Volume Title: Asset Prices and Monetary Policy

Volume Author/Editor: John Y. Campbell, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-09211-9

Volume URL: http://www.nber.org/books/camp06-1

Conference Date: May 5-6, 2006

Publication Date: September 2008

Chapter Title: Measuring the Macroeconomic Risks Posed by Asset Price Booms

Chapter Author: Stephen G. Cecchetti

Chapter URL: http://www.nber.org/chapters/c5368

Chapter pages in book: (p. 9 - 43)

1

Measuring the Macroeconomic Risks Posed by Asset Price Booms

Stephen G. Cecchetti

1.1 Introduction

We pay central bankers to be paranoid. One of their primary responsibilities is to do extensive contingency planning, preparing for every possible calamity. And when they do their job well, most of us don't even notice. In the past decade, there are numerous examples of the central bank actions that were taken in response to an increase in the probability of disaster. These include the Federal Open Market Committee's interest rate reductions in the fall of 1998 that followed the Russian government's bond default, the preparations for the century date change, the enormous liquidity injections in the immediate aftermath of the September 11, 2001 terrorist attacks in the United States, as well as the discussions that occurred as nominal interest rates and inflation approached zero simultaneously. All of these episodes demonstrate policymakers' willingness to take actions in order to reduce the chance of disaster, acting as the risk mangers for the economic and financial system.

Then Federal Reserve Board Chairman Alan Greenspan put it best in

Stephen G. Cecchetti is the Rosenberg Professor of Global Finance, International Business School, Brandeis University, and a research associate of the National Bureau of Economic Research.

An earlier version of this chapter was distributed under the title "GDP at Risk: A Framework for Monetary Policy Response to Asset Price Movements." I would like to thank the discussant Andrew Levin as well as John Campbell, Blake LeBaron, Peter Phillips, Ritirupa Samanta, Jeremy Stein, and the participants at the conference for their numerous comments and suggestions. In addition, I owe an enormous debt to various collaborators over the years on related topics, especially Michael Bryan, Hans Genberg, Stefan Krause, Róisin O'Sullivan, and Sushil Wadhwani. Anne LePard and Damir Cosic provided research assistance. Finally, I would like to thank the Bank for International Settlements (BIS) for supplying some of the data. All of the views expressed here, as well as the errors, are my own.

2003 when he said that "A central bank seeking to maximize its probability of achieving its goals is driven, I believe, to a risk-management approach to policy. By this I mean that policymakers need to consider not only the most likely future path for the economy but also the distribution of possible outcomes about that path" (Greenspan 2003, 3). Importantly, the common practice of risk management requires controlling the probability of catastrophe. For a financial intermediary, the focus is on reducing the risk of significant monetary loss. For a central banker, it means acting to reduce the chances that output or the price level will be substantially below trend.

To control risk in financial institutions, risk managers employ the concept of *value-at-risk* (VaR). Value-at-risk measures the worst possible loss over a specific time horizon, at a given probability.[1] A commercial bank might say that the daily VaR for a trader controlling $100 million is $10 million at a 0.1 percent probability. That means that, given the historical data used in the bank's models, the trader cannot take a position that has more than one chance in 1,000 of losing 10 percent in one day.

In some circumstances, VaR is all you need. For example, if it is being used to measure the probability of institutional insolvency, it doesn't really matter how insolvent you are. But policymakers care not only about VaR, they are also concerned about the expected loss given that an event is in the lower tail—something called the *expected tail loss* (ETL). That is, not only where the 5th or 10th percentile of the distribution of gross domestic product (GDP) outcomes falls, but the expected value conditional on being in the lowest 5th or 10th percentile.

Risk-management measures like VaR and ETL are computed from the lower tail of the distribution of possible outcomes, examining the worst events that could occur. This requires moving beyond simple quadratic measures of risk like variance or standard deviation. It is fairly easy to imagine circumstances where the worst possible events have become worse, but the standard deviation of the distribution of all the possibilities is the same. This is one view of the case in the fall of 1998. The point forecasts for the aggregate price level and the GDP gap, and their standard deviation stayed roughly the same. But the lower tail shifted—the probability and size of a very bad outcome—rose. Policymakers acted in response to the perception that the GDP at risk and ETL had gone up.[2]

A risk-management approach comes naturally to central bankers. It is the basis for the creation and maintenance of the lender of last resort: the policy of providing loans to private financial intermediaries that are illiquid but not insolvent helps to ensure that the payments system continues

1. See Jorion (2001).
2. Formally, this means that the central bank's loss function is not quadratic. For a recent discussion, see Surico (forthcoming).

to operate smoothly. Together with deposit insurance, central bank lending is designed to reduce the probability of bank runs to a negligible level. (The implementation of prudential regulation and supervision is the response to the moral hazard created by these policies.)

All of this makes it surprising that many central bankers are hesitant to address the potential risks created by asset price booms and crashes—what are commonly referred to as *bubbles.* The evidence is not in dispute. Bubbles increase the volatility of growth and inflation and threaten the stability of the financial system. The 2003 IMF *World Economic Outlook* estimates that the average equity price bust lasts for 2.5 years and is associated with a 4 percent GDP loss that affects both consumption and investment. While less frequent, property (or housing) busts are twice as long and are associated with output losses that are twice as large.[3]

Asset price bubbles distort decisions throughout the economy. Wealth effects cause consumption to expand rapidly and then collapse. Increases in equity prices make it easier for firms to finance new projects, causing investment to boom and then bust. The collateral used to back loans is overvalued, so when prices collapse, it impairs the balance sheets of financial intermediaries that did the lending. It is the job of central bankers to eliminate the sort of economic distress caused by asset price bubbles. Although the rhetoric has been changing slowly, especially in the case of the responses to Australian and British housing market booms several years ago, most monetary policymakers remain reluctant to act directly to manage these risks.

Any discussion of bubbles must distinguish between equity and property prices. This is true for several reasons. First, the efficient markets hypothesis is more likely to apply to equity than to property. Arbitrage in stocks, which requires the ability to short sell, is at least possible. In housing and property, it is not. Second, even in the few countries with sizeable equity markets, ownership tends to be highly concentrated among the wealthy—people whose consumption decisions are well insulated from the vicissitudes of the stock market. By contrast, home ownership is spread much further down the income and wealth distribution. Finally, in many countries, housing purchases are highly leveraged, leaving the balance sheets of both households and financial intermediaries exposed to large price declines. This suggests that the macroeconomic impact of a boom and crash cycle in property prices might be larger in countries that have more credit outstanding.[4]

In this chapter, I examine equity and housing price booms and crashes from a risk management perspective. Using equity price data from twenty-

3. See the excellent essays in Chapter II of IMF (2003) for a summary of the evidence.

4. For a somewhat more detailed discussion of the issues and the debate, see Cecchetti (2003).

seven countries and housing price data from seventeen countries, I will look at the various consequences of rising equity and housing prices for growth and inflation. I begin by examining how asset price booms influence the mean and variance of deviations in (log) output and (log) price level from their (time-varying) trends. I then proceed to measure both the GDP at risk and the price level at risk that these booms create.

The scarcity of booms and crashes, especially in property prices, means that I must pool data across countries. From what data there are, I come to the following conclusions: housing booms are bad in virtually every way imaginable; they drive the output gap down, increase its volatility, increase GDP at risk, and push the lower tail of outcomes (ETL) even lower (decreasing the expected value of the GDP gap conditional on being in the lower tail of the distribution). By contrast, equity booms have little impact on either the level or volatility of the output and price-level gaps at horizons of three years; do not change GDP at risk, but increase the risk of prices falling dramatically below trend; and drive the lower tail (ETL) even lower.

Before continuing, it is worth noting the relationship between the use of risk management and robust control in the context of monetary policy. Robust control examines policy making in the presence of model uncertainty.[5] Instead of choosing optimal policy based on the most likely economic model, it ignores the probability that any particular model of the economy is true and selects the policy that delivers the best result even when the worst model is true. That is, it computes the policy path or instrument rule that minimizes the maximum loss, regardless of how likely or unlikely that case might be. As Onatski and Stock (2002) show, in contrast to the standard case in which uncertainty breeds caution, this has the potential to yield aggressive policy responses—aggressive enough to ensure that the worst outcomes are avoided.[6]

While the risk management and robust-control approaches to policy making both go beyond simple quadratic measures of loss, they are quite different. Rather than focusing on model uncertainty ignoring the probability of particular cases, concepts like VaR and ETL are designed to help control both the size and likelihood of bad outcomes. In the case of asset price booms, that means first computing the probability distribution associated with growth and inflation outcomes conditional on seeing equity or property prices rise suddenly and then looking for ways in which policymakers can mitigate the worst possible outcome.

5. See Svensson (2007) and Dennis, Leitemo, and Söderström (2006) for discussions.

6. The simplest example involves the case of inflation control. Imagine that policymakers are unsure whether inflation follows a random walk or not—the largest root of the autoregressive representation of the inflation rate is estimated to be less than one, but there is some finite probability that it actually equals one. Because it is the worst possible, the robust control solution is for policymakers to react to shocks as if inflation were nonstationary. In most environments, that means aggressively countering virtually anything that would force inflation up.

The remainder of this chapter proceeds as follows. Section 1.2 provides overwhelming evidence that the distribution of output and price-level deviations from their trends have fat tails, implying that methods based on quadratic loss and normal approximations could be misleading. Then, in section 1.3, I characterize the distribution of output and price-level conditional on housing and equity booms. That is, I look at the mean, variance, value-at-risk, and expected lower tail of output and price level conditional on asset price booms. Overall, the results suggest that normal approximations are inadequate. Section 1.4 expands the discussion contrasting housing and equity booms.

There is a growing consensus that traditional interest rate policy is not very useful in the battle to combat the deleterious macroeconomic effects of asset price bubbles.[7] At the same time, it is clear that policymakers cannot ignore the threat that equity and housing booms and busts pose for central bankers' stabilization goals. Adopting a risk management perspective means asking whether there are institutional solutions to the problem. That is, are there ways to structure the financial system that will then inoculate the real economy from the adverse effects of bubbles? With this question in mind, I examine the relative impact of asset price booms in economies with market- versus bank-based financial systems. The results, reported in section 1.5, suggest that market-based systems have a somewhat higher GDP at risk in the aftermath of equity booms, but those systems weather housing booms equally poorly.

1.2 GDP and Prices: General Considerations

Financial economists employ concepts like value-at-risk in order to address the problems created by fat tails. That is, cases in which a normal (Gaussian) distribution provides an overly optimistic picture of the likelihood of extreme events. Equity returns are notorious for exhibiting high probabilities of extreme events in their lower tail. Because these "bad" outcomes are so important for controlling the risk of large losses, modeling them has attracted substantial attention.[8]

Aggregate output and prices share some of the properties exhibited by equity returns. The distribution of deviations of (log) output and the (log) price level from their respective trend exhibit fat tails. That is, the probability of observing a large negative realization is substantially higher than one would infer from a Gaussian distribution. To see this, I have calculated the 5th percentile of the distribution of log output and log price-level deviations from their Hodrick-Prescott (Hodrick and Prescott 1997) trends,

7. See Cecchetti (2006) for a discussion.
8. See LeBaron and Samanta (2005) for a discussion of the issues surrounding modeling fat-tailed distributions.

