Do Asset Price Drops Foreshadow Recessions?

John C. Bluedorn, Jörg Decressin, Marco E. Terrones*

May 2015

(Forthcoming International Journal of Forecasting)

Abstract

This paper examines the usefulness of asset prices in predicting the beginning of recessions in the G-7 countries. It finds that equity/house price drops have a substantial marginal effect on the likelihood of a new recession. Increased market uncertainty, a second-moment variable associated with equity price changes, is also a useful predictor of new recessions in these countries. These findings are robust to the inclusion of the term spread and oil prices. The new recession forecasting performance of our baseline model is superior to that of a similar model estimated over all recession and expansion periods, highlighting a difference between the probabilities of a new recession versus a continuing recession.

JEL Classification Numbers: E32; E37; G17

Keywords: Business cycles; Macroeconomic forecasting; Financial markets; Uncertainty; Oil Prices; Binary dependent variable models.

(International Monetary Fund. Revised version of the IMF Working Paper 13/203).

I. Introduction

Many observers have noted that asset price drops are often followed by a recession. Historical examples of this regularity include the 1929 stock market crash and the Great Depression; the sharp decline in asset values in 1973-74 and the ensuing economic downturn in the United States and United Kingdom; the early 1990s' asset price collapse and recession in Japan; the stock market downturn in the early 2000's and the 2001 recession in the United States; and the 2008 global crash in asset prices and the Great Recession. In all these episodes, equity and/or house prices fell sharply prior to or coincident with an economic downturn.

Other observers, however, have argued that asset price declines do not always precede or coincide with economic contractions. The sharp decline in the stock market in 1962, for instance, did little to unsettle the economic recovery process in the United States. Likewise, the stock market crash of October 1987 did not significantly affect economic activity in the United States, despite predictions of a severe recession in 1988. The August 2011 stock market collapse in the United States and Asia was also not followed by a recession in these economies. These observers assert that asset prices (and equity prices, in particular) are poor indicators of forthcoming recessions because they are inherently volatile. Samuelson's (1966) famous epigram that "the stock market has forecast nine of the last five recessions" cleverly summarizes this view.

In this study, we examine whether asset price drops show any link to the starts of recessions in the G-7 countries. Specifically, we assess whether equity and house price drops are reliable predictors of new recessions. As large drops of equity prices and rises in financial market uncertainty are often associated with major economic and political shocks, we also evaluate to what extent implied stock market volatility, a proxy for uncertainty, predicts new recessions.

We explicitly exclude periods where the economy is already in a recession from the estimation sample. This is a key departure from the previous literature, which has tended to pool information across both expansions and recessions, opting to estimate the probability of being in a recession at any given point in time. An important problem with this approach is that it can give a false impression of success in predicting new recessions. In most cases, these studies are reporting the probability of continuing in recession, conditional on the economy already being in recession.

The results from our baseline logistic regression model indicate that asset prices are significantly related to the beginning of a new recession in the G-7 economies. There is also evidence that the relationship between asset prices and the recession starts is highly asymmetric—the average marginal effect in the probability of a new recession of a large decline in equity/house prices is much larger in absolute value than that of an equivalent increase. Market uncertainty is found to be a useful predictor of new recessions, highlighting its possible role in the business cycle. The term spread and oil price growth, proxying for bond and commodity market conditions, are also included in the baseline model. Confirming the literature's findings, the term spread does help predict new recessions. However, we show that equity price movements have better in-sample forecasting performance than the term spread. Lastly, oil price changes do not appear to be particularly useful in predicting new recessions. These new results suggest that the early pessimistic assessment on asset prices' ability to help forecast new recessions needs to be revised.

The remainder of this paper is structured as follows. In section 2, we briefly motivate our model specification and describe our data and econometric methodology. In section 3, we report the main empirical findings of this paper. We conclude in section 4 with a brief summary of our main results and discussion of future research.

II. Model Motivation, Data, and Econometric Methodology

Why might asset price changes be useful predictors of new recessions? Two main reasons have been put forward in the literature. First, asset price declines could actually cause a downturn in economic activity by negatively affecting the net wealth, balance sheets, and confidence of households and firms.[1] Second, an asset price drop could signal a weakening of the economic outlook to the extent that asset prices are forward-looking. These two elements led many observers to treat asset prices, and equity prices, in particular, as harbingers of the business cycle. However, early empirical studies found that asset prices were only of limited use to forecast economic downturns.[2]

In addition, a new body of literature postulates that market uncertainty can be an important driver of the business cycle (Bloom, 2009). When faced with high uncertainty, firms reduce their investment demand and delay their projects as they gather new information, as investment can be costly to reverse (Bernanke, 1983, and Dixit and Pindyck, 1994). Market uncertainty, typically proxied by stock market volatility, may also respond to major economic and political shocks.

Beyond these asset price and financial market indicators, there is a well established literature that has found that the term spread and oil price changes can help predict recessions in advanced economies. Term spread narrowings or inversions have been found to be positively

[1] Recent studies have found that asset price movements that affect households' net wealth are associated with significant changes in their spending (Carroll, Otsuka, and Slacalek, 2011; Case, Quigley, and Shiller, 2013).

