Comments on "A new measure of monetary shocks: Derivation and implications"

By Christina Romer and David Romer

John H. Cochrane

July 17 2004, presented at NBER EFG meeting

I want to do something unusual at NBER discussions - I want to argue that the paper is much better and deeper than the authors think it is!

1. Two big innovations:

Historical evidence is crucial to both innovations.

1. "Intentions." This extends target data back in time. The target is surely the right variable to look at. This gives us more data, with less measurement error.
2. Greenbook forecasts are *sufficient statistics.*

$$\text{Greenbook forecast} = E(y_{t+k}|\text{Fed Info at } t)$$

Narrative evidence is important for this conclusion. I thought the Greenbook was just a staff-amusement exercise, but the evidence is strong for it.

This property results in a "maximally informative" shock, as I'll show later.

This is the key innovation, and gives us simpler, clearer, and more robust procedures.

In my discussion I'll show this and I'll show how the procedures can be simpler and clearer still.

2. From target to policy shocks.

Our first step is to take the target or intentions series and create policy shocks. Christina and David use the following regression

$$\begin{aligned}\text{RR } \Delta ff_m &= \alpha + \beta ffb_m + \sum_{i=-1}^{2} \gamma_i \Delta \tilde{y}_{mi} + \sum_{i=-1}^{2} \lambda_i \left(\Delta \tilde{y}_{mi} - \Delta \tilde{y}_{m-1,i} \right) \\ &\quad + \sum_{i=-1}^{2} \phi_i \tilde{\pi}_{mi} + \sum_{i=-1}^{2} \left(\pi_{mi} - \pi_{m-1,i} \right) + \rho \tilde{u}_{m0} + \varepsilon_t\end{aligned}$$

I'll simplify this a bit, with the following proposition.

Proposition 1*: To measure the effects of monetary policy on* **output** *it is enough that the shock is orthogonal to* **output** *forecasts. The shock does not have to be orthogonal to price, exchange rate, or other forecasts. It may be predictable from time t information; it does not have to be a shock to agent's or the Fed's entire information set.*

All the shock has to do is remove the reverse causality from output forecasts. Example: If the Fed "responded" to weather forecasts, that movement would be a fine shock for output measurement.

A more realistic example: If the Fed moves to offset the exchange rate, this counts as a "shock" to measure the *output* response as well. *Conditioned* on y forecasts, a movement to offset unemployment counts too!

This also helps to answer the "what is a shock" question. The Fed never rolls dice; every move is a response to something. The paper has a nice discussion of many changes which can count as shocks. This is perhaps the most important source of shocks– movements in response to something that has no correlation with output expectations.

$$\begin{aligned} \text{JC} \quad \Delta ff_m &= \alpha + \sum_{i=0}^{2} \gamma_i \Delta \tilde{y}_{mi} \quad (+\beta ff_{m-1} + \delta \Delta ff_{m-1}) + \varepsilon_t^y \\ \Delta ff_m &= \alpha + \sum_{i=0}^{2} \gamma_i \tilde{\pi}_{mi} \quad (+\beta ff_{m-1} + \delta \Delta ff_{m-1}) + \varepsilon_t^\pi \end{aligned}$$

(It's not necessary to add ff, but this makes the shocks serially uncorrelated, which is useful in the next step.)

Again, it's the *sufficient statistic* property of output forecasts that lets us do this. We usually have to include all variables (as in a VAR) because any variable *might* enter the Fed's forecast. With the sufficient statistic in hand, we can leave all the other stuff out.

Nothing wrong with RR, or VAR procedures; in population they all give the same answer. The advantage of simpler procedure in small samples: 1) It's good to keep as many shocks as you can. Low R^2 is *good*. This is the sense in which the sufficient statistic gives us a "maximally informative" shock. 2) It's good to estimate as few parameters as you can. RR are much better here than an 8 variable, 24 lag VAR. The JC regression is a bit more in the same direction.

This is exactly the goal stated in the paper.

> p. 12 "The goal of this regression is not to estimate the Federal Reserve's reaction function as well as possible. What we are trying to do is to purge the intended funds rate series of movements taken in response to useful information about future economic developments. Once we have accomplished this, it is desirable to leave in as much of the remaining variation as possible."

