

SECTORAL SHARES AND RISK PREMIA IN DYNAMIC EQUILIBRIUM*

Ravi Bansal[†]

Colin Ward[‡]

Amir Yaron[§]

18 November 2017

Abstract

We empirically show across several broad asset classes that sectoral wealth shares do not positively correlate with their risk premia implied by mean-variance considerations—a first-order prediction of canonical equilibrium models. We analyze these restrictions in a two-sector production economy and demonstrate that assuming imperfect substitutes and demand shocks help in matching the data, yet even with these features the model falls short of completely accounting for sectoral movements in wealth. We conclude that any models which adequately describe movements in sectoral wealth would require hedging demand to play a major quantitative role in generating these fluctuations.

*We thank Anmol Bhandari, Bob Goldstein, Harjoat Bhamra, and participants at the LAEF Economics and Business CYCLE Conference, Minnesota Macro-Asset Pricing Conference, NASMES, and UBC Winter Finance Conference for helpful comments. Yao Deng provided excellent research assistance.

[†]Fuqua School of Business, Duke University

[‡]Carlson School of Management, University of Minnesota

[§]The Wharton School, University of Pennsylvania

I. INTRODUCTION

Consider a capital market investor's asset demand equation:

$$\mathbf{w}_t = \frac{1}{\gamma} \boldsymbol{\Sigma}_t^{-1} \mathbb{E}_t[\mathbf{R}_{t+1}^e] + \left(1 - \frac{1}{\gamma}\right) \mathbf{H}\mathbf{D}_t, \quad (1)$$

where the chosen vector of wealth allocation, $\mathbf{w}_t$, depends on mean-variance considerations, a product of the inverse covariance matrix of excess returns, $\boldsymbol{\Sigma}_t^{-1}$, and the vector of excess returns, $\mathbb{E}_t[\mathbf{R}_{t+1}^e]$; and hedging demand, $\mathbf{H}\mathbf{D}_t$, the covariance between assets' excess returns and innovations of state variables which are material to investors. Thus, equilibrium wealth allocations are determined by the joint distribution of returns and state variables and not merely the first two moments of asset returns.

In equilibrium, the vector $\mathbf{w}_t$ equals the market capitalization shares of different sectors or asset classes. If sectoral wealth rises for a particular sector, then asset demand for this sector rises either due to a better mean-variance deal or from an increase in hedging demand. These considerations in allocating assets emerge in a large class of models, whether frictionless or not. Moreover, the first-order channel, namely the trade off between assets' means and variances, is at the center of a multi-sector Lucas-tree economy, a workhorse in the asset pricing literature and nested by canonical real business cycle theory.

In this paper, we begin by examining the first-order channel and find it to be profoundly rejected by the data: a sector's risk premium negatively covaries or does not covary with its wealth share. These rejections are evident in the US across households' portfolio allocations of stocks, bonds, and housing and within equity market portfolios. We discover, even more strikingly, that these deviations from mean-variance allocations are long-lived, lasting longer than the duration of an average five-year business cycle, and so potential explanations which rely on short-lived frictions are unlikely to be true.

Given the disconnect between movements in wealth shares and mean-variance dynamics, we next examine the role of hedging demands in accounting for sectoral wealth share dynamics. To do so, one requires a more comprehensive model that

goes beyond simple mean-variance theory. We therefore build a two-sector production economy, the simplest and clearest model featuring endogenous movements in wealth shares, risk premia, and hedging demand. The two dividend streams can represent flows from any characteristic-based grouping of assets, or more generally from broad asset classes like the value of corporations versus residential real estate.

We add two new elements to previous research.¹ We first allow the two streams to be imperfect substitutes. A relative price between the two goods then fluctuates endogenously, changing the distribution of risk and return across assets. We then include shocks to demand which are independent of the usual endowment or supply shocks, capturing the idea that preferences for these two types of goods fluctuates over time. In our calibration, we find that while these two elements are necessary to capture various empirical facts, observed sectoral movements are still puzzling from the vantage point of our general model which allows for both mean-variance and hedging demand considerations. In this respect, this puzzle poses a great challenge to theories in modern macroeconomics and finance as many of them share similar implications for aggregate and sectoral prices and quantities.

In analyzing the model's economic mechanisms governing sectoral risk premia and hedging demands we define the notion of an economy's *balancedness*, basically the degree to which demand is commensurate with supply for each sector across the entire economy. By calibrating solely to observed fluctuations in supply, the model cannot easily generate a negative relationship between a sector's risk premium and wealth share. The reason is that as an asset's wealth shrinks, its contribution to aggregate consumption falls, bringing shocks to its cash flows closer to being idiosyncratic and lowering its required return to the risk-free rate. Including demand shocks allows the possibility of having a high level of demand meet a scarce supply, an economic state far from balancedness. Marginal utility thus remains sensitive to a shrinking sector's wealth, keeping risk premia large.

To empirically examine the model implications we use historical data on households' portfolio allocations to tell us what the implicit process for demand shocks was by using the model's cross-equation restrictions on its distributions of capital

¹Some examples include Menzly, Santos and Veronesi (2004), Cochrane, Longstaff and Santa-Clara (2008), Eberly and Wang (2011), Martin (2013).

and wealth. Feeding the recovered process into the model's policy functions tells us that changes in hedging demand drive the majority of a sector's wealth share dynamics. Hedging demand not only constitutes the majority of an asset's allocation on average but also plays the main role in explaining shifts in its relative allocation over time. We therefore conjecture that any set of genuine economic mechanisms or frictions which adequately describe a sector's allocation dynamics must first appeal to shifts in hedging demand.

As noted earlier, our workhorse economy cannot sufficiently describe the joint behavior of wealth shares and risk premia in the data. The reason is that it requires demand shocks which generate extreme movements in expenditure shares that are noticeably counterfactual to the modest shifts observed empirically.² The virtue of our analysis is that there are few degrees of freedom in generating demand shocks and this also imposes discipline on hedging demand. While arguably important for assets in positive net supply, like equities, it is unlikely that hedging demands are the dominant influence; if they were, it would imply that the value-weighted market return, with an equity premium near six percent, is held by investors to hedge other income, such as income from labor or the non-traded corporate sector. While a possibility, this is hard to imagine. Therefore, our evidence poses a challenge to a large class of models, as even after accounting for hedging demands, it is difficult to account for the joint distribution of sectoral wealth and their risk premia.

Our model is set in a perfect neoclassical world, free of collateral and funding requirements, information processing constraints, and transaction costs. While done for tractability, these or other types of long-lived frictions could also play an important role in generating the key fact. These nontrivial extensions are outside the scope of the paper and we therefore view our study as providing a stepping stone for future work and, given the absence of imperfections, as an estimate of the quantitative hurdle faced by economists.

Our work overlaps with several strands of research. First, it builds on theoretical work that attempts to link wealth shares to expected returns. A partial list is Menzly

²Piazzesi, Schneider and Tuzel (2007) and Kongsamut, Rebelo and Xie (2001) empirically document the stability of expenditure shares across housing and non-housing wealth and across services, agriculture, and manufacturing, respectively.

et al. (2004), Cochrane et al. (2008), Eberly and Wang (2011), Martin (2013). Our model is closest to Eberly and Wang’s (2011) but extends their and the literature’s common setup to allow goods to be imperfect substitutes and the preference for individual goods to vary over time. Much of the earlier work is written in endowment economies and so not explore the puzzle empirically or quantitatively.

Our paper relates to the classic literature on intertemporal portfolio allocation: a short list is Samuelson (1969), Merton (1971), Campbell (1993), and Campbell and Viceira (2002). An important distinction of our work is that many of these models are solved in partial equilibrium, often exogenously specifying assets’ returns, while our setup is housed in a general equilibrium production economy. The adjustment of capital can be particularly important for longer investment horizons as sectoral realignments must take place.

The issue of the optimal allocation of wealth across different sectors is also important from the perspective of production-based multi-sector models. A salient margin in these economies connects holdings of sectoral wealth and the return distribution. For example, Long and Plosser (1983) and Boldrin, Christiano and Fisher (2001) consider multi-sector models but do not explore the connection between sectoral wealth and returns. Investors in these economies hold sectoral wealth largely along the mean and variance dimension, and thus the empirical puzzle we highlight in this paper provides an important challenge for this literature as well.

Our paper is written as follows. In Section II, we first document the puzzle in the data. In Section III, we present the setup of our model. In Section IV, we calibrate and analyze the model, drawing attention to the effects of demand shocks. In Section V, we infer the historical process of a demand shock implied by the model’s state variables and use it to estimate empirically the magnitude of hedging demand. We then conclude.

II. THE EMPIRICAL PUZZLE

We evaluate the puzzle by examining the joint time series behavior of risk premia and wealth shares for several groups of assets. From these series’ behaviors over time, we infer the importance of mean-variance and hedging demand in accounting

for portfolio allocations.

A concern with the approach is that we require estimating risk premia. To address this, we also study the temporal patterns of Tobin's Q and wealth shares. To the extent that an increase in an asset's risk premium lowers its Tobin's Q , then these patterns provide additional evidence of the puzzle that sidesteps estimation.

We first look at mean-variance theory's cross-sectional predictions, where an asset's wealth share should correlate with its expected return period by period. We then turn to looking at the puzzle intertemporally, as we further describe below. We present the empirical frameworks before discussing results.

A. THE PUZZLE IN THE CROSS-SECTION

We look *across* markets from the perspective of US households' collective investments in three major asset classes: debt, equity, and housing. We construct returns from market values of the three assets in the Federal Reserve's Flow of Funds and from investment income (interest, dividends and rental) earned by households recorded in the Bureau of Economic Analysis. Appendix A describes their construction in detail.³

To strengthen empirical findings, we also assess the puzzle *within* US equity markets. In this environment it is more plausible that investors face the same restrictions, if any. We separately study five industry and five book-to-market portfolios, both from Ken French's website. Industry portfolios allow us to cleanly measure Tobin's Q , but are not rebalanced overtime. Book-to-market portfolios are rebalanced every June but complicate the calculation and interpretation of Tobin's Q , and so we ignore this statistic for these portfolios.

A concise way to gauge the size of the empirical puzzle in expected excess returns and wealth shares is to scatter plot, period by period, risk premia implied by a mean-variance investor's equilibrium allocations on risk premia estimated in the data. We then look for departures from the 45-degree line, which would hold true when hedging demand is unimportant and only mean-variance considerations

³Our framework described below is set in a neoclassical world which serves as a benchmark but abstracts from some potentially important economic frictions, such as collateral constraints or rational inattention, that could be driving the empirical results.

matter.

More specifically, for every t we compare

$$\mathbb{E}_t[R^e] = \gamma \mathbf{\Sigma} \mathbf{w}_t \quad \text{with} \quad \mathbb{E}_t[R^e] = \hat{a} + \hat{b} \times D_t/P_t.$$

We calculate theoretical *share-implied risk premia* as the vector of current wealth shares, $\mathbf{w}_t$, pre-multiplied by a value of risk aversion, γ , and the constant covariance matrix of portfolio excess returns, $\mathbf{\Sigma}$. We also calculate empirically *valuation-implied risk premia* as fitted values of a regression of future annual cumulative excess returns, $\sum_{h=1}^4 R_{t+h} - R_{t+h}^f$, on the current cash flow yield, D_t/P_t , and a constant. Both the covariance matrix and regression coefficients are estimated from the full data sample.

