ESSAYS IN COMMODITIES AND FIRMS

By

Yuanzhi Gao

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

ECONOMICS

August, 2015

Nashville, Tennessee

Approved:

Professor Mario J. Crucini

Professor Kevin X.D. Huang

Professor Atsushi Inoue

Professor Joel Rodrigue

Professor Miguel Palacios

ACKNOWLEDGMENTS

I would like to thank my dissertation committee: Professor Mario J. Crucini (chair), Professor Kevin X.D. Huang, Professor Atsushi Inoue, Professor Joel Rodrigue and Professor Miguel Palacios for helpful comments and insightful suggestions. I owe great gratitude to my advisor, Mario Crucini and my co-author, Kevin Huang, for inspiration and guidance. This dissertation would not have existed without their constant input.

I would also like to express my particular appreciation to my wife, Mandy Xu and my parents, Bo Gao and Yurong Jiang, for their long-distance support over my entire Ph.D. tenure. Although I am 1,200 miles away from my Cambridge home and 8,000 miles away from my hometown in China, it is their encouragement that has motivated and pushed me forward.

I also benefit greatly from conversations with Jushan Bai, Christiane Baumeister, Bahattin Buyuksahin, Kan Chen, Steve Herscovici, Tong Li, Alejandro Molnar, Jacob Sagi, Tao Zha, Hongquan Zhu and seminar participants in Vanderbilt University, Bank of Canada, University of Minnesota, SEA 2014 Annual Meeting and Southwest Jiaotong University. Discussions with my flatmates and classmates, especially with Zhongzhong Hu, Kai Hong, Ruixuan Liu, Nam Vu and Lily Zhao are also very helpful.

Finally, I gratefully acknowledge the financial support from the Department of Economics at Vanderbilt University, the Kirk Dornbush family and the Walter Noel family.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Figure Page

CHAPTER I

Long-Run Risk, Uncertainty and the Spot Returns of Commodities

I.1 Introduction

Commodities are becoming increasingly important in both economics and finance. More countries and multi-national firms rely on globally traded commodities, and changes in commodity prices can bring significant variation to terms of trade (Backus and Crucini, 2000). On the other hand, the securitization of commodities since 2000 makes them more available for individual and institutional investors. According to the Investment Company Institute, the total assets invested in commodities have increased from one billion US dollars in 2004 to over one hundred billion in 2010 (Gargano and Timmermann 2014). Yet despite such urgent need from both economic policymakers and financial practitioners, commodities remain poorly understood in academic research. Evidence from data shows that commodity returns differ across commodities and from other, traditional asset classes (Tang and Xiong 2012). The question of the origin of the pricing anomaly remains unanswered. The purpose of this paper is to tackle this issue and provide an explanation, both theoretically and empirically.

A modified version of the Long-Run Risk is first presented in this study. In this model, the LR component as well as consumption corresponds only to steady state conditions of the economy based on the Permanent Income Hypothesis. The intuition behind such a modification is to try to align the model implications with those of studies based on macroeconomic DSGE models. In addition, the stochastic volatility can be neatly modeled as in Bloom (2009). The second purpose is to provide maximum testability for empirical analyses. Studies based on the original LRR setup usually focus on calibration instead of estimation exercises and calibrated results depend heavily on preference parameters. In this paper, I instead try to test theoretical propositions based on structural estimation of competitive equilibrium states.

In order to quantify the latent long-run component, a crucial connection between individual commodities and cross-sectional asset prices is demonstrated. Such a connection is strong and stable through a few unit-root processes shared by the constituents of the two markets. Those factors are first extracted from equity market data and are then fitted to commodities as the long-run components. By comparing the trends versus residual components as short-run "cycles," we are able to have a clear view of the magnitude of long-run risks in commodities. This finding is in sharp contrast to many previous studies that in general state that equity market correlation with commodities is close to zero, time-varying and industry-specific (Gorton and Rowenhorst 2006, Tang and Xiong 2012, Bhardwaj and Dunsby 2013). Those studies in general focus on the correlations of aggregated indexes. The contrast here indicates that aggregated indexes often carry very limited long-run information with divergent constituents, and very useful dynamics can be lost through the process of aggregation.

One important task based on this stable long-run relation is to find risk factors that can explain the risk premia of major exchange-traded commodity spot returns. Most of the attention on the returns of commodities have been given to commodity futures as they are more liquid. Spot prices deserve more attention, however, as they are more fundamental than futures. Changes in spot prices are not only likely to affect the costs of manufacturers, who rely on commodities as production inputs, but also all the financial derivatives based on them. Yet due to the intrinsic heterogeneity among commodities, previous attempts at finding equity-like factors for multiple commodities have largely been unsuccessful: Factor analyses based on CAPM as in Dusak (1973) and CCAPM in Breeden (1980) conclude that those models are incapable of modeling unique behaviors of commodities, especially with data at the monthly frequency (de Roon and Szymanowska 2010). My model indicates that, though commodities indeed behave in more heterogeneous ways than equities, there are still long-run risks shared with the equity markets and among themselves. Chiang, Hughen and Sagi (2014), in a parallel study, also find such a connection between equities and oil prices. In addition, agricultural commodities and resource commodities follow fairly distinct sets of stochastic trends in data. After separating the commodities into two groups, I am able to find risk factors that carry sizable risk premia for all commodities within groups. To the best of my knowledge, this is the first study that documents such factors for commodity spot prices.

Last but certainly not least, we are able to decompose both commodity risk premia and volatilities into three distinct components: long-run premia, short-run premia and uncertainty discount. This new method of decomposition is crucial to a better understanding of commodity price fluctuations: The long-run component measures the parts of premia and volatilities that fluctuate based on economic fundamentals in DSGE steady states. The short run on the other hand, captures the movements along impulse response functions. The uncertainty component is about the additional returns and volatilities brought by by the stochastic volatility of economic fundamentals, which are unpredictable by definition.. To the best of my knowledge, this is the first study to separate these three components for commodities. By examining and comparing components across assets and markets, it will be clear why commodity returns do not converge with each other and other assets, especially before the financialization of commodity markets.

The rest of this chapter is structured as follows: Section 2 gives a brief literature review of related studies. Section 3 presents the modified theoretical LRR model. Section 4 discusses the empirical model with estimation strategies. Section 5 discusses the data. Section 6 provides the main empirical results and section 7 concludes.

I.2 Literature Review

This chapter is mainly related to several strands of literature in both macroeconomics and finance. The theoretical model comes from the long-run risk finance literature initiated by Bansal and Yaron (2004). In that paper, the authors propose a model with a small but persistent long-run component in the economy that affects both consumption and the dividend growth rate. Calibrations based on the setup shows the LR component can explain the equity premium puzzle (Bansal, Kiku and Yaron 2010). While the equity premium results depend on the the elasticity of the intertemporal substitution parameter being bigger than 1, Campbell (1999) suggests that data may disagree. In a similar setup, Hansen, Heaton and Li (2008) formalize the long run risk under CRRA utility preference and show that long run risk still matters for both the discount factor (pricing kernel) and cash flow (dividend). The setup has also been used to study the bond market (Bansal and Shaliastovich 2012), firm level credit spread (Chen 2010), etc. Commodities, although perhaps the asset class most sensitive to long-run risk, are absent, partly due to the fact that cash flows from holding commodities are hard to measure. Therefore, because calibration for commodities may be hard to justify, a more data-focused estimation strategy is appropriate.

The literature on stochastic volatility has been in existence for some time. Studies in finance mainly focus on its effect on the Black-Scholes pricing formula and option pricing (Wiggins 1987, Heston 1993, etc). In macroeconomics, the concept was made famous by the seminal paper of Bloom (2009). The counter-cyclical uncertainty is shown to have explanatory power for consumption and firm level investment choices, thus affecting asset prices as a result. Christiano, Motto and Rostagno (2014) argue that the second moment shock on steady state firm investment can largely drive the business cycle dynamics. The problem in this strand of literature is how to measure the level of uncertainty in those models. Almost all the models employ stock market volatility (VIX index) as the proxy for the uncertainty level. While simple and straightforward, this measure is likely to overstate the true level of uncertainty as discussed by Jurado, Ludvigson and Ng (2015). In addition, if we don't separate the steady state and IRF factors, the source of such uncertainty cannot be identified. That's another reason for our model setup.

The third strand of the literature consists of studies on the correlation of commodity markets with equity markets. As discussed in the introduction, previous studies seem to agree that commodity prices move on their own. While Tang and Xiong (2012) state that financialization could potentially increase such correlation, other studies as in Buyuksahin, Haigh and Robe (2008) do not find supporting evidence in the data. Because of the general difficulty in forecasting future commodity prices (Gorgano and Timmermann 2014), it is still true that commodities are less understood than equities, and how they co-move with equities and other common risk factors is uncertain.

I.3 The Theoretical Model

As mentioned above, the theoretical model is based on the standard LRR model. Modifications are made with the particular intention of separating three components in competitive equilibrium. Commodities are understandably quite hard to price theoretically as there are no obvious "dividends" coming out of holding specific commodities. The concept of the convenience yield was created to address this issue, yet the problem is that it is not directly observable and there are no universally agreed upon ways of estimating it. The assumption being made here is that the pricing kernel on commodities is identical to that of stocks. Then, the convenience yield can be calculated by the aggregate dividend delivered by a portfolio with same capital gains as the target commodity. The pricing kernel on the other hand can be inferred from the consumption portfolio. This result will allow us to use LRR or any cash-flow based model. Subsections are divided to show clearly how those components are identified.

I.3.1 Model with Only Long-Run Risk

In this subsection, the simple model with only long run risk in the economy is presented and analytically solved. More complicated models with short run risk and uncertainty will be discussed in the following sections. Solution mechanisms are similar.

To start, the preferences of the representative agent are formulated as Epstein and Zin (1989):

$$V_t = [(1-\delta)C_t^{\frac{1-\gamma}{\theta}} + \delta(E_t[V_{t+1}^{1-\gamma}])^{\frac{1}{\theta}}]^{\frac{\theta}{1-\gamma}} \tag{I.1}$$

where σ is the time preference, C_t is consumption, γ is the risk aversion coefficient, $\theta = \frac{1-\gamma}{1-1/\psi}$ and ψ is the elasticity of intertemporal substitution. If we assume that there is no credit constraint on the agent, then the budget constraint is simply:

$$W_{t+1} = R_{c,t+1}(W_t - C_t) \tag{I.2}$$

where $R_{c,t+1}$ is the return on savings.

The long-run economic fundamentals evolve according to:

$$x_{t+1} = x_t + \sigma e_{t+1}. \tag{I.3}$$

Here, x_t is the log of long-run factors, σ is the log of long run risk and e_{t+1} is a random variable assumed to be i.i.d normal. From now on, all lowercase variables except shocks represent the natural logs. Consumption growth is assumed to follow the long run trend:

$$g_{c,t+1} = \mu_c + \phi_c x_{t+1}. \tag{I.4}$$

The assumption is in line with the Permanent Income Hypothesis as well as much empirical evidence, showing consumption-smoothing for consumers without an intertemporal budget constraint with a no-Ponzi scheme assumption. The convenience yield growth rate of commodity j can be similarly modeled as:

$$g_{j,t+1} = \mu_j + \phi_j x_{t+1}. \tag{I.5}$$

To solve this model analytically, we follow the standard approximations in Campbell and Shiller (1988) that

$$r_{c,t+1} = \kappa_0 + \kappa_1 z_{c,t+1} - z_{c,t} + g_{c,t+1} \tag{I.6}$$

$$r_{j,t+1} = \kappa_{j,0} + \kappa_{j,1} z_{j,t+1} - z_{j,t} + g_{j,t+1} \tag{I.7}$$

where $z_{c,t} = log(P_{c,t}/C_t)$, $z_{j,t} = log(P_{j,t}/Div_{j,t})$ and κ_0s are constants depending on the average value of zs.

We first need to solve for consumption in order to determine the innovations in the pricing kernel. We follow the popular undetermined coefficient method with the form $z_t = A_0 + A_1 x_t$. With the intertemporal marginal rate of substitution for Epstein and Zin preferences being:

$$m_{t+1} = \theta log\delta - \frac{\theta}{\psi} g_{c,t+1} + (\theta - 1) r_{c,t+1} \tag{I.8}$$

and the standard asset pricing Euler's equation:

$$E_t[exp(m_{t+1} + r_{c,t+1})] = 1. \tag{I.9}$$

We can solve the coefficient $A_{c,1}$ and $A_{j,1}$:

$$A_1 = \frac{\phi_c - 1/\psi}{1 - \kappa_1}, A_{j,1} = \frac{\phi_j - 1/\psi}{1 - \kappa_{j,1}}. \tag{I.10}$$

Then we can write return innovations as:

$$r_{c,t+1} - E_t[r_{c,t+1}|I_t] = \kappa_1 A_1 \sigma e_{t+1} \tag{I.11}$$

$$r_{j,t+1} - E_t[r_{j,t+1}|I_t] = \kappa_{j,1} A_{j,1} \sigma e_{t+1} \tag{I.12}$$

as well as the innovation to the pricing kernel:

$$m_{t+1} - E_t[m_{t+1}|I_t] = [(1-\theta)\kappa_{c,1} A_{c,1}]\sigma e_{t+1} \tag{I.13}$$

so that the volatility can be expressed as:

$$Var(r_{j,t+1}) = \kappa_{j,1}^2 A_{j,1}^2 \sigma^2. \tag{I.14}$$

That is, all volatilities come from the long-run risk σ. The risk premium of commodity j can be calculated by:

$$\begin{aligned} E_t[r_{j,t+1} - r_f] &= -cov[m_{t+1} - E_t(m_{t+1}), r_{j,t+1} - E_t(r_{j,t+1})] - 0.5var(r_{j,t+1}) \\ &= ([(1-\theta)\kappa_{c,1} A_{c,1}][\kappa_{j,1} A_{j,1}] - 0.5\kappa_{j,1}^2 A_{j,1}^2)\sigma^2. \end{aligned} \tag{I.15}$$

Thus, when the economy has only the long-run risk, the commodity return and volatility can be determined by the one state variable σ alone. The Sharpe ratio depends on preference parameters.

I.3.2 Adding the Short-Run Risk on Commodities

The short-run risk can be added to commodities by modifying the return growth on the commodity convenience yield:

$$g_{j,t+1} = u_j + \phi_j x_t + \varphi_j \sigma u_{t+1} \tag{I.16}$$

where u_{t+1} is an i.i.d normal shock and φ_j is a commodity-specific constant. So now instead of merely being exposed to the long-run risk in x_t, commodities are also subject to idiosyncratic zero-mean current period random shocks.

