GARCH models with dummies

A study of the impact of U.S. monetary policy on inflation

Scott Deacle
rsdeacle@temple.edu

April 26, 2006

Outline

Media, Pa.,-raised Robert F. Engle

Engle introduced the autoregressive conditional heteroscedasticity (ARCH) model in a 1982 article in *Econometrica*. He developed it to represent economic situations in which a variable's volatility at a given time is important.

Friedman's Hypothesis

When Engle introduced the ARCH model, macroeconomists were looking for ways to study inflation volatility. They were inspired in part by Milton Friedman's Nobel Prize address (1977). Friedman proposed higher inflation causes greater inflation volatility. The higher volatility then has consequences that reduce real output as economic agents devote resources to dealing with the risk of inflation, Friedman argued.

In his 1982 article, Engle used an ARCH model to to study inflation in the United Kingdom.

Differences between ARCH and OLS

A simple regression model produces a constant unconditional variance of the independent variable. The ARCH model produces both the unconditional variance and a process for the time-varying conditional variance.

The ARCH(q) model

The model is specified as follows: Let y represent an independent variable, $\boldsymbol{x}$ a vector of dependent variables, **b** a vector of parameters and ϵ the error term in a regression model. The time period is represented by t. Let σ_t^2 represent the time-varying conditional variance, q the number of lagged terms and $\alpha = (\alpha_0, \alpha_1, \ldots, \alpha_q)$ a vector of parameters in the conditional variance model. In ARCH(q) models, the conditional mean and conditional variance of the dependent variable are written:

$$y_t = \boldsymbol{x}_t' \boldsymbol{b} + \epsilon_t \tag{1}$$

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{q} \alpha_i \epsilon_{t-i}^2 \tag{2}$$

The Conditional Variance Term

- The conditional variance is not heteroscedastic with respect to x. It is heteroscedastic with respect to its q lagged values.
- The conditional variance grows and shrinks according to the magnitude of past shocks. This makes the model realistic for variables for which large shocks occur in clusters. It gives researchers a measure of the volatility of the dependent variable at each observation.

Stationarity and unconditional variance

The error process is assumed to be weakly stationary.

$$E[\epsilon_t] = 0 \tag{3}$$

$$Var[\epsilon_t] = \sigma_t^2 \tag{4}$$

$$Cov[\epsilon_t, \epsilon_s] = 0, t \neq s \tag{5}$$

It can also be shown that the unconditional variance is

$$Var[\epsilon_t] = \sigma_t^2 = \frac{\alpha_0}{1 - \sum_{i=1}^{q} \alpha_i} \tag{6}$$

Detecting ARCH

The presence of an ARCH process may be detected through visual and analytic means. For starters, one may inspect a time series plot of the dependent variable for clusters of large movements.

U.S. Inflation, 1948:2-2005:4

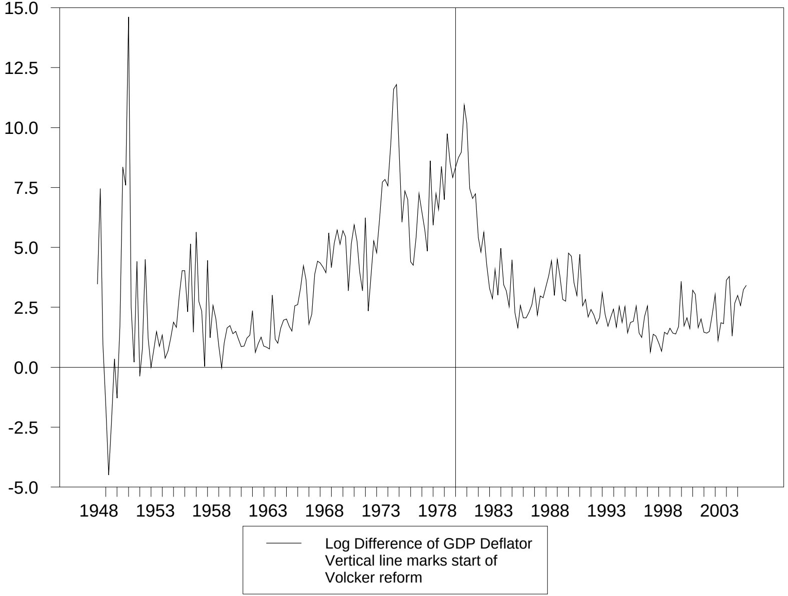

Detecting ARCH

One may also plot the squared residuals from the ordinary least squares estimate of the conditional mean model. Clusters of large values indicate the presence of ARCH.

