
Shocks and Crashes

Martin Lettau, *University of California, Berkeley, CEPR, and NBER*

Sydney C. Ludvigson, *New York University and NBER*

I. Introduction

What are the primary sources of fluctuations in real activity and financial markets? We address this question by decomposing the historical dynamics of log aggregate consumer spending, c_t , log labor earnings, y_t , and log asset wealth, a_t , into components driven by three mutually orthogonal structural disturbances, each of which we show plays a quantitatively large role in the joint dynamics of these variables in postwar data. The shocks we identify are econometrically distinguished only on the basis of their degree of persistence: two of the disturbances have permanent effects on the variables in the system, while the third has persistent but transitory effects. But we argue here for one particular economic interpretation of these shocks, and show that their relative importance has varied considerably during the two most recent expansion/recession episodes accompanied by asset market booms and subsequent crashes.

The first shock we identify is a permanent disturbance that has a long-run effect on consumption, labor earnings, and wealth, in a manner akin to a permanent (factor neutral) total factor productivity (TFP) shock in canonical frictionless stochastic dynamic general equilibrium models. A positive value for this shock quickly raises consumption to a new trend level and accounts for virtually all of its fluctuations. We argue that this shock may be plausibly interpreted as a permanent *productivity shock*.

The second shock is a permanent disturbance that moves labor earnings and wealth in opposite directions but leaves aggregate consumption unaffected. A positive value for this shock raises the stock market

component wealth and lowers labor income. This shock accounts for 77 percent of the quarterly fluctuations in labor income growth and is an important contributor to fluctuations in the level of the stock market over longer periods of time. We argue that this shock may be plausibly interpreted as a *factor shares shock*.

The third shock is a persistent but transitory innovation that accounts for the vast majority of short-run fluctuations in asset values but has a negligible impact on consumption and labor earnings, both contemporaneously and at all future horizons. These fluctuations are associated predominantly with the stock market component of wealth and have historically had smaller effects on other components of wealth, such as non-stock market financial wealth, and housing. The shock has a half-life of over four years and explains 90 percent of the quarterly variation in household net worth. We argue that this shock may be plausibly interpreted as an exogenous *risk aversion shock*.

Both the statistical identification and the economic interpretation of these mutually orthogonal structural disturbances are accomplished with three assumptions.

First, we use the restrictions implied by cointegration to distinguish disturbances on the basis of whether they have permanent or transitory effects. The presence of a cointegrating relation among c_t , a_t , and y_t follows from weak theoretical assumptions in any representative agent model where a budget constraint identity holds, and is supported by empirical evidence. Statistical tests for cointegration also imply that the vector time-series is well described by two permanent shocks and one transitory shock.

Second, we restrict the space spanned by the two permanent shocks to be orthogonal to the transitory shock. This restriction, which follows King et al. (1991) and Gonzalo and Ng (2001), may be thought of as a definition rather than an assumption. We *define* the transitory shock to be orthogonal to the space spanned by the permanent shocks, thereby allowing us to identify their independent effects. If in fact there were no transitory component in the system that had effects independent of the permanent disturbances, the econometric procedure would assign a zero role for this component in the variance decompositions of all variables in the system. This restriction, along with the previous one, allows us to completely identify the space spanned by the two permanent shocks and the single transitory shock.

The final identifying restriction allows us to distinguish the independent effects of the two permanent disturbances. For this, we make two

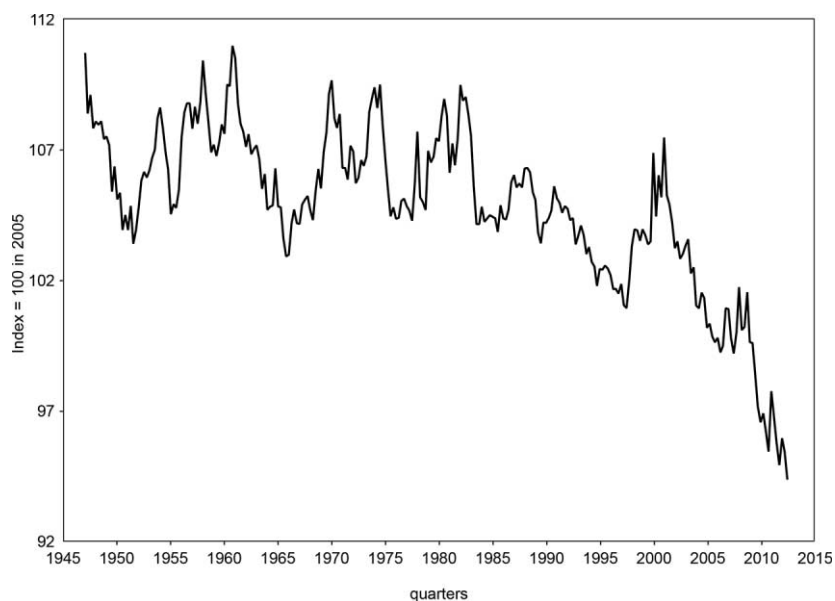


Fig. 1. BLS labor share

Notes: The figure plots the labor share constructed by the Bureau of Labor Statistics. The series is normalized to 100 in 2005:Q1. The sample is 1952:Q1 to 2012:Q3.

related assumptions. First we assume that technological innovation leads to some role for permanent productivity shocks in driving the joint dynamics of the system. Thus we assume that one of the permanent shocks in the system is a permanent TFP shock. Second, we observe that labor's share of output has fluctuated considerably over time in an extremely persistent manner (figure 1) and, as we show below, affects labor income growth independently of TFP growth. Thus we aim to recover a second permanent shock in the system that is related to shifts in factor shares. The only interpretable orthogonalization capable of recovering *both* a TFP shock and factor shares shock requires us restrict c to be ordered first in the cointegrated vector autoregression (VAR) before performing a Cholesky orthogonalization on the set of transformed innovations that reveal (potentially correlated) permanent and transitory shocks. The interpreted permanent productivity shock will then be revealed as a consumption shock, while the permanent factor shares shock will be revealed by a shock that affects labor income and wealth, with no contemporaneous movement in consumption. With this restriction and the prior two, we can completely characterize

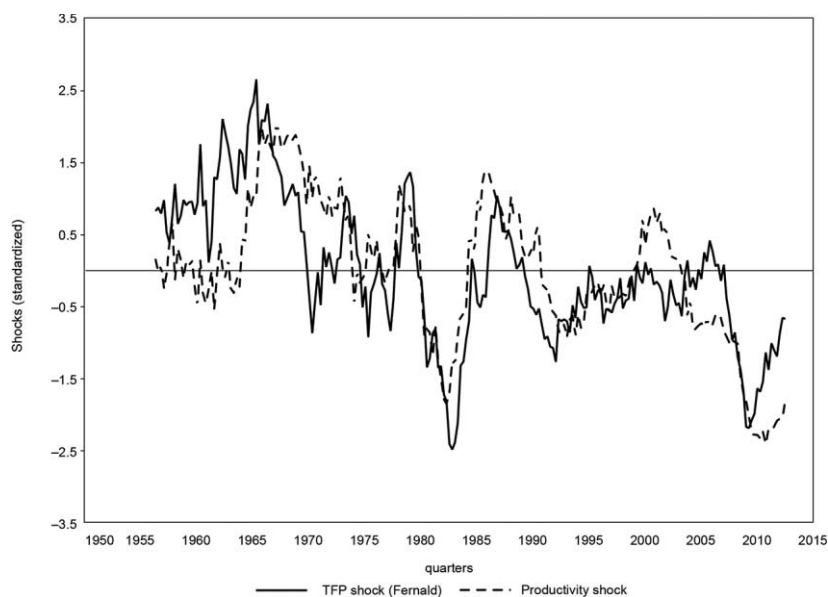


Fig. 2. Productivity shock and TFP shock

Notes: The figure shows four-year moving averages of the permanent productivity shock and TFP shocks. The TFP data are from Fernald (2009). Both series are divided by their standard deviations. The sample is 1952:Q1 to 2012:Q3.

the dynamics of consumption, labor income, and wealth as a function of three mutually orthogonal shocks.

This last orthogonalization restriction can be motivated by canonical, frictionless stochastic general equilibrium models where permanent productivity shocks determine the long-run levels of all economic variables. If there is a TFP shock in the system that affects labor income and the value of all productive capital (wealth), it should be revealed by a movement in consumption. If there is a factor shares shock in the system that is orthogonal to TFP, it must be revealed by a reallocation of rewards between y and a with no movement in consumption. Figure 2 shows that our interpreted productivity shock has cumulative effects over longer horizons that match well the fluctuations in estimated TFP provided by Fernald (2009).

The specific interpretation of these shocks as productivity, factor shares, and risk aversion shocks is formalized in a follow-up paper to this one (Greenwald, Lettau, and Ludvigson 2013). There we explic-

itly model an economy with two types of consumers (shareholders and workers) and three fundamental shocks: a permanent productivity shock that drives aggregate (shareholder plus worker) consumption, a highly persistent factor shares shock that reallocates rewards between shareholders and workers, and an exogenous shock to shareholder risk aversion that moves the stochastic discount factor pricing assets. We show that the responses of aggregate consumption, labor earnings, and asset wealth to these shocks are qualitatively very similar to the empirical responses reported here to our interpreted productivity, factor share, and risk aversion shocks.

We emphasize the following findings from this decomposition. First, the factor share shock underlies the vast bulk of quarterly fluctuations in labor income growth. This permanent disturbance cannot be explained by macroeconomic theories in which trend movements in labor income are driven solely by trend movements in factor-neutral productivity shocks that move the values of labor and productive capital in the same direction.

Second, over the last twenty-five years, the cumulative effect of the factor shares shock has persistently boosted stock market wealth and persistently lowered labor earnings. Although this disturbance has little impact on the stock market at quarterly frequencies, over long-horizons its impact on the level of the stock market is substantial and has contributed to extended periods of relatively high stock market valuation (e.g., the last twenty-five years) and relatively low stock market valuation (e.g., from the mid-1960s to mid-1980s).

Third, the transitory risk aversion shock underlies the vast bulk of quarterly fluctuations in asset values. Although transitory, these shocks are highly persistent with a half-life of over four years. We show elsewhere that these transitory fluctuations are associated with predictable variation in excess stock market returns, suggesting risk premia vary over time.¹ This finding is difficult to reconcile with modern macroeconomic models because such models typically imply constant (or nearly constant) risk premia. But this finding is also a challenge for leading asset pricing models capable of rationalizing large movements in risk premia (e.g., Campbell and Cochrane 1999). These models are consistent with the existence of a transitory component in wealth, but they cannot account for our finding that this component is unrelated to consumption.

Fourth, we evaluate the relative roles these shocks played during

the last two recessions, both accompanied by significant asset market “crashes.” We find that the “tech” crash of 2000 to 2002 and the boom that preceded it was almost entirely the result of a string of transitory risk aversion shocks, with risk aversion steadily falling in the boom years and sharply reversing at the beginning of the bust. The negative consequences for housing, consumption, and labor earnings were quite subdued in this recession, consistent with the historical pattern that these variables (especially consumption and labor income) are relatively unaffected by pure risk aversion shocks. By contrast, the “housing” asset market crash of 2007 to 2009 (also accompanied by a bust in the stock market) was characterized by large negative roles for both the transitory risk aversion shock and the permanent productivity shock. Importantly, the permanent productivity shock plays a much larger role in quarterly fluctuations of housing wealth than in stock market wealth. Thus, both consumption and housing were hard hit during the 2007 to 2009 crash, as was labor income. In contrast to the 2000 to 2002 experience, a string of negative draws for the productivity shock since 2007 is impeding the recovery in housing and weakening its medium-term outlook. At the same time, the rebound in the stock market since 2009 can be shown to be predominantly the result of a string of transitory declines in the risk aversion disturbance, an apparent reversal from the sharp increase in risk aversion during the housing bust. The cumulative effect of this latest string on stock market wealth has been sufficiently large that household net worth is now, as of 2012:Q3, farther above its long-run level than it was at the peak of the tech boom.

The rest of this paper is organized as follows. The next subsection discusses related empirical literature. Section II discusses the data and explains our econometric methodology. The objective of this methodology is to choose identifying restrictions in such a way that permits one of the two permanent structural disturbances to have effects we would expect of a permanent, factor-neutral TFP shock, and another to have effects related to shifts in factor shares. Section III presents our main findings from applying this methodology, both to the full postwar sample, and over two boom-bust subsamples that significantly affected both the real economy and asset markets during the last fifteen years. Section IV is a general discussion involving analysis of some additional data, alternatives, or refinements to our given interpretations of the identified shocks, as well as some possible weak links in our own interpretation. Section V concludes.

A. *Related Literature*

Our paper is broadly related to a series of articles that use cointegration to identify the permanent and transitory components of a system of macroeconomic variables that share at least one common stochastic trend (King et al. 1991; Cochrane 1994; Gonzalo and Granger 1995; Galí 1999; Francis and Ramey 2001; Gonzalo and Ng 2001; Lettau and Ludvigson 2004). Our study is also related to a time-honored literature that studies the sources of business cycle fluctuations as in Sims (1980) and Kydland and Prescott (1982), and more recently, Christiano, Eichenbaum, and Vigusson (2004); Fisher (2006); Smets and Wouters (2007); and Justiniano, Primiceri, and Tambalotti (2009, 2010). The study closest to the present paper is Lettau and Ludvigson (2004) (LL hereafter), who examine the trend and cyclical components of the same system of variables studied here, on earlier data. But the Lettau and Ludvigson (2004) study differs in several key ways from the present one. First, LL did not identify three mutually uncorrelated shocks in this system, instead focusing only on the space spanned by the permanent shocks. We accomplish this here by imposing additional identifying restriction on the dynamics of consumption, introduced earlier. Second, LL did not provide an economic interpretation of any of the shocks distinguished by their degree of persistence, as we do here. Third, LL did not formally relate fluctuations in the three identified disturbances to the major components of assets (stock market wealth, nonstock financial wealth, and housing) that make up household net worth.

II. **Econometric Methodology**

It is perhaps obvious that consumption, labor income, and household wealth should move together over the long term. This can be motivated more formally by considering the long-run implications of a standard household budget constraint, see Lettau and Ludvigson (2001, 2010) and LL. We refer the reader to these papers and simply note here that a cointegrating relation for log consumption, c_t , log labor income, y_t , and log asset wealth, a_t , follows from fairly weak theoretical restrictions in a broad class of models for which a household budget constraint must be obeyed.

This section first describes the data and preliminary analysis. It then describes how we isolate the permanent and transitory structural dis-

turbances of a cointegrated vector of variables, \mathbf{x}_t , that has n elements. In our application, $\mathbf{x}_t = (c_t, a_t, y_t)'$. Throughout this paper we use lower case letters to denote log variables, for example, $\ln(A_t) \equiv a_t$.

A. Data and Preliminary Analysis

The online appendix, available on the authors' web pages and at http://www.nber.org/data-appendix/c12932/scr_onlineappendix.pdf, 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. Durable goods are accounted for as part of nonhuman wealth, A_t , a component of aggregate wealth, W_t , and so are not accounted for as part of consumption.² Durables expenditures are also excluded in the definition of *flow* consumption, C_t , because they represent replacements and additions to a capital stock (investment), rather than a service flow from the existing stock. However, the total flow of consumption is unobservable because, although there is a measure of the service flow from housing, we lack observations on the service flow from the rest of the durables stock. We therefore follow Blinder and Deaton (1985) and Campbell (1987) and use the log of real, per capita expenditures on nondurables and services (excluding shoes and clothing), as a measure of c_t . From the household's budget constraint, an internally consistent cointegrating relation may then be obtained if we assume that the log of (unobservable) real total flow consumption is cointegrated with the log of real nondurables and services expenditures. The log of after-tax labor income, y_t , is also measured in real, per capita terms. Our data are quarterly and span the first quarter of 1952 to the third quarter of 2012. Table 1 presents descriptive statistics for the log differences Δc_t .

Table 1
Growth Rates

| | Consumption | Labor income | Financial net worth |
|--------------|-------------|--------------|---------------------|
| Std. dev. | 0.47% | 0.90% | 2.19% |
| Correlations | | | |
| Consumption | 1.00 | 0.31 | 0.46 |
| Labor income | 0.31 | 1.00 | 0.16 |
| Net worth | 0.46 | 0.16 | 1.00 |

Notes: The table reports descriptive statistics of log first differences of consumption, labor income, and financial net worth. The sample spans the fourth quarter of 1951 to the third quarter of 2012.

Δa_t , and Δy_t . Wealth growth is 4.7 times as volatile as consumption growth and is 2.43 times as volatile as labor income growth in quarterly data.

Let $r < n$ denote the number of cointegrating relationships in a system of n variables. The appendix presents empirical evidence supportive of a single cointegrating relationship between c_t , a_t , and y_t in quarterly postwar data. Additional tests indicate no evidence of a second cointegrating relationship (see discussion following). Given our system of $n = 3$ variables, this implies the presence of $n - r = 2$ permanent innovations and $r = 1$ transitory innovations (Stock and Watson 1988). These innovations are the structural disturbances we seek to identify.

Although statistical tests are supportive of a single trivariate cointegrating relation between c_t , a_t , and y_t , the data provide no evidence of a second linearly independent cointegrating relation (there can be at most two). In particular, bivariate log ratios of these variables appear to contain trends in our sample. Economic models with balanced growth would imply that bivariate log ratios (e.g., $y_t - a_t$) are stationary. Nevertheless, there is no evidence in our sample of cointegration for any bivariate ratio pair in our system. We therefore follow the advice of Campbell and Perron (1991) and empirically model only the single, trivariate cointegrating relation for which we find direct statistical evidence of in our sample. Campbell and Perron argue that treating the data in accordance with the stationarity properties inferred from unit root/cointegration tests results in better finite sample approximations of test statistics than does treating the data according to its asymptotic distribution that is true in population. Thus, a near-integrated but stationary data generating process is better modeled in a finite sample as a unit root variable, even though the asymptotically correct distribution is the standard one appropriate for stationary variables.

Identification of the structural disturbances is achieved in three steps. For the first step, we use restrictions implied by cointegration to identify structural disturbances distinguished by whether their effects are permanent or transitory. The procedure follows Gonzalo and Ng (2001) and is closely related to that in King et al. (1991) and Gonzalo and Granger (1995). This procedure itself has several steps, the first of which requires estimation of the cointegrating relationship(s) and the vector-error-correction model (VECM) for the cointegrated system.

We assume all of the series contained in x_t are first-order integrated, or $I(1)$, an assumption confirmed by unit root tests, available upon request. The cointegrating coefficient on consumption is normalized to

one, and we denote the single cointegrating vector for $\mathbf{x}_t = [c_t, a_t, y_t]'$ as $\boldsymbol{\alpha} = (1, -\alpha_a, -\alpha_y)'$.

The cointegrating parameters α_a and α_y are estimated using dynamic least squares, which generates “superconsistent” estimates of α_a and α_y (Stock and Watson 1993).³ We estimate $\hat{\boldsymbol{\alpha}} = (1, -0.18, -0.70)'$. The Newey and West (1987) corrected t -statistics for these estimates are 20 and 56, respectively.

The VECM representation of \mathbf{x}_t takes the form

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

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, $\boldsymbol{\Gamma}(L)$ 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.⁴ The term $\hat{\boldsymbol{\alpha}}' \mathbf{x}_{t-1}$ gives last period’s equilibrium error, or cointegrating residual, a variable we denote with $ca y_t \equiv \hat{\boldsymbol{\alpha}}' \mathbf{x}_{t-1}$. The coefficients $\boldsymbol{\gamma}$ are the vector of “adjustment” coefficients that tells us which variables subsequently adjust to restore the common trend when a deviation occurs. Throughout this paper, we use “hats” to denote the estimated values of parameters.

