

COMMODITY PRICES, CONVENIENCE YIELDS, AND INFLATION

Nikolay Gospodinov and Serena Ng*

Abstract—This paper provides evidence that the two leading principal components in a panel of 23 commodity convenience yields have statistically and quantitatively important predictive power for inflation even after controlling for unemployment gap and oil prices. The results hold up in out-of-sample forecasts, across forecast horizons, and across G7 countries. The convenience yields also explain commodity prices and can be seen as informational variables about future economic conditions as conveyed by the futures markets. A bootstrap procedure for conducting inference when the principal components are used as regressors is also proposed.

I. Introduction

MONETARY authorities seem to hold a long-standing view that commodity prices have inflationary consequences, and thus the ability to predict future commodity price movements can be important for the time path of economic policies. As the Federal Reserve chairman, Ben Bernanke, remarked,

Rapidly rising prices for globally traded commodities have been the major source of the relatively high rates of inflation we have experienced in recent years, underscoring the importance for policy of both forecasting commodity price changes and understanding the factors that drive those changes (2008).

In spite of the general view that commodity price movements have inflation implications, the formal link between inflation and commodity prices is not thoroughly understood. Some argue that commodity prices are leading indicators of inflation because they respond quickly to general economic conditions. Others believe that idiosyncratic movements in commodity prices work through the distribution channel and subsequently affect prices in general. While both explanations are plausible, the magnitude of the commodity price effect on inflation must necessarily depend on what triggers changes in commodity prices. In particular, commodity prices can change as a result of transactions, speculative demand, or precautionary demand, and each of these demand components can have a different impact on inflation.

At the empirical level, the prices of oil and various aggregate commodity price indices have been found to improve inflation forecasts on a simple autoregressive (AR) benchmark for certain countries and periods. However, the evidence

is far from robust. Hooker (2002) finds a statistically significant impact of oil prices on core U.S. inflation over the period 1962 to 1980 but not for the post-1981 period. Several other studies also report an improved ability of commodity prices for forecasting inflation until the mid-1980s but a substantial deterioration in the predictive power of commodity prices from 1985 onward.¹ Stock and Watson (2003) show that these forecasting improvements are sporadic and unstable. The problem is partly that commodity prices tend to be volatile and are themselves difficult to forecast.

This paper develops a framework for analyzing commodity prices and inflation. We use data on commodity price futures in the form of convenience yields. Although the evidence on the predictive power of futures prices for commodity and general price movements is somewhat mixed, there is a perception that commodity futures prices contain, “substantial amount of information about the supply and demand conditions that are aggregated by futures markets” (Bernanke, 2008). Our empirical approach is unique in that we extract information from a convenience yields panel using principal components. These principal components in convenience yields have strong and systematic predictive power for commodity price changes and inflation rates of the United States and other G7 countries. The results hold up even when interest rate, trade-weighted exchange rate, demand pressure variables (such as unemployment gap), and oil price are controlled for. In contrast, the IMF aggregate commodity index has little predictive power for inflation. We attribute this finding to the ability of the principal components in convenience yields to isolate variations in commodity prices that have inflationary consequences.

The rest of the paper proceeds as follows. Section II establishes a relationship between commodity prices and convenience yields from which a relation between inflation and convenience yields is obtained. The data and the construction of principal components are discussed in section III. Section IV presents a bootstrap procedure to assess the sampling error from using principal components as regressors. Predictive regressions and out-of-sample forecasts of inflation for the United States and the other G7 countries are presented in section V. The relation of convenience yields with individual commodity prices and the IMF commodity price index is studied in section VI. Supplementary results are provided in an online appendix.

II. The Determination of Commodity Prices and Inflation

Commodities share similar characteristics with money in that they can be held for everyday use, can be stored, and can be used as an asset. It is thus useful to think of the demand

¹ See Blomberg and Harris (1995) and Furlong and Ingenito (1996), among others.

Received for publication June 9, 2010. Revision accepted for publication June 23, 2011.

*Gospodinov: Concordia University and CIREQ; Ng: Columbia University.

We thank the editor, two anonymous referees, Lutz Kilian, Benoit Perron, participants at the Applied Econometrics Workshop at the Federal Reserve Bank of St. Louis, and the 4th CIREQ Time Series Conference for helpful comments. We also thank Ibrahim Jamali and Van Hai Nguyen for research assistance. N. G. gratefully acknowledges financial support from FQRSC, IFM2 and SSHRC. S. N. acknowledges financial support from the National Science Foundation (SES-0962431).

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00242.

for commodities as arising from one of three sources: (a) a component tied to current consumption and production, (b) a precautionary demand component that reflects future needs, and (c) an asset demand component that depends on its risk over the holding period relative to its potential for capital gains. As Deaton and Laroque (1992) noted, the market clearing price for commodities is a function of the availability (new production plus inventories) relative to expected total demand (current plus precautionary plus asset demand).

A. A Model for Commodity Prices and Inflation

Let S_{jt} and $F_{jt,n}$ denote the spot and futures price of commodity j for delivery at time $t+n$. Also, let $i_{t,n}$ be the nominal interest earned between period t and $t+n$. Define the basis to be the difference between the futures and the spot price, $F_{jt,n} - S_{jt}$. In the speculative storage model of commodities, inventories are held only if the expected returns are positive. Thus, S_{jt} and $F_{jt,n}$ are both functions of the level of inventories. Kaldor (1939) posits that negative basis $F_{jt,n} - S_{jt}$ consists of two components: (a) an opportunity cost of forgone interest from having to borrow and buy the commodity and (b) a (net of insurance and storage costs) convenience yield $CY_{jt,n}$. With the convention that $CY_{jt,0} = 0$, Fama and French (1987) formalize the relation between convenience yields and the basis as

$$F_{jt,n} - S_{jt} = S_{jt}i_{t,n} - CY_{jt,n}. \tag{1}$$

Kaldor (1939) uses the convenience yield to reflect the benefit from using the stored commodity whenever desired. However, in Pindyck's (1993) model of rational commodity pricing, convenience yields play the role of dividends that anticipate future changes in spot commodity prices. When the timing or the level of consumption is stochastic, holding commodity inventories acts as an insurance against unexpected price movements or demand shifts.² Thus, the convenience yield $CY_{jt,n}$ is a forward-looking variable that contains information not exclusive to future demands. The negative of $CY_{jt,n}$ is also referred to in the finance literature as interest-adjusted basis (Fama & French, 1988). Although $CY_{jt,n}$ is unobserved, it can be computed from the observed spot price, futures price, and interest rate using equation (1). If the spot price is approximated by the price of the nearest futures contract, the basis has the form of a futures price spread.

An alternative view of the basis is provided by the theory of normal backwardation according to which risk-averse investors earn a risk premium for the fluctuations in the future spot price. The basis then comprises a risk premium

²Gorton, Hayashi, and Rouwenhorst (forthcoming) provide empirical evidence that convenience yields are a decreasing and possibly nonlinear function of inventories. Alquist and Kilian (2010) show that shifts in uncertainty about future excess demand are reflected in the fluctuations of oil, which are in turn correlated with fluctuations in the precautionary demand component of the real spot price of oil.

component $\Psi_{jt,n} \equiv E_t S_{jt+n} - F_{jt,n}$ and an expected price change component $E_t S_{jt+n} - S_{jt}$ so that

$$F_{jt,n} - S_{jt} = E_t S_{jt+n} - S_{jt} - \Psi_{jt,n}. \tag{2}$$

When $\Psi_{jt,n}$ is positive, futures price is backwardated (at a discount). In the storage model, low inventory leads to low basis and, subsequently, low returns on owning the commodity.

Whereas equation (1) implicitly defines the convenience yield, equation (2) implicitly defines the risk premium. Together, equations (1) and (2) imply

$$E_t S_{jt+n} - S_{jt} = S_{jt}i_{t,n} - CY_{jt,n} + \Psi_{jt,n}.$$

Let $cy_{jt,n} = CY_{jt,n}/S_{jt}$, $\psi_{jt,n} = \Psi_{jt,n}/S_{jt}$, and $E_t \Delta^n s_{jt+n} = (E_t S_{jt+n} - S_{jt})/S_{jt}$. Then,

$$E_t \Delta^n s_{jt+n} = i_{t,n} + \psi_{jt,n} - cy_{jt,n}. \tag{3}$$

According to equation (3), the expected percentage change in commodity prices has three components: (a) a component $i_{t,n}$ related to the opportunity cost of buying and holding inventories, (b) a risk premium component $\psi_{jt,n}$, and (c) an expected marginal convenience yield component $cy_{jt,n}$. The intuition for the inverse relationship between the expected commodity price changes and convenience yield is provided by Fama and French (1988), who suggest that a permanent increase in the current and future commodity demand has a differential impact on current and expected spot prices. In particular, current spot prices should increase more than expected spot prices, especially at low inventory levels, because the demand and supply responses of consumers and producers tend to partially offset the effect of the shock on expected prices. As a consequence, higher convenience yields and lower inventory levels are associated with lower expected spot prices.

The three components in equation (3) are not mutually uncorrelated as equations (1) and (2) are alternative decompositions of the basis. It is useful to define

$$\epsilon_{jt,n} = \psi_{jt,n} - \text{Proj}(\psi_{jt,n} | i_{t,n}, cy_{jt,n})$$

to be the component of $\psi_{jt,n}$ that is orthogonal to the interest rate and the convenience yield. Equation (3) can be represented as

$$E_t \Delta^n s_{jt+n} = \alpha_{j1} cy_{jt,n} + \alpha_{j2} i_{t,n} + \epsilon_{jt,n}. \tag{4}$$

Because α_{j1} and α_{j2} are now "reduced-form" coefficients, they are not constrained to the parameter values implied by equation (3).