Since there is no change on the consumption side, all relations on the pricing kernel remain the same. For commodities, the return innovation now is:

$$r_{j,t+1} - E_t[r_{j,t+1}|I_t] = \kappa_{j,1} A_{j,1} \sigma e_{t+1} + \varphi_j \sigma u_{t+1}. \tag{I.17}$$

Therefore, the volatility can be decomposed into:

$$Var(r_{j,t+1}) = \underbrace{\kappa_{j,1}^2 A_{j,1}^2 \sigma^2}_{\text{LR Risk}} + \underbrace{\varphi_j^2 \sigma^2}_{\text{SR Risk}} \tag{I.18}$$

as well as the risk premium:

$$\begin{aligned} E_t[r_{j,t+1} - r_f] &= -cov[m_{t+1} - E_t(m_{t+1}), r_{j,t+1} - E_t(r_{j,t+1})] - 0.5var(r_{j,t+1}) \\ &= ([(1-\theta)\kappa_{c,1} A_{c,1}][\kappa_{j,1} A_{j,1}] - 0.5\kappa_{j,1}^2 A_{j,1}^2)\sigma^2 - \underbrace{0.5\varphi_j^2 \sigma^2}_{\text{SR Risk Discount}} . \end{aligned} \tag{I.19}$$

This is to say, a commodity with a larger short-run risk component will be discounted more heavily than those with a smaller SR risk. The magnitude of such premium discounting will depend on the relative size of long-run vs short-run risk, which is measured by φ_j.

I.3.3 Adding Stochastic Volatility

In order to introduce Bloom (2009) style uncertainty, the long run volatility is now time varying as in:

$$\sigma_{t+1}^2 = \sigma^2 + \sigma_w w_{t+1} \tag{I.20}$$

where σ^2 is the average long-run risk level and w_{t+1} is another i.i.d standard normal shock. The system now is as described below:

$$\begin{aligned} x_{t+1} &= x_t + \sigma_t e_{t+1} \\ g_{c,t+1} &= \mu_c + \phi_c x_{t+1} \\ g_{j,t+1} &= u_j + \phi_j x_t + \varphi_j \sigma_t u_{t+1} \\ \sigma^2_{t+1} &= \sigma^2 + \sigma_w w_{t+1}. \end{aligned} \tag{I.21}$$

In order to solve this system, now guess $z_t = A_0 + A_1 x_t + A_2 \sigma^2$. Going through a similar process again, we have A_1s that are identical to those in the previous subsections and A_2 is:

$$A_2 = 0.5\theta[\phi_c - 1/\psi + A_1^2 \kappa_1^2] \tag{I.22}$$

$$A_{j,2} = (\psi/\phi_j - \theta)A_2 + 0.5[(-\frac{\theta}{\psi} + \theta - 1)^2 + (\kappa_{j,1}A_{j,1} - (1-\theta)\kappa_1 A_1)^2 + \varphi_j^2]. \tag{I.23}$$

This result gives the variance decomposition as:

$$Var(r_{j,t+1}) = \underbrace{\kappa_{j,1}^2 A_{j,1}^2 \sigma^2}_{\text{LR Risk}} + \underbrace{\varphi_j^2 \sigma^2}_{\text{SR Risk}} + \underbrace{\kappa_1 A_{2,j} \sigma_w^2}_{\text{Uncertainty}}, \tag{I.24}$$

And the risk premium decomposition:

$$\begin{aligned} E_t[r_{j,t+1} - r_f] &= \underbrace{([(1-\theta)\kappa_1 A_1][\kappa_{j,1}A_{j,1}] - 0.5\kappa_{j,1}^2 A_{j,1}^2)\sigma^2}_{\text{LR Risk Premium}} \\ &- \underbrace{0.5\varphi_j^2\sigma^2}_{\text{SR Risk Discount}} + \underbrace{[(1-\theta)\kappa_1\kappa_{j,1}A_2 A_{j,2} - 0.5\kappa_1 A_{j,2}]\sigma_w^2}_{\text{Uncertainty Discount}} \end{aligned} \tag{I.25}$$

While models in 3.1-3.3 are simplified and idealized and results depend on the calibrated preference parameters, we are given some important implications nonetheless: First, when first and second moment shocks are independent, both premia and volatilities can be decomposed into three components. The long-run part comes from the predictable long-run economic fundamentals. Such risk carries a premium. The short-run part represents the transient market inefficiencies that are likely to go away with time. This component may carry a net risk premium, but at the same time it also has a "discount" on the LR Sharpe ratio. The uncertainty part comes from the unpredictable shock resulting from the second moment volatility σ_w. While the sign of its term in equation (25) cannot be determined without knowing the preference parameters, empirical studies in other asset classes suggest that it may carry a net risk discount.

Because the commodity portfolio is hypothetical, calibrating the model for each individual commodity is going to be very difficult, if not impossible. Structural estimation is a better approach here for the linearized nature of both premia and volatility functions. The following section discusses my estimation strategy.

I.4 Estimation Strategy

In order to empirically estimate the equilibrium from the last section, we start by stating that the long-run economic fundamentals evolve according to:

$$F_t = F_{t-1} + u_t \tag{I.26}$$

where F_ts are $k \times 1$ column vectors containing all information about economic fundamentals at time t, and u_ts are column vectors with same dimensions measuring the innovations. u_t is zero mean but may not be i.i.d. The changes in volatility can be modeled as $u_t = \sigma_t^F \varepsilon_t^F$ where σ_t^F evolve according to the process stated in model equation (20). Empirically, it can be modeled by

$$log(\sigma_t^F)^2 = \alpha^F + \beta^F log(\sigma^F)^2 + \tau^F w_t. \tag{I.27}$$

The assumption of the same pricing kernel and long-run factors across all assets means that any asset return can be decomposed into a latent long-run component from aggregating loadings on those factors and another latent short-run component of residuals. So for any commodity j,

$$logP_{j,t} = \underbrace{\Lambda_F^{j,P} F_t}_{\text{LR}} + \underbrace{v_{j,t}}_{\text{SR}}. \tag{I.28}$$

The fact that F_t is latent means that there is no uncertainty component in the underlying relation above. Since in reality we can't observe either $\Lambda^{j,P}$ or F_{t+1}, in the econometrican's view, the return of the commodity is:

$$\begin{aligned} logP_{j,t+1} &= \underbrace{E[\Lambda_F^{j,P} F_{t+1}|I_t]}_{\text{LR Risk}} + \underbrace{(\Lambda_F^{j,P} F_{t+1} - E[\Lambda_F^{j,P} F_{t+1}|I_t])}_{\text{Uncertainty from LR}} \\ &+ \underbrace{E[v_{j,t+1}|I_t]}_{\text{SR Risk}} + \underbrace{(v_{j,t+1} - E[v_{j,t+1}|I_t])}_{\text{Uncertainty from SR}}. \end{aligned} \tag{I.29}$$

In the mean time, the return of any stock i can also be similarly modeled:

$$logX_{i,t} = \Lambda_{i,X} F_t + v_{i,t} \tag{I.30}$$

By subtracting (30) with (28) multiplied by $\Lambda^{j,P}$ and move the term $\Lambda^{j,P} logP_{j,t}$ to the right hand side , we have:

$$logX_{i,t} = \Lambda^F F_t + \Lambda^{j,P} logP_{j,t} + e_{i,t} \tag{I.31}$$

where $e_{i,t} = v_{j,t} - v_{i,t}$. In putting every stock i together for commodity j, the matrix representation is then:

$$logX_t = \Lambda^F F_t + \Lambda^{j,P} logP_{j,t} + e_t \tag{I.32}$$

To model the short-run component, we use a vector error-correcting structure to capture the mean-reverting behavior of commodity prices (Casassus, Collin-Dufresne and Routledge 2005):

$$\begin{bmatrix} u_t \\ \Delta logP_t \end{bmatrix} = \lambda \begin{bmatrix} 0 & 0 \\ \Lambda^P & -1 \end{bmatrix} \begin{bmatrix} F_{t-1} \\ logP_{t-1} \end{bmatrix} + \Phi(L) \begin{bmatrix} u_t \\ \Delta logP_t \end{bmatrix} + \nu_t. \tag{I.33}$$

We omit the subscript j for expositional convenience.. λ is the coefficient to be estimated and measures the quickness of returns going back to the steady state equilibrium. Notice that unless $\lambda = 0$, the econometric restriction on short-run volatility is looser than the restriction in the long-run risk theoretical models: Here, we only need the short-run risk to be mean zero, it may be serially correlated.

The complete model for estimation is then:

$$
\begin{aligned}
\left[\begin{smallmatrix} u_t \\ \Delta logP_t \end{smallmatrix}\right] &= \lambda \begin{bmatrix} 0 & 0 \\ \Lambda^P & -1 \end{bmatrix} \left[\begin{smallmatrix} F_{t-1} \\ logP_{t-1} \end{smallmatrix}\right] + \Phi(L) \left[\begin{smallmatrix} u_t \\ \Delta logP_t \end{smallmatrix}\right] + \nu_t \\
logX_t &= \Lambda^F F_t + \Lambda^{j,P} logP_{j,t} + e_t \\
F_t &= F_{t-1} + u_t \\
log(\sigma_t^F)^2 &= \alpha^F + \beta^F log(\sigma^F)^2 + \tau^F w_t.
\end{aligned}
\tag{I.34}
$$

As the model structure is reminiscent of those in the dynamic factor model literature (except that the dimension is much larger due to the size of X_t), we can investigate existing strategies first in order to select the most efficient one for this particular exercise. There are two common ways to estimate the system: A joint Bayesian estimation using likelihood-based Gibbs sampling techniques, or a two-step estimation with a principal-component approach. The MCMC estimation techniques are surveyed in Kim and Nelson (1999), Johannes and Polson (2009), and are implemented by Kose, Otrok and Whiteman (2003) in studying international business cycle factors and by Bernanke, Boivin and Eliasz (2006) in FAVAR, among other uses. The two-step approach first estimates the factors from equations (5) and (6) and then uses estimated $\hat{F}_t$ to replace true F_t in equation (4). While the Bayesian approach estimates the whole model in one step, thus avoiding the “generated regressors”, the problem of the two-step approach, in this model it has a significant computational burden, as the dimensions of X_t and F_t are much larger than in previous studies. On the other hand, the two-step approach has much faster calculation speed and as noted in Bai (2004), the generated regressor problem problem is mitigated when the cross-sectional dimension, N, is sufficiently large relative to the number of periods, T [1]. In addition, there is no evidence showing that MCMC approaches give universally more accurate estimates. Therefore, in this study we will use the more computationally efficient approach.

For the first step, in order to (i) estimate the number of factors and (ii) estimate factors themselves, I follow Bai (2004) with all assumptions on common stochastic trends, heterogeneous factor loadings, time and cross-section dependence and heteroskedasticity and mutual independency being met. Then without the need of differencing the data, factors can be estimated consistently as $\hat{C}_t$ as well as the dimension of C_t factor space $\hat{k}$ simultaneously. Define

$$V(k) = V(k, \hat{C}_k) = min \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} (X_{it} - \lambda_i^k \hat{C}_t^k)^2 \tag{I.35}$$

as the sum of squared residuals when k trends are estimated. The task is to find a penalty function $g(N,T)$ such that the Integrated Panel Criterion function

$$IPC(k) = V(k) + kg(N,T) \tag{I.36}$$

can recover the true number of factors r[2] according to Bai (2004). Let

$$\hat{k} = argmin_{0 \leq k \leq kmax} IPC(k). \tag{I.37}$$

Then if we have $N,T \to \infty$ and (i) $g(N,T)\frac{loglog(T)}{T} \to 0$, (ii) $g(N,T) \to \infty$, r can be consis-

[1]That is to say, at any time period $t \leq T$, the number of cross-sectional observations, N_t, is large in comparison with t.

[2]Commonly used variable selection criteria include AIC and BIC. However in a large dimensional factor setting, both AIC and BIC fail on the consistency condition of estimating r.

tently estimated by $\hat{k}$. In practice, I choose the IPC function as[3]:

$$IPC(k) = V(k) + k\hat{\sigma}^2\alpha_T(\frac{N+T}{NT})log(\frac{NT}{N+T}) \tag{I.38}$$

where $\alpha_T = T/[4loglog(T)]$ and $\hat{\sigma}^2 = V(kmax)$.

So far we haven't utilized the fact that $logP_t$ is observed. In order to include it in the factor set, we need the new, partially observable factor set $Z_t = \{logP_t, F_t\}$ to span the same space as the original C_t. That is to say, we need the new set $\{logP_t, F_t\}$ to satisfy the following relation:

$$\hat{Z}_t = \begin{bmatrix} logP_t \\ \hat{F}_t \end{bmatrix} = \begin{bmatrix} \hat{M} \\ N \end{bmatrix} \hat{C}_t \tag{I.39}$$

where $\hat{Z}_t$ and $\hat{C}_t$ are $\hat{k} \times 1$ vectors. M is a $1 \times \hat{k}$ vector which can be estimated as:

$$logP_t = \sum_{i=1}^{r} m_i\hat{C}_{it} + error. \tag{I.40}$$

The error term can also be used to test whether P_t is well spanned by a linear combination of original factors C_t. N, however, is not uniquely determined, as there are various ways to rotate the system. As we want the new factor base to be orthogonal as well, the simplest way to find F_t under those criteria is: 1) Rotate P_t to any C_{it} based on the estimation in (12); 2) Rotate the rest of C_{it}s the same way as for P_t and 3) Do another round of factor extraction based on the procedure stated above for the rest of rotated C_{it}s to get F_{it}s. The component in P that can potentially change the factor structure is the error term. If P_t is a exact linear combination of C_{it}s, the new factor set Z_t is just a linear transformation of C_t. If it gives a larger number of factors, the first $\hat{k}-1$ factors are selected. Note that for part of the empirical study, this extra step of rotation is not necessary. It is mainly for consistently

[3]The choice of IPC function is of course not unique. Bai (2004) proposes three IPC functional forms satisfying all requirements. This form is chosen over the other two because, according to the Monte Carlo experiments, the estimated number of factors converge the most quickly to the true number regarding T. Such a property is particularly useful in forecasting, as often the time horizon is more limited than the cross-section.

recovering the aggregated trends from multiple factors as discussed below, where orthogonality makes factor loadings unique and allows no covariances among factors.

With estimated factors $\hat{F}_t$ in hand, the second step is to estimate the first equation in (34). Notice that the estimation of Λ^P depends on the normalization of latent factors F_t. In this study, factors are normalized by $FF'/T = I$. It is not hard to see that $\hat{\lambda}$ is invariant to factor normalization because of the zero upper block of the error correction matrix, and $\hat{\Lambda}^P$ will perfectly offset any different scaling of F. It is shown that error correction models with zero factor loading on the trend variable can be consistently estimated with generalized least squares (Gonzalo and Ng 2001). Since the key insight from the in-sample estimation is to compare the relative significance of the shared long-run risk trend component with idiosyncratic short-run cycles, I decompose commodity prices following Beveridge and Nelson (1981). When factors F_{it} are orthogonal, the real LR component can be expressed as:

$$log\hat{P}_t^g = logP_t + \sum_{i=1}^{\hat{k}} \Lambda_i^P A(I-A)^{-1} W_t. \tag{I.41}$$

Where A and W_t come from the VAR(1) representation of equation (4):

$$W_t = AW_{t-1} + e_{wt} \tag{I.42}$$

the cycle is then simply

$$logP_t^c = logP_t - logP_t^g. \tag{I.43}$$

Of course, the uncertainty comes from the imperfectly forecasted $E[A|I_t]$. As shown in (29), the uncertainty component is just

$$logP_t^u = (log\hat{P}_t^g - E[logP_t^g|I_{t-1}]) + (log\hat{P}_t^c - E[logP_t^c|I_{t-1}]) \tag{I.44}$$

which is driven by the unpredictable stochastic volatility component in F_t. The best forecasting algorithm here follows that in Chapter 2.

