

Measuring the Impact of Asset Price Booms Using Quantile Vector Autoregressions

Stephen G. Cecchetti and Hong Li*

February 2008

Preliminary Draft

Abstract

In analyzing the macroeconomic impact of asset price booms and crashes, it is the disasters that really matter. This emphasis suggests shifting the empirical study of the impact of housing and equity bubbles beyond their impact on the expected value of changes in growth and inflation. With this in mind, we develop and implement procedures for measuring the impact of equity and property booms on the extreme tails of the distribution of deviation in output and price-level from their trends. To do this, we bring together quantile regression techniques and vector autoregression methods to generate a set of tools that allow us to address questions like the following : If we see real housing prices that are ten percent above their trend, what can we infer about the 10th percentile of the distribution of (log) output deviations from its trend 1, or 2, or 3 years later ?

Using data from a broad cross-section of countries, we estimate panel-quantile-vector autoregressions to estimate point forecasts and confidence intervals for the predicted impact of asset price booms on the quantiles of the output and price-level distribution. The results suggest that at horizons of 3 years that both housing booms and equity booms significantly worsen growth and inflation prospects, creating outsized risks that real output will decline significantly below trend and that prices will rise significantly above trend.

* Barbara and Richard M. Rosenberg Professor of Global Finance, Brandeis International Business School; and Assistant Professor, Department of Economics, Brandeis University; respectively. We thank Chanont Banterngansa for research assistance and the Bank for International Settlements for providing the data on housing prices.

I. Introduction

When housing and equity markets boom, what happens? Wealth effects lead to consumption increases. Declines in the cost of firm finance and the return to residential construction bring on an investment boom. Higher government revenues give rise to rising expenditure and lower taxes. But when the boom inevitably turns into a bust, euphoria turns to despair. Consumption and investment collapse, and government revenues do, too. Households and firms go bankrupt, defaulting on loans. Balance sheet deterioration extends to financial institutions, as lenders are left holding collateral that no longer makes them whole.¹

Since the boom is gradual, while the crash is sudden, the impact of asset price bubbles would seem to be quite asymmetrical. On the way up, growth is marginally higher; on the way down, there is a real risk of catastrophe. This pattern with extended periods of regular and steady growth, punctuated by sudden and relatively short downturns, largely mirrors the time profile of economic expansions interrupted by infrequent recessions. This asymmetry means that policymakers intent on stabilizing economic activity are constantly focused on preventing downturns. In particular, for central bankers to meet their stabilization objectives, they must do their best to avoid true disaster. The specter of the Great Depression of the 1930s continues to haunt modern monetary policymakers.

In working to avoid the worst possible macroeconomic outcomes, policymakers are acting as the risk managers for the economic and financial system. Then Federal Reserve Board Chairman Alan Greenspan put it best in 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, pg. 3)

Because of their focus on behavior at the mean of the data, standard time-series macroeconomic models are not well suited to empirical implementation of the risk-management approach. Methods based on the logic of ordinary least squares regression or

¹ For a summary of the impact of asset price bubbles on the macroeconomy see Cecchetti (2006). For a more detailed discussions of the impact of housing on consumption, see Betraut (2002) and Case, Quigley and Shiller (2005). And, for the debate over the appropriate policy response, see Bernanke and Gertler (1999 and 2001) for the “con” side, and Cecchetti, Genberg, Lipsky and Wadhvani (2000) for the “pro” side.

maximum likelihood estimation provide us with an understanding of behavior at the mean of the data. The results answer questions about how the most likely outcome changes when conditioning variables (i.e. exogenous regressors) change. This provides us with information about shifts or translations in the distribution of possible outcomes. While a risk manager will be interested in these, she will also be concerned about whether the evolution of exogenous conditions has had an impact on the *shape* of the distribution. For example, does a boom in the equity or housing market affect the worst possible growth outcomes? What happens to the lowest quartile or decile of the distribution of growth outcomes following a boom? Do standard methods give the right impression, or do we need more?

We address these questions head on using quantile regression techniques. Quantile regression is designed to address the specific question we are interested in: How much does a change in a conditioning variable in a multivariate regression affect the shape of the lower or upper tail of the distribution of the dependent variable?

Since our focus is on dynamic responses, we generalize quantile methods to the case of vector autoregressions. Using quantile vector autoregression (QVARs) techniques, we are able to ask a more subtle time-series question: How does a change in the level of housing or equity prices today influence the lower tail of growth output (or the upper tail of inflation outcomes) several years in the future?² And, in order to exploit all of the information that we have available in sample of countries, we show how to apply the QVAR technique in a panel context.

Our results suggest several important conclusions. First, we find that the impact of housing and equity booms on growth and inflation differs across quantiles. That is, the effect at the center of the distribution differs from the impact on the tails in ways that standard VAR techniques are not designed to reveal. Among our various results, the most intriguing examine the impact of an asset price boom at the 90th percentile of the asset price distribution on the 10th percentile of the output growth distribution (a pretty big disaster) and the 90th percentile of the inflation distribution. Using a panel of 27 countries we find that an equity boom has a negligible effect on growth and inflation at a 4-quarter horizon, but a severe impact negative impact at a horizon of 12 quarters. That is, if we see equity in the 90th percentile today, a further increase in prices will make the worst growth and inflation outcomes significantly worse.

²This work builds a statically coherent structure for intuitive ideas introduced in Cecchetti (2008).

Using data from 17 countries we find that housing booms create analogous long-horizon effects on the tails of the output and growth distributions. But in contrast to the results for equity, housing booms also have short-horizon effects, improving the worst growth and inflation outcomes 4-quarters out.

The remainder of the paper proceeds in four sections. Section 2 presents the basic intuition of quantile regression showing how to extend the simple procedure to estimation, prediction and inference in a vector autoregression. In addition, we provide an outline of how the panel QVAR procedure is implemented. We then move on in Section 3 with a set of preliminary results based primarily on estimates from U.S. data alone. Specifically, we explain in detail the added information we obtain by estimating the models relating housing and equity to output and prices (pairwise). Section 4 moves on to the panel estimates that are based on a broad cross-section of countries. As we discuss, estimation based on 17 countries for the housing models and 27 for the equity models results in a dramatic increase in efficiency. The smaller standard errors for predictions made using the panel QVAR models allow us to draw much more precise conclusions. Section 5 summarizes the conclusions: In virtually every case, asset price booms result in the worst output and inflation outcomes becoming even worse.

II. Quantile Vector Autoregressions

Our interest is in estimating the impact of asset price booms and busts on the likelihood that growth and inflation will end up tail of their unconditional distribution. In order to address this question we will develop and implement a technique that we call “quantile vector autoregression”, or QVAR. To explain our method, we proceed in a series of three steps. Since most macroeconomists are unfamiliar with the general techniques of quantile regression, we start with the basics. Then we describe the application of quantile regression in a single-country vector autoregression. Finally, we show how to generalize our QVAR method to a panel of countries. In order to keep things manageable in the main body of the paper, we have relegated the technical details of what is new to Appendix B.

II.1 Quantile Regression: The Basics

Everyone with even the most basic understanding of econometrics knows that ordinary least squares regressions provide estimates of responses of a dependent variable to a change in a set of independent variables measured at the mean of the data. Put only slightly more technically, the coefficients in an OLS regression provide measures of the derivative of the conditional mean function. That is, for a regression model

$$(1) \quad Y_t = X_t' \beta + \varepsilon_t$$

where X_t is an $k \times 1$ vector of regressors and β is the vector of coefficients, under the assumption $E(\varepsilon_t | X_t) = 0$ (that is, there is no prediction error for a given X_t level),

$\beta_k = \frac{dE(Y_t | X_t)}{dX_t^{(k)}}$, where $X_t^{(k)}$ is the k^{th} element of the vector X_t and β_k is the corresponding

coefficient. The coefficient β_k is the answer to the question: For a one unit change in the k^{th} conditioning variable, $X_t^{(k)}$, how does the entire distribution of Y shift? Which direction does it move; to right ($\beta_k > 0$) or the left ($\beta_k < 0$); and by how much?

But the impact of conditioning variables on the central tendency of the dependent variable may not be the only quantity of economic interest. There are times when we care not only about shifts in the location of the distribution of Y – changes in the conditional mean – but also about changes in the shape of that distribution. For example, does a change in $X_t^{(k)}$ change the thickness of the upper or lower tail of the distribution of Y_t ? We submit that for the case of asset price booms and busts, this question is at least as interesting.

In order to address how changes in a set of conditioning variables influence the shape of the distribution of a dependent variable, Koenker and Bassett (1978) developed the concept of “quantile regression”. Quantile regression is designed to answer the following question: When a conditioning variable $X_t^{(k)}$ changes, what happens to the τ^{th} quantile of the distribution of Y_t ?

Koenker (2005, pg. 1) begins his book on quantile regression by quoting Mosteller and Tukey (1977), who write that “[j]ust as the mean gives an incomplete picture of a single distribution, so the regression curve gives a correspondingly incomplete picture for a set of distributions.”

In their introduction to quantile regression, Koenker and Hallock (2001) describe a variety of cases in order to illustrate how conclusions can change when one shifts away from simply modeling conditional means. One example looks at the impact of a variety of maternal and demographic characteristics on the birth weight of babies. Looking at the difference between boys and girls, Koenker and Hallock report that boys are on average 100 grams larger than girls, with a 90-percent confidence interval of roughly ± 4 grams – that is the mean effect. The authors go on to describe how this disparity is smaller at the lower quantiles than at the higher ones. At the 5th percentile of the distribution, for example, boys are only 45 grams larger than girls, while at the 95th percentile they are 130 grams larger. As Koenker and Hallock (2001, pg. 149) conclude “The conventional least squares confidence interval does a poor job of representing this range of disparities.”³

At first blush, the idea of modeling regression quantiles seems extremely complex. After all, the percentiles of a distribution are order statistics, the properties of which are not the stuff of even graduate econometrics books. In a multivariate context the problem is particularly vexing. How would you order the sample record for a dependent variable conditional on a set of explanatory variables? Fortunately, as Koenker (2005) explains, it is possible to turn the estimation of these order statistics into a straightforward optimization problem that is computationally tractable. The critical insight is to write down the right penalty or loss function.

To see how this works, recall that we can write down an ordinary least squares problem as an optimization problem where we minimize the sum of squared deviations of the fitted values for the dependent variable from the data. That is, for observations indexed by t ,

$$(2) \quad \hat{\beta}_{OLS} = \arg \min_{\beta \in B} \sum_{t=1}^T \rho(\varepsilon_t(\beta))$$

where $\rho(\cdot)$ is a weighting function, and B is the parameter space.

³ Chapter 1 of Koenker (2005) contains additional examples, including one examining the impact of unionization on wages – unionization raises the lower quantiles but has little impact on the upper one.

$$(3) \quad \varepsilon_t(\beta) = Y_t - X_t' \beta$$

is the regression error evaluated at the coefficient value β .

In the ordinary least squares case the weighting function in (2) is quadratic, so $\rho(\varepsilon_t) = \varepsilon_t^2$. But if, instead of a parabola, the function ρ is chosen to be a “V”-shape piecewise linear function (as those illustrated in Figure 1) then, depending on the exact shape of the V function, the optimization problem yields an estimate at a particular quantile of the distribution of ε .

To be more specific, consider the following weighting function for the regression error in (3),

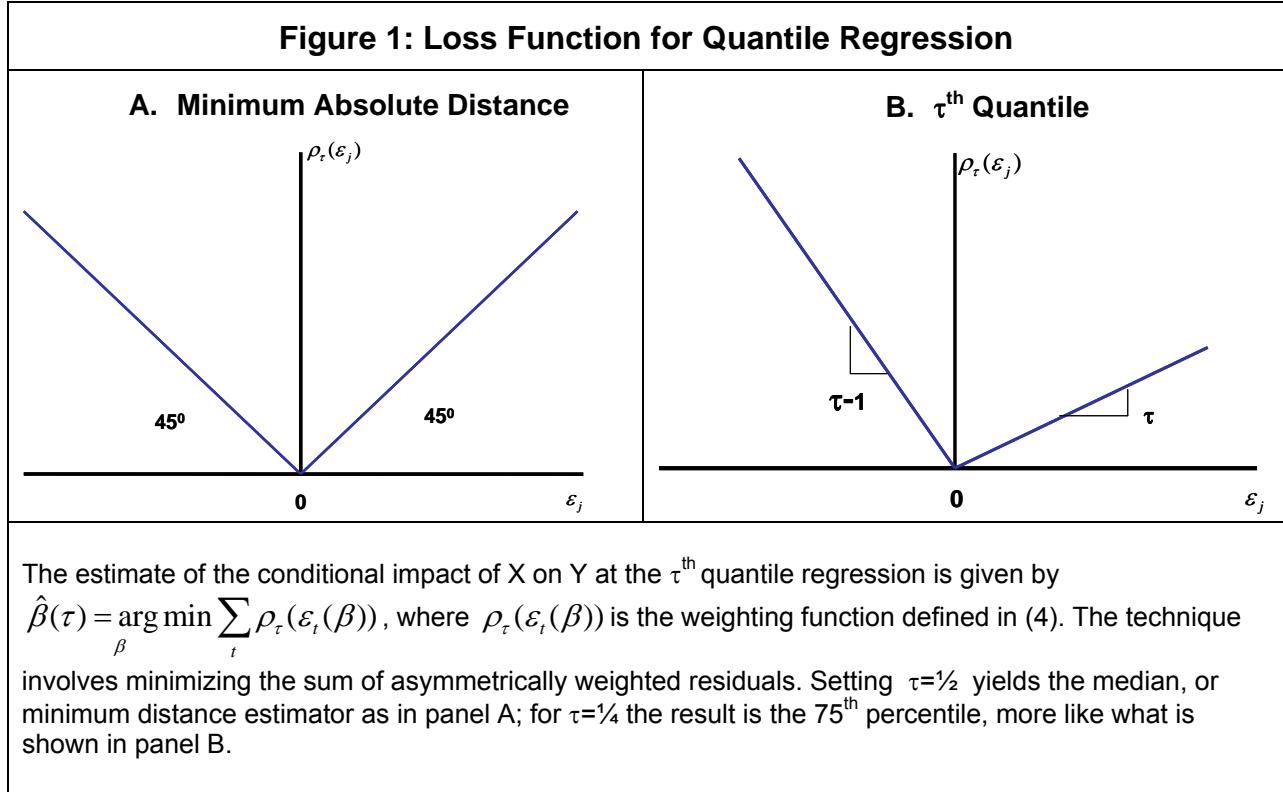
$$(4) \quad \rho_\tau(\varepsilon_t) = \begin{cases} \tau \varepsilon_t & \text{for } \varepsilon_t > 0 \\ (1-\tau) |\varepsilon_t| & \text{for } \varepsilon_t < 0 \end{cases}$$

where $|\cdot|$ denotes the absolute value. That is, for positive errors, we weight the deviations of the fitted regression line from the data by τ , and for negative errors, weight the absolute deviations by $(1-\tau)$.

According to this weighting scheme, in the case of $\tau = 1/2$, the overestimates (i.e., the positive errors) and the underestimates (i.e., the negative errors) are weighted symmetrically. Then $\hat{\beta}$ is chosen so that the probability of Y_t less than $X_t' \hat{\beta}$ equals the probability that Y_t is greater than $X_t' \hat{\beta}$. The resulting $\hat{\beta}$ is the mean-absolute-value estimator for β – the multivariate analog to the sample median in a univariate problem as shown in Panel A of Figure 1. In other words, $X_t' \hat{\beta}$ will be the sample median of Y_t .