I agree entirely; the JC regression just takes one remaining step.

We can compare the shocks in a graph,

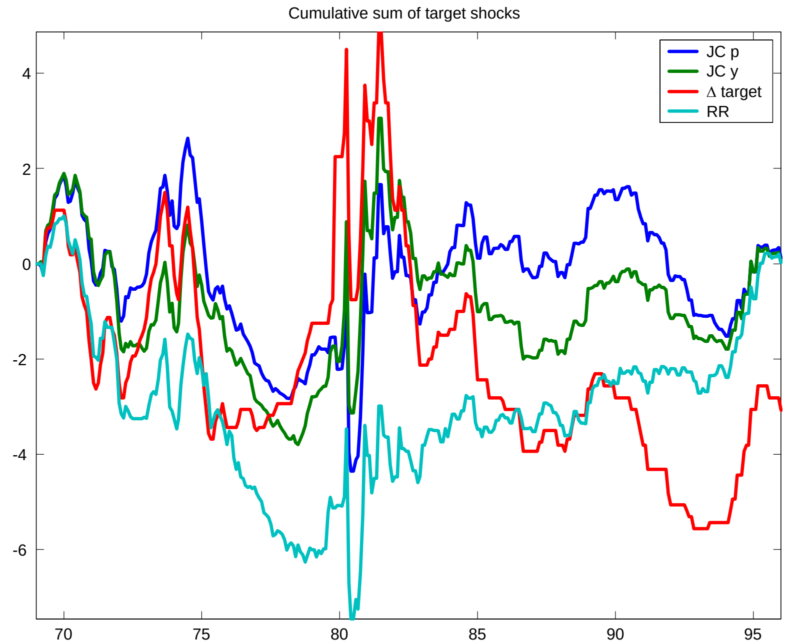

The results are about the same. The plot is cumulative sums of shocks so you can see the lines better, and the shocks are differences in the plot. .

There are some differences – you can see the unemployment and inflation levels in the differences between the lines. I allow response to u and inflation, conditioned on the y forecast, to be a shock for y. This accounts for some difference in results.

3: Measuring output response.

The second step is measuring the response of output to the monetary policy shocks. Christina and David run the following regression.

$$\text{RR} : \Delta y_t = a_0 + \sum_{k=1}^{11} a_k D_{kt} + \sum_{i=1}^{24} b_i \Delta y_{t-i} + \sum_{j=1}^{36} c_j \varepsilon_{t-j} + e_t,$$

They calculate the response by dynamic simulation; they cumulate the responses to find the response of the level of output, and they calculate standard errors by Monte Carlo.

I'll simplify again, and show the essence of the procedure with

Proposition 2 : *One can measure the output (price) response by running* **single** *regressions of output (price) on the shocks, with no other variables.*

Why? In the end we're looking for the MA representation for output.

$$\Delta y_t = \sum_j a_j \varepsilon_{t-j} + \sum_j b_j v_{t-j} = a_k \varepsilon_{t-k} + w_t$$

All right hand variables are uncorrelated, so single = multiple regressions (in population.)

Thus, I take out seasonals,$\tilde{y}$, and then run *single* regressions

$$\begin{aligned} \text{JC} \quad \tilde{y}_{t+1} - \tilde{y}_t &= a_0 + c_1 \varepsilon_t^y + e_{t+1}^1 \\ \tilde{y}_{t+2} - \tilde{y}_t &= a_0 + c_2 \varepsilon_t^y + e_{t+1}^2 \\ \tilde{y}_{t+3} - \tilde{y}_t &= a_0 + c_3 \varepsilon_t^y + e_{t+1}^3 \\ &\ldots \end{aligned}$$

$c_1, c_2, c_3 ... =$ the response of output levels directly. I calculate standard errors with usual OLS standard errors. (You need a Hansen-Hodrick or Newey-West correction for serial correlation. I use the Newey-West correction with 1.5 times as many lags as the overlap.)

Again, the procedures give no difference in population, but some advantages in samples.

