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): Mark W. Watson

Chapter URL: http://www.nber.org/chapters/c12934

Chapter pages in book: (p. 367 – 378)

Comment

Mark W. Watson, *Department of Economics, Princeton University, and NBER*

Lettau and Ludvigson's (2004) paper is a classic in the empirical cointegration literature. Motivated by an intertemporal budget constraint for a representative agent, that paper derived a cointegrating relation involving the logarithm of consumption (c), labor income (y), and asset values (a). Post–World War II quarterly US data were consistent with the cointegrating relation. Moreover, the implied vector error correction model (VECM) had a surprising structure: deviations from the long-run cointegrating relation were "corrected" by changes in asset values, not by changes in consumption or labor income. Said differently, the model's error-correction term, labeled *cay* by Lettau and Ludvigson, forecast future asset returns (Δa), but not future changes in Δc or Δy.

Lettau and Ludvigson's current paper extends their 2004 paper in two ways. First, it updates and extends the data through 2012:Q3, allowing an out-of-sample external validity check of their previous results. Second, it carries out an SVAR (structural vector autoregression) exercise to give an interpretation to the VECM's errors.

Much of my discussion will focus on the first of these extensions. I will argue that Lettau and Ludvigson's cointegration result, which the data clearly supported in their 2004 paper, is far less clear in the extended data.[1] I will then show that many, but not all, of the SVAR results they present in this paper are robust to uncertainty about cointegration.

Evidence on Cointegration in the Extended Sample

Figure C1 plots the c, a, and y data used in this paper along with the *cay* error-correction term, $c_t - 0.18a_t - 0.70y_t$. The c, a, and y series have clear upward trends, and these trends are absent in *cay*. In this sense,

978-0-226-16540-0/2014/2013-0503$10.00

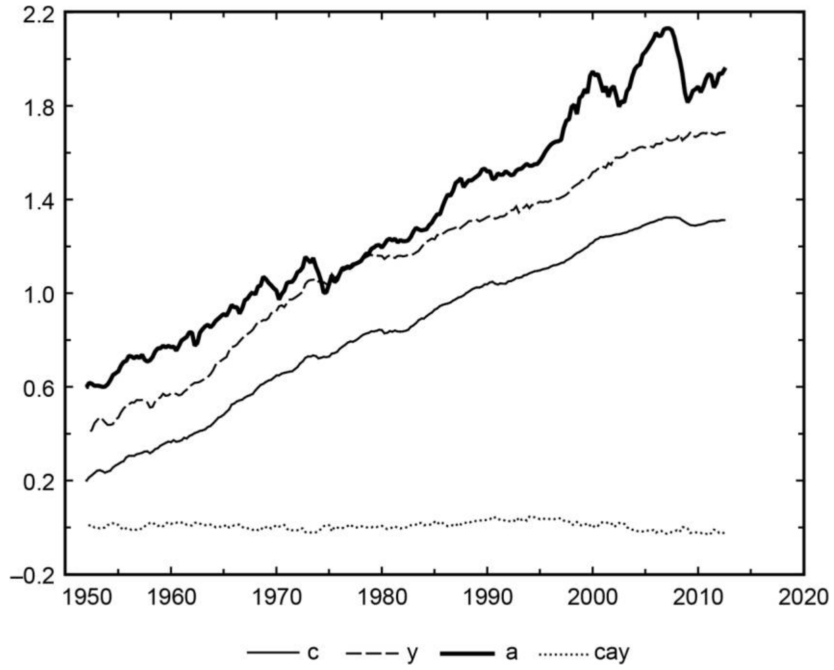

Fig. C1. *c*, *a*, *y*, and *cay*

cay is far less persistent than its constituents. Figure C2 changes the graph's scale and plots *cay* in isolation. While the large trends in *c*, *a*, and *y* are absent, *cay* is nevertheless very persistent, with three large low-frequency swings, each lasting approximately twenty years.

