

NBER WORKING PAPER SERIES

ORIGINS OF STOCK MARKET FLUCTUATIONS

Daniel L. Greenwald
Martin Lettau
Sydney C. Ludvigson

Working Paper 19818
<http://www.nber.org/papers/w19818>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2014

The authors are grateful to Jarda Borovicka, John Cochrane, Brian Kelly, Eric Swanson, and seminar participants at the Q Group 2013 fall conference, the Macro-Finance Society Fall 2014 Meeting, NYU, University of Chicago Booth School of Business, UC Berkeley Haas School, the University of Michigan, and the University of Texas at Austin, for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Daniel L. Greenwald, Martin Lettau, and Sydney C. Ludvigson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Origins of Stock Market Fluctuations

Daniel L. Greenwald, Martin Lettau, and Sydney C. Ludvigson

NBER Working Paper No. 19818

January 2014, Revised December 2014

JEL No. G0,G12

ABSTRACT

Three mutually uncorrelated economic disturbances that we measure empirically explain 85% of the quarterly variation in real stock market wealth since 1952. A model is employed to interpret these disturbances in terms of three latent primitive shocks. In the short run, shocks that affect the willingness to bear risk independently of macroeconomic fundamentals explain most of the variation in the market. In the long run, the market is profoundly affected by shocks that reallocate the rewards of a given level of production between workers and shareholders. Productivity shocks play a small role in historical stock market fluctuations at all horizons.

Daniel L. Greenwald
Department of Economics
New York University
19 West 4th Street, 6th floor
New York, NY 10012
dlg340@nyu.edu

Sydney C. Ludvigson
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10002
and NBER
sydney.ludvigson@nyu.edu

Martin Lettau
Haas School of Business
University of California, Berkeley
545 Student Services Bldg. #1900
Berkeley, CA 94720-1900
and NBER
lettau@haas.berkeley.edu

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 an important line of study because variation in the stock market price-dividend ratio is driven almost entirely by expected excess return fluctuation (i.e., forecastable movements in the equity premium).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 ex-post returns. 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* return premia) fluctuations originate.

To illustrate why, note that some economic shocks have tiny innovations but permanent or near-permanent effects on cash flows. Under rational expectations, permanent cash flow shocks have no influence on expected returns or the price-dividend *ratio*, but they can have a dramatic influence on real stock market wealth as the decades accumulate. 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 transitory expected return shocks and both permanent and transitory 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 disturbances that are associated with the vast majority (over 85%) of fluctuations in real quarterly 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 recursive orthogonalization procedure.

¹Expected 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)).

We then address the question of what these observable VAR disturbances represent economically. To do so, we provide a theoretical framework with two types of agents, shareholders and workers, and three primitive shocks. We show that, if the model generated the data, the observable VAR innovations under a specific recursive ordering would effectively recover three latent primitive shocks. Specifically, the *consumption* innovation in the empirical VAR would recover a total factor *productivity* (TFP) shock, the *labor income* innovation would recover a *factors share* shock that reallocates the rewards of production without affecting the size of rewards, and the *wealth* innovation would recover a shock to shareholder *risk aversion* that moves the stochastic discount factor pricing assets independently of stock market fundamentals or real activity such as consumption and labor earnings. We show that the dynamic responses to these mutually orthogonal VAR innovations produced from model generated data are remarkably similar to those obtained from historical data.

With this theoretical interpretation of the observable disturbances in hand, we turn to the question of how these distinct shocks have 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 and explain 75% of variation in the log difference of stock market wealth on a quarterly basis. 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. Although the factors share shock explains virtually none of the variation in the real level of the stock market over cycles of a quarter or two, 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, consumption shocks, both in the model and in the data, play a small role in the stochastic fluctuations of the stock market at all horizons. The crucial aspect of the model that makes it consistent with this finding is its heterogeneous agent specification. This finding contradicts representative agent asset pricing models in which

shocks that drive aggregate consumption play 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 was true from the mid 1960s to mid 1980s.) After removing a deterministic trend, 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 detrended stock market wealth over this period. Indeed, without these shocks, today's stock market would be roughly 10% lower than it was in 1980. An additional 38% of the increase since 1980, or a rise of 22%, is attributable to the cumulative effects of risk aversion shocks, while the cumulative effects of TFP shocks have made a negative contribution, a direct consequence of the large negative draws for consumption in the Great Recession. 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.)

Our findings are also informative about the origins of risk premia fluctuations. Shareholders in the model are 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 that generate a “flight to safety.” Even though these flights are rare and extreme, a time-varying expectation that risk tolerance could crash in the future generates plausible variation in the price-dividend ratio and empirically reasonable predictability in excess stock market returns. Time-variation in the risk premium, both in the model and the data, is revealed by the wealth shocks, which are orthogonal to movements in consumption and labor income. We find that these innovations also bear little relation to other traditional macroeconomic fundamentals such as dividends, earnings, consumption volatility, or broad-based macroeconomic uncertainty, and none of these other 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.

These findings have important implications for macroeconomic modeling. The two big empirical sources of variation that we find here are responsible for almost all of stock market variation play virtually no role in contemporary macroeconomic theories.

The rest of this paper is organized as follows. The next section discusses related literature. Section 3 describes the econometric procedure and data. Section 4 describes the theoretical model. Section 5 presents our findings. Section 6 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)). An important earlier literature 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)).

The factors share element of our paper is related to a separate macroeconomic literature that examines the long-run variation in the labor share (e.g., Karabarbounis and Neiman

(2013), and the recent theoretical study of Lansing (2014)). The factors share findings in this paper echo those from previous studies that use very different methodologies but find that returns to human capital are negatively correlated with those to stock market wealth (Lustig and Van Nieuwerburgh (2008); Lettau and Ludvigson (2009); Chen, Favilukis, and Ludvigson (2014))). Lettau, Ludvigson, and Ma (2014) focus on cross-sectional asset pricing and find that value and momentum strategies exhibit strong opposite signed exposure to low frequency fluctuations in the capital share, helping to explain why both strategies earn high average returns but are negatively correlated.

Our findings on the forecastability of excess stock market returns imply that quantitatively large component of risk premia fluctuations is acyclical, contrasting with classic earlier studies that emphasized the countercyclicality of risk premia (e.g., Fama and French (1989)). But our findings on risk premia variation are potentially consistent with other theories in which time-variation in the reward for bearing risk is largely divorced from fluctuations in macroeconomic fundamentals. Examples include the ambiguity aversion framework of Bianchi, Ilut, and Schneider (2013), models of rare events in which the probability of consumption disaster is a random variable independent of normal-times consumption shocks (e.g., Gourio (2012); Wachter (2013)), intermediary-based models in which intermediaries' risk-bearing capacity varies independently of, or in a highly nonlinear way with, macroeconomic state variables (e.g., Brunnermeier and Sannikov (2012); Gabaix and Maggiori (2013)); He and Krishnamurthy (2013); Muir (2014)). Each of these papers has a mechanism for generating acyclical fluctuations in risk premia that plays the same role as our independent risk aversion shock. The econometric evidence presented here contributes to a growing body that forms the basis of an empirical rationale for such mechanisms.

3 Econometric Analysis: Three Mutually Orthogonal Shocks

We study 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 single 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) and Lettau and Ludvigson (2010)).

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. Denote the log of real 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 also study stock market variation using as a measure of stock market wealth 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 . Our data are quarterly and span the first quarter of 1952 to the third quarter of 2012.

The model developed below is intended to focus on the implications of the empirical shocks for stock market wealth. As such, it has just one form of risky capital (equity) and a risk-free bond in zero net supply. Thus in the model, all wealth is stock market wealth, which is identically equal to total wealth and net worth: $a_t = s_t$. A question arises as to how best to connect the empirical wealth innovations in the data (which include nonstock forms of wealth) to those that arise from the model. Our approach is to construct the empirical VAR innovations using a system of variables that contains c_t , y_t and total asset wealth a_t , and then subsequently relate these innovations to stock market wealth. We do *not* construct the empirical innovations by restricting analysis to how consumption and labor income move only with stock market wealth. We do this for two reasons. First, we wish to allow for the possibility that a standard factor neutral productivity shock could affect the value of the stock market. But such a shock should affect the value of all productive capital, not just corporate equity, so a system that identifies such a shock from the data must 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 remaining shocks. Second, consumption and labor income are cointegrated with total wealth, as expected from theory (Lettau and Ludvigson

(2001)), but 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, which requires estimating a cointegrated VAR for $(\Delta c_t, \Delta a_t, \Delta y_t)$. With these results in hand, the question of how closely related the identified VAR shocks are to stock market wealth is then an empirical matter, which is the subject of an extensive investigation below. Although Δs_t may be related to these shocks, there will be an unexplained residual that in principal could be quite important.