[2] Stock and Watson (2003) report that equity prices are generally poor predictors of output growth. In contrast, virtually no study has examined the predictive content of housing prices for economic growth and recessions, partly reflecting data limitations.

associated with recessions.[3] Other work has shown that sharp rises of or persistently high oil prices have often preceded recessions.[4]

We investigate the predictive ability of these five variables for new recessions at a quarterly frequency with a logistic regression model. The dependent variable is a binary variable that takes on the value of 1 if a country has reached its cyclical peak in quarter *t*, which indicates the end of an expansion, and 0 otherwise. The cyclical peaks and troughs for the G-7 countries are obtained from Claessens, Kose, and Terrones (2012). These cyclical turning points are identified using the algorithm introduced by Harding and Pagan (2002), which generalizes the algorithm developed by Bry and Boschan (1971) for the United States.[5] Cyclical peaks or recessions starts are rare events, with just 44 occurring in the G-7 countries from 1970:Q1–2011:Q4. This means that the probability of a recession starting in the next quarter conditional on the economy being in expansion is about 4.7 percent. By contrast, these economies spent some 18 percent of their time in recession over the same period.

Equity price growth is the log difference of the country's main stock market index, where each listing's equity price is weighted with the market value of outstanding shares. The house price measure is an index of house or land prices depending on the country. These variables have been converted into real terms using the corresponding national consumer price indices (CPI). The implied or realized volatility of the S&P 500 index comes from Bloom (2009), spliced with the Chicago Board of Options' VXO index from 2006 through 2011. Lastly, the term spread is calculated as the difference between the 10-year government bond rate and the 3-month treasury bill rate (or equivalent). The underlying sources include the IMF's International Finance Statistics, OECD, BIS, Haver Analytics, Bloomberg, Global Financial Database, and various country-specific sources. To measure movements in oil prices, we employ the U.S. dollar average petroleum spot prices of West Texas Intermediate, U.K. Brent, and Dubai Fateh crude (equally weighted). This data comes from the IMF's Primary Commodity Prices database and has been converted into constant dollars using the U.S. CPI.

---

[3] For instance, Estrella and Mishkin (1998) examine the performance of the term spread, a proxy for the stance of monetary policy, as a predictor of a binary recession indicator for the United States. They find that a narrowing of the term spread helps predict recessions in the U.S. for horizons further than one quarter out. This finding has been confirmed for other economies. For instance, Duarte, Ventis, and Paya (2005) report that EMU and U.S. yield spreads are associated with EMU recessions. For a more recent contribution, see Christiansen (2013).

[4] Hamilton (2011b) found that 10 out of the last 11 recessions that the U.S. has experienced were associated with an increase in oil prices (the exception being the mild recession of 1960-1961). This result seems to hold for other economies. Engemann, Kliesen, and Owyang (2011), for instance, report that oil price shocks, in addition to the term spread, help predict recessions in seven advanced economies.

[5] This algorithm first searches for local maxima and minima of the log-level of output (y). It then makes sure that the sequence of identified maxima and minima alternate between peaks and troughs. Furthermore, the identified sequence of peaks and troughs must satisfy censuring rules which require a minimal duration for each cyclical phase and cycle.

To evaluate the predictive performance of the logit models, we utilize the receiver operating characteristic curve (ROC), which charts the true positive rate (TP(π)), that is the proportion of new recessions correctly classified as new recessions, versus the false positive rate (FP(π)), which is the proportion of ongoing expansions incorrectly classified as new recessions, for all possible classification thresholds π.[6] A logit model that is uninformative or has no discriminative ability will generate an ROC curve that coincides with the diagonal 45-degree line. The area under the ROC curve (AUC) or *c*-statistic is in this case equal to ½. In contrast, a perfectly informative logit model will generate an ROC that coincides with the left hand and top axis of the unit square, generating an AUC of 1. In general, one can use the AUC statistic as a global measure of the forecasting performance of different binary classification models—with the most accurate model showing the largest AUC and the least accurate showing an AUC close to ½.

To make the classification using the model operational, some cutoff threshold probability needs to be selected from the large set of possible thresholds characterized by the ROC curve. Because of its simplicity, we use the Youden index and its associated cutoff threshold π^* (Youden, 1950; Perkins and Schisterman, 2006). Youden's index is defined to be $J = \max\left\{TP\left(\pi^*\right) - FP\left(\pi^*\right)\right\}$, where π^* is then the cutoff threshold that maximizes the capability of the model to correctly discriminate between positives and negatives; false positives and negatives are equally weighted in the implicit loss function.

In addition to the ROC, we also utilize the Brier score (BS) to further assess the predictive performance of the models. The Brier score is just the sample mean square error between the outcome indicator and the associated probability forecast. A perfect forecaster would have a Brier score of 0, while the worst forecaster would have a Brier score of 1. Thus, lower Brier scores are preferred to higher. In some contexts, the Brier skill score (BSS), which measures the relative score over the naïve forecast, is utilized (defined as $1 - BS_1/BS_0$, where BS_0 is the Brier score for the reference model, while BS_1 is the score for the model of interest). In this case, BSS=1 for a perfect forecaster and BSS < 0 for a forecaster worse than naïve.

III. Empirical Results

We now turn to the empirical analysis of the relationship between asset prices and the starts of recessions in the G-7 countries. First, we present our logistic regression results for the start of recessions, describing the estimated effects of the regressors of interest and the in-sample forecasting ability of the model. Second, we compare the performance of our baseline model, estimated over expansions and recession starts, with one estimated over all expansion and recession periods, as is often done in the literature. We conclude the section with a comparison of the out-of-sample forcasting performances of the two models.