For the more general case where the loss function is piecewise linear but *asymmetrical*, as illustrated in panel B of Figure 1, we obtain an estimate at a quantile that depends on the relative slopes on the two sides of the origin. To paraphrase Koenker (2005, pg. 7), if an underestimate is *marginally* three times more costly than an overestimates (the case where $\tau=1/4$

in the figure), we will choose $\hat{\beta}$ so that the probability of Y_t less than $X_t'\hat{\beta}$ is three times greater than the probability that Y_t is greater than $X_t'\hat{\beta}$. The result will be the 75th percentile of Y_t .



To proceed, we need to introduce a small amount of notation some notation. It is useful to write the weighting function in (4) more compactly as

$$(4') \quad \rho_{\tau}(\varepsilon_t) = (\tau - 1_{\varepsilon_t < 0})\varepsilon_t$$

where $1_{(\cdot)}$ is an indicator function that equals one when the argument (\cdot) is true. Let function $Q_{\tau}(\cdot)$ denote the population τ^{th} quantile of the argument, so $Q_{\tau}(\varepsilon_t | X_t)$ denotes the τ^{th} quantile of ε_t conditional on the regressor vector X_t . Then under the assumption that $Q_{\tau}(\varepsilon_t | X_t) = 0$ (i.e., for a given X_t , there is no prediction mistake at the τ^{th} quantile), the quantile regression estimator, denoted by $\hat{\beta}(\tau)$, is defined as⁴

⁴ If the true model is nonlinear, then we interpret our linearized quantile regression as the best linear predictor for the conditional quantile.

$$(5) \quad \hat{\beta}(\tau) = \arg \min_{\beta \in B} \sum_{t=1}^T \rho_{\tau}(Y_t - X_t' \beta).$$

Note that the restriction on the error term in the quantile regression, that $Q_{\tau}(\varepsilon_t | X_t) = 0$, is analogous to the conditional mean restriction that $E(\varepsilon_t | X_t) = 0$ in the OLS regression. As in the standard case, this assumption provides identification in a quantile regression.

Assuming the regressors are weakly stationary and that the density of the error ε_t conditional on X_t at a given quantile τ is continuous, the first order condition (F.O.C.) for the problem in (5) is⁵

$$(6) \quad T^{-1} \sum_{t=1}^T (\tau - 1_{Y_t < X_t' \hat{\beta}(\tau)}) X_t = 0.$$

And the resulting estimator satisfies the sample quantile condition $\hat{Q}_{\tau}(Y_t | X_t) = X_t' \hat{\beta}(\tau)$. Details of the derivation of (6) can be found in Koenker and Bassett (1978) and Koenker (2005).

II.2 Single-Country Quantile Vector Autoregression: Estimation, Inference and Prediction

Since our interest is in the dynamic responses of growth and inflation to equity or housing booms, we need to reformulate the quantile procedure to allow us to consider vector autoregression. To see the approach, consider the example of a simple two-equation vector autoregression with the log of real output, y_t , and log of the real price of housing, h_t (both measured as deviations from a filtered trend). If the coefficients are constrained to be constant across quantiles, as in the standard OLS estimation, the bivariate VAR model is

$$(7) \quad \begin{bmatrix} y_t \\ h_t \end{bmatrix} = \begin{bmatrix} x_t' & 0 \\ 0 & x_t' \end{bmatrix} \begin{bmatrix} \gamma \\ \beta \end{bmatrix} + \begin{bmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{bmatrix},$$

where four lags are used in our study, so $x_t = [1 \ y_{t-1} \ y_{t-2} \ y_{t-3} \ y_{t-4} \ h_{t-1} \ h_{t-2} \ h_{t-3} \ h_{t-4}]'$.

⁵ Because of the kink in the "V", this first-order conditions cannot hold exactly. Nevertheless, as $T \rightarrow \infty$, because the single point has measure zero, the left hand side of equation (6) converges to zero. Stated more precisely, the left hand side of the first-order condition is $o_p(T^{-1/2})$.

In our study, instead, we consider the case in which the output equation is estimated at quantile τ_1 and the housing equation is estimated at quantile τ_2 . To explicitly address the dependence of the VAR model on quantiles, we rewrite (7) as

$$(8) \quad \begin{bmatrix} y_t \\ h_t \end{bmatrix} = \begin{bmatrix} x_t' & 0 \\ 0 & x_t' \end{bmatrix} \begin{bmatrix} \gamma(\tau_1) \\ \beta(\tau_2) \end{bmatrix} + \begin{bmatrix} \varepsilon_{1t}(\tau_1) \\ \varepsilon_{2t}(\tau_2) \end{bmatrix}.$$

This problem is interesting as it addresses questions like: If there is a boom in which housing prices are in 90th percentile of the data ($\tau_2=0.9$), what happens to the 10th percentile of the output gap distribution ($\tau_1=0.1$)? Or more generally, this formulation allows us to answer questions like: Conditioning on the value of the x_t vector, what is the entire distribution of the output gap y_t if the housing price h_t is at τ_2^{th} quantile?

Estimation

To proceed, assume that the errors ε_{1t} and ε_{2t} satisfy the population quantile restrictions $Q_{\tau_1}(\varepsilon_{1t}(\tau_1) | x_t) = 0$ and $Q_{\tau_2}(\varepsilon_{2t}(\tau_2) | x_t) = 0$, respectively. These restrictions imply that the population responses of the τ_1^{th} quantile of y and the τ_2^{th} of h to changes in X are characterized by:

$$(9) \quad \begin{aligned} Q_{\tau_1}(y_t | x_t) &= x_t' \gamma(\tau_1) \\ Q_{\tau_2}(h_t | x_t) &= x_t' \beta(\tau_2) \end{aligned}$$

To estimate the model (9), analogous to the single equation model in the previous subsection, we solve the following joint optimization problem:

$$(10) \quad \begin{aligned} \min_{\gamma} \sum_{t=1}^T \rho_{\tau_1}(y_t - x_t' \gamma(\tau_1)) \\ \min_{\beta} \sum_{t=1}^T \rho_{\tau_2}(h_t - x_t' \beta(\tau_2)) \end{aligned}$$

where $\rho_{\tau_i}(\cdot)$ ($i=1,2$) is the “V”-function defined in (4') for $\tau = \tau_1$ and $\tau = \tau_2$ respectively .⁶

Point estimation of $\gamma(\tau_1)$ and $\beta(\tau_2)$ are obtained by the variant of the linear programming algorithm described in Koenker and Park (1996). In practice, the two QVAR functions can be estimated on an equation-by-equation basis to yield $\hat{\gamma}(\tau_1)$ and $\hat{\beta}(\tau_2)$.

Inference

In order to conduct inference on the estimates $\hat{\gamma}(\tau_1)$ and $\hat{\beta}(\tau_2)$ we focus on the first order condition for the optimization problem in (10). The presence of more than one equation in our setup means that we must take account of the cross-equation correlation. We address this issue by stacking the two first-order conditions. Write the joint F.O.C. as

$$(11) \quad T^{-1} \sum_{t=1}^T g_t(\hat{\gamma}(\tau_1), \hat{\beta}(\tau_2)) = 0$$

where

$$g_t(\gamma(\tau_1), \beta(\tau_2)) = \begin{bmatrix} (\tau_1 - 1_{y_t < x_t' \gamma(\tau_1)}) x_t \\ (\tau_2 - 1_{h_t < x_t' \beta(\tau_2)}) x_t \end{bmatrix}.$$

The two moment functions in $g_t(\gamma(\tau_1), \beta(\tau_2))$ are obtained by applying (6) to the QVAR model in (10) on an equation-by-equation basis. Then by treating (11) as the sample moment restriction in a generalized method of moments (GMM) framework, joint asymptotic normality can be established through the standard argument, and the joint asymptotic distribution of $\hat{\gamma}(\tau_1)$ and $\hat{\beta}(\tau_2)$ can be derived that automatically takes account of the cross-equation correlation. Moreover, throughout the paper we utilize heteroskedasticity and autocorrelation consistent standard errors (HAC estimators), which is a natural extension when applying quantile regression technique to times series data. See Appendix B.2 for all derivation details.

Prediction: The Univariate Case

⁶ It is worth noting that by making one of the $\rho_{\tau}(\cdot)$ functions quadratic we can estimate one equation at the mean of the data, while estimating the other at the τ^{th} quantile. That is, we estimate one equation by OLS and the other by optimizing the function analogous to one of the equations in (7).

The nature of the questions we seek to answer requires that we formulate predictions of the endogenous variables in the QVAR several periods into the future. For example, we would like to know if a housing price boom today increases the risk of output falling well below trend 3 years in the future. This requires forecasting the impact of a move in housing prices today on the lower tail of the distribution of the output gap 12 quarters into the future. To address questions of this sort, we compute the quantile impulse response function at various horizons. To understand how we use the QVAR to do this, let us start with a simple example of a first-order univariate autoregression

$$(12) \quad Y_t = \alpha Y_{t-1} + u_t$$

where α is a parameter and u_t is an i.i.d. mean zero random variable. If we use the conventional OLS estimator for α , $\hat{\alpha}_{OLS}$, then the one-period ahead forecast for Y_{t+1} is

$$(13) \quad E(Y_{t+1} | Y_t) = \hat{\alpha}_{OLS} Y_t.$$

Iterating forward, and recalling that $E(u_{t+k} | Y_t) = 0$ for all $k > 0$, we can then compute the k -period ahead forecast as

$$(14) \quad E(Y_{t+k} | Y_t) = \hat{\alpha}_{OLS}^k Y_t,$$

where $\hat{\alpha}_{OLS}^k$ is the OLS estimate of the coefficient α raised to the k^{th} power.

Suppose, instead of the estimate of the mean of Y_{t+k} conditional on Y_t , we are interested in forecasting Y_{t+k} , conditional on Y_t , at the τ^{th} quantile. That is, for a given value of the conditioning variable, which is current Y in this case, we want to know if the shape of the tail of the distribution of Y_{t+k} has changed.

Starting again with the one-period-ahead forecast, and using the notation we established earlier, the conditional quantile function for the response of Y_{t+1} to Y_t at the τ^{th} quantile, $Q_\tau(Y_{t+1} | Y_t)$, equals

$$(15) \quad Q_\tau(Y_{t+1} | Y_t) = \hat{\alpha}(\tau) Y_t.$$

Multiple-period-ahead forecasts are more complex than in the conditional-mean case because it is possible to condition on previous forecasts at any quantile. The simplest case is where we pick one value for τ and stick with it. To understand how that works, define $\tilde{Y}_{t+1}(\tau)$ as the one-period-ahead forecast given by (15). The forecast of Y_{t+2} at the τ^{th} quantile conditional on this is then

$$(16) \quad Q_{\tau}[Y_{t+2} | \tilde{Y}_{t+1}(\tau)] = \hat{\alpha}(\tau)\tilde{Y}_{t+1}(\tau) = \hat{\alpha}(\tau)^2 Y_t,$$

which looks very similar to the case using OLS in (14).

In principle, there is no reason to condition the two-period-ahead forecast of the τ^{th} quantile of Y on the one-period-ahead forecast at the same quantile. One might, for example, chose to forecast Y_{t+1} at a quantile τ^* that is different from τ , the quantile of interest for the forecast of Y_{t+2} . In this case,

$$(17) \quad Q_{\tau}[Y_{t+2} | \tilde{Y}_{t+1}(\tau^*)] = \hat{\alpha}(\tau)\tilde{Y}_{t+1}(\tau^*) = \hat{\alpha}(\tau^*)\hat{\alpha}(\tau)Y_t.$$

The obvious conclusion is that any k-period-ahead forecast at a particular quantile depends on the entire path of the quantiles chosen for the (k-1) forecasts that precedes it. This leads to a cautionary note: Prediction using regression quantiles requires an economically sensible choice for the path of the quantiles. In what follows, we do what seems natural and predict the k-period-ahead quantile assuming that τ is fixed for the entire path.

Prediction: Using the QVAR

Turning finally to the case of prediction from the vector autoregression, we start by rewriting (8), the two-equation QVAR for output and housing, in its companion form

$$(18) \quad Z_t = A_0(\tau_1, \tau_2) + A_1(\tau_1, \tau_2)Z_{t-1} + \varepsilon_t$$

where for a two-equation four-lag QVAR, $A_0(\tau_1, \tau_2)$ and $A_1(\tau_1, \tau_2)$ are 8×1 vector and 8×8 matrix that contain coefficients $\gamma(\tau_1)$ and $\beta(\tau_2)$ at the quantiles of interest (plus zeros and ones), $Z_t = [y_t$

$y_{t-1} \ y_{t-2} \ y_{t-3} \ h_t \ h_{t-1} \ h_{t-2} \ h_{t-3}]'$; and ε_t includes ε_{1t} and ε_{2t} , as well as zeros. See appendix B.3 for details.

Suppose that we are interested in the k -period ahead forecast for the τ_1^{th} quantile of output conditional on a housing price boom of size h_t . This means that the initial level of the vector of right-hand-side variables in (17) is $Z_t=[0 \ 0 \ 0 \ 0 \ h_t \ 0 \ 0 \ 0]'$. Then the k -period ahead forecast (analogous to the impulse response) for the τ_1^{th} quantile of $y_{t+k}(\tau_1)$ is

$$(19) \quad \tilde{y}_{t+k}(\tau_1) = e_1' \left[\left(I + \hat{A}_1(\tau_1, \tau_2) + \hat{A}_1^2(\tau_1, \tau_2) + \hat{A}_1^3(\tau_1, \tau_2) \cdots \hat{A}_1^{k-1}(\tau_1, \tau_2) \right) \hat{A}_0(\tau_1, \tau_2) + \hat{A}_1^k(\tau_1, \tau_2) Z_t \right]$$

where e_1 is a selection column vector (with its first element being one and other elements being zeros) picking up the first element of the forecasted $\tilde{Z}_{t+k}(\tau_1)$, and $\hat{A}_0(\tau_1, \tau_2)$ and $\hat{A}_1(\tau_1, \tau_2)$ are matrix functions of the estimated values of the coefficients $\gamma(\tau_1)$ and $\beta(\tau_2)$.

Because the k -period-ahead forecast $\tilde{y}_{t+k}(\tau_1)$ is a complex nonlinear function of the estimated coefficient matrices $\hat{A}_0(\tau_1, \tau_2)$ and $\hat{A}_1(\tau_1, \tau_2)$ we cannot compute an exact closed-form expression for the variance of the forecast. Instead, we use the δ -method to compute an approximate variance of the forecast. Even so, a nontrivial technical problem involving matrix differentiation arises. To be more specific, we need to compute the derivative of an 8×1 vector, $(I + A_1 + A_1^2 + A_1^3 \cdots A_1^{k-1})A_0 + A_1^k Z_t$, with respect to A_0 and A_1 where A_1 is an 8×8 matrix. The solution to this problem is provided in Appendix B.3.

II.3 Panel Quantile Vector Autoregression: Estimation, Inference, and Prediction

Extending the QVAR analysis to the case of a panel data set is straightforward, so we relegate all details to Appendix B.2. Briefly, starting with the single equation panel case, we can write the output equation for country i at quantile τ_j as

$$(20) \quad y_{it} = x_{it}' \gamma(\tau_j) + \alpha_{1i} + \varepsilon_{1, it}(\tau_j).$$

In this formulation countries differ only in the additive fixed effect α_i , not in their dynamic responses, $\gamma(\tau_j)$. Following the derivations above for a single country, and restricting the fixed effect to be the same across all quantiles, we can write the conditional τ_j^{th} quantile function for output gap as

$$(21) \quad Q_{\tau_j}(y_{it} | x_{it}) = x_{it}'\gamma(\tau_j) + \alpha_i.$$

Equation (21) is based on the quantile restriction $Q_{\tau_j}(\varepsilon_{1,it}(\tau_j) | x_{it}) = 0$. Using (21), we can compute estimates of the parameters of interest from a panel of n countries, estimated over J quantiles by solving the following optimization problem:

$$(22) \quad \min_{\alpha, \gamma(\tau_1), \gamma(\tau_2), \dots, \gamma(\tau_J)} \sum_{j=1}^J \sum_{i=1}^n \sum_{t=1}^T \rho_{\tau_j} [y_{it} - x_{it}'\gamma(\tau_j) - \alpha_i].$$

Computationally, solution of (22) requires that we stack the data country by country, and then utilize the interior point method described earlier. See Appendix B.2 for more details.

