

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: NBER Macroeconomics Annual 2013, Volume 28

Volume Author/Editor: Jonathan A. Parker and Michael Woodford, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 978-0-226-16540-0 (cloth); 978-0-226-16540-0 (paper); 978-0-226-16554-7 (eISBN)

ISSN: 0889-3365

Volume URL: <http://www.nber.org/books/park13-1>

Conference Date: April 12–13, 2013

Publication Date: April 2014

Chapter Title: Comment on "Shocks and Crashes"

Chapter Author(s): John Y. Campbell

Chapter URL: <http://www.nber.org/chapters/c12933>

Chapter pages in book: (p. 355 – 366)

Comment

John Y. Campbell, Harvard University and NBER

Macroeconomics is about the dynamic behavior of aggregate variables that summarize the state of the economy. If macroeconomic models are to be tractable, they can include only a few of the thousands of variables produced by government agencies and private economy-watchers. Real business cycle (RBC) models, for example, typically include output, the capital stock, and hours worked, while New Keynesian models typically include a short-term interest rate, inflation, and a measure of the cyclical component of output. Additional variables may be added to these core sets, but this must be done with discretion to maintain the usability of the model.

Martin Lettau and Sydney Ludvigson choose a different small set of macroeconomic variables: consumption, labor income, and a broad measure of financial wealth taken from the Flow of Funds accounts. This choice of variables is motivated by an aggregate version of the household budget constraint, just as the real business cycle choice is motivated by the aggregate production function, and the New Keynesian choice is motivated by the Phillips curve and the monetary policy reaction function. Lettau and Ludvigson's variables connect their work with macro-finance research, just as real business cycle models connect with growth theory, and New Keynesian models connect with monetary economics.

The household budget constraint restricts the long-run behavior of Lettau and Ludvigson's variables. Because consumption is financed out of human and financial wealth, it should have a stable long-run relationship with these components of total wealth. If this were not the case, then consumption could drift away from the wealth that supports it, and either the intertemporal budget constraint or a transversality

condition would be violated. While human wealth is not observed, labor income can be regarded as the dividend on human wealth, which should also have a stable long-run relationship with the value of human wealth. Hence, consumption, labor income, and financial wealth should have a stable long-run relationship: that is, these time series should be cointegrated. Earlier papers by Lettau and Ludvigson (2001, 2004) have presented econometric evidence that this is the case, and have worked with the stationary linear combination of these variables that they call *cay*.

Standard models of the aggregate production function imply that human and financial wealth also have a stable long-run relationship and should be cointegrated. If this were the case, then Lettau and Ludvigson's three variables would have only one unit root—that a real business cycle model would interpret as a persistent technology shock—and all pairs of the variables would be cointegrated. However, historical US data do not suggest that this is the case. Instead there appears to be only a single cointegrating relationship, and the three variables are driven by two persistent shocks and one transitory shock.

In this paper, Lettau and Ludvigson take two further steps in their research agenda. First, they offer an economic interpretation of the shocks, labeling them as persistent shocks to technology and factor shares, and a transitory shock to the risk aversion of households. Second, they examine the contributions of these shocks to the two most recent US business cycles. They argue that the rise and fall of financial wealth during the late 1990s and the early 2000s was driven primarily by fluctuating risk aversion, and that the same was true of the boom in the mid-2000s. However, they argue that the downturn since 2007 was driven by a combination of negative technology shocks and increasing risk aversion.

Lettau and Ludvigson's framework is very rich and offers wide scope to a discussant. I will concentrate this discussion on three topics: the "wealth volatility puzzle" that financial wealth has much greater short-run volatility than macroeconomic fundamentals; the interpretation of movements in the labor income share; and the interpretation of recent boom-bust cycles in US financial markets.

The Wealth Volatility Puzzle

In the short run, the return on the aggregate stock market is much more volatile than the growth rate of consumption. While the aggregate stock

market is not necessarily an adequate proxy for total wealth, plausible corrections for bonds, real estate, human capital and other components of wealth narrow the volatility gap but by no means eliminate it. Thus, in the short run, the return on total wealth seems to be considerably more volatile than the growth rate of consumption. The same is true if we measure the growth rate of wealth rather than the return on wealth (the difference being the saving rate, which contributes relatively little to short-run volatility).