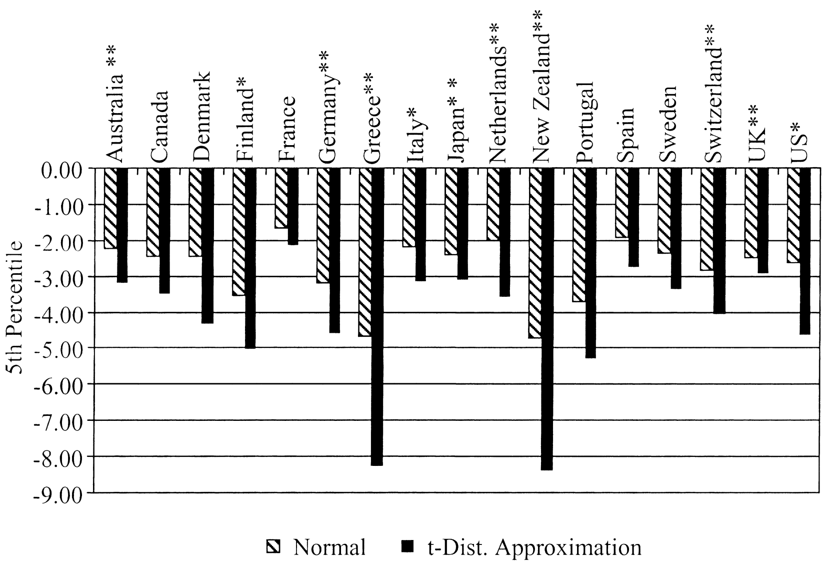

Fig. 1.1 GDP at risk, normal versus *t*-distribution approximation. GDP at risk: Normal versus fat-tailed

Notes: The *s refer to the significance level of the Jacque-Bera test for normality. A single * is for countries with a p-value of 0.10 or less, while ** signifies a p-value of 0.05 or less. The test statistic equals $(n/6)[\mu_3^2 + (\mu_4 - 3)/4]$, where μ_3 and μ_4 are the sample third and fourth moments, and n is the sample size. The statistic is distributed as χ-squared with 2 degrees of freedom. Test results are reported for the deviations of quarterly log GDP and log prices from a Hodrick-Prescott filtered trend with parameter equal to 1,600. The sample is from 1970 to 2003.

with smoothing parameter set to 1,600, for a series of countries using quarterly data from 1970 to 2003.[9] These results are plotted in figures 1.1 and 1.2. (The appendix provides a more detailed description of the data.) The figures also include results for a Jacque-Bera test for normality—these are the *s next to the country names. Normality is rejected for eleven of seventeen cases using the output gap and ten of seventeen using the price-level gap.

The figures show the results for the following calculation. For the normal distribution, this is just 1.645 times the standard deviation of the series. The alternative, which takes the fatness of the tails of the distribution into

9. I have also computed results for a shorter sample beginning in 1985 that verify the inaccuracies of the normal approximation reported in the following. In addition, the results throughout the chapter are robust to using a smoothing parameter of 9,600, rather than 1,600; to using the residuals from a four-order autoregression; and to using the residuals from the estimation of a two-equation aggregate demand—aggregate supply model based on Rudebusch and Svensson (1999) as implemented in Cecchetti, Flores-Lagunes, and Krause (2006) that includes interest rates.

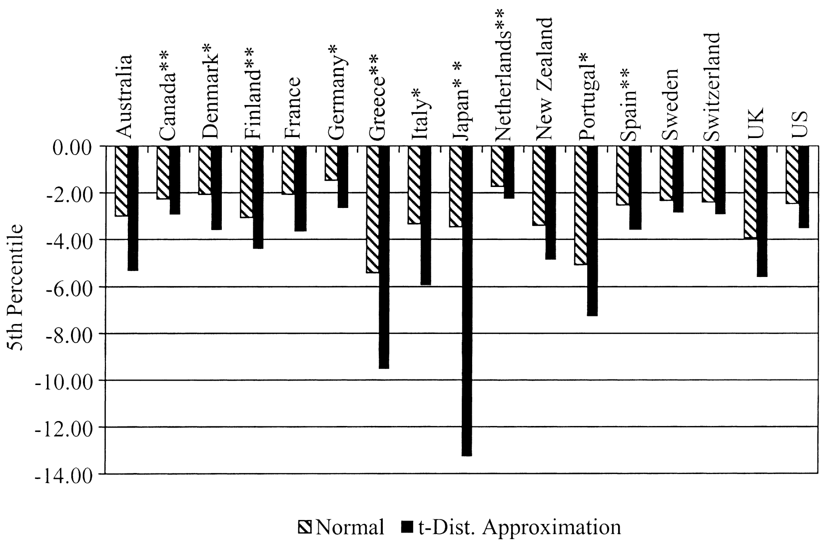

Fig. 1.2 Price level at risk, normal versus *t*-distribution approximation. Price level at risk: Normal versus fat-tailed

Note: See explanatory note for figure 1.1.

account, begins by the computation of a Hill index. As described in LeBaron and Samanta (2005), the Hill index is an estimate of the number of moments of a distribution that exists. For a normal distribution, the index is infinity. After computing the index, the tail is approximately distributed as a Student *t* with degrees of freedom equal to the Hill index value. So the *t*-distribution approximation to the 5th percentile of the deviations of log GDP or the log price level from their trend is equal to the standard deviation of the series times the 5 percent level of the *t*-distribution with degrees of freedom equal to the series' Hill index.[10]

As one would expect, in some countries the deviations of output and prices from trend—their output and price-level gaps—have fatter tails than others. But if one were to use the normal distribution, the errors would be large—averaging roughly 50 percent. For the United States, the 5th percentile of the normal distribution implies a deviation of output from trend of slightly more than –2.5 percent. Taking the fatness of the lower tail of the actual data into account yields an estimate of more than 4.5 percent. That is, the 5 percent GDP at risk for the United States (without conditioning on anything). For the price level, the estimates diverge by

10. Computation of the Hill index requires a decision about where the tail of the distribution starts. I take LeBaron and Samanta's (2005) advice and use the bottom 10 percent of the observations.

less with the normal distribution, giving a 5 percent price-level at risk equal to –2.5 percent and the t-distribution approximation yielding an estimate of –3.5 percent.

It is important to keep in mind that standard statistical and econometric procedures are designed to characterize behavior near the mean of the data, so they are particularly ill-suited to the examination of tail events. This means that when extreme events are more likely than the normal distribution implies, and we care about them, it is important to adopt techniques that explicitly account for fat tails.

1.3 Risks Created by Asset Price Bubbles

Managing risk means having information about the entire distribution of possible outcomes. That is, one needs to know not only the mean and variance, but tail probabilities as well. With that in mind, I now compute the mean, variance, value-at-risk, and expected lower tail for output and price-level deviations from their trends, all conditional on the asset price booms.

1.3.1 The Mean

How do asset price booms change the mean and volatility of output and price-level gaps? I examine this question using a series of regression, which allow straightforward statistical inference. Throughout this exercise, I treat asset price booms and busts as events that are exogenous with respect to the behavior of output and price paths several years into the future.[11]

To study the conditional mean, consider the following regression:

$$x_{it} = a + b\mathrm{d}_{\mathrm{it-k}}(\alpha) + \varepsilon_{it}, \qquad (1)$$

where x_{it} is the level of the output (or price-level) gap; $\mathrm{d}_{\mathrm{it-k}}(\alpha)$ is a dummy variable that takes on the value 1 if k periods earlier the filtered asset price data exceeds the threshold α. The coefficient b measures the impact of the asset price boom on the distribution of the gap variable.

Before continuing, let me pause to describe the procedure used to construct the data.[12] First, for each country, I take the deviation of the log of each series—real GDP, the aggregate price level, the real equity price index, and the real housing price index—from its Hodrick-Prescott filtered trend with a smoothing parameter equal to 1,600 (the results are robust to using a parameter of 9,600). All data are quarterly, and most samples are

11. The fact that the results are robust to various changes in the filtering of the data, including the use of residuals from a simple model in place of the simple filtered data as described in footnote 7, suggests that this assumption is relatively innocuous.

12. The seventeen countries in the housing price sample are Australia, Belgium, Canada, Denmark, Finland, Greece, Ireland, Israel, the Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, the United Kingdom, and the United States. The twenty-seven countries in the equity price data sample add Austria, Chile, France, Germany, Italy, Japan, Korea, Mexico, Peru, and South Africa.

from 1970 to 2003.[13] To construct the dummy variable d_{it-k}, I filter the log equity and housing price data using a Hodrick-Prescott filter with smoothing parameter equal to 3,200 (again, this is robust to increasing the parameter value). It is important to note that the use of a two-sided filter means that large positive deviations of asset prices from trend—these are the booms—must be followed by crashes. Put another way, the booms I locate cannot continue indefinitely.

Finally, taking deviations from country-specific (and time-varying) trends has the advantage in that it removes country-fixed effects. While there are surely numerous conditions that vary in these countries over the sample, this is at least a minimum condition for pooling.[14]

Returning to the results, table 1.1 reports estimates for equation (1). To read the table, take the example of the last entry in the third column under housing. That's the one where the threshold α equals 10 percent, and the lag k is twelve quarters. For this case, the estimate of b is –1.42 with a p-value of 0.00. This means that, conditional on seeing a housing boom that is 10 percent above trend, the mean of the output gap twelve quarters later is, on average, 1.42 percent. That seems like a big number, and it is precisely estimated.[15]

Overall, these results allow a number of conclusions. First, in the near term, at horizons of four quarters, both equity and housing booms lead to positive output gaps. This is for the simple reason that at a four-quarter horizon, an asset-market boom is likely to continue, adding fuel to the general economic growth. Second, housing booms create future declines in output and increases in prices while equity booms do not. And third, the bigger the housing boom, the bigger the expected drop in output and the expected increase in the price level.

1.3.2 Volatility

Next, I examine the impact of asset price booms on the volatility of output and price deviation from trend. To do this, I regress the square of the gap, that is $(x_{it})^2$ on the dummy variable $d_{it-k}(\alpha)$. That is,

$$(x_{it})^2 = a' + b' d_{it-k}(\alpha) + \upsilon_{it}. \tag{2}$$

To simplify interpretation, I standardize the data, dividing by the variance of the entire sample. This means that the coefficient is a measure of the percentage increase in the volatility. So, for example, a number like 5.28

13. While it would be interesting to look at shorter samples, there is simply not enough data to do it.

14. As in section 1.2, the results in section 1.3 are robust to use of residuals from a fourth-order autoregression and to use of residuals from a model that includes interest rates and external prices.

15. To address problems of heteroskedasiticity (throughout) and serial correlation (within each country), I have estimated the standard errors and resulting p-values using a panel version the Newey-West (Newey and West 1987) procedure with lags equal to 1.5k.

