NBER WORKING PAPER SERIES

THE STOCK MARKET CRASH REALLY DID CAUSE THE GREAT RECESSION

Roger Farmer

Working Paper 19391
http://www.nber.org/papers/w19391

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2013

This study was written while the author was a Senior Houblon-Norman Fellow at the Bank of England. I would like to thank David Rosnick of the CEPR in Washington for his comments on my previously published paper (Farmer 2012b) . I would also like to thank the Trustees of the Houblon-Norman Fund for their support and C. Roxanne Farmer for her editorial assistance. The views expressed herein do not reflect those of the Bank of England, nor those of the Monetary Policy Committee. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The Stock Market Crash Really Did Cause the Great Recession
Roger Farmer
NBER Working Paper No. 19391
August 2013
JEL No. E24,E27,E32

ABSTRACT

This note shows that a big stock market crash, in the absence of central bank intervention, will be followed by a major recession one to four quarters later. I establish this fact by studying the forecasting ability of three models of the unemployment rate. I show that the connection between changes in the stock market and changes in the unemployment rate has remained structurally stable for seventy years. My findings demonstrate that the stock market contains significant information about future unemployment.

Roger Farmer
UCLA
Department of Economics
Box 951477
Los Angeles, CA 90095-1477
and NBER
rfarmer@econ.ucla.edu

I. Introduction

In a recent paper in the *Journal of Economic Dynamics and Control*, (Farmer, 2012b), I pointed out that there are transformations of the U.S. unemployment rate, and the real value of the S&P 500 that are non-stationary but cointegrated. I provided a Vector Error Correction Model (VECM) where changes in stock market wealth cause changes in the unemployment rate. I estimated this model, using data on unemployment and the real value of the S&P 500 from 1953q1 through 1979q3, and I showed that the model provides an excellent fit to data from 1979q4 through 2011q1.

Rosnick (2013) has argued that a univariate model provides a better prediction of the unemployment rate than my published model. I show here, that although the univariate model provides more accurate out-of-sample forecasts than the VECM, a bivariate model that includes information from the stock market outperforms both alternatives. My results establish that the stock market contains significant information that helps to predict the future unemployment rate. A big stock market crash, in the absence of central bank intervention, will be followed by a major recession one to four quarters later. Further, the connection between changes in the stock market and changes in the unemployment rate has remained structurally stable for seventy years.

II. Related Literature

My previous work is related to Lettau and Ludvigson (2004) and Lettau and Ludvigson (2011) who provide a statistical model of consumption, wealth and labor earnings as non-stationary, but cointegrated, time series. The connection between stock market wealth and unemployment was recognized by Phelps (1999) who pointed out that the stock market boom of the 1990s was accompanied by a reduction in the unemployment rate. Fitoussi, Jestaz, Phelps, and Zoega (2000) found a similar correlation between the stock market and unemployment for a variety of European countries. Following Phelps (1999) and Hoon and Phelps (1992), these authors explained this connection using Phelps' (1994) structuralist model of the natural rate of unemployment. In Phelps' model, expectations of future profits cause firms to invest in customer relationships and employee training. The explanation for persistent unemployment provided in Farmer (2012b) is closer to the models of hysteresis described by Blanchard and Summers (1987, 1986) and Ball (1999) than the structuralist model of Phelps although the theoretical foundation for persistent unemployment in Farmer (2012b,a, 2013) is very different from the one provided in those of Blanchard and Summers.

III. Three Different Models

To establish my claim that the stock market helps to predict the unemployment rate, I estimate three different models on data from 1953q1 through 1979q3, and I compare their forecast performance for the sample period 1979q4 through 2011q1. Model 1 is the VECM that I reported in Farmer (2012b), model 2 is a univariate model for the unemployment rate and model 3 is a bivariate vector autoregression. Models 2 and 3 were estimated in first differences. Model 1 was estimated in first differences but includes a cointegrating vector with lagged level information.[1]

Parameter estimates for the three models are recorded below as equations (1), (2) and (3).[2] The coefficients on levels in the cointegrating equation are broken down into the loading factors, α; a 2×1 vector, and the cointegrating equation, β^T; a 1×3 vector. The symbol c stands for the constant.