Table C1 presents standard autoregressive persistence statistics for each of the series plotted in figure C1. After eliminating a deterministic linear trend, *c* and *y* remain highly persistent, with 90 percent confidence intervals for the largest AR root that exceed unity. The series *a* is somewhat less persistent; its 90 percent largest-root confidence interval ranges from 0.89 to 1.01. Importantly, the largest AR root for *cay* also appears to be large: the median unbiased estimate is 1.01 and the 90 percent confidence interval ranges from 0.96 to 1.02.[2]

Figures C3 and C4 use a less parametric approach to gauge the persistence in *cay*. These figures are based on weighted averages that summarize low-frequency variability in a series. Specifically, figure C3 plots the "cosine transformations" of *cay*

$$f_j = \sum_{t=1}^{T} \cos(j(t - 0.5)\pi T^{-1})\, cay_t, \text{ for } j = 1, \ldots, 12.$$

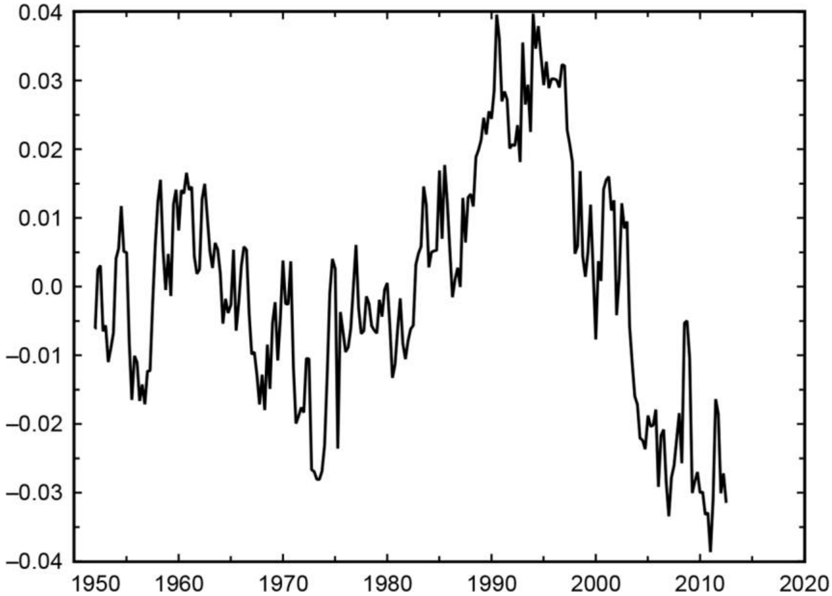

Fig. C2. *cay*

Table 1
AR persistence statistics

$$x_t = \mu_t + u_t$$
$$\mu_t = \beta_0 + \beta_1 t$$
$$\rho(L)u_t = \varepsilon_t$$

Series (x)	$\hat{\rho}^{OLS}(1)$	$\hat{\rho}^{MUB}(1)$	90% confidence interval for $\rho(1)$	*DF tstat*
c	1.00	1.01	1.01 to 1.02	–0.10
a	0.94	0.94	0.89 to 1.01	–3.19
y	0.99	1.01	1.01 to 1.02	–0.77
cay	0.95	1.01	0.96 to 1.02	–1.79

Notes: $\hat{\rho}^{OLS}(1)$ is the OLS estimate of the sum of the AR coefficients, $\hat{\rho}^{MUB}(1)$ is the median unbiased estimate, and *DF tstat* is the Dickey-Fuller t-statistic. $\hat{\rho}^{MUB}(1)$ and the 90 percent confidence interval are computed using Stock's (1991) method. The results for *cay* do not account for sampling error in the estimated cointegrating coefficients.

As discussed in Müller and Watson (2008), the set of averages $\{f_j\}_{j=1}^{k}$ capture the variability in cay_t for periods greater than $2T/k$, where T is the sample size. Thus, with T = 60 years in the Lettau-Ludvigson sample, the k = 12 points plotted in figure C3 summarize the variability in *cay* for periods greater than ten years.