Table 1.1 Impact of asset price booms on the levels (lag of asset price [*k*])

	Level of the output gap			Price level		
Threshold (α)	4	8	12	4	8	12
			Equity			
Data	*0.03*	*0.01*	0.00	–0.07	0.00	0.04
	1.00	*0.96*	0.25	0.03	0.50	0.92
4	*1.05*	*0.28*	–0.21	–0.61	0.10	0.99
	1.00	*0.99*	0.10	0.30	0.58	0.94
12	*0.92*	*0.32*	–0.15	0.04	0.54	1.32
	1.00	*0.99*	0.23	0.51	0.71	0.92
20	*0.85*	0.16	–0.07	–0.65	0.71	1.58
	1.00	0.81	0.38	0.39	0.69	0.88
			Housing			
Data	*0.06*	**–0.04**	**–0.09**	*0.04*	*0.08*	*0.07*
	1.00	**0.00**	**0.00**	*0.99*	*1.00*	*1.00*
2	*0.46*	**–0.53**	**–0.92**	*0.62*	*0.96*	*0.70*
	1.00	**0.00**	**0.00**	*1.00*	*1.00*	*1.00*
6	*0.85*	**–0.50**	**–1.28**	*0.55*	*1.14*	*0.95*
	1.00	**0.01**	**0.00**	*0.98*	*1.00*	*1.00*
10	*1.10*	–0.42	**–1.42**	0.52	*1.19*	*1.04*
	1.00	0.12	**0.00**	0.92	*1.00*	*1.00*

Notes: Table 1.1 reports the coefficient b in the regression $x_{it} = a + b\mathrm{d}_{\mathrm{it-k}}(\alpha) + \varepsilon_{it}$, where x is the deviation of either log GDP or the log price from a Hodrick-Prescott filtered trend, with parameter 1,600; and d is either a dummy variable equal to 1 if the filtered asset price exceeds the threshold (in percent), or the filtered asset price data itself. In each case, the first row of numbers is the coefficient itself, while the second row is a p-value for the test that b is strictly less than 0, computed using Newey-West standard errors with lags equal to 1.5 times k. *Italicized* values are significantly greater than 0, while **bold** values are significantly less than 0, both at the 5 percent level. Samples are described in the data appendix.

(that's the estimate for a 10 percent housing price boom at a horizon of four quarters) means a 5.28 percent increase in volatility. The results are reported in table 1.2, and they are quite stark. Housing booms increase the volatility of growth at all horizons, and that's it. Interestingly, neither housing nor equity booms have a measurable impact on the volatility of prices. And equity booms do not affect the volatility of growth—the estimates are both economically tiny and statistically irrelevant.

Focusing on the bottom-left panel of the table 1.2—the impact of housing booms on GDP volatility—we see that the bigger the boom, the bigger the impact on volatility. But the bigger impact is at short horizons where we know from table 1.1 that, on average, growth rises. So, while housing booms increase volatility, it seems to do it primarily on the upside.

1.3.3 GDP and Price Level at Risk

Next, I turn to an examination of the tails of the distribution of output and price-level outcomes, conditional on asset price booms. Are GDP at

Table 1.2 **Impact of asset price booms on volatility (lag of asset price [*k*])**

	Volatility of the output gap			Price level volatility		
Threshold (α)	4	8	12	4	8	12
			Equity			
Data	0.00	0.01	0.00	0.00	0.00	0.00
	0.14	0.07	0.30	0.38	0.08	0.18
4	0.03	0.33	0.24	0.00	0.00	0.00
	0.43	0.08	0.14	0.15	0.54	0.46
12	0.15	0.05	0.19	0.00	0.00	0.00
	0.26	0.41	0.15	0.13	0.13	0.12
20	0.39	0.20	0.11	0.00	0.00	0.01
	0.12	0.27	0.34	0.16	0.11	0.06
			Housing			
Data	**0.22**	0.12	0.05	–0.02	0.00	0.02
	0.04	0.18	0.35	0.64	0.53	0.34
2	**2.46**	**2.84**	1.43	–0.09	0.50	0.66
	0.02	**0.01**	0.12	0.57	0.12	0.17
6	**4.39**	**4.75**	1.93	0.55	0.60	0.87
	0.01	**0.01**	0.12	0.18	0.10	0.16
10	**5.48**	2.46	**5.28**	0.88	0.38	0.80
	0.04	0.07	**0.04**	0.16	0.23	0.19

Notes: Table 1.2 reports the coefficient b_2 in the regression $(x_{it})^2 = a_2 + b_2\, \mathrm{d}_{it-k}(\alpha) + \eta_{it}$, where x is the deviation of either log GDP or the log price from a Hodrick-Prescott filtered trend, with parameter 1,600; and d is either a dummy variable equal to 1 if the filtered asset price exceeds the threshold α (in percent), or the filtered asset price data itself (those are the rows labeled "data"). In each case, the first row of numbers is the coefficient itself, while the second row is a p-value for the test that b_2 is strictly greater than 0, computed using Newey-West standard errors with lags equal to 1.5 times k. **Bold** values are significantly greater than 1 at the 5 percent level.

risk and price level at risk affected by the equity or housing booms or busts? If, for example, there is a dramatic increase in equity prices, should this change our view of the possibility of bad events? And, importantly, are normal approximations likely to give the wrong signal?

Equity Bubbles

For equity booms, the answer to this question is reported in figure 1.3. The horizontal axis in the figure plots the minimum size of the equity price deviation, and the vertical axis plots the 5th percentile of the distribution of future outcomes for the GDP gap—the 5 percent GDP at risk. The two lines show the 5 percent GDP at risk four quarters ahead and twelve quarters ahead. So, for example, if equity prices are at least 10 percent above trend, the 5th percentile of the distribution of the GDP gap twelve quarters into the future is –3.6. As it turns out, this is only slight below the 5th percentile of the unconditional distribution for deviations of GDP from trend, which is 3.44, so it isn't very troubling. In other words, the GDP at risk from

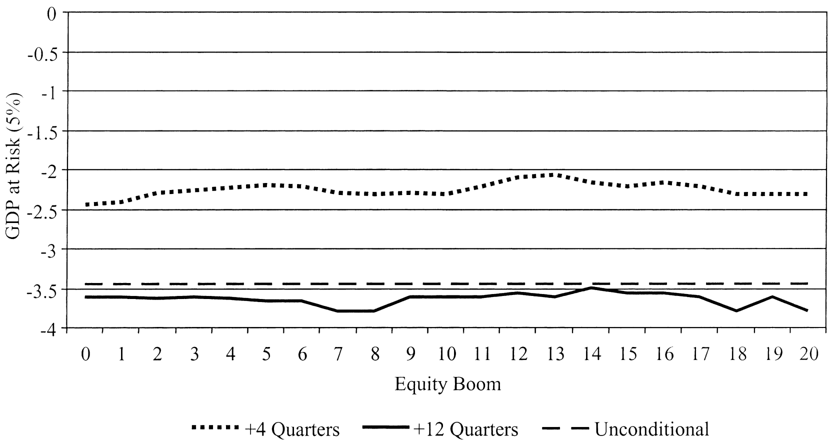

Fig. 1.3 GDP at risk following an equity boom

a 10 percent equity boom is only very slightly below the unconditional GDP at risk. The upper line in the figure, the 5 percent GDP at risk four quarters ahead, is always significantly *above* the unconditional 5th percentile of the GDP gap distribution. The reason for this is that all booms are likely to continue, so the horizon for the collapse of equity prices and GDP both is beyond four quarters.

Figure 1.4 reports the results for price level at risk following an equity boom. The price level at risk results differ quite a bit from the GDP at risk results. Since some central banks will care about prices rising while others may care more about prices falling, I report the risk results for both tails of the distribution. These are referred to as the 95 percent price level at risk. As the equity boom grows, the risk of the price level falling below trend (shown in panel A of figure 1.4) grows substantially. When real equity prices are 15 percent or more above trend, the 5th percentile of the distribution of price-level gap four quarters out is more than –9 percent. Depending on the current level of inflation, that could be a significant risk. By contrast, the risk of the extreme positive price-level gaps (in panel B of figure 1.4) goes down. Conditional on an equity boom, the distribution of price-level deviations from trend shifts down.

Housing Bubbles

Turning to housing bubbles, figures 1.5 and 1.6 report computations analogous to those reported in figure 1.3 and 1.4. The results in these two figures suggest that housing booms are followed by an increased risk of a large decline in GDP in four to twelve quarters and a decreased risk of prices falling below trend. Note from the scale that the GDP at risk is quite large. When real house prices are 5 percent or more above trend, there is a

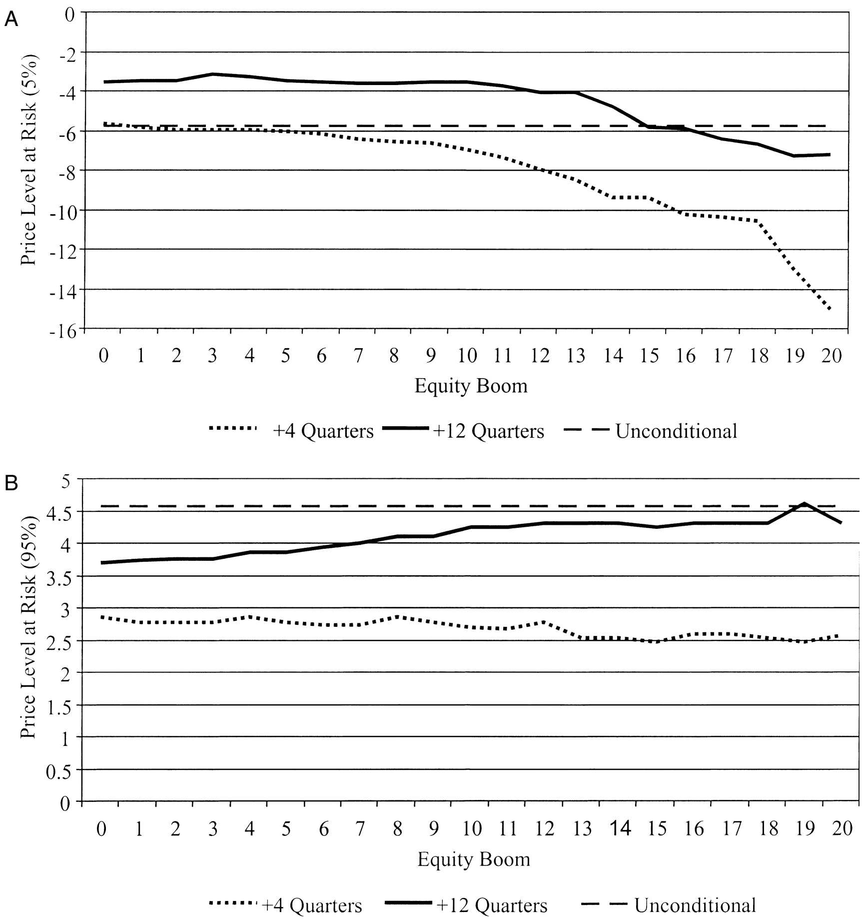

Fig. 1.4 Price level at risk following an equity boom: *A,* Risk of prices *falling* significantly *below* trend; *B,* Risk of prices *rising* significantly *above* trend

5 percent probability that twelve quarters later GDP will be at least 3.44 percent below trend—substantially below the unconditional 5th percentile of 2.86 percent.[16]

Housing booms affect the price level at risk as well. The information in figure 1.6 suggests that a housing boom has very little impact on the upper tail of the price-level distribution, but dramatically eliminates the lower tail—at least at a twelve-quarter horizon. Unconditionally, the upper tail 5 percent price level at risk twelve quarters following a 10 percent housing

16. Note that because the countries in the sample differ, the unconditional distributions for the price-level and GDP gaps are different between the equity and housing booms.