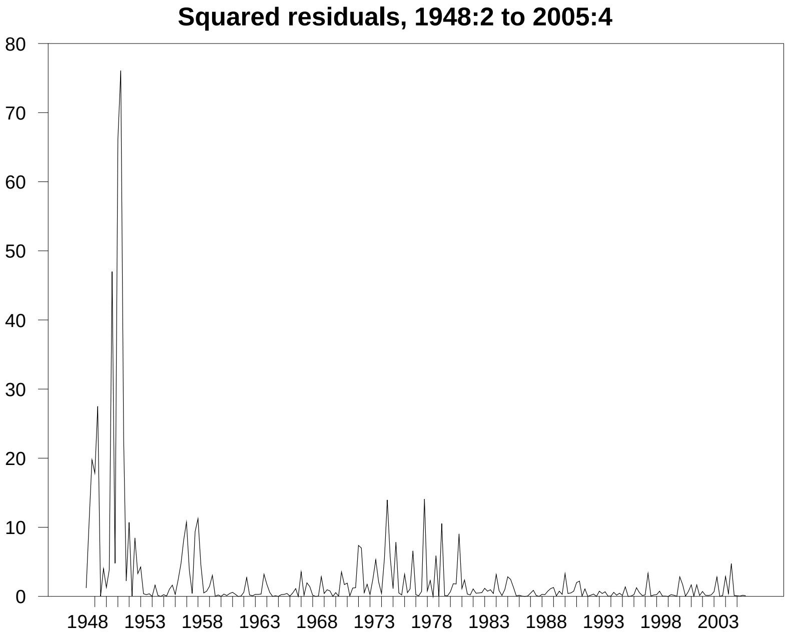
Squared residuals, 1948:2 to 2005:4
80
70
60
50
40
30
20
10
0
1948
1953
1958
1963
1968
1973
1978
1983
1988
1993
1998
2003

Analytic test for ARCH

- The OLS regression of the squared residuals on a constant and q lags provides a Lagrange Multiplier test statistic for an ARCH(q) effect.
- Under the null hypothesis of no ARCH effects, the number of observations times the coefficient of determination, TR^2, has a limiting χ^2 distribution with $q+1$ degrees of freedom.
- If the test statistic is larger than the critical value, there is evidence of ARCH effects.
- Generally, the q to be used in the model is chosen according to the highest q for which the test statistic is larger than the critical value at the chosen level of significance.

What was wrong with ARCH

ARCH models often require relatively long lags in the conditional variance equations. In early work with ARCH models, researchers often imposed arbitrarily weighted and fixed lag structures on the conditional variance equation to avoid negative unconditional variance estimates (see, for example, Engle and Kraft (1983)).

Four years after the introduction of ARCH, Engle's graduate student Tim Bollerslev addressed this issue with the generalized autoregressive conditional heteroscedasticity (GARCH) model.

The GARCH(p,q) model

In GARCH(p,q) models, the conditional variance equation is extended to include p lagged values of the conditional variance.

$$y_t = \boldsymbol{x}_t' \boldsymbol{b} + \epsilon_t \tag{7}$$

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{q} \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2 \tag{8}$$

We again assume weak stationarity in the error process.

$$E[\epsilon_t] = 0 \tag{9}$$

$$Var[\epsilon_t] = \sigma_t^2 \tag{10}$$

$$Cov[\epsilon_t, \epsilon_s] = 0, t \neq s \tag{11}$$

It can be shown that the unconditional variance is

$$Var[\epsilon_t] = \sigma_t^2 = \frac{\alpha_0}{1 - \sum_{i=1}^{q} \alpha_i - \sum_{j=1}^{p} \beta_j} \tag{12}$$

The Advantage of GARCH(p,q) models

This allows the entire history of past shocks to influence the current value of the conditional variance. Bollerslev showed a GARCH model with a small number of terms may be more efficient than an ARCH model with many terms.

