

The Stock Market Crash Really Did Cause the Great Recession*

ROGER E. A. FARMER†

†*Department of Economics, UCLA, 8283 Bunche Hall Box 951477, Los Angeles, CA 90095-1477, USA (e-mail: rfarmer@econ.ucla.edu)*

Abstract

This paper studies the connection between the stock market and the unemployment rate. I establish three facts. First, the log of the real value of the S&P 500 and the log of a logistic transformation of the unemployment rate are non-stationary cointegrated series. Second, the stock market Granger causes the unemployment rate. Third, the connection between changes in the real value of the stock market and changes in the unemployment rate has remained structurally stable over seventy years. My results establish that the fall in the stock market in the autumn of 2008 provides a plausible causal explanation for the magnitude of the Great Recession.

I. Introduction

In a recent paper in the *Journal of Economic Dynamics and Control*, (Farmer, 2012b), Roger Farmer 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. Farmer provided a Vector Error Correction Model (VECM) where changes in stock market wealth cause changes in the unemployment rate. He estimated this model, using data on unemployment and the real value of the S&P 500 from 1953q1 through 1979q3, and 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 Farmer's 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.

JEL Classification numbers: E30, C10.

*I would like to thank an anonymous referee of this journal for suggestions that considerably improved the final version of the paper. I would also like to thank C. Roxanne Farmer for her editorial assistance.

The exchange between Farmer and Rosnick raises two questions, both of which I take up in this paper. The first question is philosophical. What does it mean for one time series to cause another? I establish in section V. that the stock market causes the unemployment rate in the sense of Granger (1969, 1980) and I discuss the implications of that finding for economic policy.

The second question is more narrowly defined. If the stock market Granger causes the unemployment rate, how can a model that ignores stock market information provide a more accurate forecast than one that exploits this information to inform its prediction? I answer that question in section VI. where I draw on the results of Clements and Hendry (1988, 1999) and Hendry (2004) who show that, in the presence of a structural break, a misspecified VECM can provide misleading forecasts.

II. Related literature

The correct way to model a pair of non-stationary cointegrated time series is with a VECM (Granger, 1981; Engle and Granger, 1987). Given the causal link from the stock market to unemployment it should be possible to predict the future history of the unemployment rate using its own past and the past history of the stock market. But in the presence of structural breaks, VECMs are not robust to shifts in the underlying equilibria. Models that are overdifferenced, and therefore misspecified, are known to outperform well specified models that have undergone a structural break (Clements and Hendry, 1988, 1999; Hendry, 2006; Castle, Fawcett and Hendry, 2010). This paper illustrates the result that overdifferencing improves forecasting ability in the context of the unemployment–stock market relationship, previously studied in Farmer (2012b).

I am not the first to investigate the connection between wealth and subsequent movements in economic activity. Lettau and Ludvigson (2004, 2011) provide a statistical model of consumption, wealth and labour earnings as non-stationary, but cointegrated time series, using the methods surveyed in Hendry (2004) and Hendry and Juselius (2000, 2001). I look instead at the relationship between the real value of the stock market and the unemployment rate.

The connection between the stock market 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 and Fitoussi *et al.* (2000), who 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. My explanation for persistent unemployment (Farmer, 2010a, 2012a,b, 2013a) is closer to the models of hysteresis described by Blanchard and Summers (1986, 1987) and Ball (1999) than the structuralist model of Phelps. Nevertheless, the theoretical foundation for persistent unemployment in Farmer (2010a, 2012a,b, 2013a) is very different from the one provided in their work.

III. Properties of the data

The data I use in this analysis include the S&P 500 from Shiller (2014) and the unemployment rate from the Bureau of Labour Statistics. The stock market data are quarterly

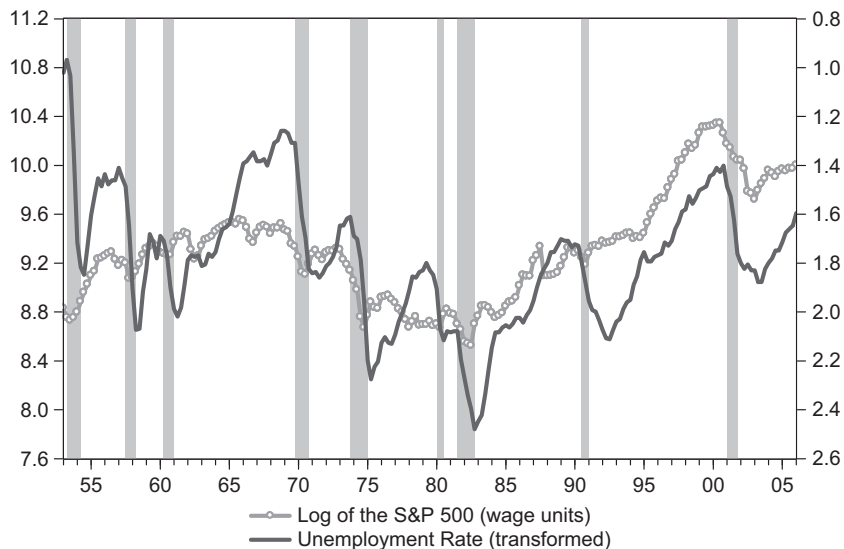


Figure 1. Data used in this study

averages of Shiller's monthly series and the unemployment rate is the quarterly average of the monthly rate. These data are graphed in Figure 1.