The short-run volatility gap is a puzzle because, over a long enough horizon, the growth rates of wealth and consumption should have equal volatilities. Any model with stationary investment opportunities and scale-independent preferences (needed to prevent trends in interest rates and risk premia as the economy grows) implies that the wealth-consumption ratio is stationary. Since aggregate time series appear to follow unit-root rather than trend-stationary processes, this implies that there is a common unit root in wealth and consumption; in other words, that these series are cointegrated. As the horizon increases over which growth rates are measured, the variances of wealth and consumption growth rates converge to equality because they are both dominated by the variance of the common unit root component, and the variance of the consumption-wealth ratio is negligible in comparison.

This logic raises the question of how the short-run volatility gap is eliminated in the long run. There are two possibilities. First, it could be that the volatility of annualized consumption growth increases with the horizon, in other words that consumption growth is persistent. Equilibrium macro-finance models with this feature include Constantinides (1990) and Bansal and Yaron (2004). Constantinides' model has an exogenous i.i.d. return to wealth and a representative consumer with habit-formation utility. The consumer chooses to smooth consumption in the short run, adjusting gradually to wealth shocks. In the longer run, as consumption adjustment proceeds, the volatility of consumption growth approaches the exogenous volatility of wealth growth. Bansal and Yaron's "long-run risk" model has exogenous positively autocorrelated consumption growth, and a representative investor with Epstein-Zin (1989) preferences and an elasticity of intertemporal substitution (EIS) greater than one. When consumption growth accelerates, the expected future dividends on total wealth increase, and with EIS greater than one there is relatively little effect on discount rates, so the value of wealth immediately reflects the long-run impact of consumption growth on the level of consumption.

Alternatively, it could be that the volatility of annualized wealth decreases with the horizon; in other words, that wealth is mean-reverting. Equilibrium macro-finance models with this feature include Campbell (1993, 1996) and Campbell and Cochrane (1999). Campbell's model has an exogenous mean-reverting return to wealth and a representative investor with Epstein-Zin preferences and an EIS close to zero. In this model consumption is almost proportional to the permanent component of wealth, so a mean-reverting shock to wealth induces a consumption response that is much less than one-for-one. Campbell and Cochrane's model has exogenous i.i.d. consumption growth and a representative investor with habit-formation utility. Rapid consumption growth temporarily increases the gap between consumption and habit, lowering risk aversion and increasing wealth even more than consumption, but this excess increase in wealth disappears as habit adjusts to the new level of consumption.

Lettau and Ludvigson's earlier papers shed light on the wealth volatility puzzle by showing that movements in *cay*, the stationary linear combination of consumption, labor income, and financial wealth, predict the growth rate of financial wealth but not the growth rates of consumption or labor income. The current paper updates this evidence, which is important because it implies that wealth, not consumption, adjusts to equate the long-run volatilities of these two series.

Other empirical research has also cast doubt on the existence of persistent fluctuations in consumption growth. Beeler and Campbell (2012), for example, point out that the longer-run autocorrelations of consumption growth tend to be negative. There is positive first-order autocorrelation, but it has little effect on longer-run consumption dynamics and may be explained by the fact that we observe time-aggregated not point-in-time consumption. Also, stock prices do not predict consumption growth as one would expect if the model of Bansal and Yaron (2004) were correct. Similarly, empirical research dating back to the late 1980s (e.g., Campbell and Shiller 1988; Fama and French 1988; Poterba and Summers 1988) has found evidence of mean-reversion in wealth.

Naturally, there are things that can be said on the other side of this debate. Stambaugh (1999) warned of the bias that can arise when predicting wealth returns using persistent variables, and recent research by Lewellen (2004), Campbell and Yogo (2006), and Cochrane (2008) has addressed this issue. In the literature on consumption, Parker and Julliard (2005) show that the excess returns on small over large stocks and value stocks over growth stocks do seem to have some predictive

power for future consumption growth. Bansal, Kiku, and Yaron (2012) argue that stochastic volatility in consumption growth is important and that it causes asset market fluctuations that disguise the ability of investors to forecast future consumption growth. It is likely that there is some predictability in consumption growth, but in my view—influenced to a large degree by Lettau and Ludvigson's results—any such predictability is modest and the more important phenomenon is mean-reversion in wealth.

If this is correct, it makes it harder to explain the equity premium puzzle. Models with persistent consumption growth imply that consumption is actually riskier than it seems in the short run, but the true level of risk is obscured by gradual adjustment. Models with mean-reverting wealth growth, on the other hand, imply that wealth is safer than it looks in the short run, because most of its fluctuations are temporary and have minimal impact on long-term investors. Unsurprisingly, the persistent-consumption-growth models of Constantinides (1990) and Bansal and Yaron (2004) fit the equity premium with much lower risk aversion coefficients than the mean-reverting-wealth models of Campbell (1993, 1996) or Campbell and Cochrane (1999). If consumption growth is not persistent, other model ingredients such as rare disasters, ambiguity aversion, or uninsurable idiosyncratic risk are needed to reconcile the equity premium with a moderate coefficient of risk aversion.