ARCH and GARCH estimation

- Maximum likelihood estimators of ARCH(q) and GARCH(p,q) models are more efficient than ordinary least squares. Because OLS estimates are consistent, they may be used as starting values for maximum likelihood estimation.
- Maximum likelihood estimation of ARCH(q) models may be performed by maximization of the log likelihood function or using a four-step procedure based on the method of scoring shown in Engle (1982). See also Dr. Buck's notes.

ARCH and GARCH estimation

- Maximum likelihood estimation of GARCH models is complicated by the presence of lagged values of the conditional variance term in the conditional variance equation. (Try setting it up sometime in MathCAD, and you'll see what I mean). Numerical methods (discussed later) are usually employed to find the parameter estimates.
- I was most successful in using RATS to estimate my GARCH models. Estima (the company that writes RATS) has some useful sample programs on its Web site.

Detecting GARCH

The same methods discussed above for detecting ARCH may be used to look for evidence of GARCH effects. If the Lagrange Multiplier test statistic gives evidence of ARCH for q of four or more, a GARCH model is probably more appropriate.

Verifying ARCH and GARCH

The Ljung-Box Q-statistic may be used to verify the appropriateness of the model's specification. After estimation of an an ARCH or GARCH model, the Ljung-Box Q-statistic is calculated by regressing the residuals on a constant and their lagged values. If the Q-statistic is greater than its χ^2 critical value, one should reject the null hypothesis of no autocorrelation in the residuals. This suggests the GARCH model is misspecified since the residuals are not weakly stationarity.

ARCH and GARCH's popularity

Variations on the basic ARCH and GARCH models have been developed to better study particular theories and types of data. ARCH and GARCH models have been especially useful for assessing the risk of an investment portfolio. As a portfolio's returns become more volatile, the conditional variance term produced by GARCH models provides wider forecast confidence intervals. This contrasts with the relatively stable ordinary least squares forecast ranges that give less information about portfolio risk.

The ability of GARCH models to account for volatility clustering also makes it useful for studying inflation and uncertainty. The next three slides show forecast confidence intervals for U.S. inflation. The OLS forecast intervals stay constant throughout the time series. The GARCH forecast intervals widen in times of high volatility and narrow in times of low volatility. This may be a good way to model uncertainty about future inflation.

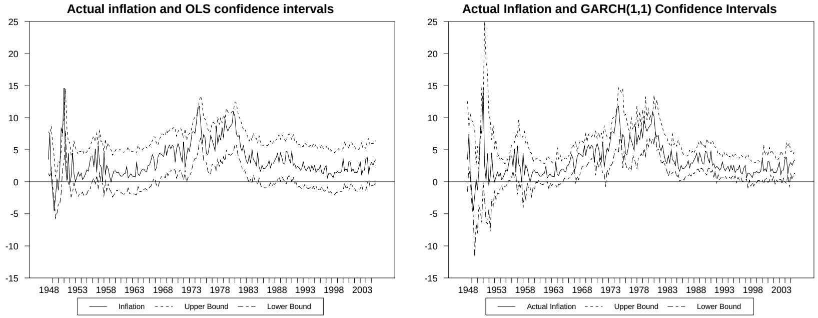
Actual inflation and OLS confidence intervals
25
20
15
10
5
0
-5
-10
-15
1948 1953 1958 1963 1968 1973 1978 1983 1988 1993 1998 2003
Inflation
Upper Bound
Lower Bound
Actual Inflation and GARCH(1,1) Confidence Intervals
25
20
15
10
5
0
-5
-10
-15
1948 1953 1958 1963 1968 1973 1978 1983 1988 1993 1998 2003
Actual Inflation
Upper Bound
Lower Bound

