Import Prices and Inflation*

James D. Hamilton
Department of Economics,
University of California, San Diego

Understanding the consequences of international developments for domestic inflation is an extremely important question for central banks. But before we claim to have measured the extent of import price pass-through, it is necessary to be clear on exactly what such a number is intended to mean. One can attempt to come up with an answer on the basis of either economic theory or empirical evidence. There are important pitfalls associated with either approach and significant benefits from combining the two.

JEL Codes: E31, F42.

1. On Measuring Inflation Pass-Through

A key question of interest to central banks is how international developments and prices of imported goods influence the domestic inflation rate. Two papers in this volume study different aspects of this question. Gianluigi Ferrucci, Rebeca Jiménez-Rodríguez, and Luca Onorante examine the effect of international food prices on domestic retail prices, while Anna Lipinska and Stephan Millard look at the consequences for the developed economies of the phenomenal growth in the emerging economies. Before getting into the details of the analyses, I think it's worth taking a moment to clarify the question that we are hoping to answer.

What effect do import prices have on the domestic inflation rate? A few examples should convince us that this is not a well-posed question. Consider first a policy experiment that we used to describe as a fully anticipated 5 percent increase in the rate of growth of the domestic money supply or, in the context of the model in Lipinska and Millard's paper, a perfectly anticipated 5 percent increase in the

*Author e-mail: jhamilton@ucsd.edu.

domestic central bank's long-run inflation target. In the absence of frictions, this would be associated with a 5 percent increase in the euro prices of imported goods and a 5 percent increase in domestic inflation. One might claim in such a setting that the "pass-through" from imported prices to domestic inflation was one for one, though in fact the true causation runs in the opposite direction.

More generally, insofar as the domestic inflation rate is at least influenced by domestic monetary policy decisions, if the central bank were to ask an academic economist, "what effect do import prices have on domestic inflation?" the academic might be tempted to answer, "whatever effect you want them to have." If central banks ask this question, there clearly must be some implicit assumptions about their own ability or willingness to offset whatever the effects might be, that is, some baseline scenario in which we're clear about the time horizon considered and exactly what we would mean by "in the absence of a response by the central bank."

A second example illustrates some further pitfalls associated with the question. Suppose that new fears of European financial stability lead to depreciation of the euro. Other things equal, we would expect this to be associated with an increase in the euro price of imported goods. Yet the broad consequences of a financial panic are predominantly deflationary, and we should not be surprised to observe the overall domestic price level actually decline despite the rising cost of imports. In such a case, we might assert that the pass-through parameter could actually be a negative number.

What these two examples have in common is that the increase in the price of imported goods was itself the result of a separate third cause. In the first example, this third cause was domestic monetary expansion, and in the second it was a domestic financial crisis. In each case, these developments not only had implications for the prices of imported goods but in addition had direct implications for domestic inflation that operate through channels other than import prices themselves. Thus if we ask the question, "what is the pass-through from import prices to domestic inflation?" the unavoidable answer is, "it depends on what caused import prices to go up in the first place."

Let us assume for the sake of discussion that we have agreed upon the necessary other conditions to have arrived at a well-posed question, namely, that we are seeking the domestic consequences of

a particular exogenous development arising outside our own country over a specified short horizon with a particular assumed response of the central bank. What tools could we use to try to answer such a question?

One approach is to use a theoretical model, which is the choice selected by Lipinska and Millard. The advantage of this approach is that we can specify exactly the question we are asking, in terms of the nature of the exogenous development, time frame of interest, assumed response of the central bank, and what is being held constant. The disadvantage is that we may not be convinced that the theoretical model accurately describes what actually happens in practice. Most of us would like to see such calculations corroborated by some evidence of what actually was observed to happen in real-world experience.

An alternative approach is empirical analysis, such as adopted by Ferrucci, Jiménez-Rodríguez, and Onorante. Here, too, there is an unambiguous question to which the analysis is giving us an answer. Whenever we do a regression of y on x we can expect to find reasonable estimates of the population linear projection of y on x, that is, a sensible answer to the following conditional forecasting question—historically, when x went up, on average by how much did y change? The disadvantage is that this historical correlation need not give us the answer to the question in which we're interested, as the two examples above illustrate.

With that as background, I now offer some specific thoughts on the two papers.

2. Lipinska and Millard

Lipinska and Millard's use of a fully specified dynamic stochastic general equilibrium model allows them to be very clear about the question they address. Specifically, the authors seek to analyze the effects of a temporary increase in productivity from emerging economies, assuming that the domestic economy adheres to a Taylor rule and using a standard calibration of wage and price stickiness. Lipinska and Millard distinguish between a headwind effect, in which strong growth from emerging economies bids up the prices that developed economies must pay for goods like oil and copper,

and a tailwind effect, in which productivity gains in the emerging economies help keep inflation low. In their benchmark parameterization (and essentially all the alternative specifications considered), the tailwind effects dominate. Lipinska and Millard conclude that strong growth in countries like Brazil, Russia, India, and China (BRIC) should have been a moderating influence on G-7 inflation rates.

