A New Investigation of the Impact of Wage and Price Controls

Charles H. Whiteman

Research Assistant
Research Department
Federal Reserve Bank of Minneapolis

How well do wage and price controls work against inflation? Past studies of the last U.S. experiment with such controls used deficient methods and don't agree on the answer to this question. To reexamine the impact of the last controls, we have used a new technique which overcomes the worst defects of these studies. Our results:

- Controls held down inflation temporarily, but they also held down production.
- Once controls ended, both prices and production rapidly moved to catch up; they rose faster than they would have had controls never been imposed—and prices ended up higher than they otherwise would have.

These wage and price controls obviously didn't work very well.

Before we look at our method and results in detail, let's briefly review the last controls program and examine the methods, results, and deficiencies of previous studies.

The Last U.S. Wage and Price Controls

Without public debate or warning, the Nixon Administration imposed a wage and price control program on August 15, 1971. While the program evolved through four phases that ended April 30, 1974, only the first two phases were really effective. During the ninety days of Phase I (August 15 through November 13, 1971), wages, prices, and rents were frozen. Phase II (which lasted until January 11, 1973) was less severe but required that wage and price increases stay within strict guidelines overseen by the Cost of Living Council and administered by the Pay Board and the Price Commission. Large firms had to get approval before any increases in wages and prices, and smaller firms had to report any increases exceeding the guidelines.

Controls were generally relaxed in the last phases of the program. Wages and prices were self-regulated during Phase III, and this caused many to believe controls had ended.[1] The resulting explosion of prices during its first months, however, led to a Presidential refreeze on June 13, 1973, in preparation for the August 12 imposition of Phase IV. Though nominally stronger than Phases I and II, both the refreeze and Phase IV were actually weaker because of progressive deregulation, pressure from unregulated prices of imports, and the catch-up which had started as Phase II ended. Therefore, though controls did not officially end until April 1974, their bite dissipated rapidly after Phase II.[2]

Past Studies of the Impact of Controls

The standard way to investigate the impact of Phases I and II is to statistically simulate what would have happened had the controls not been imposed and compare that "counterfactual benchmark" to what actually happened; the difference is assumed to be due to controls. To generate the benchmark, studies so far have primarily used two methods, both of which are flawed.

---

[1] U.S., President, *Economic Report of the President* (Washington, D.C.: Government Printing Office, 1974), p. 103.

[2] For a more complete description of the control program, see U.S., President, *Economic Reports of the President* (Washington, D.C.: Government Printing Office, 1973 and 1975), p. 51 and p. 219; Daniel J. B. Mitchell, "Phase II Wage Controls," *Industrial and Labor Relations Review*, Vol. 27, No. 3 (April 1974), p. 351; Walter Guzzardi, Jr., "What We Should Have Learned About Controls," *Fortune*, Vol. 91, No. 3 (March 1975), p. 102; and Arnold R. Weber, "The 1971 Wage-Price Freeze and Incomes Policy," *Monthly Labor Review*, Vol. 95, No. 4 (April 1972), p. 18.

One benchmark generation method builds a system of "structural" equations which are intended to describe the behavior of agents in the "wage-price sector" of the economy.[3] This structural method is essentially an attempt to identify the Phillips curve,[4] and its equations typically express current rates of wage and price inflation in terms of current and past values of variables such as GNP, the government deficit, the unemployment rate, and the rate of growth of productivity. Once the form of each equation has been determined,[5] the unknown coefficients in the relationships between variables are estimated by an appropriate statistical technique using data up to the imposition of controls. Then the equations are used to generate the sequence of values for each variable during and after the control period which presumably this sector of the economy would have generated had controls never been imposed.

The other way to generate a counterfactual benchmark doesn't rely on descriptions of economic behavior; instead, it relies on purely statistical regularities. The ARIMA (AutoRegressive Integrated Moving Average) method[6] expresses each variable in terms of linear combinations of its own past (Auto-Regressive) and random "shocks" (Moving Average).[7] To get the benchmark, an ARIMA structure is estimated using data up to the control period. Then, assuming that the shocks are zero, the structure is used to generate values for the variable for the control period and thereafter.