Lettau and Ludvigson correctly point out that their results also challenge models such as Campbell and Cochrane (1999) in which temporary fluctuations in wealth are driven entirely by contemporaneous movements in consumption. Instead, it seems that wealth fluctuates autonomously, as in the "moody investors" model of Bekaert, Engstrom, and Grenadier (2010) or the augmented Epstein-Zin model of Dew-Becker (2013).

The Labor Income Share

Traditional macro models have a stationary and only modestly variable labor income share. Older US data showed modest cyclical fluctuations in the labor share around an apparently constant mean, and this encouraged the use of the convenient Cobb-Douglas production function—with its implication of constant factor shares—in macro modeling. As Lettau and Ludvigson point out, during the last twenty years the labor income share has ceased to be stable but instead exhibits

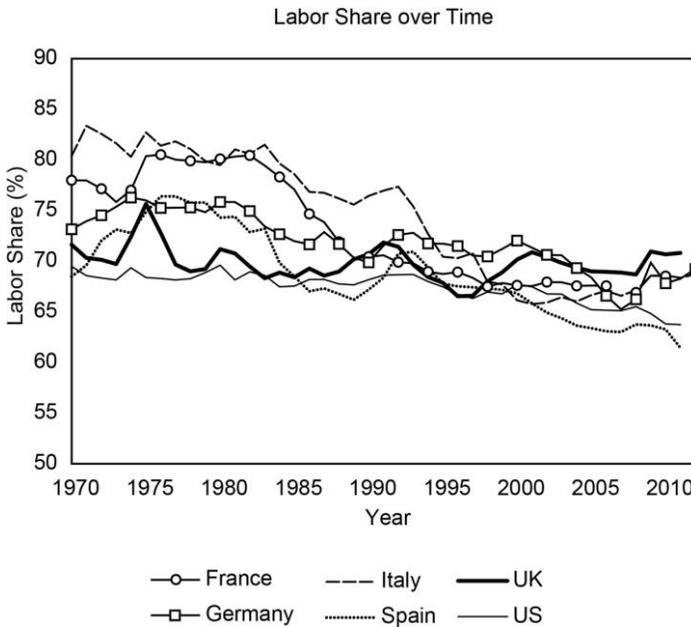


Fig. C1. Labor shares in six developed economies

a pronounced downtrend. Other recent papers, for example, Jacobson and Occhino (2012), also emphasize this fact.

Many other developed countries had volatile labor shares much earlier in the twentieth century, as emphasized by Bottazzi, Pesenti, and van Wincoop (1996). Figure C1 uses OECD (Organization for Economic Cooperation and Development) data to plot the history of the labor shares in France, Germany, Italy, Spain, the United Kingdom, and the United States since 1970. Two facts are immediately evident in this figure. First, in the late twentieth century labor shares were much more volatile in Europe than in the United States. Second, there is a general downtrend in all these countries except the United Kingdom, a downtrend that is larger and began earlier in continental Europe than in the United States.

There are of course many difficulties in measuring the labor share, notably the difficulty of separating the income received by sole proprietors and unincorporated businesses into components attributable to labor and capital inputs. However, Karabarounis and Nieman (2013) document a similar downtrend in the labor share within the corporate sectors of numerous countries.

Given the persistent decline in the US labor share, it is not surprising that Lettau and Ludvigson estimate an important role for a persistent shock that shifts labor income without moving aggregate consumption. Given that consumption is financed by labor income and capital income, such a shock must implicitly move capital income opposite to labor income, at least in the long run. Lettau and Ludvigson label this shock a factor shares shock, but do not build a structural model that could be used to understand it more precisely.

One possible interpretation of a factor shares shock, within a standard stylized real business cycle model, is a shock to the exponent on capital in a Cobb-Douglas production function, accompanied by a technology shock. A positive shock to the capital exponent raises the marginal product of capital, which stimulates investment and reduces consumption for a given level of output. To offset the effect on consumption (since Lettau and Ludvigson's factor share shock is orthogonal to current consumption), the exponent shock could be accompanied by a positive shock to technology that increases output. The combination of shocks initially has a positive effect on capital income and a slightly smaller negative effect on labor income. Over time, capital is accumulated and labor income gradually rises to a new steady state. Consumption must also increase over time, but if the elasticity of intertemporal substitution is low, this increase is extremely gradual and may be hard to detect. Such a pattern is broadly consistent with figure 4 in Lettau and Ludvigson's paper.