To link the determinants of commodity prices to inflation, let P_t denote the economy's general price index. Let $q_{jt} = \tilde{q}_{jt} - q_{jt}^*$ be the real price of commodity j defined in the currency of country k relative to its equilibrium, where

$\tilde{q}_{jt} = S_{jt}/P_t$ and q_{jt}^* is the equilibrium value of \tilde{q}_{jt} . Assuming that real commodity prices are mean-reverting as in Frankel (2006), we can write the n -period expected inflation rate $E_t \Delta^n p_{t+n} = (E_t P_{t+n} - P_t)/P_t$ as

$$E_t \Delta^n p_{t+n} = \theta_j q_{jt} + E_t \Delta^n s_{jt+n}$$

for some $\theta_j < 0$. Substituting the expression for $E_t \Delta^n s_{jt+n}$ in equation (4) yields the inflation equation

$$E_t \Delta^n p_{t+n} = \theta_j q_{jt} + \alpha_{j1} cy_{jt,n} + \alpha_{j2} i_{t,n} + \epsilon_{jt,n}. \quad (5)$$

If $\alpha_{j1} = -\alpha_{j2} = -1$, rearranging terms would imply that the disequilibrium in real commodity price q_{jt} is inversely proportional to the real carrying cost, which equals the real interest rate minus the net convenience yield.

For open economies, Frankel (2006) suggests that a monetary expansion in country k will lead to a lower interest rate, causing commodity prices in local currencies to rise until an expected depreciation restores equilibrium. Chen, Rogoff, and Rossi (2010) find that exchange rates forecast commodity prices. This motivates the open economy analog of equation (4):

$$E_t \Delta^n s_{jt+n} = \alpha_{j1} cy_{jt,n} + \alpha_{j2} i_{t,n} + \alpha_{j3} \Delta^n x_t^k + \epsilon_{jt,n}, \quad (6)$$

where x_t^k denotes the nominal exchange rate between the United States and country k . The corresponding inflation equation is then given by

$$E_t \Delta^n p_{t+n} = \theta_j q_{jt} + \alpha_{j1} cy_{jt,n} + \alpha_{j2} i_{t,n} + \alpha_{j3} \Delta^n x_t^k + \epsilon_{jt,n}. \quad (7)$$

In this model for inflation, $cy_{jt,n}$ is commodity specific by definition, and there is an exchange rate with every trading partner. Because of the heterogeneous nature of the commodities, one can expect the effects of an aggregate measure of commodity prices on inflation to be different from those of the individual commodity prices. Indeed, Mishkin (1990) notes that individual convenience yields (or basis) are too noisy to be useful predictors.

Let \bar{c}_t denote an aggregate measure of convenience yields and \bar{q}_t be an aggregate measure of disequilibrium real commodity price. Also, let \bar{x}_t be a (trade-weighted) average of exchange rates. The aggregate analog of equation (7) is

$$E_t \Delta^n p_{t+n} = \theta \bar{q}_t + \alpha_1 \bar{c}_t + \alpha_2 i_{t,n} + \alpha_3 \Delta^n \bar{x}_t + \epsilon_{t,n}. \quad (8)$$

This open-economy, commodity-based inflation model posits that the expected n -period-ahead inflation is a function of the aggregate disequilibrium in real commodity prices, aggregate convenience yield, interest, and exchange rates. Note that the use of aggregate convenience yield for predicting inflation is new and has not been previously explored in the literature. In the analysis that follows, we view convenience

yields as an informational variable that reflects future economic conditions as perceived by the commodity market participants. However, we are agnostic about the structural sources of variations in convenience yields, which could be due to production, hedging, and speculative demand for commodities.

III. Data

This section discusses the data used in the empirical work. The daily commodity price data are obtained from the Commodity Research Bureau and are available at daily frequency for March 1983 to July 2008. The data set contains spot and futures prices of 23 commodities from six commodity groups: energy (crude oil, heating oil), foodstuffs (cocoa, coffee, orange juice, sugar), grains and oilseeds (canola, corn, oats, soybeans, soybean oil, wheat), industrials (cotton, lumber), livestock and meats (cattle feeder, cattle live, hogs lean, pork bellies), and metals (copper, gold, palladium, platinum, silver). This particular choice of a cross-section of commodity prices and time period is dictated by data availability. The spot price is approximated by the price of the nearest futures contract, and the futures price is the price of the next-to-the-nearest futures contract. The time separating the nearest and next-to-the-nearest futures contracts typically differs across commodities and may not be equally spaced over the course of the year. As a result, we do not exactly match the contract maturities with the forecast horizon. Longer contracts are not used due to lack of continuous and liquid record of the corresponding maturity futures prices over our sampling period.

As consumer price data are observed at no higher than monthly frequency, monthly commodity price series are constructed from daily data by averaging the daily prices in the corresponding month. The real commodity prices q_{jt} are obtained by deflating the spot prices by the U.S. CPI (seasonally adjusted) index obtained from the Bureau of Labor Statistics (BLS). The three-month U.S. T-bill rate and the exchange rate data (U.S. dollar trade-weighted index against major currencies) are from the Federal Reserve Economic Database (FRED). We use the IMF commodity nonfuel price index to approximate the aggregate behavior of commodity spot prices.³ The detrended real commodity price q_{jt} is obtained as deviations from the Hodrick-Prescott (HP) trend estimate. The smoothing parameter for the HP filter is set to its default value for monthly data of 14,400. Results based on data detrended by a one-sided exponentially weighted moving average filter are similar.

As noted earlier, convenience yield is a forward-looking variable. We proxy this forward-looking information using

³The data for CPI, interest and exchange rates, and commodity price index can be downloaded from <http://research.stlouisfed.org/fred2/> and <http://www.imf.org/external/np/res/commod/index.asp>.

data on commodity price futures. Specifically, the percentage net convenience yield for commodity j is computed as

$$cy_{jt,n} = \frac{(1 + i_{t,n})S_{jt} - F_{jt,n}}{S_{jt}}, \quad (9)$$

using the three-month U.S. Treasury bill as $i_{t,n}$, adjusted for the time that separates the two futures contracts. Note that while commodity spot and futures prices may be volatile and subject to occasional spikes, the convenience yield (being the difference between the two series) should be less noisy. Furthermore, the data for spot and future prices tend to be highly persistent and possibly nonstationary, while convenience yields exhibit milder persistence. Using convenience yields in the empirical work also sidesteps inference issues that may arise when regressors are nonstationary.

To conserve space, we highlight only a few key features of the data.⁴ The convenience yields are persistent, with most of the first-order autocorrelation coefficients around 0.8, the largest being 0.936 for coffee. The detrended real commodity prices q_{jt} are also persistent with autocorrelation coefficients between 0.743 and 0.912. Surprisingly, the q_{jt} series in the food group are more volatile than the real prices of crude and heating oil. The convenience yields of precious metals such as gold and silver have variances that are orders of magnitude smaller than those for the other commodities. The convenience yields of sugar, oats, hogs, pork bellies, and industrials appear to be most volatile. Fama and French (1987) suggest that the differences in the variability of convenience yields could arise from seasonal variations and adjustments to demand and supply shocks. The summary statistics in the online appendix also reveal that grains and oilseeds are typically in contango, with $\text{prob}(cy_{jt,n} > 0) < 0.5$. In contrast, livestock and meats and metals (except for silver) are most of the time in backwardation, with $\text{prob}(cy_{jt,n} > 0) \geq 0.5$.

A. Constructing \bar{c}_t and \bar{q}_t by Principal Components

Our basic premise is that convenience yields contain information about commodity prices, and commodity prices anticipate inflation. The inflation model presented in section II uses aggregate measures of commodity prices and convenience yields as predictors but leaves open the question as to how the aggregate measures are to be constructed. Boughton and Branson (1991) suggest that an aggregate measure of commodity prices would not predict inflation if idiosyncratic supply shocks are not accommodated by the monetary authorities. They conjecture that future inflation may be better predicted by commodity prices, which are driven primarily by demand shocks. Kilian (2009) suggests that the macroeconomic effects of oil price shocks depend on the source of the shock. However, isolating the variations in commodity prices that have pure inflationary consequences is not a trivial exercise.

Hong and Yogo (2009) find that a simple average of the individual convenience yields helps to predict commodity prices, although it is uncorrelated with future stock and bond returns, as well as short-term interest rate, yields spread, and dividend yield. They attribute the predictive power of the aggregate convenience yield across commodity groups to its ability to pick up different types of shocks. To some extent, we share this view, though our focus here is inflation, not commodity price forecasts.

We extract principal components from our panel of 23 convenience yields. Similarly, principal components are extracted from 23 detrended real commodity prices. We denote these by $pccy_t$ and pcq_t , respectively. Prior to the computation of the principal components, the variables $cy_{jt,n}$ and q_{jt} are standardized. In brief, the first r principal components of the convenience yields are the eigenvectors corresponding to the largest eigenvalues of the $N \times N$ matrix $(NT)^{-1}cy'_n cy_n$, where cy_n is a $T \times N$ matrix of convenience yields, with $N = 23$. By construction, these principal components are mutually orthogonal but are unique only up to a column sign change.⁵

The principal components are linear combinations of the individual series constructed to best explain the total variation in the data. We use them as a statistical tool to extract the most important information in convenience yields without insisting on the presence of a factor structure. If convenience yields are indeed driven by latent common factors, then the principal components consistently estimate the space spanned by the common factors as the number of commodities tends to infinity. Pindyck and Rotemberg (1990) find excess comovement in commodity prices, while Tang and Xiong (2009) show that commodity prices have been increasingly exposed to macroeconomic shocks; a factor representation of commodity prices is defensible, though this is not crucial to the analysis.

The first two principal components $pccy_t = (pccy_t^{(1)}, pccy_t^{(2)})'$ explain 23% of the variance in convenience yields, computed as a ratio of the sum of the first two eigenvalues to the sum of all eigenvalues. There is little evidence that additional principal components are needed in the predictive regressions considered. Both $pccy_t^{(1)}$ and $pccy_t^{(2)}$ are persistent, with first-order autocorrelation coefficients of 0.93 and 0.81, respectively. The two principal components for disequilibrium commodity prices $pcq_t = (pcq_t^{(1)}, pcq_t^{(2)})'$ explain about 31% of the variance in the panel of data on q_{jt} . The first-order autocorrelation coefficients for the two principal components are 0.90 and 0.88, respectively.

In results presented in the online appendix, we find that the first element of $pccy_t$ loads heavily on corn, coffee, cotton, wheat, and crude oil, while the second principal component is highly correlated with some metals (silver, copper), as well as soybeans, cocoa, and heating oil. The first element

⁴ We refer readers to the online appendix for full details.

