^{t} e_{y,i}. \tag{43}$$

The consumption Gini data from Heathcote, Perri, and Violante (2010) are annual and run from 1980 to 2006. We therefore set $t = 1$ to 1980 and normalize $\ln f(Z_0)$ to zero. We take the average quarterly value of $e_{y,t}$ within a year as the annual observation for $e_{y,t}$. Iterating forward on (43) and applying the exponential function to the left-hand-side yields an implied series for $f(z_t)$. Finally, plugging this value into (41) yields a model-implied series for G. Because we normalize the Gini series in the plot to have zero mean, the parameter q doesn't play a role in the plotted series.

Decomposition of *cay*

This section describes how to decompose the cay_t series obtained from a VECM regression into components attributable to each of the three orthogonalized shocks. The cay series is defined by

$$cay_t \equiv \alpha' x_t - \kappa$$

where $x_t = (c_t, a_t, y_t)'$, $\alpha = (1, -\alpha_a, -\alpha_y)'$ is the cointegrating vector, and κ is a constant.

Assume that the stochastic process for x_t has a VECM representation

$$\Delta x_t = \nu + \gamma\alpha' x_{t-1} + \Gamma\Delta x_{t-1} + He_t \tag{44}$$

where $e_t = (e_{c,t}, e_{a,t}, e_{y,t})'$ are the orthogonalized shocks. Inverting the VECM, we obtain the Wold decomposition

$$\Delta x_t = \delta + D(L)e_t \tag{45}$$

where $D(L)$ is an infinite-order lag polynomial.

Theoretical Decomposition

If we let $D_c(L)$ be the column of $D(L)$ relating to the e_c shock, then we obtain the decomposition

$$\Delta x_t = \delta + D_c(L)e_{c,t} + D_a(L)e_{a,t} + D_y(L)e_{y,t}. \tag{46}$$

If we define

$$\begin{aligned} \Delta x_{c,t} &\equiv D_c(L)e_{c,t} \\ \Delta x_{a,t} &\equiv D_a(L)e_{a,t} \\ \Delta x_{y,t} &\equiv D_y(L)e_{y,t} \end{aligned} \tag{47}$$

then, cumulating up, we obtain an additive decomposition for x_t

$$x_t = \delta t + x_{c,t} + x_{a,t} + x_{y,t}. \tag{48}$$

Premultiplying (48) by α', and subtracting κ, we obtain a decomposition for cay_t:

$$cay_t = \alpha' x_t - \kappa = \alpha'\delta t + \alpha' x_{c,t} + \alpha' x_{a,t} + \alpha' x_{y,t} - \kappa \tag{49}$$

If we define

$$\begin{aligned} cay_{c,t} &\equiv \alpha' x_{c,t} \\ cay_{a,t} &\equiv \alpha' x_{a,t} \\ cay_{y,t} &\equiv \alpha' x_{y,t} \end{aligned}$$

then (49) becomes

$$cay_t = \alpha'\delta t + cay_{c,t} + cay_{a,t} + cay_{y,t} - \kappa. \tag{50}$$

Note that unlike cay_t, the components (e.g., $cay_{c,t}$) are detrended and demeaned.

Practical Decomposition

Unfortunately, we cannot fully implement (50) in practice. From our estimation procedure, we obtain estimates $\widehat{\alpha}$, and $\widehat{D}$. We also obtain estimates $\widehat{e}_t$, but only for $t = 1, \ldots, T$. If we had estimates of all $\widehat{e}_t$ for $t = T, T-1, \ldots, -\infty$, then we could evaluate estimates of each x_c, x_a and x_y term using (47) and form an additive decomposition. However, we do not, so if we estimate these terms using only the shocks from $1, \ldots, T$, then the decomposition will no longer hold exactly, as we cannot decompose what shocks were responsible for the initial conditions Δx_0 and x_0, which will have a persistent effect on the the series x_t through (44). However, since the effect of initial conditions becomes smaller as time goes on, the approximate decomposition using only shocks from $t = 1$ forward may still be of interest.

Our approximate decomposition begins with the quantities

$$\begin{aligned} \Delta\widetilde{x}_{c,t} &\equiv \sum_{j=0}^{t-1} D_{c,j}\widehat{e}_{c,t-j} \\ \Delta\widetilde{x}_{a,t} &\equiv \sum_{j=0}^{t-1} D_{a,j}\widehat{e}_{a,t-j} \\ \Delta\widetilde{x}_{y,t} &\equiv \sum_{j=0}^{t-1} D_{y,j}\widehat{e}_{y,t-j} \end{aligned}$$

where $\widehat{e}_t$ represents the estimates of the orthogonalized shocks. This decomposes the influence of the orthogonalized shocks from $t = 1$ on the various states. Cumulating these series leads to the series $\widetilde{x}_c, \widetilde{x}_a, \widetilde{x}_y$. In practice, these series can be calculated by running the VECM (44) forward (without the constant term) starting from initial condition $\Delta x_0 = x_0 = (0,0,0)'$ and applying shocks that include only the relevant entry of the estimated shocks.

An example will clarify these instructions. To evaluate $\Delta\widetilde{x}_{c,t}$, begin by setting $\Delta\widetilde{x}_{c,0} = \widetilde{x}_{c,0} = (0,0,0)'$. Given $\Delta\widetilde{x}_{t-1}$ and $\widetilde{x}_{t-1}$ for $t \geq 0$, we can obtain $\Delta\widetilde{x}_{c,t}$ by applying (44) without the constant term ν, and allowing only the $e_{c,t}$ shock to be nonzero, so that we obtain

$$\Delta\widetilde{x}_{c,t} = \widehat{\gamma}\widehat{\alpha}'\widetilde{x}_{c,t-1} + \widehat{\Gamma}\Delta\widetilde{x}_{c,t-1} + H \begin{bmatrix} e_{c,t} \\ 0 \\ 0 \end{bmatrix}. \tag{51}$$

Proceeding in this fashion, we can compute the entire series for $\widetilde{x}_{c,t}$.

To decompose the effect of the shocks on cay_t, we can simply apply the cointegrating vector to the cumulated $\widetilde{x}$ series to obtain

$$\begin{aligned} \widetilde{cay}_{c,t} &\equiv \alpha'\widetilde{x}_{c,t} \\ \widetilde{cay}_{a,t} &\equiv \alpha'\widetilde{x}_{a,t} \\ \widetilde{cay}_{y,t} &\equiv \alpha'\widetilde{x}_{y,t} \end{aligned}$$

Because this decomposition cannot account for shocks prior to $t = 1$, (49) will not hold, and we will instead end up with a "residual" term cay^*_t, such that

$$cay_t = \widetilde{cay}_{c,t} + \widetilde{cay}_{a,t} + \widetilde{cay}_{y,t} + cay^*_t \tag{52}$$

This completes the instructions for generating the decomposition. A full derivation of the decomposition, including instructions for calculating the residual term, can be found below.

Full Derivation

In order to derive (52), an additive decomposition without estimates of $e_0, e_{-1}, \ldots$, we can take advantage of the fact that the influence of $e_0, e_{-1}, \ldots$ is contained in the initial conditions Δx_0 and x_0. The first step is to define a series that represents only the effects of these unobserved shocks. To this end, define

$$e^*_t \equiv \begin{cases} e_t & \text{for } t \leq 0 \\ 0 & \text{for } t > 0 \end{cases}$$

so that e_t^* is equivalent to e_t for $t \leq 0$, but is zero for $t > 0$. Next, define

$$\Delta x_t^* \equiv \delta + D(L)e_t^* = \delta + \sum_{j=t}^{\infty} D_j e_{t-j}.$$

In other words, Δx_t^* is what Δx_t would be had all shocks from time $t = 1$ on been equal to zero. For examples, we have

$$\begin{aligned} \Delta x_1^* &= \delta + D_1 e_0 + D_2 e_1 + \ldots \\ \Delta x_2^* &= \delta + D_2 e_0 + D_3 e_1 + \ldots . \end{aligned}$$

To compute these series, we can easily obtain this series by running the VECM forward with all shocks from $t = 1$ onward set to zero. Specifically, use the initial conditions $\Delta x_0^* = \Delta x_0$ and $x_0^* = x_0$, as in the standard VECM. Then, given Δx_{t-1}^* and x_{t-1}^*, we can compute Δx_t^* using

$$\Delta x_t^* = \nu + \gamma \alpha' x_{t-1}^* + \Gamma \Delta x_{t-1}^*$$

which is just the standard VECM (including the constant term nu) but with the shocks e_t set to zero. In practice of course the coefficients of the VECM will be the estimated "hat" versions.

Next, define

$$\begin{aligned} \Delta \widetilde{x}_{c,t} &\equiv \sum_{j=0}^{t-1} D_{c,j} e_{c,t-j} \\ \Delta \widetilde{x}_{a,t} &\equiv \sum_{j=0}^{t-1} D_{a,j} e_{a,t-j} \\ \Delta \widetilde{x}_{y,t} &\equiv \sum_{j=0}^{t-1} D_{y,j} e_{y,t-j} \end{aligned}$$

so that each series corresponds to the cumulated effects of the different shocks from time $t = 1$ onward. Note that these series can be obtained by running the VECM forward (without adding the constant δ) starting from initial condition $\widetilde{x}_0 = \Delta \widetilde{x}_0 = (0, 0, 0)'$ and applying the relevant shock components one at a time, as in (51) of the previous section.

Under this definition, we have

$$\Delta \widetilde{x}_{c,t} + \Delta \widetilde{x}_{a,t} + \Delta \widetilde{x}_{y,t} = \sum_{j=0}^{t-1} D_j e_{t-j}$$

and since

$$\Delta x_t^* = \delta + \sum_{j=t}^{\infty} D_j e_{t-j}$$

we obtain

$$\Delta \widetilde{x}_{c,t} + \Delta \widetilde{x}_{a,t} + \Delta \widetilde{x}_{y,t} + \Delta x_t^* = \delta + \sum_{j=0}^{\infty} D_j e_{t-j} = x_t.$$

Since all of these components depend only on the estimated coefficients, the initial condition x_0, and the estimated shocks $\widehat{e}_1, \ldots, \widehat{e}_T$, we can calculate this decomposition using only the output from the VECM regression.