The three mutually orthogonal empirical disturbances are obtained from cointegrated VAR (or vector error correction mechanism–VECM) representation of \mathbf{x}_t taking the form

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

where $\Delta \mathbf{x}_t$ is the vector of log first differences, $(\Delta c_t, \Delta a_t, \Delta y_t)'$, \mathbf{v} , and $\boldsymbol{\gamma} \equiv (\gamma_c, \gamma_a, \gamma_y)'$ are (3×1) vectors, $\mathbf{\Gamma}(L)$ is a finite order distributed lag operator, and $\hat{\boldsymbol{\alpha}} \equiv (1, -\hat{\alpha}_a, -\hat{\alpha}_y)'$ is the (3×1) vector of previously estimated cointegrating coefficients.² The term $\hat{\boldsymbol{\alpha}}' \mathbf{x}_{t-1}$ gives last period’s equilibrium error, or cointegrating residual, a variable we denote with $cay_t \equiv \hat{\boldsymbol{\alpha}}' \mathbf{x}_{t-1}$. Throughout this paper, we use “hats” to denote the estimated values of parameters.

The results of estimating a first-order specification of (1) are presented in Lettau and Ludvigson (2013), not reported here to conserve space. 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 $\hat{\boldsymbol{\alpha}}' \mathbf{x}_{t-1}$. 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. 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.

The individual series involved in the cointegrating relation can be represented as a

²Standard 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)). We estimate $\hat{\boldsymbol{\alpha}} = (1, -0.18, -0.70)'$.

reduced-form multivariate Wold representation:

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

where \mathbf{u}_t is an $n \times 1$ vector of innovations, and where $\boldsymbol{\Omega}(L) \equiv \mathbf{I} + \boldsymbol{\Omega}_1 L + \boldsymbol{\Omega}_2 L^2 + \boldsymbol{\Omega}_3 L^3 + \dots$. 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 specify 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 $\text{Cov}(\mathbf{u}_t)$, and define

$$\mathbf{e} \equiv \mathbf{H}^{-1}\mathbf{u}_t, \quad \mathbf{C}(L) \equiv \boldsymbol{\Omega}(L)\mathbf{H}.$$

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

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

which is now a function of a vector of mutually uncorrelated innovations \mathbf{e}_t . 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. We study a particular orthogonalization by restricting the ordering of the variables as follows: Δc is first, Δy second, and Δa last in $\Delta \mathbf{x}_t$. These empirical disturbances are (i) *consumption shocks* $e_{c,t}$: unforecastable movements in Δc_t that may contemporaneously affect Δy and Δa (ii) *labor income shocks* $e_{y,t}$: unforecastable movements in Δy_t holding fixed Δc_t contemporaneously, (iii) *wealth shocks* $e_{a,t}$: unforecastable movements in Δa_t holding fixed both Δc_t and Δy_t contemporaneously.

We refer to the mutually orthogonal \mathbf{e}_t shocks as “structural” disturbances. The ordering of variables determines the specific orthogonalization, however, so providing a theoretical interpretation of these disturbances requires a theoretical rationale for the ordering, as well as a mapping between each empirical disturbance and a set of primitive economic shocks implied by theory. We use the model of the next section to provide such a theoretical rationale. We show there that, if the proposed model were true, the particular ordering chosen would be the right one for uncovering the three primitive shocks of the model.

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

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

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 unexplained by the mutually orthogonal empirical disturbances $e_{c,t}$, $e_{y,t}$, and $e_{a,t}$. We estimate the same type of relationship for the CRSP value-weighted stock price index, replacing s_t with p_t on the left-hand-side. 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.

We also decompose the log *levels* of stock market wealth into components driven by each structural disturbance. To do so, rewrite the decomposition of growth rates as

$$\begin{aligned} \Delta s_t &= \kappa_0 + \kappa_c(L) e_{c,t} + \kappa_y(L) e_{y,t} + \kappa_a(L) e_{a,t} + \eta_t \\ &\equiv \kappa_0 + \Delta s_t^c + \Delta s_t^y + \Delta s_t^a + \eta_t, \end{aligned} \quad (5)$$

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:

$$\begin{aligned} s_t &= s_0 + \kappa_0 t + \sum_{k=1}^t \Delta s_k \\ &= s_0 + \kappa_0 t + \sum_{k=1}^t \Delta s_k^c + \sum_{k=1}^t \Delta s_k^y + \sum_{k=1}^t \Delta s_k^a + \sum_{k=1}^t \eta_k \\ &\equiv s_0 + \kappa_0 t + s_t^c + s_t^y + s_t^a + \sum_{k=1}^t \eta_k, \end{aligned} \quad (6)$$

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 term $\sum_{k=1}^t \eta_k$ is the component of s_t attributable to the unexplained residual. Note that $\kappa_0 t$ is the deterministic trend in stock market wealth, which in the model below is attributable to steady state technological progress. Expressions analogous to (5) and (6) are also computed for the log stock price index p_t .

This completes our description of the empirical disturbances. Before turning to the theoretical model used to interpret these disturbances, Table 1 shows why they are good ones to consider for investigating the origins of stock market fluctuations. Taken together, these three observable disturbances explain 87% of the quarterly growth in stock market wealth and 84% of quarterly stock price growth. Thus they account for almost all of the variation. As we show below, they also explain the vast bulk of fluctuations in stock market wealth at lower frequencies, but the relative importance of the shocks changes dramatically with the horizon over which the growth in the market is measured. The next section provides a parsimonious model with which to interpret these disturbances and empirical patterns.

4 The Model

Aggregate output \mathcal{Y}_t is assumed to be governed by a constant returns to scale process:

$$\mathcal{Y}_t = \mathcal{A}_t N_t^\alpha K_t^{1-\alpha}, \quad (7)$$

where \mathcal{A}_t is a factor neutral TFP shock, and N_t and K_t are inputs of labor and capital, respectively. We 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 income 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

$$D_t = \mathcal{Y}_t - W_t N_t, \quad (8)$$

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 $f(Z_t)$:

$$W_t = [\alpha \mathcal{A}_t N_t^{\alpha-1} K_t^{\alpha-1}] f(Z_t) = \alpha \mathcal{A}_t f(Z_t). \quad (9)$$

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. A shock of this sort is required to explain the low frequency behavior of the stock market, as show 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 the prevalence of unionization. At a more basic level, it is a reduced-form way of capturing shifts in the allocation of rewards between shareholders and workers while holding fixed the size of those rewards that could occur for any number of reasons, including a change in competitiveness, a change in how labor intensive production is, factor specific technological change.

With this specification for wages, log dividends are given by

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

where $\mathbf{a}_t \equiv \ln \mathcal{A}_t$. 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), driven solely by the TFP shock:

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

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

$$\Delta \mathbf{a}_t - \mu_a = \phi_a (\Delta \mathbf{a}_{t-1} - \mu_a) + \sigma_a \varepsilon_{a,t}, \quad \varepsilon_{a,t} \sim i.i.d. (0, 1). \quad (10)$$

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

$$Z_t = \phi_z Z_{t-1} + \sigma_z \varepsilon_{z,t}, \quad \varepsilon_{z,t} \sim i.i.d. (0, 1). \quad (11)$$

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 $1 + \mu_a$, the deterministic rate of technological progress.

Worker preferences play no role in asset pricing since they hold no assets. We assume that the economy is populated by a large number of identical shareholders, leading 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. The representative shareholder in this model is akin to a large institutional investor or wealthy individual who earns income only from investments. 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. Identical shareholders maximize the function

$$U = E \sum_{t=0}^{\infty} \prod_{k=0}^t \beta_k u(C_{it}^s) \quad (12)$$

with

$$u(C_{it}^s) = \frac{(C_{it}^s)^{1-x_{t-1}}}{1-x_{t-1}}, \quad (13)$$

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. As we show below, this shifter must have low (or zero) correlation with consumption and labor income fluctuations, in order to match evidence that movements in risk premia are divorced from traditional economic fundamentals such as consumption and wage income.

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:

$$\beta_t \equiv \frac{\exp(-r_f)}{E_t \left[\frac{D_{t+1}^{-x_t}}{D_t^{-x_{t-1}}} \right]}, \quad (14)$$

where r_f is a parameter. Aggregate shareholder consumption, given by D_t , is taken as given by individual shareholders and is therefore not internalized in the individual optimization problem. The specification for the subjective time discount factor in (14) is essential for obtaining a stable risk-free rate along 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. The calibration above makes the risk-free rate is constant, given by the exponentiation of the parameter r_f .³ 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.