The results of estimating a first-order specification of (1) are presented in table 2.⁵ The estimates of the adjustment parameters in $\boldsymbol{\gamma}$ are given in the first row of table 2. An important result is that although

Table 2
VECM

| Dependent variable | Equation | | |
|--------------------|-----------------------|-----------------------|-----------------------|
| | Δc_t | Δa_t | Δy_t |
| $ca y_{t-1}$ | 0.01 (0.94) | 0.20 (2.33) | 0.05 (1.71) |
| Δc_{t-1} | 0.32 (4.82) | 0.25 (0.73) | 0.52 (3.74) |
| Δy_{t-1} | 2.00 (1.67) | -0.14 (-0.78) | -0.08 (-1.18) |
| Δa_{t-1} | 0.05 (3.92) | 0.22 (3.26) | 0.4 (1.35) |
| \bar{R}^2 | 0.26 | 0.07 | 0.08 |

Notes: The table reports the estimated coefficients from cointegrated vector autoregressions (VECM) of the column variable on the row variable; t -statistics are in parentheses. Estimated coefficients that are significant at the 5% level are highlighted in boldface. The term $c_t - \hat{\alpha}_a a_t - \hat{\alpha}_y y_t = \hat{\boldsymbol{\alpha}}' \mathbf{x}_t$ is the estimated cointegrating residual. The sample spans the fourth quarter of 1951 to the third quarter of 2012.

consumption and labor income are somewhat predictable by lagged consumption and wealth growth, they are not predictable by the cointegrating residual $\hat{\alpha}'\mathbf{x}_{t-1}$. Estimates of γ_c and γ_y are economically small and insignificantly different from zero. By contrast, the cointegrating error $ca y_t$ 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 2.3.⁶ 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.

B. Identification of Permanent and Transitory Shocks

The general identification problem is described as follows. The individual series involved in the cointegrating relation are presumed to have a reduced-form multivariate Wold representation:

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

where \mathbf{e}_t is an $n \times 1$ vector of innovations, and where $\mathbf{C}(L) \equiv \mathbf{I} + \mathbf{C}_1L + \mathbf{C}_2L^2 + \mathbf{C}_3L^3 + \dots$. The parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}$, both of rank r , satisfy $\boldsymbol{\alpha}'\mathbf{C}(1) = 0$ and $\mathbf{C}(1)\boldsymbol{\gamma} = 0$ (Engle and Granger 1987).

The "reduced-form" disturbances \mathbf{e}_t have no particular interpretation. We seek to identify $n = 3$ transformed, or structural-form, innovations distinguished by whether they have permanent or transitory effects. Denote these transformed innovations $\boldsymbol{\eta}_t \equiv (\eta_{p1,t}, \eta_{p2,t}, \eta_{T,t})'$, where two are permanent and one is transitory. Without loss of generality, shocks are ordered so that the first two have permanent effects ($\eta_{p1,t}$ and $\eta_{p2,t}$, respectively), and the third transitory effects ($\eta_{T,t}$).

In the decomposition that follows, a shock is defined to be permanent if

$$\lim_{h \rightarrow \infty} \partial E_t(\mathbf{x}_{t+h}) / \partial \eta_{pt} \neq 0. \quad (3)$$

Conversely, a shock is transitory if

$$\lim_{h \rightarrow \infty} \partial E_t(\mathbf{x}_{t+h}) / \partial \eta_{Tt} = 0. \quad (4)$$

Notice that a permanent shock under this definition differs from the long-run trend component obtained from a Beveridge-Nelson decomposition, in that here a permanent shock may contain serially correlated noise around the random walk component. Regardless of which way the permanent-transitory decomposition is defined, permanent shocks will have the same long-run effects on the variables. Moreover,

it is straightforward to decompose movements in each of the variables into components that deviate from their random walk components (see below).

Let

$$\mathbf{G} \equiv \begin{bmatrix} \boldsymbol{\gamma}'_{\perp} \\ \boldsymbol{\alpha}' \end{bmatrix}, \quad (5)$$

where $\boldsymbol{\gamma}'_{\perp}$ is a matrix of rank $n - r$ that satisfies⁷

$$\underset{(n-r) \times n}{\boldsymbol{\gamma}'_{\perp}} \underset{n \times r}{\boldsymbol{\gamma}} = \underset{(n-r) \times r}{\mathbf{0}}. \quad (6)$$

Define a new distributed lag operator

$$\mathbf{D}(L) = \mathbf{C}(L)\mathbf{G}^{-1}.$$

The structural (permanent and transitory) disturbances are given by $\boldsymbol{\eta}_t = (\eta_{P1,t}, \eta_{P2,t}, \eta_{T,t})$, where

$$\boldsymbol{\eta}_t = \mathbf{G}\mathbf{e}_t,$$

and their relation to \mathbf{x}_t is given by the Wold representation

$$\begin{aligned} \Delta \mathbf{x}_t &= \boldsymbol{\delta} + \mathbf{C}(L)\mathbf{G}^{-1}\mathbf{G}\mathbf{e}_t \\ &= \boldsymbol{\delta} + \mathbf{D}(L)\boldsymbol{\eta}_t, \end{aligned} \quad (7)$$

where $\boldsymbol{\delta}$ is a constant vector. Let $D_{ij}(L)$ denote the i, j th element of $\mathbf{D}(L)$. Through the identification of $\boldsymbol{\gamma}_{\perp}$, this decomposition imposes the following restriction on the long-run multipliers of the structural-form shocks:

$$D_{13}(1) = D_{23}(1) = D_{33}(1) = 0. \quad (8)$$

The restriction follows because the last r columns of the polynomial matrix $\mathbf{D}(L)$ are responses of $\Delta \mathbf{x}_t$ to transitory shocks, and by assumption have no influence on the variables in the long run.

This decomposition can be understood intuitively by noting that it gives the j th variable a large weight in the permanent innovations and a small weight in the transitory innovations when γ_j is small (via computation of $\boldsymbol{\gamma}'_{\perp}$). In this case, the j th variable participates little in the error-correction required to restore the series to their identified common trend, implying that it displays only small deviations from this common trend. Conversely, it gives the j th variable a small weight in the permanent innovations and a large weight in the transitory innova-

tions when γ_j is large, implying that it plays an important role in the error-correction required to restore the series to their identified common trend. In the application studied here, the elements of the adjustment vector $\boldsymbol{\gamma}$ corresponding to c_t and y_t are statistically indistinguishable from zero (table 2), implying that these variables have a large weight in the permanent innovations and a small weight in the transitory innovations. By contrast, the element of the adjustment vector $\boldsymbol{\gamma}$ corresponding to a_t is large in absolute value and strongly statistically significant, implying that a_t will have a large weight in the transitory innovations and a small weight in the permanent innovations.

The Gonzalo and Ng (2001) procedure involves restricting the values of the parameters in $\boldsymbol{\gamma}$ to zero where they are statistically insignificant at the 5 percent level: failure to do so can result in unreliable estimates of the permanent-transitory decomposition. This restriction is also imposed in other applications of this methodology (Cochrane 1994; Gonzalo and Granger 1995). In the computations that follow, we set γ_c and γ_y to zero in order to match the evidence from table 2 that these variables are small and statistically indistinguishable from zero.

The cointegration restrictions applied so far are enough to identify permanent and transitory innovations, but these innovations need not be mutually uncorrelated. To identify shocks that are mutually uncorrelated, we apply a rotation to the vector of transformed shocks $\boldsymbol{\eta}_t$. Specifically, let \mathbf{H} be a lower triangular matrix that accomplishes the Cholesky decomposition of $\text{Cov}(\boldsymbol{\eta}_t)$, and define a set of orthogonal structural disturbances $\tilde{\boldsymbol{\eta}}$ such that

$$\tilde{\boldsymbol{\eta}} \equiv \mathbf{H}^{-1}\boldsymbol{\eta}_t = \mathbf{H}^{-1}\mathbf{G}\mathbf{e}_t.$$

Also define

$$\begin{aligned}\tilde{\mathbf{D}}(L) &\equiv \mathbf{C}(L)\mathbf{G}^{-1}\mathbf{H} \\ &= \mathbf{D}(L)\mathbf{H}.\end{aligned}$$

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

$$\Delta\mathbf{x}_t = \boldsymbol{\delta} + \tilde{\mathbf{D}}(L)\tilde{\boldsymbol{\eta}}_t, \quad (9)$$

which now yields a vector of mutually uncorrelated permanent and transitory innovations $\tilde{\boldsymbol{\eta}}_t$.

We make two identification assumptions that justify the particular rotation \mathbf{H}^{-1} chosen.

First, we restrict the space spanned by the two permanent shocks to

be orthogonal to the transitory shock. As discussed above, this restriction may be thought of as a definition rather than an assumption. We *define* the transitory shock to be orthogonal to the space spanned by the permanent shocks, thereby allowing us to identify their independent effects. If in fact there were no transitory component in the system that had effects independent of the permanent disturbances, the econometric procedure would assign a zero role for this component in the variance decompositions of all variables in the system. This restriction, along with the restrictions implied by cointegration, allows us to completely identify the space spanned by the two permanent shocks and the single transitory shock.

The final identifying restriction allows us to distinguish the independent effects of the two permanent disturbances. For this, we seek an identification that allows us to recover *both* a TFP shock and a shock related to shifts in factor shares. Figure 1 shows that labor's share of output has fluctuated considerably over time in a very persistent manner. Indeed, it has plummeted at the end of our sample and is at a record low since data have been collected.

Table 3 presents results from regressions of labor income growth, Δy_t , on the growth in TFP as measured by Fernald (2009), ΔTFP_t , and two measures of labor share growth, one simply defined as employee compensation divided by compensation plus profits, ΔLS_t^{CP} , and one measured by the Bureau of Economic Analysis' Bureau of Labor Statistics (BLS), ΔLS_t^{BLS} . By itself, ΔTFP_t explains a large fraction of the variation in Δy_t (adjusted $R^2 = 21\%$), but measures of labor share have significant additional explanatory power, particularly the BLS measure: the coefficient on ΔLS_t^{BLS} is about half as large as that on ΔTFP_t , is strongly statistically significant (t -stat = 5.6), and explains an additional 6 percent of the quarterly variation in labor income growth. Although these measures are all proxies for the corresponding theoretical concepts of TFP and labor share, these results are suggestive of an important role for each in driving labor income. Controlling for a lagged value of the dependent variable, Δy_{t-1} , has no impact on these results.

Thus we aim to recover a second permanent shock in the system that is related to shifts in factor shares. The only interpretable orthogonalization capable of recovering both a TFP and factor shares shock requires us restrict c to be ordered first in the cointegrated vector autoregression (VAR) before performing a standard Cholesky orthogonalization. The shock that is then interpreted as a permanent productivity shock (the first permanent shock) will affect all variables in the system contem-

Table 3
Labor income growth regressions

| | | Dependent variable: Δy_t | | | |
|-----------------------|------------------|----------------------------------|--------------------|---------------------|-------------|
| Independent variables | Δy_{t-1} | ΔTFP_t | ΔLS_t^{CP} | ΔLS_t^{BLS} | \bar{R}^2 |
| 1 | | 0.505 | | | 0.21 |
| (<i>t</i> -stat.) | | (7.319) | | | |
| 2 | | 0.585 | 0.182 | | 0.23 |
| (<i>t</i> -stat.) | | (8.251) | (3.023) | | |
| 3 | | 0.670 | | 0.306 | 0.27 |
| (<i>t</i> -stat.) | | (9.403) | | (5.624) | |
| 4 | -0.066 | 0.521 | | | 0.21 |
| (<i>t</i> -stat.) | (-0.641) | (7.054) | | | |
| 5 | -0.052 | 0.596 | 0.177 | | 0.23 |
| (<i>t</i> -stat.) | (-0.532) | (7.917) | (2.756) | | |
| 6 | -0.083 | 0.694 | | 0.313 | 0.28 |
| (<i>t</i> -stat.) | (-0.852) | (9.265) | | (5.824) | |

Notes: *t*-statistics are obtained using Newey-West standard errors at 4 lags. Δy_t is real labor income growth obtained from the CAY data set (source: Martin Lettau and Sydney Ludvigson). ΔTFP_t is TFP growth (source: FRBSF/Fernald). ΔLS_t^{CP} is log labor share growth, where labor share is measured as (compensation)/(compensation + profit). Compensation is "Compensation of Employees: Wage and Salary Accruals" obtained from the BEA via FRED (series: WASCUR). Profit is "Corporate Profits After Tax" obtained from the BEA via FRED (series: CP). ΔLS_t^{BLS} is log labor share growth, where labor share is the series "Nonfarm Business Sector: Labor Share" obtained from the BLS via FRED (series: PRS85006173). The sample spans the period 1952:Q3 to 2012:Q2.

poraneously and be revealed as a consumption shock, implying that the second permanent shock will be revealed by a movement in labor income and wealth, with no contemporaneous movement in consumption. We argue that this second permanent shock can be interpreted as a factor share shock. With this restriction and the former two, we completely identify three mutually orthogonal shocks: two with permanent effects on the variables in the system, and one with transitory effects.

This orthogonalization can be motivated by canonical, frictionless stochastic general equilibrium models where permanent productivity shocks drive the long-run movements in the economy. If there is a TFP shock in the system that affects labor income and the value of all productive capital (wealth), it should be revealed by a movement in consumption. Thus the a factor-neutral TFP shock must contemporaneously affect consumption, as well as labor income and wealth. In addition, typically in these models consumption is not myopic and responds only to the permanent component of total (labor plus capital) income driven by the TFP shock. So if there is a factor shares shock in the sys-

tem that results in a mere reallocation of rewards between y and a with no affect on total income, it should not affect consumption. Moreover, without habits or other frictions and if the elasticity of intertemporal substitution is very small, c will be very close to a random walk if log TFP is a random walk.

As an example, figure 3 shows the dynamic responses to a permanent (log) TFP shock in a canonical real business cycle model, for three different values of the elasticity of intertemporal substitution in consumption (EIS), equal to 0.5, 0.2, and 0.1.⁸ As the EIS falls, the response of consumption is less sluggish and, for very small values of the EIS, consumption is close to a random walk. It is straightforward to show that, when the EIS $\sigma = 0$, log consumption is exactly a random walk and log labor income and log capital follow unit root processes cointegrated with consumption. Moreover, when $\sigma = 0$, the model is a general equilibrium version of the permanent income model of Hall (1978) and Flavin (1981) in which log consumption is exactly a random walk and equals “permanent income,” given by the annuity value of wealth plus the present discounted value of all future labor income.⁹ Cochrane (1994) used this reasoning to argue that consumption defined the trend in GNP, allowing the identification of a large transitory component in the latter. We can assess how well these theoretical features are captured in our data by investigating the empirical impulse responses to our interpreted TFP shocks.

An alternative to the orthogonalization just described would restrict labor income to be ordered prior to consumption in the cointegrated VAR.¹⁰ In this case the TFP shock that contemporaneously raises consumption, labor income, and the value of all productive capital would be revealed as a labor income shock, rather than a consumption shock. The difficulty with this orthogonalization is that the second permanent shock in the system, which is by construction orthogonal to the first, would have no obvious interpretation as a factor shares shock. It would contemporaneously influence consumption and wealth while leaving labor income contemporaneously unaffected.¹¹ Furthermore, with TFP shocks identified as labor income shocks, the empirical model would have no way of accounting for evidence that movements in labor’s share are substantial and have an important effect on labor income growth even after controlling for measures of TFP growth (table 3). Our aim is to find an interpretable orthogonalization that *does* recover both a TFP and factor share shocks. The y -first approach would provide no interpretable way to do so, since labor income shocks cannot both iden-

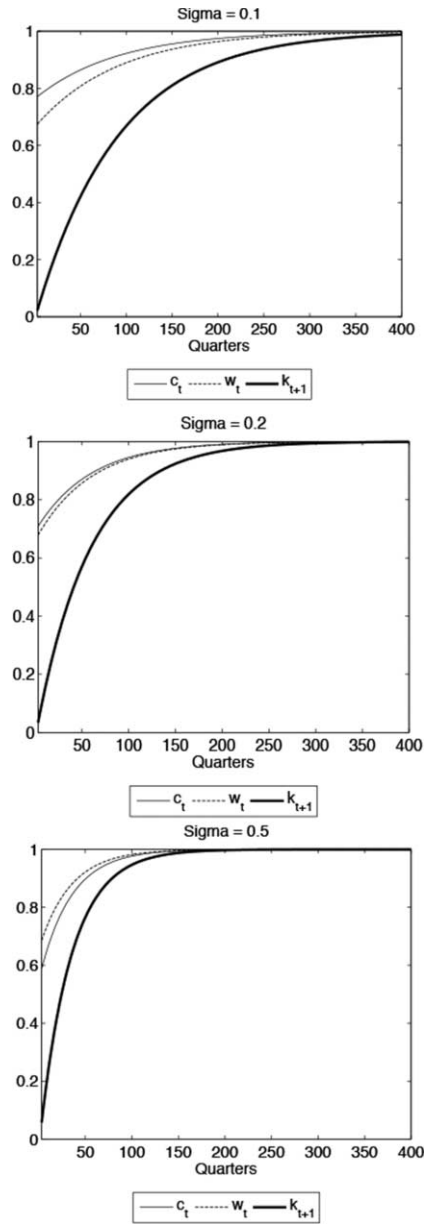


Fig. 3. Impulse response functions in real business cycle (RBC) model

Notes: Percentage response of consumption, capital, and labor income, corresponding to a 1 percent technology shock in a model with fixed labor supply (normalized to one), specified in the following equations: $Y_t = A_t^\alpha K_t^{1-\alpha}$; $K_{t+1} = (1 - \delta)K_t Y_t - C_t$; $C_t^{-\gamma} = \beta E_t[C_{t+1}^{-\gamma} R_{t+1}]$; $a_t = \phi a_{t-1} + \epsilon_t$. Lowercase letters denote log levels. w_t is log labor income (equal to real wage since labor supply is fixed at unity), c_t is log real consumption, and k_t is log capital. The parameter values are set as follows: $\phi = 1$, $r = 0.015$, $g = 0.005$, $\alpha = 0.667$, $\delta = 0.025$, and $\sigma = 1/\gamma$.

tify TFP shocks and be contemporaneously affected by factor shares shocks orthogonal to TFP. The responses to the y -first shocks would be linear combinations of the c -first responses, and therefore mixtures of interpretable TFP and factor shares responses. For this reason, we do not pursue such an orthogonalization.

A word about the interpretation of the transitory shock is in order. We interpret the transitory shock in this system to be an exogenous risk aversion shock, independent of TFP and factor shares shocks. We show elsewhere that these transitory fluctuations are associated with predictable variation in excess stock market returns, suggesting risk premia vary over time (Lettau and Ludvigson 2001, 2004, 2010). This component bears no relation to current or future measures of real activity such as consumption, labor income, dividend, or earnings growth. Thus, we conclude that this shock must be a shock to the *stochastic discount factor*, rather than to cash flows. Greenwald, Lettau, and Ludvigson (2013) model shareholder preferences so that risk aversion varies over time in a manner that is independent of these measures of real activity. The shocks to risk aversion affect asset values because they affect the rate at which future cash flows are discounted. We show there that the response of c , a , and y to a risk aversion shock in the model is very similar to the empirical responses reported below to our interpreted risk aversion shock. Modeling risk aversion as an endogenous response to, for example, TFP shocks will not work because in that model the transitory shock that is orthogonal to the permanent shocks would then (counterfactually) play little role in the dynamics of wealth.