Cumulating, and applying the α vector, we obtain

$$\begin{aligned} cay_t &= \delta t \alpha' + \alpha' \widetilde{x}_{c,t} + \alpha' \widetilde{x}_{a,t} + \alpha' \widetilde{x}_{y,t} + \alpha' x_t^* - \kappa \\ &= \widetilde{cay}_{c,t} + \widetilde{cay}_{a,t} + \widetilde{cay}_{y,t} + cay_t^* \end{aligned}$$

for

$$\begin{aligned} \widetilde{cay}_{c,t} &\equiv \alpha' \widetilde{x}_{c,t} \\ \widetilde{cay}_{a,t} &\equiv \alpha' \widetilde{x}_{a,t} \\ \widetilde{cay}_{y,t} &\equiv \alpha' \widetilde{x}_{y,t} \\ cay_t^* &\equiv \alpha' x_t^* - \kappa \end{aligned}$$

which is the desired additive decomposition.

A few final notes are in order. First, since we want an additive decomposition, it is important not to triple-count various constants. This means not including the δ constant when calculating the various $\widetilde{x}$ series, as well as not normalizing the various $\widetilde{cay}$ series by κ (although these constant terms should be used when calculating cay_t^*).

References

Abel, A. B. (1983): "Optimal Investment under Uncertainty," *American Economic Review*, 73(1), 228–233.

Acemoglu, D. (1999): "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence," *American Economic Review*, 89(5), 1259–1278.

Albuquerque, R., M. Eichenbaum, and S. Rebelo (2012): "Valuation Risk and Asset Pricing," Unpublished Paper, Northwestern University.

Andersen, T. G., N. Fusari, and V. Todorov (2013): "The Risk Premia Embedded in Index Options," Unpublished paper, Northwester Kellog.

Autor, D. H., L. F. Katz, and M. S. Kearney (2006): "The Polarization of the U.S. Labor Market," *American Economic Review: Papers and Proceedings*, 96, 189–194.

Autor, D. H., F. Levy, and R. J. Murnane (2003): "The Skill content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118(4), 1279–1333.

Bansal, R., and A. Yaron (2004): "Risks for the Long-Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59(4), 1481–1509.

Bianchi, F., C. Ilut, and M. Schneider (2013): "Uncertainty Shocks, Asset Supply and Pricing Over the Business Cycle," Unpublished paper, Duke University.

Brunnermeier, M., and S. Nagel (2008): "Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation," *American Economic Review*, 98(3), 713–736.

Brunnermeier, M., and Y. Sannikov (2012): "A Macroeconomic Model with a Financial Sector," *American Economic Review,* forthcoming.

Campbell, J. Y. (1991): "A Variance Decomposition for Stock Returns," *Economic Journal*, 101, 157–179.

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

Campbell, J. Y., A. W. Lo, and C. MacKinlay (1997): *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ.

Campbell, J. Y., and P. Perron (1991): "Pitfalls and Opportunities: What Macroeconomists Should Know About Unit Roots," in *NBER Macroeconomics Annual: 1991*, ed. by O. J. Blanchard, and S. Fischer. MIT Press, Cambridge, MA.

Chen, X., J. Favilukis, and S. C. Ludvigson (2014): "An Estimation of Economic Models with Recursive Preferences," *Quantitative Economics,* forthcoming.

Cochrane, J. H. (1991): "Explaining the Variance of Price-Dividend Ratios," *Review of Financial Studies*, 5(2), 243–280.

——— (1994): "Permanent and Transitory Components of GDP and Stock Prices," *Quarterly Journal of Economics*, 109, 241–265.

——— (1996): "A Cross-Sectional Test of an Investment-Base Asset Pricing Model," *Journal of Political Economy*, 104, 572–621.

——— (2005): *Asset Pricing, Revised Edition*. Princeton University Press, Princeton, NJ.

——— (2008): "The Dog That Did Not Bark: A Defense of Return Predictability," *Review of Financial Studies*, 21(4), 1533–1575.

Duffee, G. R. (2005): "Time Variation in the Covariances Between Stock Returns and Consumption Growth," *Journal of Finance*, 60(4), 1673–1712.

Engle, R. F., and C. W. J. Granger (1987): "Co-integration and Error Correction Representation, Estimation and Testing," *Econometrica*, 55, 251–276.

Fama, E. F., and K. R. French (1989): "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics*, 25, 23–49.

Fernald, J. (2009): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Unpublished paper, Federal Reserve Bank of San Franscisco.

Gabaix, X., and M. Maggiori (2013): "International Liquidity and Exchange Rate Dynamics," Unpublished paper, New York University.

Gonzalo, J., and S. Ng (2001): "A Systematic Framework for Analyzing the Dynamic Effects of Permanent and Transitory Shocks," *Journal of Economic Dynamics and Control*, 25(10), 1527–1546.

Guvenen, M. F. (2009): "A Parsimonious Macroeconomic Model for Asset Pricing," *Econometrica*, 77(6), 1711–1740.

Hamilton, J. D. (1994): *Time Series Analysis*. Princeton University Press, Princeton, NY.

Hansen, L. P., and R. Jagannathan (1991): "Restrictions on Intertemporal Marginal Rates of Substitution Implied by Asset Returns," *Journal of Political Economy*, 99, 225–262.

He, Z., and A. Krishnamurthy (2013): "Intermediary Asset Pricing," *American Economic Review*, 103(2), 732–770.

Heathcote, J., F. Perri, and G. L. Violante (2010): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967-2006," *Review of Economic Dynamics*, 13(1), 15–51.

Jaimovich, N., and H. E. Siu (2013): "The Trend is the Cycle: Job Polarization and Jobless Recoveries," Unpublished paper, Duke University.

Jurado, K., S. C. Ludvigson, and S. Ng (2013): "Measuring Uncertainty," Unpublished paper, Columbia University.

Kacperczyk, M., J. B. Nosal, and L. Stevens (2014): "Investor Sophistication and Capital Income Inequality," Unpublished manuscript, Columbia University.

Karabarbounis, L., and B. Neiman (2013): "The Global Decline of the Labor Share," *Quarterly Journal of Economics*, 129(1), 61–103.

Lansing, K. J. (2014): "Asset Pricing with Concentrated Ownership of Capital and Distribution Shocks," Federal Reserve Bank of San Francisco Working Paper 2011-07.

LeRoy, S., and R. Porter (1981): "The Present Value Relation: Tests Based on Variance Bounds," *Econometrica*, 49, 555–557.

Lettau, M., and S. C. Ludvigson (2001): "Consumption, Aggregate Wealth and Expected Stock Returns," *Journal of Finance*, 56(3), 815–849.

——— (2004): "Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption," *American Economic Review*, 94(1), 276–299.

——— (2005): "Expected Returns and Expected Dividend Growth," *Journal of Financial Economics*, 76, 583–626.

——— (2009): "Euler Equation Errors," *The Review of Economic Dynamics*, 12(2), 255–283.

——— (2010): "Measuring and Modeling Variation in the Risk-Return Tradeoff," in *Handbook of Financial Econometrics*, ed. by Y. Ait-Sahalia, and L. P. Hansen, vol. 1, pp. 617–690. Elsevier Science B.V., North Holland, Amsterdam.

Lettau, M., and S. C. Ludvigson (2013): "Shocks and Crashes," in *National Bureau of Economics Macroeconomics Annual: 2013*, ed. by J. Parker, and M. Woodford. MIT Press, Cambridge and London.

Lettau, M., and S. Van Nieuwerburgh (2008): "Reconciling the Return Predictability Evidence: In-Sample Forecasts, Out-of-Sample Forecasts, and Parameter Instability," *Review of Financial Studies*, 21(4), 1607–1652.

Lettau, M., and J. A. Wachter (2007): "Why is Long-Horizon Equity Less Risky? A Duration Based Explanation of the Value Premium," *Journal of Finance*, LXII(1), 55–92.

Lustig, H., and S. Van Nieuwerburgh (2008): "The Returns on Human Capital: Good News on Wall Street is Bad News on Main Street," *Review of Financial Studies*, 21, 2097–2137.

Malloy, C. J., T. J. Moskowitz, and A. Vissing-Jorgensen (2009): "Long-run Stockholder Consumption Risk and Asset Returns," *Journal of Finance*, 64, 2427–2479.

Mankiw, N. G., and S. P. Zeldes (1991): "The Consumption of Stockholders and Nonstockholders," *Journal of Financial Economics*, 29(1), 97–112.

Muir, T. (2014): "Financial Crises and Risk Premia," Unpublished paper, Yale University.

Newey, W. K., and K. D. West (1987): "A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Ng, S., and P. Perron (1997): "Estimation and Inference in Nearly Unbalanced Nearly Cointegrated Systems," *Journal of Econometrics.*

Petrick, K. A. (2001): "Comparing NIPA Profits with S& P 500 Profits," *Survey of Current Business*, pp. 16–20.

Rouwenhorst, K. G. (1995): "Asset Pricing Implications of Equilibrium Business Cycle Models," in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley, pp. 294–330. Princeton University Press, Princeton, New Jersey.

Shiller, R. J. (1981): "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?," *American Economic Review*, 71, 421–436.

Stock, J. H. (1987): "Asymptotic Properties of Least Squares Estimators of Cointegrating Vectors," *Econometrica*, 55, 113–144.

Stock, J. H., and M. W. Watson (1993): "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems," *Econometrica*, 61, 783–820.

Swanson, E. T. (2012): "Risk Aversion and the Labor Margin," *American Economic Review*, 102, 1663–1691.

Vissing-Jorgensen, A. (2002): "Limited Asset Market Participation and Intertemporal Substitution," *Journal of Political Economy*, 110(4), 825–853.

Wachter, J. (2013): "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Activity?," *Journal of Finance*, 68, 987–1035.