$$\text{VECM} \qquad \begin{bmatrix} \Delta u_t \\ \Delta p_t \end{bmatrix} = \begin{bmatrix} 0.6 & -0.27 \\ 0.13 & 0.42 \end{bmatrix} \begin{bmatrix} \Delta u_{t-1} \\ \Delta p_{t-1} \end{bmatrix} + \overset{\alpha}{\begin{bmatrix} -0.1 \\ 0.01 \end{bmatrix}} \overset{\beta^T}{\begin{bmatrix} 1 & 0.6 & -7 \end{bmatrix}} \begin{bmatrix} u_{t-1} \\ p_{t-1} \\ c \end{bmatrix} \tag{1}$$

$$\text{Univariate} \qquad [\Delta u_t] = 0.6\,[\Delta u_{t-1}] \tag{2}$$

$$\text{Bivariate} \qquad \begin{bmatrix} \Delta u_t \\ \Delta p_t \end{bmatrix} = \begin{bmatrix} 0.6 & -0.3 \\ 0.1 & 0.38 \end{bmatrix} \begin{bmatrix} \Delta u_{t-1} \\ \Delta p_{t-1} \end{bmatrix} \tag{3}$$

In all cases u_t is the logarithm of a logistic transformation of the unemployment rate and p_t is the logarithm of the S&P 500, measured in wage units.

IV. The Three Models Compared

In Figure 1, I compare the 1-step ahead forecast errors for the period 1979q4 – 2011q1.

[1] All three models were estimated in Eviews using a data set available on my website. I have also made the Matlab code available that was used to construct Figures 1 through 7 in this note.

[2] The estimates reported in Farmer (2012b, page 698) contain a sign error. The coefficient on the lagged value of the stock market in the cointegrating equation should be +0.6 as reported here and not −0.6 as reported in the published paper. I omit estimates of the constants in models 2 and 3 since they are insignificantly different from zero.

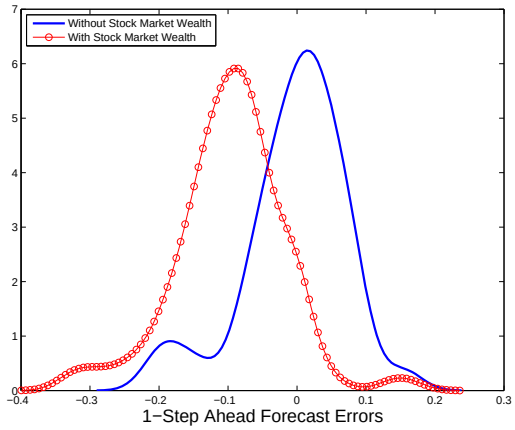

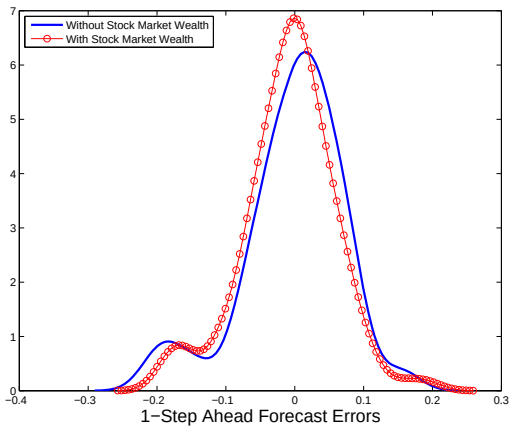

VECM vs Univariate Bivariate vs Univariate

Figure 1: 1-step ahead forecast errors for the three models

The left panel compares the VECM with the univariate model[3]; the right panel compares the bivariate and univariate models. On both panels the solid line is a smoothed histogram of prediction errors from the univariate model and the line with circles is the smoothed histogram for the comparison model. These panels show that the univariate model outperforms the VECM, but the bivariate model is better than both. Further, the distribution of univariate errors has a positive mean, indicating bias in the prediction, whereas that of the bivariate model is centered on zero, indicating that it provides unbiased estimates of unemployment out of sample. This result holds, not only for 1-step ahead forecast errors, but also at longer forecast horizons.