C. *Relating Structural Disturbances to Wealth Components*

The econometric procedure just described is applied to the system of c_t , a_t , and y_t . But we also seek to study the role of each shock in the dynamic behavior of three major components of household assets a_t : stock market wealth, non-stock market financial wealth, and housing. The category referred to as “non-stock financial” wealth includes all financial wealth outside of the stock market. So that the three components sum up to total assets, we also include nonhousing tangible assets in this category, which are a small component comprising only 10 percent of the category. For brevity, we simply refer to this component as non-stock financial wealth. On average over the period spanning the fourth quarter of 1951 to the second quarter of 2010, assets accounted for of 1.16 of net worth while liabilities accounted for 0.16. Stock market

Table 4
Wealth components

| | Net worth | Assets | Liabilities | Stock wealth | Housing | Non-stock fin. |
|-----------------------------------------|-----------|--------|-------------|--------------|---------|----------------|
| <i>Share in net worth</i> | | | | | | |
| 1951:Q4–2012:Q2 | 1.00 | 1.16 | 0.16 | 0.22 | 0.29 | 0.65 |
| 1951:Q4–1961:Q4 | 1.00 | 1.10 | 0.10 | 0.18 | 0.25 | 0.67 |
| 1999:Q3–2012:Q2 | 1.00 | 1.22 | 0.22 | 0.29 | 0.32 | 0.61 |
| Std. dev. | 2.19 | 1.92 | 1.11 | 8.84 | 1.65 | 0.73 |
| <i>Correlations of log growth rates</i> | | | | | | |
| Net worth | 1.00 | 1.00 | 0.30 | 0.90 | 0.43 | 0.32 |
| Assets | 1.00 | 1.00 | 0.37 | 0.89 | 0.45 | 0.35 |
| Liabilities | 0.30 | 0.37 | 1.00 | 0.18 | 0.51 | 0.47 |
| Stock mkt. wealth | 0.90 | 0.89 | 0.18 | 1.00 | 0.15 | 0.04 |
| Housing | 0.43 | 0.45 | 0.51 | 0.15 | 1.00 | 0.49 |
| Non-stock financial | 0.32 | 0.35 | 0.47 | 0.04 | 0.49 | 1.00 |

Notes: The table reports descriptive statistics of wealth components from the Flow of Funds. The sample spans the fourth quarter of 1951 to the third quarter of 2012.

wealth accounted for 22 percent of net worth, housing wealth 29 percent, and non-stock financial wealth 68 percent.

Table 4 summarizes the statistics for these wealth components. Stock market wealth is by far the most volatile component: the annualized standard deviation of the log difference in stock wealth is 8.84 percent. By contrast, housing wealth growth has a standard deviation of 1.65 percent and non-stock financial wealth growth just 0.73 percent. This shows that all the “action” is in stock wealth and housing. Moreover, the correlation of the log difference in net worth, $\Delta a_{i,t}$, with the log difference in stock market wealth is 90 percent. Quarterly changes in net worth are dominated by fluctuations in the stock market.

To relate these wealth components to the disturbances $\eta_{p1,t}$ (permanent productivity shock), $\eta_{p2,t}$ (permanent factor shares shock), and $\eta_{T,t}$ (transitory risk aversion shock), we estimate empirical relationships taking the form

$$\Delta z_{i,t} = A_i(L)\eta_{p1,t} + B_i(L)\eta_{p2,t} + C_i(L)\eta_{T,t} + \epsilon_{i,t}, \quad (10)$$

where $z_{i,t}$ represents the log level of the i th component of net worth (e.g., stock market wealth, non-stock market wealth, housing) and $A_i(L)$, $B_i(L)$, and $C_i(L)$ are polynomial lag operators.¹² Since $\eta_{p1,t}$, $\eta_{p2,t}$, and $\eta_{T,t}$ are mutually uncorrelated and i.i.d., we estimate these equations for each component separately by OLS (ordinary least squares) with $L = 16$

quarters. Because the disturbances $\eta_{p1,t}$, $\eta_{p2,t}$ and $\eta_{T,t}$ account for 100 percent of the variation in the log of net worth, a_t , the residuals $\epsilon_{i,t}$ are, by construction, shocks to wealth components that are orthogonal to a_t . Hence a positive innovation in one component must be met with a negative innovation in another component. Note, however, that the residuals $\epsilon_{i,t}$ also include log/level errors since the sums of the logs of the components do not equal the log of the sum of components.

D. Decomposition of Levels

To shed light on the role that each shock has played on the evolution of the levels of the variables over time, we decompose the log levels into components driven by each structural disturbance. To do so, consider the decomposition of growth rates of wealth components in (10):

$$\begin{aligned}\Delta z_{i,t} &= A_i(L)\eta_{p1,t} + B_i(L)\eta_{p2,t} + C_i(L)\eta_{T,t} + \epsilon_{i,t} \\ &\equiv \Delta z_{i,t}^{P1} + \Delta z_{i,t}^{P2} + \Delta z_{i,t}^T + \epsilon_{i,t}.\end{aligned}$$

The effect on the log levels of the variables of each disturbance is obtained by summing up the effects on the log differences (where below we drop the i subscript to denote the generic approach):¹³

$$\begin{aligned}z_t &= z_0 + \sum_{s=1}^t \Delta z_s \\ &= z_0 + \sum_{s=1}^t \Delta z_s^{P1} + \sum_{s=1}^t \Delta z_s^{P2} + \sum_{s=1}^t \Delta z_s^T + \epsilon_t \\ &= z_0 + z_t^{P1} + z_t^{P2} + z_t^T + \epsilon_t,\end{aligned}\tag{11}$$

where T is the sample size.

A similar decomposition of levels can be obtained for c_t , a_t and y_t by referring to the relevant submatrices of (9):

$$\begin{aligned}\Delta x_{i,t} &= \tilde{D}_{i,1}(L)\eta_{p1,t} + \tilde{D}_{i,2}(L)\eta_{p2,t} + \tilde{D}_{i,3}(L)\eta_{T,t} \\ &\equiv \Delta x_{i,t}^{P1} + \Delta x_{i,t}^{P2} + \Delta x_{i,t}^T,\end{aligned}$$

where $x_{i,t}$ is the i element of $\mathbf{x}_t = (c_t, a_t, y_t)'$, and where $\tilde{D}_{i,j}(L)$ denotes the scalar polynomial lag operator that is the i, j th element of $\tilde{\mathbf{D}}(L)$. There is no residual in the above because, by construction, the two permanent and one transitory shock account for all of the variance in $\Delta \mathbf{x}_t$. Summing

the first differences we again obtain the effect of each shock on the log levels:

$$\begin{aligned}
 x_{i,t} &= x_{i,0} + \sum_{s=1}^t \Delta x_{i,s} \\
 &= x_{i,0} + \sum_{s=1}^t \Delta x_{i,s}^{P1} + \sum_{s=1}^t \Delta x_{i,s}^{P2} + \sum_{s=1}^t \Delta x_{i,s}^T. \quad (12) \\
 &= x_{i,0} + x_{i,t}^{P1} + x_{i,t}^{P2} + x_{i,t}^T.
 \end{aligned}$$

III. Empirical Results

A. *Permanent and Transitory Components of Consumption, Labor Earnings, and Wealth*

Using the permanent-transitory decomposition discussed earlier, we now investigate how each of the variables in our system are related to permanent and transitory shocks.

To characterize the dynamic impact of the structural disturbances, figure 4 shows the cumulative impulse responses of Δc_t , Δa_t , and Δy_t , to a one-standard deviation innovation in each structural disturbance $\eta_{P1,t}$ (productivity), $\eta_{P2,t}$ (factor shares), and $\eta_{T,t}$ (risk aversion). Confidence intervals for these responses are presented in the online appendix. The top panel shows that a positive innovation in productivity, $\eta_{P1,t}$, leads to an immediate increase in c_t , a_t , and y_t . All three variables reach a new, higher long-run level within a few quarters in response to this shock. The second panel of figure 4 displays the empirical responses of c_t , a_t , and y_t to the factor shares shock, $\eta_{P2,t}$. Consumption is unaffected by this shock, both on impact (by assumption) and in all future periods (a result). Instead, this shock drives a_t and y_t in opposite directions. A positive value for this shock raises asset wealth a_t and lowers labor income y_t , both of which are moved to new long-run levels. The effect on labor earnings is large and immediate: labor income jumps to a new lower level within the quarter. Below we present evidence that changes in a_t resulting from this shock are predominantly driven by the stock market. The third panel of figure 4 shows that a positive transitory shock (negative risk aversion shock) leads to a sharp increase in asset wealth, but has virtually no impact on consumption and labor earnings at any future horizon. The consumption and labor income responses are economically negligible. By contrast, the effect of a risk aversion shock on a_t is strongly significant over periods from a quarter to several years,

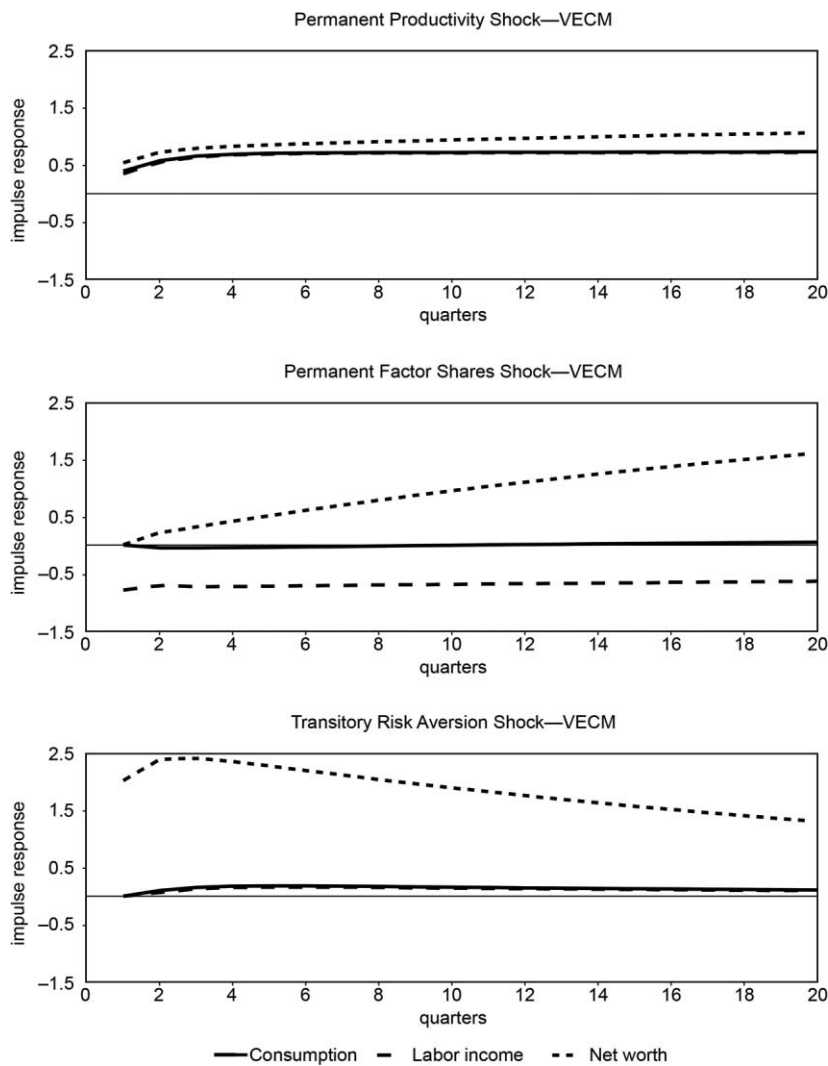


Fig. 4. Impulse response functions

Notes: The figure plots impulse response functions. The sample is 1952:Q1 to 2012:Q3.

but is eventually eliminated, as it must be, since the shock is transitory. This shock is quite persistent, however, having a half-life of over four years. Despite their persistent effect on asset values, such shocks bear virtually no relation to consumption at any future horizon.

To get a sense for how these shocks have contributed to fluctuations over time, figure 5 plots the cumulative sum over time of each

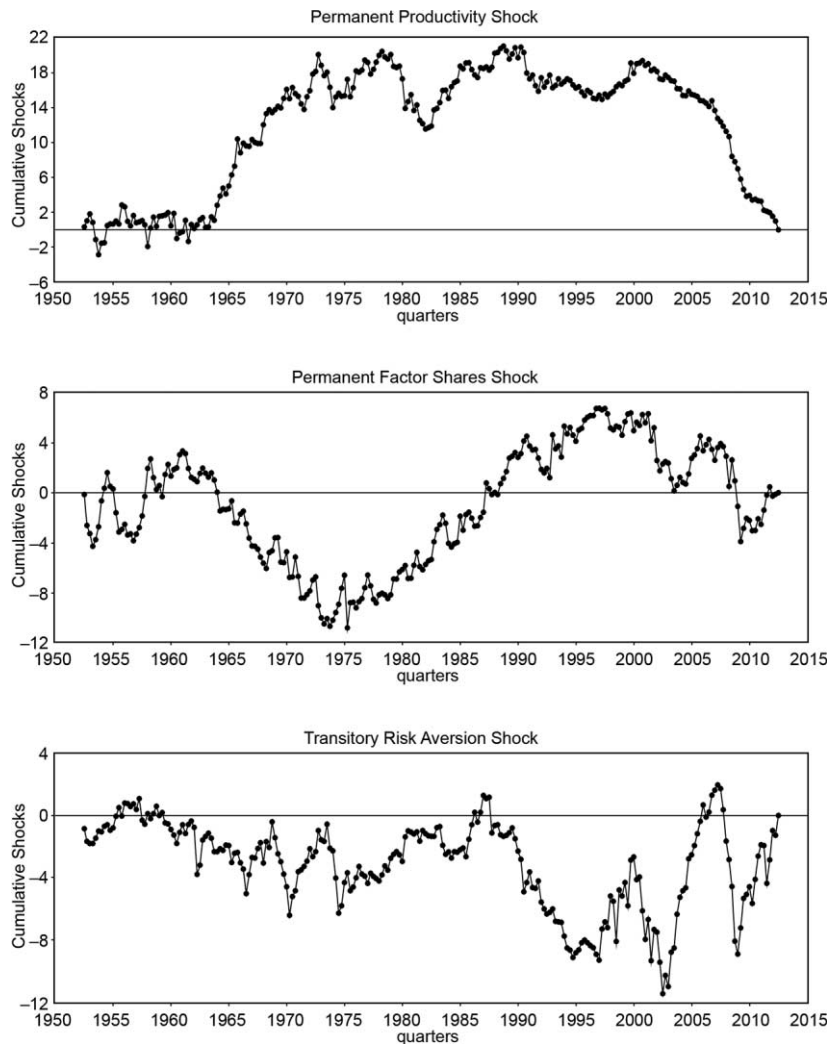


Fig. 5. Cumulative shocks

Notes: The figure plots cumulative permanent and transitory shocks identified by the PT decomposition. The sample is 1952:Q1 to 2012:Q3.

structural disturbance after having removed a deterministic trend. The permanent productivity shock (top panel) was close to average (zero) until the mid-1960s, turning positive from 1965 to 1973 and then negative again from the late 1970s to the early 1980s. The cumulative sum reaches a peak around the year 2000 and begins to decline thereafter,

falling sharply in 2007 and continuing its downward trajectory to the end of our sample, 2012:Q3.

The cumulative sum of the factor shares shocks are shown in panel 2. A positive value for this shock lowers labor income and raises asset wealth. This shock was close to average from the mid-1950s to mid-1960s, then followed by a string of negative shocks from the mid-1960s to the mid-1980s, and a string of positive shocks over the last twenty-five years. Over the last twenty-five years, the cumulative effect of the factor shares shock has persistently boosted wealth and persistently lowered labor earnings.

Finally, the cumulative sum of transitory T shocks over time shows notable above average values leading up to the peak of several asset market booms, also discussed below. This shock appears to have become more volatile since 1998.

How quantitatively important are these shocks? Table 5 displays the fraction of h -step ahead forecast error variance in the log difference of consumption, labor income, and wealth that is attributable to each shock for $h = 1$ and for $h \rightarrow \infty$, with the latter giving the portion of the variance of each variable attributable to each disturbance. To quantify the sampling uncertainty of the variance decompositions, we compute cumulative distribution functions for each variance decomposition using a bootstrapping procedure described in the online appendix (http://www.nber.org/data-appendix/c12932/scr_onlineappendix.pdf).

Table 5
Variance decomposition

| Variable | Prod. shock (%) | Fact. shares shock (%) | Risk aversion shock (%) |
|-------------------------------------------------|-----------------|------------------------|-------------------------|
| $\Delta c_{t+1} - E_t \Delta c_{t+1}$ | 100 | 0 | 0 |
| | (100, 100) | (0, 0) | (0, 0) |
| $\Delta y_{t+1} - E_t \Delta y_{t+1}$ | 16 | 84 | 0 |
| | (11, 23) | (77, 89) | (0, 0) |
| $\Delta a_{t+1} - E_t \Delta a_{t+1}$ | 7 | 0 | 93 |
| | (4, 11) | (0, 1) | (88, 96) |
| $\Delta c_{t+\infty} - E_t \Delta c_{t+\infty}$ | 93 | 1 | 6 |
| | (87, 96) | (0, 4) | (3, 11) |
| $\Delta y_{t+\infty} - E_t \Delta y_{t+\infty}$ | 22 | 77 | 1 |
| | (16, 28) | (70, 83) | (0, 3) |
| $\Delta a_{t+\infty} - E_t \Delta a_{t+\infty}$ | 7 | 3 | 90 |
| | (5, 14) | (1, 10) | (80, 91) |

Notes: The table reports the variance decomposition of consumption, labor income, and net worth. Bootstrapped 90 percent confidence intervals are in parentheses. The sample spans the fourth quarter of 1951 to the third quarter of 2012.

As table 5 shows, the permanent productivity shock explains 93 percent of the variance in the forecast error of consumption growth at long-horizons. Only 6 percent of the variation in consumption growth is attributable to the transitory shock. Simply put, *quarterly variation in consumption growth is dominated by productivity shocks.*

By contrast, the permanent factor shares shock explains 77 percent of the variance of labor income growth in our sample. The productivity shock explains 22 percent. Together, the two permanent shocks account for 99 percent of the variation in the long-run forecast error of Δy_t . This shows that consumption growth and labor income growth are dominated by permanent shocks—but they are not dominated by the *same* permanent shock. The productivity shock explains a small fraction of quarterly variation in labor earnings. *Quarterly variation in labor income growth is dominated by the factor shares shock.*

This finding is especially puzzling for canonical macroeconomic models where permanent shocks to labor earnings are driven by permanent technology shocks that move the value of labor and productive capital in the same direction. Even in models where nontechnology shocks (e.g., preference shocks, fiscal shocks, monetary shocks) play a role, they typically have only a temporary impact on the economy and do not typically reallocate rewards among factors of production.