⁵ While the signs of the principal components can be set to be consistent with a microfounded model of commodity prices and inflation, equations (4) and (5) are reduced-form models.

of pcq_t is closely associated with the commodity group of grains and oilseeds, while its second element captures the price variation of the metals group. These principal components have dynamics that are distinct from those of the energy commodities.

IV. Bootstrap Inference with Principal Components as Regressors

This section proposes a bootstrap method that can account for the fact that the principal components are generated regressors, as well as small sample distortions that may arise from time series estimation of predictive regressions. Several bootstrap procedures for factor models have been suggested (without proof of their asymptotic validity) for idiosyncratic errors that are identical and independently distributed (i.i.d.) across units and over time, which is too restrictive for the data being analyzed. As well, we want to allow for but do not want to impose a factor structure on the data. In other words, we want to leave open the possibility that the principal components are simply weighted averages of the individual convenience yields, which are meaningful predictors in their own right. Given these considerations, a bootstrap procedure for estimation using principal components as predictors seems more appropriate.

Let x be a generic $N \times T$ data matrix, where x_{it} ($i = 1, \dots, N$, $t = 1, \dots, T$) denotes the i th observed series at time t , N is the total number of variables (convenience yields, detrended real commodity prices), and T is the number of time series observations. The first r principal components of matrix x , denoted by pcx_t , are the eigenvectors corresponding to the largest eigenvalues of the $N \times N$ matrix $(NT)^{-1}xx'$.

Our interest lies in conducting statistical inference in the predictive regression

$$y_{t+h} = \beta' pcx_t + \gamma' w_t + \varepsilon_{t+h},$$

where w_t denotes a $p \times 1$ vector of other observable predictors (interest rate, exchange rate, real oil price), as well as deterministic terms and lag values of y_t , and the errors ε_{t+h} are possibly autocorrelated and heteroskedastic. The predictive regression is estimated by OLS, and the inference procedure on the estimated parameters should potentially take into account that pcx_t are “generated” regressors.

Under suitable regularity conditions (Bai & Ng, 2006), the OLS estimator $(\hat{\beta}', \hat{\gamma}')$ is root- T consistent and asymptotically normal. Furthermore, the presence of generated predictors pcx_t does not require any adjustments to the standard errors of the parameter estimates provided that $\sqrt{T}/N \rightarrow 0$. Unfortunately, in our analysis, the cross-sectional dimension N is relatively small, and the regularity conditions of Bai and Ng (2006) may not hold. As a result, we resort to bootstrap methods for inference that account for uncertainty associated with the estimation of principal components. Our proposed bootstrap algorithm is based on a moving-block resampling of the original data.

More specifically, we stack the data from both stages, $x_{1t}, \dots, x_{Nt}, y_{t+h}$ and $w_{jt}, j = 1, \dots, p$ (after truncating the last h observations of x_{it}), into the matrix

$$Z = \begin{bmatrix} x_{11} & \dots & x_{N1} & y_{h+1} & w_{11} & \dots & w_{p1} \\ x_{12} & \dots & x_{N2} & y_{h+2} & w_{12} & \dots & w_{p2} \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ x_{1(T-h)} & \dots & x_{N(T-h)} & y_T & w_{1(T-h)} & \dots & w_{p(T-h)} \end{bmatrix}.$$

The bootstrap samples $(x_{1t}^*, \dots, x_{Nt}^*, y_{t+h}^*, w_{1t}^*, \dots, w_{pt}^*)$ for $t = 1, \dots, T - h$ are then obtained by drawing with replacement blocks of $m = m_T \in \mathbb{N}$ ($1 \leq m < T$) observations from matrix Z . This ensures that the bootstrap samples preserve possible model misspecification, serial correlation, heteroskedasticity, and cross-sectional dependence in the data. The block size m is allowed to grow, but at a slower rate, with the time series dimension T .

Let z_t be the t th row of the data matrix Z above. Also, let $B_{t,m} = (z_t, z_{t+1}, \dots, z_{t+m-1})$ denote a block of m consecutive observations of z_t , $k = \lceil T/m \rceil$, where $\lceil a \rceil$ signifies the largest integer that is less than or equal to a , and $\bar{T} = km$. We resample with replacement k blocks from $(B_{1,m}, B_{2,m}, \dots, B_{\bar{T}-m+1,m})$ by drawing k i.i.d. uniform random variables $[u_1], \dots, [u_k]$ on $(1, k + 1)$. Then the bootstrap sample is given by $Z^* = [(z_1^*, z_2^*, \dots, z_m^*), (z_{m+1}^*, z_{m+2}^*, \dots, z_{2m}^*), \dots, (z_{\bar{T}-m}^*, z_{\bar{T}-m+1}^*, \dots, z_{\bar{T}}^*)] = (B_{[u_1],m}, B_{[u_2],m}, \dots, B_{[u_k],m})$.

For each bootstrap sample $\{x_{it}^*\}$ ($i = 1, \dots, N$, $t = 1, \dots, \bar{T}$), the r principal components (denoted by pcx_t^*) are reestimated as the largest eigenvalues of the $N \times N$ matrix $(N\bar{T})^{-1}x^*x^{*'}.$ These are plugged into the predictive regression for the bootstrap data,

$$y_{t+h}^* = \hat{\beta}^{*'} pcx_t^* + \hat{\gamma}^{*'} w_t^* + \hat{\varepsilon}_{t+h}^*,$$

where $\hat{\beta}^*$ and $\hat{\gamma}^*$ are OLS estimates.

Let $\hat{\beta}_l$ ($l = 1, \dots, r$) denote the estimated coefficient on the l th principal component from the observed sample and $\hat{\beta}_{l,j}^*$ be the estimated coefficient from the j th bootstrap sample with corresponding standard errors $s.e.(\hat{\beta}_l)$ and $s.e.(\hat{\beta}_{l,j}^*)$ computed using a heteroskedasticity and autocorrelation consistent (HAC) estimator. Due to the sign indeterminacy of the principal components, we set the sign of pcx_t^* to be consistent with the dynamics of pcx_t estimated from the original sample. We then construct the sequence $t_{\alpha,l,j}^* = (\hat{\beta}_{l,j}^* - \hat{\beta}_l) / s.e.(\hat{\beta}_{l,j}^*)$, sort it in ascending order, and let v_{α}^* and $v_{(1-\alpha)}^*$ denote the α th and $(1 - \alpha)$ th elements of the sorted sequence for a prespecified nominal level α . The $100(1 - \alpha)\%$ equal-tailed percentile- t bootstrap confidence interval for β_l is obtained as $[\hat{\beta}_l - s.e.(\hat{\beta}_l)v_{(1-\alpha/2)}^*, \hat{\beta}_l + s.e.(\hat{\beta}_l)v_{\alpha/2}^*]$. In addition to better approximating the small-sample distribution of $\hat{\beta}_l$, these bootstrap confidence intervals take into account the estimation uncertainty for the generated regressors (principal components). Confidence intervals for the remaining coefficients are obtained in a similar manner. Note that this bootstrap

procedure allows for possible asymmetry in the finite-sample distribution of the parameter of interest. In the empirical analysis, we set $m = 4$ although the results are similar for other values of m in the range $m \in [4, 24]$. The number of bootstrap replications B is 4,999.

Since this inference procedure is used for both the aggregate commodity price and inflation equations, in the empirical analysis, we stack the two dependent variables, $\Delta^h s_{t+h}$ and $\Delta^h p_{t+h}$, along with all regressors in the matrix Z , which is resampled as described above. We then select the appropriate columns of the bootstrap matrix Z^* to conduct the principal component analysis and estimate the corresponding predictive regressions. Some Monte Carlo simulation results on the coverage rates of the proposed bootstrap confidence intervals are presented in the online appendix.

V. Convenience Yields as Predictors of Inflation

This section analyzes the predictive ability of convenience yields for the inflation rates in the United States and the other G7 countries. We also report out-of-sample forecast results.

A. Results for the United States

Inflation is a notoriously challenging series to forecast, and it is not easy to find predictors that have systematic predictive power when augmented to an univariate autoregression. Stock and Watson (2007) find that an unobserved components model is often better at forecasting quarterly inflation as measured by the GDP deflator than models that explicitly use observed predictors. The unobserved components model implies that $\Delta p_{t+1} - \Delta p_t$ has a negative moving average component. However, Stock and Watson (2007) also report that the unobserved components model is less appropriate for quarterly inflation as measured by the (all items) CPI. Given that the commodity price data are daily, our empirical exercise focuses on forecasting monthly CPI inflation. We consider six measures of inflation based on (a) CPI all items, (b) CPI less food and energy, (c) CPI less food, (d) CPI less energy, (e) CPI food only, and (f) CPI energy only. The CPI series, taken from the BLS, are seasonally adjusted. Since it is widely documented that U.S. inflation is persistent, we start with an AR model for inflation and ask what variables have additional predictive power beyond lags of inflation. We also vary the prediction horizon, denoted by h , which may differ from the futures contract maturity n .

According to the model given in section II, inflation should be predicted by disequilibrium in commodity prices and convenience yields. Our predictive equation for h -period inflation $\Delta^h p_{t+h} = (P_{t+h} - P_t)/P_t$ is

$$\begin{aligned} \Delta^h p_{t+h} = & b + \beta_0(L) \Delta p_t + \beta_1(L)' pccy_t + \beta_2(L)' pcq_t \\ & + \beta_3(L) z_t + v_{t+h}. \end{aligned} \quad (10)$$

In the base case, z_t is empty. The Akaike information criterion (AIC) suggests that $\beta_1(L) = \beta_1$, $\beta_2(L) = \beta_2$, and $\beta_0(L)$ is a

second-order polynomial in the lag operator. This lag order is used for all in-sample inflation models and all forecast horizons. Any remaining serial correlation, such as serial correlation induced by overlapping data, is accounted for by computing HAC standard errors.⁶ The optimal lag length will likely vary across models and forecast horizons. However, the goal of these regressions is to assess which predictors are relevant, and imposing a common lag length simplifies the presentation of results. Optimal lags will be used in the subsequent out-of-sample forecasting analysis.