The intertemporal marginal rate of substitution of stockholder consumption is the stochastic discount factor (SDF) given by:

$$M_{t+1} = \frac{\exp(-r_f) \left(\frac{D_{t+1}}{D_t} \right)^{-x_t}}{E_t \left[\left(\frac{D_{t+1}}{D_t} \right)^{-x_t} \right]}. \quad (15)$$

This can be written

$$M_{t+1} = \exp[-r_f - \ln E_t \exp(-x_t \Delta d_{t+1}) - x_t \Delta d_{t+1}]. \quad (16)$$

The return on a risk-free asset whose value is known with certainty at time t is given by $R_{f,t+1} \equiv (E_t [M_{t+1}])^{-1}$.

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

³While this adjustment to the time-discount factor β may seem unusual, it is in fact a generalization to non-lognormal functions of a familiar compensating Jensen's term that appears in lognormal models of the stochastic discount factor (e.g., Campbell and Cochrane (1999)) and (Lettau and Wachter (2007)).

variable \tilde{x} that itself can take unbounded values:

$$x_t = \bar{\theta} + \frac{\bar{\theta} - \underline{\theta}}{1 + \exp(-\tilde{x}_t)},$$

$$\tilde{x}_t - \mu_{\tilde{x}} = \phi_{\tilde{x}}(\tilde{x}_{t-1} - \mu_{\tilde{x}}) + \sigma_{\tilde{x}}\varepsilon_{\tilde{x},t}, \quad \varepsilon_{\tilde{x},t} \sim i.i.d.(0, 1). \quad (17)$$

In the above, $\bar{\theta}$ and $\underline{\theta}$ are parameters that controls the maximum and minimum values, respectively, of x_t . The autoregressive parameter is restricted to $0 < \phi_{\tilde{x}} < 1$.

As a benchmark, we specify the three shocks in the model $\varepsilon_{a,t}$, $\varepsilon_{z,t}$, and $\varepsilon_{\tilde{x},t}$ to be uncorrelated. We show below that this specification allows the model to closely correspond with the empirical evidence.

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 $R_{t+1} = (P_{t+1} + D_{t+1})/P_t$. We denote the log return on equity as $\ln(R_{t+1}) \equiv r_{t+1}$, and the log excess return $\ln(R_{t+1}/R_{f,t+1}) \equiv r_{t+1}^{ex}$. From the shareholder's first-order condition for optimal consumption choice, the price-dividend ratio satisfies

$$\frac{P_t}{D_t}(\mathfrak{s}_t) = E_t \exp \left(m_{t+1} + \Delta d_{t+1} + \ln \left(\frac{P_{t+1}}{D_{t+1}}(\mathfrak{s}_{t+1}) + 1 \right) \right), \quad (18)$$

where \mathfrak{s}_t is a vector of state variables, $\mathfrak{s}_t \equiv (\Delta \mathbf{a}_t, Z_t, x_t)'$. There is no closed-form solution to the functional equation (18). We therefore solve the $\frac{P_t}{D_t}(\mathfrak{s}_t)$ 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.

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

$$RRA = \frac{-C_t^s u''(C_{t+1}^s)}{u'(C_{t+1}^s)} = E(x_t) / (1 + \mu_a),$$

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 . We refer to x_t as a *risk aversion* shock.

4.1 Calibration

The Appendix contains a table that lists all parameters and their calibrated values. The parameter α 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 $\phi_a = 0$ so that the log level of TFP, α_t , follows a unit root stochastic process with drift. The mean and standard deviation of productivity $\Delta\alpha_t$ 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 yet stationary, with $\phi_z = 0.995$, in order 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. The parameters of the risk aversion process $\sigma_{\tilde{x}}$, $\phi_{\tilde{x}}$, 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. An interesting result is that, matching (i) and (ii) simultaneously requires a risk-aversion process that is very low most of the time but highly skewed to the right, characterized by the expectation of rare states in which the market's risk tolerance implodes, leading to a "flight to safety" and a market crash.

To understand why, observe that shareholders who consume out of dividends are exposed to much greater systematic risk than would be the case for one 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, the model generates too little variation in risk premia and overshoots the mean price-dividend ratio. The density of our risk aversion process therefore has most of its mass close to zero, with the median and mode equal to unity. 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 450.⁴

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. Changing expectations of an rare spike in risk aversion generate fluctuations in the price-dividend ratio that are far less non-linear in the state than are the risk aversion dynamics itself (though fluctuations in $p_t - d_t$ are naturally largest in crisis times).

5 Results

Table 2 presents summary asset pricing statistics of the model and compares them to those in post-war data. The model closely matches the mean and standard deviation of the equity premium and price-dividend ratio. By construction, the model exactly matches the mean risk-free rate. The model also does a good job of matching the volatility of dividend growth. Because dividends are subject to the factors share shock, they are more volatile than aggregate consumption. While the model is broadly consistent with these benchmark asset pricing moments, it is limited in matching the data in one way. 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

⁴This 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.

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 to incorporate wage smoothing or stickiness (e.g., Favilukis and Lin (2013)).

Table 3 reports 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, $\sum_{j=0}^h \Delta d_{t+j+1}$. The log price-dividend ratio predicts future excess returns with statistically significant negative coefficients in the model, while the coefficients for consumption and dividend growth are statistically indistinguishable from zero. These implications are consistent with the data. The adjusted R^2 statistics for forecasting excess returns are comparable between model and 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)).

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 disturbances implied by the model, and compare them to the primitive shocks. Figure 1 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. 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 then simulating forward using the solved policy functions. Thus we plot the responses to a one standard deviation *increase* in $\varepsilon_{a,t}$, and *decrease* in $\varepsilon_{z,t}$ and $\varepsilon_{\tilde{x},t}$. 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 .⁵

⁵For this plot 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.

The key result shown in Figure 1 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 nonlinearities in the model not captured by the linear VAR. 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 $\varepsilon_{a,t}$; 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 $\varepsilon_{\tilde{x},t}$.⁶ We verify, from a long simulation of the model, that the correlation between the consumption shock $e_{c,t}$ and the productivity shock $\varepsilon_{a,t}$ 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 $\varepsilon_{\tilde{x},t}$ is 0.97.⁷

In presenting the above, we do not claim that the mutually uncorrelated VAR shocks ($e_{c,t}, e_{y,t}, e_{a,t}$) exactly equal the primitive shocks ($\varepsilon_{a,t}, \varepsilon_{z,t}, \varepsilon_{\tilde{x},t}$), 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.

5.1 Origins of Quarterly Stock Market Fluctuations

With this theoretical interpretation of the VAR disturbances in hand, we now study the role of the empirical disturbances for historical stock market data, beginning with their role in

⁶These correlations are sample correlations over the paths of IRFs of length 20Q. Perfect correlation is equivalent to the IRFs being identical up to a normalization.

⁷The innovation in Δa_t attributable to risk aversion shocks is computed as $\Delta a_t - E[\Delta a_t | \mathfrak{s}_{t-1}, Z_t, \Delta \mathbf{a}_t]$.

quarterly fluctuations.

Figure 2 shows two sets of cumulative dynamic responses of Δc_t , Δa_t , and Δy_t . The Appendix shows 90% error bands for these responses using a bootstrap procedure. 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 1, except that the responses in Figure 2 are averages across 1000 samples of size 238 quarters, rather than over one very long sample. The right column of Figure 2 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. 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. 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. Because the TFP shock in the model is the innovation to a random walk, in simulated data all three variables move immediately to a new, permanently higher level. In historical data, full adjustment does not happen entirely within one quarter, but it is still relatively fast and occurs within 3 quarters or less, close to what would occur as the result of an innovation to a random walk. 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 2 displays the dynamic responses of c_t , a_t , and y_t to the labor income shock $e_{y,t}$. Observe that, in both the model and the data, the response of consumption to this shock is economically negligible at all horizons. The zero response on impact occurs by construction as a result of our identifying assumption. But it is also true in all subsequent periods, a key finding that is not part of our identifying assumption. Instead, this shock is purely redistributive and drives a_t and y_t in opposite directions. The effect on labor earnings is large and immediate: labor income jumps to a new lower level within the quarter. The affect on wealth is also large but takes time to adjust. This sluggishness is puzzling because it suggests that the information revealed in the innovation is incorporated slowly into asset prices. Composition effects could play a role in this if, for example, an increasing fraction of firms going public over the sample employ labor-saving technologies. It could

also reflect 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. These possibilities are outside of the scope of the model. Future research is needed to formally investigate these and other possibilities.⁸

The third row of Figure 2 shows the effects of a positive wealth shock, $e_{a,t}$, driven in the model by a decline in risk aversion. 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 that is central to understanding why the risk aversion shock must be modeled as independent of consumption and labor earnings shocks. Although transitory, this disturbance's influence on a_t has a half-life of over four years in historical data.