Table 1: Parameter List and Calibration

No.	Parameter	Description	Calibration
1	α	exponent on labor in production function	0.667
2	ψ	constant term in $f(Z_t)$	0.5
3	r^f	log risk-free rate	0.0038
3	$\underline{\theta}$	minimum value x_t	1
4	$\bar{\theta}$	maximum value x_t	455
5	$\bar{a}$	mean (quarterly) $\Delta\mathfrak{a}_t$	0.0057
6	$f(\bar{Z})$	mean $f(Z_t)$	1
7	$\phi_{\mathfrak{a}}$	autocorrelation $\Delta\mathfrak{a}_t$	0
8	ϕ_z	autocorrelation Z_t	0.995
9	$\phi_{\widetilde{x}}$	autocorrelation $\widetilde{x}_t$	0.85
10	$\sigma_{\mathfrak{a}}$	standard deviation $\varepsilon_{\mathfrak{a},t}$	0.0091
11	σ_z	standard deviation $\varepsilon_{z,t}$	0.12
12	$\sigma_{\widetilde{x}}$	standard deviation $\varepsilon_{\widetilde{x},t}$	6.01

Table 2: Simulated and Data Moments

Variable	Data Mean	Data St. Dev.	Model Mean	Model St. Dev.
Δc_t	0.018	0.014	0.023	0.018
Δy_t	0.021	0.021	0.023	0.053
Δd_t	0.019	0.120	0.022	0.123
r_t^e	0.064	0.169	0.066	0.169
r_t^f	0.015	0.022	0.015	0.000
r_t^{ex}	0.049	0.165	0.051	0.169
$p_t - d_t$	3.564	0.382	3.239	0.301

Notes: Δc_t is log differences of real consumption. Δd_t is log differences of dividends obtained from CRSP, where the dividend series is extracted as the difference between the cum-dividend and ex-dividend value-weighted returns. Data for r_t^e are the cum-dividend value-weighted returns from CRSP. Data for r_t^f are the constant-maturity 1-year rate on treasury bills from FRED. r_t^{ex} is the difference between the r_t^e and r_t^f. Data for $p_t - d_t$ is obtained from the CRSP data. The variables Δd_t, r_t^f, r_t^e, and r_t^{ex} are adjusted for inflation by subtracting the log difference of realized CPI (all urban consumers) obtained from FRED. All variables are at annual frequency. For Δc_t and Δy_t, annual observations are created by summing log differences at quarterly frequency over the year. For Δd_t, annual observations are obtained by summing log differences at monthly frequency over the year. For r_t^e, r_t^f, and r_t^{ex}, annual observations are obtained by summing log levels at monthly frequency over the year. For the ratio $p_t - d_t$, annual observations are created by first summing dividends in levels over the year to obtain an annual dividend in levels. d_t is then taken to be the log of the annual dividend. p_t is taken to be the last price observation of the year (at monthly frequency). Annual observations are defined over years ending in June so that the most recent data can be included. The sample is 1953:4 - 2012:6.

Table 3: Long Horizon Predictability Regressions: $Y_{t,t+h} = a + b(p_t - d_t) + e_{t,t+h}$

DATA

$Y_{t,t+h}$:	$\sum_{j=1}^{h} \Delta c_{t+j}$		$\sum_{j=1}^{h} \Delta d_{t+j}$		$\sum_{j=1}^{h} r_{t+j}^{ex}$	
h	pd_t	$\bar{R}^2$	pd_t	$\bar{R}^2$	pd_t	$\bar{R}^2$
4	-0.125	-0.003	-0.012	0.001	$\mathbf{-0.130}$	0.068
	(-0.248)		(-0.502)		(-2.856)	
8	-0.875	0.013	-0.034	0.010	$\mathbf{-0.226}$	0.120
	(-0.776)		(-0.780)		(-2.777)	
12	-1.423	0.022	-0.055	0.019	$\mathbf{-0.272}$	0.144
	(-0.849)		(-0.857)		(-3.043)	
20	-2.719	0.050	-0.041	0.004	$\mathbf{-0.361}$	0.188
	(-1.072)		(-0.448)		(-4.023)	

MODEL

$Y_{t,t+h}$:	$\sum_{j=1}^{h} \Delta c_{t+j}$		$\sum_{j=1}^{h} \Delta d_{t+j}$		$\sum_{j=1}^{h} r_{t+j}^{ex}$	
h	pd_t	$\bar{R}^2$	pd_t	$\bar{R}^2$	pd_t	$\bar{R}^2$
4	0.006	0.011	0.076	0.034	$\mathbf{-0.266}$	0.139
	(0.007)		(1.545)		(-3.600)	
8	0.028	0.024	0.144	0.067	$\mathbf{-0.369}$	0.195
	(0.026)		(1.699)		(-3.953)	
12	0.024	0.036	0.206	0.096	$\mathbf{-0.422}$	0.223
	(0.027)		(1.863)		(-4.097)	
20	0.029	0.054	**0.315**	0.144	$\mathbf{-0.464}$	0.244
	(0.035)		(2.146)		(-4.145)	

Notes: Regressions from actual and simulated data of the variable $Y_{t,t+h}$ on $p_t - d_t$ (the log price-dividend ratio at time t) and a constant: $Y_{t,t+h} = a + b(p_t - d_t) + e_{t,t+h}$. The variable $Y_{t,t+h}$ is alternately equal to h-quarter consumption growth $\sum_{j=1}^{h} \Delta c_{t+j}$, h-quarter dividend growth $\sum_{j=1}^{h} \Delta d_{t+j}$, or h-quarter excess stock market returns, $\sum_{j=1}^{h} r_{t+j}^{ex}$. The regression coefficient b is reported along with its t-statistic, obtained as averages over 1,000 simulated regressions of 238 observations each. The t-statistics are calculated using Newey-West standard errors, with number of lags equal to the regression horizon, and the R^2 statistic is adjusted for the number of explanatory variables. Coefficients that are statistically significant at 5% level appear in bold.

Table 4: Variance Decomposition of Quarterly Log Difference in Stock Wealth

	c Shock	y Shock	a Shock	Residual
Model	0.018	0.038	0.942	0.002
	(0.002, 0.052)	(0.006, 0.098)	(0.865, 0.985)	
Data (Stock Wealth)	0.062	0.044	0.759	0.134
	(0.042, 0.135)	(0.024, 0.101)	(0.690, 0.815)	
Data (Stock Price)	0.060	0.042	0.743	0.155
	(0.043, 0.129)	(0.025, 0.101)	(0.678, 0.803)	
Campbell-Cochrane Habit	0.86	—	0.14	0.000

Notes: See Table 2. This table reports a variance decomposition of the quarterly log difference in stock market wealth using the OLS regressions of the log difference in stock wealth on contemporaneous and lagged mutually orthogonal VAR innovations. The numbers reported represent the fraction of the $h = \infty$ step-ahead forecast error in the log difference of stock wealth that is attributable to the shock named in the column heading. The last column reports the share of the variance that is attributable to the residual of the regressions. Model results are calculated as averages over 1,000 simulations of 238 observations each. The numbers in parentheses represent the 5th and 95th percentiles of these statistics from bootstrapped samples using the procedure described in the Appendix. For model results, the numbers in parentheses are the 5th and 95th percentiles of the relevant statistic over the 1,000 simulations. The row labeled "Campbell-Cochrane Habit"gives the corresponding numbers in the benchmark specification of the Campbell and Cochrane (1999) model. The historical sample spans the period 1952:Q2 - 2012:Q4.

Table 5: Long Horizon Return Regressions: $\sum_{j=0}^{h} r_{t+j+1}^{ex} = \beta' X_t + \omega_{t+1,t+h}$

DATA

X_t:	pd		cay		$\mathcal{U}_3^C$		$\mathbf{e_a}$		pd^{orth}		cay_a	
h	pd_t	$\bar{R}^2$	cay_t	$\bar{R}^2$	$\mathcal{U}_{3,t}^C$	$\bar{R}^2$	$\mathbf{e_{a,t}}$	$\bar{R}^2$	pd_t^{orth}	$\bar{R}^2$	$cay_{a,t}$	$\bar{R}^2$
4	**−0.130**	0.068	**2.722**	0.074	−0.037	0.003	**−0.401**	0.137	−0.083	0.023	**7.357**	0.128
	(−2.856)		(2.936)		(−1.030)		(7.739)		(−1.571)		(3.684)	
8	**−0.226**	0.120	**4.864**	0.133	−0.018	-0.004	**−0.673**	0.262	−0.118	0.029	**13.736**	0.261
	(−2.777)		(3.326)		(−0.299)		(12.686)		(−1.363)		(5.575)	
12	**−0.272**	0.144	**6.667**	0.197	0.014	-0.005	**−0.678**	0.314	−0.135	0.031	**16.423**	0.310
	(−3.043)		(4.098)		(0.169)		(8.419)		(−1.276)		(7.052)	
20	**−0.361**	0.188	**9.042**	0.253	−0.032	-0.003	−0.536	0.179	**−0.216**	0.061	**15.142**	0.178
	(−4.023)		(4.721)		(−0.306)		(2.245)		(−2.179)		(3.376)	

MODEL

X_t:	pd		cay				$\mathbf{e_a}$		pd^{orth}		cay_a	
h	pd_t	$\bar{R}^2$	cay_t	$\bar{R}^2$			$\mathbf{e_{a,t}}$	$\bar{R}^2$	pd_t^{orth}	$\bar{R}^2$	$cay_{a,t}$	$\bar{R}^2$
4	**−0.266**	0.139	**0.693**	0.084			**−0.283**	0.273	−0.070	0.015	**1.459**	0.274
	(−3.600)		(3.593)				(13.363)		(−0.164)		(6.279)	
8	**−0.369**	0.195	**0.984**	0.119			**−0.399**	0.351	−0.111	0.026	**2.026**	0.358
	(−3.953)		(3.643)				(16.236)		(−0.252)		(7.185)	
12	**−0.422**	0.223	**1.114**	0.129			**−0.445**	0.372	−0.143	0.034	**2.277**	0.383
	(−4.097)		(3.526)				(17.649)		(−0.337)		(7.365)	
20	**−0.464**	0.244	**1.153**	0.122			**−0.449**	0.358	−0.173	0.049	**2.428**	0.369
	(−4.145)		(3.181)				(16.328)		(−0.416)		(6.837)	