The results are quite different for asset wealth: 90 percent of the quarterly variation in the growth of asset wealth is attributable to the transitory shock; only 10 percent is attributable to permanent shocks. Because this shock bears virtually no relation with Δc_t and Δy_t , the transitory shock in this system is in effect, a wealth shock, orthogonal to these real quantities. This result can be understood intuitively by observing that, since consumption, wealth, and labor income are cointegrated, their annualized growth rates must be tied together in the very long run, and therefore so must their volatilities. Measured over quarterly horizons, however, wealth growth is far more volatile than both consumption and labor income growth (table 1). The short- and long-run properties of these variables can only be reconciled if either (a) the annualized volatility of consumption and/or labor income growth increases with the horizon over which they are measured, or (b) the annualized volatility of wealth growth decreases with the horizon over which it is measured. The second possibility implies that wealth is not a random walk, but instead displays mean-reversion and adjusts over long-horizons to match the smoothness of consumption and labor income. The evidence in table 5 suggests that the second possibility better describes

Table 6
Correlation with random walk component

| Consumption | Labor income | Financial net worth |
|-------------|--------------|---------------------|
| 0.97 | 0.99 | 0.31 |

Notes: This table reports the correlations of growth rates with the growth rates of the nonstationary random walk components constructed from the permanent/transitory identification of shocks. The sample spans the fourth quarter of 1951 to the third quarter of 2012.

US data than the first, signaling the existence of a significant transitory component in wealth that is unrelated to consumer spending and labor income. This transitory component is a reflection of the sizable forecastable component in stock market returns that is observed in US data over medium to long-horizons.¹⁴ *Quarterly wealth growth is dominated by transitory risk aversion shocks.*

The decomposition used above allows the permanent component of each variable to exhibit serially correlated “noise” around the random walk component. The random walk component is the estimated value the variable must take in the long run after all temporary shocks around it have dissipated. To give a sense of how quantitatively important this temporary noise around the random walk component is, we compute the multivariate Beveridge-Nelson decomposition for this system, which allows us to identify each variable’s random walk component. This decomposition does not rely on any particular orthogonalization of the errors and can be computed from the reduced-form errors in (2). Table 6 shows that even with the serial correlation in measured spending growth, Δc_t still displays a correlation of 97 percent with its random walk component, while Δy_t displays a 99 percent correlation with its random walk component. Thus the permanent shocks we identify are (essentially) random walks. This behavior for consumption is consistent with the canonical real business cycle model when TFP shocks are permanent and the EIS is close to zero. By contrast, asset wealth is far from a random walk, with Δa_t displaying a correlation of just 31 percent with its random walk component.

The random walk component of a_t is of interest for another reason: deviations from it tell us how far it is today from the estimated trend level it must end up at in the long run. Given the large transitory component in a_t , these deviations could be quite sizable. Let the *cyclical component* of a_t be defined as the difference between the actual series a_t and the random walk component of a_t . Note that this definition differs from

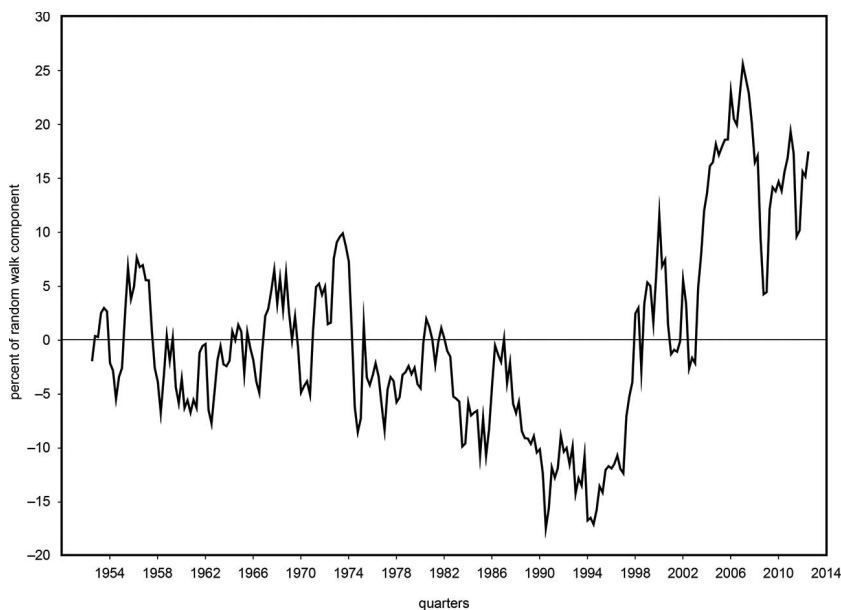


Fig. 6. Cyclical component of net worth

Notes: The figure plots log, per capita asset wealth minus its random walk component, in percent of the random walk component. The sample is 1952:Q1 to 2012:Q3.

the *transitory component* given above, which is the difference between the a_t and the permanent component of a_t .

Figure 6 plots the cyclical component of a_t , in percent of the trend component. The series displayed in the figure has been normalized so that when it is above zero, a_t is estimated to be above its long-run trend; when it is below zero, wealth is estimated to be below its long-term trend. Transitory swings in wealth are both quantitatively large and persistent. The two recent boom-bust episodes in asset markets stand out. In the tech boom, just before the 2000 to 2002 crash, the transitory component reached as high as 12.2 percent of the permanent component of wealth. Translated into dollar amounts, this implies that wealth exceeded its long-run level by as much \$19,284 per capita in 2005 dollars. The subsequent decline in stock market wealth that was predicted by this large deviation restored net worth to its long-run trend with consumption and labor earnings. Wealth was even more above its long-run trend at the peak of the housing boom; just before the 2007 to 2009 crash, the transitory component reached as high as 29.1 percent of the permanent component of wealth. Translated into dollar amounts, this

implies that wealth exceeded its long-run trend by as much \$47,986 per capita in 2005 dollars. As of the end of our sample (2012:Q3), the transitory component is 19.1 percent of the permanent component of wealth, exceeding its long-run trend by \$28,587 per capita in 2005 dollars. The rebound in the stock market since 2009 has driven household net worth farther above its long-run level than it was at the peak of the tech boom.

The large transitory component in wealth is difficult to reconcile with most modern macroeconomic models because such models typically have constant (or close to constant) risk premia. But this finding is also a challenge for leading consumption-based asset pricing models capable of rationalizing large movements in risk premia. These models are consistent with the existence of a transitory component in wealth, but they cannot account for our finding that this component is unrelated to consumption.

To illustrate, we show how our decomposition would look if the data were generated by two asset pricing models for which risk premia vary over time endogenously in response to consumption shocks. Figure 7 shows impulse responses of consumption and wealth in the Constantinides (1990) habit-formation model, a framework that can explain the high equity premium in the data without appealing to high risk aversion. The model is an endowment economy, so there is no labor income and therefore only a bivariate cointegrating relation between consumption and asset wealth. The model contains only a single primitive shock, namely the permanent shock to the invested endowment, which is a shock to log consumption. Thus there is in fact no transitory shock in the model, but because the model is nonlinear the permanent/transitory decomposition can still be computed using data simulated from the nonlinear model. The decomposition therefore produces two shocks, one labeled permanent and one labeled “transitory.” Were it not for nonlinearities in the model, the orthogonal transitory shock would explain none of the variation in either variable. Thus, perhaps not surprisingly, as figure 7 shows, the quantitative importance of the transitory shock for wealth (and consumption) is tiny. Nonlinearities alone are not enough to allow the model to explain the large transitory component in wealth found in postwar data. On the other hand, the consumption innovation in this model has effects that are very similar to the first permanent productivity shock in our decomposition.

The same calculation can be applied to the Campbell and Cochrane (1999) external habit formation model. Campbell and Cochrane (1999),

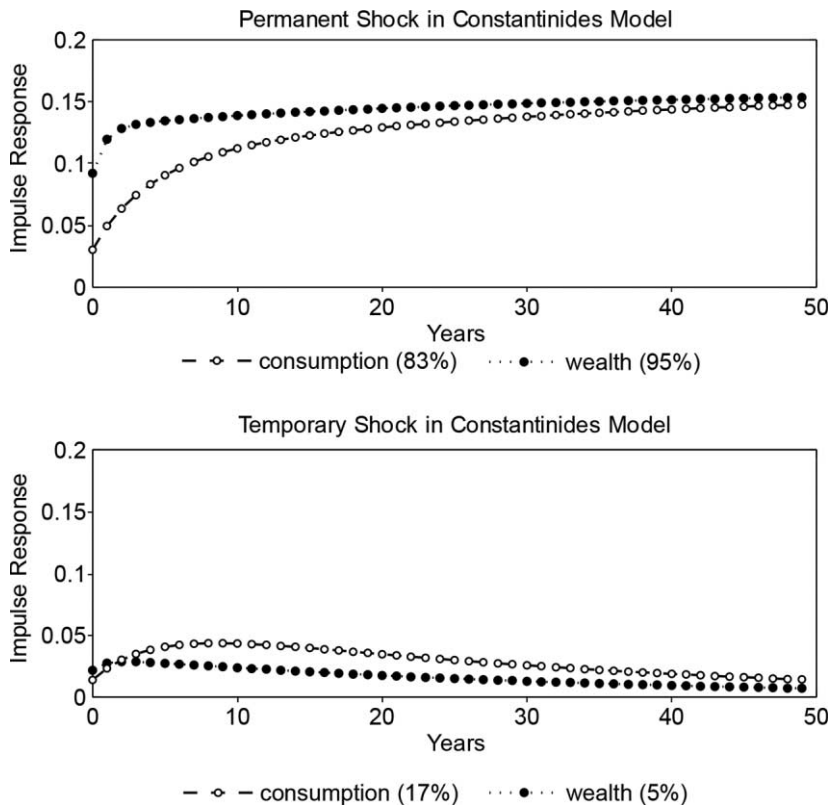


Fig. 7. PT decomposition of Constantinides habit model

Notes: The figure plots impulse response functions of consumption and wealth to permanent and transitory shocks in simulated data from Constantinides' (1990) habit model. The numbers in the legend are the fraction of the forecast error variance explained by the permanent and transitory shocks.

building on work by Abel (1990) and Constantinides (1990), showed that high stock market volatility and predictability could be explained by a small amount of aggregate consumption volatility if it were amplified by time-varying risk aversion. As for the Constantinides model, this is an endowment economy with a single permanent shock to log consumption, c_t . Also as in the Constantinides model, log consumption and wealth are cointegrated, and there is no labor income. This model differs from the Constantinides model in that the habit is external and is a highly nonlinear function of current and an infinite number of lags of past (aggregate) consumption, designed to allow the model to capture the long-horizon forecastability of stock returns by the price-dividend

ratio. Figure 8 presents impulse responses for log consumption and log wealth using simulated data from the baseline Campbell-Cochrane model. The top panel of figure 8 shows the responses of c_t and a_t to the one permanent shock in the model. Consumption jumps immediately to its new long-run level while wealth initially overshoots its long-run level. This temporary but persistent overshooting implies that wealth in the Campbell-Cochrane model deviates substantially from its random walk component, given by the log level of consumption. The model therefore implies that asset wealth has a quantitatively important transitory component that is *correlated* with the permanent consumption shock, implying that excess returns in the model have a significant forecastable component over longer horizons, as in the data. The bottom panel of this figure shows, however, that the orthogonal transitory shock represents a quantitatively small fraction of the variation in wealth, accounting for only 16 percent of its fluctuations. The model is inconsistent with evidence discussed above that orthogonal transitory shocks account for the vast majority of fluctuations in wealth.

B. *Permanent and Transitory Shocks in Wealth Components*

How do the major components of wealth respond to the structural disturbances? Figure 9 shows the cumulative impulse responses of the three major components of asset wealth (stock market wealth, housing, and non-stock market financial wealth) to a one-standard deviation innovation in each structural disturbance. The responses are constructed using the OLS estimates of (10). The productivity shock has an immediate impact on all three components (top panel), though note the scale of these responses is much smaller than that of the transitory shock (bottom panel). For housing, this shock is by far the most important, especially in the long run, as is evident from a comparison of long-horizon responses of housing to each shock. By contrast, the factor shares shock and the risk aversion shock have much smaller effects on home values, over almost every horizon.

For stock market wealth on the other hand, figure 9 shows that two of the structural disturbances are important: Over short horizons, the transitory risk aversion shock dominates variability in stock market wealth. A one-standard deviation increase in $\eta_{T,t}$ has a large and persistent effect on equity values. Figure 9 shows clearly that the risk aversion shock results in a transitory movement in net worth because it causes a transitory movement in the stock market component of wealth. But

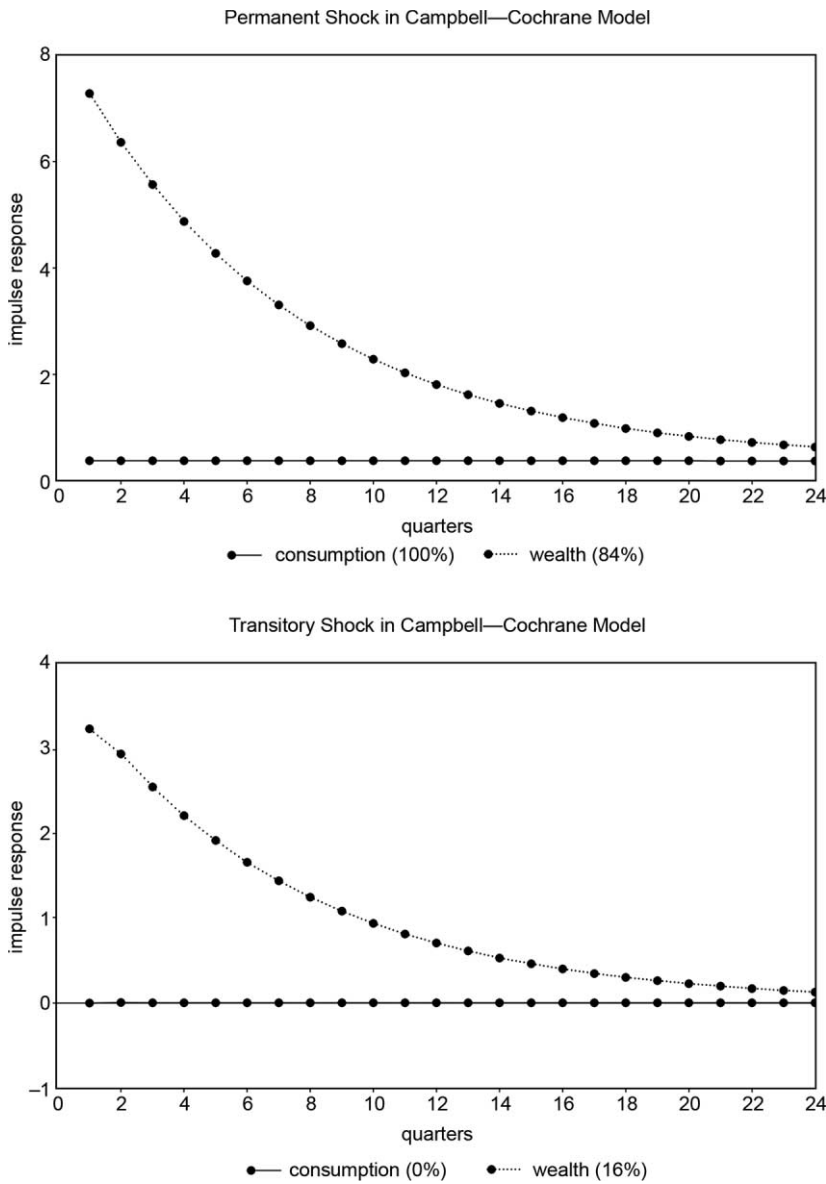


Fig. 8. PT decomposition of Campbell Cochrane habit model

Notes: The figure plots the impulse response functions of consumption and wealth to permanent and transitory shocks computed from simulated data of the Campbell and Cochrane (1999) habit model. The numbers in the legend are the fraction of the forecast error variance explained by the permanent and transitory shocks.

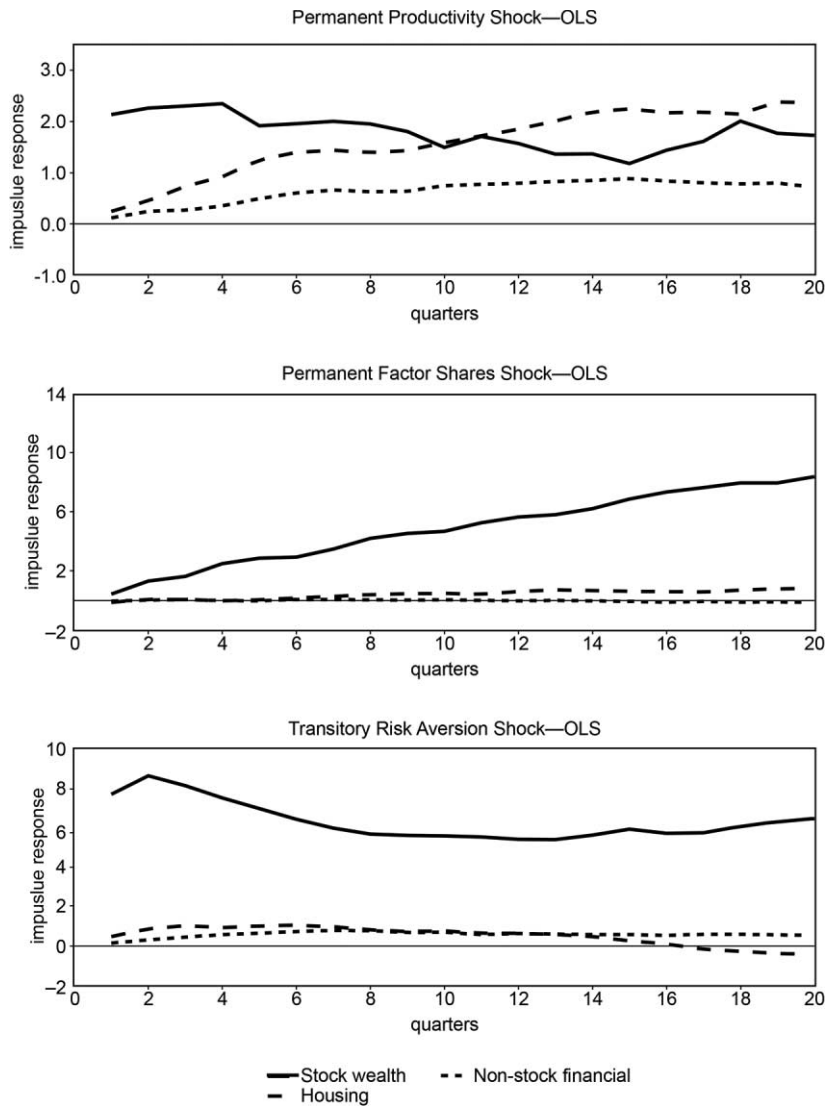


Fig. 9. Impulse response functions

Notes: The figure plots impulse response functions of wealth components to permanent and transitory shocks. The sample is 1952:Q1 to 2012:Q3.