SUMMARY: Major Steps for Estimation and Decomposition

1. Using the entire information set X_t to determine the information dimension as well as obtaining $\hat{C}_t$;

2. Rotating $\hat{C}_t$ into $(\hat{F}_t, logP_t)$;

3. Estimating equation (33) with $(\hat{F}_t, logP_t)$;

4. Calculating Beveridge-Nelson-style multivariate decomposition based on estimation results in step 3 for latent true long-run trends and short-run cycles;

5. Calculating the observed LR risk, SR risk and uncertainty components based on the true trend/cycles as well as the forecasting errors.

I.5 Data

There are two major pieces of data used in this study. On the commodity side, I use commodity prices from the Commodity Research Bureau. Data is drawn from *The CRB Commodity Yearbook 2006*, which includes entries of 36 commodities spanning all five major categories: Energy, Grain, Softs, Livestock and Metals. The data frequency is monthly. While some commodities have price entries dating from 20's, prices are sparse before 1960. In this study I use data from the period of 1965-2005, which is an adequate time frame for the purpose of this study. As discussed in Wang (2010), even though some price series contain information from two sub-commodities (e.g., the corn series contains information from No. 2 and No. 3 yellow corn), the prices of these varieties track each other very closely. Therefore, it is safe to treat all sub-commodities within one commodity category as undifferentiated goods and average them out in our framework. All data are in nominal US dollars and are not subject to any additional modification other than taking logs prior to entering our model.

Table 1 shows some summary statistics on major representative commodities in each category and compares them with commonly used aggregated indexes such as the S&P500 Index and the U.S CPI Index. Annualized returns are calculated as compounded average yearly returns in each decade. Volatility is the ratio of average fluctuations to mean returns. We see that unlike the stock index, which steadily returns around 8% to 10% annual profit with a 20% to 30% standard deviation, commodities have large differences among them, both cross-sectional and over time. As Tang and Xiong (2012) indicate, those periods date from before the financialization of commodities, and not coincidentally, we don't see much obvious correlation, either, among commodities in the data. Such heterogeneity in commodity prices makes explaining returns and volatilities difficult by using just one or a few macro indexes. There is no simple convergence or balanced growth path shown here, and this situation may be one of the reasons that previous attempts on forecasting commodities with macro factors have largely been unsuccessful.

On the equity side, I use stock information from the Center for Research in Security Prices (CRSP), which contains common stocks prices and returns on NYSE from 1926 and NASDAQ from 1972. At first sight it seems straightforward to use stock nominal prices to estimate factors. However, by doing so we ignore split and buyback events (stock splits, stock dividends, etc.) that will make nominal prices incomparable over time. To get around this issue, I discount nominal prices with the Cumulative Factor to Adjust Price (CFACPR) variable in CRSP. There is another important aspect of stocks that is ignored in stock prices alone: dividends. In order to account for them, I construct hypothetical non-dividend prices that will give dividend-paying stocks the same return as if they were non-dividend-paying. Such hypothetical prices are calculated simply by adding dividend yields to the log adjusted prices of current and all subsequent periods. Hypothetical prices constructed this way differ significantly from raw nominal prices for only a few stocks such as AT&T. For most other stocks, though, changes are minor. To avoid issues of sparse information, I restrict my information dataset to contain only large-cap and mid-cap stocks that are listed and active for at least 10 years. The idea is that the bigger the companies the more likely they are going to use more commodities purchased from the market, thus adding short-lived smaller firms provides little information. For firms that have been active for more than 10 years but are delisted by the end of 2005: If they are delisted and go out business forever, all future prices are assumed of zero value; if they are delisted because of buyout or merger, hypothetical prices are calculated based on subsequent returns of surviving entities. If firms are delisted because of privatization, I have to omit them due to the fact that no available information is recorded afterwards. After all the criteria have been applied, there are 1552 adjusted time series of log prices in total during 1965-2005. [4]

We acknowledge the fact that the dataset we construct contains about only 30% of the total 5500+ available stock price series in CRSP during the sample period. One natural question is that whether the long-run information in our dataset is representative for the whole set. In order to check the robustness, we have tried a few alternative data cleaning strategies with different requirements on capitalization and survival period. Although number of stocks vary with strategies, the most significant long-run factors extracted from alternative datasets are almost the same, which means that the omitted series from our dataset have little additional useful information. Therefore, it is safe to proceed with our cleaned data.

[4]It seems untraditional to use individual stocks in the analysis as opposed to portfolios. The argument is that constructing portfolios needs a target index for sorting. As the trend factors to be estimated are latent, no obvious sorting criteria can be found ex-ante. Inappropriate sorting will cancel out significant trend factors. On the other hand, the factor estimation technique is capable of filtering out idiosyncratic shocks. To minimize this concern, I have chosen the most restrictive Integrated Panel Criterion to make sure that factors capture only the most shared trends.

I.6 Empirical Results

In this section we will present our empirical results. Section 6.1 focuses on the correlation of commodities with long-run factors and how the correlation depends on individual commodities. 6.2 tries to exploit such a relationship and finds the common risk factors for two group of commodities. 6.3 takes the process a step further and targets the commodity pricing anomaly by decomposing premia and volatilities into three components and a comparison with equities.

I.6.1 Understanding the Heterogeneity with LR Factors

Before we consider decomposing commodities, it would be beneficial to first look at the factor structure of equity markets to see how commodities' factor basis differs from it. I first estimate factors from the entire adjusted log price dataset from 1970 to 2005, according to the first step joint-estimation. The estimation returns a total of 54 factors. Compared with the sheer number of stocks in our dataset, such data compression performs well in terms of extracting important information while reducing idiosyncratic noise. Those factors are sorted in terms of their capability to capture equity market volatilities and are indexed from C_1 to C_{54}. Although such factors are latent and cannot be directly observed, constructing factor-mimicking portfolios is not a too difficult task, given that stocks are abundant.

With these factors in hand, we are ready to examine the differences and similarities among commodities and see why commodities are sometimes considered as being "disconnected" from other asset markets. To discern this outcome, I fit and sort three top factors that can explain the most fluctuations in each and every commodity with complete price information in our sample period. I also include the S&P500 index here for comparison. Commodities are divided into two groups based on the differences in production processes: One group for commodities that are usually extracted from the ground (the resource commodities) and another group for commodities that are usually harvested (the agricultural commodities).

The differences are obvious across groups and the index as shown in Table 2. For the Index, it is not surprising to see that the top two factors in the equity markets are also the two most significant in explaining the index. The third and fourth factors have almost the same small explaining power. For resource commodities, there seems to be a large level of cross-correlation: Out of eight commodities, six of them share the fourth factor and five of them share the first one. There are also two commodities that share the fifth and fourteenth factors. For all except one, the top three factors alone can generate over 50% of realized price volatilities in those commodities. With agricultural commodities, however, there is no such uniformity. In general, there are more factors taking effect and these account for less market movement than in the resource category. We do observe, though, that six out of eight commodities share the eighth factor and four of those share the fifth. The seventh factor is the most important factor for both coffee and cocoa.

It is clear by now why previous studies that try to link commodity prices with equity indexes report close to zero betas and inconsistent correlations. The top contributing factors in the index, $\{C_1, C_2, C_4\}$, while more similar to the top factors in resource commodities $\{C_4, C_1, C_5\}$, are very different from a large group of agricultural goods that share $\{C_8, C_5, C_7\}$. In addition, there is significant variation of factor structure even within a specific group. As those factors are $I(1)$ and orthogonal, using one single or a few macro indexes can potentially capture only a very small portion of the comovements.

So what are the economic reasons behind those numbers? As again factors are measurements of long-run trends, different factor bases for different commodities imply divergent trends. This situation can arise from both the supply and demand sides of the economy. There are many ways to affect supplies in the long run such as input-specific technology advances and low reserves of certain commodities, to name a few. Commodity-specific demand changes are less common, though, because people might not suddenly change their preferences. Demand side shocks, therefore, should be more present in cyclical/SR components, as they are usually more transient. When demand shocks take on the whole economy, as in times of recessions, the cycles may become more significant than usual and temporarily be highly correlated. This result in return gives increased correlations among commodities as documented by various sources mentioned in section 1. Once such demand shocks go away, correlations should disappear, too. In this sense, commodities are not typical assets that share a common market trend as suggested by asset pricing models such as in CAPM.

The immediate question after the fact is to what extent are returns shared by commodities and equities from those LR risk factors. To answer the question, we will need to conduct a return decomposition without any uncertainty. That is to say, it is necessary to look at the true underlying processes with perfect information. The first step is to choose appropriate factors for each and every commodity: More factors can better fit the current stage trend momentum but can potentially cause an overfitting problem. At the same time, more factors can give more uncertainty by adding unnecessary dimensions. Therefore, in this paper I use the Schwarz Information Criterion (SIC) to find significant factors while minimizing the uncertainty level.

Figures 1 to 4 show the implied trend components of oil, corn, coffee, eggs and copper, each falling into one general category of commodities. Again, if equity and commodity markets have shared the LR risk component at a disaggregated level, we should expect the trend components to explain a large portion of the commodity movements. All the figures have confirmed our theory. As we can see from these, all estimated trends closely resemble the general movements of all five commodities. In particular, over 90% of the realized volatility of oil and copper are explained, while over 80% are explained for corn and coffee. For eggs, 66% is explained. These figures contrast sharply with the previous belief that commodities and equity markets are disconnected. In fact, all evidence here suggests that instead the two markets are largely connected, at least in the long run. If over 90% of the oil price can be inferred from equities, then the oil price should provide very little additional information aside from firm-level data. This outcome gives an incentive for macroeconomists to "endogenize" commodities. For example, if we have a two-country model where one country ("south") produces commodity goods and the other country ("north") produces some manufactured goods by using commodities as an intermediate input, then as long as the production and asset pricing aspects of the North country are calibrated, the commodity price will be determined simultaneously. This result will provide important implications on the wealth re-distribution of north-south trades (Gao and Huang, 2014).

For the SR cycle components, there is no evidence, at least in our basket of commodities, showing that they are correlated with each other. Table 3 shows the average deviations from long-run trends for each of those commodities: They are likely idiosyncratic, time-varying and do not follow any clear patterns. In fact, by inspecting them individually, I find that big deviations from implied trends are usually caused by commodity-specific events: The huge spike in corn prices after 1995 is likely the cause of crop failures in the midwest; the unexpected oil price hike after 1990 coincides with the Gulf War, and the cycles of coffee also seem to be closely related to crop failures in Latin America and the trade embargo. In general, agricultural commodities have more volatile cycles as their production processes are longer and more dependent on uncontrollable weather factors. On the other hand, the amount of resource commodities produced in each period is much more flexible, and so long and persistent deviations from long-run prices are less likely to happen. It will be very interesting to compare price cycles with quantity cycles for their impact on the terms of trade, which are discussed in many international trade studies such as in Cole and

Obstfeld (1991) and Berka, Crucini and Wang (2012). Due to the current data limitation however, the study on quantities is left for future studies.

I.6.2 Finding the Common Risk Factors

The second step after establishing the correlation of commodities with equity long-run returns is to find whether such a relation is useful for finding risk factors in commodity spot returns. Due to the different nature of commodities and equities, using equity market factors has proven to be insufficient. de Roon and Szymanowska (2010) finds out that two factors drawn from the classic consumption-based CAPM model can explain about half the variance at the quarterly frequency. However, such a model has almost zero fit at the monthly frequency. There are also models that try to use market convenience yields or inflations as potential factors, but results are largely time-dependent.

The correlations at the factor level found in previous subsections provide a new channel for examining this question. If trends of log prices have a stable stochastic factor structure

$$logP_t^g = \sum_1^k \Lambda^P F_t, \tag{I.45}$$

returns should have a corresponding factor representation

$$r_{pt}^g = logP_t - logP_{t-1} = \sum_1^k \Lambda^P u_t \tag{I.46}$$

as well as the risk premia as indicated in APT

$$RP_{pt}^g = \sum_1^k \Lambda_i^P RP_{it}. \tag{I.47}$$

If there are common risk factors, the LR risk premia/discount for each commodity can be divided as:

$$RP_{pt} = \underbrace{\sum_{i \in common} \Lambda_i^P RP_{it}}_{common} + \underbrace{\sum_{i \notin common} \Lambda_i^P RP_{it} + RP_{it}^c}_{idiosyncratic}. \tag{I.48}$$

The estimation procedure is similar to that in Fama and MacBeth (1973): First, the common factor loadings Λ_i^Ps are estimated by regressing factors on $logP_t^g$ and then fixing the loadings $\hat{\Lambda}_i^P$ to calculate risk premia based on (48). Newey-West standard errors are calculated, since regressing only common factors on $logP_t^g$ might result in serial correlations. The Errors-In-Variable problem as in many Fama-MacBeth analyses is minimized here, as by construction $logP_t^g$ contains only linear transformation of orthogonalized factors.

The next step is to choose the common factors. While it may be possible to find one single significant "worldwide" factor representing a universal business cycle as in Kose, Otrok and Whiteman (2003), it is not quite obvious in this case as commodities are more divergent. But again if we focus on shared factors within each group, there are a few. I pick three factors for each group, and as a result, the agricultural risk factors consist of returns in the 5th, 7th and 8th equity factors and resource risk factors are the returns in the 1st, 4th and 8th factors. They represent the factors that can explain the most volatilities in each group.

So how do those factors differ from commonly used risk factors such as the Fama-French factors? Table 4 shows the projections of those 5 factors onto Fama-French factor space. In general, they correlate very little with all three equity factors. A weak correlation with market risk premium is present at 7.55% adjusted R^2 but SMB and HML are almost orthogonal to these factors. Therefore, they should span a very different space from both groups of potential commodity risk factors. Because the common risk factors are the most important LR factors for both stocks and commodities, the Fama-French factors are more likely to account instead for the SR risk.

Table 5 shows estimated betas from agricultural factors as well as Fama-French factors. The risk premia are grouped into the first column of Table 7. The left panel clearly shows that Fama-French factors have very little power in explaining the excess returns in almost all commodities. Other than the isolated cases (SMB for butter, HML for eggs), all betas are insignificant. On the other hand, three agricultural risk factors have significant factor loadings on most of those agricultural commodities. Except for sugar and apples, agricultural risk factors explain 1% to 2.5% of annual excess returns with adjusted-R^2 ranging between 10% to over 30%. Apples and sugar seem to be more isolated from the rest of the group, as only one out of three factors is significant. As average returns on agricultural commodities are around 0.5% to 2% per year, these common agricultural risk factors seem to capture big parts of the growth in prices over time.