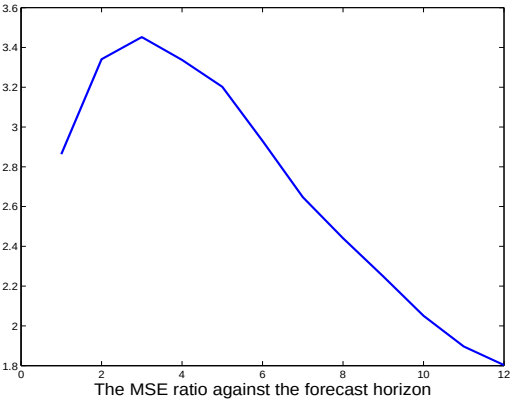

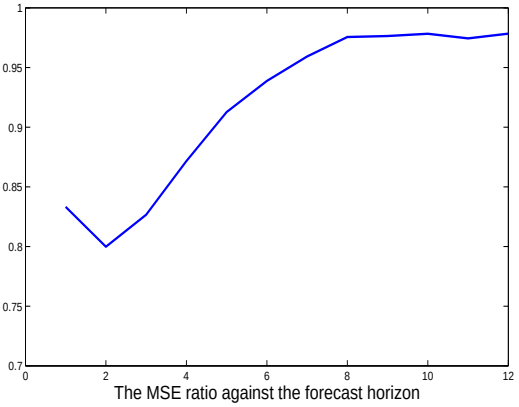

VECM vs. Univariate Bivariate vs. Univariate

Figure 2: MSE prediction errors at different forecast horizons

[3]This reproduces Figure 3 from Rosnick (2013).

Figure 2 plots the ratio of the mean squared forecast error of the comparison model to that of the univariate model, plotted as a function of the forecast horizon, for forecast horizons up to three years (12 quarters). The left panel shows that this ratio is greater than 1 at all horizons, indicating that the univariate model beats the VECM. The right panel shows that this result is reversed for the bivariate model which has a MSE ratio less than 1 at all horizons. **These results show that the stock market contains significant information that helps to predict the unemployment rate at all horizons up to and including** 12 **quarters.**

The critical observer might think that the difference between the errors from the bivariate and univariate models are small; after all, an error that is 80% of the univariate model may not be important from a policy perspective. The following section shows that this is not the case.

V. Forecasting the Great Recession

Figure 4 shows that between 2007q2 and 2009q1 the S&P 500, measured in wage units, lost 50% of its real value falling from a high of approximately 24,000 to a low of roughly 12,000. At the same time, the unemployment rate climbed from 4.5% to 10%. But could we have used the information that the stock market crashed to help forecast the Great Recession?