Table 1 presents the parameter estimates from model (10) with $h = 1, 3, 6$, and 12, along with their corresponding 90% bootstrap confidence intervals. The statistically significant coefficients at the 10% level are in bold. The second principal components of convenience yields $pccy_t^{(2)}$, and real commodity prices $pcq_t^{(2)}$ tend to be strongly significant at short horizons. Furthermore, $pccy_t^{(2)}$ remains significant at longer horizons. We attribute the significance of $pccy_t^{(2)}$ to the fact that the principal components isolate those variations, common across individual convenience yields and real commodity prices, that are relevant for predicting inflation. Further investigation reported in the online appendix reveals that of the 23 convenience yields $cy_{j,t,n}$, the estimated coefficients of cocoa, orange juice, and copper are significantly positive, while the estimated coefficients of oats, soybeans, and silver are significantly negative. The effects of the other $cy_{j,t,n}$ on inflation are not statistically significant.

In the more general specifications, we let z_t be the detrended real price of crude oil and other determinants (interest rate i_t and the log change of the trade-weighted USD exchange rate index $\Delta \bar{x}_t$) of commodity price movements. As in the baseline specification, $pccy_t^{(2)}$ appears to contain important information for predicting inflation at all forecasting horizons. The real oil price possesses incremental explanatory power and captures some sharp movements in the inflation rate for all goods and services. While the interest rate is insignificant in a model with only autoregressive dynamics, it is significant in the augmented predictive regression. However, exchange rate is not significant at all horizons. Table 2 reports results using other measures of inflation. As expected, the detrended real oil price has predictive power for the inflation of CPI energy. It also helps predict the inflation of CPI less food. However, $pccy_t$ continues to predict other inflation measures, especially those based on the CPI less food and energy and the CPI less energy.

Recall that the purpose of our new bootstrap procedure is to control for sampling uncertainty associated with the estimated principal components and to better approximate the finite-sample distribution of the t test. The equal-tailed percentile- t bootstrap method also accounts for the skewness and miscentering in the finite-sample distribution of the parameter that arise from possible estimation bias. Bootstrap test statistics tend to be more conservative than asymptotic

⁶ We use Newey-West HAC standard errors with an automatic bandwidth selection.

TABLE 1.—ESTIMATION RESULTS FOR U.S. INFLATION (CPI ALL ITEMS)

$$\Delta^h p_{t+h} = b + \sum_{k=0}^1 \beta_{0k} \Delta p_{t-k} + \sum_{k=1}^2 \beta_{1k} pccy_t^{(k)} + \sum_{k=1}^2 \beta_{2k} pcq_t^{(k)} + z_t \beta_3 + v_{t+h}$$

	h = 1		h = 3		h = 6		h = 12	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
$pccy_t^{(1)}$	-0.010 [-0.11, 0.09]	0.085 [0.02, 0.24]	-0.064 [-0.49, 0.24]	0.150 [-0.11, 0.54]	-0.177 [-1.20, 0.42]	0.149 [-0.50, 0.77]	-0.337 [-2.54, 0.95]	0.163 [-1.36, 1.46]
$pccy_t^{(2)}$	-0.074 [-0.19, -0.00]	-0.087 [-0.20, -0.02]	-0.262 [-0.83, -0.00]	-0.285 [-0.76, -0.07]	-0.538 [-1.76, -0.09]	-0.556 [-1.83, -0.10]	-1.153 [-2.90, -0.59]	-1.143 [-2.59, -0.67]
$pcq_t^{(1)}$	-0.009 [-0.06, 0.03]	-0.050 [-0.11, -0.01]	-0.089 [-0.30, 0.07]	-0.108 [-0.32, 0.07]	-0.204 [-0.66, 0.07]	-0.105 [-0.38, 0.15]	-0.228 [-0.90, 0.40]	0.128 [-0.38, 0.79]
$pcq_t^{(2)}$	0.071 [0.03, 0.14]	0.039 [-0.03, 0.13]	0.178 [0.09, 0.32]	0.143 [-0.02, 0.33]	0.125 [-0.16, 0.43]	0.167 [-0.16, 0.50]	0.084 [-0.68, 0.75]	0.336 [-0.24, 0.95]
Δp_t	0.372 [0.28, 0.59]	0.177 [0.12, 0.34]	0.165 [-0.08, 0.49]	-0.180 [-0.37, 0.06]	0.261 [-0.01, 0.61]	0.033 [-0.19, 0.36]	0.695 [0.32, 1.11]	0.603 [0.29, 0.98]
Δp_{t-1}	-0.217 [-0.37, -0.08]	-0.138 [-0.31, -0.01]	-0.176 [-0.45, 0.07]	0.004 [-0.27, 0.26]	0.060 [-0.23, 0.38]	0.288 [-0.11, 0.75]	-0.029 [-0.55, 0.55]	0.426 [-0.11, 1.10]
z_t								
$qoil_t$	-	1.184 [0.88, 1.67]	-	1.853 [1.03, 3.39]	-	1.047 [-0.05, 2.83]	-	0.094 [-1.34, 1.82]
$qoil_{t-1}$	-	-1.095 [-1.45, -0.80]	-	-2.204 [-3.39, -1.55]	-	-2.374 [-4.09, -1.43]	-	-3.540 [-5.60, -2.39]
i_t	-	0.021 [0.01, 0.03]	-	0.053 [0.02, 0.09]	-	0.090 [-0.01, 0.16]	-	0.148 [0.03, 0.26]
$\Delta \bar{x}_t$	-	0.008 [-0.01, 0.02]	-	0.008 [-0.02, 0.04]	-	0.028 [-0.01, 0.08]	-	0.031 [-0.03, 0.09]
$u_t - u^*$	-	-0.010 [-0.05, 0.12]	-	-0.153 [-0.49, 0.11]	-	-0.225 [-0.90, 0.37]	-	-0.265 [-1.01, 0.35]
\bar{R}^2	0.160	0.302	0.066	0.220	0.114	0.282	0.175	0.409

Figures in bold are statistically significant at the 10% level. The 90% bootstrap confidence intervals are reported in brackets below the parameter estimates. \bar{R}^2 denotes the adjusted R^2 .

TABLE 2.—ESTIMATES AND BOOTSTRAP CONFIDENCE INTERVALS FOR U.S. INFLATION (OTHER MEASURES)

$$\Delta^h p_{t+h} = b + \sum_{k=0}^1 \beta_{0k} \Delta p_{t-k} + \sum_{k=1}^2 \beta_{1k} pccy_t^{(k)} + \sum_{k=1}^2 \beta_{2k} pcq_t^{(k)} + \sum_{k=0}^1 \beta_{3k} qoil_{t-k} + v_{t+h}$$

	Less Food and Energy	Less Food	Less Energy	Food	Energy
h = 1					
$pccy_t^{(1)}$	-0.079 [-0.196, -0.034]	0.011 [-0.083, 0.129]	-0.070 [-0.183, -0.025]	-0.041 [-0.182, 0.095]	1.141 [0.861, 2.224]
$pccy_t^{(2)}$	-0.053 [-0.161, 0.010]	-0.087 [-0.197, -0.021]	-0.058 [-0.157, -0.001]	-0.101 [-0.236, -0.027]	0.021 [-0.761, 0.790]
$pcq_t^{(1)}$	0.022 [-0.008, 0.054]	-0.037 [-0.112, 0.008]	0.019 [-0.017, 0.050]	-0.011 [-0.151, 0.091]	-0.646 [-1.241, -0.331]
$pcq_t^{(2)}$	0.004 [-0.031, 0.035]	0.025 [-0.037, 0.093]	0.022 [-0.007, 0.062]	0.149 [0.094, 0.270]	0.228 [-0.586, 1.014]
$qoil_t$	-0.068 [-0.173, 0.067]	1.451 [1.106, 1.939]	-0.065 [-0.170, 0.081]	-0.016 [-0.523, 0.472]	14.750 [11.956, 18.417]
$qoil_{t-1}$	0.080 [-0.047, 0.199]	-1.282 [-1.638, -0.951]	0.040 [-0.079, 0.147]	-0.196 [-0.524, 0.173]	-12.941 [-15.972, -10.024]
Δp_t	0.211 [0.137, 0.332]	0.139 [0.069, 0.273]	0.258 [0.181, 0.380]	0.127 [0.030, 0.260]	0.156 [0.057, 0.294]
Δp_{t-1}	0.222 [0.118, 0.346]	-0.099 [-0.301, 0.040]	0.145 [0.043, 0.277]	-0.063 [-0.145, 0.083]	-0.198 [-0.364, -0.054]
\bar{R}^2	0.300	0.273	0.289	0.071	0.366
h = 3					
$pccy_t^{(1)}$	-0.213 [-0.545, -0.097]	-0.034 [-0.482, 0.369]	-0.213 [-0.544, -0.094]	-0.176 [-0.660, 0.149]	2.978 [1.659, 6.783]
$pccy_t^{(2)}$	-0.155 [-0.470, 0.016]	-0.259 [-0.748, 0.009]	-0.181 [-0.504, -0.034]	-0.347 [-0.753, -0.186]	0.144 [-2.780, 3.133]
$pcq_t^{(1)}$	0.064 [-0.021, 0.169]	-0.066 [-0.321, 0.148]	0.052 [-0.043, 0.163]	-0.029 [-0.366, 0.286]	-1.663 [-3.528, -0.608]
$pcq_t^{(2)}$	0.022 [-0.077, 0.128]	0.149 [-0.023, 0.364]	0.065 [-0.044, 0.198]	0.398 [0.200, 0.798]	1.169 [-0.667, 3.205]
$qoil_t$	-0.036 [-0.423, 0.327]	2.160 [1.202, 3.873]	-0.105 [-0.457, 0.239]	-0.262 [-1.059, 0.505]	23.948 [18.772, 32.798]
$qoil_{t-1}$	0.036 [-0.302, 0.389]	-2.409 [-3.740, -1.652]	-0.009 [-0.295, 0.277]	-0.574 [-1.166, 0.095]	-25.608 [-33.793, -18.417]
Δp_t	0.692 [0.495, 0.944]	-0.145 [-0.336, 0.113]	0.612 [0.417, 0.875]	0.030 [-0.152, 0.242]	-0.287 [-0.414, -0.137]
Δp_{t-1}	0.734 [0.514, 0.982]	0.122 [-0.227, 0.463]	0.599 [0.390, 0.837]	-0.108 [-0.321, 0.100]	0.030 [-0.286, 0.317]
\bar{R}^2	0.526	0.143	0.481	0.153	0.206

Numbers in bold are statistically significant at the 10% level. The 90% bootstrap confidence intervals are in brackets below the parameter estimates. \bar{R}^2 denotes the adjusted R^2 .