---

[6] The ROC curve is a graphical method that has utilized in many fields including medicine, biomedicine, psychiatry, manufacture production, and more recently economics. See for instance, Zou, O'Malley, and Mauri (2007), Berge and Jordà (2011), and Schularick and Taylor (2012).

A. Baseline Model

Table 1 shows the logistic regression results for predicting a new recession in the G-7 countries under various model specifications involving the four financial variables discussed before, plus oil prices. All models include country-fixed effects and quarterly dummies (to account for any residual seasonality). Columns (1)-(5) show the logistic regression coefficients for models based on each of these variables, taken one one-at-a-time. As can be seen, on its own, real equity price growth is highly significant and has the expected sign (negative, such that equity price rises reduce the chances of a recession start), with an AUC statistic of 0.79 (well above the 0.5 AUC of a coin toss).[7] The point estimate indicates that a one percentage point drop in equity prices increase the odds ratio for a new recession by about 12 percent.[8]

House price growth also appears to reduce the chances of a new recession (column 2), albeit not statistically significantly. Log implied S&P volatility shows a large positive and significant relationship to the onset of recessions (column 3). In terms of predictive fit, the model with implied S&P volatility with an AUC statistic of 0.76 (the second highest AUC among the univariate models) and the model with house prices with an AUC of 0.71 are both significantly higher than 0.5.

The term spread (column 4) has the expected negative and significant coefficient, indicating that spread inversions raise the estimated chance of a new recession, as reported in previous work. By contrast, real oil price growth exhibits only a small, insignificant positive relationship with recession starts (column 5).

When real equity and house price growth are jointly included (column 6), the coefficients remain roughly the same size, sign and significance as they are in the single explanatory variable models. However, when the log implied S&P volatility is also included (column 7), real house price growth becomes statistically significant. Moreover, its coefficient becomes slightly larger in size. There appears to be additional conditioning information in log implied S&P volatility that makes the estimate of the coefficient on house prices more precise. The coefficient on log implied S&P volatility drops by 60 percent compared to the model where it enters alone, but remains statistically significant.

Introducing the term spread does not markedly change the coefficients on equity prices, house prices, and implied S&P volatility (column 8). However, the coefficient on the term spread shrinks slightly in magnitude, compared to column 3. Our final baseline model, which includes real oil price growth, is shown in column 9. The estimated coefficients and AUC

---

[7] Some may wonder whether 0.79 is a high AUC statistic. This value exceeds all the AUC values reported by Jordà, Schularick, and Taylor (2011) in their analysis of financial crisis prediction.

[8] Recall that the odds ratio is defined as $P/(1-P)$, where P is the probability of a new recession. In the logistic regression case, the logarithm of the odds ratio is conveniently linear in the estimated coefficients.

statistic are very similar to those for the model shown in column 8, indicating that the inclusion of oil price growth adds basically no additional information.[9, 10]

The average marginal effects of changes in the different covariates on the predicted probability of a new recession are reported in the last column of the table. For instance, a 1 percentage point drop in equity or house price growth raises the probability of a new recession by about 0.4 percentage points, while a similarly sized drop in the term spread raises the probability by about 1.2 percentage points. Of course, these are only average marginal effects. Since the model is nonlinear, the actual impact on the predicted probability of a change in an explanatory variable depends upon the levels of all other variables.

Figures 1 and 2 illustrate graphically how the predicted probability of a new recession changes with the levels of real equity price growth and real house price growth, for different levels of contributions of the other covariates. These predicted probabilities are overlaid on a histogram showing the distribution of each of these explanatory variables in the sample. For both real asset prices growths, positive growth only affects the predicted probability at very low levels, implying tiny changes in the absolute level of the predicted probability. By contrast, negative growth is associated with much higher levels of predicted probability. Furthermore, when growth is negative, growth changes can lead to large swings in the predicted probability of a new recession, as evinced by the steeper slope of the curve. Interestingly, the predicted probabilities associated with both real asset price growths are similar in shape and magnitude. However, the variability of real house price growth is much less than that for real equity price growth, as seen by the background histograms in the figures. In practice then, it is real equity price drops that convey the stronger signal that a new recession may be imminent, since large drops are rarely seen in real house prices.

Considering in-sample model performance, the Brier score for the baseline model in column 9 is quite low at 0.037, partly reflecting the fact that recessions are rare events.[11] Even so,

[9] The results are very similar if we use Jurado, Ludvigson, and Ng (2015)'s one-quarter ahead macroeconomic uncertainty measure. The estimated model including this alternative measure has a slightly higher AUC. However, unlike the log implied S&P volatility, this measure is not readily available in real-time, requiring extensive calculations to derive it. Because of this we opt to use the readily available log implied S&P volatility.

[10] In the working paper (Bluedorn, Decressin, and Terrones, 2013), we also investigated the robustness of the baseline model to the inclusion of a number of additional explanatory variables, including measures of cross-country asset price spillovers, large equity price drops, long bond yields, real GDP growth, and distributed lags of the variables in our baseline specification. The inclusion of these variables do not substantially alter our main results, providing some reassurance that the findings here are robust.