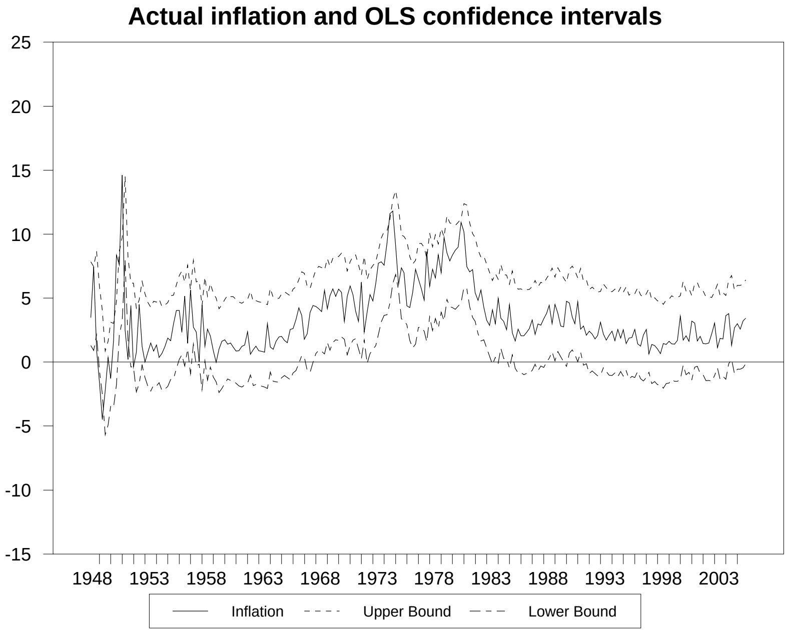

Actual inflation and OLS confidence intervals
25
20
15
10
5
0
-5
-10
-15
1948
1953
1958
1963
1968
1973
1978
1983
1988
1993
1998
2003
Inflation
Upper Bound
Lower Bound

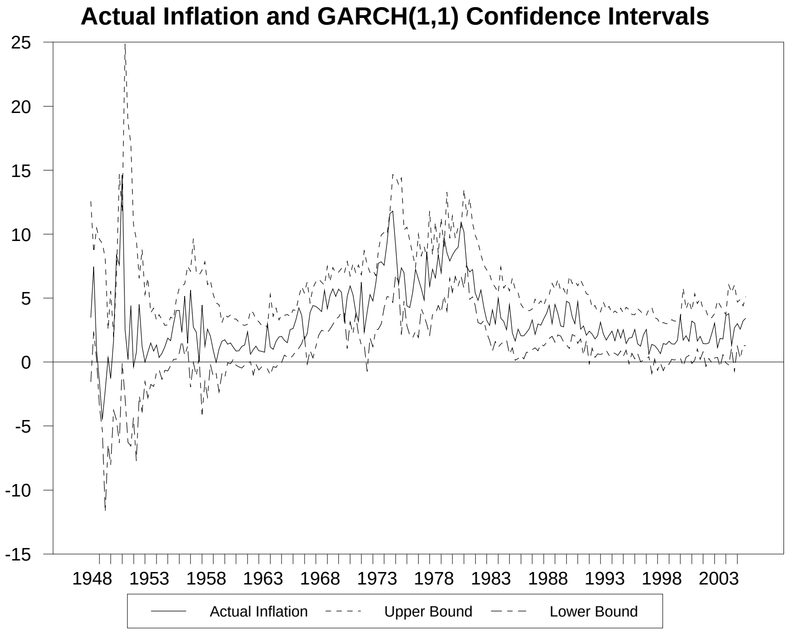
Actual Inflation and GARCH(1,1) Confidence Intervals
25
20
15
10
5
0
-5
-10
-15
1948
1953
1958
1963
1968
1973
1978
1983
1988
1993
1998
2003
Actual Inflation
Upper Bound
Lower Bound

Recommended reading

- Greene (2003), Econometric Analysis, Fifth Edition, pp. 238-246.
- Engle (1982), Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation, *Econometrica*, 50, 987-1008.
- Dr. Buck's lecture notes.
- Engle (2001), GARCH 101: The use of ARCH/GARCH Models in applied econometrics, *Journal of Economic Perspectives*
- Bollerslev (1986), Generalized autoregressive conditional heteroscedasticity, *Journal of Econometrics*, 31, 307-327. **Not freely available online through Temple**.