Figure 3 shows the cumulative 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. It is clear that the responses of stock wealth to the wealth and labor income shocks 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 two disturbances.

Returning to Table A.1, we see that the wealth shock e_a explains the largest fraction of quarterly stock wealth growth Δs_t and accounts for 76% of its quarterly variation (75% of

⁸The left column of Figure 2 shows that, even in the model, there is slight sluggish response of stock market wealth to the factors share shock, although it is less pronounced than in the data. This is a finite sample effect: the responses in Figure 2 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 estimated response is sluggish, even though the population response displays no sluggishness. This can be seen via a comparison with Figure 1, which, unlike the response in Figure 2, is computed from one very long simulation of the model, rather than from averages over many short ones.

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. The model implications in the third row are 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. 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. But note that, both in the model and in the data, aggregate consumption shocks play a very small role in quarterly stock market fluctuations. In the model, this occurs because labor’s greater average role in the production process (steady state labor share is two-thirds) means that most gains and losses from TFP shocks accrue to workers rather than shareholders, so these shocks are less important for asset pricing than are the other two. This finding is difficult to reconcile with representative agent models where aggregate consumption shocks play the key role in asset price fluctuations. As an illustration, the last row of Table 1 gives the corresponding variance decomposition numbers for the Campbell and Cochrane (1999) habit model (with no labor income), in which 84% of quarterly stock price growth is driven by consumption shocks.⁹

5.2 Origins of Long-Run Stock Market Wealth

We now turn to the question of how the sources of stock market fluctuation vary with the time horizon. To do so we first 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 $\Delta s_t - \kappa_0$ (5). Noting that $\Delta s_t - \kappa_0$ in (5) is a function of three components, Δs_t^c , Δs_t^y , Δs_t^a , plus an i.i.d. residual η_t , 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

⁹The 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.

at cycles of different lengths, in quarters. The Appendix provides additional details.

Figure 4 exhibits these decompositions for the model (top panel), and for the data using stock market wealth (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 wealth shocks is very high, close to 80%. As these cycles become longer, the fraction of variance in stock wealth explained by this shock declines and asymptotes to roughly 40%. Note that the high frequency, short horizon, variability of the stock market in post-war data is virtually unrelated to the labor income/factors share shocks. But 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 shocks. By contrast, both in the model and the data, 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 is less than 20% of the variability in the stock market at all frequencies, and asymptotes to around 10% as the horizon extends. The model captures this frequency decomposition well.

Next, we study the role each disturbance has played in driving stock market wealth at specific points in our sample using the levels decomposition of stock market wealth (6). To do so, we remove the deterministic trend and normalize the initial observation s_0 to zero in the quarter before the start of our sample. Figure 5 plots the levels decomposition for stock wealth (left column) and stock price (right column) over our sample, which are very similar. 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 show to the component attributable to the cumulation of each shock and the residual.

It is immediately clear from Figure 5 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 small. Instead, the big movers of stock market wealth are the factors share shock and the risk aversion shock. The low frequency movements in the level of stock market wealth are well tracked by the cumulative swings in the factor shares component, while shorter-lived peaks and troughs in the stock market accord well with spikes up or down in the risk aversion component.

Figure 5 also shows that the cumulative effect of the factors share shock has persistently boosted stock market wealth over the last twenty five years. 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. Figure 6 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.

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. The cumulative affects of the factors share shock 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.) An additional 38% of the increase since 1980, or a rise of 22%, is attributable to the cumulative effects of risk aversion shocks. The TFP shocks have made a negative contribution to change in stock market wealth since 1980, a direct consequence of the string of unusually large negative draws for the consumption/productivity shock in the Great Recession years from 2007-2009. These shocks accounted for -38% of the total increase since 1980. The residual accounts for the remaining -10.5% of the increase. These findings underscore 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 raise or lower all of them.

The calculations above removed a deterministic trend. As for any series that deterministically trends upward over time (stock market, GDP, consumption, etc.,) most of the increase over long periods is attributable to a deterministic trend. We can assess the quantitative importance of stochastic shocks for long-term growth inclusive of the deterministic trend by

“shutting off” the shock and studying where the level of the stock market would be today under that counterfactual that the shocks had been zero over some period. Doing so for the period since 1980, we find that the stock market would be 47% lower today than it currently is at the end of our sample had the factors share shock been set to zero.

5.3 Origins of Stock Market Predictability

Our last subject is 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.¹⁰ 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)). This section provides evidence on the question of why excess returns are predictable 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 4, with the top half showing results from historical data, and the bottom half showing results from the model.

Table 4 has several panels. The left panel reports regression results of one through three year log excess equity returns on the lagged price-dividend ratio alone. Moving rightward, the next panel reports regression results of one through three 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 a measure of stochastic consumption volatility. The panel to the right of this one, headed “ e_a only,” reports regression results of one through three year log excess equity returns on multiple lags of the i.i.d. wealth disturbances $e_{a,t}$ 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 that is *unrelated* 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 .

¹⁰For extensive reviews of this evidence see Campbell, Lo, and MacKinlay (1997), Cochrane (2005), Lettau and Van Nieuwerburgh (2008), and Lettau and Ludvigson (2010).

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^2 statistics. By contrast, time-varying stochastic consumption volatility has no predictive power for equity premia at any horizon.¹¹ These results provide no evidence that stock return predictability is driven by time-varying second moments of consumption growth or broad-based macroeconomic uncertainty.

By comparison with pd or cay , lags of the $e_{a,t}$ wealth shocks exhibit greater forecasting power than either of these variables. A Wald test strongly rejects the hypothesis that the sum of squared coefficients on the lags of these shocks is zero.¹² There is no horizon at which the wealth shocks are not strongly statistically marginally significant. A positive innovation for the wealth shock *increases* 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.

But the next columns show that 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. The pd residual components pd^{orth} that are orthogonal to the wealth shocks have no statistically significant forecasting power for equity premia. Similarly, the component of cay_t that is driven solely by the orthogonal wealth shocks, $e_{a,t}$, is responsible for all the forecasting power of cay_t . 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

¹¹A stochastic volatility model is used to estimate $E_t \left([\Delta \ln C_{t+h} - E_t [\Delta \ln C_{t+h}]]^2 \right)$, for different horizons h . The estimate is taken from Jurado, Ludvigson, and Ng (2014). 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 macroeconomic uncertainty developed in Jurado, Ludvigson, and Ng (2014) also exhibit no forecasting power for equity premia at any horizon. These results are omitted to conserve space but are available upon request.

¹²Wald tests similarly reject the hypothesis that the coefficients are jointly zero, and that the sum of coefficients is zero.

and Cochrane (1999)), since the wealth shocks that “explain” most of the forecastability of excess returns are orthogonal to movements in consumption. The opposite would be true in models with habit formation.

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 and labor income instead originate from shocks to measures of fundamental stock market value such as dividends or earnings. In summary, changes in the reward for bearing stock market risk are found to be 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.

6 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 disturbances that explain the vast majority of fluctuations since the early 1950s. We then propose a model to interpret these disturbances 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.

The results show 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.¹³ The other is a cash-flow innovation that redistributes the rewards of production between shareholders and workers with no change in 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 extends. 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. 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 total 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. The shocks responsible for big historical movements in stock market wealth are not those that raise or lower aggregate rewards, but are instead ones that redistribute a given level of rewards between workers and shareholders. We also show that predictability of excess stock market returns must 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. We argue that these findings have important implications for macroeconomic modeling: the two big sources of variation that we find here are responsible for almost all of stock market fluctuation presently play virtually no role in contemporary macroeconomic theory.

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

¹³One 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)).

environment. In future work, we plan to examine a richer model of the production side, with close attention to how important changes in the labor market over the last 30 years may have contributed to our findings on factors share shifts. Our results in this paper 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.