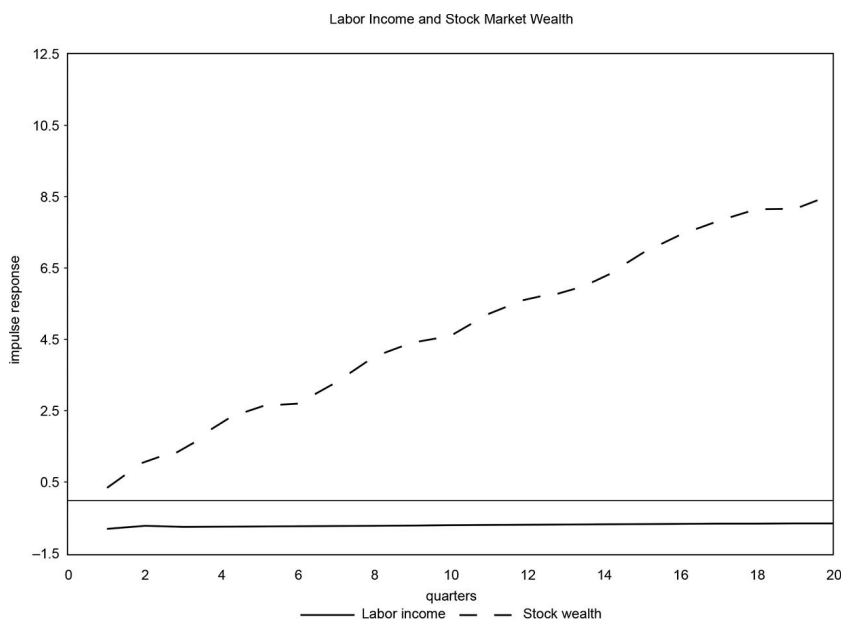


Fig. 10. Impulse response functions to factor shares shock

Notes: The figure plots impulse response functions of stock market wealth and labor income to a factor shares shock. The sample is 1952:Q1 to 2012:Q3.

over long horizons, the factor shares shock $\eta_{p2,t}$ also becomes quantitatively important. Indeed, the long-run response of stock market wealth to this shock is as quantitatively important as the short-run response to the transitory shock. It is perhaps puzzling that stock market wealth responds so sluggishly to this permanent shock, suggesting that the information revealed in the innovation is incorporated only slowly into stock prices. Alternatively, it could be that it is not the price component but rather the quantity component that responds slowly if the shock has a long-run effect on the number of firms going public. The factor shares shock has almost no effect on housing wealth or non-stock market financial wealth. Figure 9 shows clearly the factor share shock is a shock to *shareholder* wealth, not other forms of wealth.

The inverse influences of the factor shares shock on stock market wealth and labor earnings is highlighted in figure 10, which puts together the findings from figures 4 and 9. A one standard deviation increase in $\eta_{p2,t}$ leads to an immediate decline in labor earnings and a long-run increase in stock market wealth. Both stock market wealth and

Table 7
Wealth components: Variance decomposition

| Wealth component | Prod. shock (%) | Fact. shares shock (%) | Risk aversion shock (%) | Residual (%) |
|-----------------------|-----------------|------------------------|-------------------------|--------------|
| Net worth | 8 | 4 | 88 | 0 |
| Stock mkt. wealth | 7 | 6 | 74 | 13 |
| Housing | 20 | 8 | 24 | 49 |
| Non-stock fin. wealth | 19 | 7 | 24 | 50 |

Notes: This table reports the variance decomposition of wealth components based on OLS regressions on contemporaneous and lagged shocks. The numbers are the fraction of $h = \infty$ step-ahead forecast error in the log difference of the variable named in the row that is attributable to the shock named in the column heading. The last column reports that share of the variance that is due to the residual of the regressions. The sample spans the fourth quarter of 1951 to the third quarter of 2012.

labor income reach new, higher, and lower trend levels, respectively, in response to this shock.

Using (10), we characterize the relative quantitative importance of each structural disturbance (as well as of the residual) for the major components of wealth by computing a variance decomposition. Table 7 shows that the transitory risk aversion shock is most closely related to stock market wealth and accounts for 74 percent of its quarterly volatility.¹⁵ This shows that the vast bulk of quarterly variation in stock market wealth is attributable to exogenous risk aversion shocks. The two permanent shocks account for very small amounts, 7 and 6 percent for the productivity and factor shares shocks, respectively.

Housing wealth, on the other hand, is more closely related to the two permanent shocks than is stock market wealth, primarily to the first, which accounts for 20 percent of its variation. Still, the risk aversion shock accounts for a non-negligible 24 percent of the quarterly fluctuations in housing wealth growth in this sample. The remaining percentages for housing are accounted for by the factor shares shock (8 percent) and the residual $\epsilon_{i,t}$ (49 percent) in (10). The relatively large role for the residual in driving the quarterly dynamics of the log difference of housing wealth is likely to be at least in part attributable to a mechanical log/level error in (10).¹⁶

C. Decomposition of Levels

The results above decompose quarterly growth rates into components driven by the three shocks. We now wish to trace out their longer-run

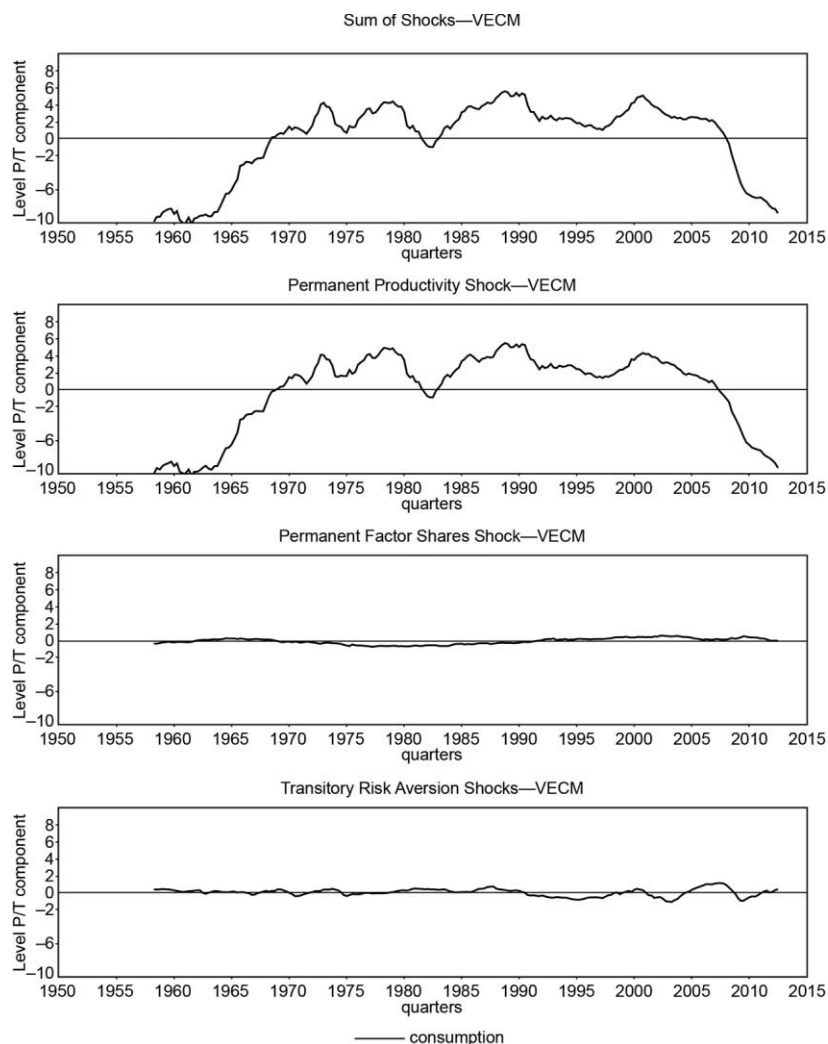


Fig. 11. VECM Level Decomposition—Consumption

Notes: The figure shows the decomposition of the log level of consumption into components driven by the productivity, factor shares, and risk aversion shocks, over time. A linear time trend is removed. The sample is 1952:Q1 to 2012:Q3.

effects on the levels of the variables. The shocks may differ in their effects on the growth rates over horizons longer than a quarter, and this should show up in the effect on the levels. Figures 11, 12, and 13 plot the levels decompositions over time for c_t , a_t , and y_t , respectively, using (12). The top panels of each figure show the sum of each component,

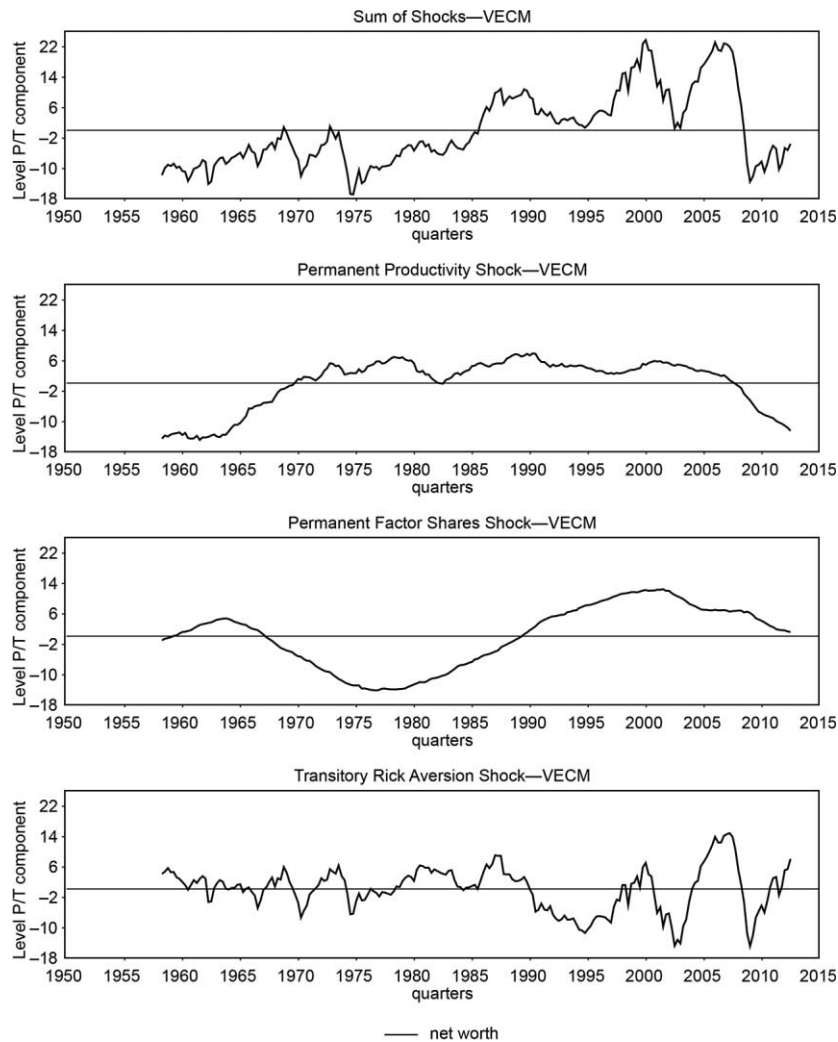


Fig. 12. VECM Level Decomposition—Net Worth

Notes: The figure shows the decomposition of the log level of net worth into components driven by the productivity, factor shares, and risk aversion shocks, over time. A linear time trend is removed. The sample is 1952:Q1 to 2012:Q3.

or the total level of the variable. The bottom panels show the decomposition in the total level attributable to the level components of each structural disturbance. We have removed the deterministic trend from the levels of each variable so that each component and its sum have a well defined mean.

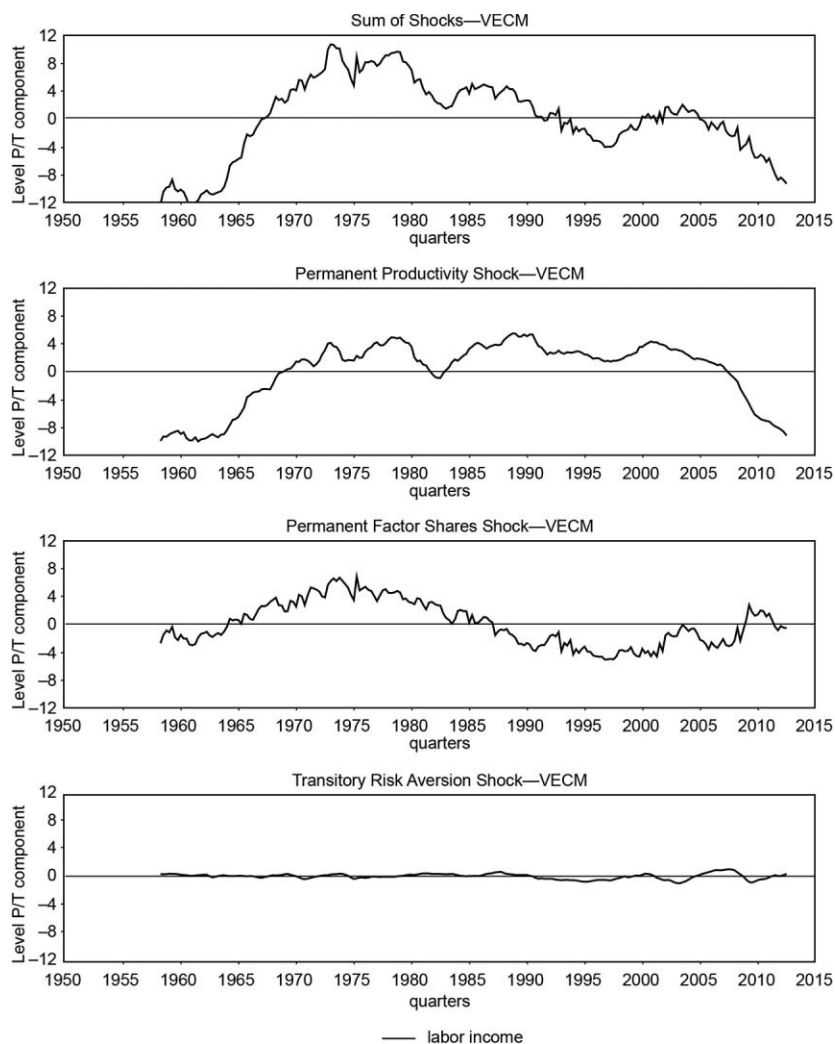


Fig. 13. VECM level decomposition—labor income

Notes: The figure shows the decomposition of the log level of labor income into components driven by the productivity, factor shares, and risk aversion shocks over time. A linear time trend is removed. The sample is 1952:Q1 to 2012:Q3.

Figure 11 shows that the movement in the log-level of consumption over time is dominated by the movement in the level of the productivity shock. The other two shocks play virtually no role in the determination of changes in the level of consumption. By contrast, shorter-term fluctuations in net worth (figure 12) are dominated by the level of the

risk aversion shock, but these short-run tendencies are shifted up or down by the cumulative effect of two permanent shocks. For labor earnings (figure 13), both the productivity and factor share shocks have influenced its level over low frequencies while the transitory risk aversion shock plays no role. In particular, the cumulation of factor shares shocks has been a persistent drag on labor earnings since the mid-1980s, while it was a persistent boon from the mid-1960s to the mid-1980s.

Figure 14 shows how these shocks have affected the levels of the three major wealth components, obtained using the OLS estimates and (11). It is clear from the top panel that all three major asset classes (stock wealth, housing, and non-stock financial wealth) are affected by about the same magnitude by the productivity shock over time. There are significant differences among these components in the roles the other two shocks play, however. For example, low frequency movements in the level of stock market wealth are dominated by the cumulative swings in the factor shares component, whereas shorter-lived peaks and troughs in the stock market have coincided with spikes up or down in the risk aversion component. But the figure also shows that the stock market experience over long periods is driven in great part by the factor shares shock. Both the factor shares shock and the risk aversion shock are much more important, quantitatively, for the level of stock market wealth than is the productivity shock (note the left-hand scales). Finally, note that housing and nonfinancial stock market wealth are little affected by the factor shares shock. For these components, the quantitatively important innovation over long and short horizons is the permanent productivity shock.

An important aspect of these results, noted earlier, is that the low frequency movements in the stock market are the inverse image of those in labor earnings. This is further illustrated by figure 15, which shows the stark inverse relationship over time between labor earnings and the stock market that is the result of the cumulative reallocative outcomes of the factor shares shock. For the last twenty-five years, the cumulative effect of this shock has persistently lowered the level of labor earnings and persistently boosted stock market wealth. By contrast, for twenty years prior to that (from the mid-1960s to the mid-1980s), the cumulative effect persistently boosted labor earnings and lowered stock market wealth.

A more formal way to examine how the behavior of the level of the variables is affected over time by the three shocks is to decompose the variance of the log level by frequency, using a spectral decomposition.

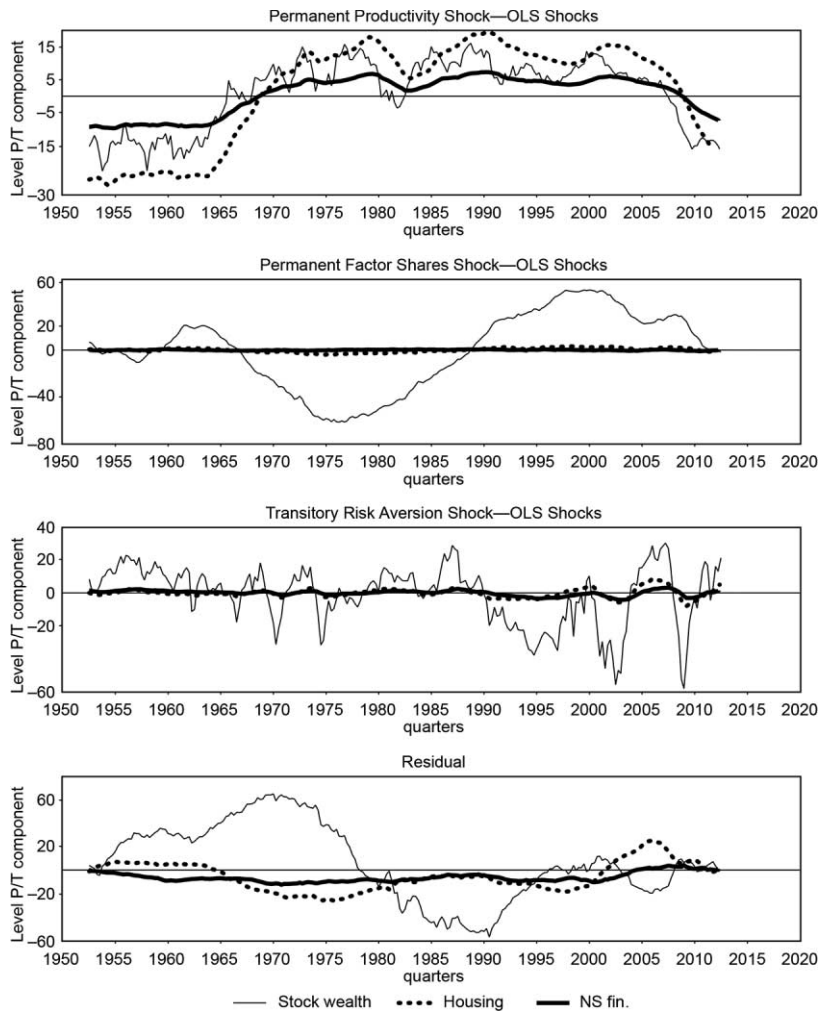


Fig. 14. Level decomposition

Notes: The figure shows the decomposition of the log level of stock wealth, housing wealth, and non-stock financial wealth into components driven by the productivity, factor shares, and risk aversion shocks over time. A linear time trend is removed. The sample is 1952:Q1 to 2012:Q3.