[11] See the working paper (Bluedorn, Decressin, and Terrones, 2013) for an investigation of a number of alternative estimation methodologies to address the rare events issue associated with recession prediction. We find little evidence of estimation bias due to the rare events problem. The estimation of nonlinear models with fixed effects can also suffer from the incidental parameters problem, leading to inconsistency of the coefficient estimates (Arellano and Hahn, 2007). In our application, we focus on logistic regression which does not have this problem. More generally, we have a large number of time observations (42 years worth) relative to the number of cross-section units which reduces any bias due to the incidental parameters problem markedly.

compared to a naïve forecast (where the probability forecast for a new recession is simply the sample frequency), the baseline model has Brier skill score of 0.171, indicating that it performs about 17 percent better on the basis of mean square probability forecast error.[12]

Selecting the optimal Youden cutoff to make point predictions would lead to a hit rate (true positive) of about 70 percent, while the false alarm rate (false positive) is about 15 percent in-sample. A 95% confidence interval for the AUC statistic of the baseline specification excludes the AUC statistics calculated for the single variable models apart from the models with equity price growth or log implied S&P volatility.

Figure 3 shows that the distribution of in-sample predicted probabilities conditional on a new recession is heavily skewed towards higher probabilities, while the distribution conditional on a continuation of the expansion is peaked near zero. This is another indication that the model does a comparatively good job within sample of separating the quarters prior to new recessions from quarters of continuing expansion.

In summary, the analysis in this section finds that equity and house price drops raise the risk of a new recession. This suggests that periods when these asset prices are both falling should be carefully monitored particularly if the underlying shocks also affect the second moments of these variables. In addition, because of the rare nature of recessions, there is evidence of asymmetry in the effects of these financial variables, where drops in equity and house prices have much larger effects on the likelihood of recession (raising it) than do favorable rises in price growth.

B. Comparison with Model Estimated Over All Observations

To determine whether or not there is a difference between predicting the start of a recession versus its continuation, we estimate our baseline model specification over a sample that includes continuing recession and expansion periods over 1970:Q1-2011:Q4. Table 2 reports the findings of this exercise.

The coefficient estimates from the all observations estimation sample (including both expansion and recession periods, whether new or continuing) tend to be larger in magnitude than those from the baseline sample, apart from real equity price growth which is slightly smaller. In terms of their statistical significance there is no substantial difference as the explanatory variables that are significant in predicting new recessions are the same as those for predicting new and continuing recessions.

The in-sample performance, however, varies across the two estimation samples. While the AUC statistic is roughly similar, at 0.82 or slightly above, the Brier score is some 60 percent lower for the model estimated over the baseline sample than that from the all observations sample. Using the optimal Youden cutoff to make in-sample predictions, the hit rate for the model estimated over all observations is about 74 percent, while that for the baseline sample

---

[12] The Brier score for the naïve forecast is given by the outcome variance, which here is 0.044.

is around 68 percent. At the same time though, the false alarm rate for the all observations model is also somewhat higher, at just over 21 percent versus about 15 percent for the baseline sample. That said, the different evaluation samples for these statistics mean that such comparisons should be taken with a pinch of salt.

Table 3 gets to grips with this, by comparing the in-sample probability predictions for the models in Table 2 over the baseline sample, assessing the ability of the two models to predict new recessions. As shown there, the model estimated over the baseline sample performs better than the model estimated over all observations along a number of metrics. Although the Brier skill score for the baseline versus the naïve forecast is 0.171, it is actually negative for the model estimated over all observations, at -0.136. In other words, the naïve forecast for a new recession is strictly preferable to using the model estimated over all observations. Figure 5 shows that the distribution of in-sample predicted probabilities conditional on a continuing expansion is humped around the empirical probability of a new recession, suggesting some greater likelihood for false alarms by the all observations model. There is less of a clear distinction between the two conditional distributions than that seen in Figure 3.

The ensemble of evidence suggests that the probability process for a new recession is not the same as that for a continuing recession. When forecasting recession starts, one does better excluding from the estimation periods when an economy is already in recession.

C. Out-of-Sample Model Evaluation and Comparison

That said, the results of the in-sample comparison are not very surprising, given that the evaluation sample coincides with the sample used to estimate the baseline. Accordingly, we undertake an out-of-sample evaluation and comparison of the two models.

We make one-step ahead predictions for a recession start by estimating the two models over an expanding window starting in 2005:Q4 and ending in 2011:Q4. This period covers the run-up to the Great Recession and the recovery from this recession. It is interesting to note that some of the G-7 countries have experienced 2 recessions in this period. The initial training samples spanning 1970:Q1-2005:Q4 for the two models differ, with the baseline model estimated over only continuing expansions and recession starts, while the all observations model also includes continuing recessions in its estimation sample.

Table 4 shows the out-of-sample evaluation of the two models over the 7 years considered. Across all performance statistics, the baseline model performs better. The Brier scores for both models are larger than those seen in the in-sample results, at 0.105 for the baseline and 0.123 for the all observations model. This translates into a Brier relative skill score for the baseline versus the all observations model of 0.146, suggesting its superiority.

To move from out-of-sample probability to point prediction with the expanding sample, the Youden cutoff threshold for the one-step ahead probability prediction at each point in time is inferred from the last set of in-sample estimates. This means that the classification threshold may change over time, as new information on the in-sample performance of the model is gained. Using this decision procedure, the hit rate from the baseline and all observations models is identical, at 70 percent. Both models successfully predict the Great Recession for all G-7 economies, apart from Canada.

However, the two models differ in their false alarm rates, with the all observations model at about 10.6 percent while the baseline model is at 8.7 percent. Over about 35 years for the G-7 sample, this difference would translate into 25 more erroneous recession calls by the all observations model. The Peirce skill score (difference between the hit and false alarm rates) is consequently larger for the baseline model. The probability of correct forecast is relatively high for both models at almost 90 percent, reflecting the rarity of recession starts, although the baseline model exhibits a slight advantage.