Table 6 shows the numbers for resource commodities as well as in the first column of Table 8. Again, Fama-French factors are not able to capture the excess returns with perhaps the sole exception of crude oil. This situation is not too surprising, as oil is the most financialized and traded commodity in the basket. Yet even for crude oil, three FF factors combined can only explain less than 3% of the return variance. Resource risk factors do a much better job. Three factors combined have at least 1% risk premia for most of the commodities and more so for rubber and oil. As returns on resource commodities are in general higher at an average around 4% per year, they are relatively less significant than agricultural risk factors. Adjusted R^2s are also at lower levels in general as well. This situation suggests that compared with agricultural commodities, they are less influenced by group specific economic events and more likely to have idiosyncratic and divergent patterns: While weather can potentially affect many agricultural commodities at the same time, it is harder to find such exogenous factors affecting multiple resource commodities.

So are there any universal risk factors for all commodities? We do observe that innovations in the 8th factor, F_8, are statistically significant in almost all commodities. The problem is that this single factor generally has less than 10% of adjusted R^2s for most of our commodities, and its economic significance is not as apparent as in common factors of equity markets and among individual countries. This observation may be the reason why the commodity market is hard to understand as it contains more unrelated and independent determinants. This characteristic makes analyses with aggregated indexes less conclusive and it is why this latent factor-based model can provide more insights.

I.6.3 Premia and Volatility Decomposition with Uncertainty

Until now our analyses has been based on the assumption that everything is perfectly predictable in the economy. However, ignoring the unpredictable component is unwise if we want to look at difficult asset pricing problems. As discussed in Bloom (2009), uncertainty can halt demand for future cash flows and in return, drive down the asset prices. This result is especially true for commodities, as their prices are very dependent on economic fundamentals. Thus an uncertain outlook of future steady state is going to affect on both the demand and supply side.

Therefore, with the addition of uncertainty, based on our theoretical model, commodity spot price and volatility can be decomposed into three observable parts: 1) the forecastable LR risk component. This part represents the part of the economy that the steady state can be unbiased forecasted. Since all factors are orthogonal and I(1), the risk associated with this component can be hedged out but not diversified. As shown in the previous subsection, if the commodity is very forecastable, this component should be big; 2) the SR risk component. By definition, this is some transient part that goes away over time. In the theoretical model, because it is i.i.d, the whole component should be unpredictable ahead of time.

However, in real data it may or may not contain a forecastable AR(1) process. As long as it is of mean zero, it will satisfy our econometric setup. 3) the uncertainty component. This component measures the return/volatility that is generated by the error in forecasting economic fundamentals, that is to say, the error in estimating F_{t+1} given the information in time t. It will also give the measurement error on estimated loadings associated with those long-run factors. Notice that in our model the short-run cash flow can also be influenced by the long-run factors, so uncertainty also comes from the measurement error in the SR component.

Table 7 shows the risk premia decomposition for agricultural commodities and Table 8 shows this decomposition for resource commodities. When we inspect both tables, the first conclusion emerging for both groups is that long-run premia account for the majority of returns over the 40 years of the sample period. This outcome is not unexpected, as the previous subsection should have implied, and the forecast algorithm we employ here has minimized uncertainty. It does look interesting, though, for the SR and the uncertainty component: Almost all short-run components carry a positive, albeit small, risk premium and for both groups, the average is about 0.3%. Since the average premium on agricultural commodities is lower, there is a larger proportion of SR premia for agricultural commodities. The number confirms our suspicion that SR risk might not be i.i.d. In fact, upon looking at patterns of each agricultural commodities individually, we find that the even these SR components are of mean zero, they are mostly skewed to the right and most of these short price hikes are produced by unforecastable events. Those price hikes are more frequent for agricultural goods, thus it is priced-in accordingly. For uncertainty, it is almost universally negative. This outcome is in line with both macroeconomic and finance literature as uncertainty is generally counter-cyclical and counter-productive. Also for the worst performers, uncertainty looks to be the culprit: For eggs, though the counter-cyclical nature causes a -0.71% discount, uncertainty has an even bigger impact at -0.8%. For copper, the

0.25% discount brings down the total risk premium to just 0.05%, and for tin, the discount is even larger at 0.65%, which also brings down the total risk premia to a level far lower than the group average. For commodities with lower growth rates, uncertainty should not be overlooked, as it is likely to be the most dominant component.

Table 9 shows the variance decomposition for all commodities. The percentages do not add up to 100% because uncertainty may not be orthogonal to either LR or SR components. The difference between the two groups is much more apparent here: For agricultural commodities, the long-run variance accounts for 66% of the total variance, while for resource commodities it is at the much higher level of 90%. SR and uncertainty variances are 22% and 13% for agricultural commodities, while they are 7% and 5% for resource commodities. Therefore even if agricultural commodities are less correlated with long-run economic fundamentals, the uncertainty caused by their stochastic volatility is actually larger. This result might be caused by the less stable long run relation for agricultural commodities as well as the short-run component's more idiosyncratic quality. Thus when long-run volatility rises, agricultural commodities are much harder to forecast than resource commodities.

Another takeaway from this story is from the simple calculation of Sharpe ratios for different components. The numbers from Table 7-9 show a distinct risk-trenching pattern for both groups of commodities: Long-run volatilities are best priced (the "prime" component) in the commodity markets with the highest Sharpe ratio. The short-run volatilities, although they are still positively priced-in, have a much smaller return-risk rate, especially for agricultural commodities. The uncertainty volatilities are priced negatively so this is the "bad risk" in the cohort, although they cannot really be eliminated in any way. Thus, we see the reason why in general agricultural commodities are less profitable: They have a bigger portion of sub-prime risk and more than double the size of uncertainty of the resource commodities. The pattern is quite different from stock portfolios as well, as in stocks the major contributor of premia and volatilities is the short-run.

I.7 Conclusion

In this chapter I have done a comprehensive investigation of commodities and their relation with equity markets. While they are definitely more heterogeneous than other typical assets, I have managed to find a set of long-run factors that are shared among themselves and with cross-sectional equities. Those factors account for large portions of returns and volatilities in most commodities. According to this demonstrated connection, a global commodity risk factor and two sets of group risk factors for agricultural and resource commodities are found. The global risk factor is weak but the group risk factors are significant. The decomposition further reveals the heterogeneity of commodities, as unlike in the equity markets, commodity markets price long-run variances higher than those in the short run. Agricultural commodities in general have more volatile SR components than the resource commodities, which leads to lower returns over time. Uncertainty is especially dominant in those worst-performing commodities.

It is also worth noting that the econometric model developed here can be readily applied to other topics in economics and finance. One immediate extension is that instead of using equity prices, we can use firm-level balance sheet data such as the long-term debt to construct the long-run factors. Although balance sheet data are not available at monthly frequency, mixed-frequency "nowcasting" techniques will be useful, as long as we have the higher frequency asset market data. The results could have important corporate finance implications as well and it is left for future research.

Table I.1: Commodity Spot Price Summary

		Y65-74		Y75-84		Y85-94		Y95-04	
Commodity	Data Availability	Return	Volatility	Return	Volatility	Return	Volatility	Return	Volatility
S&P500 Index	1965.1-2005.12	-2.42%	11.48%	8.07%	21.76%	9.84%	26.25%	9.93%	25.27%
U.S. CPI Index	1965.1-2005.12	5.22%	14.36%	7.29%	23.03%	3.56%	11.06%	2.39%	7.72%
Corn	1965.1-2005.12	10.56%	40.94%	-1.63%	15.97%	-1.79%	15.65%	-1.5%	25.17%
Crude Oil	1965.1-2005.12	14.35 %	54.13%	9.17%	41.62%	-3.94%	22.49%	9.20%	43.36%
Coffee	1965.1-2005.12	4.37%	23.09%	6.94%	29.64%	1.10%	45.98%	-5.51%	48.38%
Eggs	1965.1-2005.12	10.16%	29.80%	-0.40%	15.09%	0.95%	14.49%	1.78%	21.90%
Copper	1965.1-2005.12	7.82%	25.45%	-0.78%	16.86%	8.23%	24.84%	0.08%	32.51%

Table I.2: Top Factors in Commodities

Commodity	Top Factors			Explained Volatility
	First	Second	Third	
S&P500 Index	1	2	4	98.86%
Agriculture-Based:				
Apples	1	4	18	61.36%
Butter	2	4	3	66.20%
Cocoa	7	5	8	41.82%
Coffee	7	8	14	54.42%
Corn	8	5	3	52.13%
Cotton	8	3	5	56.46%
Eggs	8	5	4	25.59%
Sugar	8	9	11	40.43%
Resource-Based:				
Aluminum	4	2	14	68.24%
Copper	1	4	6	56.09%
Crude Oil	1	4	5	61.82%
Iron	1	8	14	66.52%
Lead	1	8	13	60.07%
Rubber	8	4	1	55.92%
Tin	7	4	5	45.71%
Zinc	2	4	9	65.89%

Table I.3: Average Deviations from Implied Trend

Commodity	Average Monthly Deviations in Nominal US $						
	Y70-74	Y75-79	Y81-84	Y85-89	Y90-94	Y95-99	Y00-05
Corn	28.86	20.14	36.42	30.16	30.20	39.32	20.18
Crude Oil	2.303	2.652	2.886	2.309	2.631	2.204	2.828
Coffee	15.46	41.99	21.46	25.42	25.66	26.08	13.35
Egg	6.832	7.374	9.151	8.341	7.162	9.178	13.25
Copper	6.439	7.321	7.053	14.10	9.179	9.258	8.982

Table I.4: Projections of Common Commodity Factor Innovations onto Fama-French Factors

Factors	u_1	u_4	u_5	u_7	u_8	Adjusted-R^2
SMB	0.797**	0.0383	0.0752	0.0378	0.0146	1.36%
HML	0.171	0.286**	0.0383	0.0235	-0.0707	4.85%
MKTRF	0.796	-.0251**	-0.205**	-0.0315	0.184	7.55%

Table I.5: Risk Loadings for Agricultural Commodities

Commodity	Fama-French Factors				Agricultural Risk Factors			
	β_{SMB}	β_{HML}	β_{MKTRF}	Adj-R^2	β_{F8}	β_{F5}	β_{F7}	Adj-R^2
Apples	0.13	-0.82	-0.34	-0.01%	0.15*	0.00	-0.01	3.1%
Butter	-0.78**	0.79	0.04	1.22%	0.38**	-0.43**	-0.14	11.48%
Cocoa	0.43	0.62	-0.12	2.15%	-0.13**	0.10**	-0.01	30.39%
Coffee	0.23	0.76	0.32	0.28%	-0.04**	0.19	-0.21**	15.97%
Corn	0.20	0.10	-0.21	0.76%	-0.14**	0.17**	-0.85**	29.04%
Cotton	-0.09	0.08	0.12	0.22%	-0.14**	0.04**	-0.13**	20.36%
Eggs	-0.07	0.86**	0.28	1.4%	-0.11**	0.04**	-0.07**	18.45%
Sugar	-0.82	-1.61	-0.72	0.03%	-0.65**	0.01	0.00	7.5%

Table I.6: Risk Loadings for Resource Commodities

Commodity	Fama-French Factors				Resource Risk Factors			
	β_{SMB}	β_{HML}	β_{MKTRF}	Adj-R^2	β_{F1}	β_{F4}	β_{F8}	Adj-R^2
Aluminum	0.85**	0.39	-0.12	0.95%	-0.37**	0.04*	-0.01	7.34%
Copper	0.37	0.09	-0.07	0%	0.13*	-0.09**	-0.09**	8.46%
Crude Oil	0.94	-1.68**	-0.83*	2.81%	0.66**	0.01	-0.16*	16.03%
Iron	0.35	0.18	-0.69**	1.96%	-0.21*	-0.11**	-0.13**	17.12%
Lead	0.22	0.11	0.22	0.04%	0.24**	0.04*	0.04**	6.01%
Rubber	0.29	0.11	0.22	0.36%	-0.06	-0.06**	-0.16**	27.04%
Tin	0.38	-0.13	-0.44	0.53 %	-0.01	0.03	-0.05**	5.52%
Zinc	-0.15	-0.04	-0.01	0.02%	-0.03	0.02*	-0.04**	11.15%

Table I.7: Annual Risk Premia Decomposition for Agricultural Commodities

Commodity	Return Decomposition				Total
	LR-Common	LR-Aggregate	Short Run	Uncertainty	Premium
Agricultural Commodities:					
Apples	0.86%	3.94%	0.18%	0%	4.12%
Butter	0.44%	1.54%	-0.21%	-0.15%	1.18%
Cocoa	2.46%	2.17%	0.35%	-0.05%	2.47%
Coffee	1.07%	4.30%	0.37%	-0.16%	4.51%
Corn	0.14%	-1.10%	1.45%	-0.11%	0.24%
Cotton	0.45%	1.13%	0.1%	-0.11%	1.12%
Eggs	0.29%	-0.71%	0.01%	-0.8%	-1.5%
Sugar	1.47%	9.37%	0.31%	0.32%	10.01%

Table I.8: Annual Risk Premia Decomposition for Resource Commodities

Commodity	Return Decomposition				Total
	LR-Common	LR-Aggregate	Short Run	Uncertainty	Premium
Resource Commodities:					
Aluminum	5.03%	8.13%	0.09%	-0.08%	8.14%
Copper	0.83%	0.06%	0.24%	-0.25%	0.05%
Crude Oil	5.44%	9.28%	0.36%	-0.28%	9.30%
Iron	1.47%	3.6%	0.10%	-0.10%	3.6%
Lead	0.16%	1.69%	0.17%	-0.08%	1.78%
Rubber	4.55%	2.16%	0.23%	-0.18%	2.21%
Tin	2.49%	0.56%	0.13%	-0.65%	0.04%
Zinc	2.75%	1.39%	0.15%	-0.12%	1.42%

Table I.9: Annual Variance Decomposition

Commodity	LR-Var	SR-Var	Uncertainty	S.D (in USD)
Agricultural Commodities:				
Apples	67.74%	22.45%	10.47%	6.50
Butter	74.89%	17.67%	9.86%	36.37
Cocoa	81.36%	10.56%	9.52%	888.38
Coffee	75.33%	12.64%	12.78%	53.71
Corn	61.19%	22.68%	16.81%	61.77
Cotton	62.87%	23.93%	13.96%	15.84
Eggs	48.67%	42.08%	11.06%	13.87
Sugar	57.12%	24.61%	20.48%	6.43
Average	66.14%	22.08%	13.12%	
Resource Commodities:				
Aluminum	87.73%	5.00%	8.41%	20.21
Copper	84.91%	9.67%	6.48%	28.23
Crude Oil	94.68%	2.52%	4.20%	11.41
Iron	87.35%	12.59%	3.29%	39.79
Lead	92.47%	7.38%	1.98%	12.77
Tin	89.38%	5.92%	5.31%	164.95
Zinc	84.98%	8.12%	7.64%	15.94
Average	90.21%	7.31%	5.33%	