Studies using these methods to investigate the impact of the last controls have produced varied results. Two studies using the structural method, for example, found that inflation was down significantly from what it would have been during Phases I and II and somewhat after that as well.[8] Another study using the same method, though, found that inflation was only slightly less than what it would have been during controls and slightly more after.[9] A study using the ARIMA method didn't look at what happened after controls, but it found still a third result for Phases I and II: depending on the price measure used, controls either had no effect on inflation or actually increased it.[10]

Who's right? That's pretty hard to say, particularly because both of the methods these studies used are defective.

The structural method may be the weakest. It requires researchers to describe economic behavior with equations, and that is not easy.[11] Although each of the structural studies referred to above, for example, presumes to identify the same "structure," their authors disagree considerably on the precise form of the equations and the appropriate price index. But even if economists could agree on a structure, this method has a more critical problem: although most economic indicators would likely be affected by wage and price controls, this method assumes some "exogenous variables" would not, and it relies on them to

---

[3]For a critique and summary of this type of study before 1973, see Walter Y. Oi, "On Measuring the Impact of Wage and Price Controls: A Critical Appraisal," *The Economics of Price and Wage Controls*, Vol. 2 of *Carnegie-Rochester Conference Series on Public Policy*, ed. Karl Brunner and Allan H. Meltzer (Amsterdam: North-Holland Publishing Co., 1976), pp. 7-64. Three studies which used this approach are Timothy W. McGuire, "On Estimating the Effects of Controls," *The Economics of Price and Wage Controls*, pp. 115-156; Robert J. Gordon, "The Response of Wages and Prices to the First Two Years of Controls," *Brookings Papers on Economic Activity-3* (Washington, D.C.: Brookings Institution, 1973), pp. 765-779; and James L. Pierce and Jared J. Enzler, "The Effects of External Inflationary Shocks," *Brookings Papers on Economic Activity-1* (Washington, D.C.: Brookings Institution, 1974), pp. 13-61.

[4]The Phillips curve is a graph of the observed inverse relationship between the inflation rate and the unemployment rate.

[5]An example of the type of equations in structural systems is McGuire's "price equation":

$$p(t) = b_0 + b_1[1/u(t)] + b_2[1/u(t-1)] + P_1^e(t-1) + v(t)$$

where $p(t)$ is the percentage change in prices at time t, $u(t)$ is the unemployment rate at time t, $P_1^e(t-1)$ is the (unobserved) expected percentage change in prices at time t, $v(t)$ is a random term, and b_0, b_1, and b_2 are coefficients to be estimated.

[6]We know of only one study to date which has used this method: Edgar L. Feige and Douglas K. Pearce, "The Wage-Price Control Experiment—Did It Work?" *Challenge*, Vol. 16, No. 3 (July-August 1973), pp. 40-44. An extended version of this paper, "Inflation and Incomes Policy: An Application of Time Series Models," appears in *The Economics of Price and Wage Controls*, pp. 273-302.

[7]An example of this is the Feige-Pearce consumer price index inflation equation:

$$p(t) = p(t-1) + p(t-4) - p(t-5) + \varepsilon(t) - 0.57\varepsilon(t-1) - 0.77\varepsilon(t-4) + 0.44\varepsilon(t-5)$$

where $p(t)$ is the rate of inflation at time t and $\varepsilon(t)$ is a random "noise" at time t.

[8]Gordon and Pierce-Enzler.

[9]McGuire.

[10]Feige-Pearce.

[11]Economists have become suspicious of the types of restrictions inherent in this procedure. See, for instance, C. A. Sims, "Macroeconomics and Reality," *Econometrica* (forthcoming).

generate what other variables would have been without controls. Two of the three structural studies referred to above, for instance, assume that the unemployment terms in their inflation equations were unaffected by controls, and the other makes a similar assumption. But economic theory and careful empirical work suggest that policies such as wage and price controls do change the behavior of agents in the labor market and hence the unemployment rate.[12] Exogeneity assumptions, then, make the results of these studies suspect.