1) Single regressions are more stable in finite samples. You avoid a 36×36 matrix inversion.

2) Most importantly in my mind, with the single regressions, the estimate is tied directly to the fact you want to measure: "What is the average change in output following a one percent change in intentions, orthogonal to output forecasts?"

4. The size and sign of output and price responses

Here are my results.

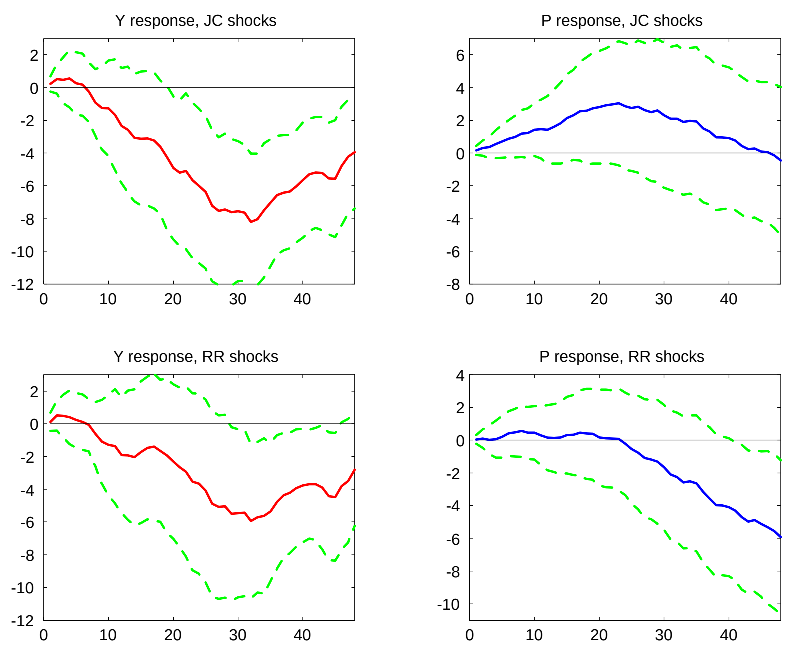

I include 2 standard error bands, so they look a bit larger than Christina and David's, but they aren't.

The top row gives my OLS regressions of output and inflation on my shocks. the bottom row gives my regressions on Christina and David's shocks. See the paper for the RR shock and method.

For output we have just about the same results. A 100 bp tightening gives a 6-8% decline at about 30 months. It's significant with a t a bit over 2.

Prices give some differences. My shocks give a bigger price puzzle. I think this is because I do not clean unemployment and output forecasts from the right side of the shock regression. Using the RR shocks I get about the same response – flat for two years and then a steady decline. However, my standard errors are much larger then theirs, and I have not tracked down why this happens.

These responses are to my eye not that much different from the VAR literature. That's fine and good. It's good to confirm what we thought we knew with much better procedures. It would be troubling if we got a revolutionary difference; say strong and significant rises in output or prices!

This is the best evidence we have, and alas it is weak. t statistics just above 2 are seductive. An advantage of the OLS approach is that we can see the regressions through scatterplots. I show here the scatterplot of the regression that measures the 3 month response of output and inflation to shocks – the regression of 30 month output or price growth on the shock.

My first scatterplot uses all the data. As you can see there are some large outliers which seem to drive the results. I looked at these by hand – they all come from the early 1980s and correspond to target movements of greater than 1 percentage point per month. We want to check that our conclusions don't come from just these datapoints, so I eliminated these 7 outliers in what follows.