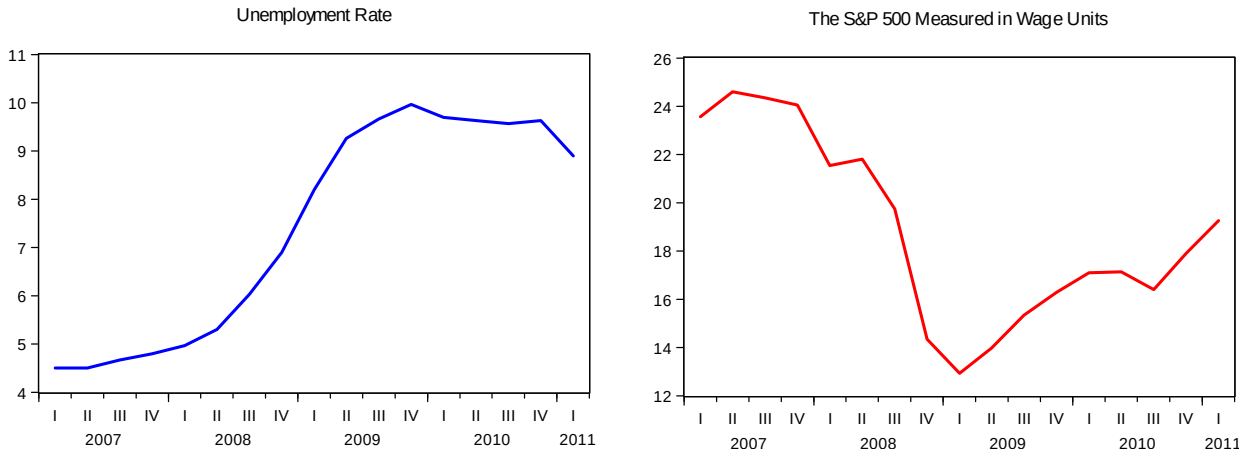

Figure 3: Unemployment and the Stock Market

Figures 4 through 6 compare univariate and bivariate dynamic predictions for the unemployment rate at three different forecast dates. In each panel, the actual path of the unemployment rate appears as a solid line. The dash-dot line is the forecast from the univariate model and the dashed line is the forecast from the comparison model. In the left panel, the comparison model is the VECM; in the right panel it is the

bivariate model. These three figures show that the bivariate model outperforms the other two and, together, they imply that the stock market has considerable predictive power if our goal is to predict the unemployment rate one to twelve quarters ahead.[4]

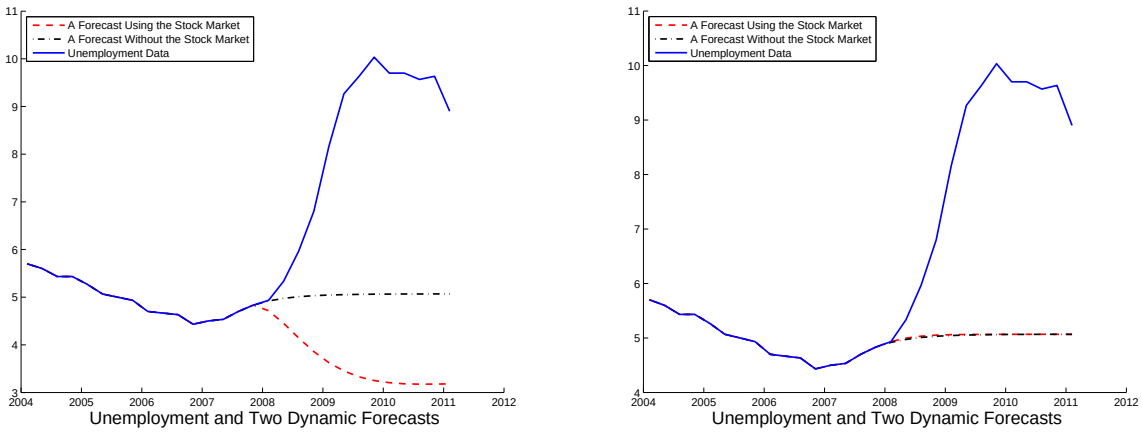

VECM vs. Univariate Bivariate vs. Univariate

Figure 4: Forecasts from 3 Models in the Fall of 2007

Figure 4 shows the dynamic forecasts that would be made by an economist, standing in the fourth quarter of 2007, using VECM, univariate and bivariate models estimated on data from 1953q1 to 1979q4. The left panel shows that the VECM does a poor job and it is apparent from this graph, that the VECM is seriously mis-specified. It predicts a large drop in the unemployment rate in 2008 in contrast to the path of unemployment that actually occurred.

The right panel shows that the problem with the VECM does not come from adding the first difference of the stock market to the univariate model. Adding the first difference of the stock market, as opposed to the level of the stock market, does not damage the univariate forecast; but it does not improve it either. That is unsurprising since, at this date, the stock market had not yet begun its spectacular decline.[5]

Figure 5 shows a dynamic forecast made in the fourth quarter of 2008. At this point, Lehman brothers had declared bankruptcy and, as is evident from Figure 3, there had been a large drop in the S&P 500. The right panel of Figure 5 shows that

[4]The left panel contains some of the same information as Figure 4 in Rosnick (2013).

[5]There was a substantial drop in housing wealth, beginning in the fall of 2006. In my view, that drop triggered a subsequent increase in the unemployment rate. But it was the precipitous crash in the stock market, beginning in the fall of 2008, that turned an otherwise mild contraction into what we now refer to as the Great Recession.

the bivariate model correctly predicts the magnitude of the Great Recession, three quarters ahead, and overshoots in the fourth quarter and beyond. In contrast, the univariate model misses the depth of the increase in the unemployment rate by two full percentage points.