Turning to the multiple-equation case – the panel QVAR – consider the following balanced panel generalization of (8):

$$(23) \quad \begin{bmatrix} y_{it} \\ h_{it} \end{bmatrix} = \begin{bmatrix} x_{it}' & 0 \\ 0 & x_{it}' \end{bmatrix} \begin{bmatrix} \gamma(\tau_1) \\ \beta(\tau_2) \end{bmatrix} + \begin{bmatrix} \alpha_{1i} \\ \alpha_{2i} \end{bmatrix} + \begin{bmatrix} \varepsilon_{1,it}(\tau_1) \\ \varepsilon_{2,it}(\tau_2) \end{bmatrix}$$

where the regressor vector x_{it} is defined as in (8). In this multiple equation case, the individual equation conditional quantile functions are still given by (20), and point estimates of the parameters can then be obtained on an equation-by-equation basis.

Inference and prediction in the panel case are generalizations of the QVAR methods described in the previous subsection. See Appendix B.2. Briefly speaking, one can concentrate out the quantile fixed effects by pre-multiplying the data matrix by a specific quantile residual matrix. After this transformation, many previous results on inference and prediction directly carry over to the panel setup.

III. Preliminaries: Data, Testing and U.S. Results

It is useful to begin with a short description of the data we use, as well as a preliminary examination of its properties. Can we justify using quantile methods on these data? Do equity and housing booms simply change the conditional mean of predicted growth and inflation, leaving the shape of the distribution unchanged? Or, as in the example of infant birth weight discussed by Koenker and Hallock (2001), do asset price bubbles affect the shape of the distribution as well? In particular, if we witness a boom in equity or housing prices, does this change our view of the outcomes in the tail of the growth and inflation distribution four or twelve quarters later? To answer this question, we test for the equality across the quantiles in the coefficients from the panel QVAR in (18) above. Next, in order to demonstrate the economic importance of moving away from conventional VAR estimates, we present results for the U.S. alone. Then, in Section IV, we turn to the panel estimates of the models.

III.1 The Data

We use the data and follow the procedures in Cecchetti (2008). Briefly, we begin with quarterly data from 1970 to 2003 on real GDP, the aggregate price level, real equity prices, and real housing prices for a cross-section of countries. For housing, we have 17 countries, and for equity there are 27.⁷ First, for each country we 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 (1997) filtered trend with a smoothing parameter equal to 1600 (the results are robust to using a parameter of 9600). Taking deviations from country-specific (and time-varying) trends has the advantage 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. It is important to note that the use of a two-sided filter means that large positive deviations of asset prices from this trend – these are the booms – must be followed by crashes. Put another way, the booms we locate cannot continue indefinitely.

⁷ The 17 countries in the housing price sample are Australia, Belgium, Canada, Denmark, Finland, Greece, Ireland, Israel, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, the U.K, and the U.S. The 27 countries in the equity price data sample add Austria, Chile, France, Germany, Italy, Japan, Korea, Mexico, Peru, and South Africa. Sources are listed in Appendix A.

III.2 Test for Equal Coefficients

Turning to the model estimates, throughout the remainder of the paper we report results for bivariate QVAR for four pairs of estimates:

Model A: Output and Housing

Model B: Inflation and Housing

Model C: Output and Equity

Model D: Inflation and Equity

To get a sense of the results, we begin with the case of model A using both U.S. data and panel data. That is, we estimate the models given by equation (8) for U.S. data and equation (23) for panel data – the bivariate QVARs for output and housing with four lags. Figure 2a plots the estimates of the coefficients for the output equation using U.S. data. Each panel of the figure plots the coefficient estimate from the 10th to the 90th quantile (with an increment of one), with its shaded 90-percent confidence interval, along the vertical axis, and the quantile along the horizontal axis. So, for example, the top left panel shows the coefficient estimate for y_{t-1} across the quantiles. For the purpose of comparison, the OLS estimate is plotted as the flat dotted line in each panel. The estimates clearly vary across the quantiles; in some cases substantially. For instance, the coefficient on output lagged one quarter, y_{t-1} , ranges from 0.8 at the 10th quantile to 1.3 at the 90th quantile, while the OLS estimate is 0.94. For coefficients of y_{t-3} and h_{t-3} , the shaded interval estimates rarely include the OLS estimate. The above observations are even more evident when panel dataset is used. See Figure 2b. These plots strongly suggest that it is worth going to the trouble of constructing these estimates.

While visual inspection of plots like the ones in Figures 2a and 2b is very informative, it is also useful to construct statistical tests for the proposition that the coefficients in the regressions are equal. Specifically, for each of the four panel QVAR models A through D we test four hypotheses. Using conventional Wald statistics, we examine

- i) equality of the estimates at the 10th and 50th quantiles,
- ii) equality of the estimates at the 50th and 90th quantiles,
- iii) equality of the estimates at the 10th and 90th quantiles, and
- iv) equality of the estimates at the 10th, 50th, and 90th quantiles together.

Figure 2a: Quantile Regression Estimates for U.S. Output-Housing Equation

$$y_t = \gamma_0 + \gamma_1 y_{t-1} + \gamma_2 y_{t-2} + \gamma_3 y_{t-3} + \gamma_4 y_{t-4} + \gamma_5 h_{t-1} + \gamma_6 h_{t-2} + \gamma_7 h_{t-3} + \gamma_8 h_{t-4} + \varepsilon_{1t}$$

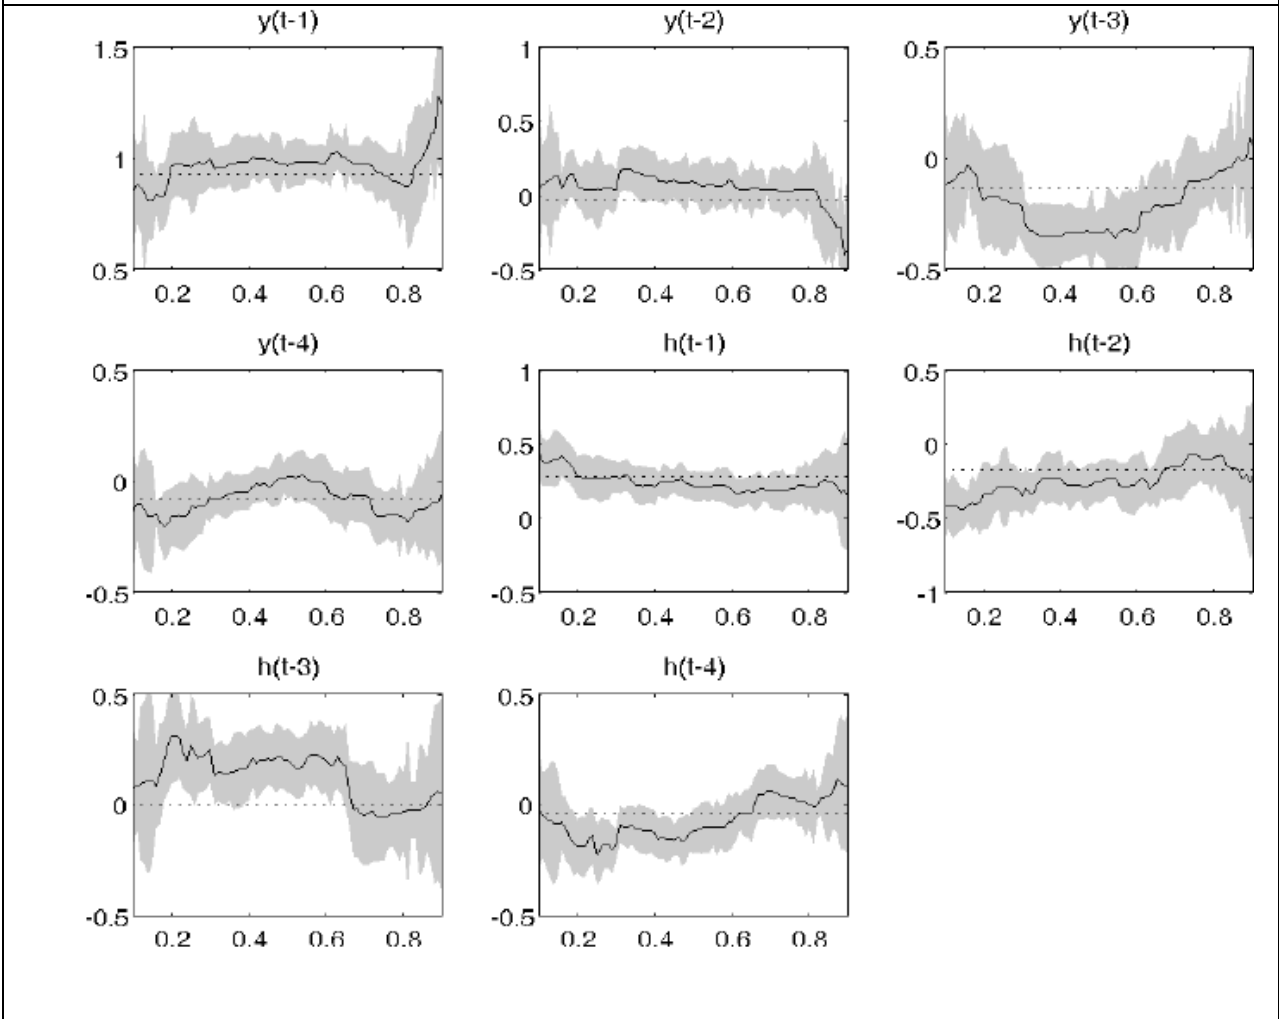


Figure 2a displays the estimates of the output-housing equation for U.S. data, over quantiles ranging from 0.10 to 0.90. The shaded regions are the 90% confidence intervals. The dotted straight line in each panel represents the OLS estimate.

We perform each of these four tests equation by equation for both U.S. data and panel data. Since each of the four models includes two equations, that means we are doing $8 \times 4 \times 2 = 64$ tests. In all cases, we are able to reject the null hypothesis of equality at very high levels of significance. The largest p-value for the test is 0.0017.⁸

⁸ We also tested a series of hypothesis associated with the symmetry of the estimated parameters. Equation by equation, following Newey and Powell (1987), we test the joint null hypothesis that $\frac{1}{2}[\beta(0.1) + \beta(0.9)] = \beta(0.5)$ and $\frac{1}{2}[\beta(0.25) + \beta(0.75)] = \beta(0.5)$, and equivalently for $\gamma(\tau)$. For the U.S.-only model, we can never reject symmetry. By contrast, for the panel-QVAR we can reject symmetry at the 5-percent significance level in 4 of the 8 cases. We find it unsurprising that the errors in the model estimated using single-country data

**Figure 2b: Quantile Regression Estimates for Panel Data
Output-Housing Equation**

$$y_{it} = \gamma_{0,i} + \gamma_1 y_{i,t-1} + \gamma_2 y_{i,t-2} + \gamma_3 y_{i,t-3} + \gamma_4 y_{i,t-4} + \gamma_5 h_{i,t-1} + \gamma_6 h_{i,t-2} + \gamma_7 h_{i,t-3} + \gamma_8 h_{i,t-4} + \varepsilon_{1,t}$$




Figure 2b displays the estimates of the output-housing equation for the panel data, over quantiles ranging from 0.10 to 0.90. The shaded regions are the 90% confidence intervals. The dotted straight line in each panel represents the OLS estimate.

Thus, exercises conducted in this subsection suggest that the macroeconomic relations of interest are not constant across quantiles. This in turn implies that OLS estimation will not be adequate for our purpose of forecasting the worst economic outcomes because OLS estimates average different quantile-specific relations. This justifies the application of QVAR technique developed in Section II to constructing predictions in output and inflation in the rest of the paper.

is more likely to be symmetric than the errors estimated in a model that pools an array of disparate countries.

III.3 Economic Significance and the Case of the U.S.

The equal coefficient tests clearly allow us to conclude that the estimates across the quantiles are statistically distinguishable. But are the differences important in a quantitative economic sense?

We address this question in two related ways. First, we calculate measures of goodness of fit – the analog to an R^2 for a quantile regression. Then we look how employing QVAR estimates changes the predictions that we obtain for the impact of housing or equity booms on output and inflation.

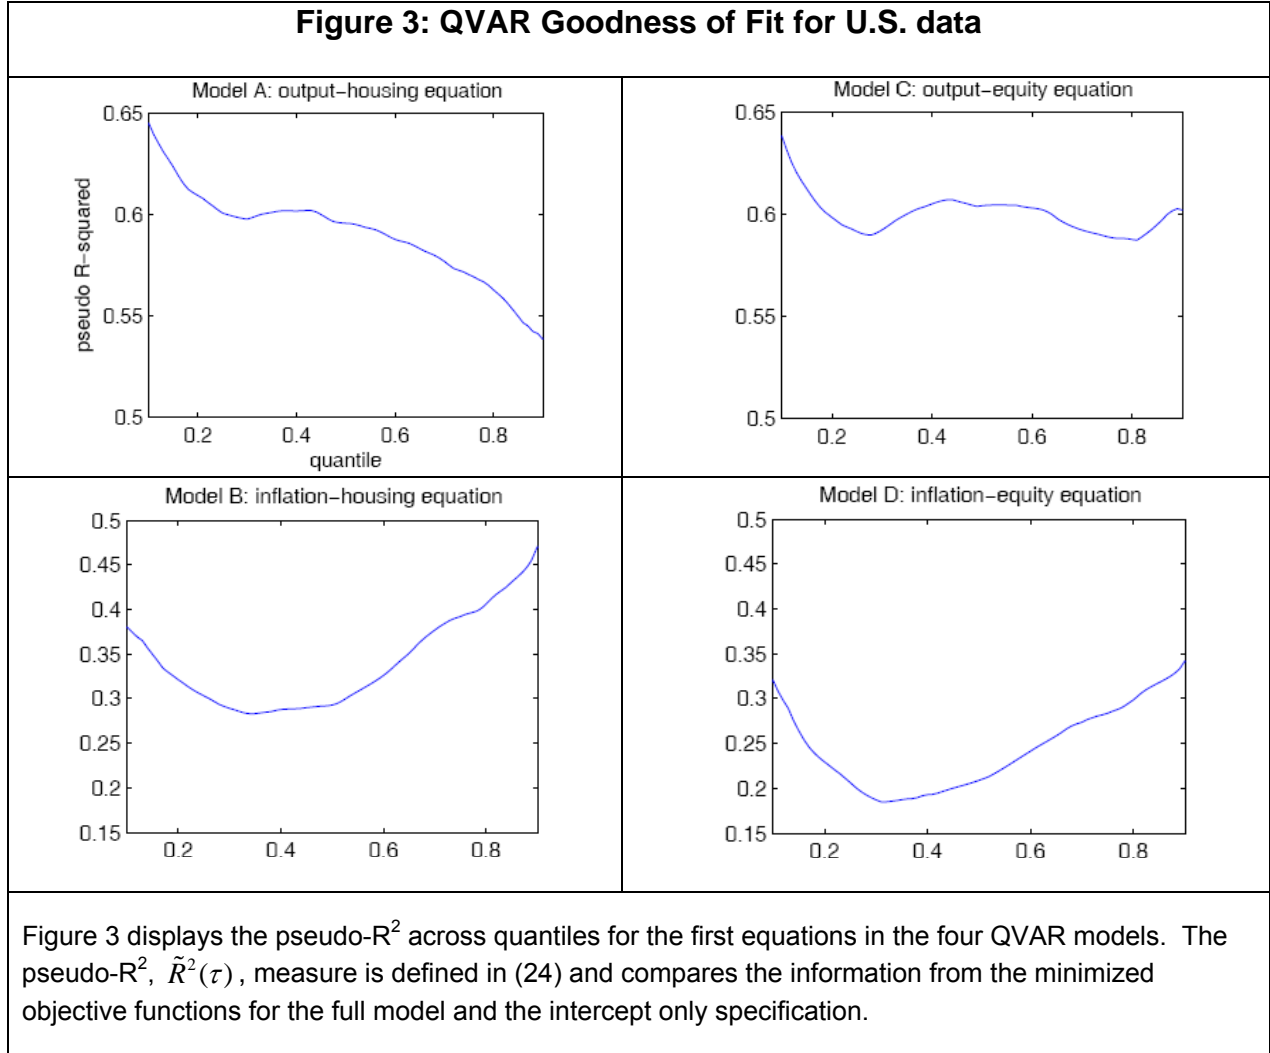
We measure the goodness-of-fit in quantile regression by computing a statistic that is analogous to the familiar R^2 of classical least squares regression. That is, we measure the improvement in the criterion function that arises from the inclusion of the regressors. Specifically, the pseudo- R^2 measure at the τ^{th} quantile, $\tilde{R}^2(\tau)$ is defined as

$$(24) \quad \tilde{R}^2(\tau) = 1 - \frac{\sum_{t=1}^T \rho_{\tau}(Y_t - X_t' \hat{\beta}(\tau))}{\sum_{t=1}^T \rho_{\tau}(Y_t - \hat{\beta}_0(\tau))}$$

where $\sum \rho_{\tau}(Y_t - X_t' \hat{\beta}(\tau))$ is the minimized unrestricted objective function at the τ^{th} quantile, and $\sum \rho_{\tau}(Y_t - \hat{\beta}_0(\tau))$ is the minimized restricted objective function at the τ^{th} quantile with only intercept and no regressors.