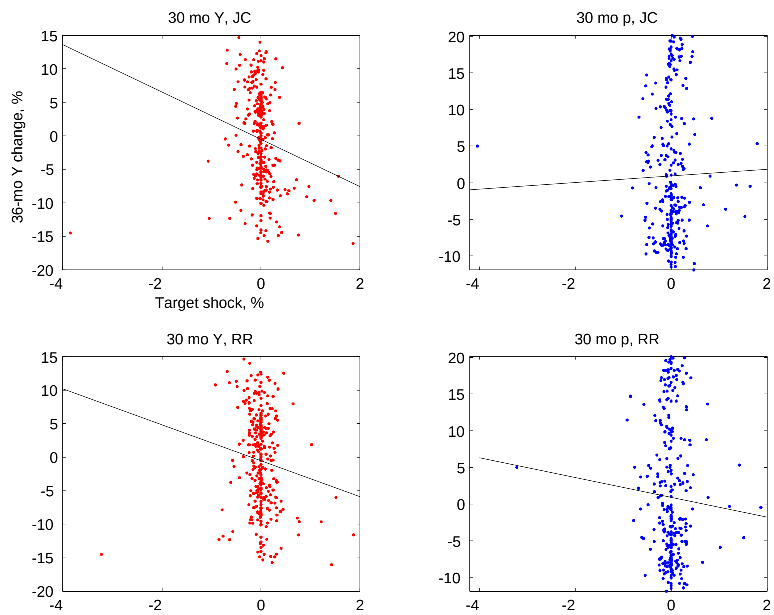

Eliminating the outliers actually made the slopes larger and more significant, and it is these results that I plotted above as impulse responses.

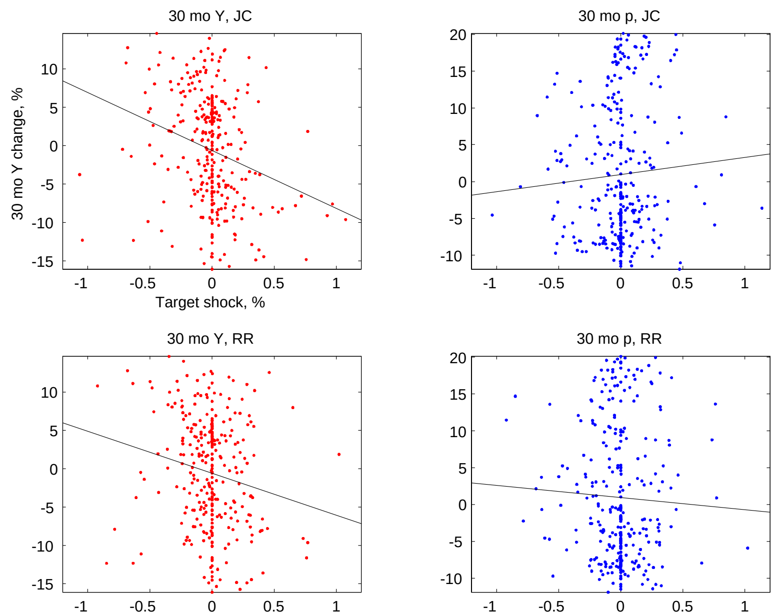

The scatterplots may seem depressing, but that's what t just above 2 in 30 years of data looks like!

Here is a plot with dates, so we can see what historical episodes drive the results – which output surges were preceded by a series of shocks.