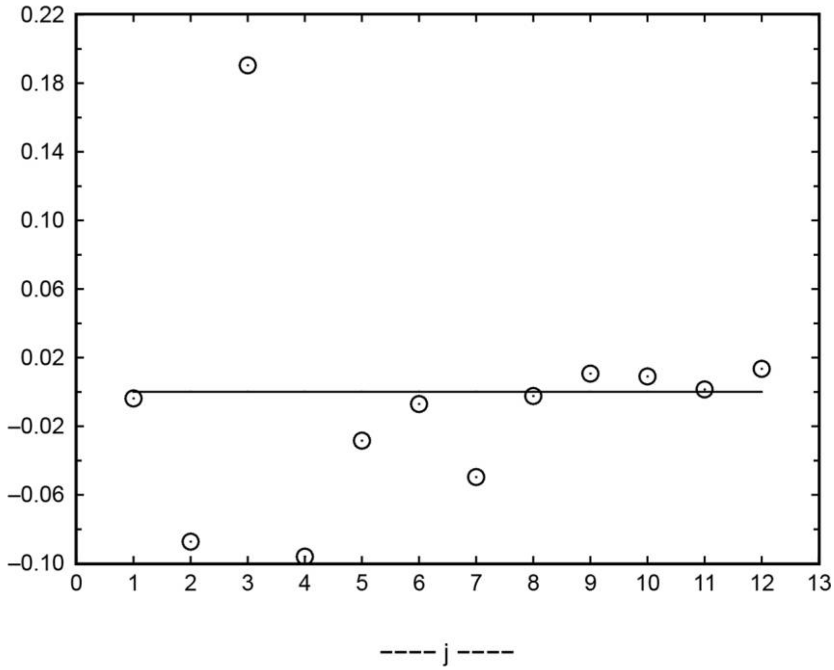

Fig. C3. Low-frequency weighted averages of *cay*

The low-frequency cosine weighted averages (f_j) of an $I(0)$ process are approximately i.i.d. $N(0,\sigma^2_{LR})$ random variables, where σ^2_{LR} is the long-run variance of the process. For a series that is more persistent than $I(0)$, the f_j random variables are heteroskedastic with higher variance for smaller values of j. (Smaller values of j correspond to lower frequencies, and more persistence translates into more low-frequency variability.) If c, a, and y were cointegrated, then the f_j values plotted in figure C3 should look like i.i.d. random variables. Instead, the figure indicates a much higher variance for f_j associated with low frequencies, so that *cay* is more persistent than an $I(0)$ process.

Figure C4 uses a parameterization from a fractionally integrated model to formalize the persistence characterized by the heteroskedasticity in figure C3. Write $(1 - L)^d cay_t = u_t$, where u_t is an $I(0)$ process. If *cay* is $I(0)$ (i.e., if c, a, and y are cointegrated) then $d = 0$, while if *cay* is $I(1)$, then $d = 1$. More generally, the fractional model allows for noninteger values of d, where values of $d > 0$ characterize processes that are more persistent than $I(0)$. Figure C4 plots the log-likelihood values for d calculated using the f_j data, normalized so that the $I(0)$ model has value of 0.

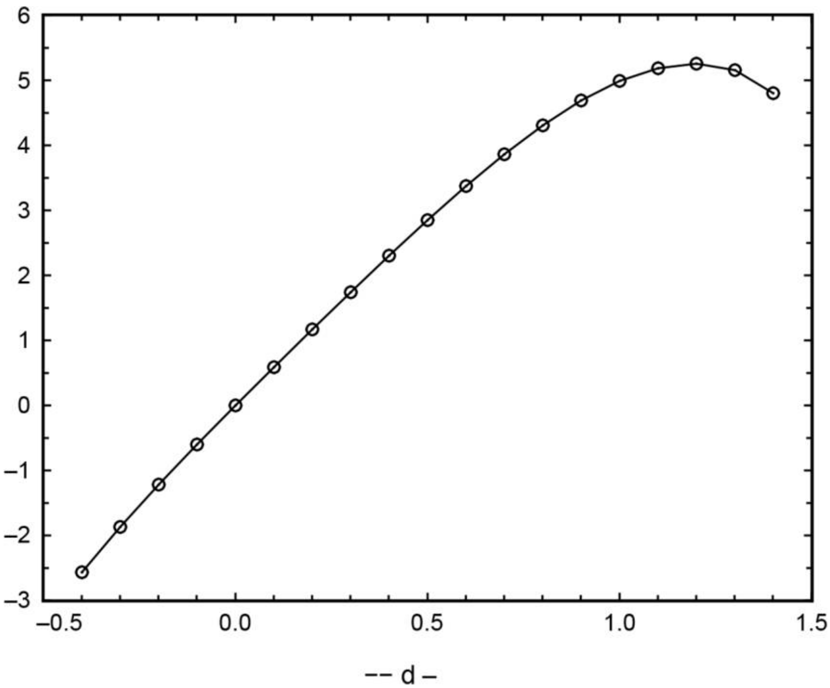

Fig. C4. Low-frequency log-likelihood values for $(1 - L)^d cay$

The log likelihood for $d = 1$ is five points larger than for $d = 0$. Thus, processes with persistence close to the $I(1)$ model fit the *cay* data much better than the $I(0)$ model.