Karabourbanis and Neiman (2013) explain the long downtrend in the labor share using a richer model in a similar spirit, in which declining prices of investment goods, driven by the advance of information technology, have induced firms to substitute from labor to capital.

Closer inspection of figure C1 suggests that shorter-term political factors have also influenced the labor share. In the United Kingdom, for example, the spike in the labor share in 1974 to 1975 was associated with political turmoil and a brief but extreme decline in the stock market. In France, the labor share was particularly high during the late 1970s and the early years of the Mitterrand presidency, declining later in Mitterrand's term after he adopted more business-friendly policies.

These observations are important for the macro-finance literature. They suggest the relevance of models in which investors are concerned about shocks that have opposite effects on capital and labor. Bottazzi, Pesenti, and van Wincoop (1996) argue that such shocks can rationalize home bias in financial investments, since the wage earners in a country

can insure themselves against local factor share shocks by holding local equities. Shocks of this sort may also be relevant to the literature on rare disasters that follows Rietz (1988) and Barro (2006). Disasters are typically interpreted as wars or financial crises that greatly reduce the total output of an economy, but stockholders may regard a negative shock to capital income as a disaster even if the shock increases labor income. The extreme fluctuations of the UK stock market in 1974 to 1975 might fit this alternative interpretation of disaster.

This notion of factor-specific risk can be further enriched if one considers heterogeneity in capital or labor. This is the idea of models by Hobijn and Jovanovic (2001), Laitner and Stolyarov (2003), and Kogan and Papanikolaou (2013). In these models, technology shocks can render old physical or human capital obsolete even while stimulating the accumulation of new capital.

Recent Boom-Bust Cycles

Lettau and Ludvigson use their methodology to analyze two recent boom-bust cycles in US financial markets: the technology boom (1994:2–2000:1) and bust (2002:2–2002:3), and the credit boom (2002:4–2007:3) and bust (2007:4–2012:3). They argue that the technology boom and bust and the credit boom were primarily driven by transient shocks to financial wealth that had little effect on consumption or labor income. The credit bust, however, is a combination of this sort of transient shock and a permanent negative shock to consumption. In this sense the downturn since 2007:4 is more serious and has more long-lasting effects on the economy than the earlier fluctuations. Lettau and Ludvigson interpret the earlier transient shocks as resulting from changes in investor risk aversion, while the recent permanent shock to consumption is interpreted within the real business cycle paradigm as a negative technology shock.

Lettau and Ludvigson's work implies that financial cycles are not all alike, but may be driven by different shocks on different occasions. The question of the similarity or dissimilarity of business cycles is an old one in macroeconomics that has particularly concerned NBER economists (see, e.g., Burns and Mitchell 1946, or Blanchard and Watson 1986). It should be noted, however, that within Lettau and Ludvigson's framework, all consumption cycles are alike because consumption fluctuations are only driven by permanent shocks to consumption.

I have studied these same two cycles in a recent paper with Stefano

Giglio and Christopher Polk. We use a different methodology, breaking aggregate stock returns into components attributable to permanent cash-flow news and temporary discount-rate news. To accomplish this decomposition we use a vector autoregression (VAR) in stock returns and predictor variables and a log-linear approximate relation between stock returns and cash flows, following Campbell and Shiller (1988) and Campbell (1991). We also consider the impact of imposing an asset pricing model, the ICAPM (intertemporal capital asset pricing model) of Merton (1973), as empirically implemented by Campbell and Vuolteenaho (2004), on the VAR and the cross-section of stock returns. The use of the ICAPM can increase the precision of the VAR estimates and the decomposition of returns, to the extent that the cross-section of stock returns is informative and the ICAPM is a correct cross-sectional asset pricing model.

The findings of Campbell, Giglio, and Polk (2013, henceforth CGP) are broadly consistent with those of Lettau and Ludvigson, which is reassuring for both papers. The main difference is that CGP finds more evidence of a cash-flow boom in the mid-2000s, based on the growth rate of corporate earnings as well as the relative performance of growth and value stocks. To illustrate these findings, figure C2 plots CGP's estimates of cash-flow and discount-rate news, smoothed using an exponentially weighted backward moving average with a decay rate of 0.08 per quarter, over the period 1995 through 2011.