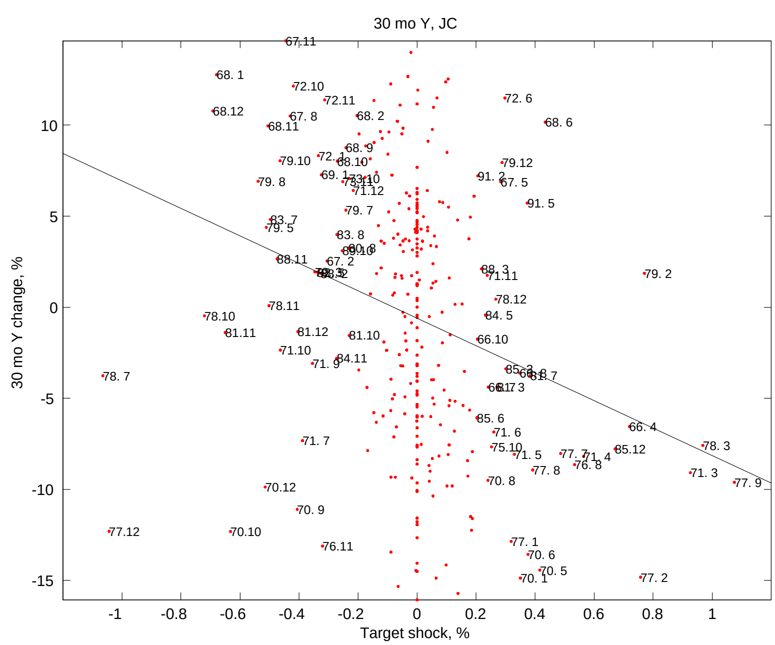

I think the plots make the point that even the output evidence is extraordinarily weak. How easily the output plot could have looked like the price plot. (I will admit that it is there, though, and somewhat more robust than it seems. A few hours of outlier-removing did not change the output sign from negative to positive.)

Methodologically, I think this is a nice way to show just what evidence we have. Hidden in the fancy response function and monte carlo is a summary of history. Here we see directly the historical experience on which a claim of output and price effects of monetary policy are built.

Reflections on responses.

These responses, like those of the VAR literature, may seem simple but to me they are profoundly troubling.

First, there is a 30 month output delay, and the output effect is still going at 48 months. this is awfully long for any theory!

Second, the *absence* of much price response – the dog that did not bark – seems to me a huge puzzle for the following reasons.

1. The data say "price puzzle." This is strong in my evidence, and the 30 month delay in the Romer-Romer shocks amounts to the same thing. This "fact" can be removed a bit by adding commodity prices to the shock regression, as shown in the paper. But the evidence for the Greenbook sufficient statistic property is strong, and if the Fed is not responding to commodity prices, we should not orthogonalize our shocks with commodity prices. Reading the minutes, there is *no* evidence of "well these nerds are wrong, commodity prices are going up, we need to tighten to stop this inflation." *This just does not happen.* Either RR give a dreadfully wrong reading of FOMC minutes, or adding commodity prices to fix the price puzzle is just a mistaken fishing expedition to make the data give the preordained right answer.

2. All theories tie output and price responses together. You can't have one without the other; its like ordering eggs sunny side down and don't turn it over. The data say otherwise. if you take the data seriously, you need a theory that interest rate policy affects y *with no effect on p at all.*

3. Current policy wisdom is that the Fed should control price level alone. Yet *there is no evidence that fed interest rate policy has any effect on the price level.* (I say "interest rate policy" deliberately. Of course the Fed could print money – though this is in fact a fiscal move. The question is whether interest rate policy of the sort practiced in sample has any effect on the price level.) *We must regard the "fact" that the Fed controls the price level by interest rate policy, and the consequent policy advice that it should do so, as matters of faith, a-priori beliefs, rather than matters supported by the historical experience. .*

5. Long-delayed responses?

There is one way around the finding of long-delayed responses. The last time I had the pleasure of discussing a Romer-Romer paper led me to think about this a lot, so I'll present the results here. Notice that after a target rise, the target keeps rising for quite a while.

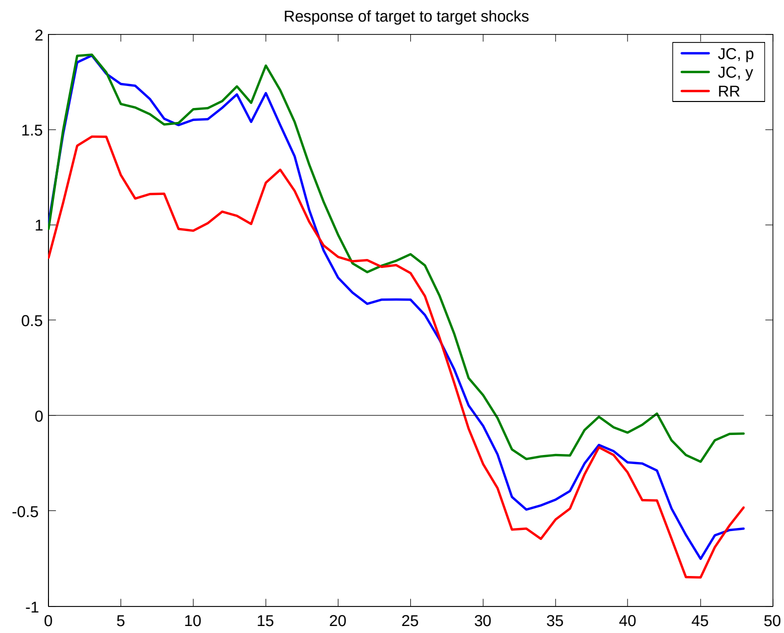

Now, does the *initial* FF shock cause output to change 2 years later, or does the *later increase* in ff, when it happens, cause output to rise? We can allow the latter view, but it requires that we believe that anticipated target rises can also affect output and prices.