IV. Conclusion

In this paper, we examined the usefulness of asset prices—equity prices and house prices—in predicting new recessions in the G-7 countries over the past forty years. Our focus on the starts of recessions differs from much of the literature, which has tended to pool recession starts and periods of ongoing recession in their estimation sample. The analysis suggests that asset price drops are significantly associated with the start of a recession in these economies. In particular, the marginal effect of an equity/price drop on the likelihood of a new recession can be substantial. As large equity price drops are observed with higher frequency than large house price drops, this supports the view that equity prices can help predict new recessions. Increased market uncertainty, a variable associated with equity price changes, is also helpful in predicting new recessions in these countries.

These findings hold even when the term spread (a popular predictor in the literature), and real oil price growth are included as explanatory variables. While confirming the usefulness of the term spread, there is no evidence that oil price movements are helpful predictors of economic contractions. Moreover, there is evidence that changes in equity price have better in-sample forecasting performance than many of the other commonly featured recession predictors, including the term spread.

Our analysis also suggests that there is a difference between predicting the start of a recession versus its continuation. Both the in-sample and out-of-sample performance for predicting new recessions is better when the model is estimated or trained only over periods of continuing expansion or new recessions. The two events are better treated separately.

Going forward, we would like to extend our analysis to all the advanced economies. This is not just an intellectual exercise but, more importantly, the development of such a framework may help policymakers to straightforwardly assess the risks of a new recession both in their own countries as well as in their financing and trading partners. In addition, we would like to study further the role of market uncertainty in predicting recessions in advanced economies by exploring additional measures of uncertainty and by examining the extent to which market uncertainty and asset price drops are interlinked.

References

Arellano, Manuel and Jinyong Hahn, 2007, "Understanding Bias in Nonlinear Panel Models: Some Recent Developments," *Advances in Economics and Econometrics*, vol. 3, pp. 381-409, ed. by R. Blundell, W. Newey and T. Persson (Cambridge University Press).

Berge, Travis J. and Òscar Jordà, 2011, "Evaluating the Classification of Economic Activity into Recessions and Expansions," *American Economic Journal: Macroeconomics*, Vol. 3, No. 2, April, pp. 246–277.

Bernanke, Ben, 1983, "Irreversibility, Uncertainty, and Cyclical Investment," *Quarterly Journal of Economics*, Vol. 98, No. 1, pp. 85–106.

Bloom, Nicholas, 2009, "The Impact of Uncertainty Shocks," *Econometrica*, Vol. 77, No.3, pp. 623–685.

Bluedorn, John C., Jörg Decressin, and Marco E. Terrones, 2013, "Do Asset Price Drops Foreshadow Recessions?," IMF Working Paper 13/203 (Washington D.C.: International Monetary Fund).

Bry, Gerhard and Charlotte Boschan, 1971, *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs*, NBER Technical Paper, No. 20 (New York: Columbia University Press).

Carroll, Christopher, Misuzu Otsuka and Jirka Slacalek, 2011, "How Large Are Housing and Financial Wealth Effects? A New Approach," *Journal of Money, Credit and Banking*, Vol. 43, No. 1, pp. 55–79.

Case, Karl E., John M. Quigley, and Robert J. Shiller, 2013, "Wealth Effect Revisited: 1975–2012," NBER Working Paper No. 18667 (Cambridge, Massachusetts: National Bureau of Economic Research).

Christiansen, Charlotte, 2013, "Predicting Severe Simultaneous Recessions Using Yield Spreads as Leading Indicators," *Journal of International Money and Finance,* Vol. 32, February, pp. 1032–1043.

Claessens, Stijn, Ayhan Kose, and Marco E. Terrones, 2012, "How Do Business and Financial Cycles Interact?" *Journal of International Economics*, Vol. 87, pp. 178–190.

Dixit, Avinash K. and Robert S. Pindyck, 1994, *Investment Under Uncertainty* (Princeton, New Jersey: Princeton University Press).

Duarte, Agustin, Ioannis Venetis, and Ivan Paya, 2005, "Predicting Real Growth and the Probability of Recession in the Euro Area using the Yield Spread," *International Journal of Forecasting*, Vol. 21, pp. 261–277.

Engemann, Kristie, Kevin L. Kliesen and Michael T. Owyang, 2011, "Do Oil Shocks Drive Business Cycles? Some U.S. and International Evidence," Federal Reserve Bank of St. Louis Working Paper 2010-007D.

Estrella, Arturo and Frederic S. Mishkin, 1998, "Predicting U.S. Recessions: Financial Variables as Leading Indicators," *Review of Economics and Statistics*, pp. 45–61.

Hamilton, James D., 2011b, "Historical Oil Shocks," in *Handbook of Major Events in Economic History,* ed. by Randall E. Parker, Robert M. Whaples (New York: Routledge).

Harding, Don, and Adrian Pagan, 2002, "Dissecting the Cycle: A Methodological Investigation," *Journal of Monetary Economics*, Vol. 49, pp. 365–381.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor, 2011, "Financial Crises, Credit Booms, and External Imbalances: 140 Years of Lessons," *IMF Economic Review*, Vol. 59, No. 2, June, pp. 340–378.

Perkins, Neil J. and Enrique F. Schisterman, 2006, "The Inconsistency of 'Optimal' Cutpoints Obtained Using Two Criteria Based on the Receiver Operating Characteristic Curve," *American Journal of Epidemiology*, Vol. 163, No. 7, pp. 670–675.