The left panels of the figure show estimates of cash-flow news, while the right panels show estimates of discount-rate news. These two news variables add up to the unexpected return on the stock market. Cash-flow news loosely corresponds to Lettau and Ludvigson's permanent consumption shock, while discount-rate news loosely corresponds to their transient shock to risk aversion. The three rows of the figure correspond to three alternative models: unrestricted VARs with two and five state variables, and a restricted VAR imposing the ICAPM. All three models display a strong hump in discount-rate news during the technology cycle of the late 1990s and early 2000s, with little evidence of a hump in cash-flow news. During the credit boom of the mid-2000s, on the other hand, all three models imply that there was unusually good news about cash flows (occurring somewhat earlier in the restricted model than in the unrestricted models). The credit downturn is attributed to a mix of bad cash-flow news and increasing discount rates, with the latter kicking in particularly during the last quarter of 2008.

These results are driven by differences in various financial and eco-

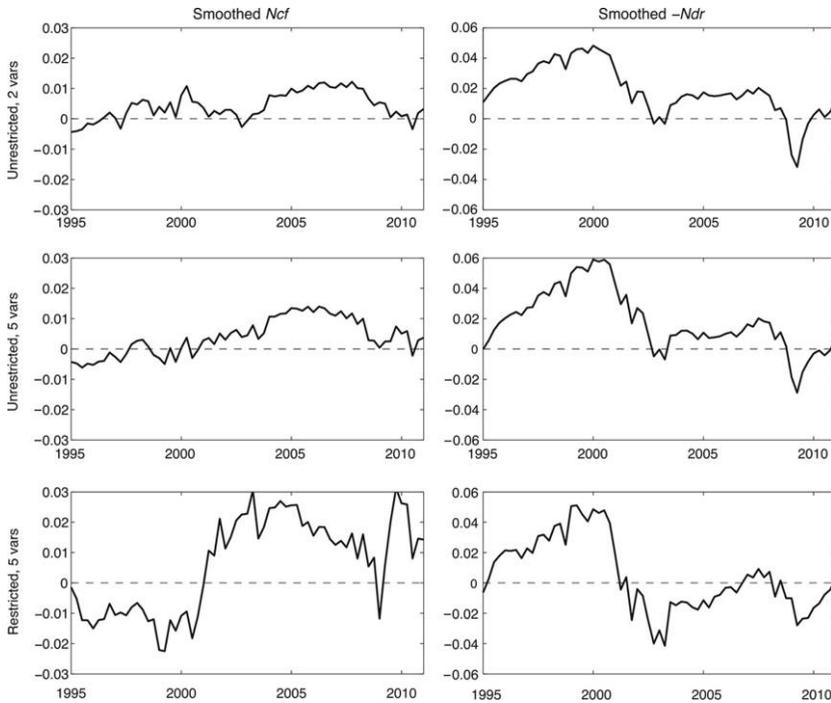


Fig. C2. Cash-flow and discount-rate shocks, from Campbell, Giglio, and Polk (2013)

nommic state variables during the two cycles. Price-earnings ratios rose much more dramatically during the technology boom, while earnings growth was stronger during the credit boom. And value stocks underperformed during the technology boom but outperformed during the credit boom. The ICAPM explanation of the value premium requires that value stocks are more sensitive to cash-flow news while growth stocks are more sensitive to discount-rate news, so the relative performance of value and growth stocks during the two cycles is informative about the shocks that were driving them.

I hope that this discussion has conveyed some of the many issues that are raised by Lettau and Ludvigson's decision to study the joint dynamics of consumption, labor income, and financial wealth. This choice of variables is an interesting alternative to the more standard focus on output, investment, and hours worked, or the interest rate, inflation, and the output gap. Lettau and Ludvigson are able to address questions that are of first-order importance in macro-finance. I look forward to

future work along these lines that develops more explicit and detailed structural models that can be confronted with the evidence.

Endnote

For acknowledgments, sources of research support, and disclosure of the author's material financial relationships, if any, please see <http://www.nber.org/chapters/c12933.ack>.