I have deflated the S&P 500 by a measure of the money wage, constructed by dividing compensation to employees from Table 1.12 of the National Income and Product Accounts by full time equivalent employees from Table 6.5 of the NIPA accounts, interpolated from the annual data. The series were further transformed as follows

$$p = \log \left(\frac{\text{S\&P 500}}{\text{Money Wage Series}} \right), \quad u = \log \left(\frac{100 \times \text{unemployment rate}}{100 - \text{unemployment rate}} \right),$$

to map them each into the real line. This latter transformation is important because I will argue that u and p , defined in this way, are non-stationary series that are unbounded above and below. All data are available on my website and the technique of deflating by the money wage is discussed in Farmer (2010a).

On Figure 1, the shaded grey areas are NBER recessions, the solid line is the unemployment rate and the line with circles represents the stock market. The unemployment rate is measured on the right axis on an inverted scale and the stock market is measured on the left axis.

Table 1 reports augmented Dickey Fuller tests for the null hypothesis that u or p have a unit root. Because there is strong evidence that many macro time series behave differently before and after the Volcker disinflation, I have split the sample in 1979q3, the date when Volcker came into office. Test statistics for the subsample from 1953q1 through 1979q3 appear in the first row, and statistics for the subsample from 1979q4 through 2011q1, appear in the second. In each case the null hypothesis of a unit root cannot be rejected at the 1% level and the smallest one-sided P -value is 29%. These results suggest that both time-series are integrated of order 1.

Table 2 presents results from two different tests developed by Johansen (1995) to test for cointegration between non-stationary series. The lag length for each test, equal to 3,

TABLE 1

Augmented Dickey-Fuller test: null is that u or p has a unit root

	u		p	
	Test statistic	P -value	Test statistic	P -value
1953q1–1979q3	–1.65	0.45	–1.38	0.59
1979q4–2011q1	–2.00	0.29	–1.44	0.56

TABLE 2

 P -values for Johansen tests: null is 0 or at most 1 CI vector

	0 CI vectors		At most 1 CI vector	
	Trace Test	Max-EV Test	Trace Test	Max-EV Test
1953q1–1979q3	0.04	0.03	0.30	0.29
1979q4–2011q1	0.02	0.03	0.15	0.15

was chosen to satisfy the Aikake Information criterion and the table presents P -values for the Johansen Trace Test and Maximum Eigenvalue Test over two different sub-periods for the null hypotheses of 0 cointegrating vectors and at most 1 cointegrating vector.

The first row of the table is for the first subsample and the second row is for the second. In all cases the null of no cointegrating vectors has a P -value of $< 5\%$. In contrast, the null of at most one cointegrating vector has P -value of 29% in the first subsample and 15% in the second. These results provide support for modelling the bivariate properties of unemployment and the stock market as a VECM with a single cointegrating vector.

IV. Granger causality tests

I have established that unemployment and the stock market are cointegrated. Farmer (2012b) claimed that the stock market *causes* the unemployment rate. To check the viability of that claim, Table 3 presents Granger causality tests for each of the two sub-periods. In each case, these tests were carried out using the correction for non-stationary time series developed by Toda and Yamamoto (1995), and implemented in Eviews with the method suggested by Giles (2011).

For both subperiods, the hypothesis that the stock market does not Granger cause unemployment has a P -value of 0. In words, there is overwhelming evidence that information contained in the stock market helps to forecast the unemployment rate one quarter later. In contrast, the hypothesis that unemployment does not help to predict the stock market has a P -value of 11% in the first subsample and 74% in the second.

I have established that the stock market Granger causes the unemployment rate conditional on an information set that contains lagged values of these two series. Perhaps, however, the stock market is reacting to some other set of information that forecasts a recession. To check for that possibility, I ran Granger causality tests on expanded data sets that included real gdp, real investment spending, the three month treasury bill rate, the CPI inflation rate and the spread of BAA bonds over ten year treasuries. With the exception of the BAA spread, none of these variables altered the conclusion that the stock market

TABLE 3
Granger causality tests

<i>Dependent variable: u</i>				<i>Dependent variable: p</i>			
<i>Excluded</i>	<i>Chi-squared</i>	<i>df</i>	<i>P-value</i>	<i>Excluded</i>	<i>Chi-squared</i>	<i>df</i>	<i>P-value</i>
Subsample: 1953q1–1979q3							
<i>p</i>	14.46	3	0.00	<i>u</i>	9.04	3	0.11
Subsample: 1979q4–2011q1							
<i>p</i>	22.64	3	0.00	<i>u</i>	6.69	3	0.74