Figure 3 plots the pseudo- R^2 measure defined in (24), $\tilde{R}^2(\tau)$, for the output and inflation equations in the four bivariate QVAR models (Models A–D), estimated using U.S. data. The results are quite striking. For both housing and equity the fit of the two output models (Models A and C) is best at the lower quantiles. That is, we are able to explain movements in the 10th quantile of the output distribution (the lower tail) better than we can match the evolution of the 90th quantile (the upper tail). For inflation, shown in the lower two panels of the figure (Models B and D), the situation is reversed. Here we see that the best fit is at the upper quantiles. Not only does the fit vary at different points in the output and inflation distribution, but in several cases the differences across the quantiles is large. For example, looking at the results for Model B,

the inflation-housing model in the bottom left of Figure 3, $\tilde{R}^2(\tau)$ varies from 0.28 to 0.47. This is further evidence of the importance of moving away from the OLS estimation to model the dynamics of the responses of output and inflation to asset price movements.



We have emphasized from the outset that our primary interest is in measuring the extent to which housing and equity booms influence the lower tail of the output distribution and the upper tail of the inflation distribution. To address this issue directly, we look at the predictions that are derived in equation (19) of Section II.2. Figure 4a plots the QVAR prediction results for using U.S. data for models A and C, output-housing and output-equity, respectively. These plots answer the following question: If we see a housing or equity boom of x-percent (measured as a deviation from the trend in real housing or equity prices), what happens to the 10th quantile of

the distribution of the deviation of log output from trend at a horizon of 4 or 12 quarters? This is the quantity that Cecchetti (2008) refers to the 10-percent GDP at Risk.⁹

The panels of Figure 4a plot the size of the boom on the horizontal axis and the change in the 10-percent GDP at risk (the 10th quantile of the GDP gap distribution) on the horizontal axis. Included are estimates for 4- and 12-quarter horizons, together with shaded 90-percent confidence intervals. The two rows of Figure 4a differ in the quantile of the housing equation used to construct the predictions. The top row reports results when the housing/equity equations are estimated at median, and the bottom panel shows estimates based on the 90th quantile for the housing/equity equation. In other words, in the context of the QVAR model in (8), Figure 4a reports the case where τ_1 is fixed at 0.1 (the 10th quantile of output), but the two rows reports prediction results for $\tau_2 = 0.5$ and $\tau_2 = 0.9$, respectively.

To better understand the role of asset price booms, we include in Figure 4a the forecasts based on a simple quantile autoregression (QAR) for output alone. Take an example, a four-lag QAR model for output is

$$y_t = \gamma_0(\tau) + \gamma_1(\tau)y_{t-1} + \gamma_2(\tau)y_{t-2} + \gamma_3(\tau)y_{t-3} + \gamma_4(\tau)y_{t-4} + \varepsilon_t(\tau).$$

Given the initial value of $y_{t-p}=0$ (for $p= 1$ to 4), the 4-period and 12-period ahead forecasts can be computed based on the quantile estimates of γ 's. Since deviations of the QAR predictions of output from zero are a consequence solely of the persistence of the output-gap process, comparing these results with those from the QVAR gives us a sense of the importance of asset price. Since the estimates from the QAR model explicitly exclude the role of housing and equity, we refer to them as the “unconditional forecasts.” These unconditional estimates are displayed as the dotted or dashed horizontal lines in Figure 4a.

For comparison, we also report the results obtained from a standard VAR model estimated using ordinary least squares. These forecasts measure shifts in the mean of the distribution of the output forecasts conditional on asset price booms and are plotted in Figure 4b. A comparison of Figure 4a and Figure 4b shows the difference between conditional quantile predictions and shifts in the conditional mean. Combining information from QVAR and OLS

⁹ The concept is analogous what financial economists call “value at risk.” See Jorion (2001).

forecasts, we will be able to draw conclusion about the shift in both central location and tail of the output and inflation distributions.

Figures 4a and 4b lead to several observations. First, one of the most important things to notice is how different the QVAR estimates are from the unconditional ones. For instance, at a four-quarter horizon, the confidence bands of the QVAR predictions in Figure 4a typically do not cover the unconditional estimates. This observation is consistent with results from Granger causality tests we performed on the output-housing and output-equity equations, which yield p-values of zeros at the 10th quantile.

Second, at a four-quarter horizon, the bigger the boom the higher the 10th percentile of the output gap distribution. For moderate to large asset booms, the output gap forecasts are positive. This is not an enormous surprise – when we see housing or equity booming, the chances are good that the boom will continue for another year.

Third, at a 12-quarter horizon the pictures are different. This is especially true when we use the 90th quantile estimates from the housing/equity equation (the second row of Figure 4a). Now we see that as the boom increases, the lower tail of the output distribution gets worse. The bigger the boom the more negative is the 10th percentile of the output gap distribution. As documented originally in Cecchetti (2008), asset price booms worsen the worst outcomes. That is, housing and equity bubbles increase in GDP at Risk.

Fourth, it is clear by inspection that the QVAR estimates are very different from the estimates by OLS and provide additional information to complement the OLS estimates. To see this, compare the second row of Figure 4a with Figure 4b. Focusing on the 12-quarter horizon, the OLS forecasts are not statistically different from zero regardless of the size of the housing or equity boom. In other words, seeing an asset price boom today does not influence the predicted mean of the output distribution three years ahead. On the other hand, the QVAR forecasts and the corresponding confidence bands (in the second row of Figure 4a) are well below zero, for any boom size. Thus although asset price booms will not shift the central location of output distribution, they very clearly change its shape.

Figure 4a: QVAR Predictions for Output using U.S. data

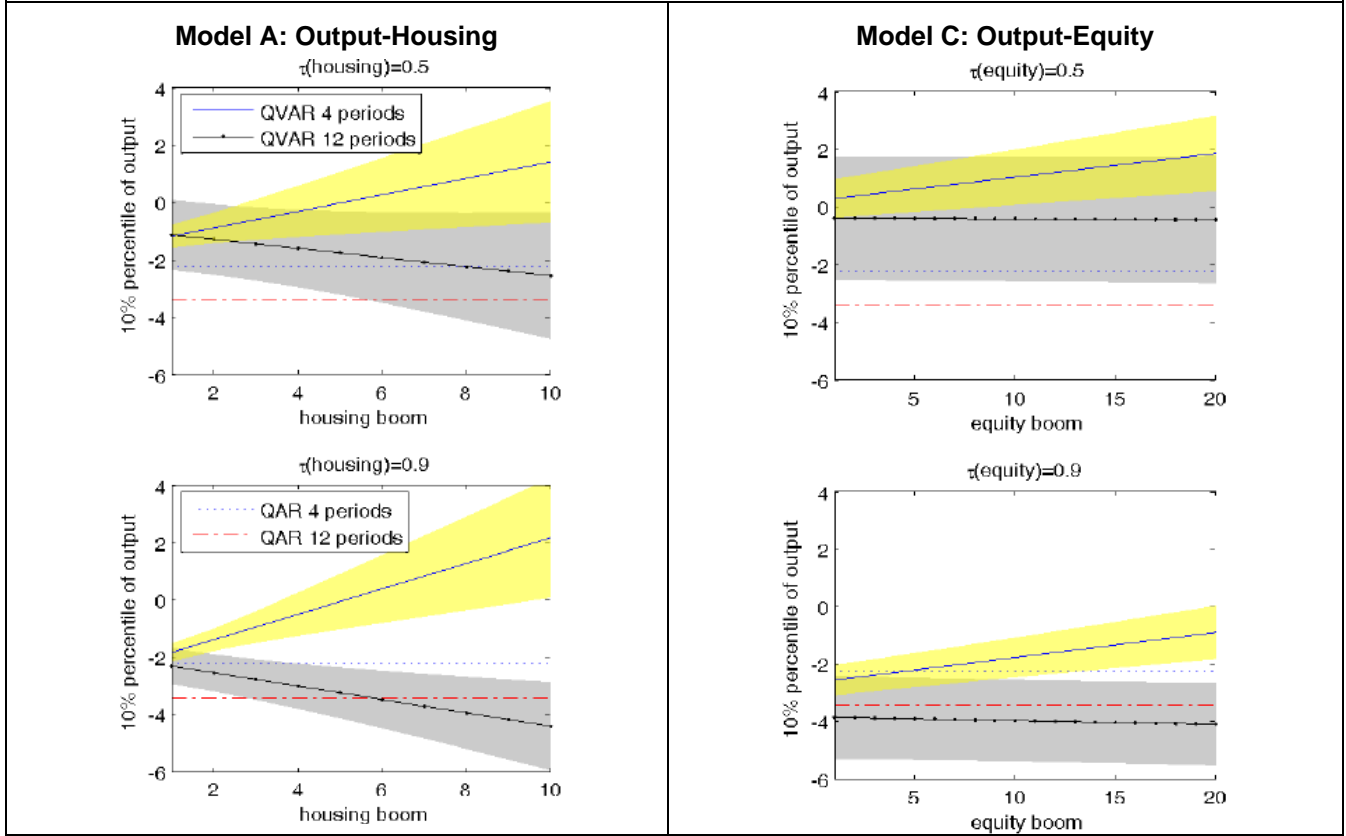


Figure 4a displays the QVAR prediction for the 10% percentile of output in the presence of housing and equity booms. Dotted blue line is the unconditional QAR forecast 4-quarters ahead. Dashed red line is the unconditional QAR forecast 12-quarters ahead. Solid blue line with yellow shade is the QVAR forecast 4-quarters ahead. Solid black line with gray shade is the QVAR forecast 12-quarters ahead.

Figure 4b: OLS-VAR Predictions for Output using U.S. data




Figure 4b displays the VAR prediction for the 10% percentile of output in the presence of housing and equity booms, calculated by OLS. Solid blue line with yellow shade is the VAR forecast 4-quarters ahead. Solid black line with gray shade is the VAR forecast 12-quarters ahead.

Figures 5a and 5b plot QVAR and OLS-VAR predictions based on the two inflation models estimated using the U.S. data. Figure 5a reports results analogous to those in Figure 4a with one important change. While we look at impact of asset price booms on the 10th quantile of log output deviations from trend, here we present results for the consequences of booms on the 90th quantile of log price level deviations from its trend. That is, we are worried about the 90-percent Price Level at Risk.

Looking at these results we draw the following conclusions. First, in the case of the 4-quarter-ahead forecasts, the quantile estimates differ markedly from the unconditional forecasts.

Second, booms tend to drive prices up. This effect is especially conspicuous for housing booms in the short run and for equity booms in the long run. Note that the impacts of a housing boom and an equity boom are not symmetric: For a housing booms impact at a 4-quarter horizon is larger than that at a 12-quarter horizon, while for an equity booms it is just the opposite – the impact at longer horizons is bigger.

Third, the rightward shift in the upper tail of the price-level distribution can be substantial. If we observe an 8 percent housing boom, then using the 90th percentile estimates for the housing equation (see the lower left panel of Figure 5a) we conclude the predicted impact on the 90th percentile of the price distribution is almost +8 at a 4-quarter horizon and more than +5 at a 12-quarter horizon. For a 15 percent equity boom, the numbers are very similar though the horizons are reversed (see the lower right panel of Figure 5a).

Fourth, comparing the QVAR inflation predictions with the OLS estimates, in almost all cases the upper tail of the inflation distribution responds more strongly to asset price booms than the mean does. That is, the primary impact of the boom is to change the shape of the distribution, not shift its location. For instance, in the context of inflation-equity relation, equity boom of any size has essentially zero impact on the conditional mean of the inflation distribution (Figure 5b, the right panel), although the upper tail is severely affected. Thus, OLS analysis may not be adequate when our concern is the worst outcome.

Figure 5a: QVAR Predictions for Inflation using U.S. data

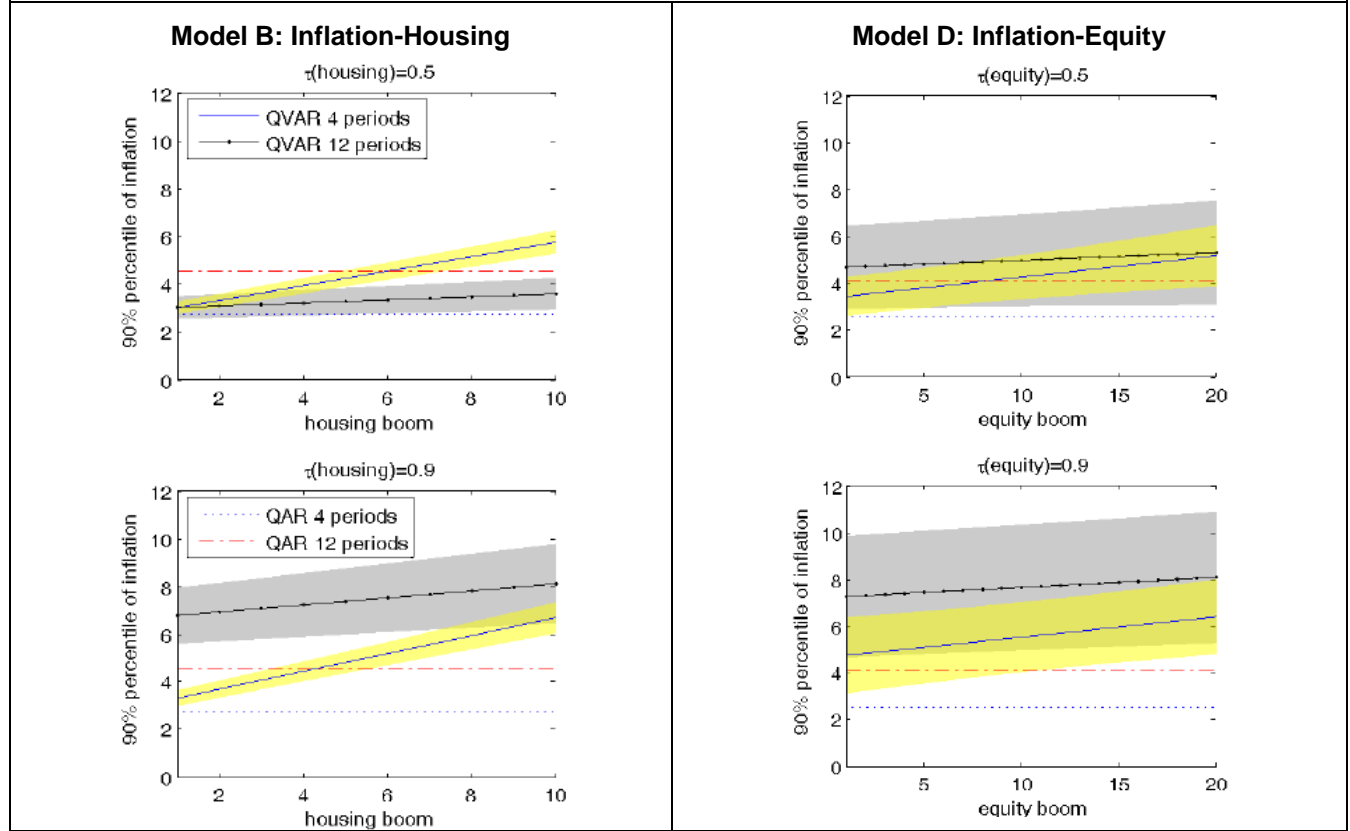


Figure 5a displays the QVAR prediction for the 90% percentile of inflation in the presence of housing and equity booms. Dotted blue line is the unconditional the QAR forecast 4-quarters ahead. Dashed red line is the unconditional QAR forecast 12-quarters ahead. Solid blue line with yellow shade is the QVAR forecast 4-quarters ahead. Solid black line with gray shade is the QVAR forecast 12-quarters ahead.

