Volume Title: Commodity Prices and Markets, East Asia Seminar on Economics, Volume 20

Volume Author/Editor: Takatoshi Ito and Andrew K. Rose, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-38689-9
ISBN13: 978-0-226-38689-8

Volume URL: http://www.nber.org/books/ito_09-1

Conference Date: June 26-27, 2009

Publication Date: February 2011

Chapter Title: Introduction to "Commodity Prices and Markets, East Asia Seminar on Economics, Volume 20"

Chapter Authors: Takatoshi Ito, Andrew K. Rose

Chapter URL: http://www.nber.org/chapters/c11855

Chapter pages in book: (1 - 12)

Introduction

Takatoshi Ito and Andrew K. Rose

The chapters of this volume were first presented at the twentieth annual East Asian Seminar on Economics, and are all focused on the theme of commodity prices and markets.

Commodity price fluctuations represent tremendous challenges and opportunities for economists and policymakers. For economists, they often represent large, plausibly exogenous shocks of tremendous importance, especially for small open economies. For societies dependent upon commodity sales, price changes are either disasters or windfalls. The statistical linkages between commodity prices—especially the price of oil—have been well-known for a long period of time, though the exact interpretation is not universally agreed upon. For policymakers of oil-importing countries, oil price increases represent an adverse supply shock, posing difficult policy options. Due to an increase in imported energy price, domestic prices tend to increase and output tends to be depressed. Monetary easing (intended to help stimulate aggregate demand) may result in further increases in prices, which may prompt workers' demand wage increases to maintain real wages. Moreover, lower interest rates may depreciate the exchange rate, further aggravating inflation. An inflationary spiral of price and wage increases may thus be ignited. Knowing this risk, an inflation targeting central bank might be reluctant to relax monetary policy to support output. On the other hand, if the central bank tightens monetary policy in fear of inflation, output activities will contract further, deepening the recession. Although a contrac-

Takatoshi Ito is professor in the Graduate School of Economics at the University of Tokyo, and a research associate of the National Bureau of Economic Research. Andrew K. Rose is the B.T. Rocca Professor of Economic Analysis and Policy at the Haas School of Business, University of California, Berkeley, and a research associate of the National Bureau of Economic Research.

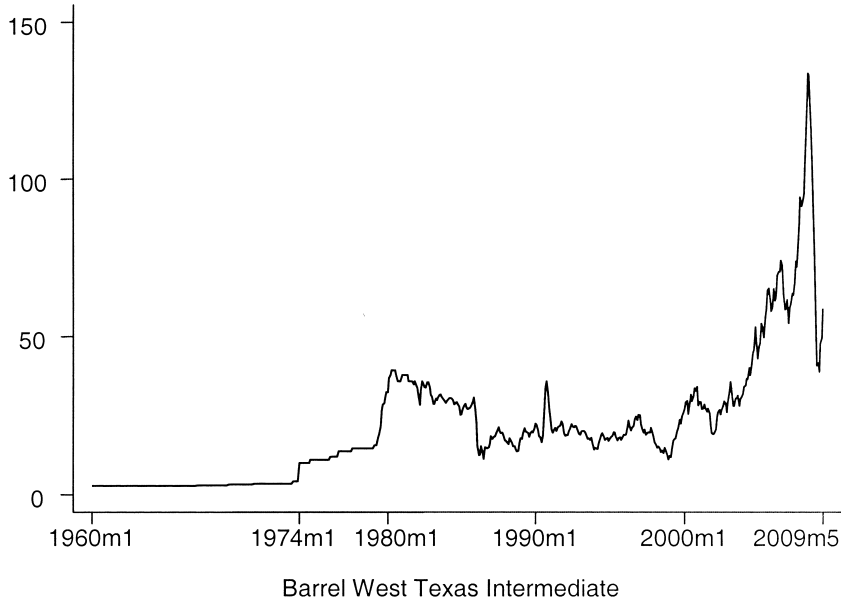

Fig. I.1 Nominal price of oil

tion in output is inevitable in the face of the adverse supply shock, and the standard of living of the oil-importing economy has to fall, it may be difficult for policymakers to convince the public of this consequence.