Robustness of the SVAR Results to Cointegration

The structural VAR estimated by Lettau and Ludvigson has the form:

$$\begin{bmatrix} \Delta c_t \\ \Delta y_t \\ \Delta a_t \end{bmatrix} = \mu + \begin{bmatrix} \phi_{11} & \phi_{12} & \phi_{13} \\ \phi_{21} & \phi_{22} & \phi_{23} \\ \phi_{31} & \phi_{32} & \phi_{33} \end{bmatrix} \begin{bmatrix} \Delta c_{t-1} \\ \Delta y_{t-1} \\ \Delta a_{t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \gamma_a \end{bmatrix} cay_{t-1} + H \begin{bmatrix} \eta_{1t}^P \\ \eta_{2t}^P \\ \eta_{1t}^T \end{bmatrix}, \quad (1)$$

where H is lower triangular, η_{1t}^P is the first permanent structural shock (labeled the *productivity shock* by Lettau and Ludvigson), η_{2t}^P is the second permanent shock (which they label as the *factor shares shock*), and η_{1t}^T is the transitory shock (labeled as the *risk aversion shock*). This formulation imposes two restrictions on the VECM: first, consistent with their

2004 specification, *cay* enters the VAR equation for Δa, but not the equations for Δc or Δy, and second, the lower triangular restriction on H means that the structural shocks are identified by a $\Delta c \to \Delta y \to \Delta a$ Wold causal ordering.

To understand the role of cointegration for identifying the structural shocks, it is useful to rewrite (1) as

$$\begin{bmatrix} \Delta c_t \\ \Delta y_t \\ \Delta cay_t \end{bmatrix} = \mu + \begin{bmatrix} \phi_{11} & \phi_{12} & \phi_{13} \\ \phi_{21} & \phi_{22} & \phi_{23} \\ \tilde{\phi}_{31} & \tilde{\phi}_{32} & \tilde{\phi}_{33} \end{bmatrix} \begin{bmatrix} \Delta c_{t-1} \\ \Delta y_{t-1} \\ \Delta cay_{t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \tilde{\gamma}_a \end{bmatrix} cay_{t-1} + \tilde{H} \begin{bmatrix} \eta^P_{1t} \\ \eta^P_{2t} \\ \eta^T_{1t} \end{bmatrix}, \quad (2)$$

where the Δc and Δy equations in (2) are identical those in (1), and the final equation in (2) replaces the equation for Δa in (1) with an equation for Δcay. (The equation for *cay* is a linear combination of the equations in [1], with coefficients given by the cointegrating coefficients.) The matrix $\tilde{H}$ in (2) retains the same lower triangular form as H in (1).

In the VECM (2), *cay* follows an $I(0)$ process (so that c, a, and y are cointegrated) when $\tilde{\gamma}_a < 0$. In contrast, *cay* follows an $I(1)$ process (so that c, a, and y are *not* cointegrated) when $\tilde{\gamma}_a = 0$. The first two equations are unaltered by the cointegration restriction (whether $\tilde{\gamma}_a = 0$ or $\tilde{\gamma}_a < 0$). Because the first two structural shocks, η^P_{1t} and η^P_{2t}, are computed from the errors in the first two equations, they are unaltered by the cointegration restriction. Thus, Lettau and Ludvisgon's *productivity shock* and *factor shares shock* are invariant to the cointegration restriction. However, because the *risk average shock*, η^T_{1t}, is computed from the third equation, it depends on the cointegration restriction.