TABLE 3.—ESTIMATES AND BOOTSTRAP CONFIDENCE INTERVALS FOR U.S. INFLATION (CPI ALL ITEMS)

$$\Delta^h p_{t+h} = b + \sum_{k=0}^1 \beta_{0k} \Delta p_{t-k} + \sum_{k=1}^2 \beta_{1k} pccy_t^{(k)} + \sum_{k=1}^6 \beta_{2k} \overline{cy}_t^{(k)} + \beta_3 \Delta s_t^{IMF} + v_{t+h}$$

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
$pccy_t^{(1)}$	-0.210 [-0.509, -0.165]	-0.622 [-1.810, -0.412]	-0.891 [-3.598, -0.244]	-1.314 [-5.366, -0.050]
$pccy_t^{(2)}$	-0.089 [-0.211, -0.011]	-0.362 [-0.908, -0.091]	-0.763 [-2.210, -0.168]	-1.412 [-3.092, -0.691]
$\overline{cy}_t^{(1)}$ (foodstuffs)	-0.000 [-0.025, 0.008]	0.006 [-0.051, 0.034]	-0.009 [-0.115, 0.030]	-0.002 [-0.174, 0.080]
$\overline{cy}_t^{(2)}$ (grains)	-0.033 [-0.062, -0.018]	-0.085 [-0.176, -0.036]	-0.083 [-0.233, 0.014]	-0.165 [-0.365, -0.008]
$\overline{cy}_t^{(3)}$ (industrials)	-0.009 [-0.022, -0.006]	-0.019 [-0.055, -0.006]	-0.025 [-0.090, 0.003]	-0.010 [-0.112, 0.024]
$\overline{cy}_t^{(4)}$ (meats)	0.002 [-0.008, 0.008]	-0.009 [-0.042, 0.016]	-0.013 [-0.069, 0.019]	0.009 [-0.071, 0.066]
$\overline{cy}_t^{(5)}$ (metals)	0.006 [-0.021, 0.025]	0.039 [-0.059, 0.097]	0.123 [-0.111, 0.233]	0.170 [-0.163, 0.313]
$\overline{cy}_t^{(6)}$ (energy)	-0.010 [-0.032, -0.002]	-0.049 [-0.126, -0.027]	-0.094 [-0.247, -0.050]	-0.135 [-0.350, -0.072]
Δs_t^{IMF}	0.008 [-0.002, 0.018]	0.017 [-0.005, 0.042]	0.027 [-0.006, 0.081]	0.025 [-0.042, 0.090]
Δp_t	0.351 [0.276, 0.562]	0.105 [-0.077, 0.354]	0.172 [-0.041, 0.502]	0.451 [0.143, 0.921]
Δp_{t-1}	-0.239 [-0.392, -0.101]	-0.237 [-0.495, 0.005]	-0.027 [-0.301, 0.334]	-0.215 [-0.687, 0.409]
\bar{R}^2 (all variables)	0.170	0.111	0.174	0.253
\bar{R}^2 (Δs_t^{IMF} , Δp_t , Δp_{t-1})	0.144	0.001	0.022	0.042

$\overline{cy}_t^{(k)}$ is a simple average of $cy_{j,t,n}$, $N_k^{-1} \sum_{j=1}^{N_k} cy_{j,t,n}$, for commodity group k . Numbers in bold are statistically significant at the 10% level. The 90% bootstrap confidence intervals are reported in brackets below the parameter estimates. \bar{R}^2 denotes the adjusted R^2 . The last line reports the \bar{R}^2 from a regression of $\Delta^h p_{t+h}$ on Δs_t^{IMF} , Δp_t , and Δp_{t-1} only.

inference, rendering the coefficients less significant or even insignificant. In spite of these stringent hurdles, the estimates on $pccy_t^{(2)}$ are still found to be statistically significant.

Given the popularity of the Phillips curve as a forecasting equation for inflation, we include in our augmented model of inflation the variable $u_t - u^*$ proxied by the deviations of the U.S. unemployment rate (produced by the BLS) from its HP trend. The estimated coefficients on this variable are insignificant at all forecast horizons. We note that $u_t - u^*$ is significant in the simple version of the Phillips curve that excludes $pccy_t$. This suggests that $pccy_t$ captures information about the economy that is typically reflected in the output gap.

Our analysis has so far focused on aggregate measures of convenience yields constructed as principal components of $cy_{j,t,n}$ of all commodities. To investigate if alternative ways of aggregating information are as effective, we compute simple averages of convenience yields for each of the six commodity groups: food, grains, industrials, meats, metals, and energy. These $\overline{cy}_t^{(j)}$ for $j = 1, \dots, 6$ are added as predictors in the inflation model that contains two lags of inflation, change in the IMF commodity index Δs_t^{IMF} , $pccy_t^{(1)}$, and $pccy_t^{(2)}$. The results in table 3 show that the average convenience yields for grains and especially for energy exhibit incremental predictive ability but did not render $pccy_t$ irrelevant. On the other hand, the changes in aggregate commodity prices are insignificant at all horizons, in line with the weak forecasting power in commodity prices for inflation found in the literature.

We also consider a regression of inflation on only Δs_t^{IMF} and two lags of inflation. The last row of table 3 shows that the adjusted R^2 for this regression is low for $h \geq 3$, with most of the predictive power coming from lagged inflation. However, replacing Δs_t^{IMF} by $pccy_t$ and \overline{cy}_t increases the adjusted R^2 substantially, reinforcing our thesis that it is the aggregate convenience yield component of commodity prices that has strong predictive power for inflation.

B. Out-of-Sample Forecast Performance

The results indicate that the principal components in convenience yields are statistically significant in-sample predictors. We now evaluate the pseudo out-of-sample predictive power of convenience yields for U.S. inflation.

We consider a recursive out-of-sample forecasting exercise by estimating the model using T_1 observations ($T_1 = T_0, T_0 + 1, \dots, T - h$) and producing h -period-ahead forecasts for $h = 1, 3, 6$, and 12. The principal components are computed with information only up to time T_1 . We use a one-sided two-year moving average filter with exponentially decreasing weights $\phi(1 - \phi)^i$ for $\phi = 0.15$ and $i = 1, \dots, 24$ to detrend real commodity prices. As a result, the initial sample uses the first 24 observations (March 1983–February 1985) for the one-sided moving average filter, while data from March 1985 to December 1997 ($T_0 = 154$ observations) are used for estimating the parameters and principal components. The pseudo out-of-sample forecast exercise starts in January 1998 and continues through July 2008.

Four forecasting models are considered. The first three models of $p_{t+h} - p_t$ have predictors: (a) $pccy_t^{(1)}$, $pccy_t^{(2)}$, Δp_t and their lags (denoted by CY), (b) $pccy_t^{(1)}$, $pccy_t^{(2)}$, $q_{oil,t}$, Δp_t and their lags (denoted by OIL), and (c) Δp_t and its lags (denoted by AR). The fourth model is an IMA(1,1) model of Δp_{t+1} whose h -period forecasts are obtained by aggregating the one-period forecast over h periods. The lags for the first three models are selected by the AIC over each recursive sample with the maximum lag set to 4 to avoid over-parameterization. The AIC tends to select lags that are more appropriate for forecasting. Stock and Watson (2007) present evidence that an IMA(1,1) model often outperforms AR and backward-looking Phillips curve models in out-of-sample predictions.

Table 4 presents the root-mean-squared forecast errors (RMSFEs) relative to that of the AR model. Numbers smaller than 1 indicate that the corresponding model outperforms the AR benchmark. Results are reported for forecast horizons of up to one year for two measures of inflation (CPI all items and CPI less food and energy). For the CPI (all items) inflation rate, model CY, which incorporates $pccy_t$, evidently dominates in terms of RMSFE. The gains generated by model CY increase with the forecast horizon and are as large as 9% compared to the best-performing competitor at the one-year horizon. For the core inflation, model CY continues to provide relatively large forecast improvements over the AR and

TABLE 4.—RECURSIVE OUT-OF-SAMPLE FORECASTS FOR U.S. INFLATION RATE (RELATIVE RMSFEs)

	All Items			Less Food and Energy		
	CY	OIL	IMA	CY	OIL	IMA
$h = 1$	1.004	0.951	1.049	1.037	1.076	0.935
$h = 3$	0.967	0.975	1.007	0.964	1.048	0.920
$h = 6$	0.939	1.077	1.037	0.911	1.019	0.916
$h = 12$	0.909	1.125	1.148	0.878	1.019	0.900

The period for pseudo-out-of-sample forecast evaluation starts in January 1998 and continues through July 2008. CY is a model of $p_{t+h} - p_t$ on $pccy_t^{(1)}$, $pccy_t^{(2)}$, Δp_t and their lags; OIL is a model of $p_{t+h} - p_t$ on $pccy_t^{(1)}$, $pccy_t^{(2)}$, $qoil_t$, Δp_t and their lags; the third model is an AR model of $p_{t+h} - p_t$ on Δp_t and its lags; and IMA is an IMA(1,1) model of inflation Δp_{t+1} whose h -period forecast is obtained by aggregating the one-step forecast over h periods. The lags for the first three models are selected using AIC. The reported RMSFEs are relative to the RMSFE of the AR model.

OIL models. However, it is dominated at short horizons by the IMA(1,1) model.

The good performance of the IMA(1,1) model for forecasting core inflation can be explained by the time series properties of this inflation series. Specifically, the estimated AR and MA parameters by fitting an ARMA(1,1) model to core inflation over the whole sample are 0.990 and -0.905 , with standard errors of 0.007 and 0.027, respectively. In contrast, the corresponding AR and MA estimates from an ARMA(1,1) model for CPI (all items) inflation are -0.041 and 0.458, with standard errors 0.132 and 0.117, respectively. As in Stock and Watson (2007), we also find that the IMA(1,1) is not an appropriate model for CPI (all items). It is precisely in this case that model CY dominates the IMA(1,1) model in out-of-sample forecasts.