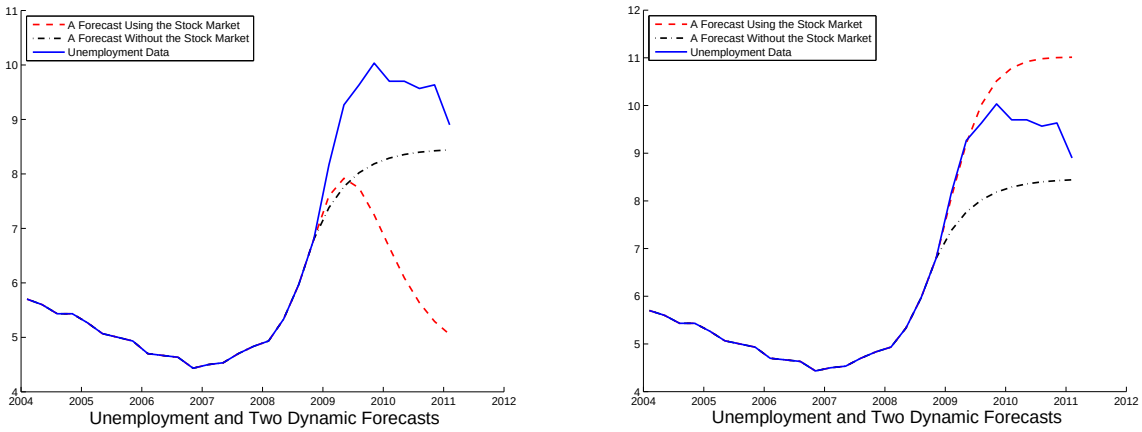

VECM vs. Univariate Bivariate vs. Univariate

Figure 5: Forecasts from 3 Models in the Fall of 2008

The left panel of Figure 5 shows that the VECM outperforms the univariate model for two quarters, but that improvement does not last long. By the third quarter, the mis-specified cointegrating equation kicks in and tries to pull the relationship between unemployment and the stock market back to its pre 1980 level.

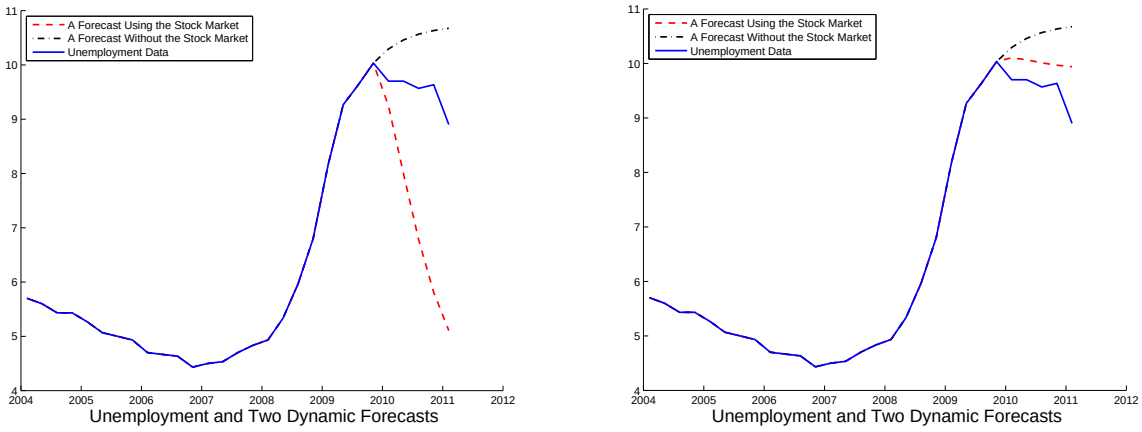

VECM vs. Univariate Bivariate vs. Univariate

Figure 6: Forecasts from 3 Models in the Fall of 2009

Finally, Figure 6 shows the dynamic forecast of the future unemployment rate using information up to and including 2009q4. At this point, the stock market had recovered quite a bit of lost ground. As a consequence, the bivariate forecast, plotted in the right panel, correctly predicts an improvement in the labor market. In contrast, the univariate model predicts that the unemployment situation will continue to deteriorate. The left panel of Figure 6 shows that, once again, the VECM performs poorly as a forecasting tool.