The ARIMA method overcomes the exogeneity problem since benchmarks generated by it depend only on the past history of the variable being forecasted, not on any other variables. Still, because of this, the ARIMA method has its own serious problem: it omits important variables. For instance, although one would almost surely expect the level of the money supply to help forecast the price level, the ARIMA method uses only past values of the price index. Omitting variables can be very costly. In the ARIMA study referred to above, for example, the estimates of the CPI inflation rate for the sixteen months before controls miss the actual values by an average of about 70 percent per month.[13]

Our New Method and Results

To investigate the impact of wage and price controls, we use a new method to generate a counterfactual benchmark: Vector-AutoRegression (VAR).[14] It shares some features of the structural and ARIMA methods but overcomes their biggest problems.

Like the ARIMA method, VAR tries to capture statistical, not behavioral regularities in data. It expresses variables in terms of past values of themselves too, but it does better than that: it expresses each variable in a list of variables (the "Vector" of "VAR") in terms of not only its own past but also the past of the others in the list (AutoRegression). Thus, important variables need not be omitted, as they are in the ARIMA method. This can really make a difference in results. Whereas the estimates by the ARIMA study mentioned earlier miss the actual inflation rates for the sixteen months before controls by about 70 percent per month,[15] our VAR estimates miss them by less than 2 percent.

Like the structural method, VAR does rely on economic theory, but only to choose the variables to include in the system. It solves the structural exogeneity problem by generating forecasts for each of the variables in the list; no variable is assumed unaffected by controls. Once the system is estimated up to the control period, the benchmark values for the next month—August 1971—are produced by forecasting all the variables using data up to that time. August's forecasted values are then used to generate September's benchmark, September's to generate October's benchmark, and so on.[16] Each VAR benchmark, therefore, is exactly what a benchmark should be: a sequence of predicted values completely unaffected by events during controls.

Our study improves on past studies in another way. Past studies have reported the impact of controls as the difference between benchmark and actual values. But benchmarks cannot be that precise: any forecast has some uncertainty around it. We explicitly take this into account by calculating and displaying a benchmark region for each variable.[17] The region represents the set of values the variable most likely would have had without controls, based on the random variation inherent in any economic variable. If the actual values of any variable fall outside its benchmark region, we can be relatively sure it's be-

---

[12]John Geweke, "Wage and Price Dynamics in U.S. Manufacturing," *New Methods in Business Cycle Research: Proceedings from a Conference* (Federal Reserve Bank of Minneapolis, Minnesota, 1977), pp. 111-156.

[13]Feige-Pearce, Table 3a, p. 285. The ARIMA was estimated using data from 1953.7 to 1970.4. Feige and Pearce do note that this is less than one standard error. Our experimentation with ARIMA specifications got two results. First, for a revised CPI (1967=100), there exists a model specification (numbers of autoregressive and moving average parameters) which fits the data better than that used by Feige-Pearce. Second, when the Feige-Pearce specification is fit to the revised CPI, the estimated parameters are not qualitatively different, and the benchmark rate of inflation becomes *negative* eighteen months after the imposition of Phase I and decreases without bound. This result is not inconsistent with the values generated and reported by Feige-Pearce for the sixteen-month period beginning in August 1971.

[14]An example of a VAR is

$$Z(t) = A_1Z(t-1) + A_2Z(t-2) + \ldots + A_kZ(t-k) + \varepsilon(t)$$

where $Z(t)$ is a vector of variables at time t, $\varepsilon(t)$ is a vector of random terms at time t, and $A_1, \ldots, A_k$ are matrices of coefficients.

[15]Feige-Pearce.

[16]This procedure is an application of the chain rule of forecasting.

[17]Because the benchmark region for the real wage (wages deflated by the price level) involves the distribution of a very complicated random variable, it is not presented.

cause of controls, not chance.[18]

We chose variables for our VAR vector based on a theory of the interaction of the markets for labor and aggregate output. In the labor market, supply and demand determine the real wage (the money wage divided by the price level) and employment. These in turn determine aggregate output, and aggregate demand then determines the absolute price level and the money wage.[19] Hence, our vector includes

these variables:	as measures of:
the industrial production index (1967=100)	output
weekly hours worked in manufacturing industries	employment
straight-time average hourly earnings in manufacturing industries	the money wage
M1 – the money supply (currency + demand deposits)	aggregate demand

We used these in VAR systems for three measures of the price level:[20]

the consumer price index for all items (CPI, 1967=100)
the consumer price index for all items except food (1967=100)
the wholesale price index (WPI, 1967=100)

Figures 1-3 show our results for the price measures, and they're all alike. Generally, controls curbed inflation temporarily: during Phases I and II,

[18]The benchmark regions displayed in the figures are two standard error bounds for the forecasts under the assumption that the estimation procedure determined the "true" matrices $A_1, \ldots, A_k$ of note 14.