References

- ABEL, A. B. (1983): “Optimal Investment under Uncertainty,” *American Economic Review*, 73(1), 228–233.
- ALBUQUERQUE, R., M. EICHENBAUM, AND S. REBELO (2012): “Valuation Risk and Asset Pricing,” Unpublished Paper, Northwestern University.
- 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 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.
- 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(1), 241–65.
- (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.
- FAVILUKIS, J., AND X. LIN (2013): “Does Wage Rigidity Make Firms Riskier? Evidence From Long-Horizon Return Predictability,” <http://personal.lse.ac.uk/FAVILUKI/WageRigidEmpirical.pdf>.
- 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.
- GOURIO, F. (2012): “Disaster Risk and Business Cycles,” *American Economic Review*, 102(6).
- 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.

- 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.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2014): “Measuring Uncertainty,” *The American Economic Review*, forthcoming.
- 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.
- (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. 617–90. Elsevier Science B.V., North Holland, Amsterdam.
- LETTAU, M., AND S. C. LUDVIGSON (2013): “Shocks and Crashes,” in *National Bureau of Economics Research Macroeconomics Annual: 2013*, ed. by J. Parker, and M. Woodford, vol. 28, pp. 293–354. MIT Press, Cambridge and London.

- LETTAU, M., S. C. LUDVIGSON, AND S. MA (2014): “Capital Share Risk and Shareholder Heterogeneity in U.S. Stock Pricing,” National Bureau of Economic Research Working Paper No. w20744.
- 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 Non-stockholders,” *Journal of Financial Economics*, 29(1), 97–112.
- MUIR, T. (2014): “Financial Crises and Risk Premia,” Unpublished paper, Yale University.
- 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.
- 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.

Figures and Tables

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

Table 1: Variance Decomposition of Quarterly Log Difference in Stock Wealth. See Table 2. This table reports a variance decomposition of the quarterly log difference in stock market wealth using. 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. 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. The historical sample spans the period 1952:Q2 - 2012:Q4.

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

Table 2: Simulated and Data Moments. Δc_t and Δd_t are log differences of real consumption and dividends; r_t^e is the cum-dividend value-weighted CRSP return; r_t^f is the constant-maturity 1-year T-bill rate. $r_t^{ex} \equiv r_t^e - r_t^f$. Data for $p_t - d_t$ are obtained from CRSP. All variables are at annual frequency. The sample is 1953:4 - 2012:6.

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

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

MODEL						
$Y_{t,t+h}$:	$\frac{\sum_{j=1}^h \Delta c_{t+j}}{pd_t}$		$\frac{\sum_{j=1}^h \Delta d_{t+j}}{pd_t}$		$\frac{\sum_{j=1}^h r_{t+j}^{ex}}{pd_t}$	
h	R^2	R^2	R^2	R^2	R^2	R^2
4	0.006 (0.007)	0.011	0.076 (1.545)	0.034	-0.266 (-3.600)	0.139
8	0.028 (0.026)	0.024	0.144 (1.699)	0.067	-0.369 (-3.953)	0.195
12	0.024 (0.035)	0.036	0.206 (2.146)	0.096	-0.422 (-4.145)	0.223

Table 3: Long Horizon Predictability Regressions. 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.

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		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	$e_{a,t}$	\bar{R}^2	pd_t^{orth}	\bar{R}^2	$cay_{a,t}$	\bar{R}^2
4	-0.130 (-2.856)	0.068	2.722 (2.936)	0.074	-0.037 (-1.030)	0.003	-0.401 (7.739)	0.137	-0.083 (-1.571)	0.023	7.357 (3.684)	0.128
8	-0.226 (-2.777)	0.120	4.864 (3.326)	0.133	-0.018 (-0.299)	-0.004	-0.673 (12.686)	0.262	-0.118 (-1.363)	0.029	13.736 (5.575)	0.261
12	-0.272 (-3.043)	0.144	6.667 (4.098)	0.197	0.014 (0.169)	-0.005	-0.678 (8.419)	0.314	-0.135 (-1.276)	0.031	16.423 (7.052)	0.310

MODEL											
X_t :	pd		cay		e_a		pd^{orth}		cay_a		
h	pd_t	\bar{R}^2	cay_t	\bar{R}^2	$e_{a,t}$	\bar{R}^2	pd_t^{orth}	\bar{R}^2	$cay_{a,t}$	\bar{R}^2	
4	-0.266 (-3.600)	0.139	0.693 (3.593)	0.084	-0.283 (13.363)	0.273	-0.070 (-0.164)	0.015	1.459 (6.279)	0.274	
8	-0.369 (-3.953)	0.195	0.984 (3.643)	0.119	-0.399 (16.236)	0.351	-0.111 (-0.252)	0.026	2.026 (7.185)	0.358	
12	-0.422 (-4.097)	0.223	1.114 (3.526)	0.129	-0.445 (17.649)	0.372	-0.143 (-0.337)	0.034	2.277 (7.365)	0.383	

Table 4: Long Horizon Return Regressions. 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 $e_{a,t} = (e_{a,t}, \dots, e_{a,t-19})$, 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.

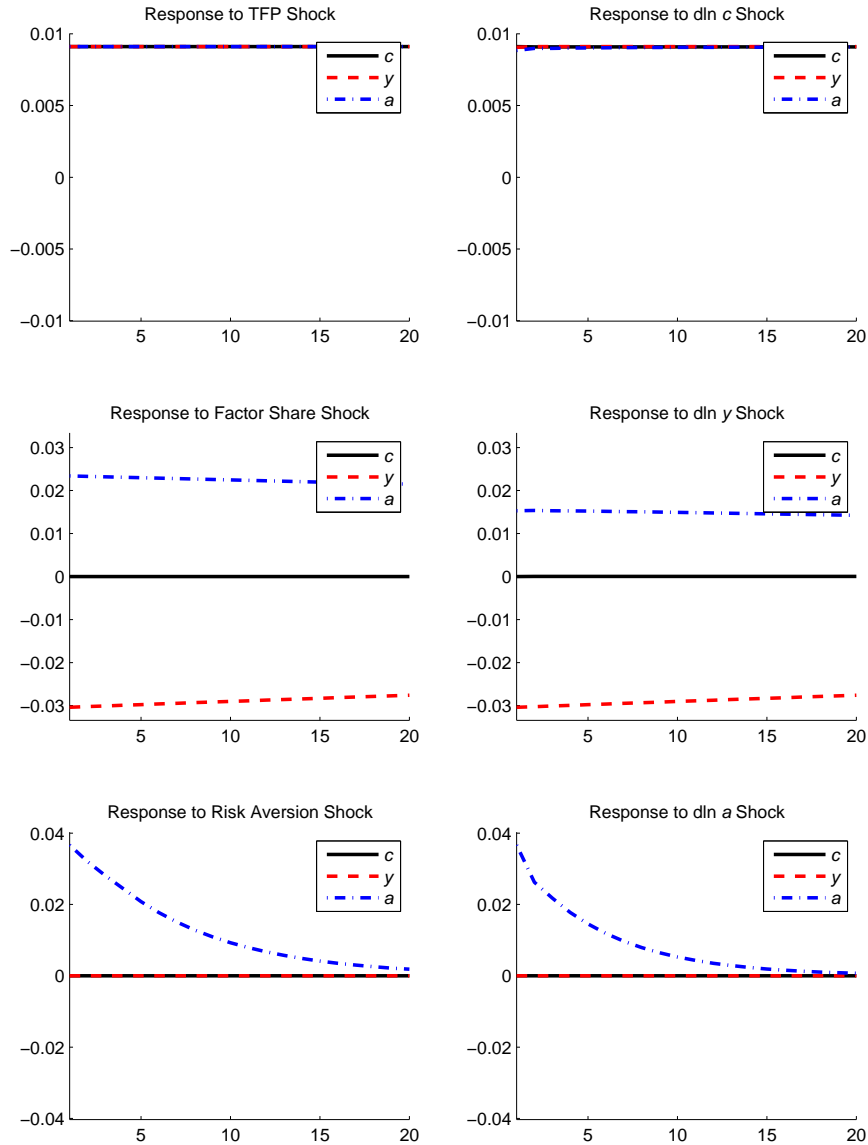


Figure 1: Model-Based Responses: Primitive vs. VAR Shocks. The left column plots the cumulated model impulse responses of the log differences of c , y , and a to the primitive shock named in the sub-graph title. The right column plots the cumulative impulse responses implied by the model from a VAR in 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.