Figure 5b: OLS-VAR Predictions for Inflation using U.S. data




Figure 5b displays the VAR prediction for the 90% percentile of inflation in the presence of housing and equity booms, calculated by OLS. Solid blue line with yellow shade is the VAR forecast 4-quarters ahead. Solid black line with gray shade is the VAR forecast 12-quarters ahead.

Before moving on to the multi-country panel estimation, we collect the results based on the U.S. data in Table 1. The conclusions reported are based on the QVAR models where the output (or inflation) equation is estimated at the 10th percentile (or 90th percentile), and the asset price equation (in either housing or equity) is estimated at the 90th percentile.

From the summary in Table 1, we draw two conclusions: (1) At a 4-quarter horizon, we predict the impact of asset booms on output to be largely positive in the sense that they reduce the probability of very low growth; and we predict the impact of booms on inflation to be largely negative in the sense that they increase the probability of very high inflation. (2) At a 12-quarter horizon, by which time the boom is likely to have ended in a bust, things are not generally rosy: in most cases asset booms increase the probability of very low growth and very high inflation outcomes.

| Table 1: QVAR Forecast Summary, U.S. only | | |
|--|-------------------------|--------------------------------------|
| $\tau_{\text{output}} = 0.1, \tau_{\text{inflation}} = 0.9, \tau_{\text{asset price}} = 0.9$ | | |
| Model | 4-quarters ahead | 12-quarters ahead |
| Output-Housing | Positive and large | Mixed and small |
| Output-Equity | Positive but small | Negative but small |
| Inflation-Housing | Negative and large | Negative but smaller than 4-quarters |
| Inflation-Equity | Negative and large | Negative and large |

Table 1 summarizes the findings reported in Figures 4a and 5a. All results are based on the bivariate QVAR models where output equation is estimated at the 10th percentile; inflation equation is estimated at the 90th percentile; and the asset price (housing or equity) equation is estimated at the 90th percentile. An asset price effect is large if the departure of the QVAR prediction is far away from the corresponding QAR prediction, the effect is small if the departure is small.

But overall, our results from examining U.S. data alone are quite imprecise. Looking at Figures 4a and 5a it is clear that the standard errors for the QVAR impulse response estimates are relatively big. For example, the 90-percent confidence bands on the estimated impact of a housing or equity boom on output (the bottom of Figure 4) are on the order of ± 1.5 percentage points. Given that the point estimates at a 12-quarter horizon run are roughly -4 percent, this is quite large. As a result, our conclusions thus far are merely suggestive of the impact that asset price booms might have on the shape of the tails of the output and inflation distributions.

We believe that the imprecision in our estimates arises from the fact that over the sample period we study there have been very few large asset price booms and busts in the United States. Our interest all along has been in the measuring the impact of asset price booms on output and inflation. Almost by definition, such “bubble” events occur infrequently. The technical consequence of sparse data in the tails of the distribution is lower power estimates. To address this shortcoming, we move on to incorporate data from other countries using the panel QVAR techniques introduced in Section II.3.

IV. Multi-Country Panel QVAR Estimation

The most natural way to expand our data set is to add more countries. With this in mind, we proceed to estimate our models by applying the panel QVAR to the full data set described in Section III.1. For housing, this means incorporating information from 17 countries; and for equity we now have data from 27 countries. As described in Section II.3, we allow for country-specific fixed effects, but constrain the dynamic responses to be the same across countries.¹⁰ This constraint allows us to dramatically improve the efficiency of our estimates, reducing the confidence intervals of the forecasts of interest by nearly a factor of four.

| Table 2: Efficiency Comparison | | | | | | |
|---|------------------------------|------------------|------------------------------|------------------|------------------------------|------------------|
| U.S.-only and Panel QVAR <u>Standard Errors</u> in the Output-Housing Regression | | | | | | |
| $y_t = \gamma_0 + \gamma_1 y_{t-1} + \gamma_2 y_{t-2} + \gamma_3 y_{t-3} + \gamma_4 y_{t-4} + \gamma_5 h_{t-1} + \gamma_6 h_{t-2} + \gamma_7 h_{t-3} + \gamma_8 h_{t-4} + \varepsilon_{1t}$ | | | | | | |
| | $\tau_{\text{output}} = 0.1$ | | $\tau_{\text{output}} = 0.5$ | | $\tau_{\text{output}} = 0.9$ | |
| Coefficient | Panel | U.S.-only | Panel | U.S.-only | Panel | U.S.-only |
| y_{t-1} | 0.06 | 0.17 | 0.03 | 0.07 | 0.07 | 0.09 |
| y_{t-2} | 0.07 | 0.23 | 0.04 | 0.09 | 0.09 | 0.13 |
| y_{t-3} | 0.09 | 0.23 | 0.04 | 0.09 | 0.09 | 0.11 |
| y_{t-4} | 0.07 | 0.13 | 0.03 | 0.07 | 0.07 | 0.09 |
| h_{t-1} | 0.02 | 0.13 | 0.01 | 0.03 | 0.03 | 0.08 |
| h_{t-2} | 0.03 | 0.25 | 0.02 | 0.04 | 0.04 | 0.13 |
| h_{t-3} | 0.03 | 0.16 | 0.02 | 0.04 | 0.04 | 0.15 |
| h_{t-4} | 0.02 | 0.10 | 0.01 | 0.03 | 0.02 | 0.09 |

Table 2 compares the QVAR standard errors of coefficients in the output-housing equation, based on the U.S. data and panel data. When computing the standard errors, the housing equation (i.e., the second equation in the output-housing QVAR) is estimated at the 90th percentile, while the output equation is estimated at the 10th, 50th and 90th percentiles.

¹⁰ Recall, however, that our filtering procedure means that we have removed country-specific time-varying means from the all of the individual time-series before estimating the model.

The large efficiency gains are evident from the estimated coefficients of the panel QVAR. To make this point, Table 2 presents a comparison of the estimated heteroskedasticity and autocorrelation consistent (HAC) standard errors for coefficients in the output-housing equation (i.e., the first equation in the output-housing bivariate QVAR) using the U.S.-only and panel data sets. Results are reported for the 10th, 50th, and 90th quantiles.

Every pair-wise comparison in the table reveals an increase in precision as measured by a decline in the standard error of the coefficient estimate. In most cases, the precision improvement is substantial. For 19 of 25 cases reported in Table 2, standard errors decrease by at least 50%. Note that the efficiency gain is particularly large at the lower tail. To see this, for the 10th output quantile, which has been the focus of much of our investigation, the shaded area in Table 2 reports an average decline in the standard errors of the estimated lagged housing coefficients of more than 80 percent. Because of the dramatic increase in precision that comes from broadening our sample to seventeen countries, and the resultant increase in the power of any hypothesis tests we wish to perform, we are inclined to put more weight on the panel estimates in drawing our conclusions.

Turning to the question of primary interest, Figures 6 and 7 present panel results matching the ones presented in Figures 4 and 5. Again we examine the impact of an asset-price boom of a given size on output and inflation 4- and 12-quarters later. As the results reported in Table 2 lead us to expect, the 90-percent confidence intervals shown by the shaded areas are quite a bit smaller when we use the entire panel of countries available to us.

Looking in more detail, and starting with the results for output in Figure 6a, we note a series of important conclusions. First, the 4-quarter effects of housing and equity booms on output are asymmetrical. The bottom left panel of the figure shows that at a 4-quarter horizon the 10th quantile of output is higher than the unconditional for any housing boom larger than 5%. For instance, instead of around -2 percent, as would be implied by the unconditional forecast, the panel QVAR implies that a 10 percent housing boom today reduces the 10th quantile of the output distribution by a more than one-half of one percentage point, to about -1.4 percent. On the other hand, for the case of equity, shown in the bottom right panel of Figure 6a, at a 4-quarter horizon the 10th percentile of output is lower than the unconditional for most boom sizes.

Figure 6a: Panel QVAR Predictions for Output




Figure 6a displays the panel QVAR prediction for the 10% percentile of output in the presence of housing and equity booms. Dotted blue line is the unconditional QAR forecast 4-quarters ahead. Dashed red line is the unconditional QAR forecast 12-quarters ahead. Solid blue line with yellow shade is the QVAR forecast 4-quarters ahead. Solid black line with gray shade is the QVAR forecast 12-quarters ahead.

Figure 6b: Panel OLS-VAR Predictions for Output




Figure 6b displays the panel VAR prediction for the 10% percentile of output in the presence of housing and equity booms, calculated by OLS. Solid blue line with yellow shade is the VAR forecast 4-quarters ahead. Solid black line with gray shade is the VAR forecast 12-quarters ahead.

Second, the real payoff to the panel QVAR method comes when we look at the 12-quarter-ahead forecasts. Here the predictions from the QVAR are predictions differ markedly from the unconditional ones. In every case, we find that the asset-price booms make things worse. Since the boom is likely to be followed by a bust that will almost surely come within the 3-year horizon of the forecast, high housing or equity prices today lead us to conclude that the 10th quantile of the output distribution will be much lower 12 quarters from now. For either a 10-percent housing boom or a 20-percent equity boom, the 10th percentile of the distribution of the output gap falls from -2 to nearly -3.5.

Third, there are large differences between the QVAR forecasts in Figure 6a and the OLS forecasts reported in Figure 6b. As was the case when we examined the U.S. data alone, the OLS forecasts at both 4- and 12-quarter horizons are small and close to zero. This is obviously not the case for QVAR forecasts, which lie below zero in all cases for all boom sizes. Once again, the data suggest that the central location of the conditional distribution of output does not shift in the presence of asset price booms. Instead housing and equity booms change the shape of the output distribution, making the worst outcomes even worse.

The panel QVAR results for inflation are summarized in Figure 7a. Looking at these, we again conclude that there are benefits to looking across the quantiles. Specifically, in all cases, large housing and equity booms shift the 90th percentile of the inflation distribution to the right, increasing the probability of a significant inflation increase by much more than the unconditional estimates imply. The problem is worse at the shorter horizons, but clearly continues as the horizon increases. Again, quantile predictions in Figure 7a provide useful information which is impossible to obtain by the OLS prediction in Figure 7b: At the mean level of inflation, asset price effects are negligible; but for upper quantile of inflation, the effects are substantial.

Finally, we summarize the forecasts based on the panel estimates reported in Figures 6a and 7a in Table 3. Comparing the panel results in Table 3 with the U.S.-only results summarized in Table 1, several conclusions are in order. First, as we expected, the improved precision of our panel QVAR estimates allow us to come to more exact conclusions than we could when we only used data from the United States. Before, we were unable to conclude much about the impact of a housing boom on the 10th quantile of the output distribution at a 12-quarter horizon. Now, we can say with some confidence that the impact is large and negative. That is, incorporating

Figure 7a: Panel QVAR Predictions for Inflation

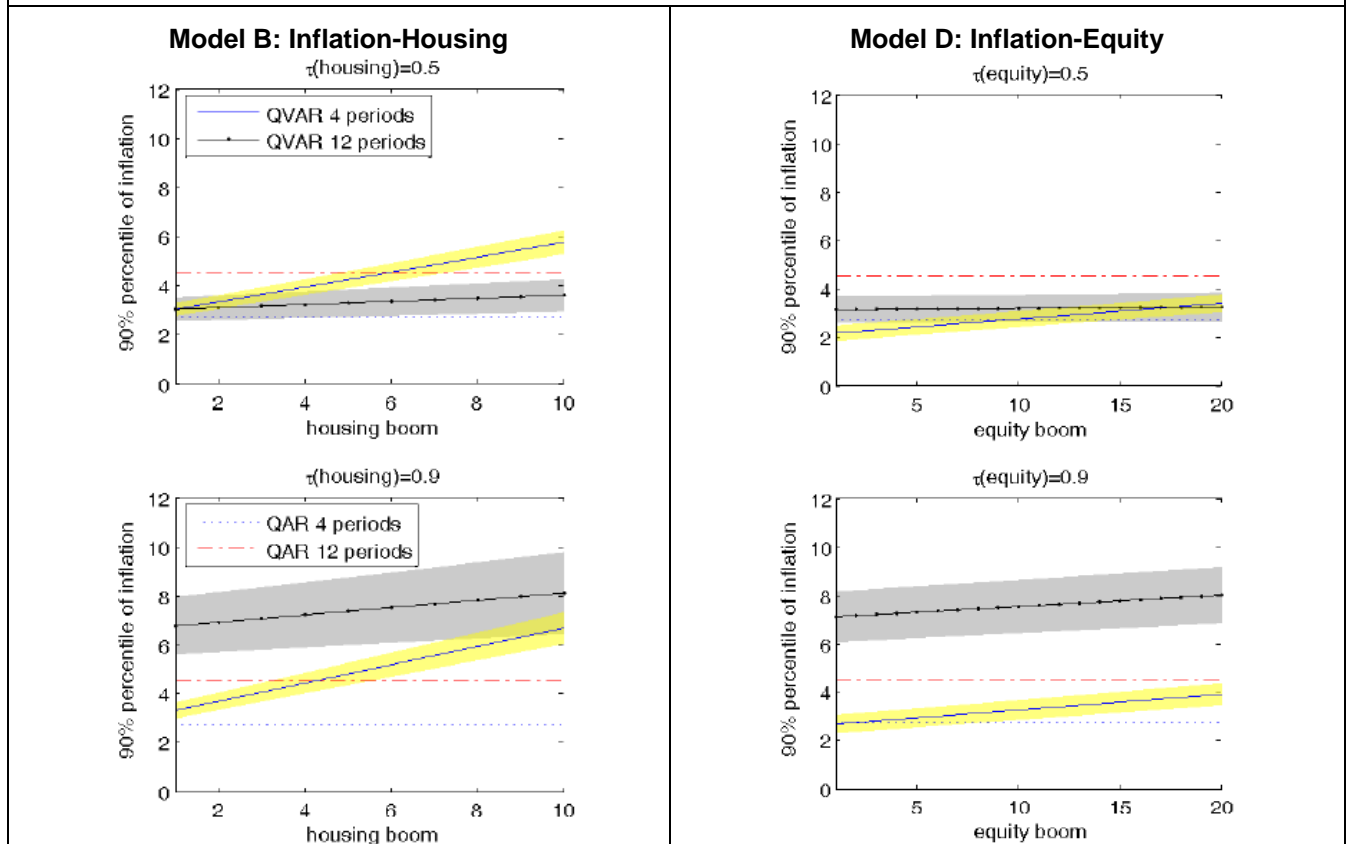


Figure 7a panel displays the QVAR prediction for the 90% percentile of inflation in the presence of housing and equity booms. Dotted blue line is the unconditional the QAR forecast 4-quarters ahead. Dashed red line is the unconditional QAR forecast 12-quarters ahead. Solid blue line with yellow shade is the QVAR forecast 4-quarters ahead. Solid black line with gray shade is the QVAR forecast 12-quarters ahead.