Notes: Regressions from actual and simulated data of the variable $\sum_{j=0}^{h} r_{t+j+1}^{ex}$ on X_t and a constant: $\sum_{j=0}^{h} r_{t+j+1}^{ex} = \beta' X_t + \omega_{t+1,t+h}$. The variable X_t is alternately pd, cay, $\mathcal{U}_3^C$, e_a, pd^{orth}, and cay_a. pd is the log price-dividend ratio (source: CRSP). cay is the consumption-wealth variable from Lettau and Ludvigson (2001), while cay_a is the component of cay driven by shocks to $e_{a,t}$. $\mathcal{U}_{3,t}^C$ is the square root of the conditional expectation of the squared consumption innovation $\Delta c_{t+1} - E_t \Delta c_{t+1}$, one quarter ahead, computed using a stochastic volatility model in Jurado, Ludvigson, and Ng (2013). The VAR innovation $e_{a,t}$ is a shock that raises Δa_t, holding fixed Δc_t and Δy_t. The variable pd_t^{orth} is the fitted residual from a regression of pd on contemporaneous and 19 lagged values of $e_{a,t}$. For $\mathbf{e}_{a,t} = (e_{a,t}, ..., e_{a,t-19})$, the coefficient reported is the sum of the individual regression coefficients $\sum_{j=0}^{19} \beta_{e_a,j}$, where $\beta_{e_a,j}$ is the coefficient on $e_{a,t-j}$, and the statistic reported in parentheses is a Wald statistic for the null hypothesis that the squared coefficients sum to zero: $\sum_{j=0}^{23} \beta_{e_a,j}^2 = 0$. For the variables pd_t and cay_t, the statistics reported in parentheses are t-statistics for the null hypothesis that the regression coefficient is zero. A constant is included in each regression even though it is not reported in the table. Bolded coefficients indicate significance at the 5 percent or better level. Test statistics are corrected for serial autocorrelation and heteroskedasticity using a Newey-West estimator with 24 lags. $\bar{R}^2$ is the adjusted R^2 statistic. For the $\mathcal{U}_{3,t}^C$ regression, the sample spans 1960:Q4 : 2012:Q1. For all other regressions, the sample spans 1952:Q4 - 2012:Q3.

Figure 1: Unconditional Distribution of Risk Aversion

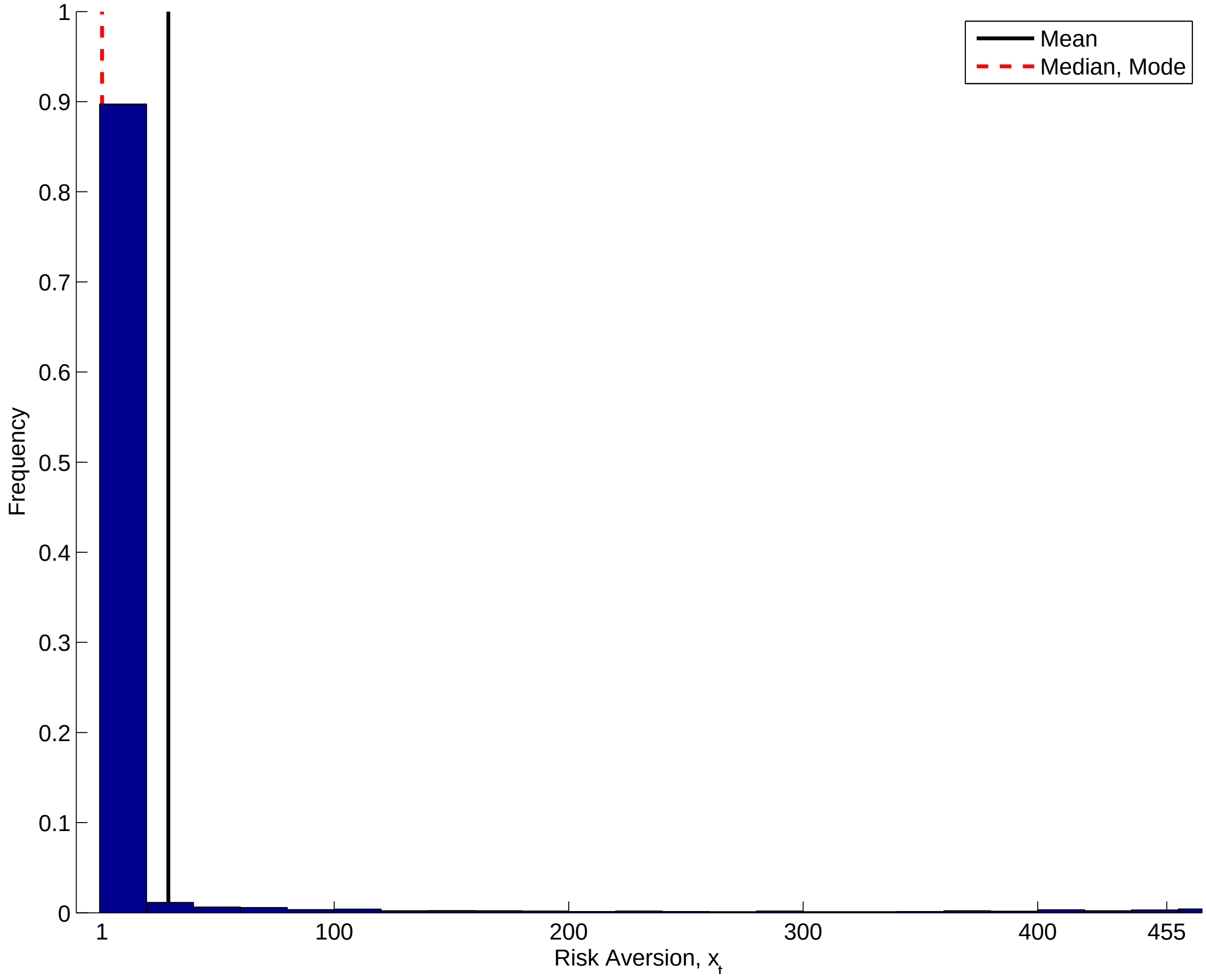

Notes: Each bar represents the frequency with which the risk aversion variable x_t falls in that specific bin. The median and mode are so close as to be indistinguishable on the plot.

Figure 2: Model-Based Responses: Primitive vs. VAR Shocks

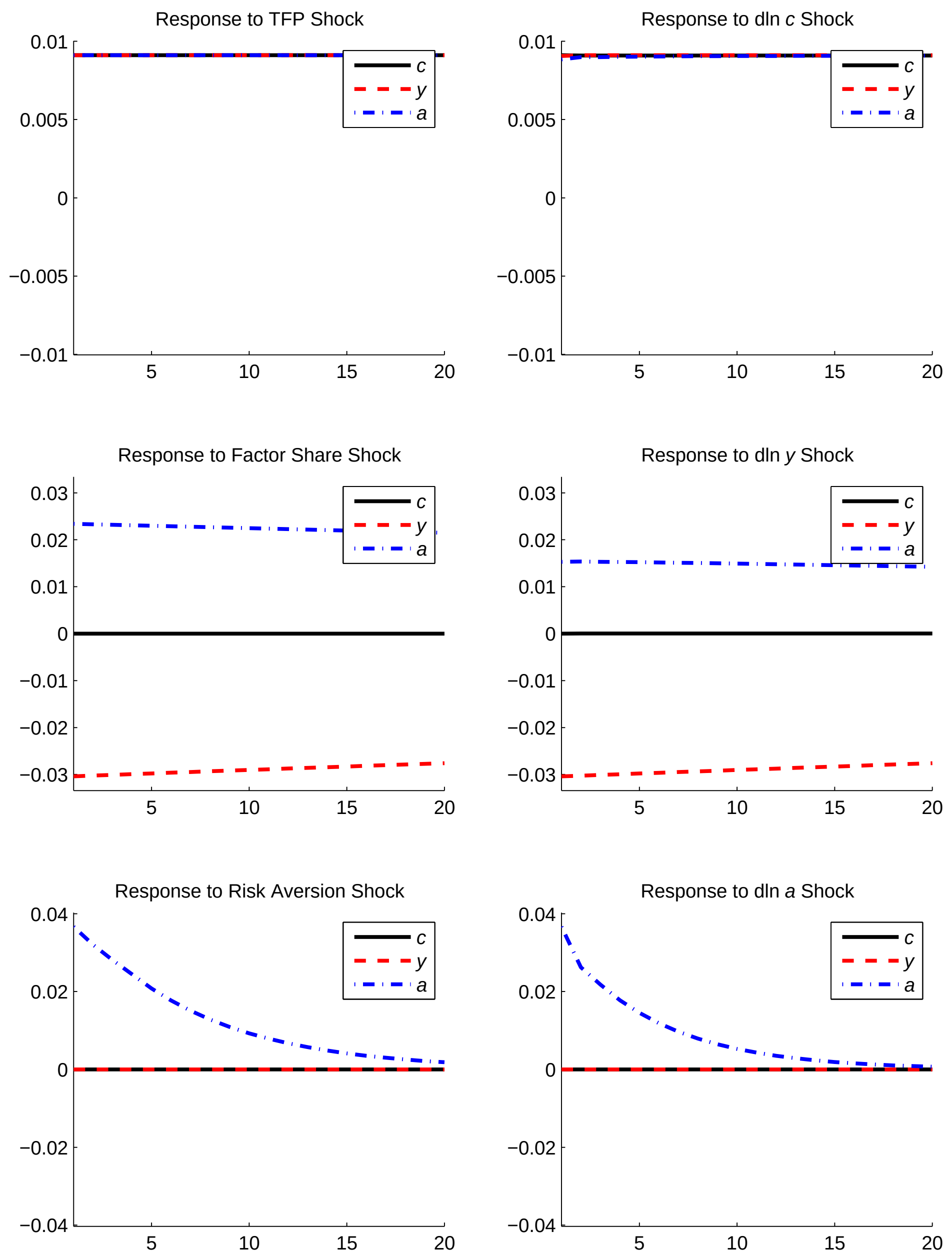

Notes: The left column plots the cummulated model impulse responses of the log differences of c, y, and a to the primitive shock named in the sub-graph title. The right column plots the cumulative impulse responses implied by the model from a VAR in the log differences of c, y, and a to the orthogonalized VAR shocks using data simulated from the model (right column). Impulse responses to primitive shocks are obtained by applying at $t = 0$ a one-standard deviation change in the direction that increases Δa_t, and simulating the model forward with all other shocks set to zero. Impulse responses to the VAR shocks are obtained by simulating the model over one very long sample (equal to 238,000 observations), estimating a cointegrated VAR in the log differenced data, inverting to Wold representation and computing the responses to orthogonalized c, y, and a shocks equal to one standard deviation changes in the direction that increases Δa_t with that ordering in the VAR. The size of the shocks are normalized so that the initial response of a variable to its own shock in the right panel is the same as the response of that variable to the corresponding primitive shock in the left panel.