Before we introduce the chapters, we first provide a brief background to the material. We begin with oil prices. Oil represents one of the most important sources of energy used in both advanced and developing economies. Other fossil fuels (e.g., coal and natural gas) are also used widely around the world, especially to produce electricity. Other important sources of energy exist, including hydro, nuclear, and renewable power such as wind and solar. Still, oil has played, and will continue to play, a vital role in a number of key transportation types (e.g., cars, trucks, ships, and airplanes) in the foreseeable future. Fluctuations in oil prices are imperfectly but substantially associated with business cycle fluctuations, especially in the United States. Thus, it is eminently reasonable to begin with a brief overview of the oil market.

Figure I.1 provides a time-series plot of a standard measure of the price of oil (in this case, the nominal dollar price of a barrel of West Texas Intermediate Crude). Several features of the data jump out, upon even casual observation. First, the price of oil was exceptionally smooth until 1973. In fact, the price of oil fluctuated between only \$2.57/bbl (oil barrel) and \$3.56/bbl from January 1948 until July 1973. In the 1950s, oil "majors" controlled the oil price, and in the 1960s the OPEC (Organization of Petroleum Exporting Countries) became influential in fixing oil prices. This period of

exceptional tranquility ended with the first oil price shock. Oil prices jumped from around $3/bbl to $12/bbl in a matter of a few months, in the aftermath of a war in the Middle East. After another period of low volatility, the price of oil again rose steeply between the summers of 1979 and 1980, a period known widely as the second oil price shock. The oil market has been relatively volatile since the second oil crisis, and has exhibited some dramatic price fluctuations. Of special note are the collapse of oil prices in the mid-1980s (possibly due to increases in new supply from the North Sea and other oil field discovery and delivery), the spike in oil prices associated with the Gulf War of 1990, and the very dramatic run-up in oil prices, which began around 2004. The price movement in the latest episode is said to be amplified by speculative money. Finally, the global recession of 2008 to 2009 coincides with an equally dramatic collapse in oil prices.

What caused these oil price changes? Hamilton has argued in a long series of papers (e.g., Hamilton 2009) that supply disruptions and discovery of the North Sea oil have characterized most of the post-war oil price shocks. However, he argues that the most recent increases in the price of oil seem to have been caused by strong demand confronting stagnant production. This characterization has been widely but not universally accepted; Kilian (2008) argues that most oil price shocks have been driven by global demand for industrial commodities (including crude oil), along with shocks to the precautionary demand for crude oil. In any case, it is unclear that the nature of the shock has important consequences for the macroeconomy of a typical oil importer.

As is well-known, the timing of the price changes seems to coincide remarkably well with macroeconomic fluctuations. The big oil price increases associated with OPEC-I, OPEC-II, the 1990 Gulf War, and 2007 to 2008 all closely lead business cycle downturns that hit a number of industrial countries at about the same time. Accordingly, much conventional macroeconomic theory takes oil price shocks as exogenous supply shocks. Significant increases in the price of oil might be expected to create both inflation and recession; this "stagflation" has been famously modeled by Bruno and Sachs (1985) in an important analysis of the first two OPEC shocks. However, it is not completely obvious that oil price shocks need cause enduring inflation. Any recessionary effects of oil prices are necessarily transient, and oil price shocks are simply a large and important relative price change (rather than an intrinsic source of persistent changes in the absolute price level). Thus some authors, notably Bernanke, Gertler, and Watson (1997), have argued that the endogenous monetary response to oil price shocks (rather than the shocks themselves) are what has caused many of the adverse consequences of oil prices. In any case, the macroeconomy seems to respond differently to oil price shocks of late than it did during the period of the large OPEC shocks of the 1970s. Blanchard and Galí (2007) argue that this is the result of a combination of better monetary policy, more flexible labor markets, a