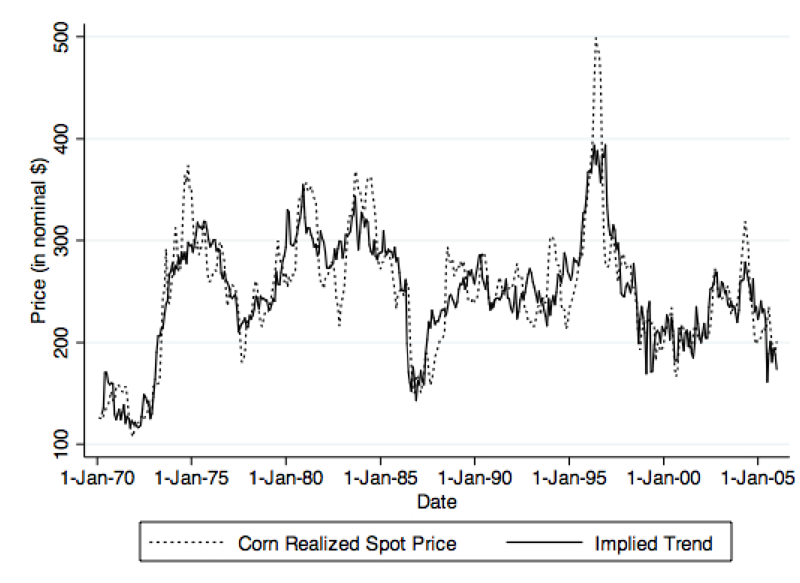

Figure I.1: Implied LR Trend in the Corn Realized Spot Price

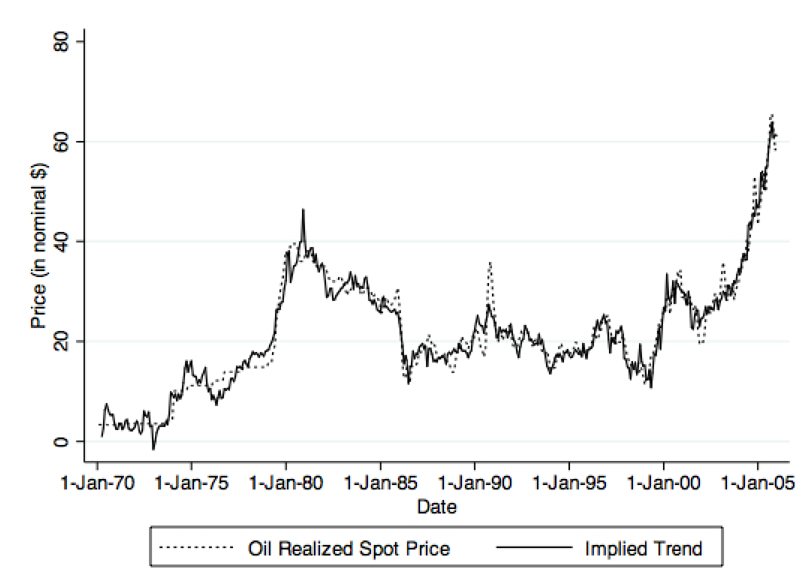

Figure I.2: Implied LR Trend in the Oil Realized Spot Price

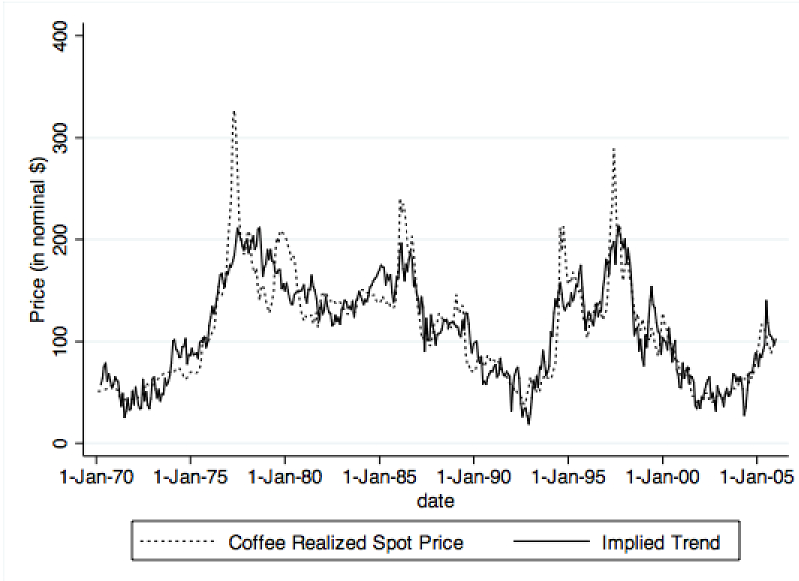

Figure I.3: Implied LR Trend in the Coffee Realized Spot Price

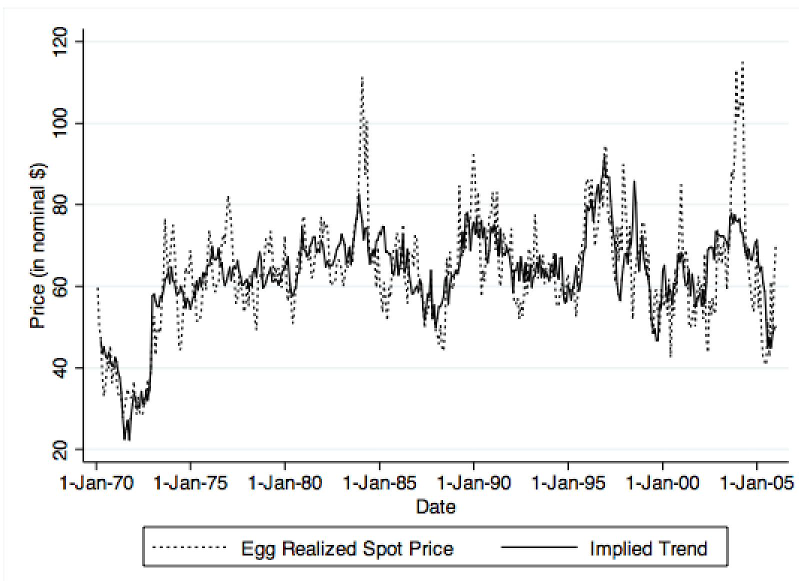

Figure I.4: Implied LR Trend in the Eggs Realized Spot Price

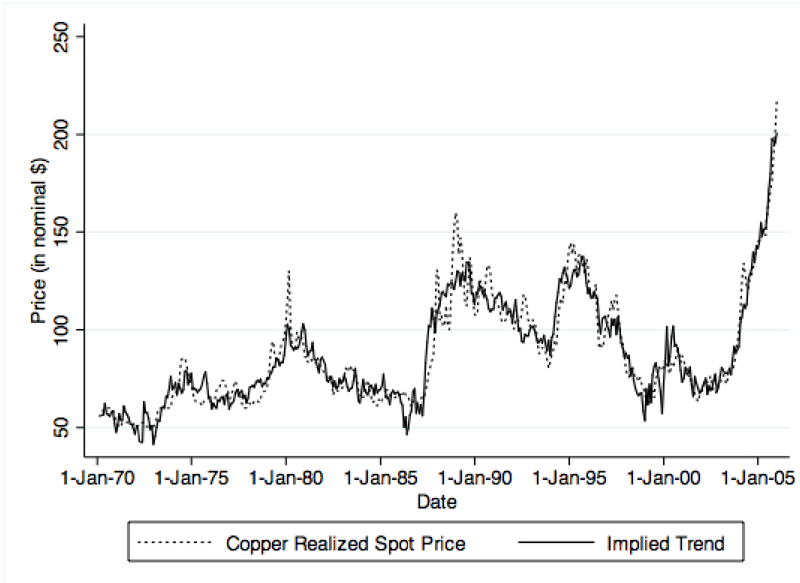

Figure I.5: Implied LR Trend in the Copper Realized Spot Price

CHAPTER II

Forecasting the Prices of Commodities in a Non-Stationary Data-Rich Environment

II.1 Introduction

In recent years, the fluctuations in commodity markets have attracted more attention from economists and practitioners alike. As globalized production is more dependent than ever on internationally traded commodities, a precise forecast of their future prices can be extremely helpful for firms deciding on investments as well as for policy makers. However, despite such urgent need, the academic literature on developing accurate forecast models for commodities is still in its infancy. Most attention has been given to the forecast of crude oil, given its unique publicity among all key commodities, yet many models are proven to be no better than the naive no-change forecast (Alquist, Kilian and Vigfusson 2011). Practitioners on the other hand have been using forward contract prices as the best forecast. Those prices are subject to market inefficiencies, and for the short-term forecast, they are essentially identical to the no change forecast when the contracted maturity dates are close. In addition, there may not be enough market depth for commodities other than oil, and thus obtaining the price corresponding to the desired forecast window is not always possible.

In this paper, we are proposing a unified two-component forecast method that is shown to have significant forecast power for many of the commodities in our data. In the model, we separate the price fluctuations into two distinct and tractable components: a long-run component driven by the permanent shared economic fundamentals and another short-run component driven by shared or idiosyncratic transient shocks. The theoretical backup of such a split can be found in the first chapter. Instead of forecasting price as a whole, the method forecasts each individual component based on different underlying price generating mechanisms: Since the long-run component is derived from the LR economic factors, it should be best forecasted from such factors. On the other hand, because the short-run component is defined as unforecastable from LR factors, a traditional forecasting method such as the no-change forecast or AR(n) forecast may be more feasible. Therefore, the cointegration relationship is forecasted by the LR component, while the idiosyncratic volatility is drawn from the SR component. We find evidence in data that such a split is economically meaningful while reducing the Mean-Squared Forecast Error (MSFE) in comparison with traditional forecast methods.

There are several key contributions of this paper. First, to my best of knowledge this is the first forecasting model that can explain the movements of multiple commodities at the same time. The heterogeneity among commodities is well known, which poses difficulty in forecasting them together. Our model overcomes the hurdle by: 1) providing a multi-dimensional long-run trend vector; and 2) forecasting an idiosyncratic cycle based on past performance. Therefore, the difference can be picked up by different factor loadings as well as the cycles. As long as the long-run relationship is stable, the out-of-sample returns can be easily calculated based on historical data. As shown by the empirical evidence, commodities indeed seem to share the long-run economic fundamentals and, as a result, our forecast can generally beat the no-change forecast as well as the future contract forecast.

Second, instead of using common macro indexes to represent the economic fundamentals as in Gargano and Timmermann (2014), I use large dimensional stock market data to estimate the underlying macro factors. That is, the LR factors are estimated similarly to the diffusion index (Stock and Watson 2002) approach, except that each and every index has to be of unit-root. The advantage of this approach is that: 1) the information is more comprehensive, since stock market data should contain information from all past, current and future projects; 2) the information is current and of high frequency. Compared to macro indexes that are usually published with lags and are subject to revisions, measurement errors can be reduced. The potential disadvantage of this approach is that individual latent factors are somewhat difficult to explain. Yet it does not mean that we cannot decide whether some specific observable variable has forecast power. One straight forward approach is to augment the latent factor space with the observable data and compare forecast performance with the original approach. If the augmented factor space has further reduced the MSFE, it is safe to conclude that such an observable factor has strong forecast power. In the data, however, no such further reduction in MSFE is observed, which means that a forecast with observable variables is unlikely to beat the model we have here.

Third, there is strong evidence across the board that expected and unexpected returns are negatively correlated. The theoretical framework for stocks is as explained in Pastor and Stambaugh (2013). The theory is confirmed in commodities as the past-year average monthly return has significant forecast power on short run cycles and the relationship is always negative for all commodities. This outcome also explains the “mean-reverting” behavior of commoditiy prices as documented in Casassus, Collin-Dufresne and Routledge (2005). It also makes the “unexpected” return somewhat forecastable from the “expected” return. As a result, the short-run forecast model in this paper can also beat the no-change forecast from last period. This is another reason for us to model the long-run and short-run separately. If we model everything in just one component, such potential forecastability is lost.

Last but not the least, our econometric setup is also helpful in the following ways: 1) It does not require detrending of raw data. There are potential forecast models such as Factor-Augmented VAR (Bernanke, Boivin and Eliasz 2005) that can also deal with a large amount of available data; however, this VAR setup requires the data to be stationary. Usually, such detrending is done by HP filtering. Yet since HP-filter is an in-sample filter, the meaning of forecasted values is questionable. Forecasting nominal prices without detrending has the advantage of clearer intepretation. 2) In-sample fit indicates good out-of-sample forecast. It is widely known that such relationship is not automatic. The advantage of trend-based forecasts is that if there is some underlying long run relation between commodities and factors, a good recursive forecast can be implied by good fit. When dealing with problems where factor loadings are time-varying, forecast can still be done on a rolling basis.

The rest of the paper is structured as follows: Section 2 provides a brief literature review on related studies. Section 3 introduces the forecast methodology. Section 4 persents the data. Section 5 presents the main empirical findings and Section 6 concludes.

II.2 Literature Review

This paper relates to several strands of literature. The studies on simultaneously forecasting all commodities is scarce to the best of my knowledge. Frankel and Rose (2009) do develop a model and try to find common risk factors of commodities but the study is in-sample. Chen, Rogoff and Rossi (2010) finds that commodity currency - US Dollar exchange rates can help forecast commodity price indexes. Gospodinov and Ng (2013) shows that common factors in commodity implicit convenience yields can also forecast the commodity indexes. Research on oil is extensive, and Alquist, Kilian and Vigfusson (2011) have a comprehensive survey on various traditional time-series forecast methods of oil future prices and show that many of them do not beat the native random walk. Baumeister and Kilian (2015) on the other hand show that by combining the traditional methods, the forecast error can be reduced in comparison with any of those stand-alone models. The implication is that traditional model forecast power is usually time-varying and scenario-dependent. This frequentist approach is also used in other studies such as Stock and Watson (2004). Baumeister, Guerin and Kilian (2014) also test the usefulness of the MIDAS approach developed by Ghysels, Santa-Clara and Valkanov (2005). They use high-frequency financial data to augment a Mixed Frequency VAR model and do not find financial data to be of additional help.

The use of common factors extracted by a PCA-like approach is popularized by Stock and Watson (2002) and Bernanke, Boivin and Eliasz (2005). All of them show that the dimension-reduction technique can help incorporate more information into the model than otherwise possible, thus avoiding the curse of dimensionality. The theoretical backup comes from Bai and Ng (2002) and Bai (2004), where they show that factors as well as the number can be consistently estimated. They also prove that the generated regressor problem is negligible as long as the dimension of cross section N is much larger than than the time period t for every period. The dicussion of variable selection issue under this predicting framework is in Bai and Ng (2008), where several methods including LASSO are presented. Inoue and Kilian (2007) show that Bagging can also help reducing MSFE with very simple pre-tests. Kim and Swanson (2014) revisit the variable selection issue with new data and conclude that none of the new techniques, including their combination, can dominate others.

While the functional form of our forecast model is closely related to the long run risk model of Bansal and Yaron (2005), the original model is rarely put into an out-of-sample test. Ferson, Nallareddy and Xie (2013) use a simplified version of the model with stock market data and conclude that the model has modest forecast power for the stock market, but not much better than the traditional CCAPM approach. The likely cause is that the study uses consumption data as the long-run proxy, and thus may not accurately recover the long-run trend in the data.