Samuelson, Paul, 1966, "Science and Stocks," *Newsweek*, September 19, p. 92.

Schularick, Moritz and Alan M. Taylor, 2012, "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles and Financial Crises, 1870–2008," *American Economic Review*, Vol. 102, No. 2, pp. 1029–61.

Stock, James H. and Mark W. Watson, 2003, "Forecasting Output and Inflation: The Role of Asset Prices," *Journal of Economic Literature*, Vol. XLI, pp. 788–829.

Youden, W. J., 1950, "Index for Rating Diagnostic Tests," *Cancer*, Vol. 3, No. 1, pp. 32–35.

Zou, Kelly, James O'Malley, and Laura Mauri, 2007, "Receiver-Operating Characteristic Analysis for Evaluating Diagnostic Tests and Predictive Models," *Circulation*, pp. 654–657.

Table 1. Explaining Recession Starts in the G-7

Baseline Sample, 1970:Q1-2011:Q4

	Logistic Regression Model									
Explanatory Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	Aver. Marg. Eff.
Real Equity Price Growth	-0.124***					-0.125***	-0.106***	-0.102***	-0.102***	-0.00384***
	(0.0216)					(0.0220)	(0.0245)	(0.0282)	(0.0293)	(0.0010)
Real House Price Growth		-0.0875				-0.0819	-0.0944*	-0.115***	-0.116***	-0.00437***
		(0.0727)				(0.0513)	(0.0562)	(0.0437)	(0.0442)	(0.0015)
Log Implied/Realized S&P Volatility			2.377***				0.947**	0.963*	0.970*	0.0365*
			(0.4460)				(0.3960)	(0.5030)	(0.5110)	(0.0186)
Term Spread				-0.360**				-0.317**	-0.317**	-0.0119**
				(0.1440)				(0.1510)	(0.1500)	(0.0053)
Real Oil Price Growth					0.00243				0.00084	0.00003
					(0.0070)				(0.0064)	(0.0002)
Observations	945	945	945	945	945	945	945	945	945	
Pseudo R-squared	0.18	0.0672	0.113	0.105	0.0628	0.185	0.191	0.223	0.223	
No. of Cases	44	44	44	44	44	44	44	44	44	
Log-Likelihood	-145.8	-166	-157.8	-159.2	-166.7	-145.1	-143.9	-138.2	-138.2	
AUC	0.794	0.71	0.763	0.739	0.691	0.799	0.806	0.825	0.825	
95% LB for AUC	0.724	0.627	0.691	0.659	0.613	0.729	0.737	0.753	0.754	
95% UB for AUC	0.864	0.793	0.835	0.818	0.769	0.87	0.875	0.896	0.896	
Brier Score	0.0381	0.0431	0.0421	0.0416	0.0431	0.038	0.038	0.0368	0.0368	
Outcome Var.	0.0444	0.0444	0.0444	0.0444	0.0444	0.0444	0.0444	0.0444	0.0444	
Forecast Var.	0.00494	0.00133	0.00249	0.00244	0.0012	0.00509	0.00543	0.00671	0.00671	
Twice Cov. of Forecast/Outcome	0.0112	0.00263	0.00483	0.00526	0.00245	0.0114	0.0118	0.0143	0.0143	
Spiegelhalter Z-Stat.	-0.257	0.00499	0.0281	-0.0688	-0.00889	-0.249	-0.192	-0.186	-0.184	
Optimal Youden Cutoff	0.0374	0.0492	0.0684	0.0356	0.0491	0.0405	0.0368	0.0674	0.068	
True Pos. Rate	75	68.18	59.09	81.82	63.64	75	79.55	70.45	68.18	
False Pos. Rate	31.41	31.41	18.65	42.18	30.85	28.75	30.85	14.98	14.65	
Link Test Z-Stat.	1.51	0.144	-0.288	0.873	0.279	1.493	1.147	1.161	1.14	

Note: The dependent variable is the Bry-Boschan/Harding-Pagan algorithm identified peak for seasonally adjusted, quarterly real GDP growth. The baseline sample includes continuing expansion periods. Heteroskedasticity and autocorrelation-within-country robust standard errors are in parentheses underneath the coefficient estimate. * indicates statistical significance at the 10% level, ** at the 5% level, and *** at the 1% level. Shown only for model (9), the average marginal effects show the average impact of a one-unit change in the explanatory variables on the probability of a new recession in the next quarter. Growth rates are log differences times 100. All models include country-specific intercepts and quarterly dummies.

Table 2. Explaining Recession Starts versus Recession Periods in the G-7

1970:Q1-2011:Q4

Explanatory Variable	Baseline	All Obs.
Real Equity Price Growth	-0.102***	-0.0766***
	(0.0293)	(0.0181)
Real House Price Growth	-0.116***	-0.274***
	(0.0442)	(0.0948)
Log Implied/Realized S&P Volatility	0.970*	1.524***
	(0.5110)	(0.5190)
Term Spread	-0.317**	-0.315***
	(0.1500)	(0.1080)
Real Oil Price Growth	0.00084	-0.00110
	(0.0064)	(0.0040)
Observations	945	1154
Pseudo R-squared	0.223	0.213
No. of Cases	44	168
Log-Likelihood	-138.2	-376.7
AUC	0.825	0.820
95% LB for AUC	0.754	0.784
95% UB for AUC	0.896	0.856
Brier Score	0.0368	0.099
Outcome Var.	0.0444	0.124
Forecast Var.	0.00671	0.027
Twice Cov. of Forecast/Outcome	0.0143	0.053
Spiegelhalter Z-Stat.	-0.184	0.161
Optimal Youden Cutoff	0.068	0.147
True Pos. Rate	68.18	73.81
False Pos. Rate	14.65	21.3
Link Test Z-Stat.	1.14	-1.366

Note: Column shows the model estimates for the baseline sample, which includes continuing expansion periods. Column 2 shows the model estimates for the full sample, which includes all observations (both expansion and recession periods). Heteroskedasticity and autocorrelation-within-country robust standard errors are in parentheses underneath the coefficient estimate. * indicates statistical significance at the 10% level, ** at the 5% level, and *** at the 1% level. Growth rates are log differences times 100. All models include country-specific intercepts and quarterly dummies.