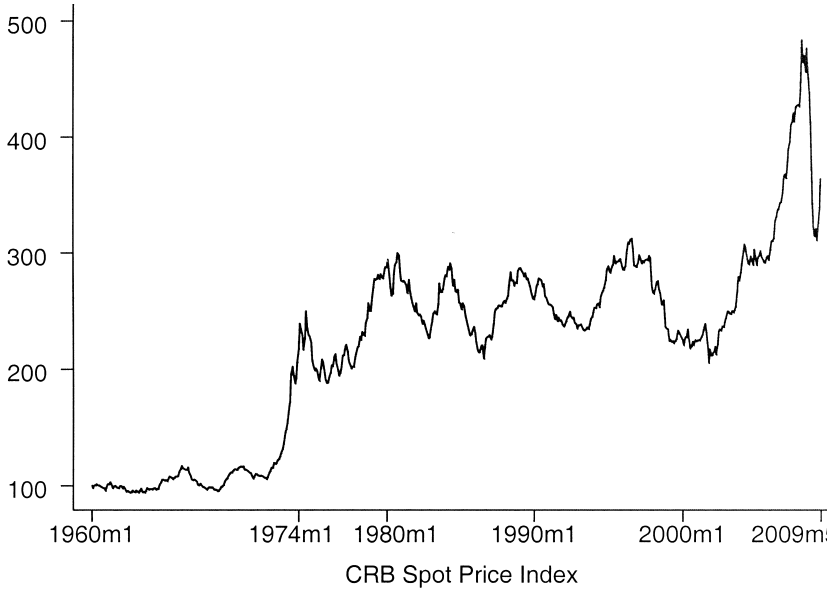

Fig. I.2 Nominal commodity prices

reduced importance of oil, and simple good luck (in that adverse oil price shocks did not coincide with other shocks).

Where figure I.1 portrays the price of oil, figure I.2 is an analogous time-series plot of a commonly used index of commodity prices (in this case, the spot market index provided by the Commodity Research Bureau including twenty-two basic commodities). One of the striking features of the broader index of commodity price index is its close similarity to the oil price series of figure I.1 (even though the index does not include petroleum products).[1] The two series share a long period of low volatility at the beginning of the sample, which ends at around the same time. After thirty years of continuing volatility, they share the same enormous run-up, collapse, and small rebound during the dramatic events of 2007 to 2009. The common and dramatic nature of the recent price movements suggests that while oil and general commodity prices may help cause business cycle fluctuations, they may also be responsive to global economic fluctuations, as argued by Kilian (2008).

With these features of the data in mind, we now turn to the chapters of *Commodity Prices and Markets.* Rather than begin with the macroeconomy, we begin with more narrow examinations of commodity prices, then gradually begin to broaden our scope.

1. See http://www.crbtrader.com/crbindex/spot_background.asp.

Forecasting Currencies with Commodity Prices

In a well-known recent working paper, Chen, Rogoff, and Rossi (2008) (hereafter CRR) argue that the exchange rates of countries that produce disproportionate amounts of commodities can be used to forecast future commodity prices themselves. Forecasting commodity prices is not only important for a large number of businesses and official agencies, but has also proven to be remarkably difficult to forecast (even via prices of commodity futures). Thus, this positive result has the potential to be quite significant, if it stands up to scrutiny. In "Commodity Prices, Commodity Currencies, and Global Economic Developments," Groen and Pesenti (chapter 1, this volume) provide exactly the sort of independent robustness analysis that is standard in other sciences, but deplorably rare in economics.

Groen and Pesenti do a very thorough job of scrutinizing the results of CRR. Where CRR focus on one commodity price index at the quarterly frequency, Groen and Pesenti exhaustively analyze ten commodity price indices using monthly data to forecast ahead at four different horizons. More importantly, CRR are interested in asking whether commodity prices can be forecast with *either* exchange rates or arbitrary combinations of other macroeconomic and commodity-relevant variables. The latter is the most intriguing part of their study, since they cast their net widely to examine a host of economic "fundamentals" that might be of relevance in forecasting commodity prices. In fact, their potential set of fundamentals are so large as to require factor-analytic procedures to reduce the "curse of dimensionality," which might otherwise eliminate all available degrees of freedom. They compare statistical models that are augmented with extra information from either exchange rates or economic fundamentals to laughably simple time-series benchmark models such as the simple random walk popularized by Meese and Rogoff. Unfortunately, despite all the judicious use of econometric technology, Groen and Pesenti consistently find only weak results. Any improvement in forecasting commodity prices over simple benchmarks tend to be ephemeral, sensitive to the exact measure of commodity prices used in the forecasting index, the precise horizon, and so forth. This cannot be considered definitive until it completely encompasses the much more positive results first found by CRR. However, it represents serious pushback for those who believe in the forecasting ability of commodity currencies, and is likely to trigger a lively debate in the reviving field of exchange rate economics.