This is clearly in the philosophy of what we are doing. Nothing in what I have done, or in what RR do, imposes the view that "only unanticipated ff matters." In particular, we have made no effort to make shocks orthogonal to agent's information sets.

Most theories allow for anticipated effects. (Though it usually is not the case that there is no difference, as we are implicitly assuming. For example, if prices are sticky, firms will start raising them ahead of an anticipated loosening.) All theories allow anticipated money to affect prices.

Thus, perhaps we should characterize the "effects of monetary policy" by "what if Fed increases ff *for only one month*" not "what if Fed increases ff and *keeps increasing it as it usually does?"*

RR have a nice discussion of this possibility but no calculation. Here is a preliminary calculation.

You can form the calculation just from the output and price responses.

$$
\begin{aligned}
y \text{ response:} \qquad & y_t = a(L)\varepsilon_t + ... \\
ff \text{ response:} \qquad & ff(t) = b(L)\varepsilon_t + ...
\end{aligned}
$$

$$y_t = a(L)b(L)^{-1}ff_t + ...$$

$$\text{simulate } b(L)y_t = a(L)ff_t$$

The fact that the responses I estimated so far are a little jagged, since they have no imposed autoregressive structure, makes the results uncomfortably jagged. (There are undoubtedly much better ways to do this!) Therefore I smoothed the responses with a centered moving average, which I present in the thick red and blue lines.

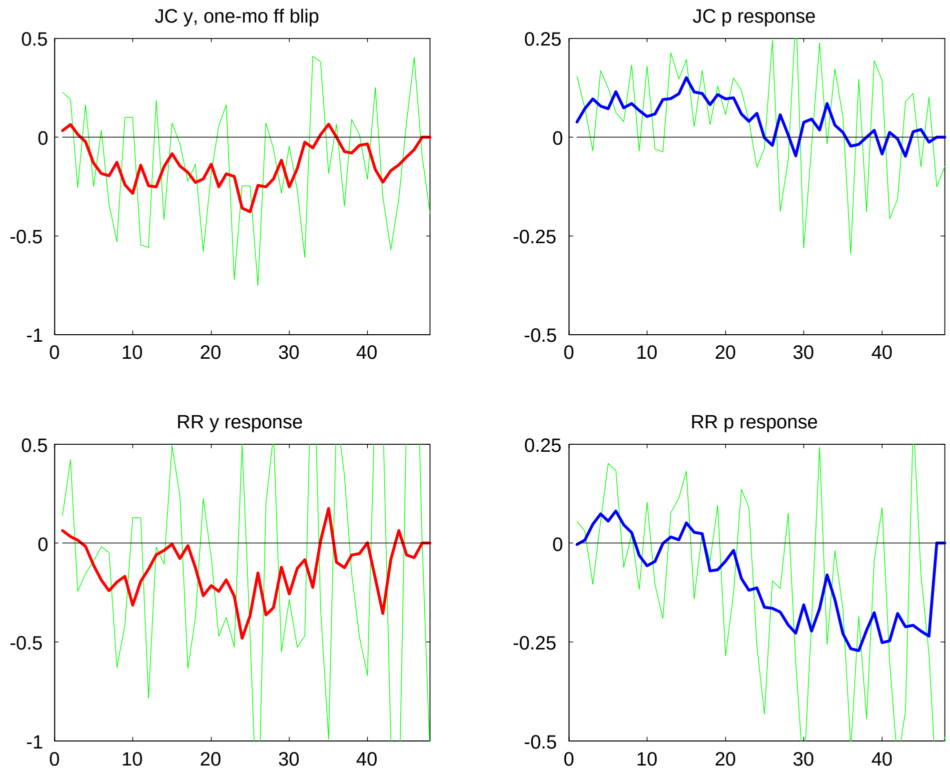

Effect of a one-month 1% increase in ff target, *not* followed by customary further increases.