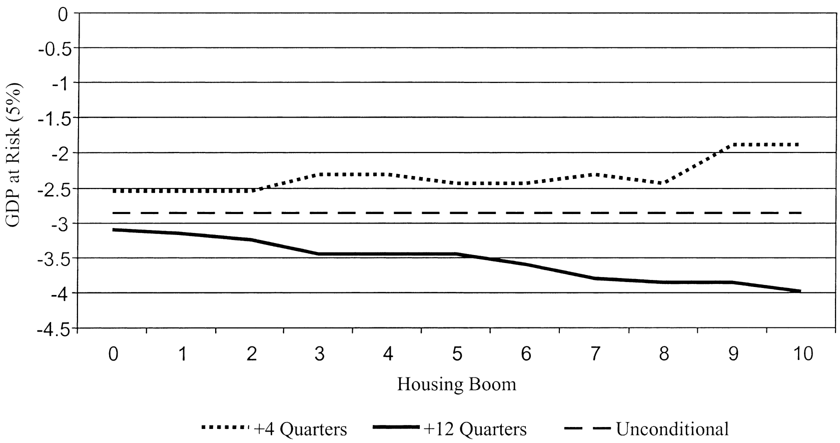

Fig. 1.5 GDP at risk following a housing boom

price boom is roughly one-quarter the unconditional 5th percentile—that is, it—1 percent as compared with –4 percent.

Comparing the Normal Approximation and the Empirical Density

It is important to ask whether there is any difference between the results in figures 1.3 and 1.5 and those from a simple normal approximation. That is, if a central banker had been looking at the –1.645 times the standard deviation of the distribution of output and price-level gaps, conditional on an equity market boom, would they have done anything differently? The results suggest that the answer to this is yes.

Figure 1.7 compares the 5th percentile for the GDP gap computed using a normal approximation with one from the empirical density. For equity booms, the normal approximation gives an overly pessimistic view of the size of the lower tail. The average distance between the two estimates of the 5th percentile of the distribution is roughly three-quarters of 1 percentage point. This particular example suggests that a policymaker using a quadratic loss would likely overestimate the importance of an equity boom.

Housing is another story. Here the normal distribution gives an overly optimistic view of the true size of the lower tail. The 5th percentile of the empirical density is, on average, 1.25 percentage points below what is implied by the normal approximation. Because the probability of extreme negative outcomes for the GDP gap is higher than suggested by a Gaussian distribution, policymakers focusing on quadratic loss will underestimate the importance of a housing boom.

In the case of price-level outcomes, normal approximations are also misleading. For example, twelve quarters following a housing boom, the 5th percentile of the upper tail of outcomes is 2.5 percentage points *smaller*

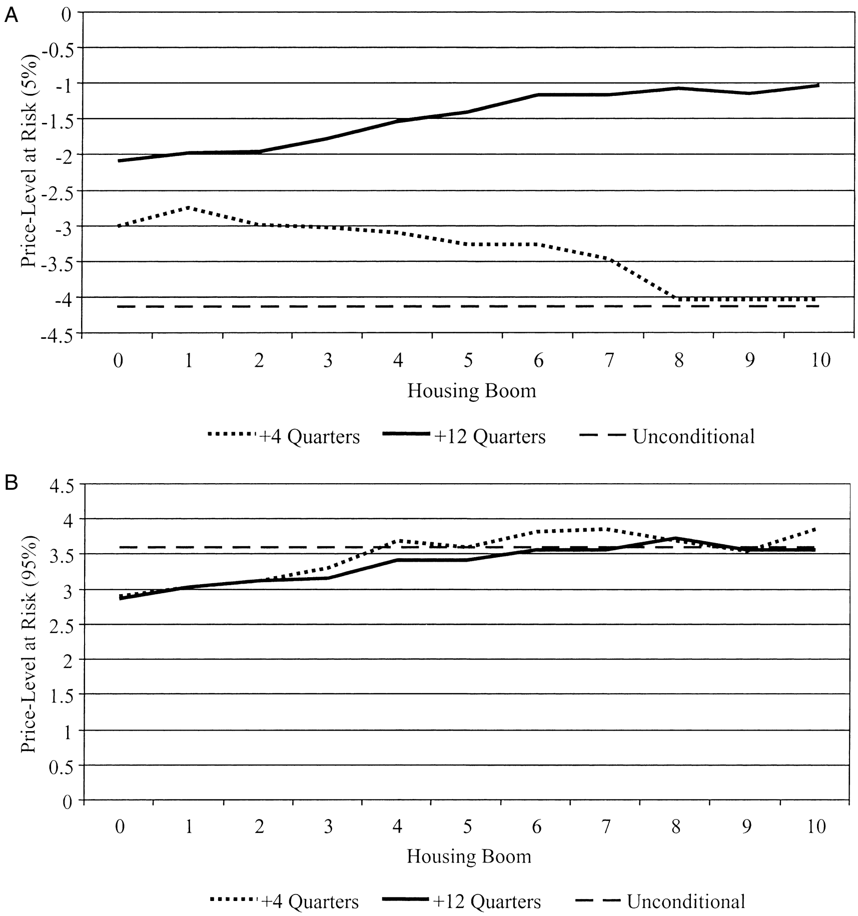

Fig. 1.6 Price level at risk following a housing boom: *A,* Risk of prices *falling* significantly *below* trend; *B,* Risk of prices *rising* significantly *above* trend

than would be implied by simply multiplying the standard deviation of the observed outcomes by 1.64.

1.3.4. Expected Lower Tail Loss

Direct statistical inference for a number like GDP at risk is difficult.[17] Instead of constructing Monte Carlo experiments that might allow confi-

17. Quantile regression, pioneered by Koenker (2005), is an alternative to the figures in section 1.3.3 and the regressions in section 1.3.4. Rather sorting the data based on arbitrarily chosen thresholds for the right-hand-side variables, quantile regression examines changes in the relationship based on the quantile of the regression residuals. Such a technique has the distinct advantage of allowing for the additional control variable in the regression. Future research will examine the robustness of these results to these alternative statistical methods.

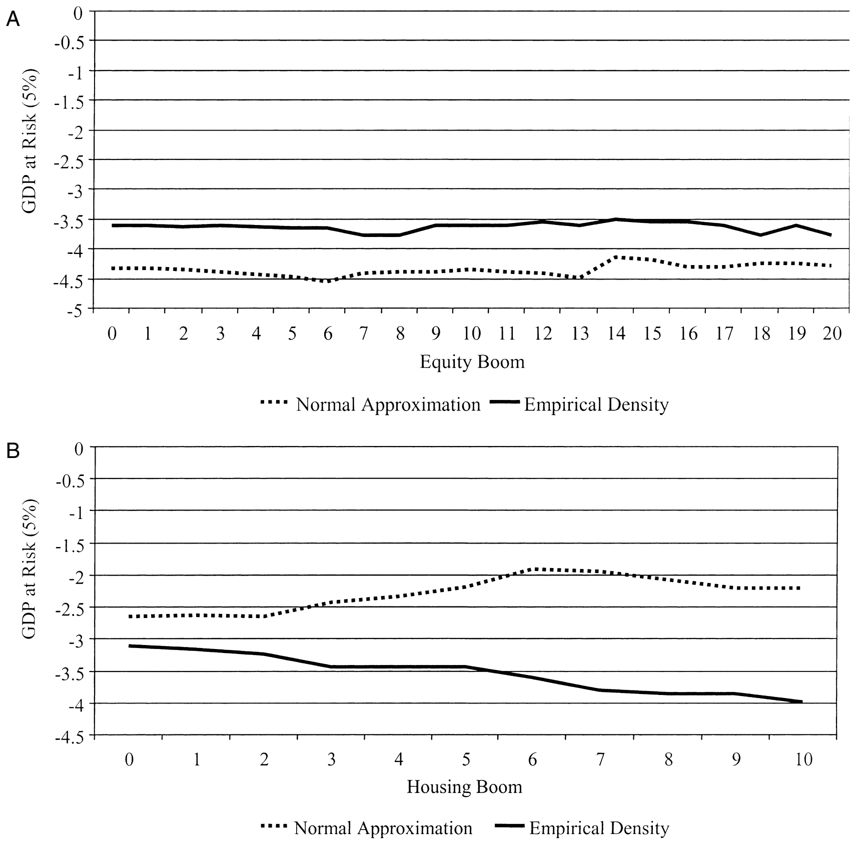

Fig. 1.7 Comparing the normal approximation with the empirical density GDP at risk at + twelve quarter horizon: *A,* Conditional on an equity boom; *B,* Conditional on a housing boom

dence interval estimation, I turn to the examination of the ETL. This is the expected value, conditional on being in the tail of the distribution. As in the case of the GDP at risk and price level at risk, here I ask whether the ETL changes when asset prices boom. In order to do inference, I run a regression similar to equation (1):

$$x_{it} = a + b_0 \mathrm{d}_{\mathrm{it-k}}(\alpha) + b_1 \mathrm{tail}(\beta)_{\mathrm{it}} + b_3\, \mathrm{d}_{\mathrm{it-k}}(\alpha) x \mathrm{tail}(\beta)_{\mathrm{it}} + \eta_{it}, \tag{3}$$

where x_{it} is the output or price-level gap; $\mathrm{d}_{\mathrm{it-k}}(\alpha)$ is a dummy variable equal to 1 if k periods earlier the filtered asset price data exceeds the threshold α; and $\mathrm{tail}(\beta)_{\mathrm{it}}$ is a dummy variable that equals 1 if x_{it} is in the β percent lower tail of the distribution of all x_{it}.

The coefficient b_3 on the interaction term in equation (3) provides an estimate of the impact of an asset price boom of size α on the ETL in the lowest β percent of the distribution of the output or price-level gap. Because of the structure of the regression, it is possible to compute standard errors that are robust to both serial correlation and heteroskedasiticity in the error term η_{it}.[18]

The results of this regression are reported in table 1.3, and they are quite striking. Asset price booms—both equity and housing—result in a fall in the expected lower tail loss. The decline is both economically and statistically significant. Put another way, equity and housing booms make it more likely that something bad will happen.

1.3.5 Summary of the Results

Table 1.4 summarizes the results of this section. The conclusion is that housing booms dramatically change the distribution of outcomes in virtually every way. By contrast, equity booms have little impact on the mean and variance of deviation from trend, but do affect the lower tail of the distribution.

1.4 The Difference between Equity and Housing Bubbles

To understand the differential impact of equity and housing bubbles, it is useful to focus on their consumption effects. Booms in either equity or property prices drive up the wealth of individuals. The natural response to an increase in wealth is to raise consumption. If you are rich, you can buy a fancy car, purchase a bigger and flatter television, go on nicer vacations, eat in expensive restaurants, and the like. And the data show that this is exactly what happens.

A useful rule of thumb is that a $1 increase in U.S. wealth generates between two and five cents of additional consumption by American households.[19] That is, the marginal propensity to consume for wealth is in the range of 0.02 to 0.05.

As Norman, Sebastia-Barriel, and Weeken (2002) note, the marginal propensity to consume is of somewhat less interest than the elasticity of consumption with respect to wealth.[20] They emphasize that we care more about the impact of a 10 percent increase in the value of wealth than we do about the number of cents or pence that consumption rises per dollar or pound of additional wealth. This is especially true of equity wealth because the size of equity markets vary so widely across countries. Bertaut (2002) reports that, at the end of 2001, total equity market capitalization

18. The estimation method is an adaptation of the Newey-West estimator to a panel in which there is serial correlation and heteroskedasiticity within a country, but no dependence between countries.

19. See, for example, Norman, Sebastia-Barriel, and Weeken (2002).

20. The elasticity of consumption with respect to wealth is equal to the marginal propensity to consume out of wealth times the ratio of wealth to consumption.