Why do Groen and Pesenti fail to find that commodity currencies do not provide significant forecasting ability for commodity prices? One possible explanation is that commodity markets are fully efficient, quickly and completely reflecting all possible sources of information. If commodity prices were fully efficient, then it might be impossible for currencies—or any other

information—to help forecast their movements. Considerable skepticism is warranted here, since futures markets for commodities have a long and ignoble history of providing awful forecasts for future spot prices of commodities. Still, the hypothesis seems worth investigating; in the (unlikely) event that it is a reasonable characterization, an efficient commodity market would lead to a number of radical conclusions concerning commodity prices. And surprisingly—at least to us—Chan, Tse, and Williams (chapter 2, this volume) examine futures markets for commodities and find considerable evidence consistent with the hypothesis of market efficiency in "The Relationship between Commodity Prices and Currency Exchange Rates."

Chan, Tse, and Williams are, like Groen and Pesenti, particularly interested in the relationship between commodity prices and currencies; they are similarly motivated by the findings of CRR. But, they depart from other investigators in looking at the price/currency relationships with higher-frequency data. Their daily data set has more observations (though a shorter span) than those of others working in the area, so that they can tease out even relatively short-lived forecasting relationships. They also depart from other researchers in using the futures data for commodities actually employed by speculators and hedgers, rather than the more illiquid spot prices. Their results are consistent with Fama's notion of "semistrong" market efficiency in that they find that a week's worth of lags in currencies returns do not, in fact, help improve forecasts of commodity returns. In this, they are consistent with the results of Groen and Pesenti (and inconsistent with those of CRR) in that they find no advantage to using currencies when forecasting commodity prices. Indeed, an (unjustified) extrapolation of their findings implies that *nothing* can be used to forecast future commodity returns.

We find this result to be narrow, but intriguing. The findings are narrow because they rely on linear Granger-causality tests for a limited number of currencies and commodities forecast at short horizons. Still, they are unusually positive in a literature often plagued by unexplained rejections of market efficiency. We look forward to future work in the area, as additional forecasting variables (such as the "order flow" used in foreign exchange markets) are applied with more general statistical techniques. If the results stand up—an outcome that is far from certain—they could represent the beginnings of a shift in one of the field's deeply held priors.

Commodity Prices, the Terms of Trade, and Exchange Rates

Commodity prices are of interest to economists for a great many reasons. One of these is that the characteristics of commodity prices are quite different from those of other prices in the economy. For instance, where many nominal prices of goods and services exhibit "stickiness" of one sort or another, it is well-known that commodity prices are volatile—indeed, scarily so. The effect of commodity price fluctuations on more traditional

prices can then be examined, and, indeed, this is done by a number of the following chapterq that estimate "pass-through" equations. We return to these issues later in the book.

Another difference between commodity and more traditional prices for goods and services is that commodity prices are often denominated in one currency—almost always the U.S. dollar—and thus show essentially full and instantaneous pass-through of exchange rate changes. For instance, if the dollar price of gold stays constant, but the euro/dollar exchange rate fluctuates, the price of oil in euros will change 1:1 with the exchange rate. This difference between commodities and a more conventional price series is exploited by Broda and Romalis (chapter 3, this volume) in a novel and interesting way. In "Identifying the Relationship between Trade and Exchange Rate Volatility," they use the difference between the negligible effects of exchange rate volatility on commodity trade with the fact that exchange rate volatility might have a large effect on trade in noncommodity sectors. This is a plausible identifying restriction that allows them to incorporate the possibility that trade—whether from commodities or other goods—might, in turn, have a feedback effect on exchange rate volatility. This clever identification strategy allows them to model the simultaneous determination of exchange rate volatility and trade in a way that is more satisfying and plausible than other approaches. Allowing for simultaneity turns out to have a big effect in practice; it dramatically lowers the much-disputed effect of currency unions on trade first identified by one of the editors.