Figure 3: VAR Impulse Responses (Model vs. Data)

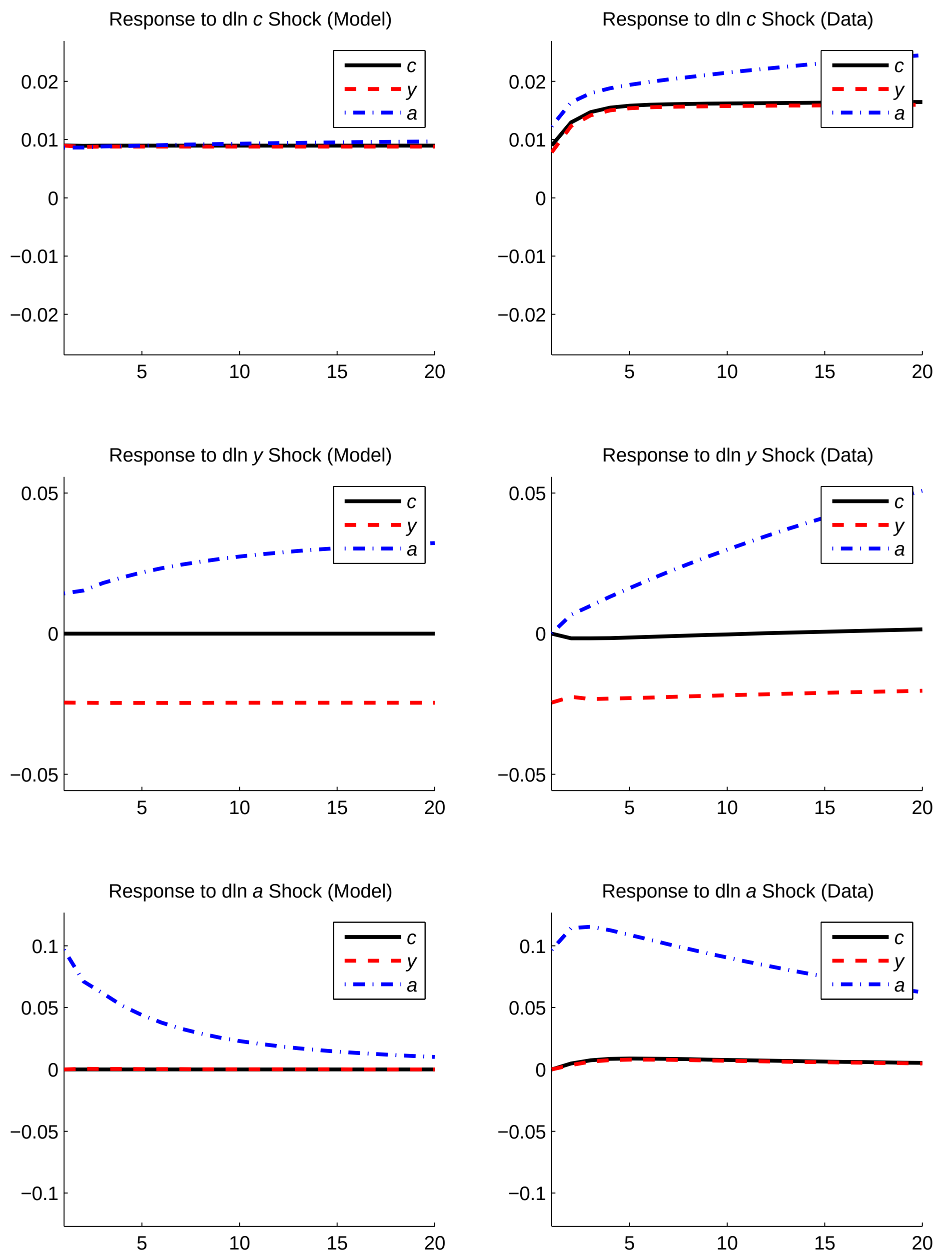

Notes: The figure plots impulse response functions to the VAR shocks in both the model (left column) and data (right column). true structural shocks to the orthogonalized shocks obtained from the VECM regression using data simulated from the model (left column), and actual data (right column). In both cases, impulse responses are obtained by estimating a cointegrated VAR in the log differenced data, inverting to Wold representation and computing the responses to orthogonalized c, y, and a shocks equal to one standard deviation changes in the direction that increases Δa_t with that ordering in the VAR. Model results are calculated as averages over 1,000 simulations of 238 observations each. The historical sample spans the period 1952:Q2 - 2012:Q4.

Figure 4: Stock Market Impulse Responses

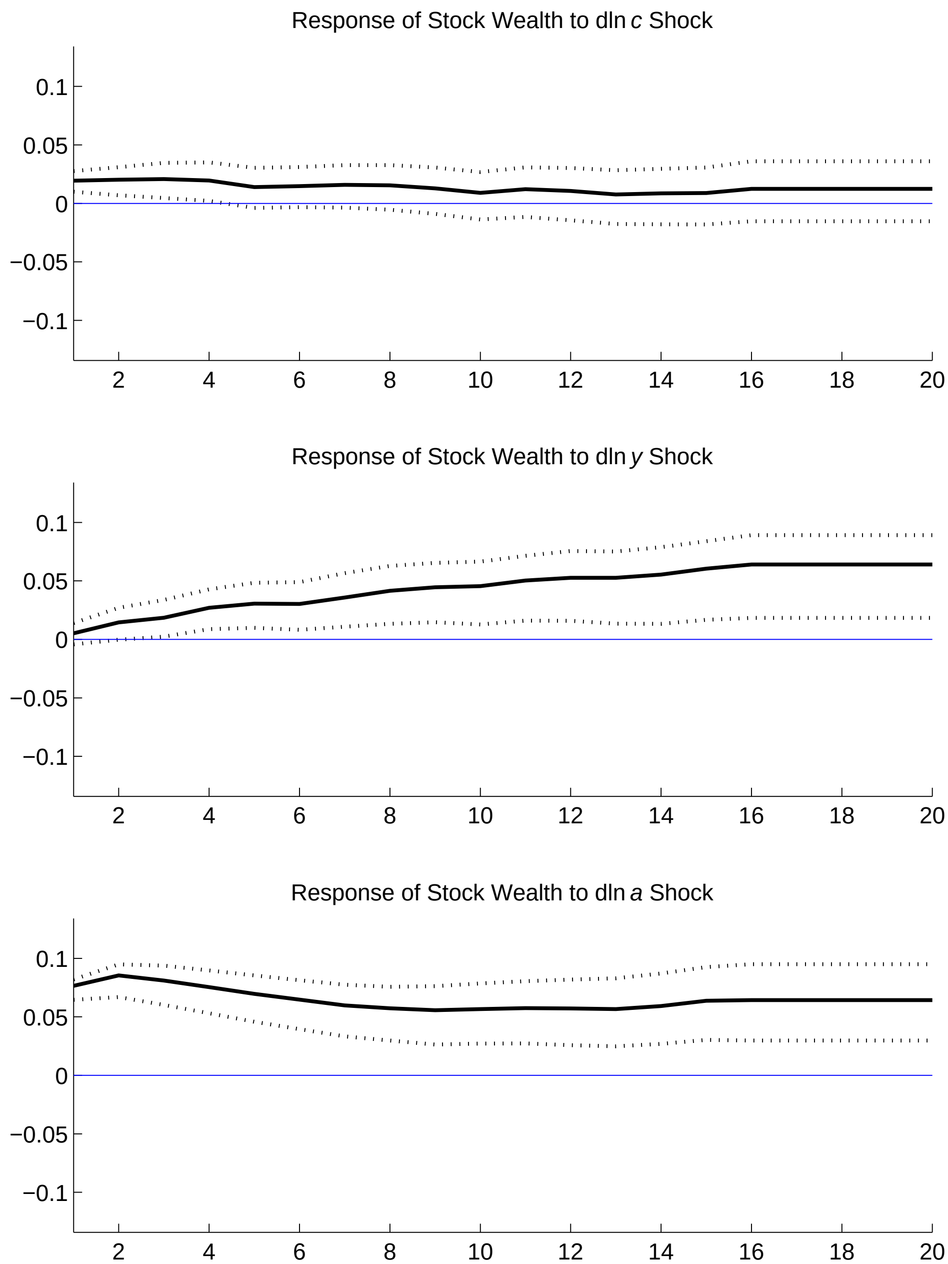

Notes: The figure plots impulse responses of stock wealth to the shocks obtained from the VECM regression. Dotted lines are 90% error bands obtained using the bootstrap procedure described in the Appendix. The historical sample spans the period 1952:Q2 - 2012:Q4.