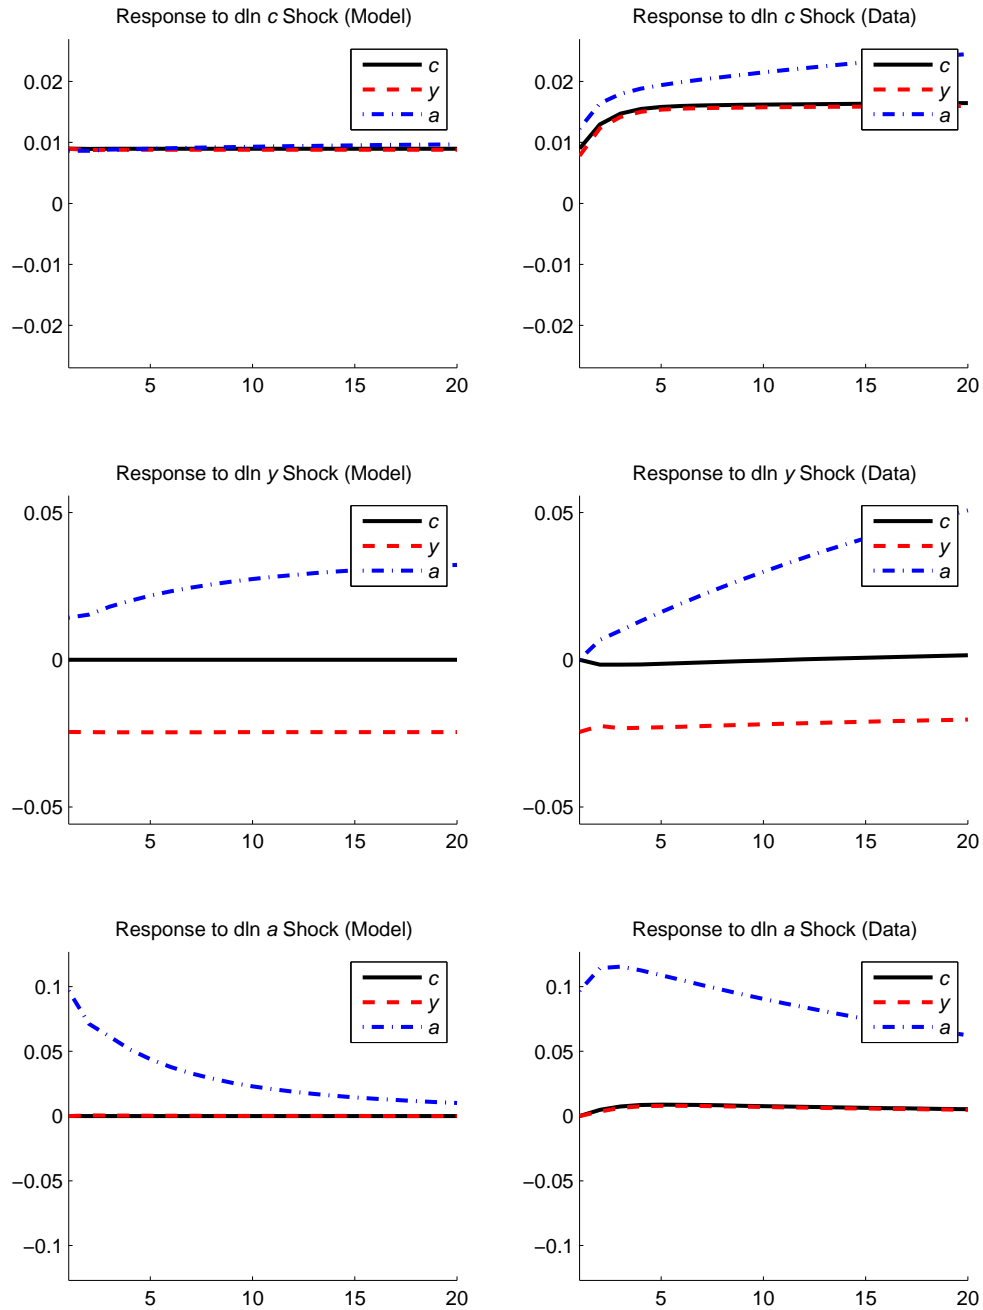


Figure 2: VAR Impulse Responses (Model vs. Data). The figure plots impulse response functions to the VAR shocks in both the model (left column) and data (right column). The structural shocks to the orthogonalized shocks obtained from the VECM regression using data simulated from the model (left column), 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 3: Stock Market Impulse Responses. 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 4: Decomposition of Spectrum (Model vs. Data). 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.

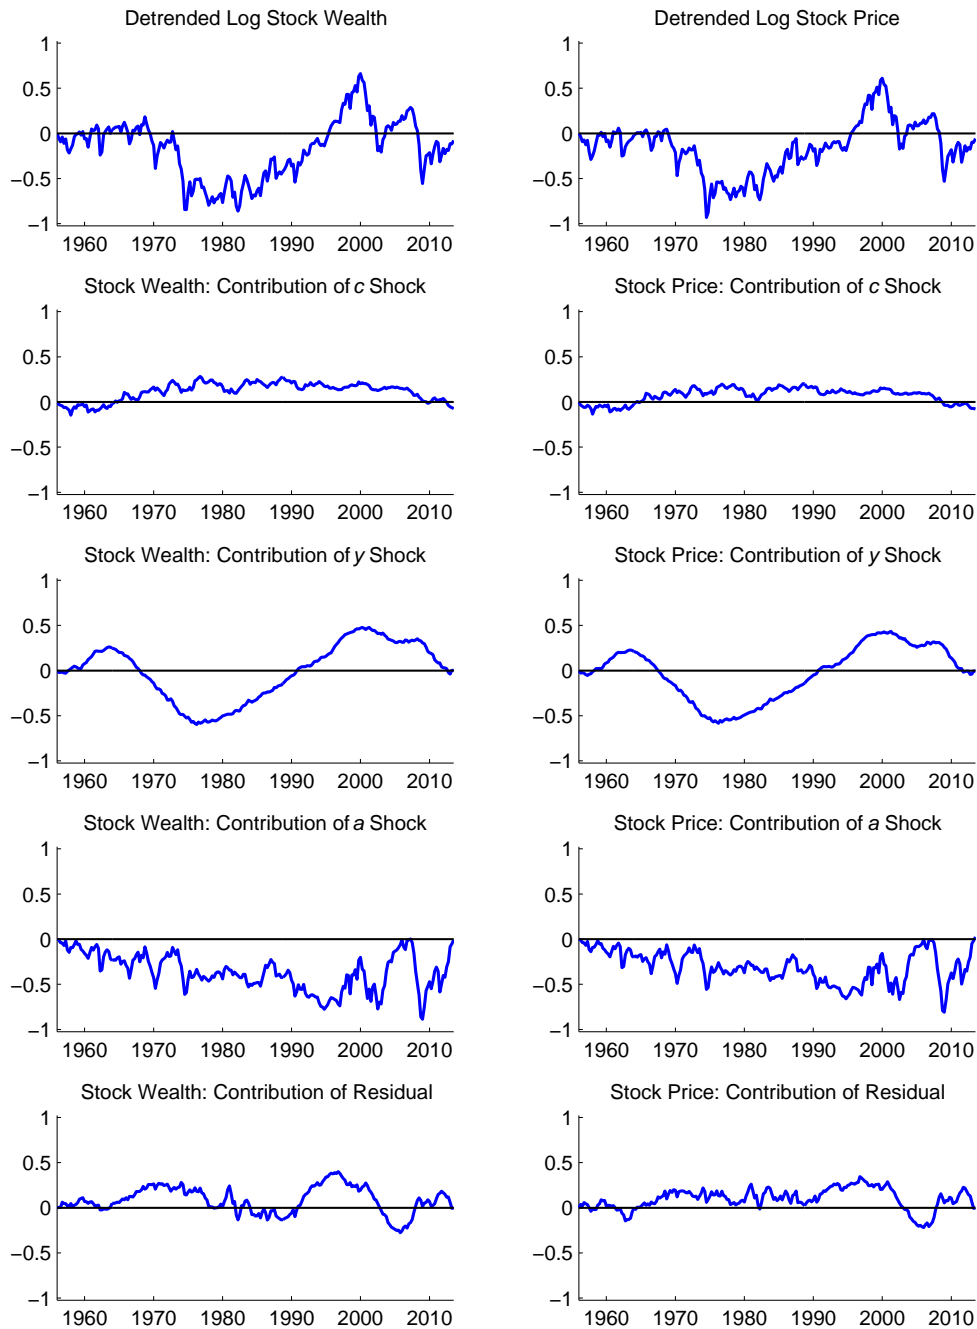


Figure 5: Level Decomposition (Data). See Table 2. The figure shows the decomposition of the log level of stock wealth into components driven by the orthogonalized c , y , and a shocks obtained from the VECM regression. A deterministic trend is removed from the log level of stock wealth by removing the 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 6: Decomposition of Labor Income and Stock Market Wealth. 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.