Table 1.3 **Impact of asset price booms on the lowest quartile (lag of asset price [*k*])**

	Output gap			Price-level gap		
Threshold (α)	4	8	12	4	8	12
			Equity			
Data	–0.03	–0.02	0.01	**–0.26**	–0.21	0.03
	0.10	0.12	0.63	**0.02**	0.06	0.59
4	**–3.81**	**–2.50**	**–1.87**	**–14.05**	**–16.12**	**–13.88**
	0.00	**0.00**	**0.00**	**0.00**	**0.00**	**0.00**
12	**–4.63**	**–1.75**	**–1.70**	**–16.38**	**–19.20**	**–16.35**
	0.00	**0.01**	**0.01**	**0.00**	**0.00**	**0.00**
20	**–5.37**	**–2.05**	**–0.85**	**–18.06**	**–20.73**	**–17.36**
	0.00	**0.02**	**0.05**	**0.00**	**0.00**	**0.00**
			Housing			
Data	–0.01	–0.03	–0.01	0.06	0.10	0.09
	0.35	0.11	0.35	0.67	0.84	0.76
2	**–1.53**	**–1.08**	**–0.69**	**–2.47**	**–4.03**	**–5.01**
	0.00	**0.00**	**0.00**	**0.00**	**0.00**	**0.01**
6	**–1.42**	**–1.15**	–0.28	**–2.89**	**–4.83**	**–9.22**
	0.00	**0.00**	0.18	**0.00**	**0.00**	**0.00**
10	**–1.16**	–0.34	–0.59	**–3.28**	**–3.91**	**–12.29**
	0.00	0.12	0.08	**0.00**	**0.00**	**0.00**

Notes: Table 1.3 reports the coefficient b_3 in the regression $x_{it} = a + b_1 \mathrm{d}_{\mathrm{it-k}}(\alpha) + b_2 \mathrm{tail}(\beta)_{it} + b_3 \mathrm{d}_{\mathrm{it-k}}(\alpha) x\ \mathrm{tail}_{it}(\beta) + \upsilon_{it}$, where x_{it} is the deviation of either log GDP or the log price from a Hodrick-Prescott filtered trend, with parameter 1,600; $\mathrm{tail}_{it}(\beta)$ is a dummy variable that equals 1 if x_{it} in the lower β-percent tail; and d is either a dummy variable equal to 1 if the filtered asset price exceeds the threshold α (in percent), or the filtered asset price data itself (those are the rows labeled "data"). In each case, the first row of numbers is the coefficient itself, while the second row is a *p*-value for the test that b_3 (the coefficient on the interaction term) is strictly less than 0, computed using Newey-West standard errors with lags equal to 1.5 times *k*. **Bold** values are significantly greater than 1 at the 5 percent level.

equaled 153 percent of GDP in the United Kingdom, but only 59 percent of GDP in Germany. To understand the importance of this, consider the impact of a 10 percent increase in equity prices on consumption in each country, assuming that the marginal propensity to consume is the same. The estimated impact in the United Kingdom the impact would be roughly three times as large as that in Germany.[21]

This highlights the importance of thinking about bubbles in housing and equity prices separately. There are two reasons for this. First, equity prices are substantially more volatile than housing prices, so the former is much less likely to be permanent than the latter. Reasonably, households respond more aggressively to changes in wealth that they perceive to be perma-

21. Careful econometric estimates show an even larger disparity. Bertuat (2002) reports that 10 percent increase in stock market creates 0.5 to 1.0 percent increase in consumption in the long run in the United States and United Kingdom, but only 0.07 in Germany where the equity is less than 60 percent of GDP.

Table 1.4 **Summary of the impact of asset price booms on the distribution of macroeconomic outcomes (lag of asset price)**

	Output gap		Price-level gap	
Moment	$k=4$	$k=12$	$k=4$	$k=12$
	Equity			
Mean	Higher	None	None	None
Variance	None	None	None	None
5% VaR	Better	None	None	Worse
25% Expected Tail Loss	Lower	Lower	Lower	Lower
	Housing			
Mean	Higher	Lower	Higher	Higher
Variance	Higher	Higher	None	None
5% VaR	Better	Worse	Better	None
25% Expected Tail Loss	Lower	Lower	Lower	Lower

Note: Table 1.4 summarizes the results in tables 1.1, 1.2, and 1.3 and figures 1.3 to 1.6.

nent.[22] Second, equity ownership tends to be concentrated among the wealthy—people who are much less likely to adjust their consumption levels. Housing ownership, by contrast, is distributed more broadly. And while the quality of housing and the concentration of ownership vary across countries, the differences are far less dramatic.

Returning to the evidence, using data from fourteen developed countries, Case, Quigley, and Shiller (2005) discuss how a 1 percent increase in housing wealth raises consumption by between 0.11 and 0.17 percent. By contrast, they find that the stock market wealth elasticity of consumption is substantially smaller, only 0.02. It is natural that the housing booms would have more of an impact on the distribution of macroeconomic outcomes than equity booms do.

1.5 Policy Responses: Risk Management and Financial Structure

Is there anything to be done about all of this? Can we provide any useful guidance on how to avoid the risks bubbles pose? Researchers have investigated myriad possible responses, including, but not restricted to, reacting only to bubbles insofar as they influence inflation forecasts; reacting only to the fallout of a bubble after it bursts; leaning against a bubble as it develops; including asset prices in the price index central bankers target; and examining various regulator solutions involving margin and lending requirements. In Cecchetti (2006), I summarize the traditional debate in

22. Kishor (2005) estimates that while 98 percent of the change in housing wealth is permanent, only 55 percent of the change in financial wealth is. This suggests that the housing wealth effect should be roughly twice the stock market wealth effect.

each of these cases. Briefly, there is a consensus building against the purely activist view. As Gruen, Plumb, and Stone (2005) discuss, the information requirements for the activism are fairly high, and thereare significant risks of costly missteps. The conclusion is that interest rates should play only a modest role in combating the destabilizing effects of asset price bubbles

From a risk management perspective, the discussion of central bank responses to asset price bubbles is unnecessarily restrictive. Why focus only on traditional monetary policy tools? Risk managers do more than simply monitor and react to developments; they build institutional structures that are unlikely to collapse when hit by large shocks. The regulators and supervisors of the financial system have built mechanisms exactly like this. Are there similar responses to bubbles? When subjected to equity and property price bubbles, are some financial systems more resilient than others?

Recent work by Dynan, Elmendorf and Sichel (2006) and Cecchetti, Flores-Lagunes, and Krause (2006) suggests that changes in the financial system have been an important source of stabilization over the past several decades. Their results suggest that enhanced household access to credit allows for increased consumption smoothing that has been a major factor in reducing the volatility of aggregate real growth.[23] This brings up the natural question: does the impact of housing and equity bubbles on GDP at risk or price level at risk depend on financial structure?

To examine this, I begin with data on financial structure taken from Demirguc-Kunt and Levine (2001). Briefly, Demirguc-Kunt and Levine have constructed a data set on financial indicators during the 1990s covering a broad cross section of countries. Included are measures of the relative size of a country's stock market and banking sector, as well as a measure of the relative efficiency of the two. Countries with "market-based financial systems" are those with bigger more efficient stock markets. I examine the relationship of this composite financial structure index and the behavior of an economy following booms in equity or housing prices.

As a first step, I reproduce figures 1.3 and 1.5 with the data for GDP at risk dividing the data based on whether it comes from a country with a predominantly market-based or bank-based financial system. The results, reported in figure 1.8, show that for countries where equity markets are important, equity booms increase GDP at risk. By contrast, GDP at risk following a housing boom is not sensitive to financial structure as characterized by this index.

To examine this a bit further, and to try to get a grasp on whether any of it is precise in a statistical sense, I add the financial structure variable to re-

23. The argument is that there is a linkage not only between financial system development and the *level* of real growth, as described in Ross Levine's (1997) survey, but also between financial development and the stability of real growth.

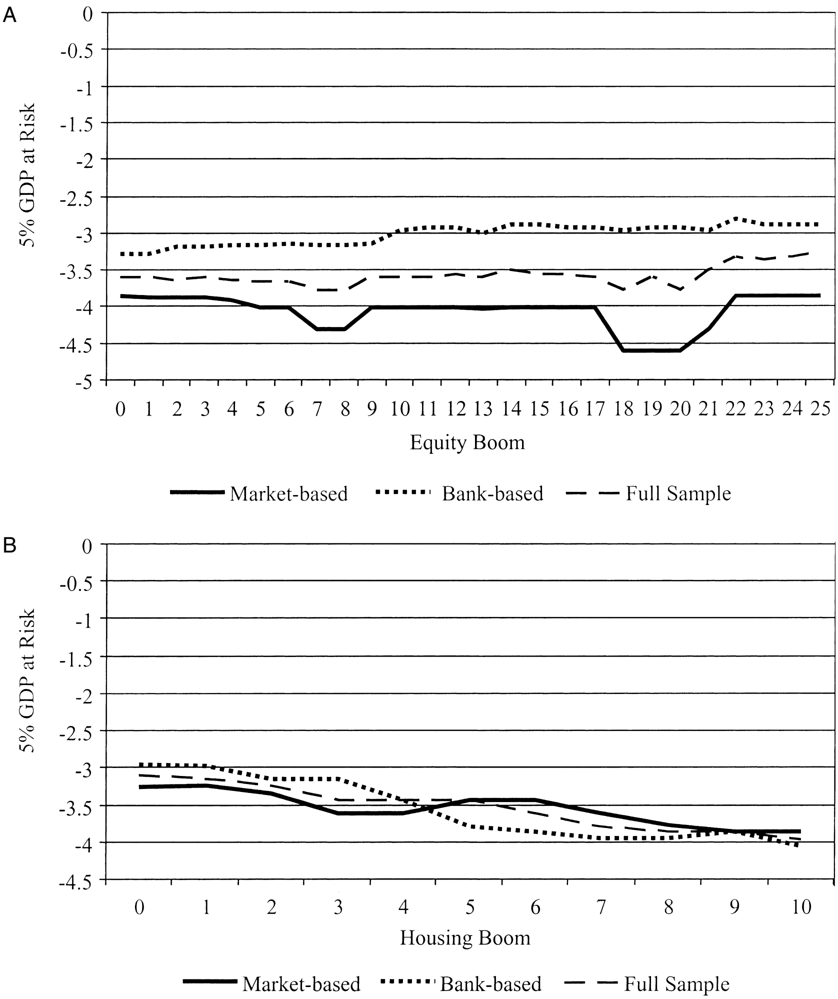

Fig. 1.8 Market- versus bank-based financial systems GDP at risk at + twelve quarter horizon: *A*, Equity booms; *B*, Housing booms

gressions (1), (2) and (3)—both as a level and interacted with the asset price boom dummy. Here's an example:

(1′) $x_{it} = a + b\mathrm{d}_{\mathrm{it-k}}(\alpha) + cf_i + df_i\mathrm{d}_{\mathrm{it-k}}(\alpha) + \varepsilon_{it},$

(2′) $(x_{it})^2 = a' + b'\mathrm{d}_{\mathrm{it-k}}(\alpha) + c'f_i + d'f_i\mathrm{d}_{\mathrm{it-k}}(\alpha) + \upsilon_{it}.$

(3′) $x_{it} = a + b_0\mathrm{d}_{\mathrm{it-k}}(\alpha) + b_1\mathrm{tail}(\beta)_{\mathrm{it}} + b_3\mathrm{d}_{\mathrm{it-k}}(\alpha)x\mathrm{tail}(\beta)_{\mathrm{it}}$

$b_4f_i + b_5f_i\mathrm{d}_{\mathrm{it-k}}(\alpha) + b_6f_i\mathrm{tail}(\beta)_{\mathrm{it}} + b_7f_i\mathrm{d}_{\mathrm{it-k}}(\alpha)x\mathrm{tail}(\beta)_{\mathrm{it}} + \eta_{it},$

where f_i is the composite structure index from the CD-ROM that is distributed with Demirguc-Kunt and Levine (2001).[24]

Table 1.5 reports the estimated coefficients on the interactions terms in each of these: $f_i \mathrm{d}_{\mathrm{it-k}}(\alpha)$ in equations (1′) and (2′), and $f_i \mathrm{d}_{\mathrm{it-k}}(\alpha) x\ \mathrm{tail}(\beta)_{\mathrm{it}}$ in equation (3′). These tell us whether differences in financial structure change the impact of an asset price boom on the mean, variance, or lower tail events in the distribution of the output gap. I report the results for a lag of four and twelve quarters. The financial structure index is positive for market-based economies and negative for bank-based ones. For example, it takes on a value of +0.17 for the United States and –0.18 for Greece.