B. THE PUZZLE AT LONG HORIZONS

In addition to analyzing cross-sectional predictions, we also examine the relationship between current wealth shares and future excess returns over long horizons. In particular, for each asset in a particular portfolio we look at regressions of the form

$$\frac{1}{H} \sum_{h=1}^H R_{t+h} - R_{t+h}^f = a_H + b_H \times w_t + \epsilon_{t+H}, \text{ for all } t \quad (2)$$

where the variable H indexes the horizon over which future excess returns are cumulated. The slope coefficient b_H measures the effect of a one percent increase in an asset's wealth share on its excess return per period over the next H periods.

Evidence of mean-variance allocations would be present if slope coefficients were positive and statistically different from zero for all horizons. If short-lived frictions exist that prevent the immediate alignment of wealth shares with mean-variance allocations but diminish at longer horizons, then slope coefficients would be initially zero yet would rise and become positive and significant as H became large. If deviations from mean-variance theory were long-lived, then coefficients would be zero or negative at both short and long horizons.

C. RESULTS

Our cross-sectional results are depicted in Figure 1. The top two panels plot the return-based empirical puzzle for the two sets of equity portfolios. The bottom-left panel repeats the exercise for households' three broad asset classes.⁴ The bottom-right panel plots changes in Tobin's Q on changes in wealth shares for the industry portfolios. We relegate the empirical relation between Q and shares for the broad asset classes to Section IV, though results are consistent with what is presented here.

All figures suggest the same phenomenon and the pattern seems robust: there is a significant departure from the mean-variance portfolio allocation. In fact, all expected return patterns suggest that mean-variance demand explains very little, as each portfolio's share-implied risk premia are effectively flat with respect to valuation-implied risk premia. The upward slopes for each industry's Tobin's Q relative to their wealth shares, moreover, corroborate the empirical fact that risk premia, all else equal, seem to decline in wealth.

We view the empirical puzzle as ubiquitous.⁵ These facts are challenging for models at the nexus of financial markets and the real economy. In standard models, risk premia generally increase in wealth shares. As the output share of a sector grows, its output contributes more to aggregate consumption, and therefore the systematic risk of the sector increases, along with its risk premium. In a standard multi-sector production economy, a sector's Tobin's Q would fall after it invests and increases its wealth share because the marginal product of capital would decline, all else equal. Neither of these common predictions are evident in the data.

We now turn to our results over long horizons and depict them in Figure 2. For each industry portfolio and the three household assets, we plot the 95 percent con-

⁴The valuation-implied expected excess return on debt may seem high. This derives from our data sample. Basically, cash flows from debt are high during the 1980s at the same time that values of bonds are low, implying high expected excess returns over this period. We nevertheless leave the debt return variable unaltered as it is consistent with our portfolio construction from the perspective of US households.

⁵In other work, Bansal, Fang and Yaron (2008) document a similar negative relationship between wealth shares and both expected returns and Sharpe ratios for the Fama-French 12 industry portfolios. There they find that considering volatility does not affect the general conclusion, so we focus on the first moment of returns here.

fidence intervals obtained in (2) over a 10-year horizon. Every asset over the entire horizon displays estimates that either do not differ from zero or else are negative. All results provide evidence that is contradictory to mean-variance theory and furthermore that these deviations are long-lived. What is striking is that equity and housing shares are negatively related to future excess returns even up to 10 years out.

Thus, not only is the puzzle ubiquitous, it is economically large and persistent. Any coherent explanation would need to address the long duration of these deviations from mean-variance theory. For example, a potential explanation which appealed to financial or informational frictions to explain household allocations would need these frictions to last for roughly a decade.

To understand the empirical results, we proceed to theoretically develop and analyze a multi-sector production economy featuring real frictions in the form of capital adjustment costs. Our theory contains insights of the Lucas-tree asset pricing model and macroeconomic Real Business Cycle theory, yet it extends these canonical paradigms to include imperfectly substitutable goods and demand shocks. For reasons related to the tractability of our production-based model, we focus on a two-sector economy, which we now describe in detail.

III. MODEL

Our depiction of the economy is shown in Figure 3. Our study centers around the portfolio choice of US households and we treat them as the owners of residential housing and the business sector, whose ownership encompasses the debt and equity of both financial and nonfinancial sectors. Each sector has capital (at market values) of qK and is financed by debt and equity (with subscripts denoting sectors when necessary). The nonfinancial sector is financed in part by bank loans L and bonds issued directly to households B_{NF} , whereas the financial sector's debt is all funded through deposits D or bond issues B_F . Borrowing from the financial sector through mortgages M or bank loans is thus netted out in our data calculations as described in Appendix A.

We abstract from financial market imperfections and so in this Modigliani and

Miller economy household leverage and collateral considerations do not play a role. We correspondingly view our exercise as a neoclassical benchmark, whose predictions should provide an assessment of the quantitative ability of frictionless models to address this puzzle.

A. PREFERENCES

The representative investor in the economy has continuation utility

$$J_t = \mathbb{E}_t \left[\int_t^\infty f(C_s, J_s) ds \right], \quad (3)$$

where C denotes aggregate consumption is a CES composite over two sectors' dividend streams D_n , $n = 1, 2$,⁶

$$C_t = \left(\Theta_t^{\frac{1}{\epsilon}} D_{1t}^{\frac{\epsilon-1}{\epsilon}} + (1 - \Theta_t)^{\frac{1}{\epsilon}} D_{2t}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}. \quad (4)$$

The aggregator takes the usual Duffie and Epstein (1992) form

$$f(C, J) = \frac{\rho}{1 - 1/\psi} \frac{C^{1-1/\psi} - ((1 - \gamma)J)^{\frac{1-1/\psi}{1-\gamma}}}{((1 - \gamma)J)^{\frac{\gamma-1/\psi}{1-\gamma}}}, \quad (5)$$

where we interpret $\rho > 0$ as the rate of time preference, $\psi > 0$ as measuring the intertemporal elasticity of substitution, and $\gamma \geq 1$ as the coefficient of relative risk aversion.

⁶All investor wealth is assumed to be tradable. Human wealth, the expected discounted value of future labor earnings, could of course affect portfolio choice. In an important paper, Bodie, Merton and Samuelson (1992) show that if labor income is riskless, then human wealth is equivalent to an investment in the riskless asset and the investor should tilt their portfolio towards risky assets holdings relative to an investor who owns only tradable assets. Campbell and Viceira (2002) further show that no matter how idiosyncratically risky labor income is, then all investors should still tilt their portfolios towards risky assets. If labor income is positively correlated with the risky asset, however, then investors should tilt their portfolios away from the risky asset. All of these results, however, apply only to one risky asset, and not a choice of how to allocate wealth across risky assets that is studied here. How assumptions of labor income affect allocations within the set of risky assets is left for future work.

The optimal consumption choice for each sector at time t is given by

$$D_{1t} = \left(\frac{p_{1t}}{p_t} \right)^{-\epsilon} C_t \Theta_t \quad \text{and} \quad D_{2t} = \left(\frac{p_{2t}}{p_t} \right)^{-\epsilon} C_t (1 - \Theta_t). \quad (6)$$

where intratemporal elasticity of substitution is $\epsilon \in (1, \infty)$, the good's relative price is p_n , and the ideal price index is $p = (p_1^{1-\epsilon} \Theta + p_2^{1-\epsilon} (1 - \Theta))^{\frac{1}{1-\epsilon}}$; with this definition we have $\sum_n p_n D_n = pC$.⁷ We normalize $p \equiv 1$ for all t granting each sector's good a time-varying relative price to the numeraire of aggregate consumption.

We specify Θ_t as a stochastic process and interpret it as a (relative) demand shock for sector one's good.⁸ Because we focus on a two-sector economy, we let Θ_t take M possible values in $0 < \Theta_1 < \Theta_2 < \dots < \Theta_M < 1$, where the collection of points form a fine grid on the unit interval. The generator matrix is $\Pi = [\pi_{\theta\theta'}]$ for $\theta, \theta' \in \{\Theta_m\}_{m=1}^M$, whose elements can simply be thought of as the probability of Θ_t moving from state θ to θ' within time Δ is approximately $\pi_{\theta\theta'} \Delta$. The process can be equivalently represented as a sum of Poisson processes:

$$d\Theta_t = \sum_{\theta' \neq \Theta_{t-}} q_{\theta'}(\Theta_{t-}) dN_t^{(\Theta_{t-}, \theta')}, \quad (7)$$

where $q_{\theta'}(\theta) = \theta' - \theta$ and $N^{(\theta, \theta')}$ are independent Poisson processes with intensity parameters $\pi_{\theta\theta'}$ with each jump in Θ_t corresponding to a change of state for the

⁷If $\epsilon \rightarrow \infty$ goods become perfect substitutes and goods' prices always equal one. If $\epsilon \rightarrow 1$, aggregate consumption becomes Cobb-Douglas in both goods. The expenditure share

$$e_{nt} = \frac{p_{nt} D_{nt}}{C_t}$$

moves with the relative quantity of good- i consumption ($\partial e_n / \partial D_n > 0$) and against its relative price ($\partial e_n / \partial p_n < 0$) if $\epsilon > 1$ —Piazzesi et al. (2007) provide evidence that housing and nonhousing expenditure are substitutes. Related, the true price elasticity of demand varies with the sector's expenditure share because we do not have a continuum of goods:

$$-\frac{\partial \ln D_n}{\partial \ln p_n} = \epsilon + (1 - \epsilon)e_n,$$

which goes to ϵ when e_n gets small, and goes to one when e_n gets large.

⁸Pavlova and Rigobon (2007) show that demand shocks can be alternatively interpreted as pure sentiment, catching up with the Joneses, or, under a special case in a multiple agent economy, as density processes reflecting heterogeneous beliefs.

Markov chain.

There are two advantages for the process for demand shocks in this way. First, it makes the solution tractable, replacing a high-dimensional partial differential equation with a system of ordinary differential equations. Second, arbitrary autocorrelations and variances of the demand shock can be set via the generator matrix, allowing us to study how the autocorrelation in demand affects returns and outcomes in the real economy.

The investor owns complete financial market claims to each sector's dividend.⁹ Complete markets allow us to highlight that our results are not generated by financial market imperfections or from some notion of background risk that is untradable. All risks can be completely hedged with portfolios of state contingent claims. All produced goods are either consumed or invested in one or any combination of the other sectors; that is, goods can be reallocated from one sector to the other.

B. TECHNOLOGY

A sector's firm produces differentiated output with a linear technology possessing constant productivity $A_n > 0$ and capital K_n

$$Y_{nt} = A_n K_{nt}. \quad (8)$$

Capital is accumulated via

$$dK_{nt} = \Phi_n(I_{nt}, K_{nt})dt + \sigma_n K_{nt}dB_{nt}, \quad (9)$$

where $\Phi_n(\cdot)$ is the sector's Penrose-Uzawa installation technology that measures the sector's efficiency in converting investment goods into capital goods and $\sigma_n > 0$ scales the variability of this efficiency which evolves with a Brownian shock B_{nt} . The shocks could be correlated with degree $\varphi \in (-1, 1)$.

The installation technology is homogeneous of degree one in investment and

⁹To complete markets we assume that the agent can trade a third asset whose return is linearly independent of the other two assets and that is held in zero net supply.

capital

$$\Phi_n(I_{nt}, K_{nt}) = \phi_n(i_{nt})K_{nt}, \quad (10)$$

where $i_n = I_n/K_n$ is the investment rate and $\phi'(\cdot) > 0$ and $\phi''(\cdot) \leq 0$. In particular, we specify adjustment costs as

$$\phi_n(i_{nt}) = i_{nt} - \delta_n - \frac{\kappa_n}{2}(i_{nt} - \delta_n)^2. \quad (11)$$

When $\kappa_n > 0$ becomes negligibly small, the economy approaches a frictionless economy in the spirit of Cox, Ingersoll and Ross (1985). When κ_n gets large, capital becomes illiquid and fixed as in the endowment economies of Cochrane et al. (2008) and Martin (2013). In steady state when $i_n = \delta_n$, adjustment costs are zero.