Appendix: For Online Publication

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 $[\text{wages and salaries}/(\text{wages and salaries} + \text{proprietors' income with IVA and CCADJ} + \text{rental income} + \text{personal dividends} + \text{personal interest income})]$ 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

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

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

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

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: $d_{t+12} - d_t = d_{t+12} - d_{t+11} + d_{t+11} - d_{t+10} + \dots + d_{t+1} - d_t$. 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 here. 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 $\ln(P_t^{Ax}/D_t^A)$. 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.

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.

ANNUALIZATION

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.

Risk Aversion Along a Balanced Growth Path

The budget constraint for the representative shareholder can be written

$$A_t = \theta_t (P_t + D_t) + B_t \tag{19}$$

$$C_t^s + B_{t+1}q_t + \theta_{t+1}P_t \leq A_t, \tag{20}$$

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 $q_t = 1/(1 + R_f)$ 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 (19) as

$$\begin{aligned}
A_{t+1} &= P_t \theta_{t+1} \left(\frac{P_{t+1} + D_{t+1}}{P_t} \right) + B_{t+1} \\
&= P_t \theta_{t+1} R_{t+1} + B_{t+1} \Rightarrow \\
P_t \theta_{t+1} &= \frac{A_{t+1}}{R_{t+1}} - \frac{B_{t+1}}{R_{t+1}}.
\end{aligned} \tag{21}$$

Plugging (21) into (20) and evaluating (20) 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:

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

or solving forward

$$A_t = \sum_{i=0}^{\infty} \left(\frac{1}{R_f} \right)^i C_{t+i}^s. \tag{22}$$

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 (20)

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

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

$$RRA_t = \frac{-A_t E_t V''(A_{t+1})}{E_t V'(A_{t+1})}. \tag{23}$$

It can be shown (see below) that

$$V'(A_{t+1}) = u'(C_{t+1}^s) \Rightarrow \tag{24}$$

$$V''(A_{t+1}) = u''(C_{t+1}^s) \frac{\partial C_{t+1}^{s*}}{\partial A_{t+1}}, \tag{25}$$

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 $G = 1 + \mu_a$ 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

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

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

$$\begin{aligned} u'(C_t^{s*}) &= \beta R_f u''(C_{t+1}^{s*}) \Rightarrow \\ u''(C_t^{s*}) \frac{\partial C_t^s}{\partial A_t} &= \beta R_f u''(C_{t+1}^{s*}) \frac{\partial C_{t+1}^{s*}}{\partial A_t} \Rightarrow \\ -\bar{x}(C_t^s)^{-\bar{x}-1} \frac{\partial C_t^{s*}}{\partial A_t} &= \frac{\exp(-r_f)}{G^{\bar{x}}} R_f \left[-\bar{x}(G C_t^{s*})^{-\bar{x}-1} \right] \frac{\partial C_{t+1}^{s*}}{\partial A_t} \Rightarrow \\ \frac{\partial C_t^{s*}}{\partial A_t} &= G^{-1} \frac{\partial C_{t+1}^{s*}}{\partial A_t}. \end{aligned} \quad (26)$$

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

$$\begin{aligned} u''(C_{t+1}^{s*}) \frac{\partial C_{t+1}^s}{\partial A_t} &= \beta R_f u''(C_{t+2}^{s*}) \frac{\partial C_{t+2}^{s*}}{\partial A_t} \Rightarrow \\ \frac{\partial C_{t+1}^{s*}}{\partial A_t} &= G^{-1} \frac{\partial C_{t+2}^{s*}}{\partial A_t} \end{aligned} \quad (27)$$

and combining (26) and (27) we obtain

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

and iterating obtain

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

Now differentiate (22) evaluated along the balanced growth path with respect to A_t :

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

implying

$$\frac{\partial C_t^{s*}}{A_t} = \frac{R_f - G}{R_f}. \quad (28)$$

Assets A_t along a non-stochastic balanced growth path are

$$\begin{aligned}
A_t &\equiv \sum_{i=0}^{\infty} \left(\frac{1}{R_f} \right)^i C_{t+i}^{s*} \\
&= \sum_{i=0}^{\infty} \left(\frac{G}{R_f} \right)^i C_t^{s*} \\
&= \frac{R_f C_t^{s*}}{(R_f - G)}.
\end{aligned} \quad (29)$$

Plugging (24), (25), (28), and (29) into (23), we obtain a value for risk aversion along a balanced growth path equal to

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

Derivation of (24). First-order condition for B_{t+1} :

$$-u'(C_t^s) q_t + \beta_t V'(A_{t+1}) \underbrace{\frac{\partial A_{t+1}}{\partial B_{t+1}}}_{=1 \text{ from (19)}} = 0. \quad (30)$$

First-order condition for θ_{t+1} :

$$-u'(C_t^s) P_t + \beta_t V'(A_{t+1}) (P_{t+1} + D_{t+1}) = 0. \quad (31)$$

Differentiate the value function

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

with respect to A_t , keeping in mind

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

and

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

We have

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

Evaluating at the optimum using (30) and (31), the terms in brackets are zero, leaving

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

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 A.4 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 A.4 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.¹⁴ 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 A.4 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.

Numerical Solution

The price-dividend ratio satisfies

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

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 $\frac{P_t}{D_t}(i, j, k)$, $i, j, k = 1, \dots, N$, 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

¹⁴This 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.

for the population spectrum for each component in (5) 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

$$\frac{S_{\Delta s^c}(\omega)}{S_{\Delta s}(\omega)}, \quad (32)$$

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 (32), 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

$$\begin{aligned} \Delta \mathbf{x}_t &= \hat{\nu} + \hat{\gamma} \hat{\alpha}' \mathbf{x}_{t-1} + \hat{\Gamma}(L) \Delta \mathbf{x}_{t-1} + \hat{H} e_t \\ \Delta s_t &= \hat{\kappa}_0 + \hat{\kappa}_c(L) e_{c,t} + \hat{\kappa}_y(L) e_{y,t} + \hat{\kappa}_a(L) e_{a,t} + \eta_t, \end{aligned} \quad (33)$$

denoted $(\hat{e}_{c,t}, \hat{e}_{y,t}, \hat{e}_{a,t}, \hat{\eta}_t)$ 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 $(\hat{e}_{c,t}, \hat{e}_{y,t}, \hat{e}_{a,t}, \hat{\eta}_t)$ with replacement. Denote the new randomly sampled (via block bootstrap) values for the residuals $(\tilde{e}_{c,t}, \tilde{e}_{y,t}, \tilde{e}_{a,t}, \tilde{\eta}_t)$ for

$t = 1, \dots, T$. The new bootstrapped sample of observable data, $(\tilde{\mathbf{x}}_t, \tilde{s}_t)$, is constructed from

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

Given this new sample of data, all parameters in (33) (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 2 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 2 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

$$G = 1 - q - \alpha f(z). \tag{34}$$

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 (34).

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

$$\Delta \ln f(Z_t) \simeq b_1 e_{y,t} \tag{35}$$

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 $\hat{e}_{y,t}$ from the empirical VAR, we can now construct an implied series for the Gini coefficient in the model. Note from (35), we have

$$\ln f(Z_t) = \ln f(Z_0) + b_1 \sum_{i=1}^t e_{y,i}. \tag{36}$$

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 (36) and applying the exponential function to the left-hand-side yields an implied series for $f(z_t)$. Finally, plugging this value into (34) 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

$$\Delta x_t = \nu + \gamma \alpha' x_{t-1} + \Gamma \Delta x_{t-1} + H e_t \quad (37)$$

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

$$\Delta x_t = \delta + D(L) e_t \quad (38)$$

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

$$\Delta x_t = \delta + D_c(L) e_{c,t} + D_a(L) e_{a,t} + D_y(L) e_{y,t}. \quad (39)$$

If we define

$$\begin{aligned} \Delta x_{c,t} &\equiv D_c(L) e_{c,t} \\ \Delta x_{a,t} &\equiv D_a(L) e_{a,t} \\ \Delta x_{y,t} &\equiv D_y(L) e_{y,t} \end{aligned} \quad (40)$$