It has long been known that shocks to commodity prices are responsible for much of the variation in a country's terms of trade, especially for developing countries that export disproportionate amounts of commodities. The terms of trade is conventionally defined as the ratio of a country's export to import prices. To measure the numerator of the terms of trade, one simply adds up the prices faced by the country's exporters, weighting by the importance of a particular good in a country's aggregate exports; the denominator is constructed analogously. Commodity-exporting countries tend to produce specialized export bundles, so that the volatility of their terms of trade tends to derive from shocks to their exports rather than their more widely diversified imports. All this is well-known and has been much studied.

Do the consumers of a country face the same terms of trade as producers? If not, then the marginal rate of substitution across goods faced by consumers differs from the marginal rate of transformation faced by producers. This inefficiency can potentially have serious consequences for welfare. It would also be interesting to understand the proximate causes for the wedges between consumer and producer prices, which are often caused by tariffs, taxes, and other aspects of government policy. These, in turn, may have political-economy origins that are also worthy of study. First, though, it is necessary to see if there are, in fact, nontrivial differences between the

traditional (producer) terms of trade and the consumers' terms of trade (hereafter CTT). This initial measurement task is the one set for themselves by Berka and Crucini (chapter 4, this volume) in "The Consumption Terms of Trade and Commodity Prices."

Berka and Crucini create consumer terms of trade using underlying price data from the Economist Intelligence Unit, which is a panel of disaggregated prices for a number of individual goods collected in a number of important global cities across time. Rather than using city-specific prices, however, the authors average prices across cities to create average consumer prices. This procedure might be problematic, given the existence of dramatic and persistent deviations from the law of one price studied by many authors (including Crucini). However, it does deliver a single set of worldwide consumer prices that can then be aggregated up to create the CTT when weighted by net export shares. (The aggregation scheme necessarily means that any good in which a country's trade is approximately balanced will contribute little to the CTT.) Since the study depends on a dubious assumption as well as a narrow panel of goods, cities, and time, we view this study as a preliminary one, necessary to determine whether a larger effort is warranted.

After creating the consumer terms of trade, Berka and Crucini then perform some intriguing empirical analysis. While their results are preliminary, at least two results are compelling. First, they find that just a few goods—oil, automobiles, and pharmaceuticals—contribute disproportionately to volatility in the CTT. Second, they find considerable differences between the traditional (producers') terms of trade and CTT. We view this work as far from definitive; the set of caveats necessary for a broad interpretation is long. However, the subject material has now passed this initial "smell test" and warrants a more comprehensive and grounded study; we look forward to more analysis from the authors in the future.

"Pass-through" of Commodity Prices to the General Price Level

Periods immediately after large increases in commodity prices—especially the price of oil—have sometimes been associated with inflation; these inflations have often been large and persistent. However, the great inflations associated with the first two OPEC price shocks of the 1970s seem now to be a thing of the past. During the last few years, oil prices have risen very dramatically without any clear and strong inflationary consequences. That is, the "pass-through" of oil prices (which are determined on global markets) to domestic prices seems to be changing. Given the extreme volatility of oil prices that the world has experienced of late, it is important to understand *why* pass-through patterns seem to be changing. Accordingly, three of the chapters in this volume examine the changing nature of pass-through.