TABLE 4
Granger causality tests – subsample: 1953q1–1979q3

<i>Dependent variable: u</i>				<i>Dependent variable: p</i>			
<i>Excluded</i>	<i>Chi-squared</i>	<i>df</i>	<i>P-value</i>	<i>Excluded</i>	<i>Chi-squared</i>	<i>df</i>	<i>P-value</i>
<i>sp</i>	17.31	3	0.00	<i>sp</i>	11.07	3	0.01
<i>p</i>	2.95	3	0.40	<i>u</i>	6.50	3	0.09
<i>All</i>	38.21	6	0.00	<i>All</i>	21.96	6	0.00

<i>Dependent variable: sp</i>			
<i>Excluded</i>	<i>Chi-squared</i>	<i>df</i>	<i>P-value</i>
<i>u</i>	1.36	3	0.71
<i>p</i>	16.35	3	0.00
<i>All</i>	19.09	6	0.00

Granger causes the unemployment rate. The result of adding the BAA spread to the mix is reported in Tables 4 and 5.

Table 4 presents the results of Granger causality tests for a three variable autoregression including the stock market, the unemployment rate and the BAA spread, *sp*, for the period 1953q1–1979q3, and Table 5 presents Granger causality tests for the second subperiod from 1979q4 through 2011q1. For each table the top left panel reports causality tests for the unemployment rate, the top right is for the stock market and the bottom left, for the BAA spread.

Table 4 shows that, when the BAA spread is added to the information set in the first subperiod, the explanatory power of the stock market disappears and its role is taken by the BAA spread. However, the spread itself is Granger caused by the stock market.

Table 5 shows that, in the second subperiod, the individual probabilities that either the stock market or the spread do not Granger cause the unemployment rate are both zero. In both subperiods, the *P*-value for the joint hypothesis that neither the spread nor the stock market Granger causes the unemployment rate is zero. In other words; there is overwhelming evidence that there is information in the financial markets that helps predict recessions.

V. Causality and control: Animal spirits or fundamentals?

The fact that information from the financial markets Granger causes the unemployment

TABLE 5

Granger causality tests – subsample: 1979q4–2011q1

<i>Dependent variable: u</i>				<i>Dependent variable: p</i>			
<i>Excluded</i>	<i>Chi-squared</i>	<i>df</i>	<i>P-value</i>	<i>Excluded</i>	<i>Chi-squared</i>	<i>df</i>	<i>P-value</i>
<i>sp</i>	33.97	3	0.00	<i>sp</i>	31.47	3	0.00
<i>p</i>	26.75	3	0.00	<i>u</i>	1.55	3	0.67
<i>All</i>	83.35	6	0.00	<i>All</i>	33.38	6	0.00

<i>Dependent variable: sp</i>			
<i>Excluded</i>	<i>Chi-squared</i>	<i>df</i>	<i>P-value</i>
<i>u</i>	2.17	3	0.54
<i>p</i>	1.19	3	0.75
<i>All</i>	4.04	6	0.67

rate does not *necessarily* imply that, if we could control the asset markets by government intervention, we would be able to control the unemployment rate. That point is made clearly by Granger (1980) who distinguished between causation and control. To make the case for control, one needs an economic model that suggests a plausible mechanism to explain the causal chain.

Consider, for example, the following two explanations for the deep recession that followed the collapse of Lehman Brothers in September of 2008.

In the first explanation, market participants received a signal in the autumn of 2008, that a fundamental event was about to occur that would depress the value of stock market earnings and increase the value of unemployment for an extended period of time. That news also increased the likelihood of corporate bankruptcies and increased the cost of credit for low-quality corporate borrowers. An example of such an event would be a court ruling that increased union bargaining power and was perceived to lead to significant future labour market disruptions and loss of output. I will call this the fundamental view of the market.

In the second explanation, there was an increase in the perceived risk of running a business. Although nothing fundamental had changed in the economy, market participants anticipated that a recession was on the horizon. This fear spread to the stock market and participants sold shares because they believed that future markets prices would be lower. As a consequence of the perceived increased risk in the financial markets, the face value of paper assets dropped and households curtailed their spending causing firms to lay-off workers. The reduced level of economic activity resulted in a self-fulfilling drop in the value of earnings per share. I will call this the animal spirits view of the market.¹

According to the fundamental view of the market, an attempt to restore confidence by treasury or central bank intervention will be self-defeating. If government buys shares or low-quality corporate bonds, paid for by borrowing, they will lose money in the long-run because asset market intervention cannot effectively counteract the fundamental cause of the market crash. According to the animal spirits view of the market, restoration of

¹ I use the case of a court ruling as the fundamental event that triggers a crisis purely as an illustration. In the context of the 2008 crisis it is difficult to find a plausible candidate for a fundamental event of any kind, and, for that reason, I personally find the animal spirits explanation more plausible.

confidence through asset market purchases is an effective way to prevent a market crash from causing a recession.² These two views cannot be distinguished *ex ante* although they clearly have different policy implications.