To understand why the VECM performs poorly, Table 1 presents estimates of the cointegrating vector for the two subsamples. This is the vector β^T in Equation (1).

First Sub-sample				Second Sub-sample			
$1953.1 - 1979.3$	u	p	c	$1979.4 - 2011.1$	u	p	c
coefficient	1	0.6	-7.4	coefficient	1	0.36	-5.3
standard error		0.25	2.3	standard error		0.09	0.87
t-statistic		(2.47)	(-3.21)	t-statistic		(3.86)	(-6.06)

Table 1: Estimates of the Cointegrating Vector

This table shows that there was a structural break in the cointegrating vector between the first and second subsamples. The coefficient on the stock market is estimated to be 0.6 in the first subsample and 0.36 in the second. Similarly, the constant in the cointegrating vector moves from -7.4 to -0.53. Although these estimates are within two standard error bounds of each other, the poor out-of-sample fit suggests that the differences are statistically important. Failing to account for a break in the cointegrating vector causes the VECM, as opposed to the bivariate differenced model, to perform badly as a forecasting tool. But that does not allow us to infer that the stock market can be ignored. **As shown in Figures 4 through 6, changes in the stock market have a large and statistically significant impact on changes in the future unemployment rate.**

VI. Simulating the Great Recession

Where does this leave my claim that the stock market crash of 2008 caused the Great Recession? Figure 7 presents the result of simulating the effect of a one quarter shock of 30% to the S&P. Thereafter, the log of the S&P follows a first order AR model in differences with a coefficient of 0.36. The log of the unemployment rate follows a bivariate VAR with a coefficient of 0.6 on the lagged log difference of unemployment

and a coefficient of -0.3 on the lagged log difference of the stock market.[6] I assume that there are no further shocks after the first quarter drop in the value of the S&P.

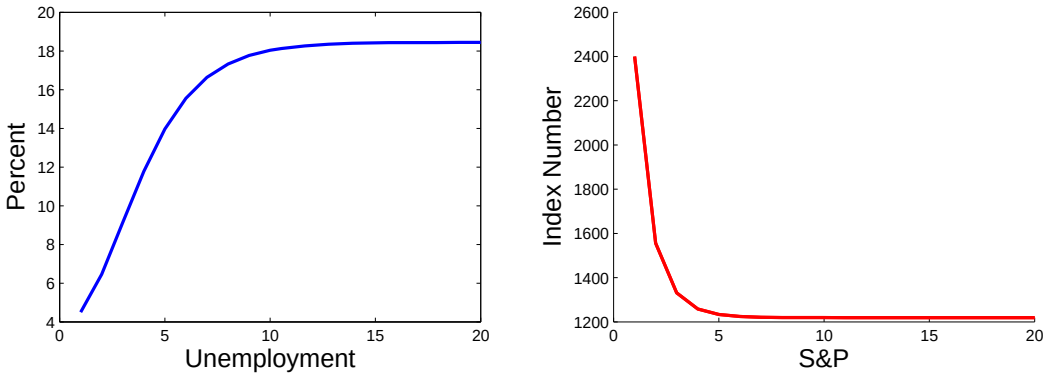

Figure 7: Simulating a Stock Market Crash

In my simulation, a once and for all one quarter shock of 30% to the value of the stock market causes the market to fall further over time, from 24,000 to 12,000. This drop mimics the realized fall in the U.S. data and it generates an increase in the unemployment rate from 4.5% to 18%, a number that is closer in magnitude to the Great Depression than the 10% peak that actually occurred. There are two reasons for that discrepancy. First, the stock market drop that actually occurred was not a single shock of 30%, followed by a smooth downward decline: it was a sequence of positive and negative shocks. In my view, the stock market recovered, in large part, because of the policy of Quantitative Easing pursued by the Federal Reserve. In the absence of that policy response, I conjecture that the path of unemployment depicted in the simulation is a good forecast of what might otherwise have occurred. **These results demonstrate that the fall in the stock market in the fall of 2008 provides a plausible quantitative explanation for the magnitude of the Great Recession.**