When $\tilde{\gamma}_a < 0$, c,a, and y are cointegrated, *cay* is $I(0)$, and η^T_{1t} is "transitory" in the sense that its effect of c_{t+h}, a_{t+h}, and y_{t+h} decays exponentially as h grows large. In contrast, when $\tilde{\gamma}_a = 0$, c, a, and y are not cointegrated, *cay* is $I(1)$, and η^T_{1t} is "permanent" in the sense that its effect c_{t+h}, a_{t+h}, and y_{t+h} remains bounded away from zero as h grows large. The estimated value of $\tilde{\gamma}_a$ is –0.059, so that η^T_{1t} is transitory in the estimated model (see, e.g., Lettau and Ludvigson's figure 4).[3] Figure C5 contrasts the impulse responses of η^T_{1t} in the model with $\tilde{\gamma}_a = 0$ with those in the estimated model from Lettau and Ludvigson. The impulse responses are similar over short horizons, say, two to three quarters, but are substantially different over longer horizons.

Because nearly all of the variability of c and y (see Lettau and Ludvigson's figure 4) is explained by η^P_{1t} and η^P_{2t}, the identification of η^T_{1t} has little effect on the paper's conclusions about c and y. In contrast, much

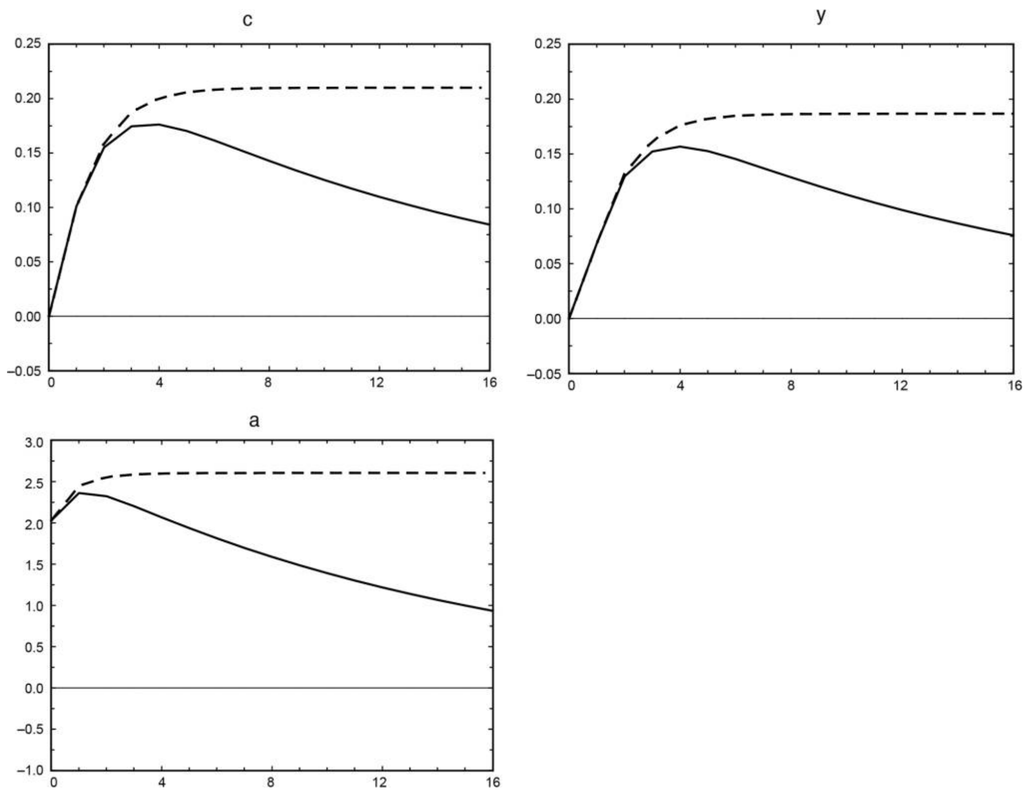

Fig. C5. Impulse responses for η_{1t}^{T} ("risk aversion" shock) ($I(0)$ is solid, $I(1)$ is dashed)