VI. Forecasting with structural breaks

In this section, I turn my attention to Rosnick's finding that a univariate ARMA process is a better predictor of the future unemployment rate than a regression that includes lagged values of the stock market. To establish my claim that the stock market *does* help to predict the unemployment rate, I estimated three different models on data from 1953q1 through 1979q3, and I compared their forecast performance for the sample period 1979q4 through 2011q1. Model 1 is the VECM 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.³

$$\begin{bmatrix} \Delta u_t \\ \Delta p_t \end{bmatrix} = \begin{bmatrix} 0.6 & -0.27 \\ (0.06) & (0.10) \\ 0.13 & 0.42 \\ (0.06) & (0.09) \end{bmatrix} \begin{bmatrix} \Delta u_{t-1} \\ \Delta p_{t-1} \end{bmatrix} + \begin{bmatrix} \alpha \\ -0.1 \\ (0.02) \\ 0.01 \\ (0.19) \end{bmatrix} \begin{bmatrix} \beta^T \\ 1 \\ (0.23) \\ -7 \\ (2.07) \end{bmatrix} \begin{bmatrix} u_{t-1} \\ p_{t-1} \\ c \end{bmatrix}, \quad (1)$$

$$[\Delta u_t] = \begin{bmatrix} 0.6 \\ (0.07) \end{bmatrix} [\Delta u_{t-1}], \quad (2)$$

$$\begin{bmatrix} \Delta u_t \\ \Delta p_t \end{bmatrix} = \begin{bmatrix} 0.49 & -0.36 \\ (0.06) & (0.09) \\ 0.13 & 0.40 \\ (0.06) & (0.08) \end{bmatrix} \begin{bmatrix} \Delta u_{t-1} \\ \Delta p_{t-1} \end{bmatrix}. \quad (3)$$

Parameter estimates for the VECM the univariate and bivariate models are reported above as equations (1), (2) and (3).⁴ 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. 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.

²The case for central bank intervention in the asset markets is made in Farmer (2010b, 2013b, 2014).

³All three models were estimated in Eviews. The data set and the Matlab code used to construct Figures 2–7 are available on my website. Parameter estimates for model 3 use robust regression (Huber, 1973) to correct for outliers. Standard errors are in parentheses.

⁴The estimates reported in Farmer (2012b, p. 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.

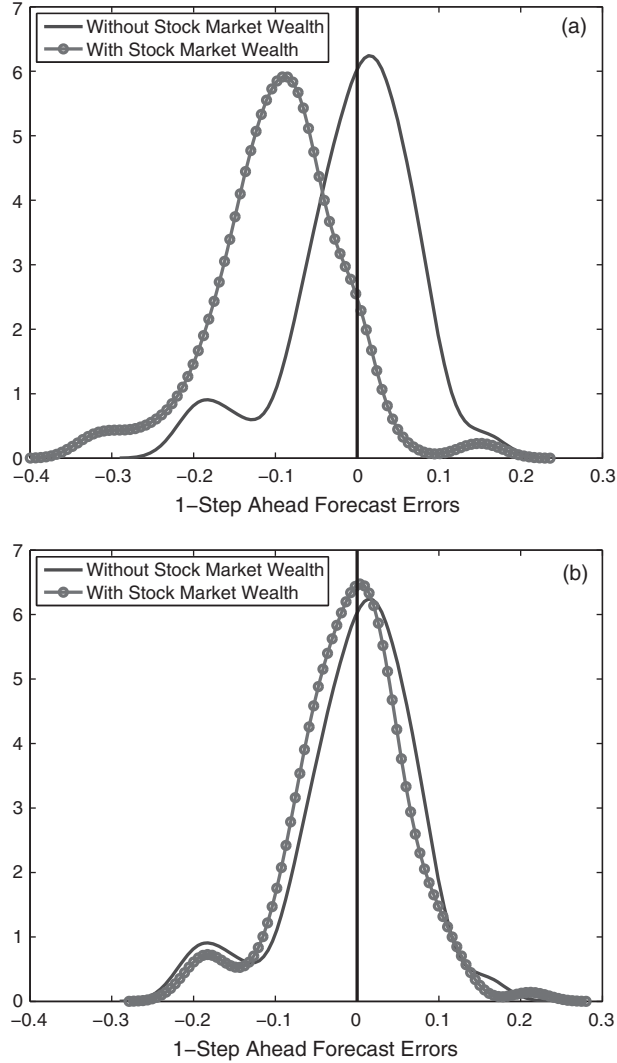


Figure 2. 1-step ahead forecast errors for the three models

VII. The three models compared

In Figure 2, I report smoothed histograms of the 1-step ahead forecast errors of the unemployment rate for the period 1979q4–2011q1. The upper panel compares the VECM with the univariate model;⁵ the lower 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

⁵This reproduces Figure 3 from Rosnick (2013).

mean, indicating bias in the prediction, whereas that of the bivariate model is centred 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.

Figure 3 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 upper panel shows that this ratio is > 1 at all horizons, indicating that the univariate model beats the VECM. The lower panel shows that this result is reversed for the bivariate model which has a MSE ratio < 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.