In "Pass-Through of Oil Prices to Japanese Domestic Prices," Shioji and Uchino (chapter 5, this volume) study pass-through issues for Japan. They

begin by quantifying the size and nature of pass-through of oil prices to domestic in inflation for Japan. However, their real interest lies in understanding why the effects of oil price shocks seems to have declined so much over the last few decades. They use time-series techniques (vector autoregressions, or VARs), which allow for variation over time in their coefficients, and examine disaggregated price data. They compare a number of different and plausible explanations for declining pass-through. It might be the case that the Bank of Japan is slowly asserting its recent monetary independence and establishing greater credibility for monetary policy. If so, a more credible nominal anchor for the Japanese economy means that a given oil price shock is viewed increasingly as a shift in relative prices that need not result in greater inflation. Alternatively, more flexible Japanese labor markets might also result in lower pass-through. Another possible explanation is simply that energy is less important to the Japanese economy, in part because of the sharp increases in energy prices of the 1970s. Based upon a close examination of input-output tables for Japan, this turns out to be the most plausible explanation. It states, essentially, that Japanese firms reacted to the first two OPEC oil price shocks by changing their cost structures. For instance, switching to less oil-intensive production structures over time lessens the effects of subsequent oil price shocks (such as the one currently being experienced). The contrast between Japanese and American automotive firms is not only implicit, but striking. Still, Shioji and Uchino leave much room for future work. It is unclear which mechanism led to the changing responses of Japanese firms to energy price signals. Did the exchange rate play a role? Were public policies—especially energy taxes—important? How were consumers affected by all this? We enjoyed reading this step along the path toward a complete understanding of the effects of oil price shocks, and look forward to more.

Fukunaga, Hirakata, and Sudo (chapter 6, this volume) broaden the range of inquiry in "Effects of Oil Price Changes on the Industry-Level Production and Prices in the United States and Japan" in two key ways. They are interested in the effects of oil price shocks on prices (the focus of Shioji and Uchino) and also output. This allows them to characterize the consequences for the real economy both at a macroeconomic level, and at a more disaggregated industry level. They also do comparable work for both the United States and Japan. The choice of this pair of countries is natural: they constitute the two largest economies in the world, and have reacted quite differently to large changes in energy prices. Like Shioji and Uchino, Fukunaga and colleagues also use time-series data in their VAR-based empirical analysis. However, they diverge in exploiting a set of identification assumptions recently popularized by Kilian (2008).

Kilian's identification scheme relies on splitting the economy into three separate blocks, which are identified recursively. A global oil market depends only on itself, but also affects the aggregate macroeconomy. Both of these

markets spill out to affect individual industries. Oil price shocks can then be identified to be either demand or supply in nature, and one can analyze their effects on different sorts of industries in both countries. We find this identification scheme eminently plausible, though the authors' modeling of the domestic macroeconomies is narrower than we would prefer (omitting, for instance, the exchange rate and all measures of monetary policy). In any case, the empirics deliver plausible results concerning the effects of different oil shocks, particularly for the United States, on oil and nonoil industries. As one would expect, the effects of oil price changes depend on the exact nature of the shock and of the industry being affected. More interestingly, the transmission of different shocks is very different for the United States and Japan. The United States, for example, is oil-intense and less export-dependent compared with Japan. Perhaps most intriguing is the fact that there seem to be weaker, or even positive effects of unexpected oil (supply) price increases on Japan, a result that stands in contrast with strong negative effect on the United States. The authors provide a plausible explanation for their findings, since the effect of oil price increases can potentially be positive because Japan produces energy-efficient goods (e.g., fuel-efficient Japanese cars). While we do not believe that the authors have presented enough evidence to be completely persuasive, we consider this finding fascinating, and well worth further study.

Kuo and Peng (chapter 7, this volume) are also interested in pass-through issues and study them using Taiwanese data in "Price Pass-Through, Household Expenditure, and Industrial Structure: The Case of Taiwan." Theirs is a primarily descriptive analysis that characterizes Taiwan, reasonably enough, as an emerging economy that is not yet rich enough to have the low pass-through effects that characterize a typical advanced economy. They find that when global commodity prices rise, only around a fifth of these increases eventually show up in domestic consumer prices of energy and food; these are two of the most volatile parts of consumer prices, which are substantially driven by commodity prices. These components are also worthy of study since their volatility means that they are often omitted from inflationary measures that focus on core underlying inflation. Kuo and Peng analyze the effects of these price changes on household expenditures using coarsely disaggregated Taiwanese data and the "Almost Ideal Demand System" developed by Deaton and Muellbauer. The causes, however, of these relatively small responses are left unmodeled, so it is difficult to know how to interpret these effects or forecast them in the future.