Figure 5: Decomposition of Spectrum (Model vs. Data)

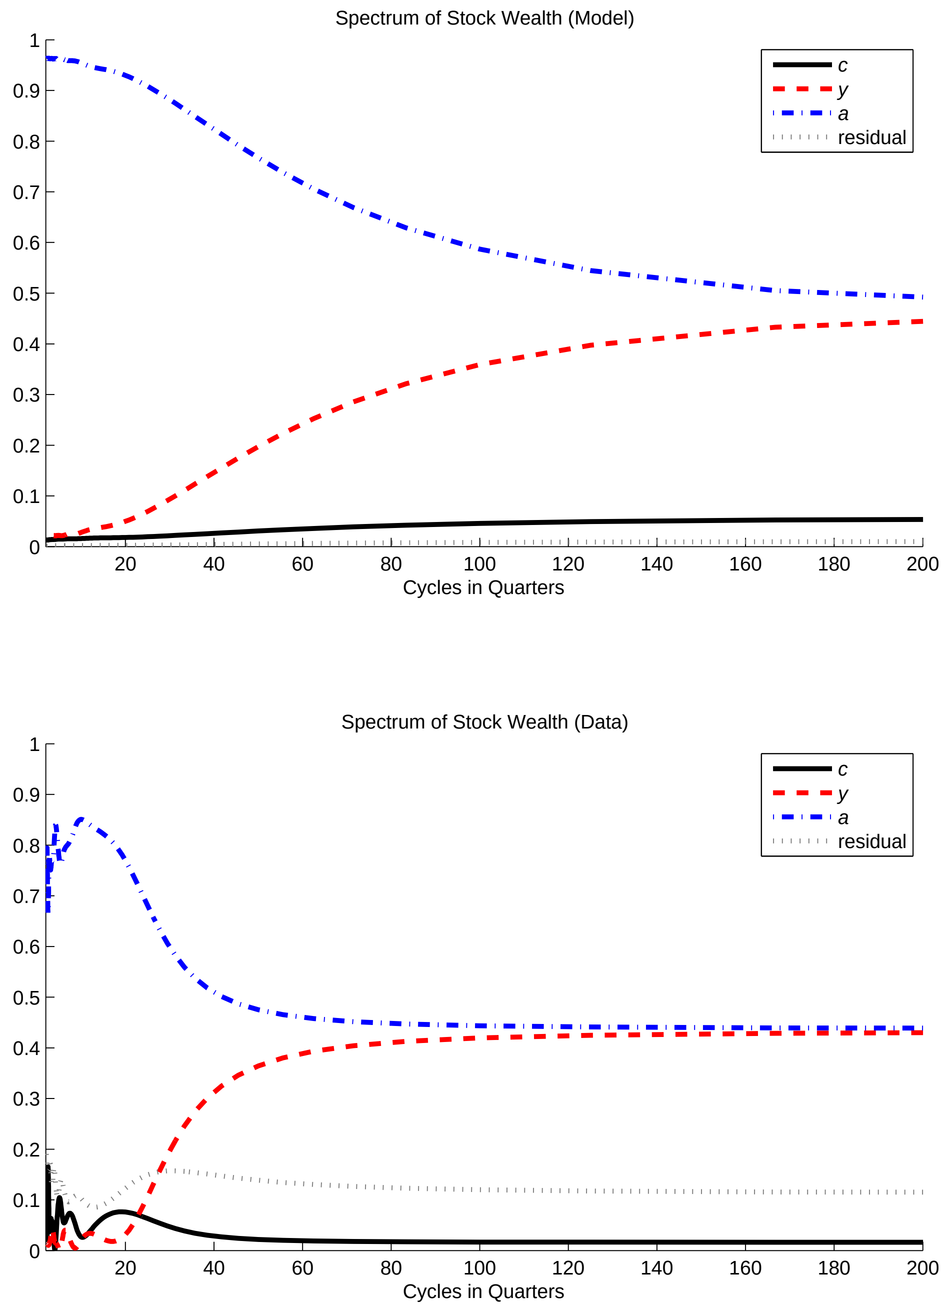

Notes: The figure shows the decomposition of spectra at different frequencies into components driven by each of the orthogonalized shocks of the consumption, labor income and wealth VAR. Results for the model are computed from averages over 1,000 simulations of 238 observations. Results from historical data appear in the two bottom panels. The historical sample spans the period 1952:Q2 - 2012:Q4.

Figure 6: Level Decomposition (Data)

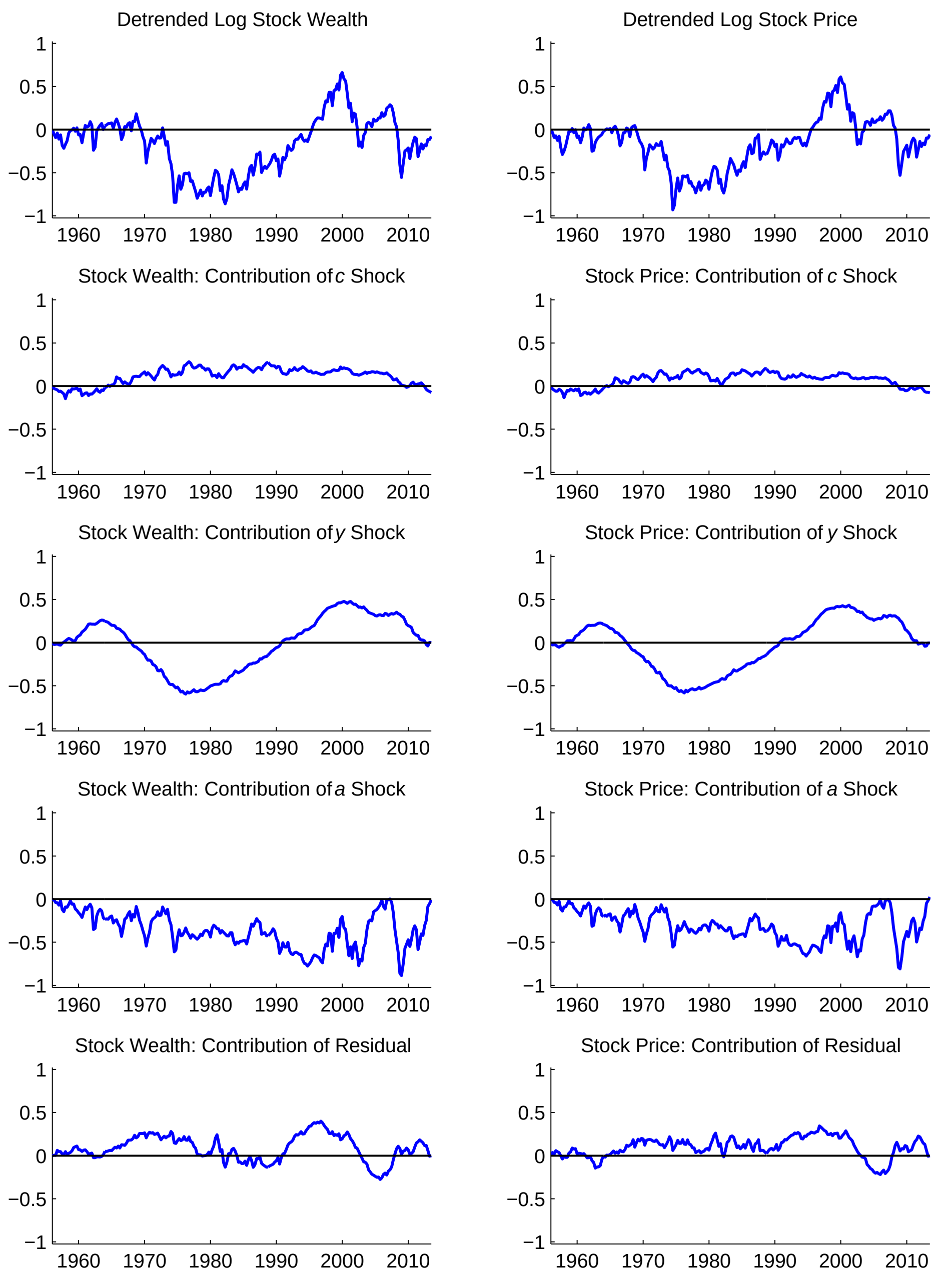

Notes: See Table 2. The figure shows the decomposition of the log level of stock wealth into components driven by the orthogonalized c, y, and a shocks obtained from the VECM regression. A deterministic trend is removed from the log level of stock wealth by removing the mean in log differences before cumulating. The components plus the residual sum to the log level of detrended stock wealth and detrended stock price in the left and right panels, respectively. Each component and the sum are normalized so that the value in 1952:Q1 is zero. The sample spans the period 1952:Q2 - 2012:Q4.

Figure 7: Decomposition of Labor Income and Stock Market Wealth

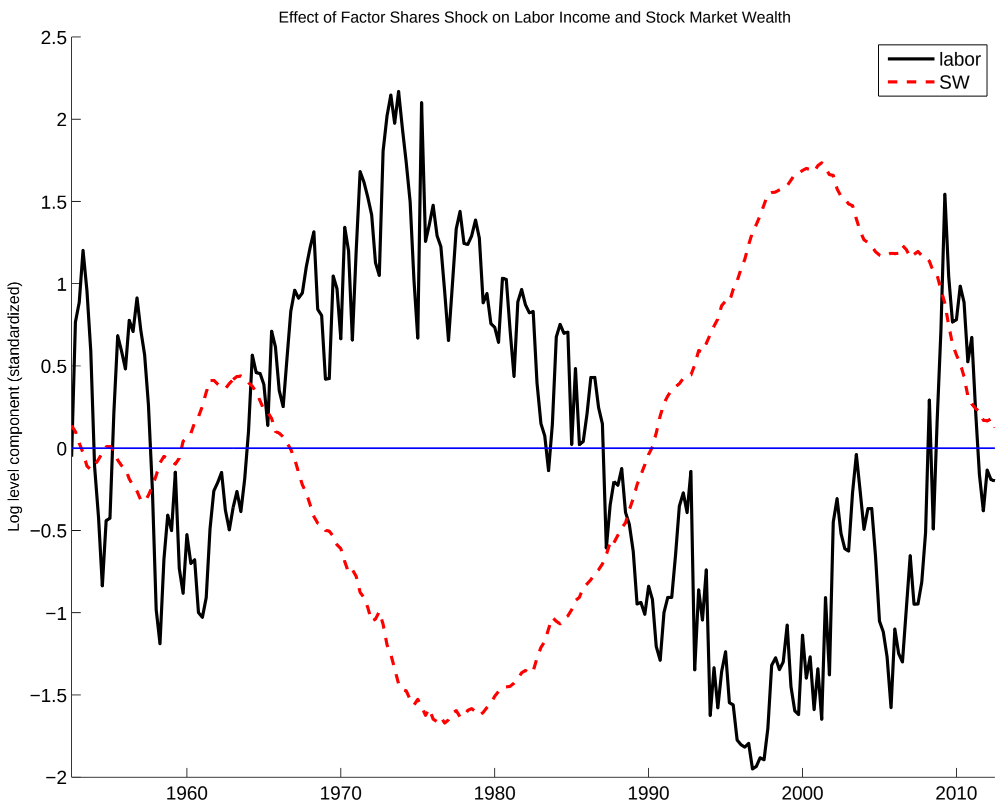

Notes: The figure shows the component of the log levels of stock market wealth and labor income that is attributable to the factor shares shock over time. The effect of the factors share shock $e_{y,t}$ on the log level of each series is obtained by summing up the estimated effects of $e_{y,t}$ on the log differences over time. Both series are demeaned and divided by their standard deviations. The sample spans the period 1952:Q2 - 2012:Q4.