In the next section we are going to present our main forecast methodology that will improve upon the existing models.

II.3 Forecast Methodology

In this section we present our main forecast methodology. The first subsection describes the model for forecasting the LR component and the second subsection presents possible methods for forecasting the less structured SR component. The third section is to show the validity of our LR forecast under various simulated scenarios by comparing forecast performance with popular benchmarks.

II.3.1 Forecasting the LR Trend

To forecast the trend component, we first start with the empirical model in chapter 1:

$$\begin{aligned}
\left[\begin{smallmatrix} u_t \\ \Delta logP_t \end{smallmatrix}\right] &= \lambda \begin{bmatrix} 0 & 0 \\ \Lambda^P & -1 \end{bmatrix} \left[\begin{smallmatrix} F_{t-1} \\ logP_{t-1} \end{smallmatrix}\right] + \Phi(L) \left[\begin{smallmatrix} u_t \\ \Delta logP_t \end{smallmatrix}\right] + \nu_t \\
logX_t &= \Lambda^F F_t + \Lambda^{j,P} logP_{j,t} + e_t \\
F_t &= F_{t-1} + u_t \\
ln[(\sigma_t^F)^2] &= \alpha^F + \beta^F ln[(\sigma^F)^2] + \tau^F w_t
\end{aligned} \tag{II.1}$$

where u_t is the vector innovations in the long run factors F_t, P_t is the scalar price for any commodity, X_t is the matrix containing all stock information until time t. Λs are vector factor loadings while λ, α^F, β^F, τ^F are scalar coefficients. The system is driven by v_t, e_t, u_t and w_t. If we rewrite the first equation into a VAR(1) representation:

$$W_t = AW_{t-1} + e_{wt} \tag{II.2}$$

then the real LR trend component can be calculated by:

$$lnP_t^T = lnP_t + \sum_{i=1}^{\hat{k}} \Lambda_i^P A(I-A)^{-1} W_t. \tag{II.3}$$

Since the upper block of the error correction matrix in the first equation of (1) is zero, the trend can be written in a linear fashion similar to Crucini and Shintani (2014):

$$ln\hat{P}_t^T = \sum_{i=1}^{k} \hat{D}_i(L)F_{i,t} \tag{II.4}$$

where $D_i(L)$ is the decomposition lag operator corresponding to the ith factor F_i. To update the equation (4) for period $t+1$,

$$ln\hat{P}_{t+1}^T = \sum_{i=1}^{k} \hat{D}_i(L)[F_{i,t} + G_{i,t}(L)F_t] = \sum_{i=1}^{k}[\hat{D}_i(L)(1+G_{i,t+1}(L)]F_{i,t} \tag{II.5}$$

where $u_{i,t+1} = G_{i,t+1}(L)F_t$. That is to say, $G_{i,t}$ is the forecast matrix for $F_{i,t}$ that as long as we know exactly how it looks, we are able to forecast $F_{i,t+1}$ exactly. This leads to our predictive equation for the LR component by assuming that the decomposition operator $D(L)$ is independent with the forecast operator $G(L)$:

$$ln\tilde{P}_{t+1}^T|I_t = \sum_{i=1}^{k}[E[D_i(L)|I_t](1+E[G_{i,t+1}(L)|I_t,\sigma_t^F])]F_{i,t} \tag{II.6}$$

The above formula implies that the forecastability comes from two parts: the degree of stability of the decomposition lag operator $D_i(L)$s as well as the forecastability of innovations in F_t as in $G_{i,t+1}$. If the long-run decomposition is totally stable, the best forecast of $E[D_i(L)|I_t]$ is simply the estimates from last period $\hat{D}_{i,t}(L)$. If the innovations in F_t are of constant volatilities, the best predictor for the forecast matrix should also just be the in-sample estimate. For commodities, though there is strong evidence of stochastic volatilities in the data as shown in the first chapter, and therefore the predictor shall be adjusted accordingly. For the functional form of risk shock as in our model, the historical average volatility should be used. On the other hand, if $F_{i,t}$ is totally unforecastable, the forecast will shrink to a no-change forecast. Careful readers must have realized that it is not necessarily true that our forecast is always better than the no-change forecast. If the intrinsic long run decomposition is not stable, the error introduced by the first term can make our forecast perform worse than the naive random walk. It is then better to look for alternative approaches.

II.3.2 Forecasting the SR Cycle

By construction, the SR cycle is less structured than the LR trend. However it does not mean that we should treat the entire SR cycle as "unforecastable". As the model indicates, the SR cycle should be unforecastable by the factors F_t, therefore using LR factors to forecast an SR cycle would not work, as $E[P^C_{i,t+1}|F_t]$ will only give out spurious relations when $P^C_{i,t+1}$ and F_t are orthogonal. Instead, I consider the following forecasting candidates:

Zero Forecast: Regardless of past performance, the SR cycle is always assumed to be zero for the next period. This is essentially a forecast with only the trend forecast. This method is likely to work well when the LR is a dominant part and SR consists of mainly white noise.

No Change Forecast: This is the other extreme of the zero forecast, that the SR cycle forecast is simply the realized last period value. While simple and straightforward, it performs well in many forecast exercises. In fact it is the best forecast in some fields, as in the forecast of exchange rates (the Meese-Rogoff puzzle).

The above two are the benchmark models. The zero Forecast is useful in Monte-Carlo experiments since the SR can be muted and then it is possible to test only on the validity of the LR predictor. The no-change forecast is used in comparison with more sophisticated models as described below:

Volatility Shock Forecast: This forecast is used test whether changes in realized volatilities can help determine the size of future SR components. It is well known that return volatilities suppress risk premia expectation. Therefore, it might be the case that such relation is reflected in the SR component for the next period.

Uncertainty Aversion Forecast: The assumption here is that the uncertainty level may be correlated with the SR component. As discussed in Bloom (2009), the uncertainty makes people hold back from assets that look unsure to them. So a big forecast error in the past might have a negative influence on its future return.

Auto-Regressive Forecast: This is a standard statistical forecast with the simple assumption that there is some consistent persistency in the SR component. If it is true, we can extract the part of such auto-regressive forecastability and incorporate it with other baseline forecast methods to improve forecast performance.

GARCH Forecast: The GARCH forecast is similar to the AR forecast, except that the volaility now is time-varying as well. While the GARCH forecast has been quite successful in forecasting stock market portfolios, we have to make sure that modeling volatility is possible and necessary for the SR component, as over-fitting or poorly-fitting can both lead to increase in MSFE. Therefore before using GARCH, we should employ some pre-test to see whether such complexity over simpler AR forecast is justified.

Ultimately, the pattern of the data will determine the most appropriate method for the SR component. Finally, our forecast strategy can be summarized as:

SUMMARY: Major steps for out-of-sample forecasting strategy

1. To start, "train" the system with information set X_0 and $logP_0$ and obtain $\hat{F}_0$ and estimation of LR parameters as if in sample.

2. Forecast $ln\tilde{P}_1^T$ with $ln\hat{P}_1^T$ using parameters obtained in step 1; forecast $log\tilde{P}_1^C$ with the best SR strategy.

3. When actual data X_1, P_1 is observed, re-estimate LR equation and $\hat{F}_1$ with the addi-

tional data.

4. Forecast $ln\tilde{P}_2^g$, $ln\tilde{P}_2^c$ and so forth.

II.3.3 Monte-Carlo Experiments

In this subsection we test the validity of the long-run predictor with four different sets of Monte Carlo experiments. In every scenario, the short-run component is set to white noise and the forecast method consequently is the zero forecast. By this construction, we are making the trend predictor an unbiased predictor of the future price under ideal circumstances. For each set of experiments, simulations are repeated 100 times. I then compare the trend-only predictors with HP filter trends and naive random walk predictors. I also compare the predicted volatility with realized volatility. For all experiments, I set the number N of cross-sectional observations of X_t to be 100 and the time horizon T to be 50. There are 50 additional time periods generated before $t = 0$ to train my model prior to forecasting. Information is assumed to be at the monthly frequency. The HP filter is the only one that has an information adantage as it always uses the full set of data. For each set, I pick one representative simulation and discuss it below. The complete Monte Carlo results are shown in Table 1. Note that since our experiments run on a small set of factors, we would like to see how increasing the number of factors affects the predictive power. Thus in addition, I report what happens when the number of factors is doubled.

Case 1: Standard Dynamic Factors. Suppose that in the economy, the informational data follows the true data generating process:

$$X_t = \beta_{c1}C_{1t} + \beta_{c2}C_{2t} + \beta_{c3}C_{3t} + e_{ct} \tag{II.7}$$

where C_{1t}, C_{2t} and C_{3t} are the only three factors that drive the macro economy and e_{ct} is the idiosyncratic shock on X_t. β_{ci}s are assumed to be i.i.d and follows $N(0,1)$. The C_t factors are assumed to be generated in the following way:

$$\begin{aligned} C_{1t} &= C_{1,t-1} + \mu_{c1t} \\ C_{2t} &= C_{2,t-1} + \mu_{c2t} \\ C_{3t} &= C_{3,t-1} + \mu_{c3t} \end{aligned} \tag{II.8}$$

so each macro factor is $I(1)$. The economic "shocks" as modeled in μ_ts, are assumed to follow AR(1) processes:

$$\begin{aligned} \mu_{c1t} &= \beta_{\mu_1}\mu_{c1,t-1} + e_{1t} \\ \mu_{c2t} &= \beta_{\mu_2}\mu_{c2,t-1} + e_{2t} \\ \mu_{c3t} &= \beta_{\mu_3}\mu_{c3,t-1} + e_{3t} \end{aligned} \tag{II.9}$$

where e_{1t}, e_{2t} and e_{3t} are i.i.d and distributed standard normal. P_t is observed and is generated based on:

$$lnP_t = \beta_{p1}C_{1t} + \beta_{p2}C_{2t} + \beta_{p3}C_{3t} + \omega_t \tag{II.10}$$

where ω is a zero mean, $I(0)$ process that we assume to be i.i.d $N(0,1)$ as well. However, in a general sense, it does not need to be i.i.d as long as it does not contain additional trend information (i.e., another $I(1)$ process not spanned by the factor base C_t). The case that lnP_t does not fully lie in the C_t factor space will be considered shortly.

To simulate the economy, I let all βs that are not in (10) to be sampled from a uniform $U(-1,1)$ distribution. For illustration purposes, I assign the following values: $\beta_{\mu 1} = 0.6$, $\beta_{\mu 2} = -0.5$, $\beta_{\mu 3} = 0.4$ and $\beta_{p1} = 0.1$, $\beta_{p2} = 0.2$, $\beta_{p1} = 0.9$. Notice that although we have assigned some negative numbers to the coefficients, there is no cyclicality assumption on lnP_t. The co-movement relation with all three factors can only be determined after u_{it}s and ω_t are realized.

We start by looking at HP filtered results in Figure 1 with $\lambda = 1600$. The solid line represents the generated lnP_t, while the dashed line is the HP-filtered trend and the dotted line is the cycle. As we can see, the HP filter does a poor job of capturing the true trends in this economy. The cycle contributes to the vast majority of fluctuations in lnP_t. In fact, over 90% of the realized volatility of lnP_t is from the unpredictable cycle component. In terms of real world application, such a feature is very unappealing.

Figure 2 shows the one-period forecast results from our forecast method. To obtain this result, we first estimate the number of C_t factors in the informational set X with $kmax = 10$. The program returns an estimate of $\hat{k} = 3$ consistent with the DGP and the estimated $\hat{C}_t$s based on the algorithm described in the above section. Then we test which C_t is in the closest proximity to P_t. It turns out that P_t is the closest proxy to $\hat{C}_1$, which means that P_t is responsible for the largest portion of variances in X. Then we rotate the rest of the factors $\hat{C}_2$ and $\hat{C}_3$ to form $\hat{F}_1$ and $\hat{F}_2$. Finally, we can bring our new factors $[lnP_t, \hat{F}_{1t}, \hat{F}_{2t}]$ into the extraction function for each period from $t = 0$ to $t = 49$. Since the factors are normalized in estimation and not to scale, the estimated coefficients are of little meaning except that the signs of elements in A relating factors and lnP_t DO show cyclicality. In this case, both signs are negative and lnP_t is counter-cyclical to both macro factors $\hat{F}_1$ and $\hat{F}_2$.

It is obvious, in comparing two figures that our method has tremendously increased forecast accuracy. In our method, the trend more accurately reflects the true economy-driving forces, and the unpredictable forecasting error is down to 16% of the realized volatility, a significant decline over the HP filter.

If we repeat the same process 100 times and summarize FMSEs and volatilities, we obtain the first section of Table 1. In this ideal setting, our forecast method outperforms the random walk and HP trend by a wide margin. When the economy has 3 macro factors, our predictor in general returns an average error of 1.75 with slightly inflated predicted volatility vs realized volatility comparing 1.85 for the random walk and 1.92 for the HP trend. Predicted volatility is reduced by adding the random walk cycle predictor as discussed. The performance gap further widens for the six-factor case as measured by RMSE. It shows that when data is well-behaved, our forecast method matches both level prediction and volatility very well. The random walk method, though by design will return 100% predicted volatility, will return larger errors on average.

Case 2: When X_t contains lagged factors. It is not always true that factors have only a contemporaneous effect on the economy. Some factors may take time to be effective. To capture this feature, we modify our observed data generating process to become

$$X_t = \Lambda^F(L)F_t + \Lambda^P lnP_t + e_t \tag{II.11}$$

and in simulation, we alter the true data generating process to be

$$X_t = \beta_{c1}C_{1t} + \beta_{c1l}C_{1,t-1} + \beta_{c2}C_{2t} + \beta_{c3}C_{3t} + e_{ct} \tag{II.12}$$

with β_{c1l} assumed to be i.i.d and standard normally distributed. All other conditions are the same as in the previous subsection. We don't consider the case where there are lagged factors in lnP_t, as it is easy to verify that such a structure is no more than an increased error term $\tilde{\omega}_t$.

Figures 3 and 4 show the dynamics from HP filtering and our forecast method, respectively. The interesting feature of this scenario is that the DGP gives a fairly moderate fluctuation in lnP_t for about the first 20 periods and volitility then dramatically increases. This result is usually caused by a dominant role from factor(s) that have more a persistent effect than others. Indeed, the simulated data shows that such a big fluctuation is generated mostly from $\hat{C}_1$. In this case, the problem with the HP filter and other statistical filters is that they are not capable of capturing the sudden, big ups and downs as shown in the second half of our time series. Those big jumps are usually "smoothed out" in this class of filters.

Our method is able to predict the general momentum of lnP_t: In the first half of data the predicted trend matches the data almost perfectly, and in the second half, the directions of movement are predicted correctly. The unpredictable volatility from the cycle component consists of about 41% of the total realized volatility from my model, compared with 53% under the HP filter.

Looking at Monte Carlo results, we see that in general our method is still the most accurate. Although in the three-factor case, such an advantage is not as obvious as in the standard case. Our method returns a slightly flattened predicted series compared with the realized data. Prediction accuracy improves in the six-factor case, but that may be because of the decreased significance of C_1, as random walk prediction also improves.