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

$$x_t = \delta t + x_{c,t} + x_{a,t} + x_{y,t}. \quad (41)$$

Premultiplying (41) by α' , and subtracting κ , we obtain a decomposition for *cay*_{*t*}:

$$cay_t = \alpha' x_t - \kappa = \alpha' \delta t + \alpha' x_{c,t} + \alpha' x_{a,t} + \alpha' x_{y,t} - \kappa \quad (42)$$

If we define

$$cay_{c,t} \equiv \alpha' x_{c,t}$$

$$cay_{a,t} \equiv \alpha' x_{a,t}$$

$$cay_{y,t} \equiv \alpha' x_{y,t}$$

then (42) becomes

$$cay_t = \alpha' \delta t + cay_{c,t} + cay_{a,t} + cay_{y,t} - \kappa. \quad (43)$$

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

Practical Decomposition

Unfortunately, we cannot fully implement (43) in practice. From our estimation procedure, we obtain estimates $\hat{\alpha}$, and \hat{D} . We also obtain estimates \hat{e}_t , but only for $t = 1, \dots, T$. If we had estimates of all \hat{e}_t for $t = T, T-1, \dots, -\infty$, then we could evaluate estimates of each x_c, x_a and x_y term using (40) and form an additive decomposition. However, we do not, so if we estimate these terms using only the shocks from $1, \dots, T$, 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 series x_t through (37). 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

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

where \hat{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 $\tilde{x}_c, \tilde{x}_a, \tilde{x}_y$. In practice, these series can be calculated by running the VECM (37) 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 $\Delta\tilde{x}_{c,t}$, begin by setting $\Delta\tilde{x}_{c,0} = \tilde{x}_{c,0} = (0, 0, 0)'$. Given $\Delta\tilde{x}_{t-1}$ and \tilde{x}_{t-1} for $t \geq 0$, we can obtain $\Delta\tilde{x}_{c,t}$ by applying (37) without the constant term ν , and allowing only the $e_{c,t}$ shock to be nonzero, so that we obtain

$$\Delta\tilde{x}_{c,t} = \hat{\gamma}\hat{\alpha}'\tilde{x}_{c,t-1} + \hat{\Gamma}\Delta\tilde{x}_{c,t-1} + H \begin{bmatrix} e_{c,t} \\ 0 \\ 0 \end{bmatrix}. \quad (44)$$

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

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

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

Because this decomposition cannot account for shocks prior to $t = 1$, (42) will not hold, and we will instead end up with a “residual” term cay_t^* , such that

$$cay_t = \widetilde{cay}_{c,t} + \widetilde{cay}_{a,t} + \widetilde{cay}_{y,t} + cay_t^* \quad (45)$$

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 (45), an additive decomposition without estimates of e_0, e_{-1}, \dots , we can take advantage of the fact that the influence of e_0, e_{-1}, \dots 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

$$\Delta x_1^* = \delta + D_1 e_0 + D_2 e_1 + \dots$$

$$\Delta x_2^* = \delta + D_2 e_0 + D_3 e_1 + \dots$$

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 ν) 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

$$\Delta \tilde{x}_{c,t} \equiv \sum_{j=0}^{t-1} D_{c,j} e_{c,t-j}$$

$$\Delta \tilde{x}_{a,t} \equiv \sum_{j=0}^{t-1} D_{a,j} e_{a,t-j}$$

$$\Delta \tilde{x}_{y,t} \equiv \sum_{j=0}^{t-1} D_{y,j} e_{y,t-j}$$

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 $\tilde{x}_0 = \Delta\tilde{x}_0 = (0, 0, 0)'$ and applying the relevant shock components one at a time, as in (44) of the previous section.

Under this definition, we have

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

and since

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

we obtain

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

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

Cumulating, and applying the α vector, we obtain

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

for

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

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 \tilde{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^*).

Additional Figures and Tables

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) $\Delta \mathbf{a}_t$	0.0057
6	$f(\bar{Z})$	mean $f(Z_t)$	1
7	$\phi_{\mathbf{a}}$	autocorrelation $\Delta \mathbf{a}_t$	0
8	ϕ_z	autocorrelation Z_t	0.995
9	$\phi_{\tilde{x}}$	autocorrelation \tilde{x}_t	0.85
10	$\sigma_{\mathbf{a}}$	standard deviation $\varepsilon_{\mathbf{a},t}$	0.0091
11	σ_z	standard deviation $\varepsilon_{z,t}$	0.12
12	$\sigma_{\tilde{x}}$	standard deviation $\varepsilon_{\tilde{x},t}$	6.01

Table A.1: Parameter List and Calibration. This table lists the parameters of the model and their baseline calibrated values.




Figure A.1: Consumption Shocks and TFP. 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. 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.

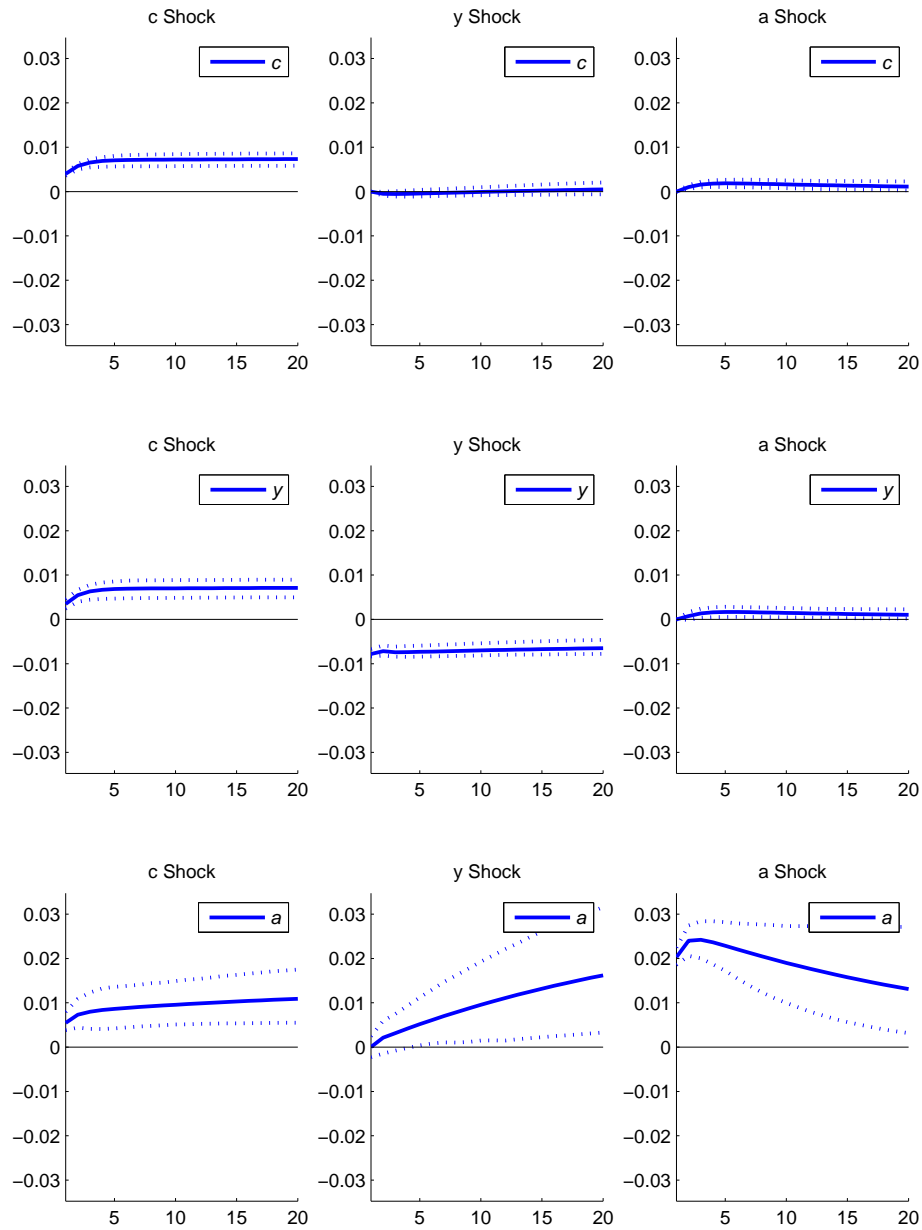



Figure A.3: VAR Impulse Responses with Error Bands. The figure plots impulse response functions to the VAR shocks obtained from the VECM regression using data. Student Version of MATLAB Impulse responses are obtained by estimating a cointegrated VAR, inverting to Wold representation and computing the 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.



Student Version of MATLAB

Figure A.4: Gini Coefficient and Cumulated e_y Shock. 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.




Figure A.5: Model Gini Coefficient. 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.