While attaining a maximally diversified economy is ideal to the representative investor, the act of reallocation consumes resources. This reallocation cost thus generates a tradeoff between diversification and growth (Eberly and Wang (2011)). Production also allows our agent to alter the future distribution of capital and risk in the economy, which is a potentially important consideration for our long-lived investor. In addition, the presence of adjustment costs allows us to relate demand shocks to Tobin's Q .

Each sector's firm takes the equilibrium stochastic discount factor as given and maximizes firm value. Each sector's dividends are $D_n = Y_n - I_n$ and their market values, in their own units, are the product of their capital stock and marginal q :

$$P_{nt} = q_{nt}K_{nt} \quad \text{and} \quad q_{nt} = \frac{1}{\phi'_n(i_{nt})} = \frac{1}{1 - \kappa_n(i_{nt} - \delta_n)}. \quad (12)$$

Following Hall (2001), we can interpret κ for sector n as a doubling time. If, for example, the sector's q doubles from one to two, then i exceeds δ by $1/(2\kappa)$; if this investment rate continues for an interval of time of length 2κ , then capital is expected to double if $\phi(i) \approx i - \delta$, when adjustment costs are small. This doubling time interpretation allows us to parameterize adjustment costs by using data only

net investment rates:

$$\kappa_n \approx \frac{\log(2)}{2(i_n^* - \delta_n)}. \quad (13)$$

C. SOLUTION

The second welfare theorem allows us to solve a Planner's problem to get allocations and then to decentralize it using prices to examine the behavior of assets. By the principle of optimality, we have a nested system of M ordinary differential equations

$$\begin{aligned} 0 = \max_{I_1, I_2} & \left\{ f(C, J(K_1, K_2, \theta)) + \sum_{\theta' \neq \theta} \pi_{\theta\theta'} (J(K_1, K_2, \theta') - J(K_1, K_2, \theta)) \right. \\ & + \sum_{n=1}^2 \left(J_n(K_1, K_2, \theta) \Phi_n(I_n, K_n) + \frac{1}{2} J_{nn}(K_1, K_2, \theta) \sigma_n^2 K_n^2 \right) \\ & \left. + J_{12}(K_1, K_2, \theta) \varphi \sigma_1 K_1 \sigma_2 K_2 \right\}, \text{ for } \theta \in \{\Theta_m\}_{m=1}^M \end{aligned} \quad (14)$$

where J_n , J_{nn} , and J_{12} denote the value function's first, second, and cross-partial derivatives with respect to sector n 's capital holding θ fixed.

In Appendix B, we describe features of the solution as well as steps to simplify it that exploit the model's homogeneity. Definitions to know for what follows are the parameter $\eta \equiv (\epsilon - 1)/\epsilon$ and variables weighted capital $K \equiv (K_1^\eta + K_2^\eta)^{1/\eta}$ and the first sector's capital share $k \equiv (K_1/K)^\eta$, an important continuous state variable summarizing the distribution of the supply of capital in the economy.

D. SOURCES OF RISK

The stochastic discount factor (SDF), which is the present value of an extra unit of aggregate consumption (the numeraire) at time t , takes the form given by Duffie

and Skiadas (1994):

$$\begin{aligned}\Lambda_t &= \exp \left\{ \int_0^t f_J(C_s, J_s) ds \right\} f_C(C_t, J_t) \\ &= \exp \left\{ \int_0^t \frac{\rho}{1 - 1/\psi} \left[\left(\frac{c(k_s, \Theta_s)}{v(k_s, \Theta_s)} \right)^{1-1/\psi} (1/\psi - \gamma) - (1 - \gamma) \right] ds \right\} \rho K_t^{-\gamma} e^{\zeta(k_t, \Theta_t)},\end{aligned}\tag{15}$$

where $\zeta(k, \theta) = (1/\psi - \gamma) \log v(k, \theta) - 1/\psi \log c(k, \theta)$ is a function of $v(\cdot)$ and $c(\cdot)$, which are respectively continuation utility and consumption per unit of weighted capital and defined in Appendix B. Ito's lemma then implies that our pricing factor inherits the following dynamics:

$$\begin{aligned}\frac{d\Lambda}{\Lambda} &= -r(k, \theta)dt - \lambda_1(k, \theta)\sigma_1 dB_1 - \lambda_2(k, \theta)\sigma_2 dB_2 \\ &\quad + \sum_{\theta' \neq \theta} \left(\frac{e^{\zeta(k, \theta')}}{e^{\zeta(k, \theta)}} - 1 \right) \left(dN^{(\theta, \theta')} - \pi_{\theta\theta'} dt \right).\end{aligned}\tag{16}$$

The risk-free rate r in the numeraire's units is given in Appendix B and responds to the usual features of expected growth and precautionary savings. There are three sources of risk. The first two are diffusive supply shocks:

$$\begin{aligned}\lambda_1(k, \theta) &= \gamma k - \frac{\partial \zeta(k, \theta)}{\partial k} \eta k (1 - k), \\ \lambda_2(k, \theta) &= \gamma (1 - k) + \frac{\partial \zeta(k, \theta)}{\partial k} \eta k (1 - k).\end{aligned}\tag{17}$$

Looking at the first sector's price of risk, $\lambda_1 \sigma_1$, the first piece increases linearly with k from 0 to $\gamma \sigma_1$ as k goes from zero to one. The sign of the second piece depends on $\zeta(k, \theta)$ and η . In our calibration, the first piece dominates and the first sector's price of risk effectively increases monotonically with k . A similar description holds for the second supply shock.

This monotonicity lies behind the counterfactual result of standard two-sector economies producing a positive relationship between risk premia and wealth shares (Cochrane et al. (2008)). Calibrations which feature supply shock volatility near

3 percent, an empirically reasonable value, face difficulty in matching the negative pattern observed in the data.

The third source of risk arises from demand shocks and equals the growth rate of $\zeta(\cdot)$ while holding k fixed. These shocks alter the pricing kernel through their effects on the growth rates of consumption and continuation utility. Our calibration sets the price of risk on continuation utility ($\gamma - 1/\psi > 0$) greater than on consumption ($1/\psi$), implying that our agent cares more about the smoothness of continuation utility (the long-run path of consumption) than the variability of current consumption.

There is a benefit of specifying demand shocks θ to be orthogonal to supply shocks: specifically, demand shocks allow the model to generate great variation in marginal utility through $\zeta(k, \theta)$ even when k is small. This important margin helps the model reconcile the key patterns that we see in the data.

With our pricing kernel in hand, we can now study the risk and return properties of the assets and distinguish the sources of asset demand.

E. RISK PREMIA AND PORTFOLIO CHOICE

For what follows, it is convenient to define a sector's Tobin's Q in units of the numeraire as $Q_n \equiv p_n q_n$. The return on a stock with price P_n in units of the numeraire is then

$$\begin{aligned} dR_n &= \frac{d(p_n P_n)}{p_n P_n} + \frac{D_n}{P_n} dt = \frac{dQ_n}{Q_n} + \frac{dK_n}{K_n} + \frac{D_n}{P_n} dt + \text{second-order terms} \\ &+ \sum_{\theta' \neq \theta} \left(\frac{Q_n(k, \theta')}{Q_n(k, \theta)} - 1 \right) dN^{(\theta, \theta')}. \end{aligned} \quad (18)$$

The dividend yield, $D_n/P_n = (A_n - i_n)/q_n$, is unit-free, but the expected capital gains component depends on the quantity and value of capital. Holding fixed a sector's capital stock, a demand shock affects its return through Tobin's Q .

The risk premium is a function of a sector's capital share and the relative de-

mand for its good. For the first sector, it takes the form

$$\begin{aligned}
\left(\frac{1}{dt} \mathbb{E}_t[dR_1] - r \right) &= -\frac{1}{dt} \mathbb{E}_t \left[\frac{d\Lambda}{\Lambda} dR_1 \right] = \lambda_1 \sigma_1 \underbrace{\left(\frac{\partial Q_1 / \partial k}{Q_1} \eta k (1 - k) (\sigma_1 - \varphi \sigma_2) + \sigma_1 \right)}_{\equiv \beta_{11}} \\
&+ \lambda_2 \sigma_2 \underbrace{\left(\frac{\partial Q_1 / \partial k}{Q_1} \eta k (1 - k) (\varphi \sigma_1 - \sigma_2) + \varphi \sigma_1 \right)}_{\equiv \beta_{12}} + \sum_{\theta' \neq \theta} \pi_{\theta\theta'} \underbrace{\left(\frac{e^{\zeta(k, \theta')}}{e^{\zeta(k, \theta)}} - 1 \right)}_{\equiv \lambda_{Demand}} \underbrace{\left(1 - \frac{Q_1(k, \theta')}{Q_1(k, \theta)} \right)}_{\equiv \beta_{1Demand}},
\end{aligned} \tag{19}$$

where for betas the first subscript indexes the sector and the second indexes the risk source. The second sector's premium has a similar representation and is omitted for brevity. If a demand shock occurs whereby marginal utility rises while a sector's Q falls, then the agent will require positive risk compensation to hold that sector. Furthermore, if marginal utility could potentially rise especially high when a sector's capital share is small, then the small sector could command high risk premia, a pattern that is consistent with the data.

Next, the first sector's wealth share is

$$w = \frac{p_1 q_1 K_1}{p_1 q_1 K_1 + p_2 q_2 K_2}, \tag{20}$$

leaving the other sector with share $1 - w$ of total wealth $W = \sum_n p_n q_n K_n$. We collect these wealth shares in the vector, $\mathbf{w} = [w, 1 - w]'$, and define the agent's mean-variance portfolio demand as

$$\mathbf{MV} = \frac{1}{\gamma} \times \begin{pmatrix} \text{var}_t(dR_1) & \text{cov}_t(dR_1, dR_2) \\ \text{cov}_t(dR_1, dR_2) & \text{var}_t(dR_2) \end{pmatrix}^{-1} \begin{bmatrix} \mathbb{E}_t[dR_1] - r dt \\ \mathbb{E}_t[dR_2] - r dt \end{bmatrix}. \tag{21}$$

We compute hedging demands as the difference between wealth shares $\mathbf{w}$ and the

mean-variance demands¹⁰

$$\mathbf{HD} = \mathbf{w} - \mathbf{MV}. \quad (22)$$

Hedging demand measures how much the agent holds of the sector's market value in excess of its myopic mean-variance tradeoff. An investor with a long horizon is averse to news that future returns will be lower, because their wealth or consumption will be. The investor will therefore bid up the prices of stocks that do well on such news, hedging this risk. Thus, equilibrium expected returns will depend not only on covariation with the current market return, but also on covariation with the news of future returns.