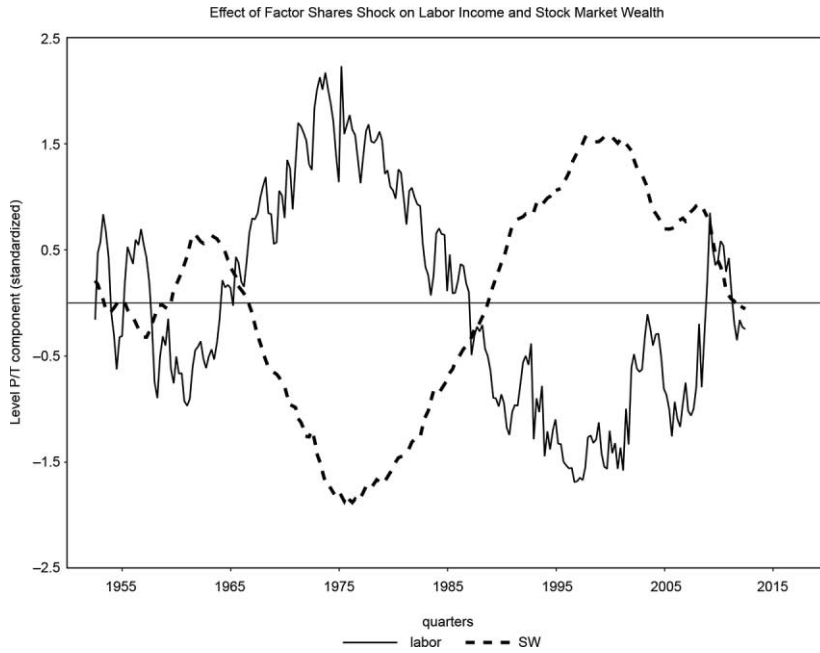


Fig. 15. Decomposition of labor income and stock market wealth

Notes: The figure shows the component of the log levels of stock market wealth and labor income that is attributable to the cumulative effect of the factor shares shock over time. Both series are demeaned and divided by their standard deviations. The sample is 1952:Q1 to 2012:Q3.

To do so, we estimate the spectrum for the log difference in each variable and then apply a filter to infer the level spectrum, allowing us to compute the fraction of the variance in the log level of each variable that is attributable to cycles of different lengths, in quarters. Figure 16 exhibits this decomposition for c_t , a_t , and y_t ; Figure 17 exhibits the same decomposition for the major components of wealth.

Figure 16 shows that the variance in the level of consumption is dominated at all frequencies by the productivity shock (top panel). The variance of labor income, on the other hand, is dominated by factor shares shock for cycles of one to sixteen quarters, where the latter roughly corresponds to the length of a typical (median) NBER business cycle as measured from cycle peak to cycle peak. For cycles between six and thirty-two quarters, the factor shares shock and productivity shock explain about the same amount of the variance in the level of labor earnings; for very long cycles, the productivity shock explains a little more

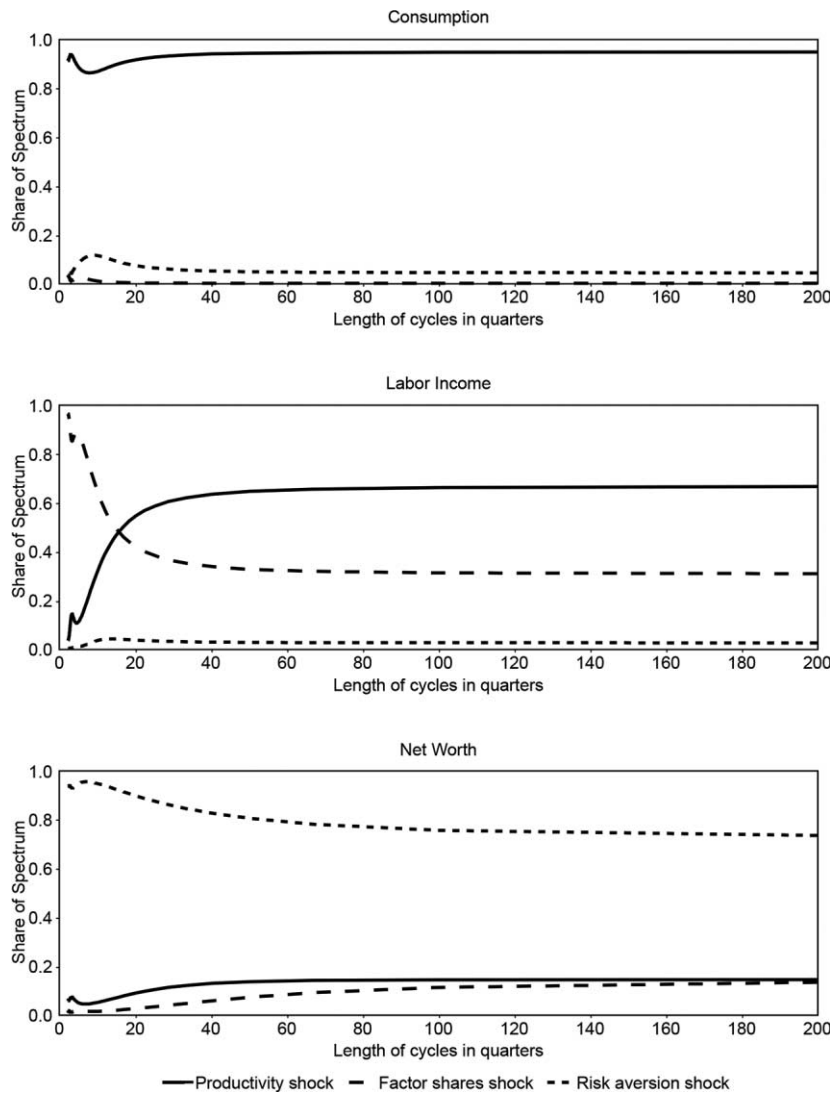


Fig. 16. Decomposition of spectra

Notes: The figure shows the decomposition of spectra at different frequencies. Spectra are estimated in first differences and converted into spectra for levels: $S_x(\omega) = (1 - \exp(-i\omega))^{-1} (1 - \exp(-i\omega))^{-1} S_{\Delta x}$. The sample is 1952:Q1 to 2012:Q3.

than 60 percent, while the factor shares shock explains a little less than 40 percent. The transitory risk aversion shock plays no role in the variance of labor income at any cycle. Yet this shock is the most important contributor to the volatility of total net worth at all frequencies, a fact that reflects the persistent nature of the transitory wealth shock (bottom panel).

Figure 17 shows that the large role for the transitory component in the variance of a_t is reflected in all the major components of a_t , especially at frequencies corresponding to cycles between two and forty quarters. For stock market wealth, however, the factor shares disturbance plays an increasingly important role as the length of the cycle increases, eventually explaining over 25 percent of the variance at very low frequencies, while the importance of the risk aversion shock declines from a high of almost 90 percent at high frequencies to 40 percent at low frequencies. For housing wealth, the transitory risk aversion shock plays an important role over cycles of short lengths while the productivity shock plays the most important role over long cycles.

D. A Comparison of Two Asset Market Cycles

How does the recent behavior of these shocks relate to the observed volatility in asset values and the real economy in recent business and asset market cycle episodes? During the last twenty years, the US economy has experienced two recessions, accompanied by two asset market “crashes.” In the “tech bust,” from March 2000 to September 2002, the S&P 500 index declined 39 percent, accompanied by only a modest drop in real activity during the 2001 recession. In fact, real, per capita consumption and housing wealth rose from the first quarter of 2000 through third quarter of 2002, increasing 3.6 and 23 percent, respectively. During the boom years leading up to this contraction, the S&P 500 index rose 221 percent (March 1994 to March 2000) while real, per capita housing wealth rose a more modest 22 percent. By contrast, the 2007 to 2009 recession was associated with a decline in stock market wealth of similar magnitude (the S&P 500 fell 44 percent from September 2007 to March 2009), but a much larger decline in real activity and housing wealth. In the “housing bust” from the third quarter of 2007 through the first quarter of 2009, consumption fell 1.67 percent and housing wealth declined 26 percent on a real, per capita basis. Prior to this crash, the S&P 500 rose 75 percent (September 2002 to September

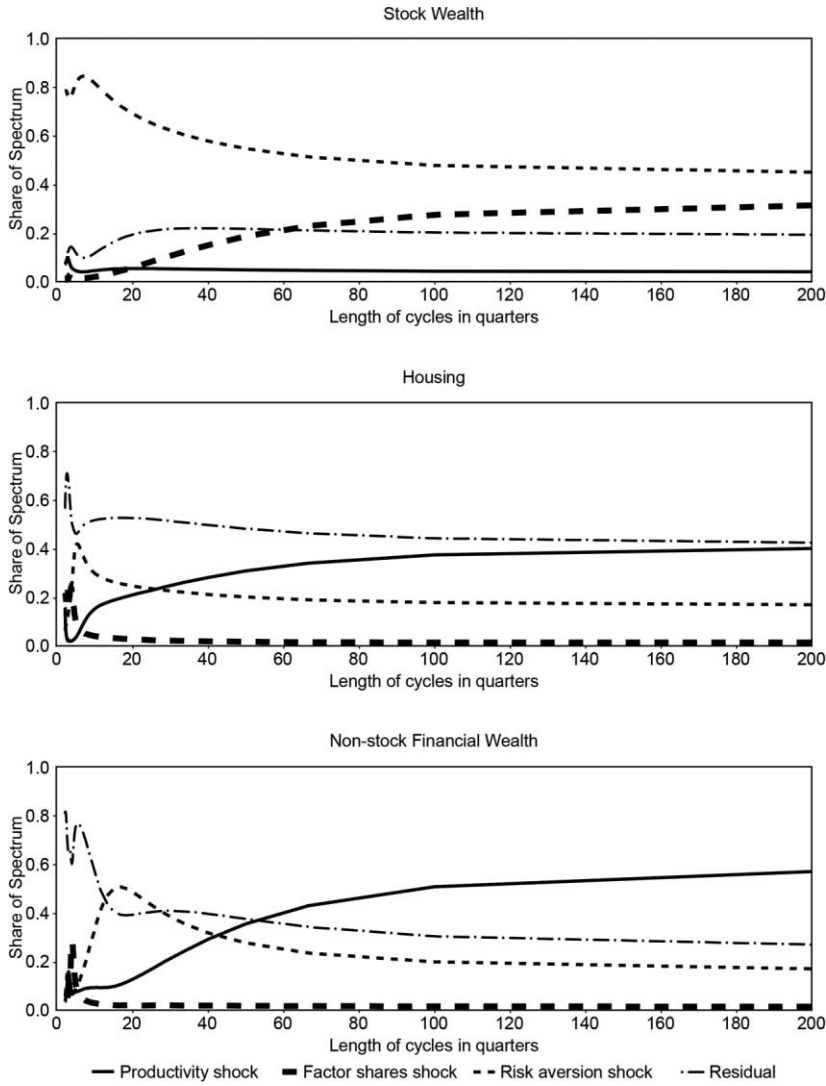


Fig. 17. Decomposition of spectrum

Notes: The figure shows the decomposition of spectra at different frequencies. Spectra are estimated in first differences and converted into spectra for levels: $S_x(\omega) = (1 - \exp(-i\omega))^{-1} (1 - \exp(-i\omega))^{-1} S_{\Delta x}$. The sample is 1952:Q1 to 2012:Q3.

2007), while housing wealth rose 38 percent (third quarter 2002 to first quarter 2006).

The next figures explore the role of each structural disturbance in household net worth and its major asset components over these two boom/bust cycles. Figures 18 through 20 show the same level decompositions described above, but focused on the time period of these two episodes. The three vertical lines in each plot divide the episodes into four subpanels, two boom and two bust periods, with the first boom period measured from 1994:Q1 to 2000:Q1, followed by a bust from 2000:Q1 to 2002:Q3, followed by a boom from 2002:Q3 to 2007:Q3, followed by a bust from 2007:Q3 to the end of our sample, 2010:Q2.

Figure 18 exhibits the role each shock played in driving the level of total net worth during these episodes. The figure shows that the tech bust from 2000 to 2002 (and the boom that preceded it) was almost entirely the result of a string of transitory risk aversion shocks: declining risk aversion shocks in the tech boom leading to big transitory increases in wealth, and rising risk aversion shocks in the bust leading to big transitory decreases. The boom in net worth in the first subpanel on the left and the decline in the second subpanel mirror closely the boom-bust pattern in the cumulative effects of our identified transitory shock. Indeed, the two permanent shocks had little effect on the level of net worth: the cumulative effect of these shocks was essentially flat over these two subperiods. Figure 19 shows that this same pattern during the tech boom and bust is evident in stock market wealth but there is no similar boom-bust pattern in housing wealth (figure 20). The asset market downturn of 2000 to 2002 and the boom that preceded it was almost entirely the result of a string of transitory risk aversion shocks that affected only the stock market component of wealth.

By comparison, the housing market boom from 2002:Q2 to 2007:Q3 (shown in the third subpanel from the left) was the result of a culmination of positive transitory shocks (negative risk aversion shocks) that affected all three major wealth components. The asset market *crash* of 2007 to 2009, on the other hand, was characterized by large negative roles for *both* the risk aversion shock and the permanent productivity shock, with the latter having especially important effects on housing wealth.

The last subpanel of each plot includes the level forecasts for each component, extended out past the end of our sample, 2012:Q3. The forecasts are computed from (10) by assuming all future innovations (past the end of our sample) are equal to their population means of zero, and

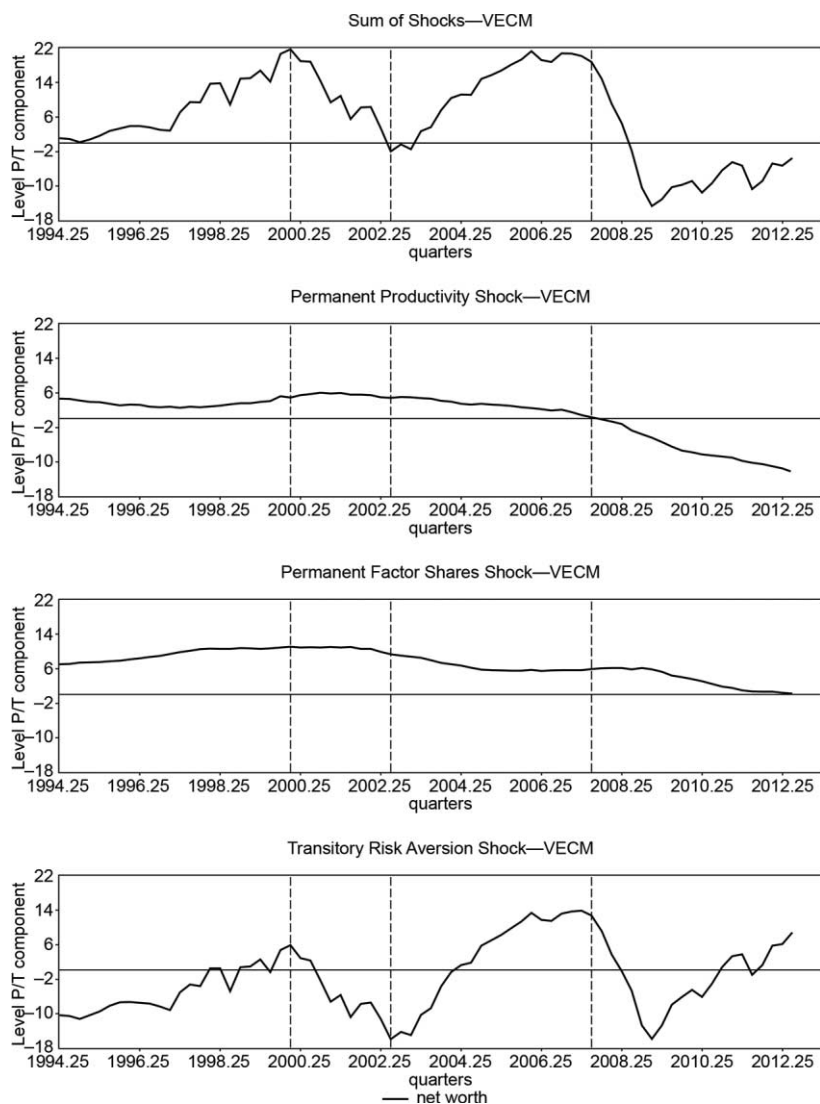


Fig. 18. VECM level decomposition—net worth: Subsample 1994:Q1–2012:Q3

Notes: The figure shows the decomposition of the log level of net worth into components driven by the productivity, factor shares, and risk aversion shocks for the 1994:Q2–2000:Q4 subsample. The vertical lines divide the sample into two boom periods (1994:Q2 to 2000:Q1 and 2002:Q4 to 2007:Q3) and two bust periods (2000:Q2 to 2002:Q3 and 2007:Q4 to 2012:Q3). A linear time trend is removed.

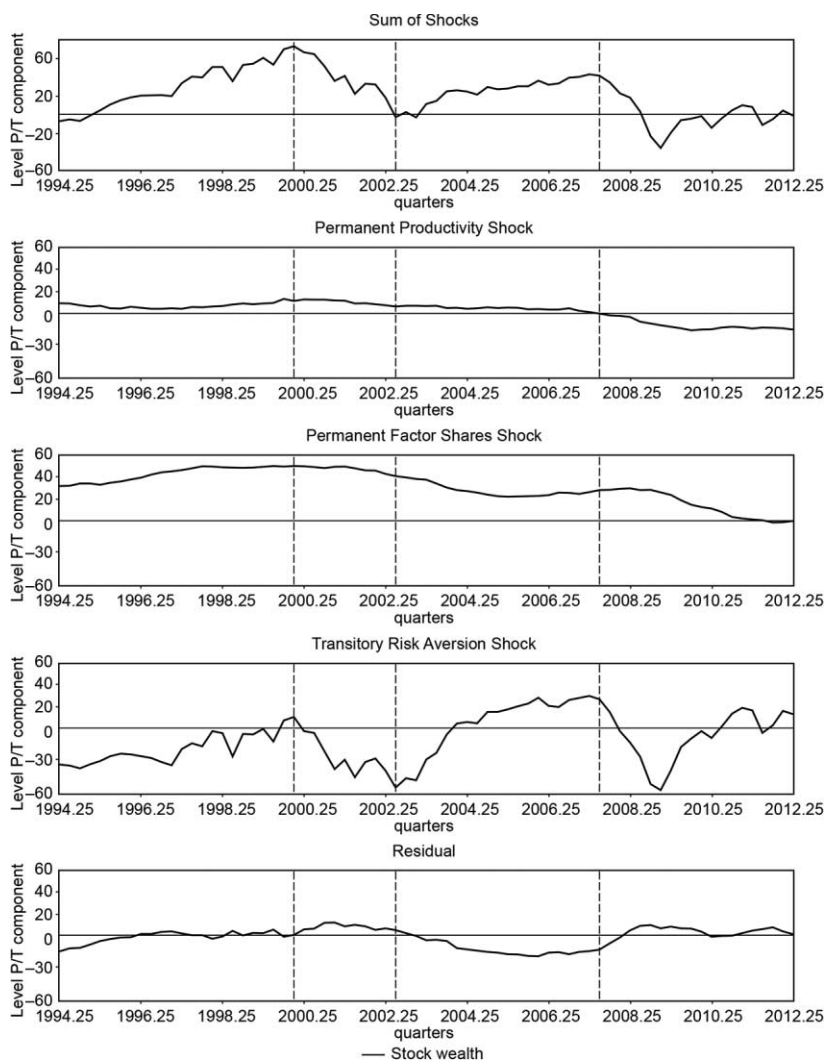


Fig. 19. OLS level decomposition—Stock market wealth: Subsample 1994:Q1–2012:Q3
Notes: The figure shows the decomposition of the log level of stock market wealth into components driven by the productivity, factor shares, and risk aversion shocks for the 1994:Q2–2000:Q4 subsample. The vertical lines divide the sample into two boom periods (1994:Q2 to 2000:Q1 and 2002:Q4 to 2007:Q3) and two bust periods (2000:Q2 to 2002:Q3 and 2007:Q4 to 2012:Q3). A linear time trend is removed.