While their model suggests that an increase in BRIC productivity would reduce G-7 inflation, the same simulations imply that BRIC inflation rates should decline by even more than do G-7 rates. But while real GDP growth rates have indeed been observed to have been remarkably higher in the BRICs than in the G-7, the BRIC inflation rates have also been substantially higher, not lower (see the authors' figures 1 and 2). The authors attribute the differences in GDP growth in figure 1 to a temporary increase in BRIC productivity growth, though this shock is modeled as exhibiting a quarterly autocorrelation coefficient of 0.99, implying extremely slow return to steady-state growth. By contrast, the differences in inflation in figure 2 are evidently modeled as permanent, reflecting an unexplained preference of the BRIC central banks for permanently higher inflation rates. Although one might argue on the basis of the history of countries like Brazil that such a specification is not unwarranted, we would also want to entertain the possibility that the observed differentials in inflation and GDP growth could in fact be related. Of course it's unreasonable to attribute BRIC growth entirely to a monetary-induced boom—if a central bank can produce China's growth rate by being sufficiently expansionary, then by all means let's all do it! But the possibility that credit expansion associated with the productivity boom is also related to some of the observed inflation should not be dismissed. The fact that China has largely pegged its currency to the dollar is another reason one might see a relation between their observed growth and inflation rates. However, the authors investigate this in one of their alternative scenarios and find in their figure 7 that pegging the currency would mitigate the size of the predicted inflation drop for the BRICs but not reverse the sign. In any case, the asymmetry with which the observed differences in growth rates are viewed as temporary while the observed differences in inflation rates are interpreted as a permanent and unrelated phenomenon is potentially troubling.

Another point worth noting in the Lipinska and Millard simulations is that the potential headwind effect implied by their parameterization appears implausibly small—the large increase in BRIC productivity only leads to a 0.5 percent increase in oil prices in their baseline scenario. By contrast, the real price of oil was actually observed to increase 150 percent between January 2005 and June 2008. Kilian (2009) and Kilian and Hicks (2011) attributed this price increase primarily to growth in demand, particularly from the emerging economies, while Hamilton (2009) attributed it to the interaction of strong emerging-economy demand with stagnating global production levels. The fact that the increase in global oil prices and BRIC inflation was substantially different from that predicted by the Lipinska and Millard simulations must leave us wondering how much confidence to place in the model's predictions for inflation rates in the G-7.

3. Ferrucci, Jiménez-Rodríguez, and Onorante

In contrast to the theoretical approach in Lipinska and Millard, the paper by Ferrucci, Jiménez-Rodríguez, and Onorante relies on an empirical investigation of the historical correlations. The question posed in their paper is what effects changes in the prices of internationally traded foods have on retail food prices in the euro area. Their answers are based on regressions of the change in consumer prices for a particular food item on its own lags, lags of changes in corresponding producer price index for the item, and lags of measures of raw commodity prices for the item. Similar regressions were estimated with the producer price index or raw commodity prices on the left-hand side to get a complete vector autoregression (VAR).

Ferrucci, Jiménez-Rodríguez, and Onorante make two key innovations relative to the previous literature. First, for their measure of raw commodity prices they used European Union (EU) internal market prices for commodities that are produced in the EU. These differ from the price of the same commodity traded internationally as a result of factors such as subsidies, tariffs, price supports, and minimum prices. They found, not surprisingly, that EU consumer prices tracked EU commodity prices more closely than they did international commodity prices, and suggest that use of international prices

rather than EU internal prices may account for the weak evidence of pass-through obtained by earlier researchers.

Which explanatory variable one wants to use, international prices or internal EU prices, depends on what question one is trying to answer. If the interest is in characterizing the details by which changes in bulk prices get passed on to retail customers, the specification of Ferrucci, Jiménez-Rodríguez, and Onorante is obviously to be preferred. On the other hand, from the perspective with which I began my comments above—what is the short-run effect of prices of imported goods on domestic inflation—one could make a case for using the international prices instead. Indeed, Ferrucci, Jiménez-Rodríguez, and Onorante have documented how EU price stabilization policies helped insulate consumers from fluctuations in the prices of internationally traded commodities, or at least did prior to 2005, when EU prices were as often below as they were above the international prices, helping smooth out the volatility (see their figure 2). On the other hand, EU prices followed international prices much more closely subsequently, so that for the most recent data, the Ferrucci, Jiménez-Rodríguez, and Onorante estimates may provide a better prediction of what would have happened as a result of the increase in international agricultural prices in 2007–08.