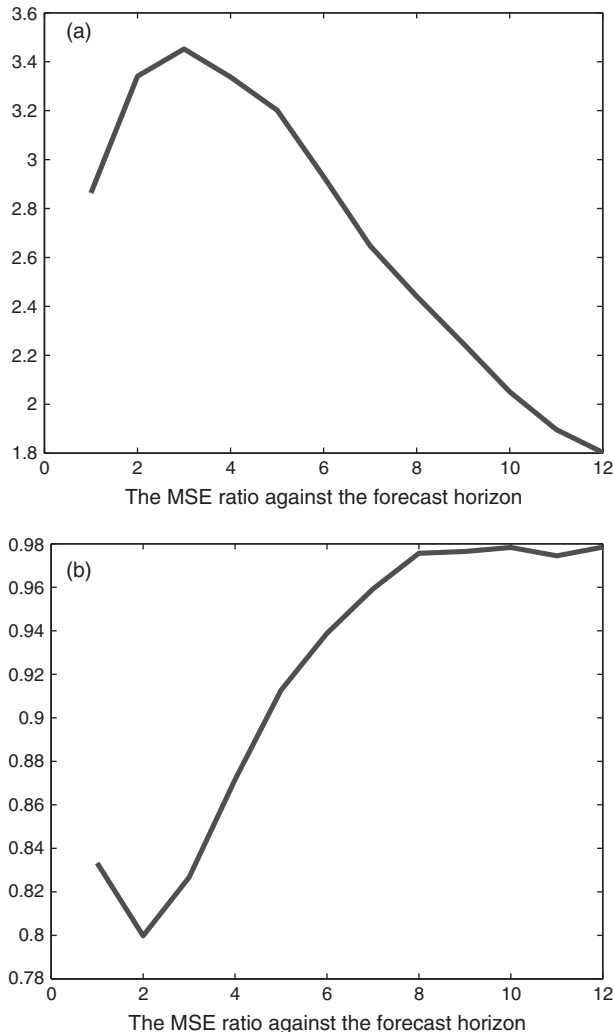


Figure 3. MSE prediction errors at different forecast horizons

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. The following section shows that this is not the case.

VIII. 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?

Figures 5–7 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 line with circles is the forecast from the univariate model and the line with crosses is the forecast from the comparison model. In the upper panel, the comparison model is the VECM; in the lower 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 12 quarters ahead.⁶

Figure 5 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 upper panel of this figure shows that the VECM does a poor job and it is apparent from this graph, that the VECM is seriously misspecified. It predicts a large drop in the unemployment rate in 2008 in contrast to the path of unemployment that actually occurred. The lower panel of Figure 5 compares the univariate model with the bivariate VAR estimated in first differences. This figure shows that adding the first difference of the stock market to a prediction made in the fourth quarter of 2007 makes little or no difference to the univariate

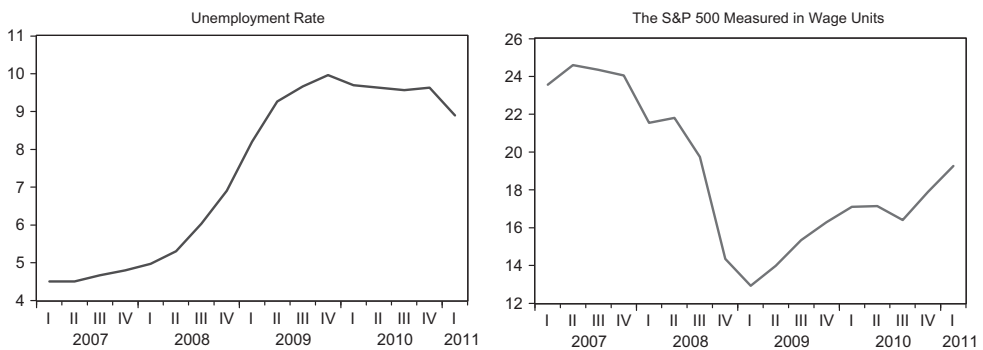


Figure 4. Unemployment and the Stock Market

⁶The upper panel contains some of the same information as Figure 4 in Rosnick (2013).