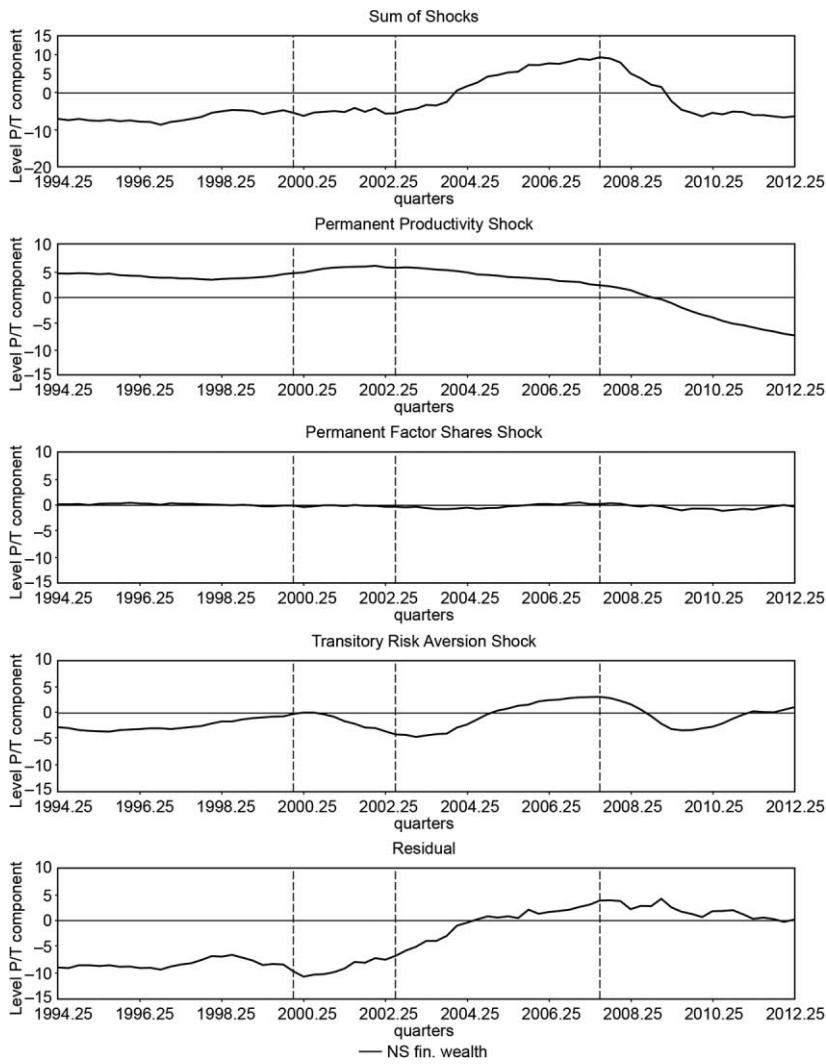


Fig. 20. OLS level decomposition—non-stock financial wealth: Subsample 1994:Q1–2012:Q3

Notes: The figure shows the decomposition of the log level of non-stock financial wealth into components driven by the productivity, factor shares, and risk aversion shocks for the 1994:Q2–2000:Q4 subsample. The vertical lines divide the sample into two boom periods (1994:Q2 to 2000:Q1 and 2002:Q4 to 2007:Q3) and two bust periods (2000:Q2 to 2002:Q3 and 2007:Q4 to 2012:Q3). A linear time trend is removed.

then rolling the computation for the levels forward (see [11]). In contrast to the 2000 to 2002 experience, the model forecasts imply persistently low home values going forward from the end of our sample, the result of a string of negative draws to the permanent productivity shock since 2007. The cumulative effects of the other two shocks play no role in this gloomy forecast for housing. At the same time, the rebound in the stock market since 2009 can be traced to a string of transitory innovations (figure 19, third panel): the only component of stock market wealth moving up during the period since 2009:Q1 is the transitory component interpreted as a rebound in risk tolerance from the drastic declines that occurred during the housing bust. Returning to Figure 6 we see that stock market wealth has once again driven household net worth to values significantly above its long-run level by the end of the sample.

These patterns are similar to findings reported in Campbell, Giglio, and Polk (2013), who argue that, in 2000 to 2002, stock prices fell primarily because discount rates increased (implying a transitory decline in stock wealth), while the 2007 to 2009 crash was attributable to both worsening cash flow prospects (implying a permanent decline in stock wealth) and to higher discount rates.

We left out any separate analysis of the twin recessions in the quarters between 1980:1 and 1982:4. There was no significant decline in the stock market during this period—indeed, the S&P 500 index rose 33 percent between 1980:Q1 and 1982:Q4—though housing wealth deteriorated. But it is worth noting that this recession was characterized by a string of negative productivity shocks, with no major contribution from risk aversion or factor shares shocks. This reinforces the previously found pattern that housing is particularly hard hit by negative productivity shocks but not by risk aversion shocks or factor shares shocks, while the opposite is true for stock market wealth.

IV. Discussion

We now discuss some alternative possible interpretations for our shocks, along with some weak links in our own interpretation. We start with our own weak links.

It is important to remember that, econometrically, we have decomposed the historical variation in c , a , and y into shocks distinguished only by their degree of persistence, and then followed by assigning each shock an economic interpretation. The advantage of this is that—given only a hypothesis of cointegration—we can completely charac-

terize the dynamics of the system with this decomposition. We have argued that the interpretations we have assigned are plausible, but one should bear in mind that we have not directly identified shocks that have those interpretations. It follows that any specific interpretation will be imperfect.

Our identification of the TFP shock (for example) is obviously imperfect if one takes estimates of Fernald (2009) as accurate. Figure 2 shows that the long-run tendencies in these series are quite similar, but there are some notable gaps in certain periods. We have also plotted the capacity *unadjusted* TFP series. Whatever is actually removed to obtain the adjusted series, we found that it causes measured TFP shock to display a far weaker relation with consumption and labor income growth than is the case for the unadjusted TFP series. To the extent that the adjustment accurately reflects changes in capacity utilization, our interpreted productivity shock should then *not* be thought of as a pure technological innovation, but rather as an innovation in productivity that includes shocks independent of changes in the technological frontier. Of course TFP is itself unobservable, and both our measure and Fernald's are proxies for the true value of this latent primitive shock.

Our interpretation of consumption shocks as driven by permanent TFP shocks rules out any role for a second mutually uncorrelated permanent shock to contemporaneously affect consumption, and it rules out a contemporaneous role for possible other transitory shocks (driven by, for example, monetary or fiscal shocks) that could temporarily affect consumption growth through a change in the expected real interest rate if the EIS is nonzero. We emphasize, however, that if such additional shocks are present, they do not appear to be quantitatively important for consumption: both the second permanent shock and the transitory shock in our system explain a negligible fraction of the variance of consumption growth at all horizons, not just in the zero-th horizon in which the restriction is imposed. Our interpretation of consumption shocks is consistent with the findings of Justiniano, Primiceri, and Tambalotti (2009), who show that the comovement of consumption with the rest of the economy is driven mainly by factor-neutral technology shocks.

Besides TFP, there are two other shocks in our system, one or the other of which is predominately responsible for quarterly variation in labor income growth or asset wealth growth. We are not the first to conclude that factor-neutral technology shocks are only one source of variation in systems of economic aggregates (see, e.g., King et al. 1991; Galí

1999; Christiano, Eichenbaum, and Vigusson 2004; Fisher 2006; Smets and Wouters 2007; Justiniano, Primiceri, and Tambalotti 2009, 2010).¹⁷

Other researchers have focused primarily on output, investment, and hours, and on their variability at business cycle frequencies, whereas we have focused on economic aggregates motivated by the household budget constraint, namely consumption, asset wealth, and labor earnings, and have scrutinized both long- and short-run tendencies in these variables.

Others have also examined the role of shocks to labor's share in the production process, both in the data and in a standard real business cycle model (e.g., Ríos-Rull and Santaaulalia-Llopis 2009). Still others have argued that variation in the labor share explains an important fraction of inflation variation, as in Galí and Gertler (1999). Ríos-Rull and Santaaulalia-Llopis (2009) find empirically that shocks to labor's share tend to be negatively correlated with measures of productivity, something we confirm using Fernald's measure of TFP and the BLS measure of labor share. But since we have assumed our two permanent shocks are orthogonal in order to identify their independent effects, our decomposition is silent on the role played by the correlated component of the factor share and TFP shocks. Figure 21 plots our measure of the factor shares component of labor income and the BLS labor share over time. The two series generally follow similar long-run trends, but there are some noticeable gaps in the mid-1950s and at the end of the sample.

We have shown previously that the transitory risk aversion shock is unrelated to consumption or labor income. Even if our interpretation of this shock as risk aversion disturbance exogenous to consumption were correct, it does not follow that it must be unrelated to other measures of real activity, such as investment. A simple Q -theory of investment would imply that such a shock should have affects on investment because it affects the rate at which future cash flows are discounted (e.g., Cochrane 1991; Abel 1983; Abel and Blanchard 1986). Consistent with this, empirical evidence suggests that cay_t forecasts both investment growth and excess stock returns, in opposite directions, over long horizons (Lettau and Ludvigson 2002).

Figure 22 shows the dynamic response of various measures of investment expenditure to the three structural disturbances in our system. The responses are computed in the same way that the responses for the wealth components are computed, by summing the growth rates implied by an OLS regression of investment on distributed lags of $\eta_{p1,t'}$, $\eta_{p2,t'}$, and $\eta_{T,t'}$, as in (11). An increase in the transitory shock raises all

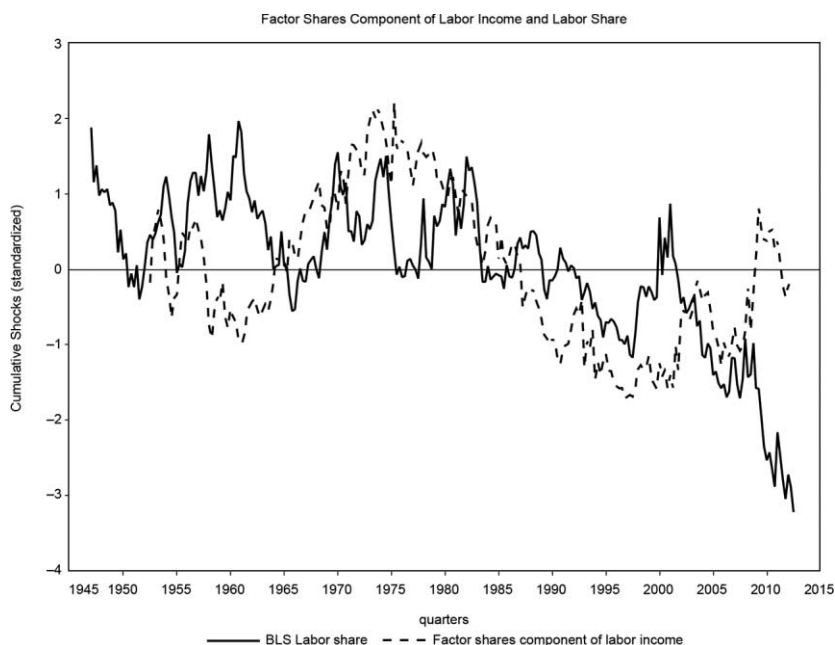


Fig. 21. Factor shares component and labor share

Notes: The figure plots the labor share constructed by the Bureau of Labor Statistics and the component of the log level of labor income that is attributable to the cumulative effect of the factor shares shock. Both series are demeaned and divided by their standard deviations. The sample is 1952:Q1 to 2012:Q3.

four types of investment (private investment, nonresidential investment, structures, and equipment) within a few quarters' time and then eventually lowers investment, consistent with the interpretation of the transitory shock as a discount rate shock driven by risk aversion. Since the transitory shock is unrelated to consumption, however, the source of this discount rate shock cannot be a consumption innovation, as in many consumption-based asset pricing models. Note also that an increase in the permanent productivity shock has a positive effect on all investment categories. This is not surprising since a positive productivity shock increases the expected value of marginal profits (cash-flows), and therefore optimal investment in a Q model with adjustment costs. Finally, the figure shows that a positive value for the factor shares shock leads to a decrease in investment. This could occur if the shift in factor shares is driven by a persistent shock to uncertainty or by price mark-ups, as discussed next.

Let us now entertain several possible theoretical explanations for the

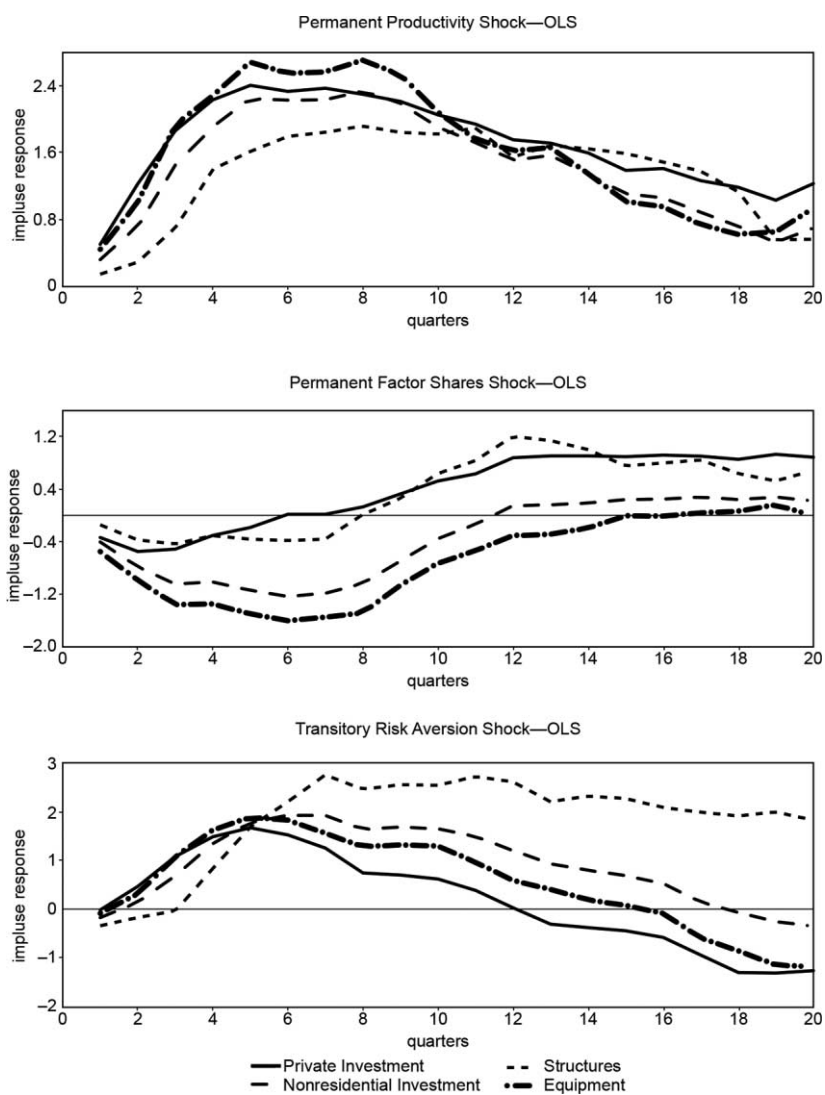


Fig. 22. IRF of investment to shocks

Notes: The figure plots impulse response functions of different measures of (log level) investment in response to the productivity, factor shares, and risk aversion shocks. The sample is 1952:Q1 to 2012:Q3.

second permanent shock that offer potential refinements of our interpretation. One could involve “uncertainty” shocks of the type considered by Bloom (2009). A positive uncertainty shock makes the distribution of future dividends right-skewed and raises the price-dividend ratio (Pastor and Veronesi 2006). At the same time, macroeconomic researchers have explored the interactions of financial frictions with increases in uncertainty about firm-level growth rates and have found that increases in uncertainty generate a fall in both output and labor income and a rise in interest rates (e.g., Arellano, Bai, and Keno 2010). These models may also imply that positive uncertainty shocks interacting with financial frictions lead to a reduction in investment, as in Gilchrist, Sim, and Zakrajsek (2010), consistent with the effect of the factor shares shock on investment discussed above.

Jurado, Ludvigson, and Ng (2013) develop empirical measures of aggregate uncertainty over different forecast horizons of h periods. Two different measures capture uncertainty across a large number of economic indicators: common firm-level uncertainty factors and common macro uncertainty factors. The first measures common variation in forecast error variance across a large number of firm-level profit growth rates; the second measures the same for forecast error variance across a large number of macroeconomic time series. Figure 23 shows the dynamic responses of these uncertainty measures to our factor shares shock, computed in the same way that the responses for the wealth components are computed, by summing the growth rates implied by an OLS regression of uncertainty factors on distributed lags of $\eta_{P1,t}$, $\eta_{P2,t}$ and $\eta_{T,t}$, as in (11). The top panel shows that common firm-level uncertainty does rise in response to our second permanent shock. The bottom panel shows a more muted response for the macro uncertainty measures. These results hint at a possible role for uncertainty shocks in driving the second permanent shock that shifts factor shares, but more work is needed to fully flesh out this link.

The factor shares shock could instead be the result of *directed* technological change. Researchers have postulated models of directed technological change to explain why the returns to some factors deviate persistently from others (e.g., Acemoglu 2002). This suggests that, in principle, a model of technological change biased *in favor* of capital and *against* labor might explain the simultaneous high rewards to the stock market and low rewards to labor effort resulting from the cumulative effect of the factor shares shock over the last twenty-five years, while

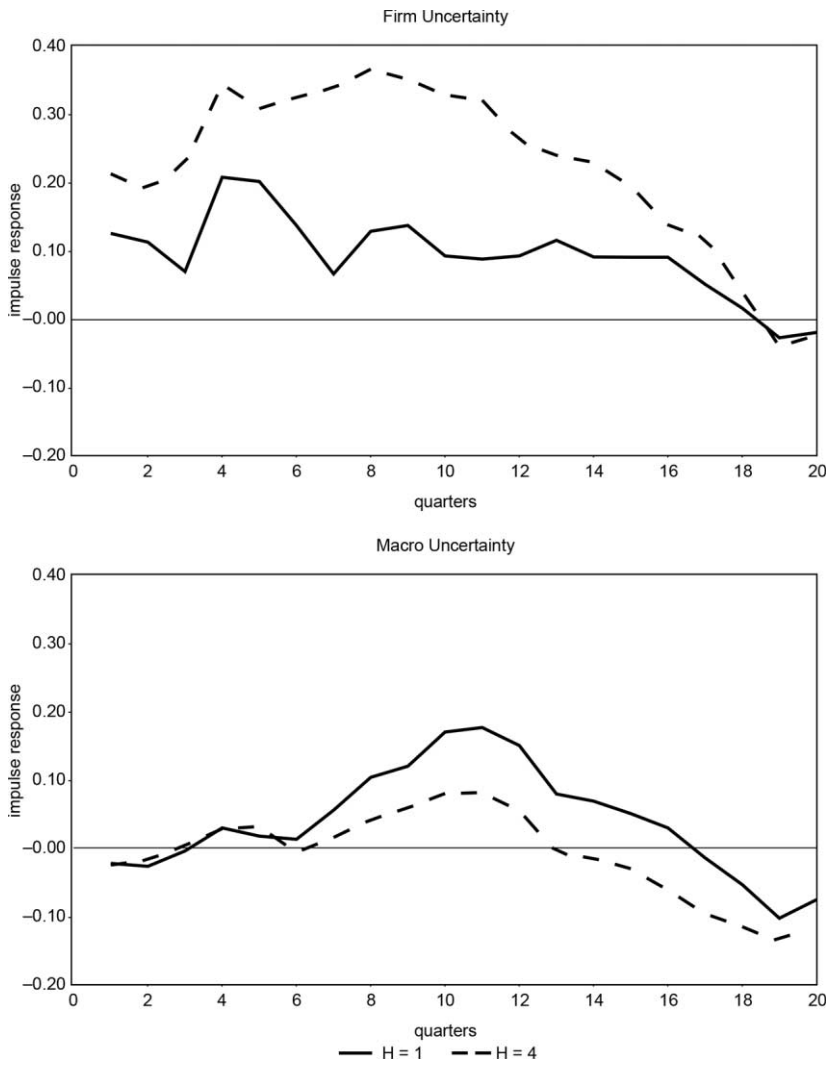


Fig. 23. IRF of uncertainty to factor shares shock

Notes: The figure plots impulse responses of different measures of aggregate uncertainty from Jurado, Ludvigson, and Ng (2013) in response to the factor shares shock. The top panel plots the response of common firm-level uncertainty factors for uncertainty horizons $h = 1$ quarter and $h = 4$ quarters. The bottom panel plots the response of common macro uncertainty factors for the same uncertainty horizons.

technological change biased *in favor* of labor and *against* capital might explain the converse over the twenty years prior.