Case 3: When *lnP* contains information outside *X*. In previous scenarios, lnP is always "carefully chosen" so that it does not contain $I(1)$ information outside of C_t. In the real world though, lnP may not be chosen as perfectly. In this exercise, we consider the case in which lnP contains additional information that is not in X. To proceed, we alter the data generating process of P_t to become:

$$lnP_t = \beta_{p1}C_{1t} + \beta_{p2}C_{2t} + \beta_{p3}C_{3t} + \beta_{\omega}\omega_{1t} + \omega_{2t}. \tag{II.13}$$

where ω_{2t} corresponds to the i.i.d zero mean shock as in previous cases and ω_{1t} is the new information and follows an $I(1)$ process:

$$\omega_{1t} = \omega_{1t-1} + e_{\omega t} \tag{II.14}$$

we assign $\beta_{\omega} = 0.8$ for the purpose of illustration.

Figures 5 and 6 again show the comparison of the HP filter and our method. The HP filter is unaffected by contaminated lnP as it does not rely on the such information and it suffers the same problem as discussed above. Figure 6 shows that our method still returns good predictions under mild information "overflow" in lnP. The predicted trend, though, does seem unable to correct itself as fast as possible as in case 1 where lnP is generated "clean". The cycle component does show signs of unit-root, especially towards the end of the time span. A Dicky-Fuller test confirms the existence of a unit-root (although the unit-root test is not a sufficient test on whether P has more unit roots than X). This fact will make the forecast error larger. In this case, the forecast error is 31% of the realized volatility in lnP, which is still significantly below 59% under the HP filter.

Turning to Monte Carlo, since the contamination is targeted at disconnecting X_t and lnP_t, we see that while our model still predicts the trend well, its RMSE is bigger (but not by much) than the random walk when the contamination is large in the three-factor setting. Accuracy significantly improves when such contamination becomes milder as in the six-factor case. The random walk is understandably tougher to beat here as it does not rely on the assumption that X_t and P_t are somehow driven by (mostly) the same factor set. However we see that as long as such contamination is not too large, our forecast method still gives better predictions than the random walk.

Case 4: When lnP is noisy. The last part specifies the case when P has an $I(1)$ outside of X. This setup gives an advantage to the random walk model since my forecast model cannot adjust efficiently to shocks originated from that $I(1)$. In empirical studies, there is often another way of introducing errors such that while unit roots of P can be correctly captured by X, there is a very significant white noise component that is only in P. While white noise in general creates difficulty in forecasting, it should hit the random walk model particularly hard. In this scenario, it is our forecast method that has the theoretical advantage.

In order to see the Monte Carlo comparison, I take the data-generating process in case 1, except that now ω_t has a standard deviation three times as large as in the original case. The result summarized in section 4 of Table 1 has confirmed our initial guess. Such noise, especially in the three-factor model, hits prediction results universally across all methods, since white noise is by definition unpredictable. We see that, however, with more factors coming into the system, our forecast method improves its prediction much faster than the random walk model does.

In summary, simulations have indicated that my model is capable of producing good forecasts under various conditions. It also almost always dominates the random walk model. This outcome is important in the sense that for applications using real data, we don't know the true data-generating process. We would like our forecast to still be as accurate as possible, even with potential model mis-specification or data contamination.

II.3.4 Comparison with Bernanke-Boivin-Eliasz's FAVAR

It is natural to compare this method with FAVAR, as both methods extract factor information from a large informational dataset. There are a few distinctions that are worthy of mention: First, since FAVAR models generally adopt an unrestricted VAR framework, they require the data to be stationary. In empirical research, the common practice is to either use differenced data or HP filtered data if the data is non-stationary. While such adjustments are fine in economy-wide aggregated time series, we need to be cautious when using commodities and other microeconomic time series. For one thing, forecasting HP filtered cycles is different from forecasting the original variable, as the HP filter itself provides little forecast power. On the other hand, Bai and Ng (2002) as used in FAVAR tend to give inconsistent estimates on factors for differenced data: If P_t is of the form $P_t = \lambda_1 F_t + \lambda_2 F_{t-1} + e_t$, differencing with P_{t-1} gives $\Delta P_t = \lambda_1 u_t + \lambda_2 u_{t-1} + \Delta e_t$. FAVAR will have two factors while we only need one. In the world of micro data, more factors mean fewer degrees of freedom and potential over-fitting, which tranlates to bigger forecast errors when historical performance is limited. In our model, data are in log level form and no adjustment is necessary before being analyzed. Secondly, since there is no long-run restriction imposed in the FAVAR, the trend-cycle decomposition is not possible so there are no long-run implications. Thirdly, the ordering of VAR is critical in FAVAR, as the "slow-moving" vs "fast-moving" time series assumption is used for identification. While many stylized facts are documented in terms of macro indexes, for commodities there are few. Therefore I have avoided this FAVAR's identification assumption on the information dataset. In a general sense, FAVAR and my model really target on different issues with similar modeling features.

II.4 Data

Data used in this chapter are similar to what is in the first chapter. The stock market data are from CRSP and the data cleaning process is as described in the data section of that chapter. The commodity spot prices are taken from the Commodity Research Bureau Yearbook 2006. There are some minor combinations of commodity sub-categories but they do not affect the price levels by much. The future contract data is from Thomson Reuter's Datastream. Other policy related variables are taken from FRED, the database maintained by the Federal Reserve Bank of St. Louis. For a full description, including summary statistics, please refer to the data section in Chapter 1.

II.5 Empirical Results

As discussed in the last section, there is no universal method for the SR component. Therefore, it is better to pre-test for an optimal model for SR only before moving on to forecasting the combined price series.

Table 2 and Table 3 show the summary of four potential forecasting models. The four models correspond to those mentioned in the methodology section. The testing strategy is to employ linear predictive regressions in sample and test for parameter significance over the entire sample period. To avoid over-fitting the model and introducing unnecessary MSFE when performed out-of-sample, I set all predictors in all four models to be yearly measures in the past 12 months instead of monthly. In other words, the interpretation of those models is whether the average movements of those predictors in the past year have any potential predictive power on the SR cycle component. The forecastability from monthly frequency is likely to be picked up by augmenting with a no-change forecast. We test with multi-year lags, but since none of those lags are significant, in the tables I report only the coefficients for the first lags. Each model is tested individually. Therefore, the concept of this pre-test is similar to using the t-test as in Inoue and Kilian (2007). There is no guarantee that such a pre-test is the best test, but its testing power meets our need here.

The first column in the two tables shows the results for the volatility shock forecast. In general, there is only modest evidence that change in the last year's volatilities can forecast the future SR return. Out of 15 commodities, only three of them show some significant correlation. For apples, a one percent standard deviation change is likely to cause a 0.28% increase in the short run for the next period. For coffee the number is 0.22%. The only resource commodity that shows a similar pattern is crude oil, where a 0.12% return increase is accompanied by a 1% increase in volatility. In general however, such correlation is not evident for many commodities.

The second column shows the effect of uncertainty on SR returns. In recent years, multiple studies have been focused on the effect of uncertainty on economies and asset prices (Bekaert, Engstrom and Xiong 2005; Christiano, Motto and Rostagno 2014, etc). The results here seem to suggest that uncertainty does not play a strong rule in determining commodity prices, especially since uncertainty by definition is not a LR economic factor, and therefore its only possible effect should be on the SR components. For agricultural commodities, only the price of eggs shows strong uncertainty aversion: A one percent increase in the uncertainty level from the past year is likely to decrease the return of the next period by 0.9%. On the resource commodity side, only the crude oil shows modest statistical significance and even for that, the number is much smaller than that of eggs at 0.17%. For all the commodities, the predicting power of the uncertainty level looks weak.

The third column tries to measure the behavior of "mean-reverting" in the SR components. As shown in the data, such behavior is very evident. To the sharp contrast of the first two models, all coefficients are statistically significant at the 5% level, and all of them are with negative sign. That is to say, a high average SR return in the past 12 months is likely to cause the next period SR return to decline. The scale of such decline depends on individual commodities: For agricultural commodities, the numbers are mostly in the -0.5% to -0.7% range, whereas for resource commodities, the range is even narrower at -0.46% to -0.6%, except for iron. This evidence confirms that commodity traders do trade with "equilibrium" prices in mind: If some commodities trade at higher prices than the market equilibrium can justify, the price is likely to be corrected in the future. Also, it might be possible that overcorrection exists, provided the historical SR component is sufficiently large. However, in the long term, the prices will eventually revert to the levels where economic fundamentals support without further SR shocks. Given the significance and the stability of those autoregressive coefficients, this is the better model to use for commodities.

The fourth column tests on whether we should use a more complicated GARCH approach instead of the simple AR approach. A straightforward way to do so is to test whether volatilties are serially correlated. The numbers in the two tables show mixed results: For agricultural commodities, the stochastic volatility assumption does not hold for most. Only sugar has a significant coefficient. For resource commodities on the other hand, about half of the commodities are significant at the 5% level: while there is no obvious evidence for oil, lead and tin, copper, iron and zinc all show strong pre-test results that suggest a GARCH approach might be better. This outcome probably explains in part why the GARCH forecast can be successful for a few selected commodities. Since our goal here is to find an unified forecast model for all commodities, we will proceed with the auto-regressive approach.

After selecting the forecast method, the next step is to test it in an out-of-sample environment. In order to also account for possible persistency, we augment the AR model with a no-change predictor from last month. If the AR structure is really helpful in forecasting, we should see reduced MSFE in comparison with the simple random walk. The linear predictive regression consists of a one-period lag predictor (the no change predictor) plus the historical 12-month return predictor (the AR predictor). Results are summarized in Tables 4 and 5. The first two columns are coefficients associated with the predictors that are estimated in the last period of forecasting, and the right two columns are the average forecast error throughout the forecasting period. It is quite obvious that the single forecasting model works remarkably well for all commodities, as all of them show a positive correlation with the last period return while having a negative correlation with the last 12-period average return. For agricultural commodities, the 1-period "persistency" parameters range from 0.71 to 0.88, while the "reversal" parameters are within -0.30 to -0.12. For resource commodities, they range from 0.65 to 0.88 and -0.34 to -0.31 respectively. However upon further inspection, except for apples, eggs and iron, almost all the commodities have a persistency parameter around 0.85 and a reversal parameter around -0.15 to -0.2. The fact that such a consistent trading pattern can be identified across so many commodities in a 40-year time span is very interesting and can readily be used to better forecast the short-run components that are based more on trader confidence, market momentum, etc. The MSFE comparison also tells us that our forecast is superior to the naive random walk forecast: In general, we observe around a 10% MSFE reduction from the baseline model and some of them (eggs, iron, etc) can have a reduction of more than 15%. This result is especially important for agricultural commodities where SR components are likely to be more significant.

Now that we have a good SR forecast, it is time to combine it with the LR forecast. As shown in the simulation section, the LR forecast is also almost always better than the no-change forecast. Conceptually, the combined forecast should then also be a better forecast. In comparison with the fake data experiment, there is one additional task here that we have to consider: the variable selection, as here we have many more LR factors to choose from. Since it is unlikely that all of them can have an effect on every commodity, it is essential to pick up the most useful ones. There are multiple ways of doing that as described in Bai and Ng (2008). In this study, we find that the bagging (bootstrap aggregation) method first introduced by Breiman (1996) and formalized in Inoue and Kilian (2007) to be the most effective. It is a sub-sampling technique to stabilize predictors and minimize MSFE. For our exercise, I use the popular 12-period block replacement so to take care of potential seasonality in our data. For the pre-test, I find the simple t-test to work better than the more commonly used Schwarz Information Criterion (SIC). I use the data from 1965 to 1970 as the training set and report forecast performance for 1971 to 2005.

Figures 4 to 7 show the continuous one-month-ahead forecasted trends of corn, oil, coffee, eggs and copper. At this point, because the SR forecast are muted, the forecast plots are trends only. In general, my predicted trends track the real trends well with perhaps the sole exception of eggs. For eggs, while the general momentum is correctly predicted, short-term volatilities are also very large. It is understandable, though, as the in-sample estimation the volatilities of eggs also comes mostly from the cycles. For corn, the forecast goes well from the very beginning. It does not capture a few temporary price hikes as well as in in-sample analysis. For oil, while the forecast for the first few years is a little off, it afterwards performs much better. It also successfully captures the oil booms during the 80s and after 2000. A similar pattern exists in the copper forecast as well. It makes economic sense, as those commodities are facing much stronger demand from emerging economies, especially from China, and such increased demand should show its impact on the equity prices of related firms. On the other hand, the midwestern drought of 1995-1996 that caused the food price hike, including an increase in corn prices, could not be foreseen from other economic activities. For the price of coffee, while it is well forecasted, it is also subject to significant regulation influence (Ponte 2002), and those uncaptured hikes coincide with an unexpected trade tariff and embargo changes as well as crop failures overseas.

By adding back the SR forecast, we will be able to compare full model forecastability with benchmark models. Here I employ two baseline forecasts: The first forecast is simply the forecast based on future contract prices that are used extensively in the related literature. The second is a so-called Factor-Augmented Random Walk (FARW) model. It is a no-change forecast adjusted by the linear combination of the same factors in our LR forecast. The idea of doing this is to test whether our LR-SR split has any implied economic meaning, as both models are based on the same information. If our forecast has less MSFE, it means that not only just the factors, but also the functional form we propose are useful in the forecast.

Table 6 shows the forecast error comparison among the three models with multiple commodities in both the agricultural and resource groups. The errors are presented in the mean-squared forecast error (MSFE). To test whether our model is better than the other two, we use the Diebold-Mariano (1995) testing technique with the uniform kernel. In agricultural commodities, our model performs best for coffee and cotton with significant MSFE reductions relative to both FARW and future-based models. For both commodities, the future-based forecasts are the worst of the three. The cost of storage is another likely cause of poor performance beyond the previously discussed convenience yield. For corn, there is no statistical advantage in modeling with the long-run trend in comparison with the unrestricted FARW model, despite the MSFE being slightly smaller for my model. For cocoa, the future-based forecast is marginally better, though the advantage is insignificant based on the Diebold-Mariano test. For resource commodities, our model is still the best among the three in general, although the edge seems smaller. For aluminum and crude oil, the differences are obvious but less so for copper and iron. Another observation from resource commodities is that the FARW model performs as well as future prices, so if without the long run restriction, forecasts based on factors only are less likely to perform better than future price predictors.

Another task before we conclude is to test whether other documented observable factors that are shown to have predictive power can further reduce the MSFE. If they can, that will be strong evidence that the forecast power is robust and those predictors contain information that is not available from the equity markets. If not, it means that the information from these predictor is likely to be already contained. To accomplish this task, we use US Consumer Price Index, US Producer Price Index, effective Federal Funds Rate, and exchange rates of CAD, AUD and NZD to USD as if they were one of the LR factors. To avoid not being selected as a factor by our algorithm, they are not subjected to the forecast pre-test. Then we compare the MSFE of this slightly modified model to our original model. We also test on the SR component where we add the past 12-month average to the AR model. Table 7 shows the improvement check results. While some of those have modest forecast power on the SR components (especially for oil and copper), they cannot further improve on the long run component. It is therefore safe to conclude that their predictive power is mostly contained in our LR factors, and the commodities do share the information with equity markets, at least in the long run.