Figure 7b: Panel OLS-VAR Predictions for Inflation




Figure 7b displays the panel VAR prediction for the 90% percentile of inflation in the presence of housing and equity booms, calculated by OLS. Solid blue line with yellow shade is the VAR forecast 4-quarters ahead. Solid black line with gray shade is the VAR forecast 12-quarters ahead.

data from a broad cross-section of countries, we are able to conclude that housing booms worsen the worst growth outcomes, increasing GDP at risk.

Second, more generally, based on panel estimates, all 12-quarter-ahead effects of both types of asset price boom are uniformly negative with large magnitudes. So three years ahead, asset prices booms pose significant risks, making the tails of the distributions of output and inflation even worse than they already were. This result is much stronger than that based on U.S. data.

Third, notice that even at a 4-quarter horizon, panel QVAR forecasts produce more negative pictures than U.S.-only forecasts in Table 1. Based on the panel QVAR analysis, in three out of four cases, asset prices booms make the worst outcomes more likely. The single exception is the impact of housing booms on output deviations from trend at a short horizon.

| Table 3: Panel QVAR Forecast Summary | | |
|--|-------------------------|--------------------------|
| $\tau_{\text{output}} = 0.1, \tau_{\text{inflation}} = 0.9, \tau_{\text{asset price}} = 0.9$ | | |
| Model | 4-quarters ahead | 12-quarters ahead |
| Output-Housing | Mixed and small | Negative and large |
| Output-Equity | Negative but small | Negative and large |
| Inflation-Housing | Negative and large | Negative and large |
| Inflation-Equity | Negative and small | Negative and large |

Table 3 summarizes the results reported in Figures 6a and 7a. All results are based on the panel bivariate QVAR models where output equation is estimated at the 10th percentile; inflation equation is estimated at the 90th percentile; and the asset price (housing or equity) equation is estimated at the 90th percentile. An asset price effect is large if the departure of the QVAR prediction is far away from the corresponding QAR prediction, the effect is small if the departure is small.

V. Conclusions

When considering the impact of asset price booms on output and inflation, it is the worst outcomes that matter. Just as risk-managers in a financial institution concern themselves with reducing the probability of significant loss, macroeconomic policymakers worry about mitigating the possibility of growth falling well below trend or inflation rising well above trend. Because of their focus on the conditional mean of the quantities being models, traditional time-series

econometric tools are ill-equipped to address the questions that are foremost in the minds of policymakers who adopt such a risk management perspective.

In this paper, we have fashioned and applied a new set of tools designed to help address this deficiency. First, we have incorporated the methods of quantile regression to vector autoregression. The resulting procedure allows us to investigate the dynamic impact of a change in a conditioning variable on the tail of the distribution of an independent variable over a particular horizon. So, for example, we are able to estimate the impact of an asset price boom today on the lower tail of the distribution of output deviation from trend several years ahead. We show how to estimate quantile vector autoregressions (QVARs), how to use the QVARs for prediction, and how to conduct inference on these estimates and predictions.

With these methods in hand, we proceed to examine the impact of housing and equity booms on the tails of the distribution of (log) output deviations from trend and (log) price-level deviations from trend. We conclude that in virtually every case, asset price booms worsen the worst outcomes significantly. At horizons of 3 year years, both housing and equity booms lower the 10th percentile of the output gap distribution and raise the 90th percentile of the price-level gap distribution.

References

- Bernanke, Ben and Mark Gertler (1999). "Monetary Policy and Asset Price Volatility". In *New Challenges for Monetary Policy: A Symposium Sponsored by the Federal Reserve Bank of Kansas City*. Federal Reserve Bank of Kansas City, pp. 77-128.
- _____ and _____. (2001). "Should Central Banks Respond to Movements in Asset Prices?" *American Economic Review*, May.
- Bertaut, Carol "Equity Prices, Household Wealth, and Consumption Growth in Foreign Industrial Countries," International Finance Discussion Papers 724. Washington: Board of Governors of the Federal Reserve System, 2002.
- Case, Karl E., John M. Quigley and Robert J. Shiller, "Comparing Wealth Effects: The Stock Market versus the Housing Market," *Advances in Macroeconomics*, vol. 5, no. 1 (2005).
- Cecchetti, Stephen G. "Measuring the Macroeconomic Risks Posed by Asset Price Booms," in John Y. Campbell, editor, *Asset Prices and Monetary Policy*, Chicago, Ill.: University of Chicago Press for NBER, 2008.
- _____, 'The Brave New World of Central Banking: Policy Challenges Posed by Asset Price Booms and Busts,' *National Institute Economic Review*, vol. 196, April 2006.
- _____, Hans Genberg, John Lipsky and Sushil Wadhvani. *Asset Prices and Central Bank Policy*. Geneva Reports on the World Economy, No. 2, International Center for Monetary and Banking Studies and Centre for Economic Policy Research, July 2000.
- Greenspan, Alan, "Risk and Uncertainty in Monetary Policy," *American Economic Review*, volume 94, no. 2, May 2004, pg. 33-40.
- Hodrick, Robert J. and Edward C. Prescott, "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit, and Banking*, 29 (1), February 1997, 1-16.
- Huber, P. J., "The Behavior of Maximum Likelihood Estimates under Nonstandard Conditions." Proceedings of the Fifth Berkeley Symposium 1, 1967, 221-33. University of California Press.
- Jarociński, Marek and Frank Smets, "House Prices and the Stance of Monetary Policy," presented at the 32nd annual economic policy conference, Federal Reserve Bank of St. Louis, October 2007.
- Jorion, Phillipe. *Value at Risk*. 2nd edition. New York: McGraw-Hill, 2001.
- Koenker, Roger. *Quantile Regression*. Cambridge, U.K. 2005.
- Koenker, Roger and Kevin F. Hallock, "Quantile Regression," *Journal of Economic Perspectives*, vol. 15, no. 4, Fall 2001, 143-156.
- Koenker, Roger and Beum J. Park, "An Interior Point Algorithm for Nonlinear Quantile Regression," *Journal of Econometrics*, vol. 71, no. 1-2, March-April 1996 pg. 265-283.

Leamer, Edward E., "Housing Is the Business Cycle," NBER Working Paper No. 13428, September 2007.

Newey, Witney K. and Kenneth D. West, "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix" *Econometrica*, vol. 55, no. 3, May 1987, 703–708.

Surico, Paolo, "U.S. Monetary Policy Rules: The Case for Asymmetric Preferences," *Journal of Economic Dynamics and Control*, forthcoming.

Appendix A: Data

Price Data: Computed for consumer price inflation data was obtained from the *International Financial Statistics* on line and the OECD Economic Outlook No. 76, December 2004.

GDP data was obtained from the *International Financial Statistics CDRom* (December 2004) and the OECD Economic Outlook No. 76, December 2004.

Equity Prices are from the *International Financial Statistics* on line.

Housing Prices: Data for Australia, Belgium, Canada, Denmark, Finland, Ireland, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, U.K, and U.S. 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, on line. Data for Japan are from Goldman Sachs. Data for New Zealand are from the Reserve Bank of New Zealand.

Appendix B: Derivations

B.I: Distribution of single-country QVAR estimator

In this appendix we derive the asymptotic distribution of the coefficient vector in the quantile vector autoregression in text equation (11). We then proceed to construct the estimator for the covariance matrix of the estimated coefficient vector in the QVAR.

Assuming the regressors are weakly stationary and that the density of the error ε_t conditional on X_t at a given quantile τ is continuous, it can be shown that the approximate first order condition (F.O.C.) of the first objective function in (10), the output equation, is

$$(B.1) \quad T^{-1} \sum_{t=1}^T (\tau_1 - 1_{y_t < x_t' \hat{\gamma}(\tau_1)}) x_t = 0,$$

which contains K equations, $K = \dim(x_t)$. $1_{(\cdot)}$ is an indicator function defined following text equation (4'). Similarly, the approximate F.O.C. of the second objective function in (10), the housing equation, is

$$(B.2) \quad T^{-1} \sum_{t=1}^T (\tau_2 - 1_{h_t < x_t' \hat{\beta}(\tau_2)}) x_t = 0.$$

Because we wish to derive the joint asymptotic distribution of $(\hat{\gamma}(\tau_1)' \hat{\beta}(\tau_2))'$, we stack the two

sets of F.O.C.s in (B.1) and (B.2),

$$(B.3) \quad T^{-1} \sum_{t=1}^T \begin{bmatrix} \tau_1 - 1_{y_t < x_t' \hat{\gamma}(\tau_1)} \\ \tau_2 - 1_{h_t < x_t' \hat{\beta}(\tau_2)} \end{bmatrix} \otimes x_t = 0$$

where " \otimes " denotes the Kronecker product. It can be shown that the stacked F.O.C.s, as specified above, implies a moment function that fits into the GMM framework. To streamline

notation, define the joint moment function as $g_t(\gamma(\tau_1), \beta(\tau_2)) = \begin{bmatrix} \tau_1 - 1_{y_t < x_t' \gamma(\tau_1)} \\ \tau_2 - 1_{h_t < x_t' \beta(\tau_2)} \end{bmatrix} \otimes x_t$. Then,

$$(B.4) \quad T^{-1} \sum_{t=1}^T g_t(\hat{\gamma}(\tau_1), \hat{\beta}(\tau_2)) = 0,$$

which is text equation (11). It is straightforward to show that under certain regularity conditions $E[g_t(\hat{\gamma}(\tau_1), \hat{\beta}(\tau_2))] = 0$. This establishes the validity of $g_t(\cdot, \cdot)$ as a moment function. The GMM framework can then be used to establish asymptotic normality of the coefficient estimates. The resulting joint distribution of $(\hat{\gamma}(\tau_1)' \hat{\beta}(\tau_2))'$ takes the form of

$$(B.5) \quad \sqrt{T} \begin{pmatrix} \hat{\gamma}(\tau_1) - \gamma(\tau_1) \\ \hat{\beta}(\tau_2) - \beta(\tau_2) \end{pmatrix} \Rightarrow \mathcal{N}(0, \Gamma_0^{-1} \Gamma_1 \Gamma_0^{-1})$$

The expressions of Γ_0 and Γ_1 in (B.5) are derived as follows. As in the standard GMM setup, the matrix Γ_0 is the expected first derivative of the moment function. It can be shown that

$$(B.6) \quad \Gamma_0 = \frac{\partial E[g_t(\gamma(\tau_1), \beta(\tau_2))]}{\partial(\gamma', \beta')} \\ = \begin{bmatrix} E(f_{\varepsilon_{1t}(\tau_1)}|x_t(0))x_t x_t' & 0 \\ 0 & E(f_{\varepsilon_{2t}(\tau_2)}|x_t(0))x_t x_t' \end{bmatrix}$$

Where $f_{\varepsilon_{1t}(\tau_1)}(0)$ and $f_{\varepsilon_{2t}(\tau_2)}(0)$ are conditional densities of $\varepsilon_{1t}(\tau_1)$ and $\varepsilon_{2t}(\tau_2)$, conditional on evaluating the regressor vector x_t at zero. Note that if the densities $f_{\varepsilon_{it}(\tau_i)}(0)$ ($i = 1, 2$) depend on the set of regressors, then we refer to this as "heteroskedasticity" in the context of quantile regressions.¹¹

¹¹ Note that if the conditional densities in (B.2), the $f(\cdot)$'s, are independent of the regressors, then they can be factored out of the expectation and the $\beta(\tau)$'s will all be the same and the quantile regression collapses to OLS.

Matrix Γ_1 in (B.5) results from a Central Limit Theorem of the moment restriction,

$$(B.7) \quad T^{-1/2} \sum_{t=1}^T g_t(\gamma(\tau_1), \beta(\tau_2)) \Rightarrow \mathcal{N}(0, \Gamma_1)$$

So Γ_1 is the asymptotic variance of the scaled sample moment condition. In the case where the quantile regression errors, the $\varepsilon_{it}(\tau)$'s are serially uncorrelated, Γ_1 is simply the variance of $g_t(\gamma(\tau_1), \beta(\tau_2))$ given by

$$(B.8) \quad \begin{aligned} \lambda_0 &= E[g_t(\gamma(\tau_1), \beta(\tau_2)) g_t(\gamma(\tau_1), \beta(\tau_2))'] \\ &= \begin{bmatrix} \lambda_0^{11} & \lambda_0^{12} \\ \lambda_0^{21} & \lambda_0^{22} \end{bmatrix} \end{aligned}$$

where

$$(B.9) \quad \begin{aligned} \lambda_0^{11} &= E [(\tau_1 - 1_{y_t < x_t' \gamma(\tau_1)})^2 x_t x_t'] \\ &= E [(\tau_1 2 + (1 - 2\tau_1) 1_{\varepsilon_{1t} < 0}) x_t x_t'] \\ &= \tau_1 (1 - \tau_1) E(x_t x_t') \\ \lambda_0^{22} &= E [(\tau_2 - 1_{h_t < x_t' \beta(\tau_2)})^2 x_t x_t'] \\ &= \tau_2 (1 - \tau_2) E(x_t x_t') \\ \lambda_0^{12} &= E [(\tau_1 - 1_{y_t < x_t' \gamma(\tau_1)}) (\tau_2 - 1_{h_t < x_t' \beta(\tau_2)}) x_t x_t'] \\ &= E [(\tau_1 \tau_2 - \tau_1 1_{\varepsilon_{2t} < 0} - \tau_2 1_{\varepsilon_{1t} < 0} + 1_{\varepsilon_{1t} < 0} 1_{\varepsilon_{2t} < 0}) x_t x_t'] \\ &= E [(F_{\varepsilon_{1t}(\tau_1), \varepsilon_{2t}(\tau_2) | x_t}(0, 0) - \tau_1 \tau_2) x_t x_t'] \end{aligned}$$

where in (B.9), we use the law of iterated expectations together with the following results:

$$(B.10) \quad \begin{aligned} E(1_{\varepsilon_{1t}(\tau_1) < 0} | x_t) &= \int_{-\infty}^0 f_{\varepsilon_{1t}(\tau_1) | x_t}(\eta) d\eta = \tau_1 \\ E(1_{\varepsilon_{2t}(\tau_2) < 0} | x_t) &= \int_{-\infty}^0 f_{\varepsilon_{2t}(\tau_2) | x_t}(\eta) d\eta = \tau_2 \\ E(1_{\varepsilon_{1t}(\tau_1) < 0} 1_{\varepsilon_{2t}(\tau_2) < 0} | x_t) &= \int_{-\infty}^0 \int_{-\infty}^0 f_{\varepsilon_{1t}(\tau_1), \varepsilon_{2t}(\tau_2) | x_t}(\eta_1, \eta_2) d\eta_1 d\eta_2 = F_{\varepsilon_{1t}(\tau_1), \varepsilon_{2t}(\tau_2) | x_t}(0, 0) \end{aligned}$$

where $F_{\varepsilon_{1t}(\tau_1), \varepsilon_{2t}(\tau_2) | x_t}(0, 0)$ is the conditional joint CDF of $\varepsilon_{1t}(\tau_1)$ and $\varepsilon_{1t}(\tau_2)$, conditional on evaluating the regressor vector x_t , at zeros.