As you can see, the effect of a one-month 1% increase in ff target, *not* followed by customary further target increases, is much shorter than the usual response. It *ends* at $t = 30$, rather than *peaking* at $t = 30$.More importantly, it's 1/10 smaller, peaking at at –0.4-0.5% not -6-8%

If we do not believe that "only unanticipated money matters" this seems a more natural summary of "the effects of m policy". It tells you the effect of decision *today*, not a path of decisions. A theorist might have a much easier time getting a model to match these shorter and smaller responses than the long drawn out response above.

This is a preliminary calculation of course. I think the right answer is shorter still. My target response ends sharply at 24 months, while the output response peaks at 30 months, so some dynamics are still necessary. I'm not convinced the sudden end of the target response is real; if it were drawn out we could account for the long output response with a much shorter structural output response.

The next step here is clear. Theories all *do* specify anticipated / unanticipated distinctions, and they impose strong cross-equation restrictions linking the output and price responses. Imposing those seriously is the right way to go.

Note also that the anticipated/unanticipated question has more impact on the results than most other specification questions!

6. What next?

Perhaps we are asking the wrong question. The *rule*, not shocks and responses, may be more interesting. After all, this is the end of a 40 year quest starting with the "St. Louis Fed" regressions, and monetary theory has changed since then. I would love for next year's

Romer-Romer paper to tell us more about the policy rule. There are two reasons.

1) (Minor) The weakest link in the chain here is the regression that defines shocks. The regression works over the whole sample, despite our strong suspicion that the rule changed over time. Narrative evidence on the rule would allow us to avoid full-sample regression and uncertainty in turning Δintentions into shocks.

Ideally, I dream that Christina and David could summarize a FOMC meeting by "we'll raise 50 basis points to fight this weak dollar, and another 25 to slow down inflation a bit. The economy is going just fine right now." Then we could assign a 50 basis point price-measurement shock and a 75 basis point output-measurement shock with no regressions at all.

2) (Major.) For many questions the rule is more important anyway. Most importantly, Taylor rule economics is all focused on the rule:

$$\begin{aligned} ff_t &= a + \phi(\pi_t - \bar{\pi}) + b(y_t - \bar{y}) \\ \phi &> 1? \end{aligned}$$

$\phi > 1$ is the crucial question. Clarida, Gali and Gertler and others run regressions and say the conquest of inflation comes from a change from $\phi < 1$ to $\phi > 1$.

However, I think the following is true (this is something I'm working on right now)

Conjecture 3. *Regressions cannot answer the determinacy question. In an economy with* $\phi > 1$, *regressions will show* $\phi < 1$.

(Hansen-Sargent, Farmer, Upcoming JC)

Here's a simple example. Suppose there is a structural model of the "good" sort. Nominal rate rise more than one for one with inflation, and if real rates are constant this means inflation is explosive:

$$\begin{aligned} \pi_{t+1} &= 1.1\pi_t + v_{t+1} \\ v_{t+1} &= 0.9v_t + \varepsilon_{t+1} \end{aligned}$$

Following the Taylor-rule philosophy, we solve the unstable root forward

$$\pi_t = E_t \sum_{j=1}^{\infty} \frac{1}{1.1^j} v_{t+j} = \sum_{j=1}^{\infty} \frac{0.9^j}{1.1^j} v_t = 4.5v_t$$

then, the reduced form is

$$\pi_{t+1} = 0.9\pi_t + 4.5\varepsilon_{t+1}$$

Notice that you *measure* the stable root, not the unstable root.

More deeply, these models work by specifying *off-equilibrium* threats (explosive inflation) that we never see *in* equilibrium. We can never measure off-equilibrium threats

Here the narrative method could really shine; it can try to measure a parameter of crucial, fundamental importance, that *cannot* be measured by econometric means.