IV. CALIBRATION AND ANALYSIS

Our chosen parameters are in Table I. We begin by choosing household tastes that are consistent with prior literature in production economies with recursive preferences (Croce (2014) and Kung and Schmid (2015)) and also empirical literature

¹⁰Further analytical progress on the determinants of hedging demand is unobtainable in our jump-diffusion setup. To understand why, the evolution of wealth can be written as

$$dW = W \left(\sum_{n=1}^2 w_n (dR_n - r dt) \right) - C dt + r W dt,$$

where the wealth shares are defined according to (20). The optimal portfolio choice for a solution $J(W)$ requires the joint solution of a two-equation system,

$$\begin{aligned} 0 = & W J'(W) \left(\frac{1}{dt} \mathbb{E}_t[dR_n^c] - r \right) + J''(W) W^2 \frac{1}{dt} (w_n \text{var}_t(dR_n) + w_{m \neq n} \text{cov}_t(dR_n, dR_m)) \\ & + \sum_{\theta' \neq \theta} \pi_{\theta\theta'} J' \left(W \sum_{n=1}^2 w_n \frac{Q_n(\theta')}{Q_n(\theta)} \right) W, \text{ for } n = 1, 2, \end{aligned}$$

where dR_n^c is the continuous part of the return process in (18). The last term makes the problem nonlinear and unable to be solved in closed-form except in special cases (Ait-Sahalia, Cacho-Diaz and Hurd (2009)). While the closed-form could be in essence achieved by specifying a Brownian-driven continuous-time process that approximates our discrete jump process for demand shocks—we thank Harjoat Bhamra for this suggestion—we leave these calculations for other work. Instead, we proceed to intuitively relate systematic risk, risk premia, and hedging demand with depictions of the solution to the calibrated model.

on estimates of the value of intertemporal elasticity of substitution for stockholders (Attanasio and Vissing-Jorgenson (2003)): Risk aversion is $\gamma = 10$ and IES is $\psi = 1.1$. The rate of time preference is set to $\rho = 0.02$. For the important parameter ϵ that determines the magnitude of relative price variation and the importance of demand shocks, we choose a value of $\epsilon = 3$, which is within the range of estimates given by Ogaki and Reinhart (1998), who find $\epsilon \in (2.9, 4.0)$ across nondurable and durable goods when using a cointegration approach on long series of annual data (1929-1990), and Ravn, Schmitt-Grohe and Uribe (2006), who estimate the value $\epsilon \in (2.5, 5.3)$ in their deep habits model.

We then calibrate the economy's technologies that drive the capital state variables to fit the data. Specifically, we take observed investment, depreciation, and real growth rates of capital stocks from the BEA's long data sample and use them to pin down parameters corresponding to average investment and depreciation rates, supply shock volatility, and the curvature of the adjustment cost function. We define k to measure the share of business sector's capital stock in the economy, leaving $1 - k$ for housing's.

For what follows, we refer respectively to the business sector first and the housing sector second. Average depreciation rates δ are calculated and simply set to 7.5 and 2 percent. The volatilities σ are equated to the data's estimate of the volatility of the growth rate: 3.2 and 3.4 percent; the shocks are left uncorrelated as is common in the literature for clarity, $\varphi = 0$. We set adjustment cost parameters based on net investment rates via (13), putting κ at 12 and 14 for business and housing, reflecting a doubling time of about 24 and 28 years. These parameters are close to the average estimates obtained by directly calculating the doubling times using overlapping data of the real capital stocks, which suggest κ should be 10 and 12. Finally, we calibrate productivity A to 0.15 and 0.1 to match the average investment rates of each sector in their respective one-sector economies via (B8).

To calibrate the transition matrix and values of the demand shock, we use the quadrature method of Tauchen and Hussey (1991) to approximate an annual AR(1) with a mean of 0.5, an unconditional standard deviation of 0.1, and a persistence of 0.9. We chose the mean of 0.5 for symmetry and because we do not have a strong prior for the long-run average share of household expenditure on the housing sector

relative to the business sector. The unconditional standard deviation was chosen so that $\{\Theta_m\}$ is contained in the unit interval. The choice of persistence pins then down the conditional volatility of the shock, $0.1 \times \sqrt{1 - 0.9^2} = 4.36$ percent per year, which is about one percentage point higher than the average supply shock volatility. We convert the discrete-time Markov chain to continuous time using Jarrow, Lando and Turnbull's (1997) approximation that is based on assumption of more than one change of state is close to zero within the period. We tabulate the state space of Π and its ergodic distribution in Table II.

A. OUTLINE OF ANALYSIS

We split the analysis into three parts. In the first part, we discuss the intuition of the model, in particular the trade-off between growth and the investor's desire to achieve a balance in the economy. We next examine risk premia and Tobin's Q at the sector level, drawing a contrast between a model with both types of shocks and one with only supply shocks, the literature's usual benchmark. We finally decompose risk premia, highlighting the demand shocks's contribution and the role hedging demand plays in reconciling the main empirical patterns that we see in the data.

We plot several of the model's policy functions below with k on the horizontal axis. Because k and θ are orthogonal, holding θ fixed, a shift *along* a policy function changes only the distribution of capital; holding k fixed, a shift in demand is a shift *across* policy functions. Of course, an economy that only has supply shocks will feature different policy functions than the ones below because it would exclude risk premia and all the general equilibrium effects attributed to demand shocks.

To complement the theory's qualitative predictions in our analysis, we do the following. We simulate the model 5,000 times for 480 quarters each, burning in the first half and leaving 60 years of quarterly data, a length comparable to our data sample. Based on these simulations we compute statistics which allow us to quantitatively evaluate the model. Summary statistics on commonly-studied macroeconomic and financial variables are tabulated in Table III.

Adding demand shocks, in short, helps the model generate dynamics that are consistent with the data. Without them, cash flow volatility is basically equal capital stock volatility, in stark contrast to the data. Demand shocks add variation to goods'

relative prices while keeping the supply side of the economy unchanged. Of course, adding another shock raises the average excess return and return volatility of each sector, but the introduction of demand shocks, as we will discuss, adds more than simply including another otherwise identical supply shock.¹¹

One restriction faced when choosing ϵ and a process for demand shocks is that consumption volatility increases as relative prices become more volatile. As is standard in production-based models, it is hard to accurately model both aggregate consumption volatility and sectors' dividend volatilities. Since we are primarily interested in sectoral movements, we choose to focus on the dynamics of sectoral cash flows.

A data measure theoretically close to what our “retired investor” would consume would be cash flows earned from the housing and business sectors ($C = p_1 D_1 + p_2 D_2$). To calculate the growth rate of this measure in the data, we simply construct a time series of $dC/C = d(p_1 D_1)/(p_1 D_1) \times s + d(p_2 D_2)/(p_2 D_2) \times (1 - s)$, where $s = p_1 D_1/(p_1 D_1 + p_2 D_2)$ is the business sector's cash flow share of total cash flows, which is calculated period by period. The volatility of this empirical consumption growth rate is 6.8 percent, whereas the model produces a volatility of 5 percent.¹²

Within the model's 99 percent confidence interval of (2.8, 7.8), moreover, the model's standard deviation of consumption falls in line with the usual measure of consumption based on services and nondurables expenditure, estimated to be 2.9 percent annually over the period from 1929 until 2015. Taken together, we believe our current calibration does not substantially overstate consumption volatility.

¹¹Note that the high average return and Sharpe ratio of the business sector is due to two effects. First, the portfolio's return is a market-weighted average of debt and equity. Second, we take the data as is from the BEA for interest income and from the Flow of Funds for market values. The average cash flow yields over our sample on debt and equity are 11 and 5 percent, respectively.

¹²Related, Vissing-Jorgensen (2002) estimates the volatility of the consumption expenditure growth of stockholders to be between 4 and 10 percent using data based on the US Consumer Expenditure Survey. Additionally, our model aggregates consumption expenditure from both business goods and services and housing services. Part of this expenditure would therefore be spent on durable goods, not only on nondurables and services. Gomes, Kogan and Yogo (2009) report that the volatility of durable goods ranges from 8 to 17 percent.

B. BALANCEDNESS

The representative agent faces a trade off between growth and the desire to achieve a balanced economy.¹³ A balanced economy occurs when a sector has a capital share that is close in size to the demand for its good; for the first (business) sector for example, when $k_t \approx \Theta_t$. For brevity we'll refer to this property simply as *balancedness*. The economy features unbalancedness when either a small capital share produces to meet a high demand, or vice versa, or more generally when there simply exists a sectoral imbalance of supply and demand. Furthermore, the economics differ when demand is high and supply is low for a sector than from when demand is low and supply is high, so asymmetries exist.

When the economy is balanced, the market risk premium, market volatility, and the demand for precautionary savings reach local lows, leading the agent to consume a lot out of wealth. Figure 4 plots these variables on the business sector's capital share k for three levels of demand for the sector's good, $\theta \in \{\Theta_1, \Theta_5, \Theta_9\}$. Each (k, θ) pair has a maximal point of balancedness. The three minima of the market risk premium across k for each θ align closely with these points.

Our representative agent is long-lived and has recursive utility, so the consumption-wealth ratio depends on the expected discounted value of all future market returns. Because our calibration has $\psi > 1$, the substitution effect dominates the wealth effect, leading the agent to prefer postponing consumption for investment when future market returns are expected to be high. During these high market return periods, then to a first-order the consumption-wealth ratio is low.

The agent prefers a balanced economy and can invest to attain this in our production economy. But adjustment is costly as a sector's output is diverted to invest-

¹³This desire originates from the consumption-to-weighted capital ratio $c(k, \theta)$ in (B4). Loosely speaking, when k_t and Θ_t are high, then i_{1t} would be chosen to be low, making the dividend-to-capital ratio, $(A_1 - i_{1t})$, large. (In our two-sector economy this is equivalent to $1 - k_t$ and $1 - \Theta_t$ being both small, so analogous arguments hold for the other sector.) This choice has two effects. First, $c(k, \theta)$ would be high because a sector with abundant capital produces a good in high demand. Second, the volatility of consumption would be low because the Cobb-Douglas function of Θ_t and $(A_1 - i_{1t})$ would feature small marginal "products", reducing the volatility of these components. These are therefore good times. In contrast, if k_t were high but Θ_t low, then a small dividend-to-capital ratio would be chosen and similar arguments would imply a low $c(k, \theta)$ and a high volatility of dC/C .

ment at an increasing marginal cost. There is therefore a tradeoff between growth and restoring balancedness, exemplified by the drift of the business sector's capital share, $\mathbb{E}[dk]$. It is either positive or negative parabola, or a sine curve, depending on the level of demand. The economy is therefore mean-reverting and in the long-run features a nondegenerate ergodic distribution. When demand for the business good increases, the agent chooses to consume more of it at the expense of investing in production, eventually forcing $\mathbb{E}[dk]$ negative and k towards zero, thereby unbalancing the economy even further.

The bottom-right panel shows that when demand increases for the business sector, its wealth share increases relative to its capital share; a similar pattern holds for housing. This plot shows a sector's wealth share closely tracks its capital share but diverges depending on the level of demand.

To preview the asymmetry in the model, contrast how an increase in business demand (Θ rises) were to affect the economy conditional on initially there being excess supply ($\Theta_t < k_t$) versus excess demand ($\Theta_t > k_t$): in the former, the agent would choose to consume relatively more out of wealth than before, as demand catches up with supply, balancedness is restored, and aggregate volatility declines; the latter would lower relative consumption, as the agent would prefer to invest to restore balancedness. These two cases differentially affect the consumption-savings decision, marginal utility, and the pricing of risky assets in the economy, as we'll now discuss.

C. RISK PREMIA AND TOBIN'S Q

Demand shocks not only affect the marginal utility by altering the consumption-savings decision, they also affect Tobin's Q by shifting a sector's relative price p and marginal value q . Figure 5 shows that holding θ fixed, a sector's Tobin's Q decreases with its capital share, simply because a marginal unit of capital is valued less as the sector's capital base grows. With only supply shocks, then, it is likely there is always a negative relationship between a sector's Q and its wealth share, a fact at odds with the positive empirical relation.