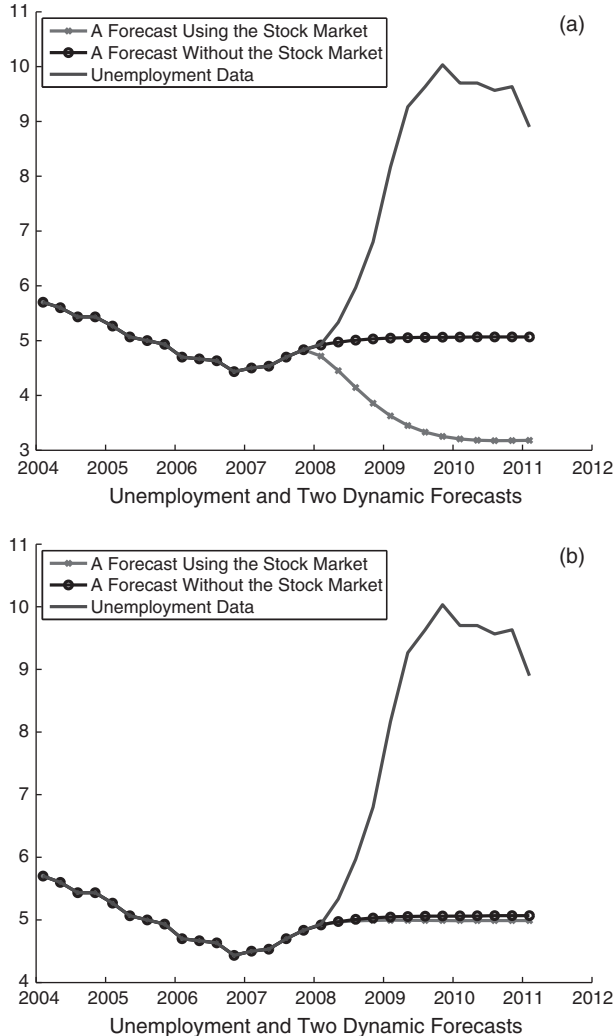


Figure 5. Forecasts from 3 models in the autumn of 2007

forecast. That is unsurprising since, at this date, the stock market had not yet begun its spectacular decline.⁷

The failure of the VECM as a forecasting device is what one would expect if there had been a shift in the cointegrating relationship between the stock market and the unemployment rate. Castle *et al.* (2010) suggest a number of responses to the problem of shifts in a cointegrating equation including updating, intercept corrections and differencing. Taking first differences of the VECM leads to a bivariate VAR in unemployment and the stock market that ignores the estimated cointegrating equation. As I will show below, the bivariate model has significantly better forecasting ability than either the misspecified

⁷There 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.

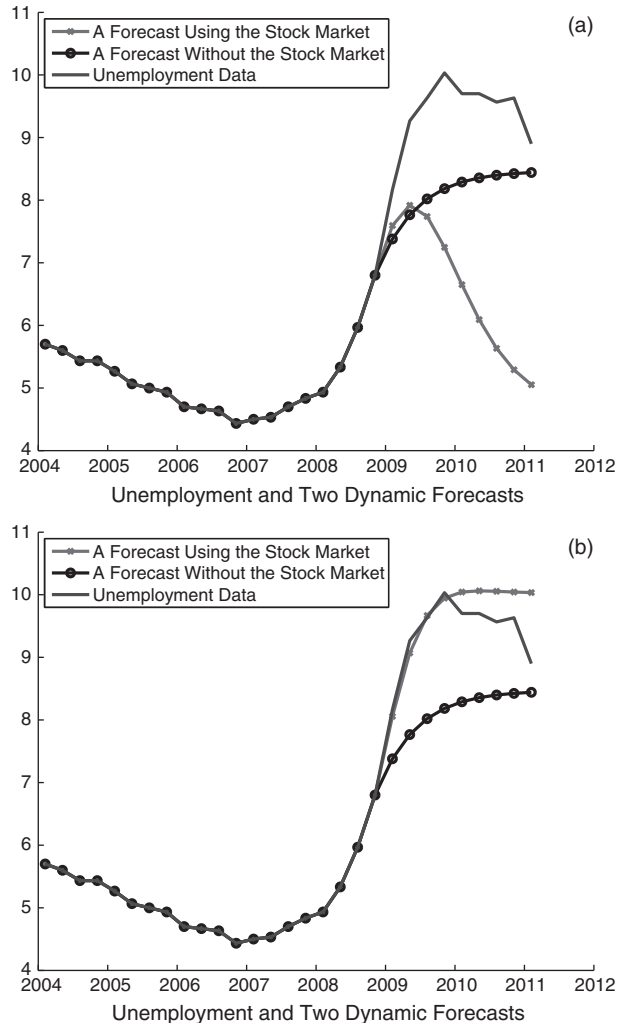


Figure 6. Forecasts from 3 models in the autumn of 2008

VECM or the univariate VAR for forecasting periods that include information from the stock market crash in the autumn of 2008.

Figure 6 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 4, there had been a large drop in the S&P 500.

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

The lower panel of Figure 6 shows that the bivariate model correctly predicts the magnitude of the Great Recession four quarters ahead, but overshoots in the fifth quarter and beyond. In contrast, the univariate model misses the depth of the increase in the unemployment rate by two full percentage points.