Table 3. Explaining Recession Starts in the G-7
Comparison by Estimation Sample

1970:Q1-2011:Q4

Statistic	Baseline	All Obs.
AUC	0.825	0.797
95% LB for AUC	0.754	0.717
95% UB for AUC	0.896	0.878
Brier Score	0.0368	0.050
Outcome Var.	0.0444	0.044
Forecast Var.	0.00671	0.019
Twice Cov. of Forecast/Outcome	0.0143	0.019
Spiegelhalter Z-Stat.	-0.184	-5.902
Optimal Youden Cutoff	0.068	0.213
True Pos. Rate	68.18	63.69
False Pos. Rate	14.65	13.08
Link Test Z-Stat.	1.14	-0.228

Note: The first column shows the performance of the model estimated over the baseline sample for predicting recession starts. The second column shows the performance of the model estimated over all observations (including all expansion and recession periods) for predicting recession starts.

Table 4. Predicting Recession Starts Out-of-Sample
One-Step Ahead Model Performance

2006:Q1-2011:Q4

Statistic	Baseline	All Obs.
Brier Score	0.105	0.123
Pierce Skill Score	0.613	0.594
Odds Ratio Skill Score	0.922	0.904
Probability of Correct Forecast	0.895	0.877
True Pos. Rate (Hit Rate)	70.00	70.00
False Pos. Rate (False Alarm Rate)	8.65	10.58
Conditional Miss Rate	3.06	3.13
False Alarm Ratio	56.25	61.11

Note: The first column shows the performance of the model estimated over the baseline sample for predicting recession starts. The second column shows the performance of the model estimated over all observations (including all expansion and recession periods) for predicting recession starts.

Figure 1.

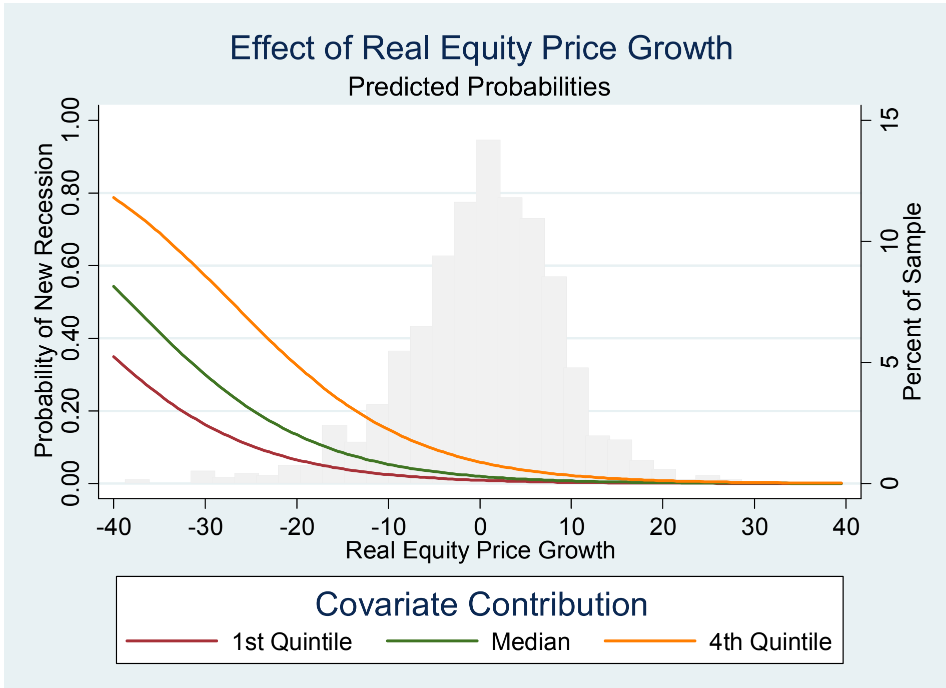

Figure 2.