[19]This is an equilibrium or classical explanation of the interaction of these two markets. The disequilibrium or Keynesian interpretation would not change the list of variables used in the VAR.

[20]The VAR which generated the results below was chosen for its ability to predict the period 1971.1-1971.6. It had nine lags (the "k" of note 14) and was estimated over the period 1964.1-1971.7. The system had no statistical priors on the rapidity with which the lag coefficients died out. The estimation used the VAR computer program, PREDICT, developed by Robert Litterman, Research Assistant, Research Department, Federal Reserve Bank of Minneapolis.

Figure 1

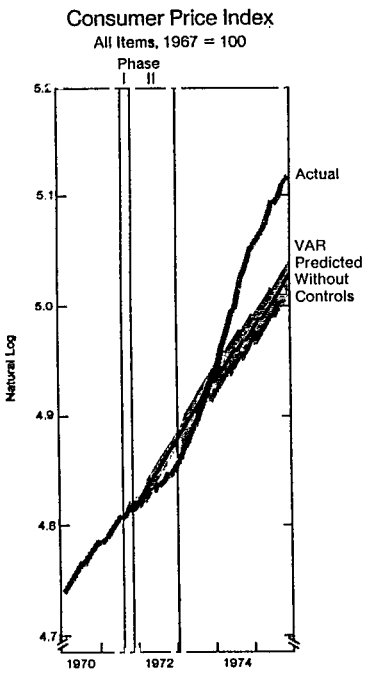

Source of basic data U.S. Department of Labor

Figure 2

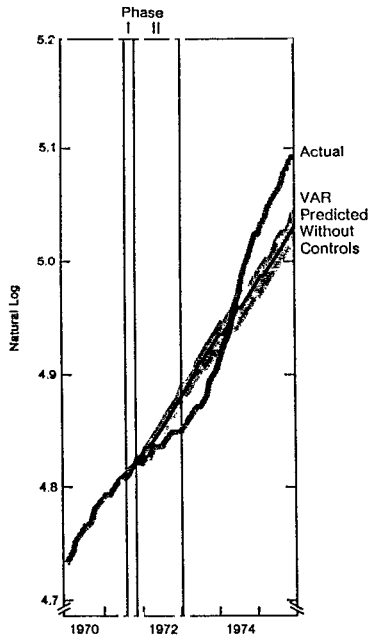

Figure 3

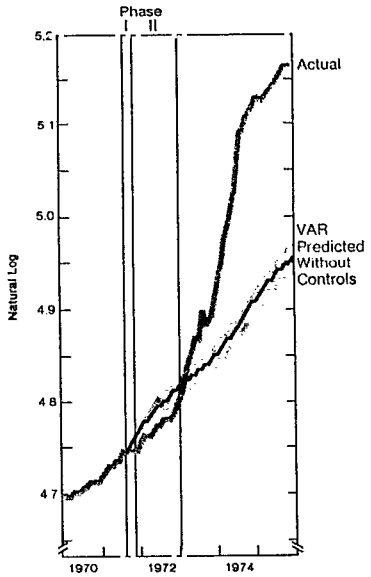

controls held prices somewhat below the benchmark region, that is, below what they would have been without controls. After Phase II, though, prices rose faster than the benchmark region and ended up far above it—higher than they would have been without controls.

A look at price inflation rates verifies the temporary impact of controls. Figure 4 shows the differences between the actual rate of increase in the CPI and what our VAR says it would have been without controls. Again we can see that during Phases I and II, the inflation rate was lower than it would have been. But this clearly did not last long. After Phase II, inflation shot up to a rate much higher than it would have been otherwise.

Controls apparently had little or no effect on wage inflation, though. As Figure 5 shows, wages were held slightly below what they would have been during the few months of Phase I. Through Phase II and for nearly two years thereafter, however, actual and benchmark values for average hourly earnings essentially coincide.