The October 1979 Reform of U.S. Monetary Policy
and
Inflation Targeting

Cape May, N.J., native Paul A. Volcker

Paul Volcker took office as Fed chairman in August 1979 with a reputation as an inflation fighter. At the time, inflation was on the minds of many Americans. By historical standards, it was at a relatively high level.

Using the Phillips Curve, the Fed's policymakers responded to economic slowdowns in the 1960s and 1970s by lowering targets for interest rates. In the view of many economists, the public came to believe the Fed would adopt such inflationary policies in response to future economic downturns. As a result, firms and workers increasingly contributed to inflation by raising prices and demanding higher wages in anticipation of future inflation.

Changes in monetary policy

Volcker called emergency meetings of the open market committee and Board of Governors on October 9, 1979. The governors voted unanimously to raise target interest rates one percent. They also changed their tactics for influencing the money supply. Public statements following the meeting emphasized the Fed's determination to lower the inflation rate.

In the years following the October 1979 policy change, estimates of U.S. inflation dropped from a range between 5 and 12 percent to a range between 1 and 4 percent, where they have stayed for the last 15 years.

Inflation targeting

Many other countries have tried a different tactic to control inflation. Led by New Zealand in 1990, these countries' central banks adopted publicly announced inflation rate targets.

The fundamental argument in support of inflation targeting is that targets will cause the public to believe the government will try to keep inflation low. The public will then raise prices and wages at relatively low rates, contributing to overall price stability. Some proponents of inflation targeting say the practice improves economic performance by reducing uncertainty about inflation.

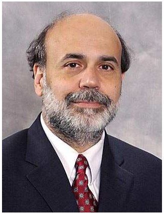

Current Fed chairman Ben S. Bernanke (who used to live in Princeton, N.J.,) supports inflation targeting.

Cons of inflation targeting

Depending on how they are implemented, inflation targets may limit central banks' ability to respond to economic crises. This is the main reason they are controversial. Former Fed chairman Alan Greenspan opposes inflation targeting. (So does Paul Volcker.)

Do we need inflation targeting when we have Paul Volcker?

Kontonikas (2004) uses several variations on the basic GARCH model to study the impact of inflation targeting on inflation and inflation uncertainty in the United Kingdom. He concludes that after the announcement of inflation targeting, U.K. inflation became substantially less persistent and less variable. He also finds a significant negative impact from inflation targeting on long-run uncertainty as measured by the GARCH conditional variance.

One of the goals of this paper is to use similar methods to study the effect on the U.S. of the 1979 reforms and the ensuing years of anti-inflationary policy. Did Volcker's reforms and Greenspan's policies have an effect similar to that Kontonikas found from inflation targeting? And is there a negative relationship in the U.S. between inflation variance and the level of inflation?

Data

An implicit GDP deflator was calculated from quarterly Bureau of Economic Analysis data (1947:1 to 2005:4, 236 observations) on real and nominal GDP.

$$Deflator = \frac{NominalGDP}{RealGDP} * 100 \tag{13}$$

Quarterly inflation rates (235 observations) were calculated by taking the log differences of successive quarterly deflators.

A dummy variable was also created, taking the value zero from 1947:1 to 1979:4 and one from 1980:1 to 2005:4. The variable's values correspond to the periods before and after the October 1979 reform.