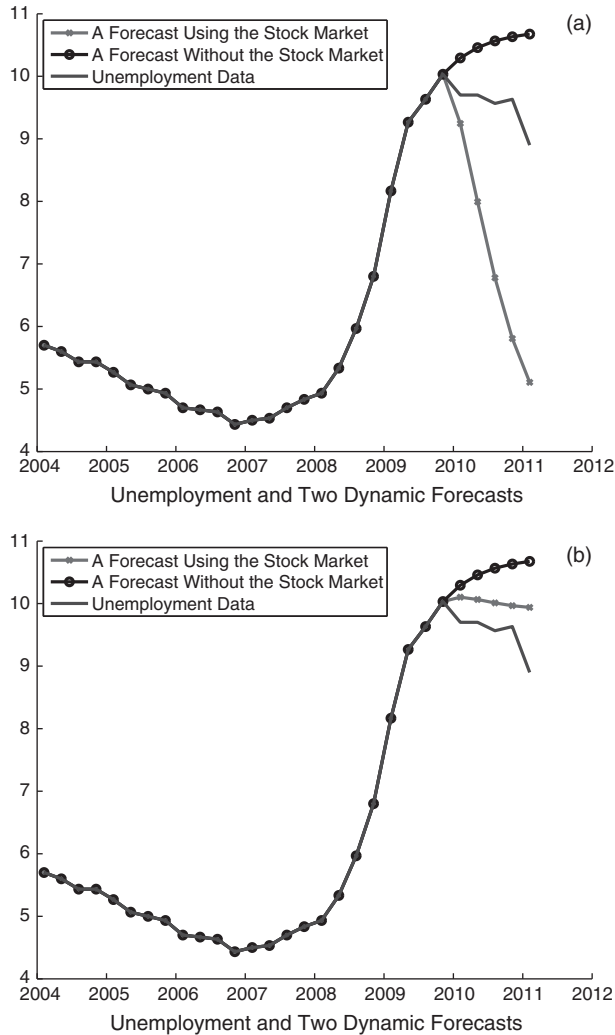


Figure 7. Forecasts from 3 models in the autumn of 2009

Figure 7 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 lower panel, correctly predicts an improvement in the labour market. In contrast, the univariate model predicts that the unemployment situation will continue to deteriorate. The upper panel of Figure 7 shows that, once again, the VECM performs poorly as a forecasting tool.

To understand why the VECM performs poorly, Table 6 presents estimates of the cointegrating vector for the two subsamples.

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 -5.3 . Although these estimates are within two standard error

TABLE 6
Estimates of the cointegrating vector

First sub-sample				Second sub-sample			
1953q1–1979q3	<i>u</i>	<i>p</i>	<i>c</i>	1979q4–2011q1	<i>u</i>	<i>p</i>	<i>c</i>
Coefficient	1	0.6	−7.4	Coefficient	1	0.36	−5.3
SE		0.25	2.3	SE		0.09	0.87
<i>t</i> -statistic		(2.47)	(−3.21)	<i>t</i> -statistic		(3.86)	(−6.06)




Figure 8. The cointegrating vectors

bounds of each other, the poor out-of-sample fit suggests that the differences are statistically important.

Figure 8 plots the cointegrating vector estimated over the period from 1953q1 to 1979q3 as the solid line. The line with circles is the cointegrating vector estimated over the subsample from 1979q4 through 2011q1. This second cointegrating vector tracks the first until 1985q1. After that period it begins to diverge and this divergence appears to be growing over time. It is this divergence that is the cause of the failure of VECM to generate accurate forecasts.

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 2–7, changes in the stock market have a large and statistically significant impact on changes in the future unemployment rate.

IX. Simulating the great recession

Where does this leave the claim that the stock market crash of 2008 caused the Great Recession? Figure 9 presents the result of simulating the effect of a one quarter shock of




Figure 9. Simulating a stock market crash

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.⁸ I assume that there are no further shocks after the first quarter drop in the value of the S&P.

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 US 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. One might be concerned that the model predicts a larger increase in the unemployment rate than we observed. However, the simulation depicted in Figure 9 ignores the impact of subsequent shocks to the stock market. In the data, these shocks were large and unprecedented

In November of 2008 the Federal Reserve more than doubled the monetary base from eight hundred billion dollars in October to more than two trillion dollars in December: And over the course of 2009 the Fed purchased eight hundred billion dollars worth of mortgage backed securities.⁹ According to the animal spirits explanation of the recession (Farmer, 2010a, 2012a,b, 2013a), these Federal Reserve interventions in the asset markets were a significant factor in engineering the stock market recovery.

The animal spirits theory provides a causal chain that connects movements in the stock market with subsequent changes in the unemployment rate. If this theory is correct, the path of unemployment depicted in Figure 8 is an accurate forecast of what would have occurred in the absence of Federal Reserve intervention. These results support the claim, in the title of this paper, that the stock market crash of 2008 really did *cause* the Great Recession.

⁸ 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.

⁹ See Farmer (2013b, Chart 6) for a graph that illustrates the coincidence in timing of the recovery in the stock market with the purchase by the Fed of mortgage backed securities.

X. 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 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 (Clements and Hendry, 1988, 1999; Hendry, 2006).

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, 2010a, 2012a, 2013a), the stock market matters for the unemployment rate: and it matters a lot.