References

- Bansal, Ravi, Dana Kiku, and Amir Yaron. 2012. "An Empirical Evaluation of the Long-Run Risks Model for Asset Prices." *Critical Finance Review* 1: 183–221.
- Bansal, Ravi, and Amir Yaron. 2004. "Risks for the Long Run." *Journal of Finance* 59:1481–509.
- Barro, Robert J. 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics* 121:823–66.
- Beeler, Jason, and John Y. Campbell. 2012. "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment." *Critical Finance Review* 1:141–82.
- Bekaert, Geert, Eric Engstrom, and Steven R. Grenadier. 2010. "Stock and Bond Returns with Moody Investors." *Journal of Empirical Finance* 17:867–94.
- Blanchard, Olivier, and Mark Watson. 1986. "Are Business Cycles All Alike?" In *The American Business Cycle*, edited by Robert J. Gordon, 123–80. Cambridge, MA: NBER.
- Bottazzi, L., P. Pesenti, and E. van Wincoop. 1996. "Wages, Profits, and the International Portfolio Puzzle." *European Economic Review* 40:219–54.
- Burns, A., and W. Mitchell. 1946. *Measuring Business Cycles*. New York: NBER.
- Campbell, John Y. 1991. "A Variance Decomposition for Stock Returns." *Economic Journal* 101:157–79.
- . 1993. "Intertemporal Asset Pricing without Consumption Data." *American Economic Review* 83:487–512.
- . 1996. "Understanding Risk and Return." *Journal of Political Economy* 104:298–345.
- Campbell, John Y., and John H. Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107:205–51.
- Campbell, John Y., Stefano Giglio, and Christopher Polk. 2013. "Hard Times." *Review of Asset Pricing Studies* 3:95–132.
- Campbell, John Y., and Robert J. Shiller. 1988. "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors." *Review of Financial Studies* 1:195–228.
- Campbell, John Y., and Tuomo Vuolteenaho. 2004. "Bad Beta, Good Beta." *American Economic Review* 94:1249–75.
- Campbell, John Y., and Motohiro Yogo. 2006. "Efficient Tests of Stock Return Predictability." *Journal of Financial Economics* 81:27–60.
- Cochrane, John H. 2008. "The Dog That Did Not Bark: A Defense of Return Predictability." *Review of Financial Studies* 21:1533–75.

- Constantinides, George. 1990. "Habit Formation: A Resolution of the Equity Premium Puzzle." *Journal of Political Economy* 98:519–43.
- Dew-Becker, Ian. 2012. "A Model of Time-Varying Risk Premia with Habits and Production." Unpublished Paper, Federal Reserve Bank of San Francisco.
- Epstein, Lawrence, and Stanley Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57:937–69.
- Fama, Eugene F., and Kenneth R. French. 1988. "Dividend Yields and Expected Stock Returns." *Journal of Financial Economics* 22:3–25.
- Hobijn, Bart, and Boyan Jovanovic. 2001. "The Information-Technology Revolution and the Stock Market: Evidence." *American Economic Review* 91:1203–20.
- Jacobson, M. M., and F. Occhino. 2012. "Labor's Declining Share of Income and Rising Inequality." *Federal Reserve Bank of Cleveland Economic Commentary*, September 12.
- Karabarbounis, Loukas, and Brent Neiman. 2013. "The Global Decline of the Labor Share." NBER Working Paper no. 19136, Cambridge, MA.
- Kogan, Leonid, and Dimitris Papanikolaou. 2013. "Growth Opportunities, Technology Shocks, and Asset Prices." Forthcoming, *Journal of Finance*.
- Laitner, John, and Dmitriy Stolyarov. 2003. "Technological Change and the Stock Market." *American Economic Review* 93:1240–67.
- Lettau, Martin, and Sydney Ludvigson. 2001. "Consumption, Aggregate Wealth, and Aggregate Stock Returns." *Journal of Finance* 56:815–49.
- . 2004. "Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption." *American Economic Review* 94:276–99.
- Lewellen, Jonathan. 2004. "Predicting Returns with Financial Ratios." *Journal of Financial Economics* 74:209–35.
- Merton, Robert C. 1973. "An Intertemporal Capital Asset Pricing Model." *Econometrica* 41:867–87.
- Parker, Jonathan, and Christian Julliard. 2005. "Consumption Risk and Cross-Sectional Returns." *Journal of Political Economy* 113:185–222.
- Poterba, James M., and Lawrence H. Summers. 1988. "Mean Reversion in Stock Prices: Evidence and Implications." *Journal of Financial Economics* 22:27–59.
- Rietz, Thomas A. 1988. "The Equity Risk Premium: A Solution." *Journal of Monetary Economics* 22:117–31.
- Stambaugh, Robert F. 1999. "Predictive Regressions." *Journal of Financial Economics* 54:375–421.