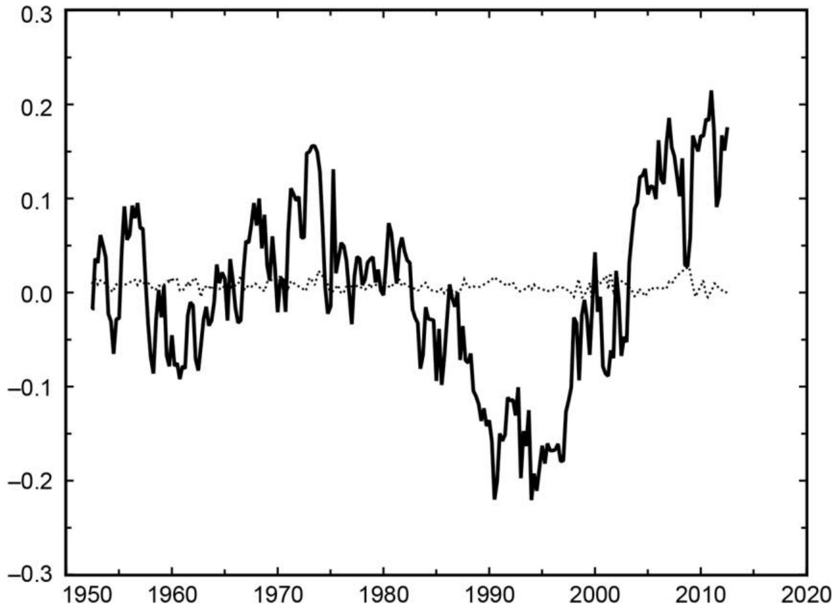

Fig. C6. Transitory component of $a_t a_t - E(a_{t+k} - k\mu_a \mid I_t)$ ($I(0)$ is solid, $I(1)$ is dashed)

of the variation in a is explained by η_{1t}^{T}, so that some of the paper's conclusions about a depend importantly on the cointegration restriction. In particular, the low-frequency characteristics of a depend on the cointegration restriction. I provide two examples.

Figure C6 plots the Beveridge-Nelson (1981) transitory component of asset values: $a_t - \tau_t$, where $\tau_t = \lim_{k\to\infty} E_t(a_{t+k} - k\mu_a)$ and μ_a is the mean of Δa for the estimated cointegrated model. Two plots are shown: the solid curve is for the VECM model estimated by Lettau and Ludvigson, and the dashed curve is for the non-cointegrated model (i.e., the estimated model imposing $\tilde{\gamma}_a = 0$).

Because much of the variability in a is explained by η_{1t}^{T}, whether this shock is transitory (so that it is absent from τ) as in the VECM model, or permanent (so that it is included in τ) as in the non-cointegrated model, has an important effect of the estimated transitory component of asset values. In turn, this has an important effect on forecasts of future values of asset values.

If much of the variance of a arises from η_{1t}^{T}, and η_{1t}^{T} is transitory, then there will be large predictable "mean reversion" in a. (Where by "mean

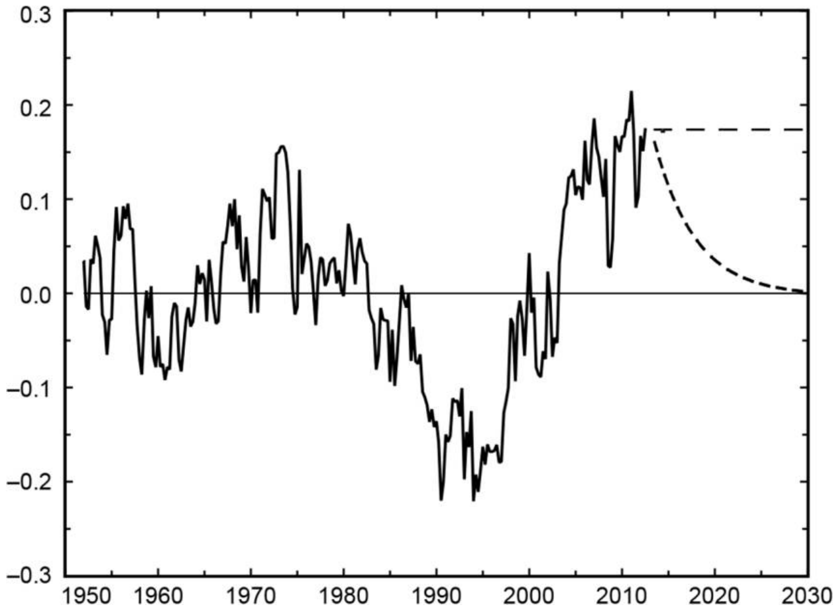

Fig. C7. Forecasts of $-\beta_a^{-1}cay_t$ (actual is solid, $I(0)$ forecast is short dashes, $I(1)$ is long dashes)

reversion" I mean reversion to the stochastic trend components in c and y.) By contrast, if η_{1t}^T is permanent, there will be far less predictability in a.