On the other hand, when fixing supply and varying demand, Tobin's Q will reflect shifts in risk premia. A change in demand has different effects on risk premia

and Q depending on the economy's distribution of capital. When a sector is small, an increase in demand further imbalances the economy and lowers Q as the sector's risk premium rapidly rises; when it's large, increasing demand could restore balancedness, lowering the risk premium while raising Q .

We'll decompose the SDF, risk exposures, and risk premia in the next section to shed more light on the economics underlying the patterns, but for now it suffices to know that demand shocks could resolve the two main empirical facts. To quantify this, we simulate our model and, as in the data, run time series regressions of a sector's excess returns or Q on its wealth share. Specifically, in each simulation we run the regressions of

- (i.) $\frac{1}{dt}\mathbb{E}_t[dR_{nt}] - r_t = a_n + b_n \times w_{nt} + \epsilon_{nt}$, for all t ; and
- (ii.) $Q_{nt} = a_n + b_n \times w_{nt} + \epsilon_{nt}$, for all t ,

for each sector n . We also run regressions on the variables' first-differences. Furthermore, to isolate the effects of demand demand, we also run these regressions when the model is simulated only with supply shocks. We tabulate the distribution of slope coefficients in Table IV and compare them to point estimates from the data.

As expected, for regressions of Q on wealth, supply shocks alone always generate a negative relation. Demand shocks are required to match the data. The model generates this by having a small level of demand catch up with a large supply of the asset. When a sector has a large capital share, an increase in demand for that good is good news for the agent, lowering macroeconomic risk and the required risk premium. A falling risk premium pushes up Q .

A model featuring only supply shocks also cannot replicate the negative data pattern between risk premia and wealth. The reason is that as an asset's wealth shrinks its contribution to aggregate consumption falls, bringing shocks to its cash flows closer to being idiosyncratic and lowering its required return to the risk-free rate. Therefore a positive relationship holds.¹⁴

¹⁴Martin (2013) is able to break this monotonicity in a CRRA endowment economy with a mix of high risk aversion ($\gamma \geq 5$) and high exogenous dividend volatility (10 percent). Our model does not generate the same effect because we calibrate to only supply shock and not cash flow volatility, so the amount of total risk in the economy is small. There are, however, two problems with this intuition

This idiosyncratic cash flow effect is still present in a model with demand shocks, but introducing the possibility of imbalance between supply and demand creates a new source of risk to the economy over and above additional shock volatility. More specifically, demand shocks matter more, and affect marginal utility and risk premia more, when supply is scarce, and the level of risk premia reflect this heightened sensitivity. Imbalance can thus generate high risk premia at low k , and if supply then chases demand, one would see a negative relationship between risk premia and wealth shares which is observed in the data.

In summary, supply shocks alone generate counterfactuals with respect to the two main regressions because systematic risk premia shrink as a sector becomes small. Demand shocks break this counterfactual by ensuring marginal utility stays sensitive to even a small sector's output. Further, our regression results suggest that the magnitude of our demand shocks, given our specification of preferences and technology, are capable of generating a distribution of regression coefficients that are comparable with what we observe in the data. We now turn to decomposing risk premia.

D. DECOMPOSING RISK PREMIA

Since in our calibration both sectors' Q 's and risk premia largely have the same graphs with respect to k or $1 - k$, we decompose only the first (business) sector's risk premium for brevity. We begin by looking separately at the pieces of the prices of risk and betas defined in (19) in Figure 6. To plot the pieces related to demand we take expectations conditional the current (k, θ) pair; that is, $\mathbb{E}[\lambda_{Demand}|k, \theta] = \sum_{\theta' \neq \theta} \pi_{\theta\theta'} \left(\frac{e^{\zeta(k, \theta')}}{e^{\zeta(k, \theta)}} - 1 \right)$ and $\mathbb{E}[\beta_{1Demand}|k, \theta] = \sum_{\theta' \neq \theta} \pi_{\theta\theta'} \left(1 - \frac{Q_1(k, \theta')}{Q_1(k, \theta)} \right)$. The product of these expectations is *not* the demand shock's risk premium, but are computed and displayed for intuition.

The market prices of risk, $\lambda_1\sigma_1$ and $\lambda_2\sigma_2$, effectively increase linearly with k

and result holding up in a more general model. First, a low IES combined with a reasonable dividend growth volatility induces counterfactually large variation in the risk-free rate. Second, when moving from an endowment to a production economy, a small asset that has a high price-dividend ratio would also have a high Tobin's Q , indicating profitable investment. If investment were allowed, it would reduce the sector's valuation, its potential for comovement with the market, and therefore its risk premium.

and $(1 - k)$, respectively, according to (17). The pattern on betas for each of these shocks is also straightforward. The business asset's exposure to the first shock β_{11} increases in both k and θ , all else equal, as the asset's return becomes more and more exposed to the sector's fortunes as it grows in wealth. The beta on the second sector β_{12} declines in both k and θ for a similar reasoning. Both of these effects are present in the perfect substitutes models of Cochrane et al. (2008) and Martin (2013).

What is new in our paper can be seen in the bottom two subpanels that display the approximations to the price and quantity of demand risk. For over a large range of k and for extreme values of θ (when θ is either near 0 or 1), λ_{Demand} is negative—it has a negative price of risk. The interpretation is that when the economy is unbalanced at either low k and high θ or high k and low θ , marginal utility is expected to *fall*, arising from the demand shock being mean-reverting and restoring balancedness. It is thus natural to think about demand shocks' economics as how they affect balancedness and not simply as a positive or negative realization in absolute terms.

Fixing $\theta = \Theta_9$, as k shrinks, the economy becomes more unbalanced and the risk of a demand shock requires greater compensation, increasing the magnitude of λ_{Demand} . As k rises, the magnitude of λ_{Demand} shrinks and eventually crosses zero and becomes positive. It becomes positive because when the first sector has a large wealth share and is basically the whole economy, a decline in demand for the good would increase marginal utility because effectively consumption $C \approx p_1 D_1$ has fallen. Similar to this, when $\theta = \Theta_5$, then the price of risk is positive as the agent needs to be compensated for exposure to this additional risk source as in standard models.

The first asset's exposure to the demand shock is also a novel finding: again for a large range of k and for extreme values of θ , $\beta_{1Demand}$ is negative, implying that the asset is held to hedge the demand shock. If a high demand is expected to fall at low k , then Q rises as the market risk premium falls. Thus, a small asset is held as protection against a future decline in the investment opportunity set, as proxied by the level of the market risk premium. Conversely, if a low demand is expected to increase at high k , then Q rises again as the market risk premium falls, so hedging

is not confined to the small asset. As unbalancedness in general gets worse, this hedge-beta gets larger in magnitude. On the other hand, when k increases enough, then the economy becomes more balanced and the beta of the asset crosses zero and becomes positive.

The presence of these large hedge betas naturally leads to a hedging-demand interpretation that we discuss in the next section. But first, the product of the negative beta and the negative price of risk suggests that demand shocks increase the demanded risk premia in the economy. To understand how important demand shocks are, we decompose the three sources of risk premia in Figure 7.

The figure shows several things. First, by far demand shocks account for the majority of the variation in risk premia, even when $\theta = \Theta_5$. Nonetheless, at $k = \Theta_5$, that is, when the economy is at balance, the magnitude of the risk premium attributed to demand shocks is small (as shown in the two bottom subpanels of Figure 6). Generally, as the economy becomes more unbalanced, demand's contribution to risk premia grows even more and ultimately dwarfs supply's contribution. Finally, from the lens of our production economy featuring constant returns to scale, the variation in Tobin's Q should then be primarily driven by demand shocks and hedging considerations attributed to them. Therefore, our results suggest that hedging demand is potentially important to consider when trying to understand the dynamics of investment.

E. SOURCES OF ASSET DEMAND

An ICAPM argument would say that an average investor would bid up the price of an asset that acts as a hedge against adverse realizations of state variables. A long-lived investor would be averse to news that future returns are lower, because their long-term wealth or consumption will be lower. Assets that pay off when the path of the expected market return falls (or the path of market volatility is expected to increase) would therefore act as hedges against a deterioration in the (future) investment opportunity set. The average investor would therefore bid up these assets' prices, lowering the share of their market values attributed to myopic demand while raising hedging demand's share.

Figure 8 plots each sector's mean-variance demand and hedging demand as a

function of its wealth w . The two dashed black lines are the zero line and the 45-degree line. By construction, the sum of each sector's mean-variance demand and hedging demand equals the sector's wealth share. When the economy is balanced, hedging demand is close to zero and for each θ reaches a local minima across k . Moving away from these minima in general increases hedging demand.

To understand why, consider a case of excess demand ($k_t < \Theta_t$) for the *business* good. During this period, the agent holds *housing* increasingly as a hedge. Why? Eventually balancedness will occur, from either a rising k_t or a falling Θ_t , and when it does the expected return on wealth will be low; ie, future returns face reinvestment risk. Housing pays off in either scenario: when Θ_t falls, the relative price of housing rises simply because the demand for it, $(1 - \Theta_t)$, does; when k_t rises, the relative supply of housing falls, again raising its relative price for a given demand—it is simply a good hedge against the looming reinvestment risk. A similar story analogously holds for the excess supply case in our two-sector economy.

Because the economy features unbalancedness the level of hedging demand is magnitudes more than in a model that only has supply shocks. For example as Figure 8 shows, near a 40 percent wealth share and $\theta = \Theta_1$, the level of business sector wealth is *entirely* attributed to hedging demand. To emphasize that demand shocks create these phenomena, we record the level of hedging demand for each sector during our simulations and compare them with those obtained in a simulated supply-shock-only economy (we resolve the model holding $\Theta = 0.5$). As seen in Figure 9, the inclusion of demand shocks, and the potential for imbalance of supply and demand, generate a markedly different pattern of the magnitude of hedging demand. In contrast, the model that only has supply shocks generates a small amount of hedging demand on the order of plus-minus 8 percent.

The analysis in the previous two sections suggests that demand shocks, risk premia, Tobin's Q and hedging demand all intertwine and thus jointly modeling their dynamics is the key to understanding the empirical puzzle. In the next section, we now try to figure out these important economic variables in our particular data sample.

V. APPLICATION: INVERTING DEMAND SHOCKS

As a final analysis, we use the model framework to infer the level of demand Θ_t from our observed data. We do this in the following steps:

1. We take our data series for both the business sector's capital share $\{k_t\}$ and wealth share $\{w_t\}$ over the period 1951Q4 until 2015Q4.
 - The capital share is the real stock of nonresidential fixed assets over itself plus the real stock of residential fixed assets, taken from BEA data.
 - The wealth share is the business sector's market capitalization share of total household wealth, taken from Flow of Funds data.
2. For every quarter in the data we map exactly the data's k_t into the model's k .¹⁵
3. Given k_t , we then, for each quarter, infer Θ_t by matching the data's w_t as close to our model's policy function $w(k_t, \Theta_t)$ by minimizing the distance between the two measures.

Thus, the model's choice of Θ_t is restricted by first exactly fitting the capital distribution k and then best fitting the wealth distribution w . With our bivariate time series of $\{k_t, \Theta_t\}$, we can then plot salient model objects over time. The output is depicted in Figure 10. The top panel shows the time series of $\{k_t\}$. As constructed, the model matches the data exactly quarter by quarter. The bottom panel plots the business sector's wealth share $\{w_t\}$ on the left axis and the implied level of demand for the business sector's good $\{\Theta_t\}$ on the right axis.