Substituting (B.9) into (B.8), we get the final expression for the covariance matrix of the quantile regression coefficient estimates, $\hat{\gamma}(\tau_1)$ and $\hat{\beta}(\tau_2)$,

$$(B.11) \quad \lambda_0 = \begin{bmatrix} \tau_1(1 - \tau_1)E(x_t x_t') & E[(F_{\varepsilon_{1t}(\tau_1), \varepsilon_{2t}(\tau_2)}|x_t(0,0) - \tau_1\tau_2)x_t x_t'] \\ E[(F_{\varepsilon_{1t}(\tau_1), \varepsilon_{2t}(\tau_2)}|x_t(0,0) - \tau_1\tau_2)x_t x_t'] & \tau_1(1 - \tau_1)E(x_t x_t') \end{bmatrix}$$

where F is defined in equation (B.10).

However, when time series data are used, it is likely that the moment function will exhibit serial correlation. Ignoring this would lead to bias and inconsistency in the estimator for the covariance matrix of the quantile regression coefficient estimates.. To remedy this, we need to Γ_1 as the long-run variance of the moment condition. That is,

$$(B.12) \quad \begin{aligned} \Gamma_1 &= LRvar(g_t(\gamma(\tau_1), \beta(\tau_2))) \\ &= \lim_{T \rightarrow \infty} \left(T^{-1/2} \sum_{t=1}^T g_t(\gamma(\tau_1), \beta(\tau_2)) \right) \left(T^{-1/2} \sum_{t=1}^T g_t(\gamma(\tau_1), \beta(\tau_2)) \right)' \\ &= \lambda_0 + \sum_{j=1}^{\infty} (\lambda_j + \lambda_j') \end{aligned}$$

where λ_0 is given in (B.11). Similarly, it can be shown that λ_j (for $j = 1, 2, \dots$), which captures the autocorrelation of the moment conditions, is

$$(B.13) \quad \begin{aligned} \lambda_j &= E \left(g_t(\gamma(\tau_1), \beta(\tau_2)) g_{t-j}(\gamma(\tau_1), \beta(\tau_2))' \right) \\ &= E \left(\begin{bmatrix} F_{\varepsilon_{1t}(\tau_1), \varepsilon_{1t-j}(\tau_1)}|x_t, x_{t-j}(0,0) - \tau_1^2 & F_{\varepsilon_{1t}(\tau_1), \varepsilon_{2t-j}(\tau_2)}|x_t, x_{t-j}(0,0) - \tau_1\tau_2 \\ F_{\varepsilon_{1t-j}(\tau_1), \varepsilon_{2t}(\tau_2)}|x_t, x_{t-j}(0,0) - \tau_1\tau_2 & F_{\varepsilon_{2t}(\tau_2), \varepsilon_{2t-j}(\tau_2)}|x_t, x_{t-j}(0,0) - \tau_2^2 \end{bmatrix} \otimes x_t x_{t-j}' \right) \end{aligned}$$

where $F_{\varepsilon_{mt}(\tau_m), \varepsilon_{ns}(\tau_n)}|x_t, x_s(0,0)$ ($m = 1, 2$ and $n = 1, 2$) denotes the conditional joint CDF of the errors $\varepsilon_{mt}(\tau_m)$ and $\varepsilon_{ns}(\tau_n)$ conditional on regressor vectors x_t and x_s , evaluated at (0,0).

Combining (B.6) and (B.12), the joint asymptotic distribution in (B.5) becomes

$$(B.14) \quad \sqrt{T} \begin{pmatrix} \hat{\gamma}(\tau_1) - \gamma(\tau_1) \\ \hat{\beta}(\tau_2) - \beta(\tau_2) \end{pmatrix} \Rightarrow \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \text{avar}(\hat{\gamma}(\tau_1)) & \text{acov}(\hat{\gamma}(\tau_1), \hat{\beta}(\tau_2)) \\ \text{acov}(\hat{\beta}(\tau_2), \hat{\gamma}(\tau_1)) & \text{avar}(\hat{\beta}(\tau_2)) \end{bmatrix} \right)$$

where the population long-run variance and covariance matrices have the following

expressions:

$$\begin{aligned} \text{avar}(\widehat{\gamma}(\tau_1)) &= \Gamma_{\gamma 0}^{-1} \Gamma_{\gamma 1} \Gamma_{\gamma 0}^{-1} \\ \Gamma_{\gamma 0} &= E[f_{\varepsilon_{1t}(\tau_1)|x_t}(0) x_t x_t'] \\ \Gamma_{\gamma 1} &= \tau_1(1 - \tau_1) E[x_t x_t'] \\ &\quad + \sum_{j=1}^{\infty} E[(F_{\varepsilon_{1t}(\tau_1), \varepsilon_{1t-j}(\tau_1)|x_t, x_{t-j}}(0, 0) - \tau_1^2)(x_t x_{t-j}' + x_{t-j} x_t')] \end{aligned}$$

$$\begin{aligned} \text{avar}(\widehat{\beta}(\tau_2)) &= \Gamma_{\beta 0}^{-1} \Gamma_{\beta 1} \Gamma_{\beta 0}^{-1} \\ \Gamma_{\beta 0} &= E[f_{\varepsilon_{2t}(\tau_2)|x_t}(0) x_t x_t'] \\ \Gamma_{\beta 1} &= \tau_2(1 - \tau_2) E[x_t x_t'] \\ &\quad + \sum_{j=1}^{\infty} E[(F_{\varepsilon_{2t}(\tau_2), \varepsilon_{2t-j}(\tau_2)|x_t, x_{t-j}}(0, 0) - \tau_2^2)(x_t x_{t-j}' + x_{t-j} x_t')] \end{aligned}$$

$$\begin{aligned} \text{acov}(\widehat{\gamma}(\tau_1), \widehat{\beta}(\tau_2)) &= \Gamma_{\gamma 0}^{-1} \Gamma_{\gamma \beta} \Gamma_{\beta 0}^{-1} \\ \Gamma_{\gamma \beta} &= E[(F_{\varepsilon_{1t}(\tau_1), \varepsilon_{2t}(\tau_2)|x_t}(0, 0) - \tau_1 \tau_2) x_t x_t'] \\ &\quad + \sum_{j=1}^{\infty} \{ E[(F_{\varepsilon_{1t}(\tau_1), \varepsilon_{2t-j}(\tau_2)|x_t, x_{t-j}}(0, 0) - \tau_1 \tau_2) x_t x_{t-j}'] \} \\ &\quad + \sum_{j=1}^{\infty} \{ E[(F_{\varepsilon_{1t-j}(\tau_1), \varepsilon_{2t}(\tau_2)|x_t, x_{t-j}}(0, 0) - \tau_1 \tau_2) x_{t-j} x_t'] \} \end{aligned}$$

In practice, the above estimators of the variances and covariance matrices can be obtained by replacing the population moments by their sample analogs. Since expressions for $\Gamma_{\gamma 1}$, $\Gamma_{\beta 1}$ and $\Gamma_{\gamma \beta}$ involve infinite sums, estimators of these matrices can be computed by employing the Newey-West method commonly used in conditional mean regressions. Throughout our study, the Bartlett kernel, with $p = 0.75T^{1/3}$, is used to determine the bandwidth. Thus, we have:

$$\begin{aligned}
 \widehat{\Gamma}_{\gamma 0} &= T^{-1} \sum_{t=1}^T \widehat{f}_{\varepsilon_{1t}(\tau_1)|x_t}(0) x_t x_t' \\
 \widehat{\Gamma}_{\beta 0} &= T^{-1} \sum_{t=1}^T \widehat{f}_{\varepsilon_{2t}(\tau_2)|x_t}(0) x_t x_t' \\
 \widehat{\Gamma}_{\gamma 1} &= \tau_1(1 - \tau_1) T^{-1} \sum_{t=1}^T x_t x_t' \\
 &\quad + \sum_{j=1}^{p-1} (1 - j/p) \left\{ (T - j)^{-1} \sum_{t=j+1}^T \left[(1_{\widehat{\varepsilon}_{1t}(\tau_1) < 0} 1_{\widehat{\varepsilon}_{1t-j}(\tau_1) < 0} - \tau_1^2) (x_t x_{t-j}' + x_{t-j} x_t') \right] \right\} \\
 \widehat{\Gamma}_{\beta 1} &= \tau_2(1 - \tau_2) T^{-1} \sum_{t=1}^T x_t x_t' \\
 &\quad + \sum_{j=1}^{p-1} (1 - j/p) \left\{ (T - j)^{-1} \sum_{t=j+1}^T \left[(1_{\widehat{\varepsilon}_{2t}(\tau_2) < 0} 1_{\widehat{\varepsilon}_{2t-j}(\tau_2) < 0} - \tau_2^2) (x_t x_{t-j}' + x_{t-j} x_t') \right] \right\} \\
 \widehat{\Gamma}_{\gamma\beta} &= T^{-1} \sum_{t=1}^T (1_{\widehat{\varepsilon}_{1t}(\tau_1) < 0} 1_{\widehat{\varepsilon}_{2t}(\tau_2) < 0} - \tau_1 \tau_2) x_t x_t' \\
 &\quad + \sum_{j=1}^{p-1} (1 - j/p) \left\{ (T - j)^{-1} \sum_{t=j+1}^T \left[(1_{\widehat{\varepsilon}_{1t}(\tau_1) < 0} 1_{\widehat{\varepsilon}_{2t-j}(\tau_2) < 0} - \tau_1 \tau_2) x_t x_{t-j}' \right] \right\} \\
 &\quad + \sum_{j=1}^{p-1} (1 - j/p) \left\{ (T - j)^{-1} \sum_{t=j+1}^T \left[(1_{\widehat{\varepsilon}_{1t-j}(\tau_1) < 0} 1_{\widehat{\varepsilon}_{2t}(\tau_2) < 0} - \tau_1 \tau_2) x_{t-j} x_t' \right] \right\}
 \end{aligned}$$

where the residuals $\widehat{\varepsilon}_{1t}(\tau_1)$ and $\widehat{\varepsilon}_{2t}(\tau_2)$ are computed by $\widehat{\varepsilon}_{1t}(\tau_1) = y_t - x_t' \widehat{\gamma}(\tau_1)$,

$\widehat{\varepsilon}_{2t}(\tau_2) = h_t - x_t' \widehat{\beta}(\tau_2)$; and estimation of the conditional densities follows Koenker (2005). To be

specific, the densities $\widehat{f}_{\varepsilon_{1t}|x_t}(0)$ and $\widehat{f}_{\varepsilon_{2t}|x_t}(0)$ are estimated as

$$\widehat{f}_{\varepsilon_{i,t}(\tau_i)|x_t}(0) = \frac{2h_i}{x_t'(\widehat{\gamma}(\tau_i + h_i) - \widehat{\gamma}(\tau_i - h_i))}$$

for $i = 1, 2$. h_i 's are the optimal bandwidths given by $h_i = 0.6T^{-1/5} \left[4.5 \frac{\phi 4(\Phi^{-1}(\tau_i))}{(2(\Phi^{-1}(\tau_i))^2 + 1)2} \right]^{1/5}$

where ϕ and Φ are the Gaussian density and CDF respectively. See Koenker (2005) for this choice of optimal bandwidth. The estimated density in (B.16) might not be positive, due to "crossing" of the estimated conditional quantile plane. (But in practice this problem occurs only infrequently and in the most extreme regions of the design space.) Thus the refined estimator of Hendricks and Koenker (1991) is used. In our context, that means replacing the estimated densities in (10) by their positive

parts,

$$(B.17) \quad \hat{f}_{\varepsilon_{i,t}(\tau_i)|x_t}^+(0) = \max \left\{ 0, \frac{2h_i}{[x_t'(\hat{\gamma}(\tau_i + h_i) - \hat{\gamma}(\tau_i - h_i))] - \eta} \right\}$$

where η is a small tolerance number intended to avoid dividing by zero in the rare cases.¹²

We note that the methods used here to derive the distribution of the QVAR estimator allow for both dependence of the distributions of error terms on the regressor vector, and serial correlations of the moment function. Furthermore, the estimator for the covariance matrix of the coefficient estimates is heteroskedasticity and autocorrelation consistent (HAC).

B.2: Extension to panel QVAR

To extend the techniques of Appendix B.1 to the panel case, we stack the data country by country and introduce the fixed effects by including a constant matrix. We restrict attention to balanced panel so that $T_i = T$ for $i = 1, \dots, n$. Take the first QVAR equation, the output equation, as an example. For a single quantile τ , the panel regression in matrix form can be written as

$$(B.18) \quad \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} = \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_n \end{bmatrix} \gamma(\tau) + \begin{bmatrix} \iota_T & 0 & \cdots & 0 \\ 0 & \iota_T & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \iota_T \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n \end{bmatrix} + \begin{bmatrix} \varepsilon_1(\tau) \\ \varepsilon_2(\tau) \\ \vdots \\ \varepsilon_n(\tau) \end{bmatrix}$$

where y_i is an $T \times 1$ vector of output for country i , X_i is a $T \times K$ matrix, $\varepsilon_i(\tau)$ is an $T \times 1$ residual vector, and ι_T is an $T \times 1$ vector of ones. We write this stacked equation system more compactly as

$$(B.19) \quad y = X\gamma(\tau) + Z\alpha + \varepsilon(\tau)$$

where $Z = I_n \otimes \iota_T$, y is $nT \times 1$, X is $nT \times K$, Z is $nT \times n$ and $\varepsilon(\tau)$ is $nT \times 1$.

¹² In principle it would be possible to implement an estimation technique in which the estimate quantiles never crossed. This is topic for future research.

To estimate the model for J quantiles simultaneously, we further stack the model (B.19) quantile by quantile and constrain the fixed effect to be constant across quantiles,

$$(B.20) \quad \begin{bmatrix} y \\ y \\ \vdots \\ y \end{bmatrix} = \begin{bmatrix} X & 0 & \cdots & 0 \\ 0 & X & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & X \end{bmatrix} \begin{bmatrix} \gamma(\tau_1) \\ \gamma(\tau_2) \\ \vdots \\ \gamma(\tau_J) \end{bmatrix} + \begin{bmatrix} Z \\ Z \\ \vdots \\ Z \end{bmatrix} \alpha + \begin{bmatrix} \varepsilon(\tau_1) \\ \varepsilon(\tau_2) \\ \vdots \\ \varepsilon(\tau_J) \end{bmatrix},$$

or equivalently,

$$(B.21) \quad (\iota_J \otimes y) = (I_J \otimes X)\gamma(\tau_1, \dots, \tau_J) + (\iota_J \otimes Z)\alpha + \varepsilon(\tau_1, \dots, \tau_J)$$

where ι_J is an $J \times 1$ vector of ones, $\gamma(\tau_1, \dots, \tau_J) = [\gamma(\tau_1)' \ \gamma(\tau_2)' \ \dots \ \gamma(\tau_J)']'$ and $\varepsilon(\tau_1, \dots, \tau_J) = [\varepsilon(\tau_1)' \ \varepsilon(\tau_2)' \ \dots \ \varepsilon(\tau_J)']'$.

From (B.21) we see that implementing the method of Koenker and Park (1996) for panel data requires only that we redefine the response vector to be the $JnT \times 1$ vector $(\iota_J \otimes y)$ and the regressor matrix to be the $JnT \times (JK + n)$ matrix $[(I_J \otimes X) : (\iota_J \otimes Z)]$. In this way, the panel QVAR estimation problem becomes exactly the same as the one we examine in appendix B1.