Together with results reported earlier, the predictive power of $pccy_t$ seems robust both in- and out-of-sample. Figure 1 plots the actual and forecast values from model CY of twelve-month-ahead (all items and core) inflation. When $pccy_t$ is used as the predictor, the forecast tracks the local trends and turning points in the actual inflation dynamics closely.

C. G7 Inflation Rates

In order to see if the predictability of inflation by the convenience yields and real commodity prices also holds for other countries, we use data on CPI (all items), provided by the OECD, to construct inflation rates for the rest of the G7 countries: Canada, Japan, Germany, France, Italy, and the United Kingdom. Bilateral exchange rates are used to convert the commodity prices into local currencies. Interest and exchange rate data for these countries are also obtained from the OECD database.

For each of these countries, we convert all commodity price variables in domestic currency using the market exchange rate. Each country's interest rate is used in equation (9) to construct the country-specific convenience yield. Similarly, the real commodity prices are converted in local currency and deflated by the corresponding CPI index. For each country, we estimate a model that includes two principal components of convenience yields and detrended real commodity prices, as well as two lags of the detrended real oil price and the dependent variable. Note that the conversion of convenience

yields into local currency may change the ordering and interpretation of the principal components computed above with U.S. data.

The estimation results for each country at horizons $h = 1$ and 12 are presented in table 5. At the one-month horizon, $pccy_t^{(1)}$ and $pccy_t^{(2)}$, as well as the real oil price, appear to be statistically significant for almost all countries. For example, $pccy_t^{(1)}$ is significant at the 10% level for all countries even after accounting for estimation uncertainty using bootstrap confidence intervals. The predictive power of $pccy_t^{(1)}$ remains significant for all countries at $h = 12$, but the effects of the other determinants (real commodity and oil prices) are substantially diminished. The in-sample predictive power of $pccy_t^{(1)}$ is robust across countries in spite of the volatile dynamics of some of these monthly inflation series.

To see if the predictive ability of the aggregate convenience yields continues to hold out-of-sample for the G7 countries, table 6 presents relative RMSFEs from the same four forecasting models considered in table 4. The CY model consistently dominates the competing models for Canada, Japan, France, and the United Kingdom with forecasting gains of more than 19% for Japan and the United Kingdom at longer horizons. On the other hand, the IMA(1,1) model outperforms the CY model for Germany and Italy. We attribute this finding to the negative MA component in the inflation data for Germany and Italy. For example, when an ARMA(1,1) model is fitted to Italy's inflation, the AR and MA estimates are 0.973 and -0.850 with standard errors of 0.010 and 0.036, respectively.⁷ These results reinforce our earlier observation that the CY model yields systematic forecast improvements when the negative moving average component in the inflation series is not strong.


VI. Commodity Prices and Convenience Yields

Convenience yields appear in our inflation model because they are related to commodity prices. This section provides evidence in support of this link. The dynamics of commodity price movements are of interest in their own right for a variety of reasons. For developing countries that depend heavily on exports of commodities for revenue, fluctuations in commodity prices are the main cause of income volatility. In other countries, commodity price movements are often tied to real exchange rate appreciation and depreciation, depending on whether the country is an exporter or an importer of commodities. For small open economies like Canada, New Zealand, and Australia, fluctuations in commodity prices are sufficiently important for the design of economic policies that these central banks produce their own commodity price indices to appropriately reflect the commodities they produce.

Identifying robust predictors of commodity prices turns out to be a challenging problem since these prices tend to

⁷Ng and Perron (2001) document that a large autoregressive root and a large negative moving average component are a characteristic of the inflation series (GDP deflators) of many G7 countries.

FIGURE 1.—ACTUAL AND FORECAST VALUES OF TWELVE-MONTH INFLATION: ALL ITEMS (TOP) LESS FOOD AND ENERGY (BOTTOM)



The forecast values are obtained recursively from model CY with $pccy_t^{(1)}$, $pccy_t^{(2)}$, Δp_t , and their lags as predictors.

TABLE 5.—ESTIMATES AND BOOTSTRAP CONFIDENCE INTERVALS FOR G7 INFLATION RATES

$$\Delta^h p_{t+h} = b + \sum_{k=0}^1 \beta_{0k} \Delta p_{t-k} + \sum_{k=1}^2 \beta_{1k} pccy_t^{(k)} + \sum_{k=1}^2 \beta_{2k} pcq_t^{(k)} + \sum_{k=0}^1 \beta_{3k} q_{oil,t-k} + v_{t+h}$$

	Canada	Japan	France	Germany	Italy	United Kingdom
<i>h</i> = 1						
<i>pccy_t</i> ⁽¹⁾	-0.136 [-0.272, -0.001]	-0.157 [-0.255, -0.058]	-0.143 [-0.229, -0.056]	0.123 [0.019, 0.318]	-0.116 [-0.157, -0.042]	-0.234 [-0.391, -0.086]
<i>pccy_t</i> ⁽²⁾	-0.057 [-0.307, 0.155]	0.001 [-0.201, 0.208]	0.064 [-0.114, 0.373]	-0.000 [-0.220, 0.200]	0.052 [-0.083, 0.221]	0.091 [-0.188, 0.438]
<i>pcq_t</i> ⁽¹⁾	-0.078 [-0.212, -0.005]	-0.109 [-0.175, -0.047]	-0.058 [-0.099, -0.017]	-0.078 [-0.135, -0.009]	-0.013 [-0.046, 0.034]	-0.158 [-0.250, -0.076]
<i>pcq_t</i> ⁽²⁾	0.117 [0.012, 0.247]	-0.027 [-0.124, 0.112]	-0.046 [-0.146, 0.049]	0.007 [-0.099, 0.162]	-0.019 [-0.097, 0.066]	-0.090 [-0.061, 0.228]
<i>q_{oil,t}</i>	0.855 [0.360, 1.395]	0.564 [0.081, 1.153]	0.518 [0.132, 0.836]	0.465 [0.132, 0.872]	0.270 [0.076, 0.509]	0.626 [0.076, 1.200]
<i>q_{oil,t-1}</i>	-0.687 [-1.267, 0.059]	-0.344 [-0.779, 0.062]	-0.344 [-1.048, -0.430]	-0.526 [-0.839, -0.187]	-0.249 [-0.444, -0.079]	-0.968 [-1.548, -0.400]
Δp_t	0.004 [-0.110, 0.152]	0.091 [-0.030, 0.227]	0.187 [0.120, 0.346]	-0.052 [-0.162, 0.097]	0.316 [0.248, 0.461]	0.065 [0.031, 0.179]
Δp_{t-1}	0.053 [-0.025, 0.168]	-0.308 [-0.382, -0.191]	-0.061 [-0.162, 0.109]	0.054 [-0.016, 0.173]	0.194 [0.104, 0.358]	-0.054 [-0.115, 0.026]
\bar{R}^2	0.084	0.143	0.188	0.031	0.410	0.101
<i>h</i> = 12						
<i>pccy_t</i> ⁽¹⁾	-1.851 [-2.459, -1.142]	-1.725 [-2.644, -0.925]	-1.013 [-1.571, -0.138]	1.176 [0.326, 3.125]	-1.470 [-2.027, -0.585]	-2.602 [-3.812, -1.332]
<i>pccy_t</i> ⁽²⁾	-0.232 [-1.317, 0.920]	-0.253 [-2.374, 1.458]	0.596 [-0.510, 3.045]	0.817 [-0.685, 3.496]	-0.597 [-3.287, 1.062]	0.794 [-2.028, 5.101]
<i>pcq_t</i> ⁽¹⁾	-0.167 [-0.967, 0.306]	-0.570 [-1.119, -0.241]	-0.058 [-0.363, 0.330]	-0.037 [-0.478, 0.521]	-0.392 [-0.901, -0.062]	-0.560 [-1.126, 0.210]
<i>pcq_t</i> ⁽²⁾	0.197 [-0.532, 0.827]	0.089 [-0.627, 0.954]	-0.543 [-1.453, 0.095]	-0.105 [-0.844, 0.685]	-0.390 [-1.194, 0.702]	-0.880 [-1.894, 0.059]
<i>q_{oil,t}</i>	1.198 [-0.635, 3.379]	1.450 [-0.195, 4.247]	-2.219 [-4.207, -0.567]	0.632 [-1.594, 3.717]	-0.439 [-2.331, 1.610]	-1.877 [-4.035, 0.396]
<i>q_{oil,t-1}</i>	-2.365 [-4.422, -0.967]	-1.333 [-4.214, 0.040]	-0.169 [-1.590, 1.185]	-0.665 [-3.442, 1.570]	-1.773 [-3.395, -0.694]	0.197 [-2.203, 2.153]
Δp_t	0.242 [-0.265, 0.793]	0.084 [0.061, 0.391]	1.408 [0.823, 2.347]	0.633 [0.267, 1.118]	2.699 [2.382, 3.546]	0.598 [0.304, 1.126]
Δp_{t-1}	0.252 [-0.195, 0.753]	0.218 [0.095, 0.516]	1.572 [0.967, 3.071]	0.607 [0.188, 1.103]	2.640 [2.124, 3.792]	0.557 [0.242, 1.213]
\bar{R}^2	0.326	0.374	0.422	0.234	0.688	0.450

Numbers in bold are statistically significant at the 10% level. The 90% bootstrap confidence intervals are reported in brackets below the parameter estimates. \bar{R}^2 denotes the adjusted R^2 .