Figure 8: Earnings Impulse Responses

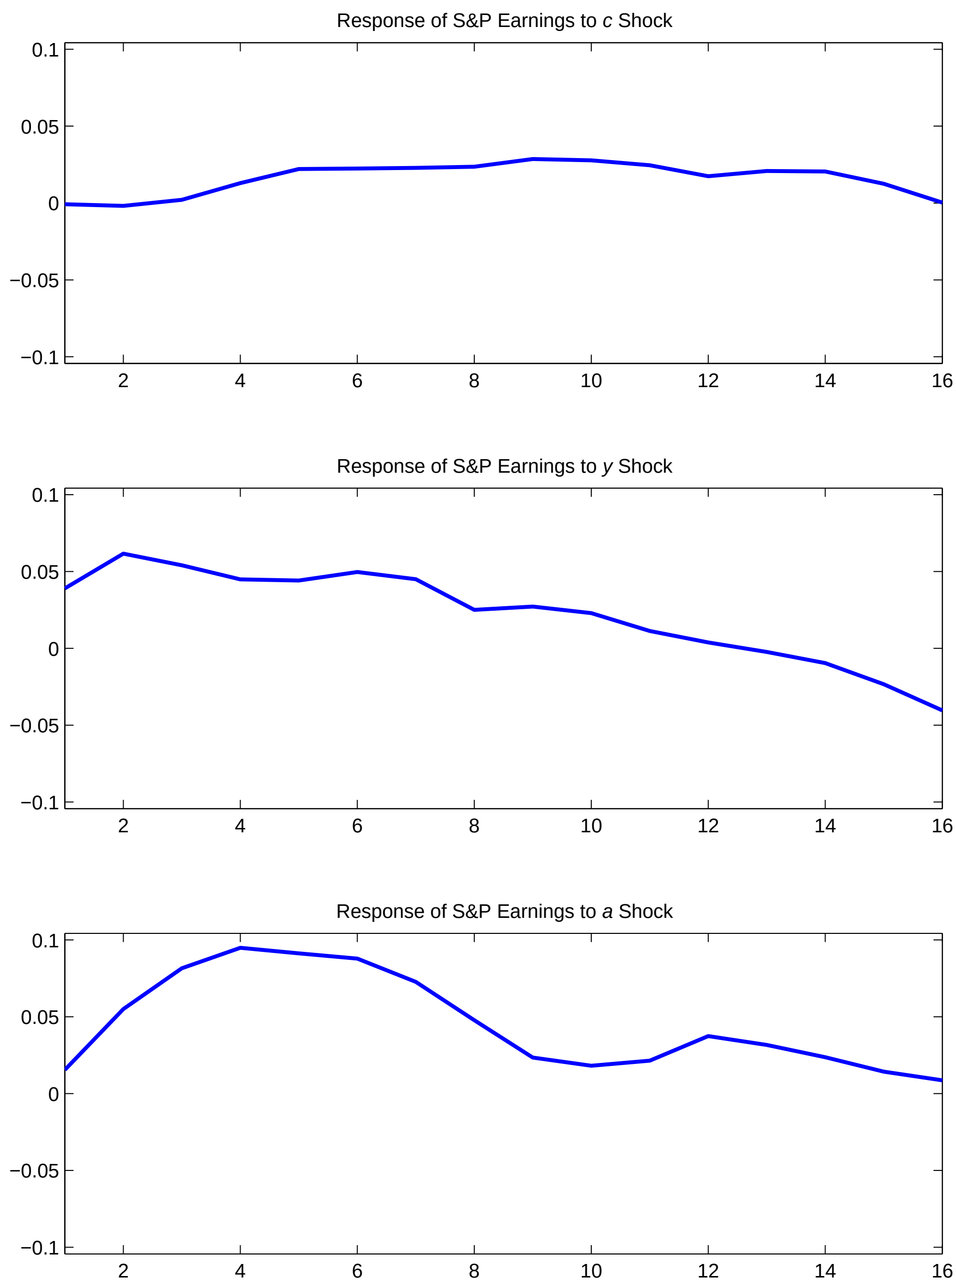

Notes: The figure plots the responses of 16Q real per-capita S&P 500 earnings growth to contemporaneous and lagged values of the shocks obtained from the VECM regression. The IRFs are computed as the coefficients A, B, and C from the earnings regression

$$\sum_{j=0}^{15} \Delta \ln(EA_{t-15+j}) = \alpha_0 + A(L)e_{c,t} + B(L)e_{y,t} + C(L)e_{a,t} + \nu_t.$$

Each lag polynomial includes sixteen terms total (one contemporaneous and fifteen lag terms). EA_t is the four quarter moving average of real reported S&P earnings obtained from the website of Robert Shiller. The sample spans the period 1952:Q2 - 2012:Q4.

Figure 9: Gini Coefficient and Cumulated e_y Shock

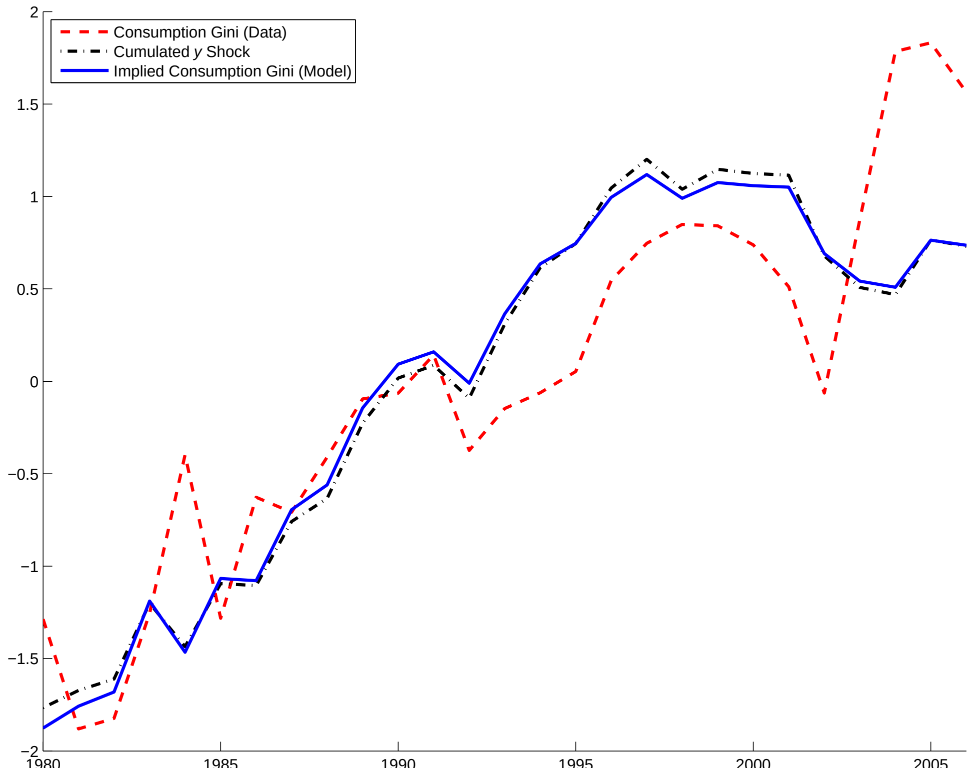

Notes: The consumption Gini (data) series is the Gini coefficient of inequality for nondurable consumption (source: Heathcote, Perri and Violante (2010)). The cumulated y shock series is the running total of all the e_y shocks to date: $\sum_{j=1}^{t} e_j$. The implied consumption Gini (model) series uses the model to calculate the implied Gini coefficient for consumption based on the observed sequence of y shocks (see appendix for details). All series are presented at annual frequency. For the cumulated y shock and model-implied Gini series, annual observations are averages over the calendar year. All three series are normalized to have zero mean and unit standard deviation in the sample.

A Additional Figures

Figure A.1: Consumption Shocks and TFP

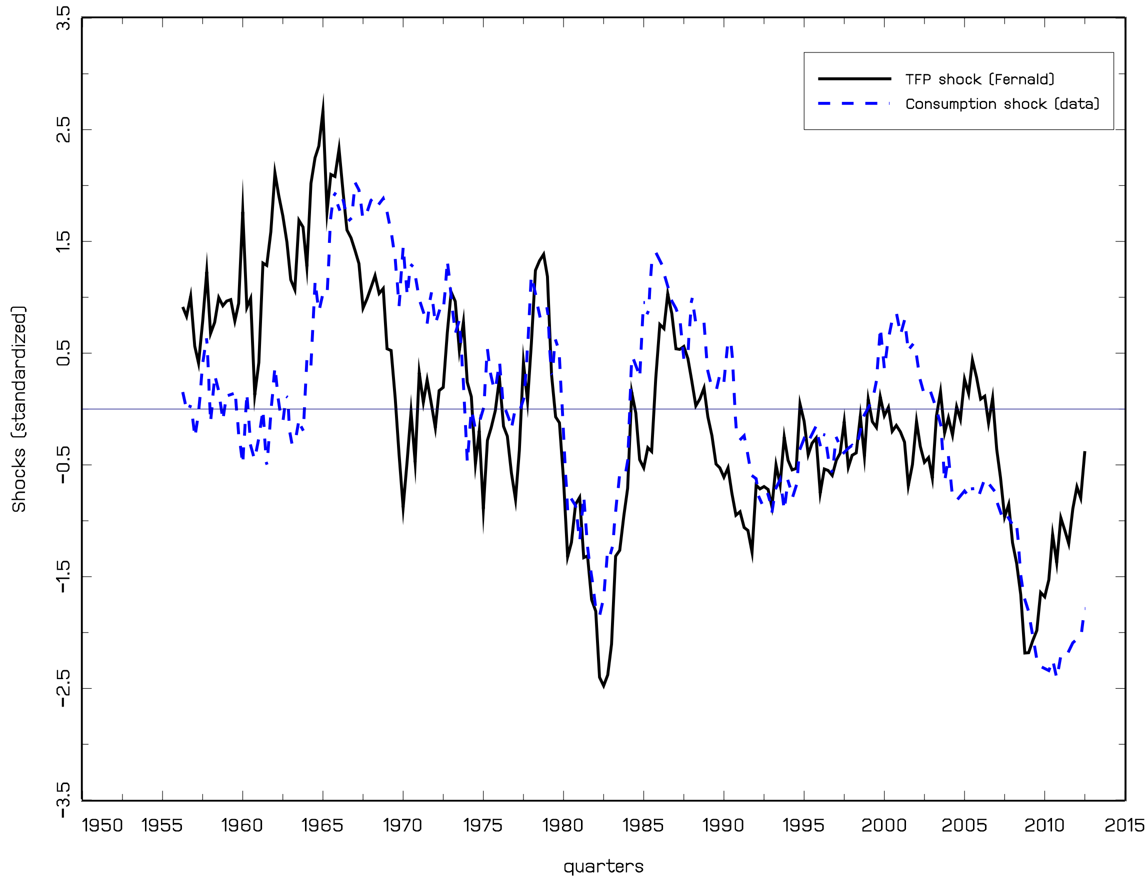

Notes: The TFP shock series is differenced Business Sector TFP (source: Fernald). The consumption shock series is taken from the VECM. The sample spans the period 1947:Q2 - 2013:Q3.