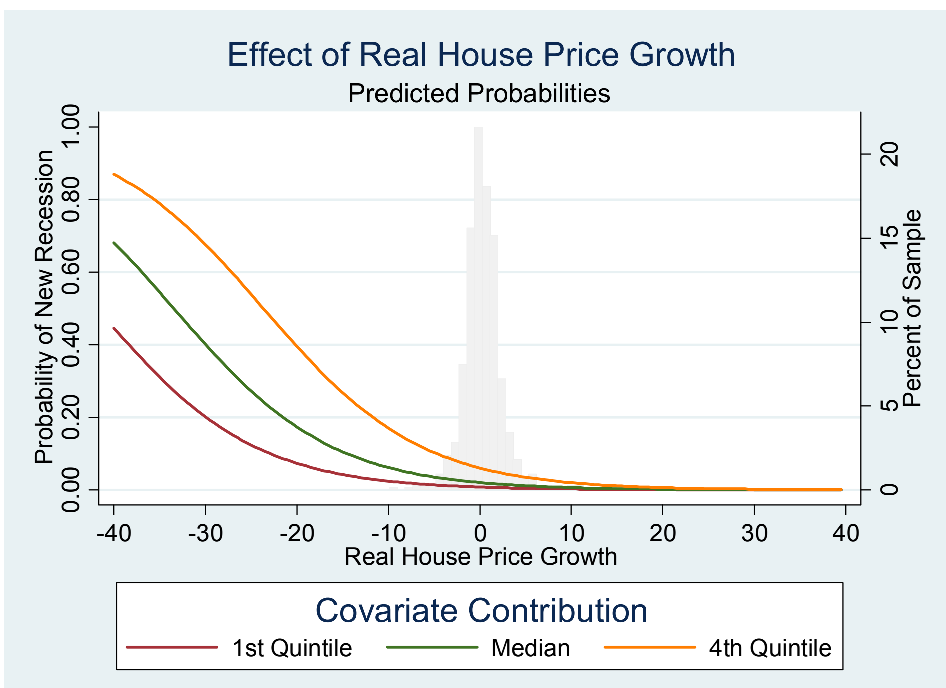

Figure 3.

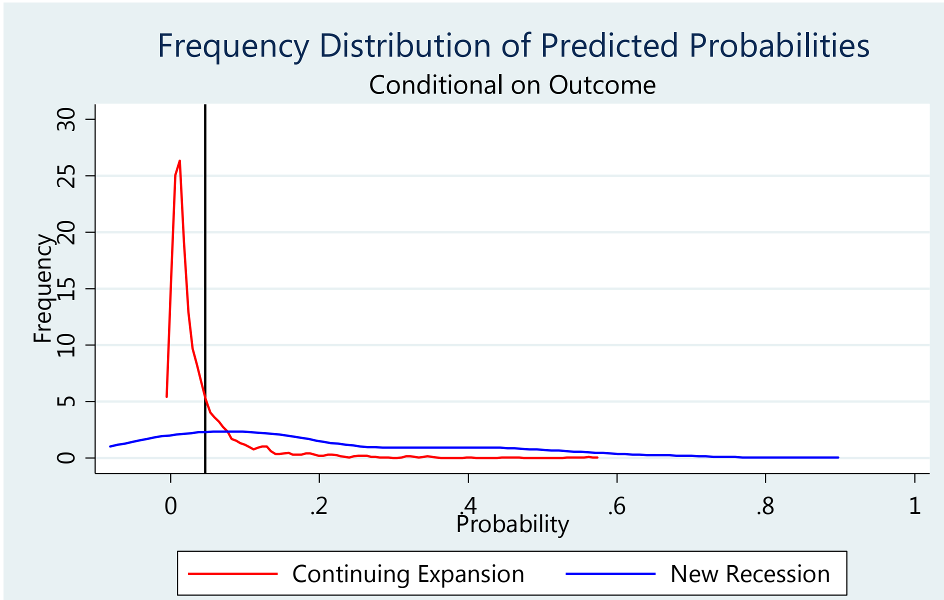

Note: These are the conditional densities of the predicted probabilities from the baseline model estimated over the sample of continuing expansions and recession starts (Table 4, column 1). The unconditional probability of a new recession (vertical line) is 0.047.

Figure 4.

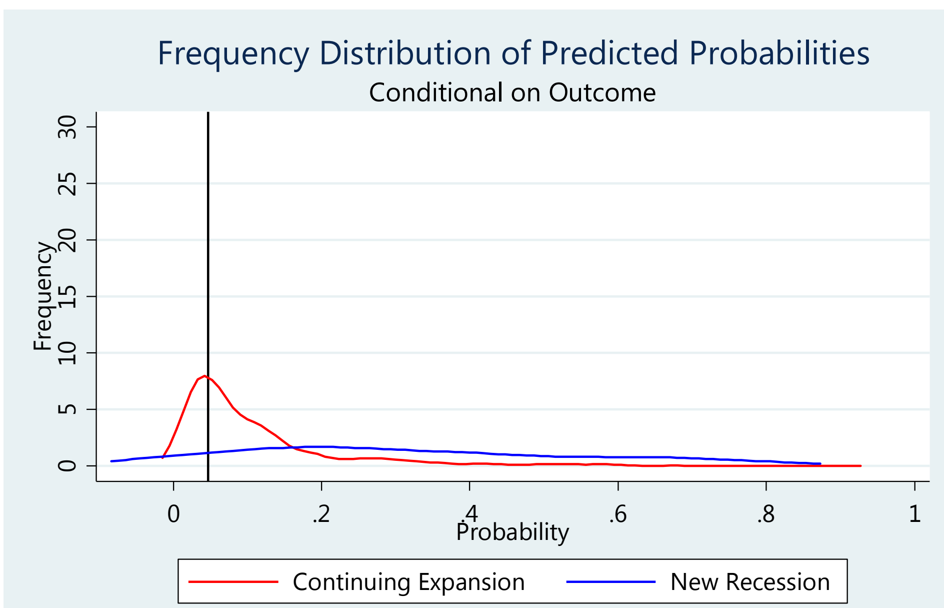

Note: These are the conditional densities of the predicted probabilities from the baseline model estimated over the sample of all expansion and recession periods (Table 4, column 2). The unconditional probability of a new recession (vertical line) is 0.047.