TABLE 6.—RECURSIVE OUT-OF-SAMPLE FORECASTS FOR G7 INFLATION RATES (RELATIVE RMSFEs)

	Canada			Japan			France		
	CY	OIL	IMA	CY	OIL	IMA	CY	OIL	IMA
<i>h</i> = 1	0.996	1.003	1.007	0.884	1.065	1.015	0.966	0.950	1.027
<i>h</i> = 3	0.958	1.049	1.009	0.850	1.110	1.048	0.923	1.010	1.050
<i>h</i> = 6	0.973	1.036	1.013	0.805	1.155	1.049	0.862	0.992	1.030
<i>h</i> = 12	0.972	1.004	0.989	0.813	1.289	1.029	0.871	1.070	1.033
	Germany			Italy			United Kingdom		
	CY	OIL	IMA	CY	OIL	IMA	CY	OIL	IMA
<i>h</i> = 1	0.979	0.994	0.967	1.013	1.109	0.933	0.849	1.001	1.006
<i>h</i> = 3	0.945	1.000	0.961	0.977	1.124	0.853	0.865	1.056	1.120
<i>h</i> = 6	1.114	1.058	0.938	0.954	1.142	0.744	0.811	1.099	1.135
<i>h</i> = 12	1.217	1.199	0.876	1.017	1.089	0.683	0.879	1.168	1.290

The period for pseudo-out-of-sample forecast evaluation starts in January 1998 and continues through July 2008. CY is a model of $p_{t+h} - p_t$ on $pccy_t^{(1)}$, $pccy_t^{(2)}$, Δp_t , and their lags; OIL is a model of $p_{t+h} - p_t$ on $pcq_t^{(1)}$, $pcq_t^{(2)}$, $q_{oil,t}$, Δp_t , and their lags; the third model is an AR model of $p_{t+h} - p_t$ on Δp_t and its lags; and IMA is an IMA(1,1) model of inflation Δp_{t+1} whose *h*-period forecast is obtained by aggregating the one-step forecast over *h* periods. The lags for the first three models are selected using AIC. The reported RMSFEs are relative to the RMSFE of the AR model.

be extremely volatile and have proven difficult to forecast (Chen et al. 2010). The ability of convenience yields to predict inflation raises the question of whether the convenience yields can also forecast commodity prices. We begin with predictive regressions for the individual commodity price changes:

$$\Delta s_{jt+1} = a + \alpha_{j0}(L) \Delta s_{jt} + \alpha_{j1}(L) cy_{jt,n} + \alpha_{j2}(L) i_t + \alpha_{j3}(L) \Delta \bar{x}_t + e_{jt+1}$$

For most commodities, the AIC selects $\alpha_{j1}(L) = \alpha_{j1}$, $\alpha_{j2}(L) = \alpha_{j2}$, $\alpha_{j2}(L) = \alpha_{j3}$, and $\alpha_{j0}(L)$ is a second-order polynomial in the lag operator. We adopt the same lag order for all commodities and analyze the predictive ability of $cy_{jt,n}$, i_t , and $\Delta \bar{x}_t$, conditional on the presence of the remaining two determinants. Table 7 shows that the convenience yields have highly statistically significant predictive power for most commodity prices, but the interest and exchange rates do not. The

TABLE 7.—ESTIMATES AND STANDARD ERRORS

$$\Delta s_{jt+1} = a + \sum_{k=0}^1 \alpha_{j0k} \Delta s_{jt-k} + \alpha_{j1} c y_{jt,n} + \alpha_{j2} i_t + \alpha_{j3} \Delta \bar{x}_t + e_{jt+1}$$

<i>j</i>	$\hat{\alpha}_{j1}$	$\hat{\alpha}_{j2}$	$\hat{\alpha}_{j3}$	\bar{R}^2
Foodstuffs				
Cocoa	-0.297 (0.169)	-0.084 (0.181)	-0.143 (0.222)	0.033
Coffee	-0.117 (0.143)	-0.194 (0.190)	-0.043 (0.347)	0.035
Orange juice	-0.445 (0.146)	0.272 (0.189)	-0.132 (0.174)	0.098
Sugar	-0.495 (0.109)	-0.476 (0.233)	-0.243 (0.295)	0.125
Grains				
Canola	-0.533 (0.233)	-0.003 (0.164)	0.266 (0.158)	0.169
Corn	-0.634 (0.118)	0.216 (0.144)	0.063 (0.158)	0.210
Oats	-0.339 (0.098)	-0.359 (0.215)	0.226 (0.172)	0.060
Soybeans	-0.958 (0.249)	-0.260 (0.129)	0.166 (0.138)	0.187
Soybean oil	-0.433 (0.235)	-0.111 (0.194)	0.200 (0.188)	0.138
Wheat	-0.393 (0.092)	0.071 (0.140)	0.023 (0.189)	0.067
Industrials				
Cotton	-0.642 (0.111)	0.362 (0.158)	-0.071 (0.201)	0.289
Lumber	-0.422 (0.103)	-0.273 (0.187)	-0.021 (0.258)	0.090
Meats				
Cattle, feeder	-0.506 (0.124)	-0.026 (0.085)	0.081 (0.108)	0.154
Cattle, live	-0.484 (0.055)	0.061 (0.104)	0.027 (0.119)	0.242
Hogs	-0.530 (0.041)	0.240 (0.133)	0.056 (0.196)	0.326
Pork bellies	-0.914 (0.106)	-0.046 (0.204)	0.061 (0.244)	0.226
Metals				
Copper	-0.100 (0.158)	-0.250 (0.127)	0.017 (0.176)	0.054
Gold	0.867 (0.869)	-0.284 (0.069)	-0.118 (0.127)	0.040
Palladium	-0.361 (0.176)	0.062 (0.243)	0.422 (0.248)	0.030
Platinum	0.281 (0.237)	-0.252 (0.118)	-0.357 (0.152)	0.031
Silver	0.186 (0.916)	-0.593 (0.152)	0.167 (0.185)	0.052
Energy				
Crude oil	-0.512 (0.303)	-0.144 (0.216)	0.043 (0.277)	0.076
Heating oil	-0.788 (0.171)	-0.023 (0.184)	0.023 (0.244)	0.123

Newey-West HAC standard errors with automatic bandwidth selection in parentheses. Numbers in bold are statistically significant at the 10% level. \bar{R}^2 denotes the adjusted R^2 .

TABLE 8.—ESTIMATES AND BOOTSTRAP CONFIDENCE INTERVALS

$$\Delta^h s_{t+h}^{IMF} = a + \sum_{k=0}^1 \alpha_{10} \Delta s_{t-k}^{IMF} + \alpha_{11} pccy_t^{(1)} + \alpha_{12} pccy_t^{(2)} + \alpha_2 i_t + \alpha_3 \Delta \bar{x}_t + e_{t+h}$$

	<i>h</i> = 1	<i>h</i> = 3	<i>h</i> = 6	<i>h</i> = 12
$pccy_t^{(1)}$	0.874 [0.079, 2.332]	3.218 [0.748, 8.184]	6.318 [1.372, 17.225]	5.795 [-3.127, 23.340]
$pccy_t^{(2)}$	-1.259 [-3.035, -0.761]	-3.155 [-9.664, -0.679]	-3.660 [-15.855, 3.103]	-3.946 [-16.917, 5.351]
i_t	-0.069 [-0.187, 0.040]	-0.237 [-0.573, 0.112]	-0.568 [-1.316, 0.259]	-1.770 [-2.699, -0.424]
$\Delta \bar{x}_t$	0.017 [-0.150, 0.177]	-0.181 [-0.451, 0.070]	-0.435 [-0.924, -0.019]	-0.585 [-1.375, -0.039]
Δs_{t-1}^{IMF}	0.205 [0.088, 0.323]	0.270 [0.046, 0.492]	0.163 [-0.197, 0.593]	0.192 [-0.327, 0.699]
Δs_{t-1}^{IMF}	0.034 [-0.032, 0.092]	-0.012 [-0.200, 0.190]	-0.047 [-0.361, 0.268]	0.119 [-0.333, 0.707]
\bar{R}^2	0.147	0.219	0.259	0.304

Numbers in bold are statistically significant at the 10% level. The 90% bootstrap confidence intervals are reported in brackets below the parameter estimates. \bar{R}^2 denotes the adjusted R^2 .

exceptions are the commodities in the metals group (precious metals).

The *h*-period predictive regression for IMF commodity price index is given by

$$\Delta^h s_{t+h}^{IMF} = a + \alpha_0(L) \Delta s_t^{IMF} + \alpha'_1 pccy_t + \alpha_2 i_t + \alpha_3 \Delta \bar{x}_t + e_{t+h}. \tag{11}$$

Table 8 presents results for *h* = 1, 3, 6, 12. The variable $pccy_t^{(1)}$ is statistically significant (at the 10% level) for all

TABLE 9.—RECURSIVE OUT-OF-SAMPLE FORECASTS OF COMMODITY PRICE CHANGES (RELATIVE RMSFES)

	CY	ER/IR	RW
<i>h</i> = 1	0.955	1.013	1.180
<i>h</i> = 3	0.915	0.986	1.207
<i>h</i> = 6	0.893	0.946	1.266
<i>h</i> = 12	0.903	0.850	1.035

The period for pseudo out-of-sample forecast evaluation starts in January 1998 and continues through July 2008. CY is a model of $s_{t+h}^{IMF} - s_t^{IMF}$ on $pccy_t^{(1)}$, $pccy_t^{(2)}$, Δs_t^{IMF} and their lags; ER/IR is a model of $s_{t+h}^{IMF} - s_t^{IMF}$ on i_t , $\Delta \bar{x}_t$, Δs_t^{IMF} and their lags; the third model is an AR model of $s_{t+h}^{IMF} - s_t^{IMF}$ on Δs_t^{IMF} and its lags; and RW is a no-change model, that is, $E(s_{t+h}^{IMF} - s_t^{IMF}) = s_t^{IMF} - s_{t-h}^{IMF}$. The lags for the first three models are selected using AIC. The reported RMSFEs are relative to the RMSFE of the AR model.