To obtain the asymptotic distribution of the panel QVAR, we extend the approach used by Koenker (2004) in his analysis of the single equation panel model in Koenker (2004). Let

$$(B.22) \quad M_{\tilde{Z}(\tau_j)} = I - Z(Z'F(\tau_j)Z)^{-1}Z'F(\tau_j)$$

where $F(\tau_j) = \text{diag}(f_{\varepsilon_{it}(\tau_j)|x_{it}}(0))$ in which $f_{\varepsilon_{it}(\tau_j)|x_{it}}(\cdot)$ is the conditional density at quantile τ_j evaluated at zero. $M_{\tilde{Z}(\tau_j)}$ is the residual matrix with respect to the fixed effects in a quantile regression setup. Pre-multiplying the regressor matrix X by $M_{\tilde{Z}(\tau_j)}$ will partial out the fixed effects:

$$(B.23) \quad \tilde{X}(\tau_j) = M_{\tilde{Z}(\tau_j)}X.$$

Then, the joint asymptotic distribution of $\hat{\gamma}(\tau_1)$ and $\hat{\beta}(\tau_2)$ is given by

$$(B.24) \quad \sqrt{T} \begin{pmatrix} \hat{\gamma}(\tau_1) - \gamma(\tau_1) \\ \hat{\beta}(\tau_2) - \beta(\tau_2) \end{pmatrix} \Rightarrow \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \text{avar}(\hat{\gamma}(\tau_1)) & \text{acov}(\hat{\gamma}(\tau_1), \hat{\beta}(\tau_2)) \\ \text{acov}(\hat{\beta}(\tau_2), \hat{\gamma}(\tau_1)) & \text{avar}(\hat{\beta}(\tau_2)) \end{bmatrix} \right)$$

where

$$\begin{aligned} \text{avar}(\widehat{\gamma}(\tau_1)) &= \Gamma_{\gamma 0}^{-1} \Gamma_{\gamma 1} \Gamma_{\gamma 0}^{-1} \\ \Gamma_{\gamma 0} &= E[f_{\varepsilon_{1it}(\tau_1)} | x_{it}(0) \tilde{x}_{it}(\tau_1) \tilde{x}_{it}(\tau_1)'] \\ \Gamma_{\gamma 1} &= \tau_1 (1 - \tau_1) E[\tilde{x}_{it}(\tau_1) \tilde{x}_{it}(\tau_1)'] \\ &\quad + \sum_{j=1}^{\infty} E \left[(F_{\varepsilon_{1it}(\tau_1), \varepsilon_{1it-j}(\tau_1)} | x_{it}, x_{it-j}(0, 0) - \tau_1^2) (\tilde{x}_{it}(\tau_1) \tilde{x}_{it-j}(\tau_1)' + \tilde{x}_{it-j}(\tau_1) \tilde{x}_{it}(\tau_1)') \right] \end{aligned}$$

$$\begin{aligned} \text{avar}(\widehat{\beta}(\tau_2)) &= \Gamma_{\beta 0}^{-1} \Gamma_{\beta 1} \Gamma_{\beta 0}^{-1} \\ \Gamma_{\beta 0} &= E[f_{\varepsilon_{2it}(\tau_2)} | x_{it}(0) \tilde{x}_{it}(\tau_2) \tilde{x}_{it}(\tau_2)'] \\ \Gamma_{\beta 1} &= \tau_2 (1 - \tau_2) E[\tilde{x}_{it}(\tau_2) \tilde{x}_{it}(\tau_2)'] \\ &\quad + \sum_{j=1}^{\infty} E \left[(F_{\varepsilon_{2it}(\tau_2), \varepsilon_{2it-j}(\tau_2)} | x_{it}, x_{it-j}(0, 0) - \tau_2^2) (\tilde{x}_{it}(\tau_2) \tilde{x}_{it-j}(\tau_2)' + \tilde{x}_{it-j}(\tau_2) \tilde{x}_{it}(\tau_2)') \right] \end{aligned}$$

$$\begin{aligned} \text{acov}(\widehat{\gamma}(\tau_1), \widehat{\beta}(\tau_2)) &= \Gamma_{\gamma 0}^{-1} \Gamma_{\gamma \beta} \Gamma_{\beta 0}^{-1} \\ \Gamma_{\gamma \beta} &= E[(F_{\varepsilon_{1it}(\tau_1), \varepsilon_{2it}(\tau_2)} | x_{it}(0, 0) - \tau_1 \tau_2) \tilde{x}_{it}(\tau_1) \tilde{x}_{it}(\tau_2)'] \\ &\quad + \sum_{j=1}^{\infty} \left\{ E \left[(F_{\varepsilon_{1it}(\tau_1), \varepsilon_{2it-j}(\tau_2)} | x_{it}, x_{it-j}(0, 0) - \tau_1 \tau_2) \tilde{x}_{it}(\tau_1) \tilde{x}_{it-j}(\tau_2)' \right] \right\} \\ &\quad + \sum_{j=1}^{\infty} \left\{ E \left[(F_{\varepsilon_{1it-j}(\tau_1), \varepsilon_{2it}(\tau_2)} | x_{it}, x_{it-j}(0, 0) - \tau_1 \tau_2) \tilde{x}_{it-j}(\tau_1) \tilde{x}_{it}(\tau_2)' \right] \right\} \end{aligned}$$

where $\tilde{x}_{it}(\tau_j)$ and $\tilde{x}_{it-j}(\tau_j)$ are the (i, t) and $(i, t - j)$ elements of the transformed matrix $\tilde{X}(\tau_j)$. A feasible version of the variance-covariance matrix can be constructed by substituting the population means by their sample analogs, for instance, $E[\tilde{x}_{it}(\tau_j) \tilde{x}_{it}(\tau_j)']$ is replaced by $(nT)^{-1} \sum_{i=1}^n \sum_{t=1}^T \tilde{x}_{it}(\tau_j) \tilde{x}_{it}(\tau_j)'$.

Appendix B.3: Uncertainty of the forecasts

In this final appendix, we show how to construct estimates of the variance of the forecasts obtained from the quantile vector autoregression. To begin, consider equation (18) of Section II.2. Take the two-equation four-lag QVAR in output and housing as an example. Suppose the VAR takes the form

$$\begin{aligned} \text{(B.25)} \quad y_t &= \gamma_0 + \gamma_{11}y_{t-1} + \dots + \gamma_{14}y_{t-4} + \gamma_{21}h_{t-1} + \dots + \gamma_{24}h_{t-4} + \varepsilon_{1t} \\ h_t &= \beta_0 + \beta_{11}y_{t-1} + \dots + \beta_{14}y_{t-4} + \beta_{21}h_{t-1} + \dots + \beta_{24}h_{t-4} + \varepsilon_{2t} \end{aligned}$$

Furthermore, suppose that output equation and housing equation are estimated at quantiles τ_1 , τ_2 respectively. Then the estimated $A_0(\tau_1, \tau_2)$ and $A_1(\tau_1, \tau_2)$ in equation (18) are

$$(B.26) \quad \widehat{A}_0(\tau_1, \tau_2) = \begin{pmatrix} \widehat{\gamma}_0(\tau_1) \\ 0 \\ 0 \\ 0 \\ \widehat{\beta}_0(\tau_2) \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

and

$$(B.27) \quad \widehat{A}_1(\tau_1, \tau_2) = \begin{pmatrix} \widehat{\gamma}_{11}(\tau_1) & \widehat{\gamma}_{12}(\tau_1) & \widehat{\gamma}_{13}(\tau_1) & \widehat{\gamma}_{14}(\tau_1) & \widehat{\gamma}_{21}(\tau_1) & \widehat{\gamma}_{22}(\tau_1) & \widehat{\gamma}_{23}(\tau_1) & \widehat{\gamma}_{24}(\tau_1) \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ \widehat{\beta}_{11}(\tau_2) & \widehat{\beta}_{12}(\tau_2) & \widehat{\beta}_{13}(\tau_2) & \widehat{\beta}_{14}(\tau_2) & \widehat{\beta}_{21}(\tau_2) & \widehat{\beta}_{22}(\tau_2) & \widehat{\beta}_{23}(\tau_2) & \widehat{\beta}_{24}(\tau_2) \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}$$

$\widehat{A}_0(\tau_1, \tau_2)$ and $\widehat{A}_1(\tau_1, \tau_2)$ are used in equation (19) to compute the k -period forecast for the τ_1 -th quantile of output conditional on a housing boom. From (19), \widetilde{y}_{t+k} is a nonlinear function of the coefficient matrices A_0 and A_1 . The δ -method can be used to compute the variance and standard error of the forecast.

But here a nontrivial technical issue is that it involves matrix differentiation. The techniques of matrix differentiation and their applications in econometrics are documented in Magnus and Neudecker (1999). Applying here, the variance of \widetilde{y}_{t+k} is computed as follows. To streamline notation, denote \widetilde{y}_{t+k} by \widetilde{y}_k and let

$$(B.28) \quad g(A_0, A_1) = (I + A_1 + A_1^2 + \dots + A_1^{k-1})A_0 + A_1^k Z_0.$$

The forecast in text equation (18) can be rewritten as

$$(B.29) \quad \widetilde{y}_k = e_1' g(\widehat{A}_0, \widehat{A}_1).$$

Then the variance of the forecast in (B.29) can be computed as

$$(B.30) \quad \text{var}(\widetilde{y}_k) = e_1' \text{var}(g(\widehat{A}_0, \widehat{A}_1)) e_1$$

where

$$(B.31) \quad \text{var}(g(\hat{A}_0, \hat{A}_1)) = \begin{bmatrix} D_{g0}(\hat{A}_0, \hat{A}_1) & D_{g1}(\hat{A}_0, \hat{A}_1) \end{bmatrix} \text{var} \begin{pmatrix} \text{vec} \hat{A}'_0 \\ \text{vec} \hat{A}'_1 \end{pmatrix} \begin{bmatrix} D'_{g0}(\hat{A}_0, \hat{A}_1) \\ D'_{g1}(\hat{A}_0, \hat{A}_1) \end{bmatrix}$$

where $\text{vec}(\cdot)$ denotes placing the second column below the first, the third column below the previous two, and so on. $D_{g0}(\hat{A}_0, \hat{A}_1)$ and $D_{g1}(\hat{A}_0, \hat{A}_1)$ are two matrix partial derivatives.

Next we provide the detailed derivation of D_{g0} and D_{g1} . To start, let $f(A_1) = \sum_{q=1}^{k-1} A_1^q$. Then,

$$(B.32) \quad g(A_0, A_1) = f(A_1) \times A_0 + A_1^k \times Z_0.$$

The first differential of the column vector $g(A_0, A_1)$ can be decomposed into

$$(B.33) \quad dg(A_0, A_1) = df(A_1) \times A_0 + f(A_1) \times dA_0 + dA_1^k \times Z_0$$

or, equivalently written in vec form,

$$(B.34) \quad \begin{aligned} & (dg(A_0, A_1)) \\ &= \text{vec}(df(A_1) \times A_0) + \text{vec}(f(A_1) \times dA_0) + \text{vec}(dA_1^k \times Z_0) \\ &= \text{vec}(A'_0 \times df(A_1)) + \text{vec}(dA'_0 \times f(A_1)) + \text{vec}(Z'_0 \times d(A_1^k)) \end{aligned}$$

where the second equality uses the following fact: if $A \times B$ is a column vector, then

$$\text{vec}(A \times B) = \text{vec}(A \times B)' = \text{vec}(B' \times A').$$

Next, compute the three right-hand-side terms in (B.34). The first term is

$$(B.35) \quad \begin{aligned} & \text{vec}(A'_0 \times df(A_1)) \\ &= (I \otimes A'_0) \times \text{vec}(df(A_1)) \\ &= (I \otimes A'_0) \times \frac{\text{vec}(df(A_1))}{(\text{vec}(dA_1))'} \times \text{vec}(dA_1) \end{aligned}$$

where the first order derivative in the middle of (B.35) has the expression

$$\begin{aligned}
 \text{(B.36)} \quad \frac{\text{vec}(df(A'_1))}{(\text{vec}(dA'_1))'} &= (A_1^{k-1} + A_1^{k-3} + A_1^{k-4} + \dots + A_1 2 + A_1 + I) \otimes I \\
 &\quad + (A_1^{k-3} + A_1^{k-4} + \dots + A_1 2 + A_1 + I) \otimes A'_1 \\
 &\quad + (A_1^{k-4} + \dots + A_1 2 + A_1 + I) \otimes (A'_1)^2 \\
 &\quad + \dots \\
 &\quad + (A_1 + I) \otimes (A'_1)^{k-3} \\
 &\quad + I \otimes (A'_1)^{k-2}
 \end{aligned}$$

The second term in (B.34) is

$$\text{(B.37)} \quad \text{vec}(dA'_0 \times f(A'_1)) = f(A_1) \times d(\text{vec}A'_0)$$

The last term in (B.34) is

$$\begin{aligned}
 \text{(B.38)} \quad &\text{vec}(Z'_0 \times d(A'_1)^k) \\
 &= (I \otimes Z'_0) \times \text{vec}(d(A'_1)^k) \\
 &= (I \otimes Z'_0) \times \frac{\text{vec}(d(A'_1)^k)}{(\text{vec}(dA'_1))'} \times \text{vec}(dA'_1)
 \end{aligned}$$

where the first order derivative in the middle of (B.38) is

$$\text{(B.39)} \quad \frac{\text{vec}(d(A'_1)^k)}{(\text{vec}(dA'_1))'} = A_1^{k-1} \otimes I + A_1^{k-2} \otimes A'_1 + \dots + A_1 \otimes (A'_1)^{k-1} + I \otimes (A'_1)^{k-1}$$

Combining (B.35), (B.37) and (B.38),

$$\begin{aligned}
 \text{(B.40)} \quad &dg(A_0, A_1) \\
 &= \left[(I \otimes A'_0) \times \frac{\text{vec}(df(A'_1))}{(\text{vec}(dA'_1))'} + (I \otimes Z'_0) \times \frac{\text{vec}(d(A'_1)^k)}{(\text{vec}(dA'_1))'} \right] \times \text{vec}(dA'_1) + f(A_1) \times dA_0
 \end{aligned}$$

From (B.40), we get

$$\begin{aligned}
 \text{(B.41)} \quad \frac{\partial g(A_0, A_1)}{\partial (\text{vec}A'_0)'} &= f(A_1) \\
 \frac{\partial g(A_0, A_1)}{\partial (\text{vec}A'_1)'} &= (I \otimes A'_0) \times \frac{\text{vec}(df(A'_1))}{(\text{vec}(dA'_1))'} + (I \otimes Z'_0) \times \frac{\text{vec}(d(A'_1)^k)}{(\text{vec}(dA'_1))'}
 \end{aligned}$$

Finally, substituting (B.36) and (B.39) into (B.41), we get

$$\begin{aligned}
 D_{g0}(A_0, A_1) &= \frac{\partial g(A_0, A_1)}{\partial (\text{vec} A_0)'} = \sum_{q=0}^{k-1} A_1^q \\
 D_{g1}(A_0, A_1) &= \frac{\partial g(A_0, A_1)}{\partial (\text{vec} A_1)'} \\
 \text{(B.42)} \quad &= (I \otimes A_0') \times \sum_{q=0}^{k-2} \left[\left(\sum_{i=0}^{k-2-q} A_1^i \right) \otimes (A_1')^q \right] + (I \otimes Z_0') \times \sum_{q=0}^{k-1} \left[A_1^{k-1-q} \otimes (A_1')^q \right]
 \end{aligned}$$

where $A_0 = A_0(\tau_1, \tau_2)$ and $A_1 = A_1(\tau_1, \tau_2)$. In practice, vector A_0 and matrix A_1 can be replaced by the consistent estimators in (B.26) and (B.27).