The factor share shock could be driven by price mark-up shocks. A price mark-up creates a wedge between prices and marginal cost, so an increase in the mark-up increases profits and decreases the real wage. Justiniano, Primiceri, and Tambalotti (2009) find that price mark-up shocks explain 31 percent of the variance of wages at business cycle frequencies, while factor-neutral productivity shocks explain 40 percent. Unreported results in Justiniano, Primiceri, and Tambalotti (2009) also imply that a positive price markup shock lowers investment, as in the response of investment to a factor share shock. Given the cumulative values of the factor shares shock displayed in figure 5, this interpretation implies that the economy has become much less competitive in the last twenty-five years.

We close this discussion by noting one aspect of the data on consumption, asset wealth, and labor income that has changed over time and that makes the statistical interpretation of the model more ambiguous today than it has been in the past: the estimated cointegrating residual cay_t has become much more persistent. This can be seen in the half-life to the transitory shocks that affect a_t . In Lettau and Ludvigson (2004), these shocks were estimated to have a half-life of three years; in this update of the data, they have a half-life of four years. Thus deviations from the common trend for these variables are more long-lasting, and the forecastable component of asset returns operative at longer horizons, compared to previous samples. On a purely statistical basis, this increase in persistence means it is now more difficult to econometrically distinguish a stationary specification for cay_t from a unit root specification, although it should be stressed that we still find no evidence against the null of cointegration among c_t , a_t , and y_t , even in these most recent data. There are well-known difficulties in finite samples with distinguishing a highly persistent but stationary series from a genuinely non-stationary one.

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

that the system is characterized by three independent random walks. If one believes that these variables must share a common long-run trend, and given the evidence that c and y have never been even weakly forecastable by cay_t (even in samples where the latter was much less persistent), we argue that it is more reasonable to conclude that it is Δa_t (rather than Δc_t or Δy_t) that deviates substantially from a random walk, even if those deviations appear far more persistent today than they have in the past. This reasoning rules out rational bubbles. But bubbles create volatility in asset values without creating predictability in returns, and we have shown elsewhere that almost all of the variation in Δa_t is attributable to forecastable variation in excess returns. Every major asset crash in our postwar sample has been preceded by an unusually low level of cay_t , and an unusually high value for the cyclical component of a_t .

One possible explanation for this shift in the estimated persistence of cay_t is that there has been a structural change in the long-run relationship among c , a , and y , as there has been for the mean of the log dividend-price ratio, which appears to have shifted downward in the last fifteen years. If such a shift has occurred but we have ignored it, the estimated cointegrating residual could appear to have a unit root, even though it is actually stationary in each regime. We intend to investigate these possibilities in future work.

V. Conclusion

In this study we have decomposed the postwar variation in consumption, labor income, and household wealth into components driven by three mutually orthogonal structural disturbances. The disturbances we identify are econometrically distinguished only on the basis of whether their effects are permanent or transitory. But we have argued here for a specific economic interpretation for these shocks, with one interpreted as a permanent productivity shock, another interpreted a permanent factor shares shock, and the third interpreted as an exogenous risk aversion shock.

We find that the productivity shock underlies the vast bulk of quarterly fluctuations in consumption growth, the factor shares shock underlies the vast bulk of quarterly fluctuations in labor income growth, and the risk aversion shock underlies the vast bulk of quarterly fluctuations in asset wealth growth. The first finding is straightforward to reconcile with canonical, dynamic stochastic general equilibrium models, but the last two present more of a challenge. Quarterly labor in-

come is found to be dominated by shocks that merely redistribute the rewards of production, rather than raise or lower all of them. Over the last twenty-five years, the cumulative effects of such shocks have persistently boosted shareholder wealth and persistently lowered worker income. Short-run fluctuations in wealth, on the other hand, are dominated by shocks that have no affect on consumption or labor income, though they do affect investment. This finding is difficult to square with most macro models and even with leading asset pricing models that generate large movements in risk premia, because those models do so by counterfactually linking innovations in risk premia to innovations in consumption.

We recognize that any econometric exercise involving the precise naming of shocks is fraught with potential peril because our models are mere simplified representations of the world, because any empirical “observations” on primitive inputs are almost always gross approximations of the theoretical concepts we wish to measure, and because we must always assume that some component of each of these disturbances is independent in order to distinguish their effects econometrically. We therefore close by emphasizing three empirical features of the data that stand out from our findings, and that we argue any economic model should fit, however one wishes to interpret our identified disturbances.

First, log real, per capita consumption is (essentially) a random walk. Consumption does not adapt sluggishly to shocks nor is it correlated with the large transitory movements in household wealth driven by the volatile stock market component. These statements do not depend on any orthogonalization of shocks. We argue that any model should be consistent with this basic fact. Models in which consumption adapts sluggishly to shocks (as in some habit models), or moves significantly with transitory fluctuations in income or wealth will be at variance with this fact.

Second, factor shares vary in a persistent fashion over time and have a significant impact on labor income. Over the last ten years, labor’s share has plummeted (reaching an historic low by the end of our sample) while stock prices relative to measures of fundamental value have exhibited a structural break upward.¹⁸ We argue that any economic model must fit these twin facts, which suggest the rewards from production have shifted over time, most recently away from workers and toward shareholders. Models of competitive labor markets where factors of production earn constant income shares are unlikely to fit this fact.

Third, household wealth is about five times as volatile as consumption and two-and-a-half times as volatile as labor income, yet the three share a common trend. These simple observations, combined with the further observation that deviations from this common trend display virtually no relation to consumption or labor income but are instead related to future excess stock market returns, suggests that the compensation for bearing risk changes over time in a manner quite independent from fluctuations in real activity. We argue that any economic model should fit this fact. Models that make no allowance for significant forecastable excess stock market returns, or that do so by generating an endogenous change in risk premia from a shock to consumption are unlikely to fit this fact.

Endnotes

Paper prepared for the NBER's Twenty-eighth Macroeconomics Annual. The authors are grateful to John Y. Campbell and Mark Watson for their detailed discussions, and to seminar participants at the Macroeconomics Annual conference (April 2013), UC Berkeley Haas, Columbia GSB, NYU, Stanford GSB, and to Mark Gertler, Yuriy Gorodnichenko, Jonathan Parker, Giorgio Primiceri, Eduard Schaal, and Gianluca Violante for helpful discussion. We thank Daniel Greenwald, David Kohn, Peter Gross, and Michael Weber for excellent research assistance. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <http://www.nber.org/chapters/c12932.ack>.

1. See Lettau and Ludvigson (2001, 2004, 2010).
2. Treating durables purchases purely as an expenditure (by, for example, removing them from A_t and including them in C_t) ignores the evolution of the asset over time, which must be accounted for by multiplying the stock by a gross return. (In the case of many durable goods this gross return would be less than one and consist primarily of depreciation.)
3. We use eight leads and lags of the first differences of Δy_t and Δa_t in the dynamic least squares regression. Monte Carlo simulation evidence in both Ng and Perron (1997) and our own suggested that the DLS procedure can be made more precise with larger lag lengths.
4. Standard errors do not need to be adjusted to account for the use of the "generated regressor," $\alpha'x$, 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).
5. This first-order lag length was chosen in accordance with the Akaike and Schwarz criteria.
6. We also find that the four-quarter lagged value of the cointegrating error strongly predicts asset growth. This shows that the forecasting power of the cointegrating residual for future asset growth cannot be attributable to interpolation procedures used to convert annual survey data to a quarterly housing service flow estimate, part of the services component of c_t .
7. There are many such matrices γ'_t that satisfy (6), but γ'_t is typically normalized so that it contains (as elements) as many zeros and ones as possible while still satisfying (6).
8. For simplicity, the calculations presented in figure 2 assume that labor input is fixed. The results are unchanged qualitatively if labor supply is instead elastically supplied and leisure appears nonseparably in the utility function.
9. The response of capital, k_{t+1} , is positive as in the data, but as we will see is more sluggish than the response of a_t to our interpreted productivity shock in the data. This

sluggishness follows from a well-known limitation of the canonical, frictionless business cycle model for explaining stock market behavior: the consumption-good value of a unit of installed capital is fixed at unity, so the only way wealth can adjust to an innovation in TFP is through a change in the quantity of capital, K_{t+1} , which evolves slowly over time according to an accumulation equation $K_{t+1} = (1 - \delta)K_t + I_t$, where δ is a depreciation rate and I_t is a flow of investment.

10. The ordering of a relative to c and y has no influence on the identification of the two permanent shocks. The reason is that the space spanned by the two permanent shocks is entirely summarized by c and y since these variables play no role in the error-correction necessary to restore c , a , and y to their common long-run trend. In this case, the permanent shocks will be a linear combination of c and y , and any orthogonalization designed to identify their independent components will depend only on the ordering of c relative to y .

11. Recall that we apply the Cholesky decomposition to the transformed shocks $\mathbf{G}\mathbf{e}_t$, where \mathbf{G} depends on $\boldsymbol{\gamma}'_1$. Because the elements of $\boldsymbol{\gamma}$ corresponding to the Δc_t and Δy_t equations of the VECM are statistically insignificant and therefore constrained to zero, the space spanned by the permanent shocks will be entirely driven by Δc_t and Δy_t . As a consequence, the Cholesky decomposition will render the two permanent shocks orthogonal by requiring that the second permanent shock be revealed by a movement in whichever variable among Δc_t and Δy_t is ordered second, with no contemporaneous movement in the first variable. The transitory shock will be revealed by a movement in wealth with no contemporaneous movement in c_t or y_t .

12. In principle, one could directly estimate permanent and transitory components of a larger system of variables that would include c , y , and the various components of wealth separately. We do not pursue this approach here, however, because there is no statistical evidence of cointegration in any of these larger systems where wealth components are included separately.

13. We remove the deterministic trend from the log level prior to summing over each component.

14. Recent summaries of the evidence on stock return predictability are provided in Lettau and Ludvigson (2010) and Kojien and Van Nieuwerburgh (2010).

15. The transitory shock also accounts for the majority of fluctuations in non-stock financial wealth, but this component is so stable that it plays little role in the volatility of net worth (table 4).

16. Since stock wealth has a much larger transitory component than does housing, most of the quarterly variation in net worth is driven by transitory shocks that change the *share* of stock market wealth in total net worth. Because the log of net worth is only approximately equal to a (constant) share-weighted average of the logs of the components of net worth, shocks that change the wealth shares would show up in the residual of (10), implying (for example) that a large positive transitory shock that increased the share of stock market wealth in total net worth would show up as a large negative residual in housing (and a smaller positive residual in stock market wealth). Consistent with this, the residuals in (10) for these two wealth components are highly negatively correlated.

17. Estimates of asset pricing models that suggest the return to human capital is negatively correlated with stock market returns; for example, see Chen, Favilukis, and Ludvigson (2007) and Lustig and Van Nieuwerburgh (2008).

18. For evidence of the upward break in the price-dividend ratio for the US stock market, see Lettau, Ludvigson, and Wachter (2008); Lettau and Van Nieuwerburgh (2008).

References

- Abel, A. B. 1983. "Optimal Investment under Uncertainty." *American Economic Review* 73 (1): 228–33.
- . 1990. "Asset Prices under Habit Formation and Catching up with the Jones." *American Economic Review* 80:38–42.

- Abel, A. B., and O. Blanchard. 1986. "The Present Value of Profits and Cyclical Movements in Investment." *Econometrica* 54 (2): 246–73.
- Acemoglu, D. 2002. "Directed Technical Change." *Review of Economic Studies* 69:781–809.
- Arellano, C., Y. Bai, and P. Kehoe. 2011. "Financial Markets and Fluctuations in Uncertainty." Federal Reserve Bank of Minneapolis Research Department Staff Report.
- Blinder, A. S., and A. Deaton. 1985. "The Time Series Consumption Function Revisited." *Brookings Papers on Economic Activity* 2:465–511.
- Bloom, N. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77:623–85.
- Campbell, J. Y. 1987. "Does Saving Anticipate Declining Future Labor Income? An Alternate Test of the Permanent Income Hypothesis." *Econometrica* 55:1249–73.
- 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* 107:205–51.
- Campbell, J. Y., S. Giglio, and C. Polk. 2013. "Hard Times." *Review of Asset Pricing Studies*, 2013; 3 (1): 95–132..
- Campbell, J. Y., and G. Mankiw. 1989. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." In *NBER Macroeconomics Annual: 1989*, edited by O. Blanchard, and S. Fischer, 185–216. Cambridge, MA: MIT Press.
- Campbell, J. Y., and P. Perron. 1991. "Pitfalls and Opportunities: What Macroeconomists Should Know About Unit Roots." In *NBER Macroeconomics Annual: 1991*, edited by O. J. Blanchard and S. Fischer. Cambridge, MA: MIT Press.
- Chen, X., J. Favilukis, and S. C. Ludvigson. 2007. "An Estimation of Economic Models with Recursive Preferences." Unpublished Paper, New York University.
- Christiano, L. J., M. Eichenbaum, and R. Vigusson. 2004. "What Happens After a Technology Shock?" Unpublished Paper, Northwestern University.
- Cochrane, J. H. 1991. "Production-Based Asset Pricing and the Link between Stock Returns and Economic Fluctuations." *Journal of Finance* 46 (1): 209–37.
- . 1994. "Permanent and Transitory Components of GDP and Stock Prices." *Quarterly Journal of Economics* 109:241–65.
- Constantinides, G. M. 1990. "Habit-Formation: A Resolution of the Equity Premium Puzzle." *Journal of Political Economy* 98:519–43.
- Engle, R. F., and C. W. J. Granger. 1987. "Co-Integration and Error Correction Representation, Estimation and Testing." *Econometrica* 55:251–76.
- Fernald, J. 2009. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Unpublished Paper, Federal Reserve Bank of San Francisco.
- Fisher, J. D. M. 2006. "The Dynamic Effect of Neutral and Investment-Specific Technology Shocks." *Journal of Political Economy* 114 (3): 413–51.
- Flavin, M. A. 1981. "The Adjustment of Consumption to Changing Expectations about Future Income." *Journal of Political Economy* 89:974–1009.
- Francis, N., and V. A. Ramey. 2001. "Is the Technology Driven Real Business Cycle Dead?" Unpublished Paper, University of California San Diego.
- Galí, J. 1999. "Technology, Employment and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89: 249–71.

- Gali, J., and M. Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44:195–222.
- Gilchrist, S., J. W. Sim, and E. Zakrajsek. 2010. "Uncertainty, Financial Frictions, and Investment Dynamics." Unpublished Manuscript, Boston University.
- Gonzalo, J., and C. Granger. 1995. "Estimation of Common Long-Memory Components in Cointegrated Systems." *Journal of Business and Economic Statistics* 13:27–35.
- 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): 1527–46.
- Greenwald, D., M. Lettau, and S. C. Ludvigson. 2013. "The Origins of Stock Market Fluctuations." Unpublished Paper, New York University.
- Hall, R. E. 1978. "Stochastic Implications of the Life-Cycle Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy* 86:971–87.
- Jurado, K., S. C. Ludvigson, and S. Ng. 2013. "Measuring Uncertainty." Unpublished Paper, Columbia University.
- Justiniano, A., G. E. Primiceri, and A. Tambalotti. 2009. "Investment Shocks and Business Cycles." *Journal of Monetary Economics* 57 (2): 132–45.
- . 2010. "Investment Shocks and the Relative Price of Investment." *Review of Economic Dynamics*, 14 (1): 102–21.
- King, R., C. Plosser, J. H. Stock, and M. W. Watson. 1991. "Stochastic Trends and Economic Fluctuations." *American Economic Review* 81:819–40.
- Koijen, R. S. J., and S. Van Nieuwerburgh. 2010. "Predictability of Returns and Cash Flows." Unpublished Paper, New York University.
- Kydland, F., and E. C. Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica* 50:343–60.
- Lettau, M., and S. C. Ludvigson. 2001. "Consumption, Aggregate Wealth and Expected Stock Returns." *Journal of Finance* 56 (3): 815–49.
- . 2002. "Time-Varying Risk Premia and the Cost of Capital: An Alternative Implication of the Q Theory of Investment." *Journal of Monetary Economics* 49:31–66.
- . 2004. "Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption." *American Economic Review* 94 (1): 276–99.
- . 2010. "Measuring and Modeling Variation in the Risk-Return Trade-off." In *Handbook of Financial Econometrics*, vol. 1, edited by Y. Ait-Sahalia, and L. P. Hansen, 617–90. Amsterdam: Elsevier Science B.V., North Holland.
- Lettau, M., S. C. Ludvigson, and J. A. Wachter. 2008. "The Declining Equity Premium: What Role Does Macroeconomic Risk Play?" *Review of Financial Studies* 21 (4): 1653–87.
- 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): 1607–52.
- 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* 21:2097–137.
- Newey, W. K., and K. D. West. 1987. "A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55:703–08.
- Ng, S., and P. Perron. 1997. "Estimation and Inference in Nearly Unbalanced Nearly Cointegrated Systems." *Journal of Econometrics* 79:54–81.

- Pastor, L., and P. Veronesi. 2006. "Was There A Nasdaq Bubble in the Late 1990s?" *Journal of Financial Economics* 81:61–100.
- Ríos-Rull, J., and R. Santaeulalia-Llopis. 2009. "Redistributive Shocks and Productivity Shocks." Unpublished Paper, University of Minnesota.
- Sims, C. A. 1980. "Macroeconomics and Reality." *Econometrica* 48:1–48.
- Smets, F., and R. Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian Approach." *American Economic Review* 97 (3): 586–606.
- Stock, J. H. 1987. "Asymptotic Properties of Least Squares Estimators of Cointegrating Vectors." *Econometrica* 55:113–44.
- Stock, J. H., and M. W. Watson. 1988. "Testing for Common Trends." *Journal of the American Statistical Association* 83:1097–107.
- . 1993. "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems." *Econometrica* 61:783–820.