My GARCH(1,1)-M model, with a dummy variable

I used a simple four-period autoregressive model also used in Engle and Kraft (1983) and Bollerslev (1986) with one variation. I include the conditional variance in the conditional mean model. This makes the model a GARCH-in-Mean model. The conditional variance proxies for inflation uncertainty.

$$\pi_t = b_0 + b_1\pi_{t-1} + b_2\pi_{t-2} + b_3\pi_{t-3} + b_4\pi_{t-4} + \delta\sigma_t^2 + \epsilon_t \qquad (14)$$

To account for the possible effects on inflation uncertainty of the October 1979 reforms, a dummy variable was added to the conditional variance model:

$$\sigma_t^2 = \alpha_0 + \alpha_1\epsilon_{t-1}^2 + \beta_1\sigma_{t-1}^2 + \gamma_1 D_{Reform} \qquad (15)$$

Lagrange Multiplier test

The Lagrange Multiplier test statistic, TR^2, was calculated from the ordinary least squares estimation of the squared residuals regressed on their own q-period lags.

q	1	2	3	4	5	6	7	8
TR^2	81.23	83.24	83.57	83.62	80.38	84.64	86.21	88.17
$\chi^2(5\%)$	5.991	7.81	9.49	11.07	12.59	14.07	15.51	16.92

Maximum likelihood estimation

The conditional mean and conditional variance parameters were obtained through maximum likelihood estimation. The log-likelihood function for the GARCH(1,1)-M model is

$$lnL = -\frac{1}{2}\sum_{t=1}^{T} ln(\sigma_t^2) - \frac{1}{2}\sum_{t=1}^{T} ln(\frac{\epsilon_t^2}{\sigma_t^2}) \qquad (16)$$

where T is the total number of observations. ϵ_t^2 and σ_t^2 are defined in equations (13) and (14).

The BHHH algorithm

The presence of a recursive term in the conditional variance equation complicates the optimization of the log likelihood function. Following Bollerslev (1986), the optimization was performed with the Berndt, Hall, Hall and Hausman (BHHH) iterative algorithm. Let l denote the likelihood function and $\theta^{(i)}$ denote the parameter estimates after the ith iteration. The BHHH algorithm calculates the estimators according to

$$\theta^{(i+1)} = \theta^{(i)} + \lambda_i (\sum_{i=1}^{T} \frac{\partial l_t}{\partial \theta} \frac{\partial l_t}{\partial \theta'})^{-1} \sum_{i=1}^{T} \frac{\partial l_t}{\partial \theta} \tag{17}$$

where $\frac{\partial l_t}{\partial \theta}$ is evaluated at $\theta^{(i)}$ and λ_i is a variable step length chosen to maximize the likelihood function in the given direction.

Results

The effect of conditional variance, or uncertainty, represented by δ, on the level of inflation is negative but not significant.

Table: Conditional Mean Parameter Estimates

b_0	b_1	b_2	b_3	b_4	δ
0.518^{*}	0.468^{**}	0.183^{*}	0.277^{**}	-0.001	-0.065
(1.98)	(6.89)	(2.34)	(4.37)	(-1.03)	(-0.93)

$LL = -165.92$, $Q(1) = 5.688^{**}$, $Q(6) = 9.247^{**}$, $Q(12) = 30.552$

*, ** denotes significant at 5% level, 1% level

Results

The effect of the reform dummy variable, represented by γ_1, is negative but not significant. The Q statistics indicate the model eliminates autocorrelation of residuals up to six periods in the past.

Table: Conditional Variance Parameter Estimates

α_0	α_1	β_1	γ_1
0.564^{**}	0.467^{**}	0.350^{*}	-0.318
(2.60)	(2.96)	(2.28)	(-1.79)

*, ** denotes significant at 5% level, 1% level

Results

The federal government imposed price controls during the Korean War. The price controls likely brought a sharp reduction in inflation that may have been expected. The government also imposed price controls during parts of the 1960s and early 1970s. The price controls may cause the model to overstate the conditional variance of inflation during those times. This in turn may cause the reduction in inflation volatility from the Volcker reforms to be understated.

An alternative model

To correct for this in an early ARCH study of U.S. inflation, Engle (1983) added dummy variables to the conditional variance model during the approximate periods of price controls: 1951:2 to 1953:2, 1962:1 to 1968:4 and 1971:2 to 1973:2. Dummy variables were added to the GARCH(1,1)-M model's conditional variance equations for the same time periods to give

$$\sigma_t^2 = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 + \gamma_1 D_{Reform} + \gamma_2 D_{PC1} + \gamma_3 D_{PC2} + \gamma_4 D_{PC3} \tag{18}$$

where D_{PC1}, D_{PC2} and D_{PC3} represent respectively the dummy variables for the first, second and third price control eras.