Despite the fact that their wages weren't affected much, workers were first made better off, then worse off by what controls did to prices. This is evident in Figure 6, which compares the actual and benchmark values of the real wage (again, wages deflated by the price level). During Phases I and II, with prices but not wages held down, the real wage rose—eventually to a level higher than it otherwise would have. After Phase II, prices rose much faster than expected, but wages didn't: the real wage dropped sharply and ended up lower than it would have been without controls.

Macroeconomic theory suggests that firms will respond to increases (decreases) in the real wage, other things held constant, by decreasing (increasing) employment and hence output. That is exactly what the VAR detected for the control period. As can be seen in Figures 7 and 8, during Phase I, when the real wage increased, the hours worked in manufacturing industries dropped and industrial production fell

Figure 4

The Impact of Controls on Inflation

Actual CPI Inflation Rate — VAR Predicted CPI Inflation Rate

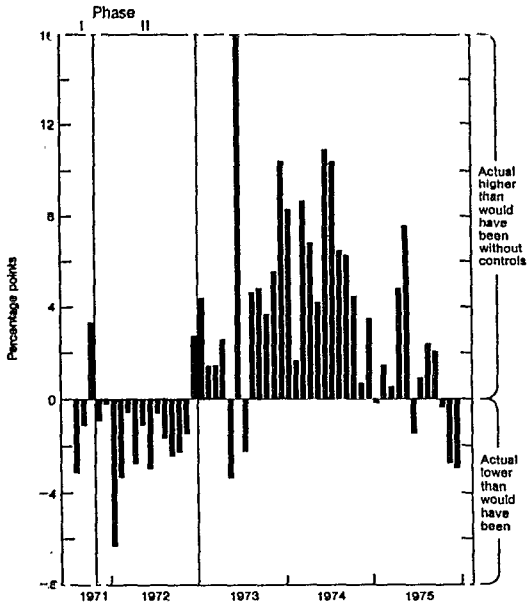

Sources of basic data: U.S. Department of Labor, Board of Governors of the Federal Reserve System

Figure 5

Average Hourly Earnings in Manufacturing Industries

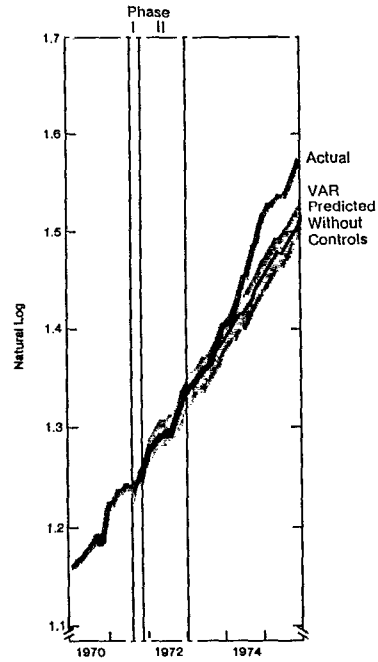

Figure 6

The Real Wage

Average Hourly Earnings in Manufacturing Industries ÷ Consumer Price Index

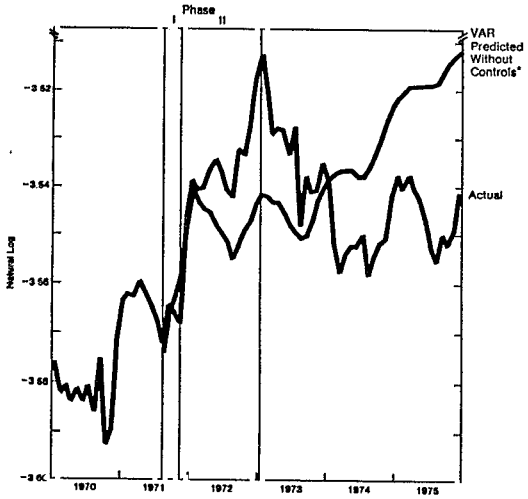

*Because the benchmark region for the real wage involves the distribution of a very complicated random variable, it is not displayed here.

below its benchmark. After controls were lifted, when prices exploded and the real wage plummetted, both employment and output more than recovered. In the long run they both settled back to where they would have been without controls.