Macroeconomic Effects of Oil Prices

As we noted earlier, the periods of time after oil price increases have historically been associated with recessions in richer economies, which tend to be oil importers (with a few exceptions, such as Norway and the United

Kingdom, after the discovery of the North Sea oil field). This might be because of the direct consequences of higher oil prices for production and consumption. Alternatively and plausibly, it may well be the indirect result of endogenous monetary policy actions that are induced by oil prices. However, both monetary policy and the apparent effects of oil prices seem to have changed of late. Monetary authorities around the world have slowly gained inflation credibility, often by adopting inflation-targeting regimes; simultaneously, the effects of higher oil prices on the economy seem to have shrunk. Is this a coincidence? In "Oil and the Macroeconomy," Lee and Song (chapter 8, this volume) investigate these possibly interrelated issues and find that in the case of Korea, monetary policy is indeed managed close to optimally, at least when it comes to the responsiveness of monetary policy to oil price shocks.

Lee and Song first establish that oil price rises seem to have a different effect on the Korean economy recently than they did in the 1970s, which were days of the large shocks. They do this through means of a conventionally identified VAR model of the Korean economy estimated on two separate periods of time (before the Asian crisis of 1997, and after 2000). They find that the adverse effects of oil prices on Korean real gross domestic product (GDP) have indeed shrunk during the later period. This result is consistent with intuition, as well as the well-known results of Blanchard and Galí (2007). Still, the real question is why exactly has the adverse effect of these supply shocks diminished over time? Since one of the most obvious answers is the response of monetary policy, Lee and Song construct a conventional dynamic stochastic general equilibrium (DSGE) model with both oil and monetary policy built in. After estimating the model, they are able to compare the actual response of monetary policy to oil shocks, with the optimal response delivered by their model. Since the two are relatively close, they conclude that the Bank of Korea seems to have used its monetary independence well (though some caution is necessary since some of the estimates are quite imprecise).

A completely different take on the role of oil in the Korean economy is provided by An and Kang (chapter 9, this volume) in "Oil Shocks in a DSGE Model for the Korean Economy." Like Lee and Song, the authors provide a state-of-the-art modern macroeconomic model of the Korean economy. The model is a relatively conventional one of a small open economy, but it has been augmented to allow for oil imports to be used either for consumption or production purposes. The authors fit the model to fifteen years of recent Korean data, and then use their model to make a set of somewhat unusual comparisons. Specifically, they are interested in understanding how much worse the model performs if oil is excluded entirely from (a) Korean production or (b) Korean consumption. Precisely what one is to make of this comparison is not completely clear to us, though we find it interesting that removing the oil inputs to production has a much more substantive

effect than removing the consumption channel. Whether this results from the absence of capital in the model (which limits the amount of factor substitutability in production) is also unclear. We think of this as a very hypothetical set of thought experiments estimated in a sophisticated manner (though with relatively little data); an answer, in other words, awaiting a suitable question.

References

Bernanke, B. S., M. Gertler, and M. Watson. 1997. Systematic monetary policy and the effects of oil price shocks. *Brookings Papers on Economic Activity* 1: 91–142.

Blanchard, O. J., and J. Galí. 2007. The macroeconomic effects of oil shocks: Why are the 2000s so different from the 1970s? NBER Working Paper no. 13368. Cambridge, MA: National Bureau of Economic Research, September.

Bruno, M., and J. Sachs. 1985. *Economics of worldwide stagnation.* Oxford: Basil Blackwell.

Chen, Y.-C., K. Rogoff, and B. Rossi. 2008. Can exchange rates forecast commodity prices? NBER Working Paper no. 13901. Cambridge, MA: National Bureau of Economic Research, March.

Hamilton, J. 2009. Causes and consequences of the oil shock of 2007–08. NBER Working Paper no. 15002. Cambridge, MA: National Bureau of Economic Research, May.

Kilian, L. 2008. The economic effects of energy price shocks. *Journal of Economic Literature* 46 (4): 871–909.