Unsurprisingly, the strongest results are those for the mean. In countries with market-based financial systems, which is to say places where equity markets are important, the first and second column of the top panel in table 1.5 shows that equity price increases lead to bigger short-horizon booms and bigger long-horizon crashes (although the latter are imprecisely estimated). Analogously, for bank-based economies, housing booms lead to bigger short-horizon GDP booms, but smaller long-horizon crashes. (These are the results in the first and second column of the bottom panel of the table.)

Turning to the volatility, there is no measurable impact on financial structure. The point estimates reported in the fourth and fifth columns of table 1.5 are all small and the p-values are never below 0.2 or above 0.8.

Finally, looking at the far right columns of table 1.5, the results from estimating equation (3′), there is some weak evidence that market-based economies fare somewhat worse at longer horizons when hit with equity price booms. Again, this is really no surprise.

In the end, these results are disappointing. While we may believe that financial structure plays role in the real economic impact of asset price booms, the data available do not show much evidence of it.

1.6 Conclusion

Stability is the watchword for central bankers. Listen to most modern monetary policymakers speak about their goals, and you are likely to hear about the desire for low, stable inflation and high, stable growth. They will explain how they raise and lower their short-term interest rate target in order to meet their stability-oriented objectives. But listen closely, and you

24. The index average of deviations from the mean of (1) stock market capitalization divided by deposit money bank assets (relative size of stock market compared to banking sector), (2) total value traded in stock market divided by claims on private sector by deposit money banks (relative activity of stock market compared to banking sector), and (3) total value traded in stock market as a share of GDP divided by banking overhead costs as a share of total assets (relative efficiency of stock market compared to banking sector). The actual data are column EQ in the file called "request8095.xls." These data are the same as those.

Table 1.5 **Financial structure and the impact of asset price booms (lag of asset price in quarters)**

	Mean		Variance		Lowest Quartile	
Threshold (α)	4	12	4	12	4	12
			Equity			
Data	*0.07*	–0.03	0.00	0.00	0.06	**–0.07**
	1.00	0.15	0.62	0.32	0.84	**0.10**
8	*1.50*	–0.60	0.01	–0.01	1.41	–1.20
	0.95	0.31	0.31	0.65	0.69	0.25
16	*2.31*	–0.87	0.01	0.00	3.68	–2.03
	0.97	0.26	0.30	0.45	0.74	0.11
20	*2.88*	–1.82	0.00	0.02	5.25	**–3.57**
	0.98	0.12	0.42	0.19	0.78	**0.06**
			Housing			
	4	12	4	12	4	12
Data	**–0.14**	–0.04	0.00	0.00	–0.02	0.05
	0.07	0.36	0.65	0.54	0.44	0.64
4	–1.40	–1.60	0.03	–0.03	–1.06	0.59
	0.16	0.14	0.37	0.63	0.21	0.65
8	–1.90	–0.02	–0.09	–0.03	–0.50	0.19
	0.21	0.50	0.71	0.59	0.40	0.54
10	–0.83	0.04	–0.20	–0.04	–0.10	0.89
	0.38	0.51	0.77	0.58	0.49	0.61

Notes: Table 1.5 reports the regression coefficients from the interaction of the financial structure measure with the asset price boom dummy variable in equations (1′), (2′), and (3′). The more positive financial structure, the more market-based a country's financial system; the more negative, the more bank-based it is. In each case, the first row of numbers is the coefficient itself, while the second row is a p-value for the test that is strictly less than zero, computed using Newey-West standard errors with lags equal to 1.5 times k. *Italicized* values are significantly greater than zero at the 5 percent level and boldfaced values are significantly less than zero at the 10 percent level.

will realize that the statements are more nuanced. While stability is the ultimate objective, it is the possibility of catastrophe that keeps central bankers awake at night. They want to ensure that nothing really bad happens, and to do this, they are looking at the entire distribution of possible outcomes.

In analyzing the macroeconomic impact of asset price booms and crashes, it is the disasters that are the true concern. This suggests a different approach to risk, one based on keeping the probability of output deviating from its trend (or price level deviations from its target trend) over some time horizon below some fixed threshold. Policy responses should be built in order to keep the lower tail of the distribution—as measured by value-at-risk or the ETL—sufficiently small.

In this chapter, I use data from a broad cross section of countries to examine the mean, variance, and lower tail risks arising from booms and crashes in equity and housing markets. The conclusion is that housing bubbles change the entire distribution of macroeconomic outcomes. By contrast, equity bubbles tend to make the worst events even worse, leaving the mean and variance of the distributions roughly unchanged. The strong conclusion is that approximations that use the normal distribution, and analyses based on quadratic loss functions, have the potential to be extremely misleading. Looking further, I present weak evidence suggesting that those countries with market-based financial systems, where stock market capitalization is relatively large, weather housing booms somewhat better and equity booms somewhat worse than countries with bank-based financial systems.

In closing, it is important to emphasize one critical implication of adopting a risk management view. As mentioned earlier, econometric modeling tends to provide characterizations of what happens near the mean of the data. In fact, in order to improve the quality of estimates, researchers have a tendency to remove outliers. This is sometimes done in the guise of sensitivity analysis and other times using limited-influence estimation that explicitly truncates tail observations. This means that standard modeling strategies provide virtually no information about the behavior of the economy when it is under stress. As a result, evaluating the problems posed by extreme events, which is at the core of risk management, necessarily requires judgment. And to quote Chairman Greenspan (2004) one final time: "Such judgments, by their nature, are based on bits and pieces of history that cannot formally be associated with an analysis of variance."

Appendix

Data Appendix

Price Data

Price data computed for consumer price inflation data was obtained from the *International Financial Statistics* online and the OECD Economic Outlook no. 76, December 2004.

GDP

Gross domestic product data was obtained from the *International Financial Statistics* CD-ROM (December 2004) and the OECD Economic Outlook no. 76, December 2004.

Equity Prices

Equity prices are from the *International Financial Statistics* online.

Housing Prices

Data for Australia, Belgium, Canada, Denmark, Finland, Ireland, the Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, the United Kingdom, and United States are all from the BIS. Data for Hong Kong are from the Hong Kong Monetary Authority, Census and Statistics Department, Monthly Digest of Statistics, table 5.9, column (6). Data for Israel are from the Israel Central Bureau of Statistics, online. Data for Japan are from Goldman Sachs. Data for New Zealand are from the Reserve Bank of New Zealand.

References

Bertaut, Carol. 2002. Equity prices, household wealth, and consumption growth in foreign industrial countries. International Finance Discussion Paper no. 724. Washington: Board of Governors of the Federal Reserve System.

Case, Karl E., John M. Quigley, and Robert J. Shiller. 2005. Comparing wealth effects: The stock market versus the housing market. *Advances in Macroeconomics* 5 (1). Article 1.

Cecchetti, Stephen G. 2003. What the FOMC says and does when the stock market booms. In *Asset prices and monetary policy*, ed. A. Richards and T. Robinson, 77–96. Proceedings of the Research Conference of the Reserve Bank of Australia. Canberra: McMillan.

———. 2006. The brave new world of central banking: Policy challenges posed by asset price booms and busts. *Economic Review of the National Institute of Economic and Social Research* 196 (May): 107–19.

Cecchetti, Stephen G., Alfonso Flores-Lagunes, and Stefan Krause. 2006. Assessing the sources of changes in the volatility of real output. *The Economics Journal* 116 (4): 408–33.

Demirguc-Kunt, Asli, and Ross Levine. 2001. Bank-based and market-based financial systems: Cross-country comparisons. In *Financial structure and economic growth,* ed. A. Demirguc-Kunt and R. Levine, 81–140. Cambridge, MA: MIT Press.

Dennis, Richard, Kai Leitemo, and Ulf Söderström. 2006. Methods for robust control. IGIER-Universitá Bocconi Working Paper no. 307. Milan: Innocenzo Gasparini Institute for Economic Research.

Dynan, Karen E., Douglas W. Elmendorf, and Daniel E. Sichel. 2006. Can financial innovation explain the reduced volatility of economic activity? *Journal of Monetary Economics* 53 (1): 123–50.

Greenspan, Alan. 2003. Opening remarks. In *Monetary policy and uncertainty: Adapting to a changing economy,* ed. Gordon H. Sellon, Jr.: 1–7. Proceedings of the Federal Reserve Bank of Kansas City Symposium. Kansas City, MO: Federal Reserve Bank of Kansas City.

———. 2004. Risk and uncertainty in monetary policy. *American Economic Review* 94 (2): 33–40.

Gruen, David, Michael Plumb, and Andrew Stone. 2005. How should monetary policy respond to asset price bubbles. *International Journal of Central Banking* 1 (3): 1–31.

Hodrick, Robert J., and Edward C. Prescott. 1997. Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit, and Banking* 29 (1): 1–16.
International Monetary Fund (IMF). 2003. *World economic outlook.* Washington, DC: IMF, April.
Jorion, Phillipe. 2001. *Value at risk.* 2nd edition. New York: McGraw-Hill.
Kishor, N. Kundan. 2005. Does consumption respond more to housing wealth than to financial wealth? If so, why? University of Washington, Department of Economics. Unpublished Manuscript.
Koenker, Roger. 2005. *Quantile regression.* Cambridge, UK: Cambridge University Press.
LeBaron, Blake, and Ritirupa Samanta. 2005. Extreme value theory and fat tails in equity markets. Brandeis University, International Business School. Unpublished Manuscript.
Levine, Ross. 1997. Financial development and economic growth: Views and agenda. *Journal of Economic Literature* 35 (2): 688–726.
Newey, Whitney K., and Kenneth D. West. 1987. A simple, positive definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55 (3): 703–8.
Norman, Ben, Maria Sebastia-Barriel, and Olaf Weeken. 2002. Equity wealth and consumption—The experience of Germany, France and Italy in an international context. *Bank of England Quarterly Review* 42 (Spring): 78–85.
Onatski, Alexei, and James H. Stock. 2002. Robust monetary policy under model uncertainty in a small model of the U.S. economy. *Macroeconomic Dynamics* 6 (1): 85–110.
Rudebusch, Glenn D., and Lars E.O. Svensson. 1999. Policy rules for inflation targeting. In *Monetary policy rules,* ed. J. B. Taylor, 203–46. Chicago: University of Chicago Press.
Surico, Paolo. Forthcoming. U.S. monetary policy rules: The case for asymmetric preferences. *Journal of Economic Dynamics and Control.*
Svensson, Lars E.O. 2007. Robust control made simple: Lecture notes. Princeton University. Unpublished Manuscript.