To see the relationship between predictability in a and in *cay*, note that $cay_t = c_t - \beta_a a_t - \beta_y y_t$, so that $a_t = -\beta_a^{-1}cay_t + \beta_a^{-1}(\beta_y y_t - c_t)$. Because there is relatively little forecastable variation in y and c, essentially all of the forecastability of a_t arises from cay_t. Figure C7 plots $-\beta_a^{-1}\, cay_t$ (the "*cay*" component of assent values) over the 1947 to 2012 sample period along with its forecast values through 2030. Two forecast paths are plotted: the first is the path implied by Lettau and Ludvigson's estimated VECM and the second is for the estimated non-cointegrated model. In the cointegrated model, *cay* is $I(0)$, so that its large 2012 value is predicted to disappear over the 2012 to 2030 period, yielding fall in asset values of approximately 20 percent. In contrast, under the $I(1)$ model, there is little predictable variation in *cay* and therefore little forecastible change in the *cay* component of asset values.

Of course, both the $I(1)$ and $I(0)$ models are constrained models of persistence. Even limiting attention to the $I(d)$ model discussed earlier, figure C4 suggests that there are many values of d beyond those implied by the $I(0)$ and $I(1)$ models that may characterize the data. In terms of figure C6, these lead to a wide range of forecast paths for future values of the *cay* component of a. Figure C7 summarizes the uncertainty about future values of this component by plotting 90 percent prediction sets for average returns: $\overline{r}_{t:t+h} = -\beta_a^{-1}h^{-1}(cay_{t+h} - cay_t)$. These prediction intervals are constructed so that they contain the realized value of $\overline{r}_{t:t+h}$ with 90 percent probability. The dashed lines show the prediction set for the cointegrated model (so that cay_t is $I(0)$) using the estimated parameters from the Lettau and Ludvigson VECM model. Under this model, the large value of *cay* in 2012:Q4 (the end of the sample) is largely the result of transitory shocks, so that *cay* is predicted to decline with high probability over the forecast period. Returns are therefore forecast to be large and negative with high probability. Thus, the 90 percent prediction set for average returns from 2012 to 2020 ranges from roughly –3 percent to –2 percent, and the 90 percent prediction set for average returns from 2012 to 2030 ranges from –1.2 percent to –0.8 percent. Evidently, if asset values follow the VECM model estimated by Lettau and Ludvigson, then with very high probability returns will average roughly 1 percent over the next two decades.

Figure C8 also shows an alternative 90 percent prediction set that is much wider. This prediction is constructed using methods outlined in Müller and Watson (2013) that incorporate uncertainty about the stochastic process for *cay*.[4] Because, for example from figure C4, *cay* may be more persistent than an $I(0)$ process, the prediction is much wider. This 90 percent prediction set suggests that average returns are far less forecastible and far less certain: the 90 percent prediction set ranges from roughly –2 percent to +3.5 percent over the prediction period.

Summary

Lettau and Ludvigson's paper both updates their important 2004 study and provides an interpretation for the shocks in their VECM. The data extended through 2012:Q3 suggest more persistence in *cay* than was evident in the earlier data, and this calls into question the cointegration restriction in their estimated VECM. That said, many of Lettau and Ludvigson's conclusions—those related to what they call *productivity* and *factor share* shocks—are robust to non-cointegration.