Monetary policy seemed to be reacting as usual to the changes in output and prices. As Figure 9 shows, when output and inflation were depressed early in the control period, the money supply was not expanded more rapidly than would have been expected; if anything, it increased less rapidly. Then, when output and inflation started their catch-up late in Phase II, the actual M1 values started to rise faster than the benchmark region.

Summary, Hedges, and Conclusion

In summary, then, our last wage and price controls didn't give us much relief from inflation, and they did affect production. Prices were below what they would have been during Phases I and II, but except for a brief initial period, wages were not. Thus, the

Figure 7

Weekly Hours

in Manufacturing Industries

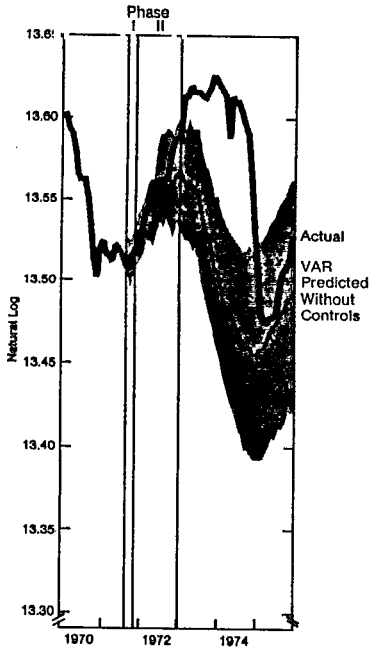

Figure 8

Industrial Production Index

1967=100

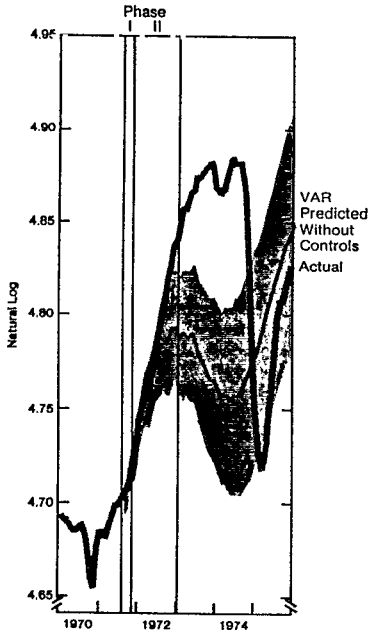

Figure 9

Money Supply

M1 = Currency + Demand Deposits

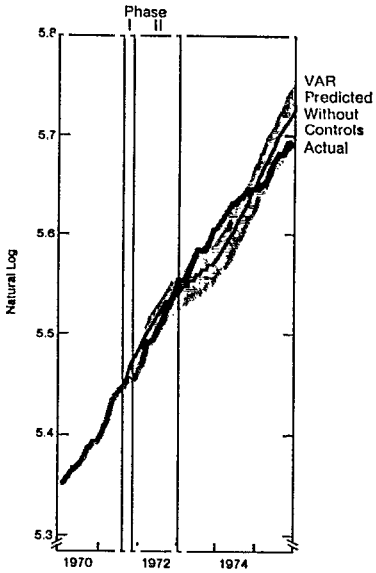

real wage increased more than it would have, and employment and output dropped. As Phase II ended, though, prices began to catch up. Because controls had little effect on wages, this catch-up lowered the real wage, and firms responded by hiring more workers and boosting production. Evidently, controls caused no long-run output distortions, since industrial production eventually returned to where it would have been without controls. Probably because of an accommodating monetary policy, however, prices ended up higher than they would have.

As with any empirical investigation of economic phenomena, our results should be interpreted carefully. All the differences between actual values and benchmark regions may not be due to the controls program; there may have been shocks to the price level we could not have accounted for. Also, because we considered only measured, published prices, we had to ignore true transaction prices—prices which during the control program would have reflected increased waiting times due to product shortages, attempts at evasion through disguised price increases, and the resources used up by government in enforcement programs designed to punish such evasion.

Still, the new method we used makes this study of the impact of the last wage and price controls a considerable improvement over past studies. So policy makers should take note: wage and price controls are not likely to provide lasting relief from inflation.

The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.