Here the model cannot match the wealth series exactly because it does not generate enough variation in wealth shares separately from variation in the capital shares (see the discussion regarding the bottom-right panel of Figure 4). Decreasing the

¹⁵The theoretically correct share is $p_1 K_1 / (p_1 K_1 + p_2 K_2)$, not simply k . We chose k because it sidesteps the need for relative prices and the effects of demand shocks, thereby allowing for a direct mapping to supply side of the economy. Moreover, given the modest fluctuation of k_t in the data, the discrepancy between the theoretically correct share and simply k is small.

elasticity across goods ϵ would shrink the gap between data and model, as relative price shifts would be amplified, but would come at the cost of increasing consumption volatility above what a typical calibration would target.

More importantly, the required movements in demand shocks vastly overstate variation in consumption expenditure shares seen in the data which are largely stable (compare the model's restrictions in (6), which are discussed further in footnote 4, with the evidence in Piazzesi et al. (2007) and Kongsamut et al. (2001)). We conclude that although the introduction of demand shocks significantly improves the model implications, it still falls short of truly reconciling the movements observed in the data. Nevertheless, we proceed to analyze other model quantities and prices in hope of identifying potential important features to consider in future work.

Given our time series of $\{k_t, \Theta_t\}$ we plot other model functions in Figure 11. The top and middle plots track Tobin's Q and the risk premium for the business sector, and the bottom plot shows both sectors' hedging demand as a percentage of each sector's wealth. All model series are smoothed with a 12-quarter moving average for clarity.

During the 1970s the model captures the decline in the business sector's wealth share with a fall in demand for its good. The economy becomes unbalanced and risk premia rise, lowering Tobin's Q . Anticipating a reversion in demand, the business sector becomes the primary hedge asset, where its share of hedging demand to its wealth more than doubles from 40 to 100 percent. Driven by a expansion in demand, the recovery of business's wealth share during the 1990s restores the supply-demand balance and the dot-com bubble manifests itself as a boom in Q with a compression of risk premia. During this episode, housing becomes the primary hedge asset—it is expected to pay off when business demand falls in the future.

The bottom panel depicts the model's requirement for hedging demand fluctuations in the time series. It makes up to about 73 percent and 75 percent of the business and housing sector's wealth allocations on average. At times, each sector's hedging demand contributes effectively 100 percent to the sector's wealth. The massive empirical literature that studies only asset returns and variances is therefore focusing on a relatively small fraction of the determinant of investor asset demand.

The majority of a sector's wealth is held for hedging, not for mean-variance considerations. We conjecture that any set of genuine economic mechanisms or frictions which adequately describe the dynamics of wealth shares must first appeal to shifts in hedging demand.

VI. CONCLUSION

A large class of economic models, whether frictionless or not, squarely place the trade off between assets' means and variances at the heart of an allocation decision. Investors should be compensated to hold an expanding asset's share by receiving some combination of a greater expected excess return or a smaller variance. In this paper, we find this first-order implication to be profoundly rejected by the data, where sectors' wealth shares negatively relate to their risk premia but positively correlate with their Tobin's Q s. These patterns emerge both within equity markets and across households' broad asset allocations, and, even more strikingly, these mean-variance deviations appear to be long-lived, lasting longer than the duration of an average business cycle.

We believe that these findings are great challenges to theories in modern macroeconomics and finance. As a first step towards an understanding, we develop a model that partially reconciles these patterns by subjecting imperfectly substitutable goods to demand shocks. This extension creates the possibility of marginal utility remaining sensitive to demand shocks when supply is scarce, and therefore it generates variation in risk premia and Tobin's Q unrelated to the distribution of capital.

But our honest modeling attempt at a reconciliation falls short of truly reconciling the joint fluctuations of wealth shares and risk premia that we see in the data. This is because the model requires demand shocks that generate extreme movements in expenditure shares which are counterfactual to the modest changes observed empirically.

A nontrivial extension of the frictionless neoclassical economy we consider is outside the scope of this paper and we therefore view our study as providing a stepping stone for future work. In particular, our analysis suggests that any set of economic mechanisms which can adequately describe these fluctuations will need

to first quantitatively appeal to changes in hedging demand. In our sample, we find that on average 73 and 75 percent of business and housing wealth, respectively, arises from the demand to hedge. Shifts in hedging demand, moreover, account primarily for the changes in these allocations over time.

We are left with questions. From where does hedging demand originate? And what exactly does a particular asset hedge? The ICAPM predicts that it arises from investors' desires to hedge marginal utility against changes in state variables. If an asset's return is good for bad realizations of a state variable, this raises the desirability and thus overall allocation to the asset. People will prefer to own stocks if they pay off when housing falls. Some of these state variables are omitted here. Investor-specific risk ranges from holding a job to running a small business. Workers in business-cycle sensitive industries will prefer to hold assets that do well in recessions. Entrepreneurs, for some reason, retain large fractions of their wealth in their own enterprises, bear a great deal of idiosyncratic risk, and hold largely undiversified portfolios.

In sum, macrofinancial theories should examine in more detail the nature of hedging demand; in particular, delineating which assets are held to hedge against what risks and why and how these time-series and cross-sectional patterns relate. Microempirical work needs to uncover how these shocks to state variables determine an individual or institutional investor's portfolio decisions and to what extent market imperfections play a role. We leave these important topics to future work.

REFERENCES

- Ait-Sahalia, Yacine, Julio Cacho-Diaz, and T.R. Hurd**, “Portfolio Choice with Jumps: A Closed-Form Solution,” *The Annals of Applied Probability*, 2009, 19 (2), 556–84.
- Attanasio, Orazio P. and Annette Vissing-Jorgenson**, “Stock Market Participation, Intertemporal Substitution and Risk Aversion,” *American Economic Review, Papers and Proceedings*, May 2003, pp. 383–91.
- Bansal, Ravi, Ed Fang, and Amir Yaron**, “Equity Capital: A Puzzle?,” March 2008. Working paper.
- Bodie, Zvi, Robert C. Merton, and William Samuelson**, “Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model,” *Journal of Economic Dynamics and Control*, 1992, 16, 427–49.
- Boldrin, Michele, Lawrence Christiano, and Jonas Fisher**, “Habit Persistence, Asset Returns, and the Business Cycle,” *American Economic Review*, 2001, 90, 149–66.
- Campbell, John Y.**, “Understanding Risk and Return,” *Journal of Political Economy*, 1993, 104 (2), 298–345.
- and **Luis M. Viceira**, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press, 2002. Clarendon Lectures in Economics.
- Cochrane, John H., Francis Longstaff, and Pedro Santa-Clara**, “Two Trees,” *Review of Financial Studies*, 2008, 21 (1), 347–85.
- Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross**, “A Theory of the Term Structure of Interest Rates,” *Econometrica*, 1985, 53 (2), 385–407.
- Croce, Mariano M.**, “Long-run Productivity Risk: A New Hope for Production-Based Asset Pricing?,” *Journal of Monetary Economics*, September 2014, 66, 13–31.

- Duffie, Darrell and Costis Skiadas**, “Continuous-Time Security Pricing: A Utility Gradient Approach,” *Journal of Mathematical Economics*, 1994, 23, 107–131.
- and **Larry G. Epstein**, “Stochastic Differential Utility,” *Econometrica*, 1992, 60 (2), 353–94.
- Eberly, Janice and Neng Wang**, “Reallocating and Pricing Illiquid Capital: Two Productive Trees,” May 2011. Working paper.
- Gomes, Joao F., Leonid Kogan, and Motohiro Yogo**, “Durability of Output and Expected Stock Returns,” *Journal of Political Economy*, October 2009, 117 (5), 941–986.
- Hall, Robert E.**, “Struggling to Understand the Stock Market,” *American Economic Review*, May 2001, 91 (2), 1–11.
- Jarrow, Robert A., David Lando, and Stuart M. Turnbull**, “A Markov Model for the Term Structure of Credit Risk Spreads,” *Review of Financial Studies*, 1997, 10 (2), 481–523.
- Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie**, “Beyond Balanced Growth,” *Review of Economic Studies*, 2001, 68, 869–882.
- Kung, Howard and Lukas Schmid**, “Innovation, Growth, and Asset Prices,” *Journal of Finance*, 2015, 70 (3), 1001–1037.
- Long, John and Charles Plosser**, “Real Business Cycles,” *Journal of Political Economy*, 1983, 91, 39–69.
- Martin, Ian**, “The Lucas Orchard,” *Econometrica*, January 2013, 81 (1), 55–111.
- Menzly, Lior, Tano Santos, and Pietro Veronesi**, “Understanding Predictability,” *Journal of Political Economy*, February 2004, 112 (1), 1–47.
- Merton, Robert C.**, “Optimum Consumption and Portfolio Rules in a Continuous-Time Model,” *Journal of Economic Theory*, 1971, 3, 373–413.

- Newey, Whitney K. and Kenneth D. West**, “A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, May 1987, 55 (3), 703–708.
- Ogaki, Masao and Carmen M. Reinhart**, “Measuring Intertemporal Substitution: The Role of Durable Goods,” *Journal of Political Economy*, 1998, 106 (5), 1078–98.
- Pavlova, Anna and Roberto Rigobon**, “Asset Prices and Exchange Rates,” *Review of Financial Studies*, 2007, 20 (4), 1139–81.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel**, “Housing, Consumption, and Asset Pricing,” *Journal of Financial Economics*, 2007, 83, 531–69.
- Ravn, Morten, Stephanie Schmitt-Grohe, and Martin Uribe**, “Deep Habits,” *Review of Economic Studies*, 2006, 73, 195–218.
- Samuelson, Paul A.**, “Lifetime Portfolio Selection by Dynamic Stochastic Programming,” *Review of Economics and Statistics*, 1969, 51, 239–46.
- Tauchen, George and Robert Hussey**, “Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models,” *Econometrica*, 1991, 59 (2), 371–96.
- Vissing-Jorgensen, Annette**, “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” *Journal of Political Economy*, 2002, 110 (4), 825–853.

A. DATA CONSTRUCTION

Our main household data are from the Federal Reserve’s Flow of Funds and cover the period 1951Q4–2015Q4 for household wealth data, although we use the CPI from the Bureau of Labor Statistics to deflate the cash flows constructed below. We supplement this data with a long annual time series over 1929–2015 from the Bureau of Economic Analysis when calculating statistics of macroeconomic aggregates: output and consumption are deflated by their appropriate price indices and fixed assets, at current cost, are made real by their corresponding investment price deflators. When computing Tobin’s Q we linearly interpolate the annual fixed asset values to a quarterly series. In addition, the equity portfolio data—returns, dividend yields, and market values—are downloaded from Ken French’s webpage.

Households and nonprofit organizations have about 70 percent of total assets in business capital and 30 percent in nonbusiness capital, which includes housing. Because we are interested in modeling households’ portfolio decisions we focus on a subset of total assets that would be more likely held as a result of an active allocation decision, one where the investor is free to make marginal changes to their investment portfolio to optimize over an Euler equation. In particular, we exclude pension entitlements (about 20 percent of total assets), equity in noncorporate business (10 percent), consumer durable goods (5 percent), and other small asset classes, leaving our coverage at about 63 percent of households’ total asset universe as of 2015Q4.