VII. Conclusion

What should the policy maker take away from the three simulations presented in this paper? First, the data on unemployment and the stock market are non-stationary but cointegrated. Second, although the coefficients on the lagged first differences of unemployment and the stock market are remarkably stable over seventy

[6]This version of the bivariate model uses the log of unemployment, instead of the log of the logistic transformation. The model performance is comparable with that which uses a logistic transformation for the unemployment rate. Because the coefficients are elasticities, they are easier to interpret.

years, there have been important structural breaks in the cointegrating relationship. Third, although the existence of structural breaks means that a VECM does a poor job of forecasting the future unemployment rate, a bivariate model using differenced data, can be relied upon as an accurate forecasting tool.

What should we take away from the existence of structural breaks in the cointegrating equation? In my view, it would be unwise to infer that low frequency movements of the stock market do not matter for the real economy. The failure of the VECM model as a forecasting tool does not imply that we should ignore the cointegrating relationship between unemployment and the stock market when formulating economic policy. When there are occasional breaks in cointegrating equations, models specified in first differences are known to generate more accurate forecasts, even if the data generating process is a VECM.[7]

It would be a mistake to assume, that because the cointegrating relationship has shifted since 1979, that long-run movements in the stock market do not matter for the long-run level of the unemployment rate. A safer inference would be that the models we use to inform policy decisions are not always the same ones we should use to make short-term predictions. As I have argued elsewhere (Farmer, 2010, 2012a, 2013), the stock market matters for the unemployment rate: and it matters a lot. As in Farmer (2012b), I conclude that the stock market crash of 2008 really *did* cause the Great Recession.

[7]Clements and Hendry (1988).

References

Ball, L. (1999): "Aggregate Demand and Long-Run Unemployment," *Brookings Papers on Economic Activity*, 2, 189–236.

Blanchard, O. J., and L. H. Summers (1986): "Hysterisis and the European Unemployment Problem," in *NBER Macroeconomics Annual*, ed. by S. Fischer, vol. 1, pp. 15–90. National Bureau of Economic Research, Boston, MA.

——— (1987): "Hysterisis in Unemployment," *European Economic Review*, 31, 288–295.

Clements, M., and D. Hendry (1988): *Forecasting Economic Time Series*. Cambridge University Press, Cambridge, England.

Farmer, R. E. A. (2010): *Expectations, Employment and Prices*. Oxford University Press, New York.

——— (2012a): "Confidence, Crashes and Animal Spirits," *Economic Journal*, 122(559).

——— (2012b): "The Stock Market Crash of 2008 Caused the Great Recession: Theory and Evidence," *Journal of Economic Dynamics and Control*, 36, 697–707.

——— (2013): "Animal Spirits, Financial Crises and Persistent Unemployment," *Economic Journal*, 123(568), 317–340.

Fitoussi, J.-P., D. Jestaz, E. S. Phelps, and G. Zoega (2000): "Roots of the Recent Recoveries: Labor Reforms or Private Sector Forces," *Brookings Papers on Economic Activity*, 1, 237–312.

Hoon, T., and E. S. Phelps (1992): "Macroeconomic Shocks in a Dynamized Model of the Natural Rate of Unemployment," *American Economic Review*, 82(4), 889–900.

Lettau, M., and S. C. Ludvigson (2004): "Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption," *American Economic Review*, 94(1), 276–299.

——— (2011): "Shocks and Crashes," *NYU Mimeo*.

Phelps, E. S. (1994): *Structural Slumps*. Harvard University Press, Cambridge, Mass.

——— (1999): "Behind this Structural Boom: the Role of Asset Valuations," *American Economic Review Papers and Procedings*, 89(2), 63–68.

Rosnick, D. (2013): "Farmer's Folly: Bringing the Nikkei Godzilla to America," http://www.cepr.net/index.php/blogs/cepr-blog/farmers-folly-bringing-the-nikkei-godzilla-to-america.

UCLA