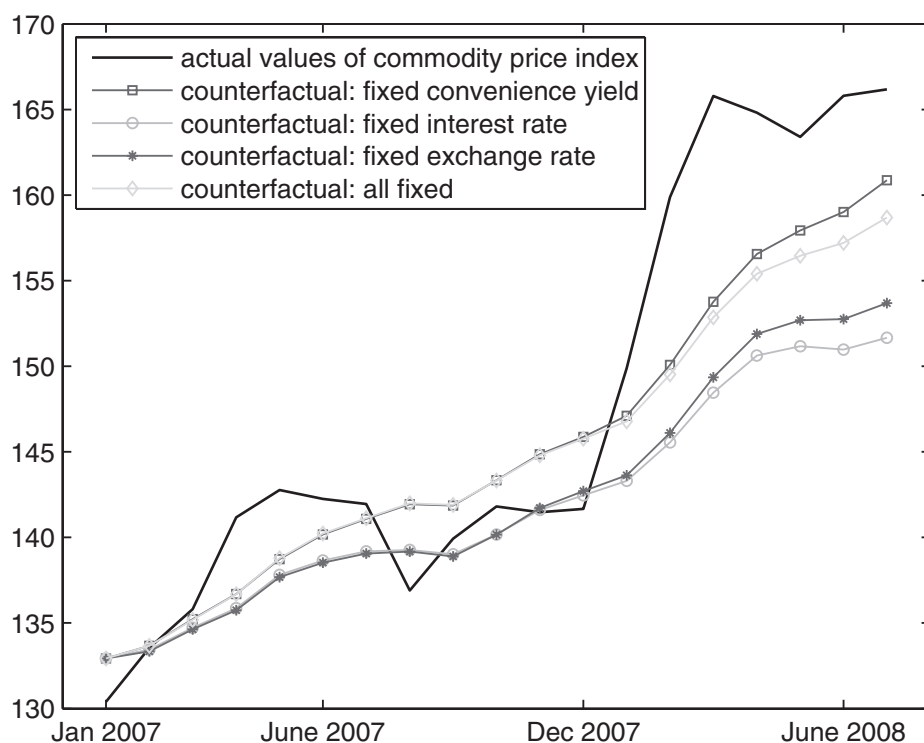
horizons except for *h* = 12. The second principal component is significant at one- and three-month horizons but is insignificant at longer horizons. Chen et al. (2010) suggest that foreign exchange values of commodity currencies can help predict the prices of commodities they export. While the change of the U.S. trade-weighted exchange index is not significant at short horizons (one and three months), it becomes significant at six- and twelve-month horizons. A similar pattern is observed for the interest rate so that the exchange and interest rates are the only statistically significant predictors at the one-year horizon.

To assess the out-of-sample predictive power of $pccy_t$ for Δs_{t+h}^{IMF} , we consider four models: CY is a model of $s_{t+h}^{IMF} - s_t^{IMF}$ on $pccy_t^{(1)}$, $pccy_t^{(2)}$, Δs_t^{IMF} and their lags; ER/IR is a model of $s_{t+h}^{IMF} - s_t^{IMF}$ on i_t , $\Delta \bar{x}_t$, Δs_t^{IMF} and their lags; the third model is an AR model of $s_{t+h}^{IMF} - s_t^{IMF}$ on Δs_t^{IMF} and its lags; and the last model (RW) is a no-change model, $E(s_{t+h}^{IMF} - s_t^{IMF}) = s_t^{IMF} - s_{t-h}^{IMF}$. The lags for the first three models are selected by AIC with the maximum lag set to 4. Table 9 reports relative RMSFEs with the AR model being the benchmark model. The results show that the largest forecasting gains of the CY model occur for one- to six-month forecast horizons, likely because the convenience yields are constructed using futures prices with one to five months to maturity. The interest and exchange rates also possess some out-of-sample forecast ability, especially at six- and twelve-month horizons, in line with our in-sample results.

Many explanations have been advanced for the surge in commodity prices between 2006 and 2007. Specifically, the IMF commodity price index rose by 30% between 2006 and 2007 and by another 10% in 2007. Given that $pccy_t$ predicts commodity prices, it is of interest to ask whether commodity price increases can be explained by the aggregate convenience yields. To this end, we use the baseline model of Δs_{t+1}^{IMF} (see table 8) to ask what would have been the level of the IMF commodity price index if the convenience yields from January 2007 onward were held at the level of December 2006 but with i_t and $\Delta \bar{x}_t$ at their actual levels. This yields the counterfactual commodity price denoted by $\hat{s}_{t+1|pccy}^{IMF}$. Similarly, i_t and $\Delta \bar{x}_t$ are held at the level of December 2006 one at a time to yield $\hat{s}_{t+1|i}^{IMF}$ and $\hat{s}_{t+1|\Delta \bar{x}}^{IMF}$, respectively. A final exercise holds all three variables fixed at their December 2006 level, giving $\hat{s}_{t+1|pccy,i,\Delta \bar{x}}^{IMF}$.

Figure 2 shows that the hypothetical prices $\hat{s}_{t+1|pccy}^{IMF}$ and $\hat{s}_{t+1|pccy,i,\Delta \bar{x}}^{IMF}$ constructed by holding $pccy$ fixed, closely follow

FIGURE 2.—ACTUAL AND COUNTERFACTUAL VALUES OF THE IMF COMMODITY PRICE INDEX, JANUARY 2007–JULY 2008



The counterfactuals hold the two principal components of convenience yields, interest rate, changes of USD trade-weighted index, and all of these variables fixed at their December 2006 values, respectively.

the trend of the actual index, while $\hat{s}_{t+1|i}^{IMF}$ and $\hat{s}_{t+1|\Delta\bar{x}}^{IMF}$ exhibit much slower growth. This suggests that some of the increases in 2007 and 2008 appear to be driven by interest rate and exchange rate changes, as Hamilton (2008) and Frankel (2006, 2008) suggested. However, there is no evidence that expectations about future economic conditions, as reflected in the convenience yields, were behind the substantial commodity price increases in 2007 and 2008.⁸

VII. Conclusion

We find that the principal components of individual convenience yields incorporate information useful for predicting of both inflation and commodity prices. The fact that $pccy_t$ explains both inflation and the IMF commodity price index but that the latter has only weak predictive power for inflation underscores the point that commodity prices have multiple sources of variation and not every one has inflationary consequences.

In spite of the predictability of $pccy_t$ for inflation and commodity prices, a formal economic interpretation of why the coefficients are significant is beyond what our reduced-form analysis can offer because predictive regressions are not structural equations. Nonetheless, the convenience yields of cocoa, orange juice, and copper have a positive effect on

one-period-ahead inflation, while those of soybeans, oats, and silver have a negative effect; the coefficients on other convenience yields are not statistically significant. Accordingly, we form two new measures of aggregate convenience yields by averaging the convenience yields of (a) copper, orange juice, and copper and (b) soybeans, oats, and silver. We find that both variables have predictive of inflation for all forecast horizons and all inflation measures similar to the two principal components. The two principal components seem to be picking up information about inflation and commodity prices in these two groups of commodities.

REFERENCES

- Alquist, Ron, and Lutz Kilian, "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Econometrics* 25:4 (2010), 539–573.
- Bai, Jushan, and Serena Ng, "Confidence Intervals for Diffusion Index Forecasts and Inference for Factor Augmented Regressions," *Econometrica* 74:4 (2006), 1133–1150.
- Bernanke, Ben, "Outstanding Issues in the Analysis of Inflation," speech at the Federal Reserve Bank of Boston's 53rd Annual Economic Conference, Chatham, MA, June 9, 2008.
- Blomberg, Brock, and Ethan Harris, "The Commodity-Consumer Price Connection: Fact or Fable?" *FRBNY Economic Policy Review* (October 1995), 21–38.
- Boughton, James, and William Branson, "Commodity Prices as a Leading Indicator of Inflation" (pp. 305–338), in K. Lahiri and G. Moore (eds.), *Leading Economic Indicators* (Cambridge: Cambridge University Press, 1991).
- Chen, Yu-Chin, Kenneth Rogoff, and Barbara Rossi, "Can Exchange Rates Forecast Commodity Prices?" *Quarterly Journal of Economics* 125:3 (2010), 1145–1194.

⁸ A similar counterfactual analysis for U.S. inflation reveals that keeping the $pccy_t$ unchanged for the last nineteen months of the sample underpredicts the actual price level, but incorporating the rise in commodity prices over this period brings the predicted CPI closer to its actual level.

- Deaton, Angus, and Guy Laroque, "On the Behaviour of Commodity Prices," *Review of Economic Studies* 59:1 (1992), 1–23.
- Fama, Eugene, and Kenneth French, "Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage," *Journal of Business* 60:1 (1987), 55–73.
- "Business Cycles and the Behavior of Metals Prices," *Journal of Finance* 43:5 (1988), 1075–1093.
- Frankel, Jeffrey, "The Effect of Monetary Policy on Real Commodity Prices," (pp. 291–327), in J. Campbell (ed.), *Asset Prices and Monetary Policy* (Chicago: University of Chicago Press, 2006).
- "An Explanation for Soaring Commodity Prices" (2008), <http://www.voxeu.org/index.php?q=node/1002>.
- Furlong, Fred, and Robert Ingenito, "Commodity Prices and Inflation," *FRBSF Economic Review* 2 (1996), 27–47.
- Gorton, Gary, Fumio Hayashi, and K. Geert Rouwenhorst, "The Fundamentals of Commodity Futures Returns," *Review of Finance*, forthcoming.
- Hamilton, James, "Commodity Prices and the Fed," *Econbrowser* (March 2008).
- Hong, Harison, and Motohiro Yogo, "Digging into Commodities," Princeton University working paper (2009).
- Hooker, Mark, "Are Oil Shocks Inflationary? Asymmetric and Nonlinear Specifications versus Changes in Regime," *Journal of Money, Credit and Banking* 34:2 (2002), 540–561.
- Kaldor, Nicholas, "Speculation and Economic Stability," *Review of Economic Studies* 7:1 (1939), 1–27.
- Kilian, Lutz, "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review* 99:3 (2009), 1053–1069.
- Mishkin, Federick, "Can Futures Market Data Be Used to Understand the Behavior of Real Interest Rates?" *Journal of Finance* 45:1 (1990), 245–257.
- Ng, Serena, and Pierre Perron, "Lag Length Selection and the Construction of Unit Root Tests with Good Size and Power," *Econometrica* 69:6 (2001), 1519–1554.
- Pindyck, Robert, "The Present Value Model of Rational Commodity Pricing," *Economic Journal* 103:418 (1993), 511–530.
- Pindyck, Robert, and Julio Rotemberg, "The Excess Co-Movement of Commodity Prices," *Economic Journal* 100:403 (1990), 1173–1189.
- Stock, James, and Mark Watson, "Forecasting Output and Inflation: The Role of Asset Prices," *Journal of Economic Literature* 41:3 (2003), 788–829.
- "Why Has Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39:Suppl. s1 (2007), 3–33.
- Tang, Ke, and Wei Xiong, "Index Investing and Financialization of Commodities," Princeton University working paper (2009).