Business capital is the sum of debt and equity capital across the nonfinancial and financial sectors. Debt includes currency and deposits including money market fund shares, debt securities, loans, and shares of bond mutual funds. Equity comprises directly- and indirectly-held securities, such as those held through life insurance companies, pension plans, retirement funds. To avoid double-counting mortgages we calculate total business wealth as

$$\text{Business Capital} = q_{NF}K_{NF} + q_FK_F = D + B_{NF} + B_F + E_F + E_{NF} - M.$$

We view this sector’s capital stock as the fixed, private, nonresidential capital stock (BEA Table 1.1, Line 4) that produces goods and services unrelated to housing. Our measure of Tobin’s Q is therefore household’s business wealth over the

total nonresidential capital stock.

Cash flows from business capital are the sum of interest earned (BEA Table 2.1, Line 14) and dividends received (BEA Table 2.1, Line 15) by households. For each quarter, this cash flow then subtracts the total change in equity for nonfinancial and financial corporations multiplied by household's share of total equity to get a dollar amount for total issuance or repurchases absorbed or earned by households.

Household's ownership of housing is for all rented and owner-occupied real estate, including vacant land and mobile homes at market value (Flow of Funds LM115035035.Q and LM155035015.Q), and in our model this simply represents $q_H K_H$. The stock of housing is the private, residential fixed assets dwelled in by households (BEA Table 5.1, Line 7).

The flow of rental income (BEA Table 2.1, Line 12) consists of rental income of tenants and the imputed income of owners' housing services, which is net of "space rent" less expenses such as depreciation, maintenance, property taxes, and mortgage interest and is consistent with our treatment of mortgages being held by the financial sector.

B. DETAILS OF SOLUTION

It is convenient to write $\eta = (\epsilon - 1)/\epsilon$, which ranges between zero when the goods are Cobb-Douglas substitutes and one when they are perfect substitutes. *Weighted capital* is defined as

$$K \equiv (K_1^\eta + K_2^\eta)^{1/\eta} \quad (\text{B1})$$

and the first sector's capital share is

$$k \equiv \frac{K_1^\eta}{K_1^\eta + K_2^\eta}. \quad (\text{B2})$$

Applying Ito's lemma to (B2), we obtain the dynamics for this process

$$\begin{aligned}
dk = & k(1-k)\eta \left[\phi_1(i_1) + \frac{1}{2} ((\eta-1) - 2\eta k) \sigma_1^2 \right. \\
& - \phi_2(i_2) - \frac{1}{2} ((\eta-1) - 2\eta(1-k)) \sigma_2^2 + \eta(2k-1)\varphi\sigma_1\sigma_2 \left. \right] dt \\
& + k(1-k)\eta [\sigma_1 dB_1 - \sigma_2 dB_2].
\end{aligned} \tag{B3}$$

If the installation technologies were symmetric and $\sigma_1 = \sigma_2$, the drift would be shaped like a sine-curve that equals zero when k equals 0, $1/2$, or 1. When $0 < k < 1/2$ there would be a positive drift; when $1/2 < k < 1$, a negative drift; it therefore would mean-revert to $k = 1/2$. It contrast to endowment economies, it features endogenous growth from investment governed by each sector's installation technology.

With these definitions we can redefine aggregate consumption in (4) as a product of weighted capital and the consumption-to-weighted capital ratio

$$C = K \underbrace{(k\theta^{1-\eta}(A_1 - i_1)^\eta + (1-k)(1-\theta)^{1-\eta}(A_2 - i_2)^\eta)^{1/\eta}}_{\equiv c(k,\theta)}. \tag{B4}$$

In contrast to the one-sector *iid* case, consumption-to-capital now varies over time with the each sector's capital share and dividend, and the demand for its good. When a sector has a high capital share, the agent prefers to have a high demand for its good, and will choose the sector's production schedule to maximize consumption of that good.

We conjecture that $J(\cdot)$ has the homogeneity property in weighted capital

$$J(K_1, K_2, \theta) = \frac{1}{1-\gamma} [Kv(k, \theta)]^{1-\gamma}, \tag{B5}$$

where $v(k, \theta)$ is a function to be determined. Using this property to rewrite (14)

gives¹⁶

$$0 = \max_{i_1, i_2} \left\{ \frac{\rho}{1 - 1/\psi} \left(\left(\frac{c(k, \theta)}{v(k, \theta)} \right)^{1-1/\psi} - 1 \right) + \sum_{\theta' \neq \theta} \frac{\pi_{\theta\theta'}}{1 - \gamma} \left(\left(\frac{v(k, \theta')}{v(k, \theta)} \right)^{1-\gamma} - 1 \right) \right. \\ \left. + \sum_{n=1}^2 \left(\phi_n(i_n) \mathcal{J}_n + \frac{1}{2} \sigma_n^2 \mathcal{J}_{nn} \right) + \varphi \sigma_1 \sigma_2 \mathcal{J}_{12} \right\}, \text{ for } \theta \in \{\Theta_m\}_{m=1}^M. \quad (\text{B6})$$

The first term in the equation relates to the agent's flow utility related to the ratio of consumption to continuation utility. The second term relates to changes in continuation utility that occurs from demand shocks. The last summation describes changes in supply that affect the agent's choice of investment and the agent's value from diversification.

The first-order conditions for i_1 and i_2 jointly solve

$$\rho \left(\frac{c(k, \theta)}{v(k, \theta)} \right)^{-1/\psi} p_1 = \frac{\phi'_1(i_1) \mathcal{J}_1(k, \theta) v(k, \theta)}{k^{1/\eta}}, \\ \rho \left(\frac{c(k, \theta)}{v(k, \theta)} \right)^{-1/\psi} p_2 = \frac{\phi'_2(i_2) \mathcal{J}_2(k, \theta) v(k, \theta)}{(1 - k)^{1/\eta}}. \quad (\text{B7})$$

Finally, the one sector economy defines the boundaries of the two-sector economy as one sector becomes negligibly small. In these limits, the sector's capital stock K_n is the single state variable. The stochastic growth rates of capital, con-

¹⁶The components of the HJB equation are

$$\mathcal{J}_1(k, \theta) = \left(k + \eta \frac{v'(k, \theta)}{v(k, \theta)} k(1 - k) \right), \\ \mathcal{J}_2(k, \theta) = \left((1 - k) - \eta \frac{v'(k, \theta)}{v(k, \theta)} k(1 - k) \right), \\ \mathcal{J}(k, \theta) = \frac{v''(k, \theta)}{v(k, \theta)} \eta^2 k^2 (1 - k)^2, \\ \mathcal{J}_{11}(k, \theta) = (1 - \eta) \left(k^2 - \mathcal{J}_1(k, \theta) \right) - \gamma \mathcal{J}_1(k, \theta)^2 + \mathcal{J}(k, \theta), \\ \mathcal{J}_{22}(k, \theta) = (1 - \eta) \left((1 - k)^2 - \mathcal{J}_2(k, \theta) \right) - \gamma \mathcal{J}_2(k, \theta)^2 + \mathcal{J}(k, \theta), \\ \mathcal{J}_{12}(k, \theta) = (1 - \eta) \left(k(1 - k) - \eta \frac{v'(k, \theta)}{v(k, \theta)} k(1 - k)(2k - 1) \right) - \gamma \mathcal{J}_1(k, \theta) \mathcal{J}_2(k, \theta) - \mathcal{J}(k, \theta).$$

sumption, investment, and output are equal. Because growth rates and returns are *iid*, the ratios of consumption-, investment-, and output-to-capital are constant, as is Tobin's Q . The optimal investment rate solves

$$A_n - i_n^* = \frac{1}{\phi_n'(i_n^*)} \left[\rho + (1/\psi - 1) \left(\phi_n(i_n^*) - \gamma \frac{\sigma_n^2}{2} \right) \right], \quad (\text{B8})$$

and the coefficient that solves the value function is given by

$$b_n = \frac{\rho}{\phi_n'(i_n^*)} \left[1 + \frac{1/\psi - 1}{\rho} \left(\phi_n(i_n^*) - \gamma \frac{\sigma_n^2}{2} \right) \right]^{\frac{1/\psi}{1/\psi - 1}}. \quad (\text{B9})$$

Using these two coefficients we define the boundaries of (B6) $v(0, \theta) = b_1$ and $v(1, \theta) = b_2$, for all $\theta \in \{\Theta_m\}$.

A. DERIVATION OF SDF AND RISK-FREE RATE

From (15), Ito's lemma shows

$$\begin{aligned} \frac{d\Lambda}{\Lambda} = & f_J(C, J)dt - \gamma \frac{dK}{K} + \frac{1}{2} \gamma (1 + \gamma) \left(\frac{dK}{K} \right)^2 - \gamma \frac{dK}{K} \frac{\partial \zeta(k, \theta)}{\partial k} dk \\ & + \frac{\partial \zeta(k, \theta)}{\partial k} dk + \frac{1}{2} \left(\left(\frac{\partial \zeta(k, \theta)}{\partial k} \right)^2 + \frac{\partial^2 \zeta(k, \theta)}{\partial k^2} \right) (dk)^2 \\ & + \sum_{\theta' \neq \theta} \left(\frac{e^{\zeta(k, \theta')}}{e^{\zeta(k, \theta)}} - 1 \right) dN^{(\theta, \theta')}, \end{aligned} \quad (\text{B10})$$


where the dynamics for weighted capital in (B1) are

$$\frac{dK}{K} = k \frac{dK_1}{K_1} + (1 - k) \frac{dK_2}{K_2} + \frac{1}{2} (\eta - 1) k (1 - k) (\sigma_1^2 + \sigma_2^2 - \varphi \sigma_1 \sigma_2) dt. \quad (\text{B11})$$

The risk-free rate in units of consumption good $-\mathbb{E}_t[d\Lambda/\Lambda]$ is therefore


$$\begin{aligned}
r(k, \theta) = & \rho + \rho \left(\frac{1/\psi - \gamma}{1 - 1/\psi} \right) \left(1 - \left(\frac{c(k, \theta)}{v(k, \theta)} \right)^{1-1/\psi} \right) + \gamma \frac{1}{dt} \mathbb{E}_t \left[\frac{dK}{K} \right] \\
& - \frac{1}{2} \gamma (1 + \gamma) (k^2 \sigma_1^2 + (1 - k)^2 \sigma_2^2 + 2k(1 - k) \varphi \sigma_1 \sigma_2) \\
& + \gamma \frac{\partial \zeta(k, \theta)}{\partial k} \eta k(1 - k) (\sigma_1^2 k - \sigma_2^2 (1 - k) + (1 - 2k) \varphi \sigma_1 \sigma_2) - \frac{\partial \zeta(k, \theta)}{\partial k} \frac{1}{dt} \mathbb{E}_t[dk] \\
& - \frac{1}{2} \left(\left(\frac{\partial \zeta(k, \theta)}{\partial k} \right)^2 + \frac{\partial^2 \zeta(k, \theta)}{\partial k^2} \right) \eta^2 k^2 (1 - k)^2 (\sigma_1^2 + \sigma_2^2 - 2\varphi \sigma_1 \sigma_2) \\
& - \sum_{\theta' \neq \theta} \pi_{\theta\theta'} \left(\frac{e^{\zeta(k, \theta')}}{e^{\zeta(k, \theta)}} - 1 \right). \tag{B12}
\end{aligned}$$

Figure 1: The Puzzle Within and Across Asset Classes



The top-left, top-right, and bottom-left panels plot share-implied risk premia ($\mathbb{E}[R^e]$) on valuation-implied risk premia for the five book-to-market portfolios, five industry, and our construction of US household portfolio wealth, respectively. We calculate share-implied risk premia as the vector $\mathbb{E}_t[R^e] = \gamma \Sigma \mathbf{w}_t$, for all t , where γ is a value of risk aversion, Σ is the estimated constant covariance matrix of portfolio excess returns and $\mathbf{w}_t$ is the vector of current wealth shares from the data. Valuation-implied risk premia are the fitted values of a regression of future annual cumulative excess returns on the current dividend yield and a constant. The bottom-right corner plots the change in Tobin's Q on the change in the industry's wealth share, with the five color-coded regression lines overlayed. Debt, equity, and housing data are quarterly, constructed from the BEA, the Flow of Funds, and the BLS, and cover from 1951Q4 until 2015Q4, as does data on Tobin's Q , which interpolates Compustat's annual series of PPENT for its denominator. Equity portfolio returns and wealth shares data are quarterly and are from Ken French's website and cover from 1928Q1 until 2015Q4.