Results

In the new model, the conditional variance demonstrated a statistically significant negative relationship with the inflation rate. But the Q-statistics lead to the rejection of the null hypothesis of no serial correlation in the residuals. This suggests the model is misspecified.

Table: Conditional Mean Parameter Estimates with Post-Reform and Price Control Dummy Variables

b_0	b_1	b_2	b_3	b_4	δ
0.698^{**}	0.461^{**}	0.164^{*}	0.301^{**}	-0.002	-0.148^{**}
(3.05)	(6.38)	(2.28)	(6.02)	(-1.69)	(-3.41)

$LL = -162.74$, $Q(1) = 20.717$, $Q(6) = 37.358$, $Q(12) = 66.42$

*, ** denotes significant at 5% level, 1% level

The post-1979 reform dummy variable, represented by γ_1, is negative and nearly significant at the 5 % level. It is significant at the 5.61 % level.

Table: Conditional Variance Parameter Estimates with Post-Reform and Price Control Dummy Variables

α_0	α_1	β_1	γ_1	γ_2	γ_3	γ_4
0.805^*	0.440^{**}	0.337^*	-0.549	-0.448	-0.444	-0.108
(2.41)	(3.13)	(2.19)	(-1.91)	(-0.14)	(-1.51)	(-0.07)

*, ** denotes significant at 5% level, 1% level

But the Q-statistics lead to the rejection of the null hypothesis of no serial correlation in the residuals. This suggests the model is misspecified.

Table: Conditional Mean Parameter Estimates with Post-Reform and Price Control Dummy Variables

b_0	b_1	b_2	b_3	b_4	δ
0.698^{**}	0.461^{**}	0.164^{*}	0.301^{**}	-0.002	-0.148^{**}
(3.05)	(6.38)	(2.28)	(6.02)	(-1.69)	(-3.41)

$LL = -162.74$, $Q(1) = 20.717$, $Q(4) = 36.105$, $Q(12) = 66.42$

*, ** denotes significant at 5% level, 1% level

Conclusions

- The first GARCH(1,1)-M model did not produce evidence that lower inflation volatility is related to a lower inflation rate. Nor did it produce evidence that inflation was significantly less volatile after the reforms.
- After controlling for possible distortions from price controls, the results changed. Inflation volatility was found to have a significant negative effect on inflation. The era after the 1979 reform was found to have lower inflation volatility at a nearly conventional level of statistical significance. The Ljung-Box statistic, however, suggests the second model is misspecified.

Results

- Combined with Kontonikas's results, this study suggests inflation targeting is more effective at lowering inflation and inflation uncertainty than a firmly anti-inflation discretionary policy.
- The lack of signifcant results could be due to model specification. Kontonikas relied heavily on Akaike-Schwarz information criteria to determine the lags he used in his autoregressive models. He also used the component GARCH-M model of Engle and Lee (1993) to separate inflation volatility into short-term and long-term components. He finds inflation targeting had a more significant negative effect on long-term uncertainty than short-term.

Results

- This study could by extended by performing estimates using monthly data as well as other measures of prices such as the consumer and producer price indices.
- Alternative specifications of the conditional mean may provide a better model. Including variables such as wages, import prices and oil prices could affect the results.
- It is also worth examining the possibility that the model of U.S. inflation should change over time.

A few closing thoughts

It is also worth noting how American public attitude toward inflation has changed since the 1970s. A recent newspaper column (Samuelson 2004) quotes polling expert Daniel Yankelovich, who wrote in 1979, "For the public today, inflation has the kind of dominance that no other issue has had since World War II. . . . It would be necessary to go back to the 1930's and the Great Depression to find a peacetime issue that has had the country so concerned and distraught."

Today, inflation is rarely mentioned as a serious complicating factor in business transactions or as an issue that needs the attention of the government. This suggests, regardless of the findings of these econometric models, the public has more certain expectations about inflation now than when Paul Volcker took office. On the other hand, these results could show that the American public was too concerned about inflation in the late 1970s.