Comment Andrew Levin

This chapter addresses a crucial topic for monetary policymakers, namely, does an asset price boom substantially raise the likelihood of a subsequent macroeconomic crisis? In this context, the chapter introduces the terms *gross domestic product* (*GDP*) *at risk* and *price level at risk* to characterize the lower tail of the distribution of each variable and then seeks to quantify

Andrew Levin is assistant director of the Division of Monetary Affairs and chief of the Monetary Studies Section at the Board of Governors of the Federal Reserve System.

I am grateful to Steve Cecchetti for providing me with the data set used in his analysis. I also appreciate helpful discussions with Morris Davis, Benson Durham, and Roberto Perli and outstanding research assistance provided by Arshia Burney and Ben Johannsen. The views expressed in this comment are solely those of the author and should not be interpreted as representing the views of the Board of Governors of the Federal Reserve System nor of anyone else associated with the Federal Reserve System.

the extent to which these risks are exacerbated by booms in either equity prices or house prices. While it would be ideal if one could consider the marginal impact on truly extreme events (such as the U.S. Great Depression of the 1930s), broad indexes of asset prices are only available for a substantial cross section of industrial economies over the post-1970 period; thus, the results reported here reflect the incidence and severity of the various recessions that actually occurred within the sample. Despite these statistical challenges, the chapter obtains significant evidence that a boom in house prices is associated with a subsequent reduction in real economic activity.

Measuring GDP at Risk

The analysis of the chapter begins by presenting evidence that output fluctuations exhibit a heavy lower tail, similar to the distribution commonly observed for equity prices. In particular, for eleven of the seventeen countries under consideration, the Jarque-Bera (*J-B*) test rejects the null hypothesis of a Gaussian distribution at a confidence level of 95 percent, and density approximations based on a *t*-distribution imply a distinctly larger magnitude of output contractions at the bottom 5th percentile of the distribution. Nevertheless, several important issues should be considered in interpreting these results.

Positive versus Negative Outliers

The *J-B* test is designed to detect skewness or excess kurtosis but does not necessarily indicate a heavy *lower* tail of the distribution. Indeed, as shown by the histograms in figure 1C.1 of this comment, positive outliers account for six cases in which the *J-B* test rejects the Gaussian null hypothesis. These outliers are concentrated in the early 1970s for four countries (namely, Italy, Japan, Switzerland, and the United Kingdom), while the Finnish outliers are associated with the 1989 to 1990 boom, and the German outliers correspond to the postreunification period of 1991 to 1992. The incidence of positive outliers also underscores the challenges in constructing measures of the output gap in the absence of any structural model; for example, a sequence of positive outliers in Hodrick-Prescott (HP)-detrended GDP could reflect either a cyclical boom or a spurt in potential output.

Transitory versus Persistent Outliers

Figure 1C.2 of this comment depicts histograms for the remaining five countries for which the *J-B* test rejects the null hypothesis of a Gaussian distribution for detrended output. Even in these cases, it is important to distinguish instances of deep recession—that is, lasting several quarters or more—from transitory fluctuations that might reflect a brief period of political turmoil or natural disaster. For example, the two outliers for

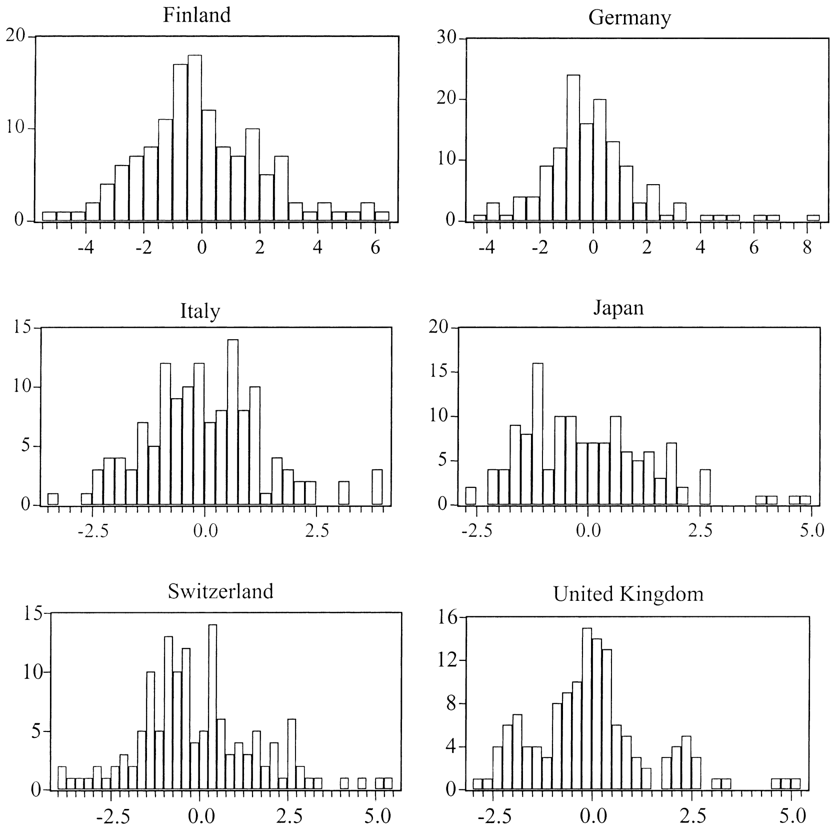

Fig. 1C.1 Positive versus negative outliers in detrended output

Note: For each country, the x-axis indicates the range of values of the HP-detrended output gap over the period 1970Q1 to 2003Q4, while the y-axis indicates the relative frequency of outcomes.

Australia reflect a single sharp recession that lasted from late 1982 through the end of 1983, whereas the data for Greece contains a single isolated outlier in the third quarter of 1974. The detrended output series for New Zealand also exhibits transitory outliers—both positive and negative—during 1973 to 1974, presumably reflecting the impact of highly volatile commodity prices. By contrast, the Netherlands's output trajectory has been remarkably stable over the past three decades, with a standard deviation of only 1.25 percent; in this case, one negative "outlier" reflects a single period in 1979, while the other two occurred during a more persistent contraction (lasting about a year) in 1982 to 1983. Finally, the cluster of negative outliers in the U.S. data correspond to the recession of 1981 to

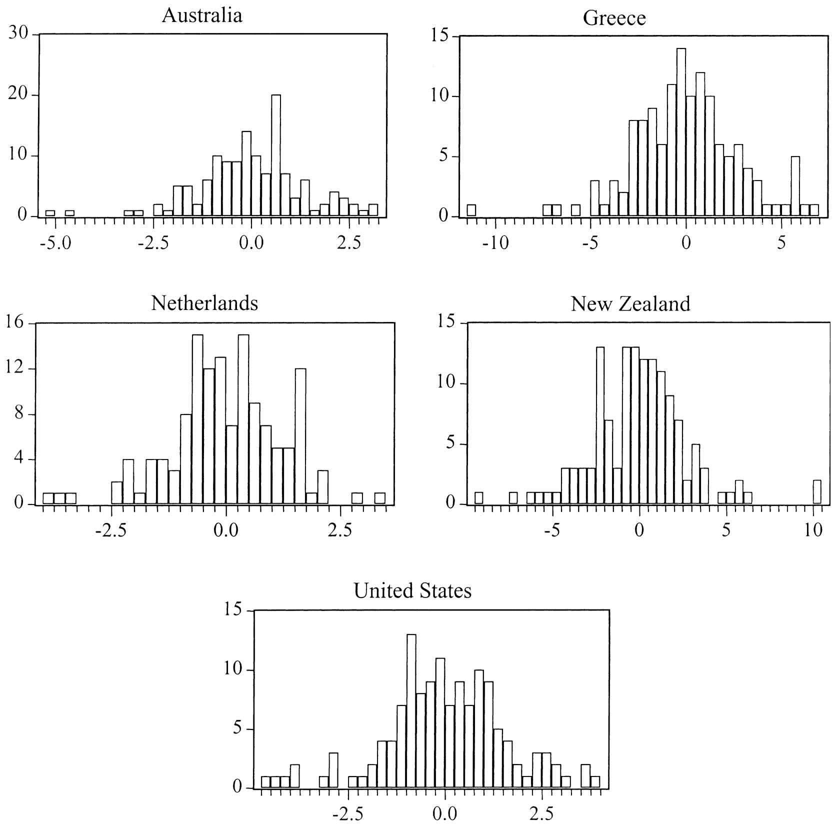

Fig. 1C.2 Transitory versus persistent outliers in detrended output

Note: For each country, the x-axis indicates the range of values of the HP-detrended output gap over the period 1970Q1 to 2003Q4, while the y-axis indicates the relative frequency of outcomes.

1983, which continues to be the largest contraction in U.S. economic activity since the Great Depression of the 1930s.

The Impact of House Price Booms

The chapter obtains substantial evidence that house price booms—measured by the deviation of each country's aggregate house price index from its HP-filtered trend—are associated with subsequent reductions in real economic activity. To shed further light on these results, it is helpful to focus on the specification that yields the highest level of statistical significance, namely, the extent to which a house price boom of at least 10 percentage points is associated with a substantial decline in the HP-detrended output gap twelve quarters later. The panel data set contains 103 observations (encompassing thirteen of the seventeen countries for which the relevant data is available) that satisfy this threshold for a house price boom, while the remaining control group of nearly 1,600 observations can be used to compute the distribution of output gaps that are *not* preceded by a house price boom at a twelve-quarter horizon.

Summary statistics regarding the distribution of output gaps—conditional on either the presence or absence of a house price boom twelve quarters earlier—are reported in table 1C.1 of this comment. Evidently, the occurrence of a house price boom systematically reduces the subsequent level of output: the mean of the conditional distribution is shifted downward by about 1.4 percent, a value that matches the regression estimate reported for this specification in table 1.1 of the chapter. The magnitude of this decline is also virtually identical to the impact shown in figure 1.5 of the chapter, which depicts the bottom 5th percentile of the unconditional distribution of output gaps in comparison with the same percentile conditional on the existence of a house price boom twelve quarters earlier.

Nevertheless, these results indicate that house price booms are *not* associated with substantially greater dispersion in the subsequent path of output. In particular, as shown in figure 1C.3 of this comment, the occurrence of a boom causes a downward shift in the entire distribution for detrended output but does not induce a heavier tail of adverse outcomes. This visual impression is confirmed by the summary statistics in the table: a house price boom has only modest effects on the standard deviation and the degree of skewness, while the degree of excess kurtosis is not affected at all. Thus, while house price booms do seem to generate some downside risk for the subsequent path of output, the magnitude of this risk appears to be fairly limited, relative to the effects of other macroeconomic disturbances.

Table 1C.1 The distribution of detrended output conditional on a house-price boom twelve quarters earlier

	No boom (1)	Boom (2)
No. of observations	1581	103
Mean	0.1	–1.3
Median	0.1	–1.2
Standard deviation	1.7	1.9
Skewness	0.4	–0.4
Excess Kurtosis	3.2	3.1

Notes: Table 1C.1 reports summary statistics regarding the conditional distribution of Hodrick-Prescott detrended output gaps for seventeen industrial countries over the period 1970Q1 to 2003Q4. Column (1) provides results for the set of observations that are preceded by a house-price boom twelve quarters earlier, while column (2) reflects all other observations.