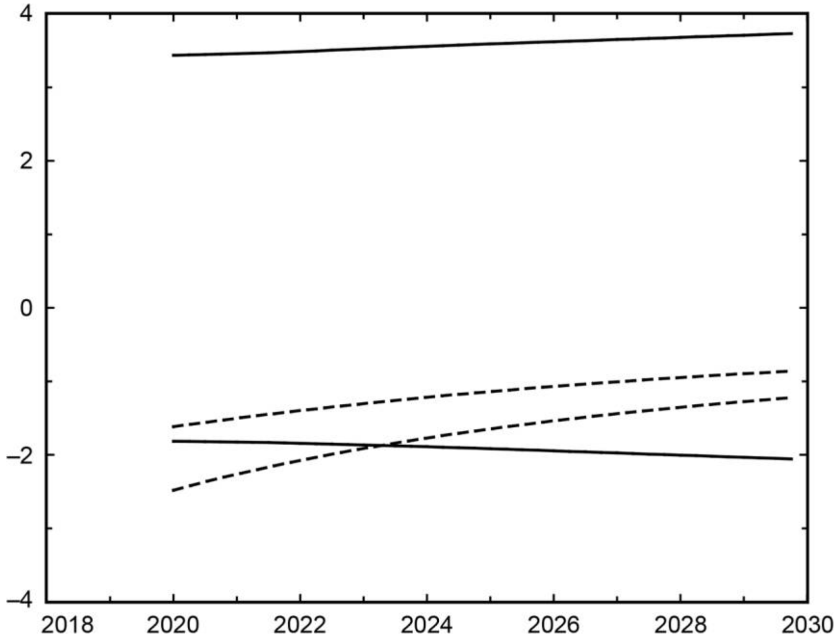

Fig. C8. 90 percent predictions sets for average returns from 2012:Q3 to $T + h$ – ($I(0)$ model is dashed. General persistence model is solid.)

Endnotes

For acknowledgments, sources of research support, and disclosure of the author's material financial relationships, if any, please see http://www.nber.org/chapters/c12934.ack.

1. Lettau and Ludvigson provide some discussion of this in section V of the current paper. Out-of-sample fragility of cointegrating relations is not unique to Lettau and Ludvigson (2004). Other examples that come to mind are the "great ratios" in Klein and Kosobud (1961), the balanced growth and other relations investigated in King et al. (1991), and the money demand relations investigated in Stock and Watson (1993).

2. The calculations for *cay* do not account for sampling error in the estimated cointegrating coefficients, which impart spurious mean reversion into the series. For example, the 5 percent one-sided critical value for the Dickey-Fuller t-statistic falls from –3.41 to –3.80 when sampling error is taken into account. (See Stock and Watson 2011, chapter 16.)

3. The estimated value of $\tilde{\gamma}_a$ is –0.059, with an estimated standard error of 0.026. Because of persistence in *cay*, the usual t-statistic will not have an approximate normal distribution and reliable inference about the true value of $\tilde{\gamma}_a$ requires nonstandard methods familiar from the literature on unit roots and near unit roots.

4. This is the "MN" frequentist set for the "*abd*" model described in Müller and Watson (2013). They describe the Bayes set, which coincides the Bayes prediction set computed for the low-frequency $I(d)$ model with uniform prior to d in (–4 to 1.4). See Pástor and Stambaugh (2012) for related calculations about long-run returns using parametric Bayes methods.

References

Beveridge, S., and C. R. Nelson. 1981. "A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the Business Cycle." *Journal of Monetary Economics* 7:151–74.

King, R., C. Plosser, J. H. Stock, and M. W. Watson. 1991. "Stochastic Trends and Economic Fluctuations." *American Economic Review* 81:819–40.

Klein, L. R., and R. F. Kosobud. 1961. "Some Econometrics of Growth: Great Ratios of Economics." *Quarterly Journal of Economics* 75 (2): 173–98.

Lettau, M., and S. C. Ludvigson. 2004. "Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption." *American Economic Review* 94:276–99.

Müller, U. K., and M. W. Watson. 2008. "Testing Models of Low-Frequency Variability." *Econometrica* 76 (1): 155–74.

———. 2013. "Measuring Uncertainty about Long-Run Predictions." Unpublished Manuscript, Princeton University.

Pástor, L., and R. F. Stambaugh. 2012. "Are Stocks Really Less Volatile in the Long Run?" *Journal of Finance* 87 (2): 431–77.

Stock, J. H. 1991. "Confidence Intervals for the Largest Autoregressive Root in US Economic Time Series." *Journal of Monetary Economics* 28:435–60.

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

———. 2011. *Introduction to Econometrics*, 3rd ed. Boston: Addison Wesley.