Figure A.2: 16Q Moving Averages of Labor Share and TFP

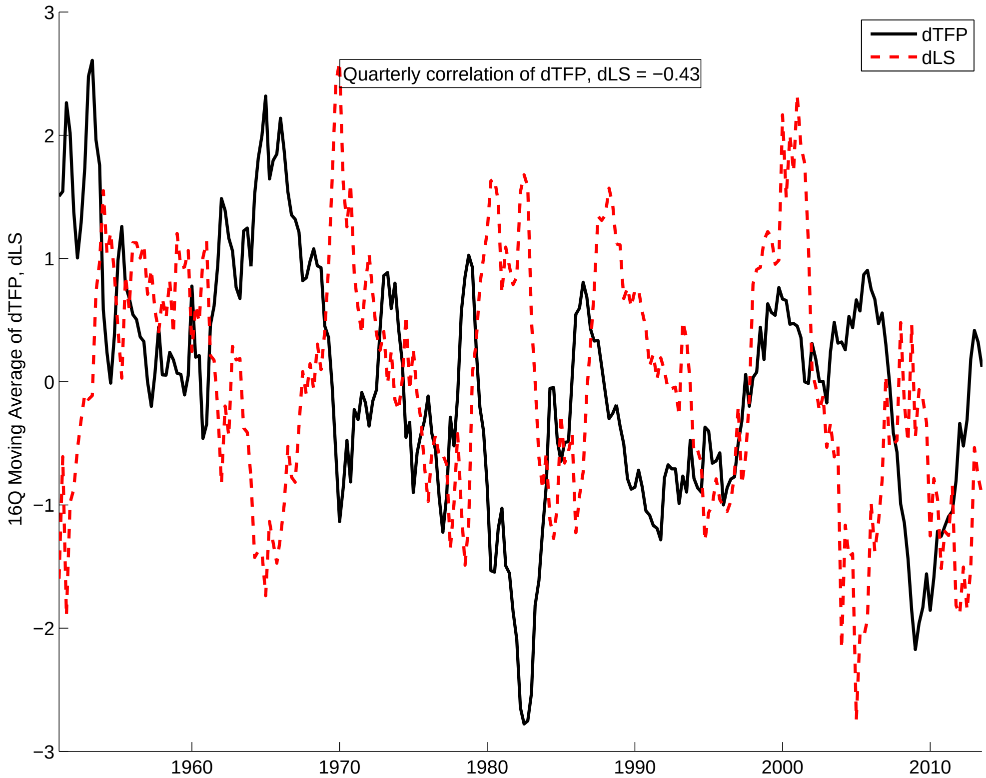

Notes: The series are 16Q moving averages of each series in log differences. TFP is Business Sector TFP (source: Fernald). LS is Nonfarm Business Sector Labor Share (source: BLS). The sample spans the period 1947:Q2 - 2013:Q3.

Figure A.3: VAR Impulse Responses with Error Bands

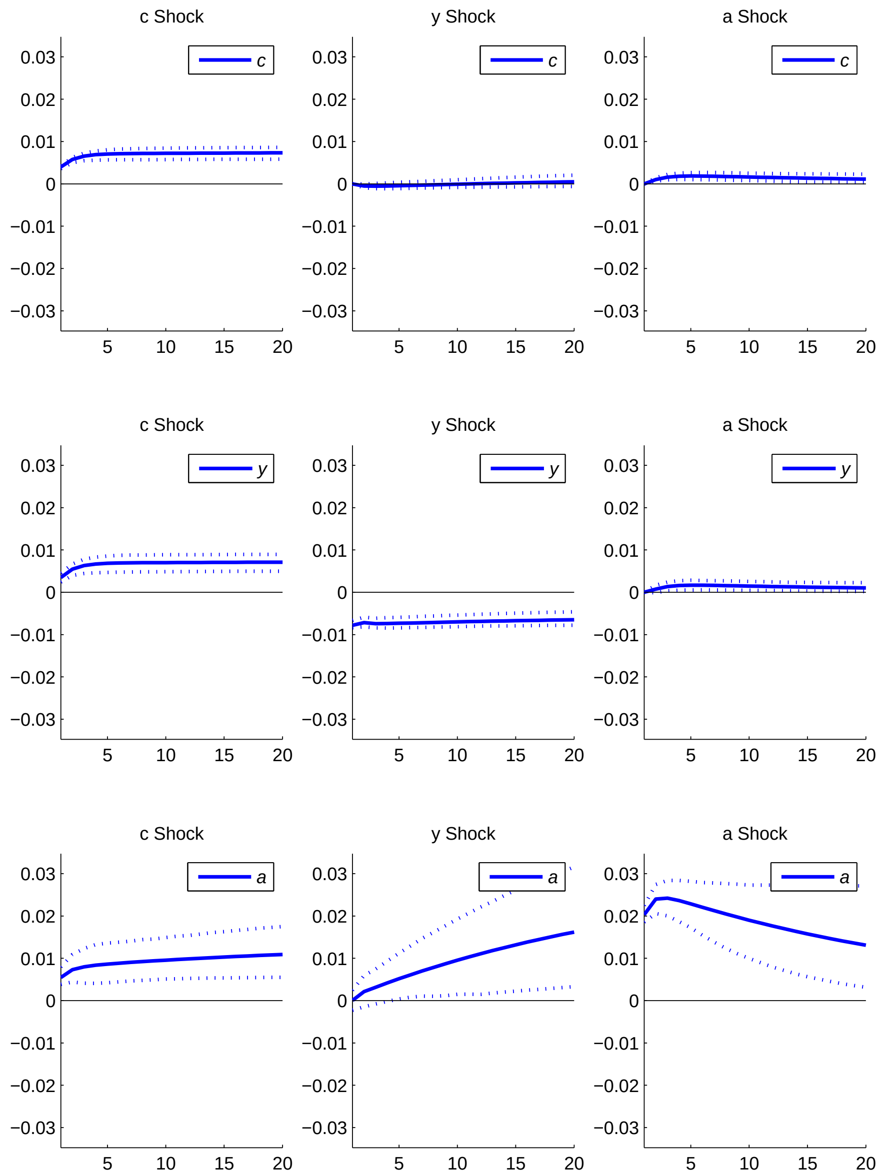

Notes: The figure plots impulse response functions to the VAR shocks obtained from the VECM regression using data. Impulse responses are obtained by estimating a cointegrated VAR, inverting to Wold representation and computing the responses to orthogonalized c, y, and a shocks with that ordering in the VAR. The dotted lines are 90% error bands obtained using the bootstrap procedure described in the Appendix. The historical sample spans the period 1952:Q2 - 2012:Q4.

Figure A.4: Model Gini Coefficient

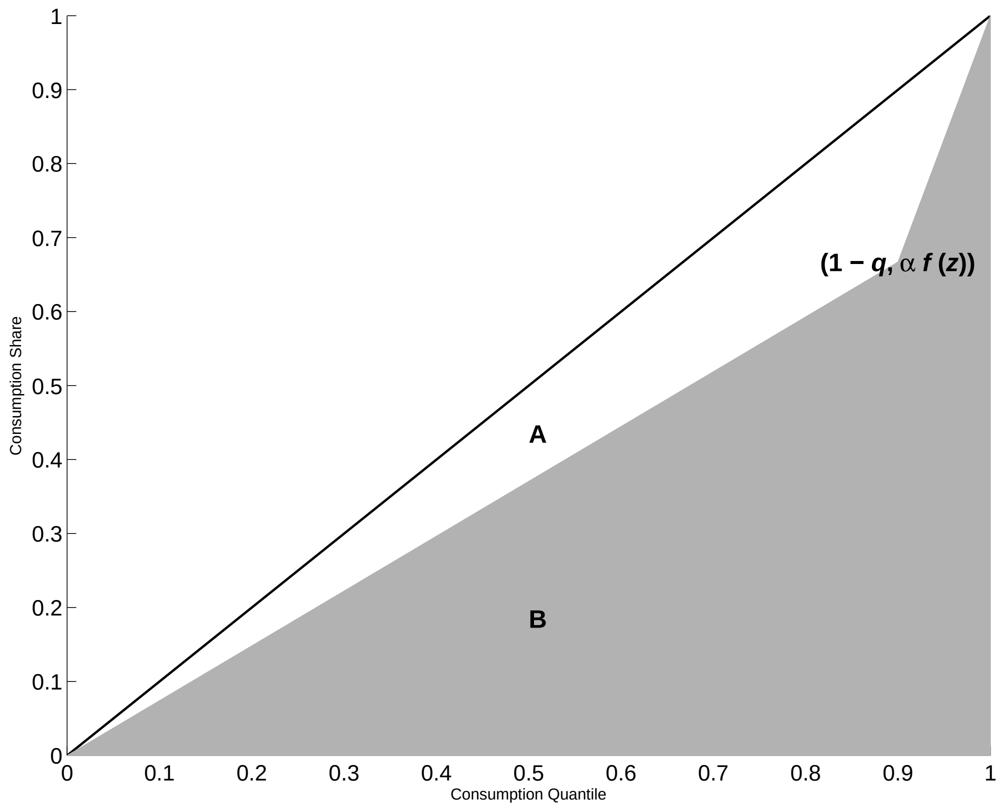

Notes: This diagram plots the distribution of consumption in the model. q is the proportion of shareholders. The Gini coefficient of consumption is defined by $G = A/(A + B)$, where A and B are the areas of the relevant regions.