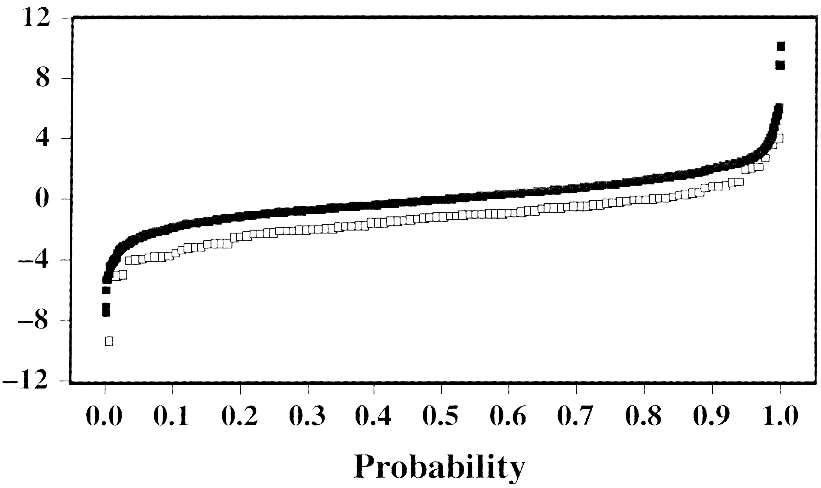

Fig. 1C.3 Housing booms and the distribution of output gaps

Note: This figure depicts the cumulative distribution of HP-detrended output gaps for seventeen industrial countries over the period 1970Q1 to 2003Q4. The hollow boxes denote this distribution for observations that are preceded by a house price boom twelve quarters earlier, while the solid boxes denote the distribution for all other observations.

Simulations of the Federal Reserve Board (FRB)/US Model

While these results reflect statistical patterns for a broad panel data set of industrial economies, simulations of the FRB/US model—which has been estimated using U.S. aggregate data and is used in ongoing policy analysis at the Federal Reserve Board—yield very similar implications regarding the macroeconomic effects of a sharp drop in house prices.[1] For example, figure 1C.4 of this comment depicts a scenario in which the aggregate house price index falls 10 percent during the second half of 2006 and by an additional 5 percent during 2007. As a result, U.S. real GDP declines to about 0.75 percent below baseline by the end of 2008.[2] The shaded regions in the figure indicate 70 and 90 percent confidence intervals obtained from stochastic simulations of the model, with shocks drawn from the set of estimated residuals over the period 1988 to 2004; these confidence intervals highlight the extent to which the model implies that even a steep drop in U.S. house prices would only have modest consequences for real GDP, at least in the absence of any other major disturbances.

1. See Brayton et al. (1997) for an overview of the specification and empirical properties of the FRB/US model.

2. Although not shown in the figure, the scenario assumes that movements in the federal funds rate are determined by Taylor's rule, which prescribes a gradual reduction to about 100 basis points below baseline by the end of 2008.

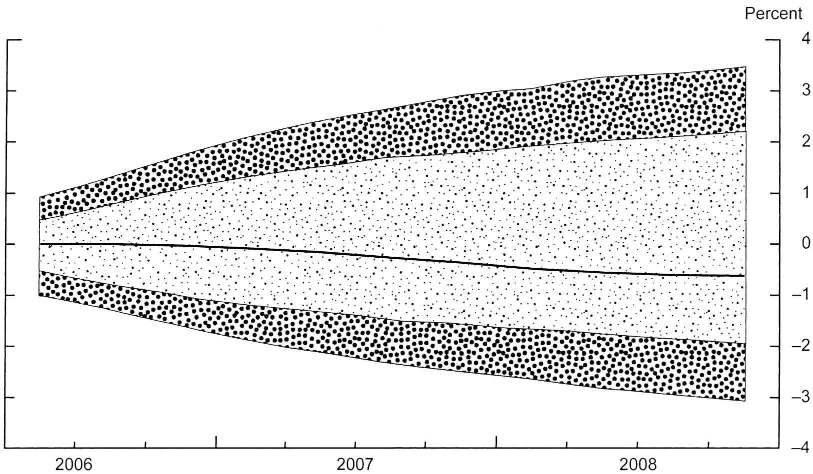

Fig. 1C.4 Simulations of the FRB/US model

Note: This figure depicts an FRB/US model simulation of a scenario in which the aggregate house price index declines by 10 percent during the second half of 2006 and an additional 5 percent during 2007; the response of U.S. real GDP (relative to baseline) is indicated by the solid line, while the shaded regions denote 70 and 90 percent confidence intervals obtained from stochastic simulations of the model, with shocks drawn from the set of estimated residuals over the period 1988 to 2004.

Real-Time Assessment of House Price Fluctuations

Finally, while the cross-country empirical analysis of this chapter has been conducted using HP-filtered aggregate house price indixes, it should be noted that financial market indicators may be useful for providing real-time information about market perceptions regarding the likelihood of a sharp decline in house prices.[3] For example, figure 1C.5 of this comment depicts the evolution of KMV-Moody measures of one-year-ahead expected default probabilities for twelve U.S. homebuilding firms over the period 1990 to 2006; while providing an early signal of downside risks to the residential construction industry prior to the onset of each of the past two recessions, this measure has not given any recent indications of a substantial near-term probability of a collapse in the housing market. This outlook is consistent with recently available information from housing futures and options, which began trading on the Chicago Mercantile Exchange in late May 2006 and suggested that market participants were anticipating a

3. Durham (2006) provides detailed analysis of asset prices related to the U.S. homebuilder industry along with an overview of Chicago Mercantile Exchange (CME) housing futures and options, while Campbell et al. (2006) analyze the extent to which recent trends in U.S. house prices can be interpreted in terms of movements in rents, real interest rates, and risk premia.

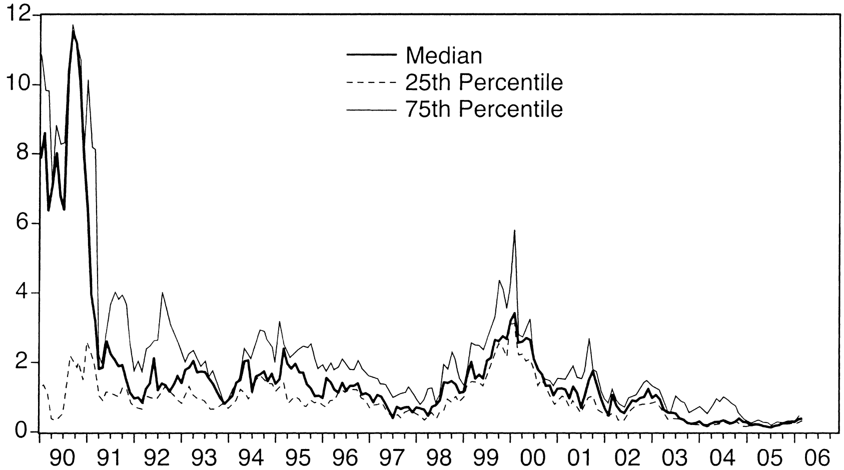

Fig. 1C.5 Expected default probabilities in the U.S. residential construction industry
Note: This figure depicts the quartiles of the distribution of one-year-ahead expected default probabilities (as estimated by KMV-Moodys) for twelve large publicly traded U.S. firms in the residential construction industry.

sharp slowing in the growth rate but apparently not a substantial decline in the level of U.S. house prices over the subsequent few quarters.

References

Brayton, F., A. Levin, R. Tryon, and J. Williams. 1997. The evolution of macro models at the Federal Reserve Board. *Carnegie-Rochester Conference Series on Public Policy* 42:115–67.

Campbell, S., M. Davis, J. Gallin, and R. Martin. 2006. What moves housing markets: A trend and variance decomposition of the rent-price ratio. Finance and Economics Discussion Series Paper no. 2006-29. Washington, DC: Board of Governors of the Federal Reserve System.

Durham, J. 2006. What do financial asset prices say about the housing market? Finance and Economics Discussion Series Paper no. 2006-32. Washington, DC: Board of Governors of the Federal Reserve System.

Discussion Summary

Lars E.O. Svensson questioned whether Alan Greenspan's "risk-management approach to policy" need be associated with special treatment of extreme events. Svensson suggested that everything Greenspan had said or written was consistent with Bayesian minimization of an expected loss function. In thinking about low-probability extreme events, it was not clear that you should ignore information from the rest of the distribution of events. Standard quadratic loss functions, for example, penalize bad outcomes. In fact, Charles Goodhart had even argued for an absolute-value loss function on the grounds that quadratic loss penalizes extreme outcomes too severely. For the same reason, Margaret Bray and Goodhart had argued for loss functions that are bounded from above.

John Williams said that he did think there was a useful distinction between risk management and minimizing expected loss and that he thought that Greenspan had distinguished between the two. He also argued that all central bankers adopt a risk management approach to some extent and that this causes econometric problems for the analysis of monetary policy because their adoption of such an approach means there are not many sharp economic downturns in the data. Levin suggests that it would be useful to look at periods of greater macroeconomic instability—such as the 1970s—to investigate the extent and causes of policy failures in these situations.

Jordi Galí asked why the analysis lumped together price run-ups that were followed by crashes with price run-ups that were not followed by crashes. He suggested that only the cases with collapses were interesting for the analysis. Cecchetti answered that he had wanted to condition only on the information that there was a price run-up. He did not want his analysis to have to rely on the knowledge that there was a subsequent collapse.

John Y. Campbell pointed out that the Hodrick-Prescott filter (used in the chapter to calculate deviations from trend) is two-sided. He suggested that this meant that the analysis was implicitly conditioning on a collapse because run-ups that were not followed by a collapse would be attributed to the trend.

Lars E. O. Svensson asked for clarification: what sorts of preferences imply an interest in risk management? What are the implications? Does the avoidance of extreme events act as a constraint, in the sense that the policymaker should minimize expected loss subject to some gross domestic product (GDP) at risk or price level at risk constraint? Cecchetti replied that he considered that he had observed policy actions that were driven by higher moment considerations. He believed that these actions were rational but could not be explained by minimization of a quadratic loss function. Svensson said that his prior was that these were, in fact, driven by quadratic loss. If not, he wondered whether Cecchetti thought that such actions were normatively sensible. Cecchetti answered that the central bank should care about extreme events although it was not clear that adjustment of interest rates would always be the appropriate policy response.

John Y. Campbell said that he thought that expected tail loss, discussed in section 1.3.4 of Cecchetti's chapter, was more relevant for policymakers than GDP at risk or price level at risk. Gross domestic product at risk is defined analogously to value at risk (VaR): that is, it specifies the loss that will be incurred over a given horizon at some chosen percentile of the distribution of outcomes. It does not consider the distribution of outcomes at still lower percentiles. Ignoring the distribution of extreme outcomes may make sense for financial institutions that become insolvent if extreme outcomes occur, but it is not appropriate for central banks conducting monetary policy.

Andrew Levin suggested that the discussion was relevant to the zero lower bound on nominal rates. Choosing higher inflation is like making an insurance payment against the possibility of hitting the zero lower bound and entering a liquidity trap. Taking out such insurance might or might not be worthwhile, depending on the size of the inflation cushion needed.