Final Manuscript Received: January 2015

References

- Ball, L. (1999). 'Aggregate demand and long-run unemployment'. *Brookings Papers on Economic Activity*, Vol. 2, pp. 189–236.
- Blanchard, O. J. and Summers, L. H. (1986). 'Hysteresis and the european unemployment problem', in Fischer S. (ed.), *NBER Macroeconomics Annual*, Vol. 1, Boston, MA: National Bureau of Economic Research, pp. 15–90.
- Blanchard, O. J. and Summers, L. H. (1987). 'Hysteresis in unemployment', *European Economic Review*, Vol. 31, pp. 288–295.
- Castle, J. L., Fawcett, N. W. and Hendry, D. F. (2010). 'Forecasting with equilibrium error corection models during structural breaks', *Journal of Econometrics*, Vol. 158, pp. 25–36.
- Clements, M. P. and Hendry, D. F. (1988). *Forecasting Economic Time Series*. Cambridge University Press, Cambridge.
- Clements, M. P. and Hendry, D. F. (1999). *Forecasting Non-stationary Economic Time Series*, Cambridge, MA: MIT Press.
- Engle, R. J. and Granger, C. W. J. (1987). 'Cointegration and error correction: representation, estimation and testing', *Econometrica*, Vol. 55, pp. 251–276.
- Farmer, R. E. A. (2010a, March). *Expectations, Employment and Prices*. New York: Oxford University Press.
- Farmer, R. E. A. (2010b). 'How to reduce unemployment: a new policy proposal', *Journal of Monetary Economics: Carnegie Rochester Conference Issue*, Vol. 57, pp. 557–572.
- Farmer, R. E. A. (2012a, March). 'Confidence, crashes and animal spirits', *Economic Journal*, Vol. 122, pp. 155–172.

- Farmer, R. E. A. (2012b). 'The stock market crash of 2008 caused the great recession: theory and evidence', *Journal of Economic Dynamics and Control*, Vol. 36, pp. 697–707.
- Farmer, R. E. A. (2013a, May). 'Animal spirits, financial crises and persistent unemployment', *Economic Journal*, Vol. 123, pp. 317–340.
- Farmer, R. E. A. (2013b). 'Qualitative easing: a new tool for the stabilisation of financial markets', *Bank of England Quarterly Bulletin*, Vol. Q4, pp. 405–413.
- Farmer, R. E. A. (2014). 'Financial stability and the role of the financial policy committee', *The Manchester School*, Vol. 82, pp. 35–43.
- Fitoussi, J.-P., Jestaz, D., Phelps, E. S. and Zoega, G. (2000). 'Roots of the recent recoveries: labor reforms or private sector forces', *Brookings Papers on Economic Activity*, Vol. 1, pp. 237–312.
- Giles, D. (2011). Granger causality. Econometrics Beat, Blog Post. Available from: <http://davegiles.blogspot.ca/2011/04/testing-for-granger-causality.html> [last accessed 13 February 2015].
- Granger, C. W. (1969). 'Investigating causal relations by econometric models and cross-spectral methods', *Econometrica*, Vol. 37, pp. 424–438.
- Granger, C. W. (1980). 'Testing for causality: a personal viewpoint', *Journal of Economic Dynamics and Control*, Vol. 2, pp. 329–352.
- Granger, C. W. J. (1981). 'Some properties of time series data and their use in econometric model specification', *Econometrica*, Vol. 16, pp. 121–130.
- Hendry, D. F. (2004). 'The nobel memorial prize for clive w.j. granger', *Scandinavian Journal of Economics*, Vol. 106, pp. 187–213.
- Hendry, D. F. (2006). 'Robustifying forecasts from equilibrium-correction systems', *Journal of Econometrics*, Vol. 135, pp. 399–426.
- Hendry, D. H. and Juselius, K. (2000). 'Explaining cointegration analysis: part 1', *Energy Journal*, Vol. 21, pp. 1–42.
- Hendry, D. H. and Juselius, K. (2001). 'Explaining cointegration analysis: part ii', *Energy Journal*, Vol. 22, pp. 75–120.
- Hoon, T. and Phelps, E. S. (1992). 'Macroeconomic shocks in a dynamized model of the natural rate of unemployment', *American Economic Review*, Vol. 82, pp. 889–900.
- Huber, P. J. (1973). 'Robust regression: asymptotics, conjectures and monte carlo', *Annals of Statistics*, Vol. 1, pp. 799–821.
- Johansen, S. (1995). *Likelihood Based Inference in Cointegrated VAR Models*. Oxford: Oxford University Press.
- Lettau, M. and Ludvigson, S. C. (2004, March). 'Understanding trend and cycle in asset values: reevaluating the wealth effect on consumption', *American Economic Review*, Vol. 94, pp. 276–299.
- Lettau, M. and Ludvigson S. C. (2011). *Shocks and Crashes*, NBER Working Paper No. 16996.
- Phelps, E. S. (1994). *Structural Slumps*. Cambridge, MA: Harvard University Press.
- Phelps, E. S. (1999). 'Behind this structural boom: the role of asset valuations', *American Economic Review Papers and Proceedings*, Vol. 89, pp. 63–68.
- Rosnick, D. (2013, April). Farmer's folly: Bringing the nikkei godzilla to America. Available from: <http://www.cepr.net/index.php/blogs/cepr-blog/farmers-folly-bringing-the-nikkei-godzilla-to-america> [last accessed 13 February 2015].
- Shiller, R. (2014). Onlinedata source for 'irrational exuberance'. Available from: <http://www.econ.yale.edu/shiller/data.htm> [last accessed 13 February 2015].
- Toda, H. Y. and Yamamoto, T. (1995). 'Statistical inference in vector autoregressions with possibly integrated processes', *Journal of Econometrics*, Vol. 66, pp. 225–250.