A second key innovation in their study is exploration of the possibility of a non-linear dependence of retail prices on the raw commodity prices, reflecting such factors as asymmetries (prices may increase faster than they decrease) and volatility or threshold effects. Here I would like to add some comments on the recent critique of these kinds of non-linear dynamic relations offered by Kilian and Vigfusson (2011). For illustration, consider the case in which price increases (c_t^+) are distinguished from price decreases (c_t^-),

$$c_t^+ = \begin{cases} c_t & \text{if } c_t > 0 \\ 0 & \text{otherwise} \end{cases}$$

$$c_t^- = \begin{cases} c_t & \text{if } c_t \leq 0 \\ 0 & \text{otherwise.} \end{cases}$$

Suppose that increases in raw commodity prices are allowed to have a different effect on consumer prices (π_t) than do decreases,

$$\pi_t = \kappa + \alpha_1\pi_{t-1} + \cdots + \alpha_p\pi_{t-p} + \beta_1 c^+_{t-1} + \cdots + \beta_p c^+_{t-p} + \gamma_1 c^-_{t-1} + \cdots + \gamma_p c^-_{t-p} + \delta_1 q_{t-1} + \cdots + \delta_p q_{t-p} + \varepsilon_t, \quad (1)$$

where q_t denotes producer prices. The point I would like to emphasize is that the observations by Kilian and Vigfusson in no way undermine the validity of using ordinary least squares (OLS) to estimate dynamic relations like (1). If ε_t in (1) were conditionally $N(0, \sigma^2)$ (and there is nothing in the specification that suggests it could not be), then OLS estimation of (1) will yield asymptotically efficient estimates of the true conditional expectation of consumer price inflation given historical values of the various explanatory variables. Furthermore, the various tests of (1) reported by Ferrucci, Jiménez-Rodríguez, and Onorante will give a valid inference about the form that the non-linearity would take.

Instead, the Kilian-Vigfusson concerns arise for estimation of the other equations of the VAR in which the non-linear transformation appears as the dependent variable:

$$c^+_t = k + a_1\pi_{t-1} + \cdots + a_p\pi_{t-p} + b_1 c^+_{t-1} + \cdots + b_p c^+_{t-p} + g_1 c^-_{t-1} + \cdots + g_p c^-_{t-p} + d_1 q_{t-1} + \cdots + d_p q_{t-p} + e_t. \quad (2)$$

Here it is not possible for e_t to have expectation zero conditional on any value for the right-hand variables. The optimal forecast of c^+_t could not be a linear function of the information set, since we know that c^+_t has to be positive. This matters potentially when we seek to use the VAR system built from elements like (1) and (2) to form impulse response functions. An additional issue is the interpretation of these impulse response functions. Because c^+_t is necessarily positive, an increase in food prices that is less than the average value of c^+_t would be associated with a negative value for the shock e_t, which may not correspond to the question the researcher has in mind.

There are three possible approaches to this problem. The first, which is adopted in most of the literature and in the body of Ferrucci, Jiménez-Rodríguez, and Onorante's paper, is to ignore these issues, be satisfied with (2) as a linear projection that may not differ too much from the true forecast, and use standard linear VAR methods to calculate dynamic responses. A second approach is to employ the methods proposed by Kilian and Vigfusson. This the authors do in a separate appendix to their paper, and they find that the results are

quite similar to those in the main body, providing some reassurance that the usual approximations work reasonably well in the current setting. A third alternative applied in my recent paper (Hamilton 2011) to a related problem adapts the suggestion of Jordà (2005), which is simply to estimate a separate forecasting equation of the form of (1) for each horizon s of interest:

$$\pi_{t+s} = \kappa_s + \alpha_{1,s}\pi_{t-1} + \cdots + \alpha_{p,s}\pi_{t-p} + \beta_{1,s}c^{+}_{t-1} + \cdots + \beta_{p,s}c^{+}_{t-p} + \gamma_{1,s}c^{-}_{t-1} + \cdots + \gamma_{p,s}c^{-}_{t-p} + \delta_{1,s}q_{t-1} + \cdots + \delta_{p,s}q_{t-p} + \varepsilon_{t+s,s}. \quad (3)$$

Again, OLS for any s provides optimal estimates of the $(s+1)$-period-ahead conditional expectation. The fundamental questions one would seek to answer with a standard impulse response function (namely, how does new information about the value of c_{t-1} cause me to revise my forecast of π_{t+s}?) can be answered directly from the data using simple OLS regressions.

Ferrucci, Jiménez-Rodríguez, and Onorante's conclusion from their exercise is that estimated pass-through elasticities align reasonably closely with factor shares. And although Lipinska and Millard do not summarize their results this way, that may not be a bad takeaway from their paper either. That's also the rough rule of thumb that a policymaker might have used if there were no guidance at all from either theory or historical experience—if the price of imported commodities increase 10 percent and imported commodities account for 20 percent of the value of final spending, then we might expect this to push domestic inflation up by 2 percentage points. The different issues encountered by these two papers in trying to approach the question from either theory or evidence remind us of how difficult it can be to improve on such rules of thumb, and give us some comfort when they arrive at a similar answer.

References

Hamilton, J. D. 2009. "Causes and Consequences of the Oil Shock of 2007–08." *Brookings Papers on Economic Activity* (Spring): 215–61.

———. 2011. "Nonlinearities and the Macroeconomic Effects of Oil Prices." *Macroeconomic Dynamics* 15 (Supplement 3).

Jordà, Ò. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–82.

Kilian, L. 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review* 99 (3): 1053–69.

Kilian, L., and B. Hicks. 2011. "Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003–2008?" Working Paper, University of Michigan.

Kilian, L., and R. J. Vigfusson. 2011. "Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?" *Quantitative Economics* 2 (3): 419–53.