Figure 2: The Puzzle at Long Horizons



This figure plots the 95 percent confidence interval of various slope coefficients from predictive regressions of the form $\frac{1}{H} \sum_{h=1}^H R_{t+h} - R_{t+h}^f = a_H + b_H \times w_t + \epsilon_{t+H}$ for each asset. The slope coefficients, b_H , range over quarterly horizons $H = 4, 8, \dots, 40$. The horizontal axis, Horizon, is measured in years. The estimate of the slope is a solid line; the confidence interval is defined by the dashed lines; and the line across zero is dotted. Newey and West (1987) standard errors are used with a lag length equal to the forecast horizon in quarters. Debt, equity, and housing data are quarterly, constructed from the BEA, Flow of Funds, and BLS, and cover from 1951Q4 until 2015Q4. Equity portfolio returns, risk-free rate returns, and wealth shares data are quarterly and are from Ken French's website and cover from 1928Q1 until 2015Q4.

Figure 3: Overview of Economy






Figure 4: Balancedness and Growth




This figure plots model policy functions as a function of k and three levels of $\theta = \{\Theta_1, \Theta_5, \Theta_9\}$. The lightest green line corresponds with the lowest level of demand for the business sector's good; increasing darkness indicates greater demand. The conditional drift of the first (business) sector's capital share is $E[dk]$. The variable $w - k$ plots the wealth share of the business sector in excess of its capital share.

Figure 5: Tobin's Q and Risk Premia




This figure plots policy functions from the model as a function of k and three levels of $\theta = \{\Theta_1, \Theta_5, \Theta_9\}$. The lightest green line corresponds with the lowest level of demand for the business sector's good; increasing darkness indicates greater demand. The left-side plots are for the business sector; the right-side, the housing sector. Both Tobin's Q and the risk premia are measured in units of the numeraire.

Figure 6: Decomposition of Business Sector's Risk Premium




This figure plots the decomposition of the risk premium for the first (business) sector: $(\frac{1}{dt}\mathbb{E}_t[dR_1] - r) = \lambda_1 \sigma_1 \beta_{11} + \lambda_2 \sigma_2 \beta_{12} + \sum_{\theta' \neq \theta} \pi_{\theta\theta'} \lambda_{Demand} \beta_{Demand}$. Functions from the model are displayed as a function of k and three levels of $\theta = \{\Theta_1, \Theta_5, \Theta_9\}$. The lightest green line corresponds with the lowest level of demand for the business sector's good; increasing darkness indicates greater demand. Approximations to the demand shock's beta and risk price are obtained by taking expectations of the functions conditional on every (k, θ) pair: $\mathbb{E}[\lambda_{Demand}|k, \theta] = \sum_{\theta' \neq \theta} \pi_{\theta\theta'} \left(\frac{e^{\zeta(k, \theta')}}{e^{\zeta(k, \theta)}} - 1 \right)$ and $\mathbb{E}[\beta_{1Demand}|k, \theta] = \sum_{\theta' \neq \theta} \pi_{\theta\theta'} \left(1 - \frac{Q_1(k, \theta')}{Q_1(k, \theta)} \right)$.

Figure 7: Business Sector's Risk Premium Contribution by Source, in Percent




This figure plots a decomposition of risk premia. Decompositions are displayed as a function of k and three levels of $\theta = \{\Theta_1, \Theta_5, \Theta_9\}$, for which each has a subplot. Supply 1 equals $\sigma_1 \lambda_1 \beta_{11}$; Supply 2 equals $\sigma_2 \lambda_2 \beta_{12}$ and Demand equals $\sum_{\theta' \neq \theta} \pi_{\theta\theta'} \left(\frac{e^{\zeta(k, \theta')}}{e^{\zeta(k, \theta)}} - 1 \right) \left(1 - \frac{Q_1(k, \theta')}{Q_1(k, \theta)} \right)$. The sum of all three sources totals the sector's risk premium for each (k, θ) pair.

Figure 8: Decomposition of Asset Demand



This figure plots each asset's mean-variance demand and hedging demand as a function of the business's wealth and three levels of $\theta = \{\Theta_1, \Theta_5, \Theta_9\}$. The lightest green line corresponds with the lowest level of demand for the business sector's good; increasing darkness indicates greater demand. The left-side plots are for the business sector; the right-side, the housing sector. The dashed black lines are a 45-degree line and a line at zero. By construction, the sum of mean-variance demand and hedging demand equals the sector's wealth.

Figure 9: Hedging Demand from Both Shocks and Only Supply Shocks




This figure plots histograms of hedging demand for 5,000 simulations of the full model (Both Shocks) and 1,000 simulations for the model having only supply shocks (Supply Shocks). The Supply Shocks simulation solves the model fixing $\Theta = 0.5$. The horizontal axis measures hedging demand $HD_n = w_n - MV_n$ in percentage points.

Figure 10: Recovering Θ from the Data



This figure plots the inverted state variables (k, θ) from the data on the business's sector capital share k and wealth share w . The capital share of the data is matched exactly to the model's first sector's capital share k , quarter-by-quarter, and is plotted in the top subpanel. Given the match of k and our calibration, the distance between the data's w and the model's counterpart is minimized for each data point. The corresponding best fit of w and the implied θ are plotted in the bottom subpanel on the left and right axes, respectively. Business and housing sector data are quarterly, constructed from the BEA, Flow of Funds, and BLS, and cover from 1951Q4 until 2015Q4.

Figure 11: Model Recovery of Tobin's Q , Risk Premia, and Hedging Demand



This figure plots model-implied time series from the history of inverted state variables. Given data on the business's sector capital share k and wealth share w , the model counterparts are exactly fitted to k and best fitted to w , jointly implying a quarter-by-quarter (k, θ) pair. This time series pair is then used by the model's policy functions to plot Tobin's Q , risk premia, and hedging demand in the top, middle, and bottom subpanel respectively. Both sectors' hedging demands are plotted in the bottom subpanel. All model time series are smoothed with a 12-quarter moving average. Debt, Equity, and Housing data are quarterly, constructed from the BEA, Flow of Funds, and BLS, and cover from 1951Q4 until 2015Q4, as does data on Tobin's Q , which interpolates Compustat's annual series of PPENT for its denominator.

Table I: Calibration (Annual)		
Parameter	Value	
<i>Preferences</i>		
ρ	0.02	
γ	10	
ψ	1.1	
ϵ	3	
<i>Technology</i>		
	Business	Housing
A	0.15	0.10
κ	12	14
σ	0.032	0.034
δ	0.075	0.020
φ	0	

Table II: Ergodic Distribution of Π

Θ_m	Value	Density
Θ_1	0.049	0.040
Θ_2	0.179	0.081
Θ_3	0.292	0.127
Θ_4	0.398	0.163
Θ_5	0.500	0.177
Θ_6	0.602	0.163
Θ_7	0.708	0.127
Θ_8	0.821	0.081
Θ_9	0.951	0.040

Table III: Summary Statistics

		<i>Data</i>		<i>Both Shocks</i>		<i>Supply Shocks</i>	
		Mean	Stdev	Mean	Stdev	Mean	Stdev
<i>Aggregate</i>	Output	3.1	4.9	3.5	6.9	3.4	2.8
	Consumption (expenditure)	3.1	2.9				
	Consumption (cash flow)	3.7	6.8	3.5	5.0	3.4	2.6
<i>Business</i>	Capital stock	3.4	3.2	3.5	3.5	3.6	3.2
	Cash flow	4.1	8.0	3.6	19.5	3.6	3.3
	Excess return	14.0	12.0	6.7	8.3	2.0	3.3
<i>Housing</i>	Capital stock	2.4	3.4	3.5	3.6	2.6	3.0
	Cash flow	3.0	12.9	3.6	15.5	2.6	3.1
	Excess return	4.5	5.7	6.8	7.1	-0.6	3.4

This table summarizes means and standard deviations of variables from the data and the model subject to both shocks and to only supply shocks. Cash flow and return data are quarterly from 1951Q4 until 2015Q4. Aggregate quantities and capital stocks are annual from 1929-2015. Sources: BEA, Federal Reserve Flow of Funds, and BLS for the CPI price deflator. Values are real where applicable. Consumption (expenditure) is the real growth rate of personal consumption expenditure on nondurables and services. Consumption (cash flow) is a cash flow-weighted growth rate of the business and housing cash flow. Simulation data are calculated with and without a demand shock. When there are no demand shocks, we solve the model fixing $\Theta = 0.5$. We simulate each model 5,000 times for 480 quarters, burning in the first half and leaving 60 years of quarterly data. We calculate both the mean and standard deviation of each variable for each simulation. We then average across both of these statistics to report the columns titled Mean and Stdev.

Table IV: Finite Sample Distribution of Regression Slope Coefficients

Panel A: Risk Premium on Wealth Share							
	$\hat{b}_n$	<i>Business</i>			<i>Housing</i>		
		1%	Mean	99%	1%	Mean	99%
	Data	-0.622			-0.303		
<i>Levels</i>	Both Shocks	-0.677	-0.015	0.512	-0.449	0.025	0.614
	Supply Shocks	0.003	0.004	0.004	0.003	0.004	0.006
	Data	-0.463			-0.016		
<i>Changes</i>	Both Shocks	-1.891	-0.416	0.209	-0.242	0.365	1.727
	Supply Shocks	0.003	0.004	0.004	0.003	0.004	0.006
Panel B: Tobin's Q on Wealth Share							
	$\hat{b}_n$	<i>Business</i>			<i>Housing</i>		
		1%	Mean	99%	1%	Mean	99%
	Data	0.031			0.012		
<i>Levels</i>	Both Shocks	-0.068	-0.006	0.086	-0.249	-0.077	0.010
	Supply Shocks	-0.020	-0.012	-0.008	-0.024	-0.016	-0.010
	Data	0.040			0.001		
<i>Changes</i>	Both Shocks	-0.077	0.054	0.166	-0.465	-0.124	0.058
	Supply Shocks	-0.020	-0.012	-0.008	-0.024	-0.016	-0.011

This table reports point estimates of regressions in the data and finite sample distributions of regression coefficients from simulated data. The regressions specifications are identical across data and model. In each simulation we run the Levels regressions of $\frac{1}{dt}\mathbb{E}_t[dR_{nt}] - r_t = a_n + b_n \times w_{nt} + \epsilon_{nt}$ in Panel A and $Q_{nt} = a_n + b_n \times w_{nt} + \epsilon_{nt}$ in Panel B for each sector n . We also run the regressions after differencing the regressors and regressand in the subpanels marked Changes. We report the mean and the 1st and 99th percentile of slope coefficients $\hat{b}_n$ across simulations. Wealth shares and risk premia are measured in percent. Debt and Equity (Business), and Housing data are quarterly and cover from 1951Q4 until 2015Q4, as does data on Tobin's Q , which interpolates Compustat's annual series of PPENT as the denominator.