II.6 Conclusion

In this paper we propose an unified 2-component forecasting method and use it to forecast spot returns of many commodities. The LR components are forecasted by unit-root latent LR factors that are extracted from the high-dimensional stock market data and are selected by bagging. The forecast of SR components is done by augmenting a no-change forecast with past 12-month average SR returns. Therefore, LR captures the long run cointegration relations, while the SR is modeled as transient deviations from the long run. The two-component approach has the superior forecast performance for both agricultural and resource commodities, as reflected by up to a 40% reduction in the mean-square forecast error in comparison with naive benchmark models. In addition, once the equity market LR information is taken into account, other macro variables that are shown to have predictive power cannot further improve on the performance. Also, as our model is forecasting at level, no detrending of the data is necessary. Compared with stationary forecasting methods, forecasted numbers from our model are easier to interpret as they are the predicted values of future raw data before any in-sample filtering. Of course, the framework is not only for forecasting commodities and its suitability for other economics and finance topics is left for future research.

Table II.1: Monte Carlo Simulations

Case	Method	Three Factors		Six Factors	
		MSFE	$\frac{Std(\tilde{P}_t)}{Std(P_t)}$	MSFE	$\frac{Std(\tilde{P}_t)}{Std(P_t)}$
1.Standard	Trend-Predictor	1.743	106%	1.211	102%
	Random Walk	1.848	100%	1.881	100%
	HP-Filtered Trend	1.919	45%	2.307	39%
2. Lagged Factors in X_t	Trend-Predictor	1.876	69%	1.671	73%
	Random Walk	1.921	100%	1.809	100%
	HP-Filtered Trend	1.861	56%	1.889	56%
3. P_t with More I(1)	Trend-Predictor	2.295	86%	1.747	87%
	Random Walk	2.108	100%	1.914	100%
	HP-Filtered Trend	2.231	85%	2.163	81%
4. Noisy P_t	Trend-Predictor	3.010	44%	1.503	52%
	Random Walk	4.348	100%	2.171	100%
	HP-Filtered Trend	3.110	25%	1.493	28%

Table II.2: Short-Run Cycle Forecast-Model Comparison: Agricultural Commodities

Commodity	Model Specifications			
	Vol.Shock	Uncer.Aver	AR(1)	GARCH(1,1)
Agricultural Commodities:				
Apples	0.28**	0.23	-0.68**	0.22
	(0.13)	(0.31)	(0.15)	(0.85)
Butter	0.18	0.40	-0.59**	0.41
	(0.33)	(0.30)	(0.21)	(0.25)
Cocoa	0.02	-0.3	-0.52**	0.25
	(0.07)	(0.31)	(0.21)	(0.19)
Coffee	0.22*	0.09	-0.56**	-0.21*
	(0.13)	(0.39)	(0.15)	(0.12)
Corn	-0.11	0.31	-0.69**	0.09
	(0.11)	(0.35)	(0.12)	(0.21)
Cotton	0.01	-0.11	-0.51**	0.12
	(0.06)	(0.82)	(0.17)	(0.11)
Eggs	-0.12	-0.90**	-0.39**	0.54
	(0.19)	(0.41)	(0.08)	(0.45)
Sugar	0.08	-0.38	-0.64**	0.61**
	(0.07)	(0.37)	(0.18)	(0.29)

Specifications:

Vol.Shock: Regressing current year SR cycle on the change in realized volatility in the past 12 months: $r_t^{SR} = \beta_0 + \beta \sum_{i=0}^{11} |\sigma_{t-i}^2 - \sigma_{t-i-1}^2| + e_t$.

Uncer.Aver: Regressing current year SR cycle on the level of realized uncertainty in the past 12 months: $r_t^{SR} = \beta_0 + \beta (\sum_{i=1}^{12} \text{Uncertainty}_{t-i}) + e_t$.

AR(1): Regressing current year SR cycle on the average level of SR cycle in the past 12 months. $r_t^{SR} = \beta_0 + \beta (\frac{1}{12} \sum_{i=1}^{12} r_{t-i}^{SR}) + e_t$.

GARCH(1,1): Regressing current period SR volatility on the level of SR volatility in the past month. $\sigma_t^{SR} = \beta_0 + \beta \sigma_{t-1}^{SR} + e_t$.

Table II.3: Short-Run Cycle Forecast-Model Comparison: Resource Commodities

Commodity	Model Specifications			
	Vol.Shock	Uncer.Aver	AR(1)	GARCH(1,1)
Resource Commodities:				
Aluminum	0.08	-0.9	-0.54**	0.22*
	(0.16)	(0.73)	(0.08)	(0.14)
Copper	0.12	0.21	-0.55**	0.35**
	(0.13)	(0.31)	(0.15)	(0.17)
Crude Oil	0.12*	-0.17*	-0.46**	0.22
	(0.07)	(0.25)	(0.16)	(0.18)
Iron	0.04	-0.88	-0.72**	0.88**
	(0.19)	(0.74)	(0.17)	(0.18)
Lead	-0.03	-0.36	-0.55**	0.23
	(0.16)	(0.89)	(0.17)	(0.18)
Tin	0.13	0.08	-0.60**	0.03
	(0.11)	(0.21)	(0.13)	(0.15)
Zinc	-0.07	-0.25	-0.52**	0.48**
	(0.15)	(0.40)	(0.21)	(0.18)

Specifications:

Vol.Shock: Regressing current year SR cycle on the change in realized volatility in the past 12 months: $r_t^{SR} = \beta_0 + \beta \sum_{i=0}^{11} |\sigma_{t-i}^2 - \sigma_{t-i-1}^2| + e_t$.

Uncer.Aver: Regressing current year SR cycle on the level of realized uncertainty in the past 12 months: $r_t^{SR} = \beta_0 + \beta (\sum_{i=1}^{12} \text{Uncertainty}_{t-i}) + e_t$.

AR(1): Regressing current year SR cycle on the average level of SR cycle in the past 12 months. $r_t^{SR} = \beta_0 + \beta (\frac{1}{12} \sum_{i=1}^{12} r_{t-i}^{SR}) + e_t$.

GARCH(1,1): Regressing current period SR volatility on the level of SR volatility in the past month. $\sigma_t^{SR} = \beta_0 + \beta \sigma_{t-1}^{SR} + e_t$.

Table II.4: Short-Run Cycle Forecast Performance

Commodity	Forecast Coefficients		Forecast Error in MSFE	
	Lag Predictor	His. Return	Full Spec.	RW
Agricultural Commodities:				
Apples	0.74**	-0.23**	40.03	44.89
	(0.04)	(0.13)	89.17%	
Butter	0.81**	-0.25**	16.52	18.05
	(0.06)	(0.07)	91.52%	
Cocoa	0.86**	-0.18*	18.59	20.20
	(0.05)	(0.12)	92.03%	
Coffee	0.88**	-0.16**	24.29	25.92
	(0.03)	(0.05)	93.71%	
Corn	0.88**	-0.12**	9.96	10.52
	(0.03)	(0.06)	94.7%	
Cotton	0.86**	-0.27**	10.67	11.80
	(0.03)	(0.05)	90.42%	
Eggs	0.71**	-0.25**	26.03	30.31
	(0.04)	(0.11)	85.88%	
Sugar	0.87**	-0.30**	50.17	56.09
	(0.06)	(0.07)	88.81%	
Resource Commodities:				
Aluminum	0.83**	-0.23**	5.42	6.02
	(0.03)	(0.06)	90.03%	
Copper	0.82**	-0.23**	9.31	10.23
	(0.09)	(0.08)	91.01%	
Crude Oil	0.79**	-0.20**	12.19	13.97
	(0.03)	(0.10)	87.26%	
Iron	0.65**	-0.34**	22.71	26.99
	(0.06)	(0.15)	84.14%	
Lead	0.81**	-0.19**	7.01	7.88
	(0.03)	(0.05)	88.96%	
Tin	0.83**	-0.14*	7.74	8.46
	(0.02)	(0.08)	91.49%	
Zinc	0.88**	-0.18**	7.72	8.18
	(0.03)	(0.06)	94.37%	

Table II.5: 1-Month Ahead Forecast Error Comparison

Commodity	Forecast Errors (MSFE) [a]			Diebold-Mariano Test	
	1.LR+SR Forecast	2. FA-Random Walk	3. Future Price	1 over 2?	1 over 3?
Agricultural Commodities:					
Corn	88.4%	93.0%	100%	No	Yes
Coffee	59.5%	85.9%	100%	Yes	Yes
Cocoa	100.3%	124.3%	100%	Yes	No
Cotton	42.1%	61.1%	100%	Yes	Yes
Resource Commodities:					
Aluminum	75.5%	97.3%	100%	Yes	Yes
Crude Oil	79.1%	101%	100%	Yes	Yes
Copper	91.9%	104.7%	100%	Yes	No
Iron	103.3%	99.9%	100%	No	No

[a] Note: All MSFEs are normalized based on future price predictors.

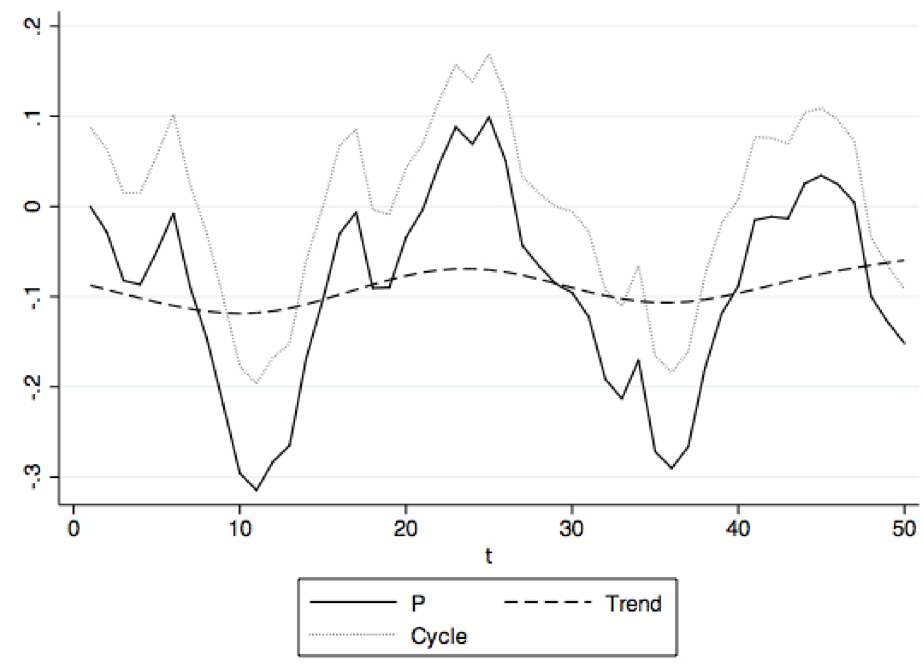

Figure II.1: Simulated HP Trend and Cycle with Standard Dynamic Factors

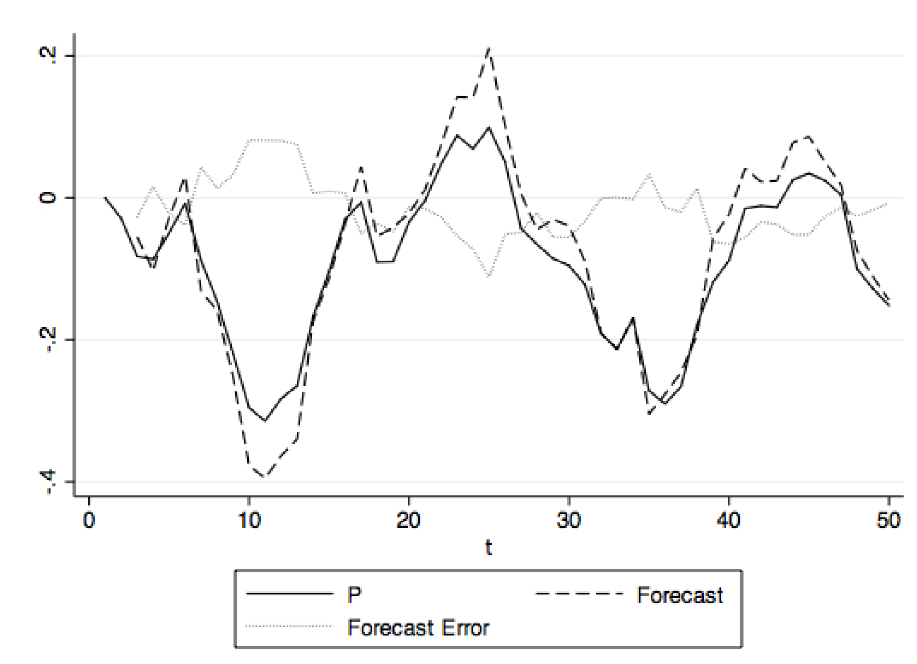

Figure II.2: Simulated One-Step Ahead Forecasted Trend with Standard Dynamic Factors

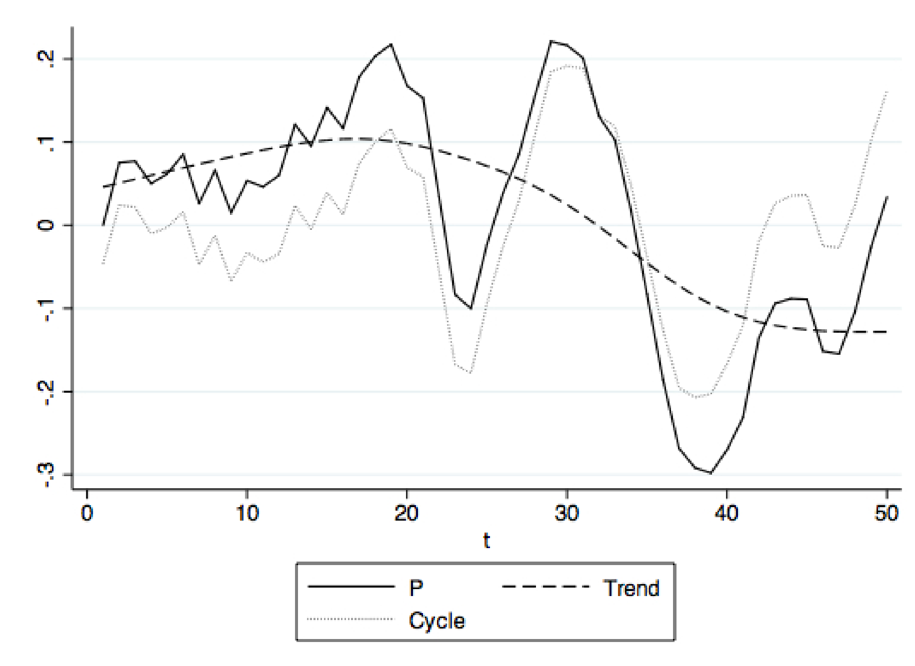

Figure II.3: Simulated HP Trend and Cycle with Lagged Factors

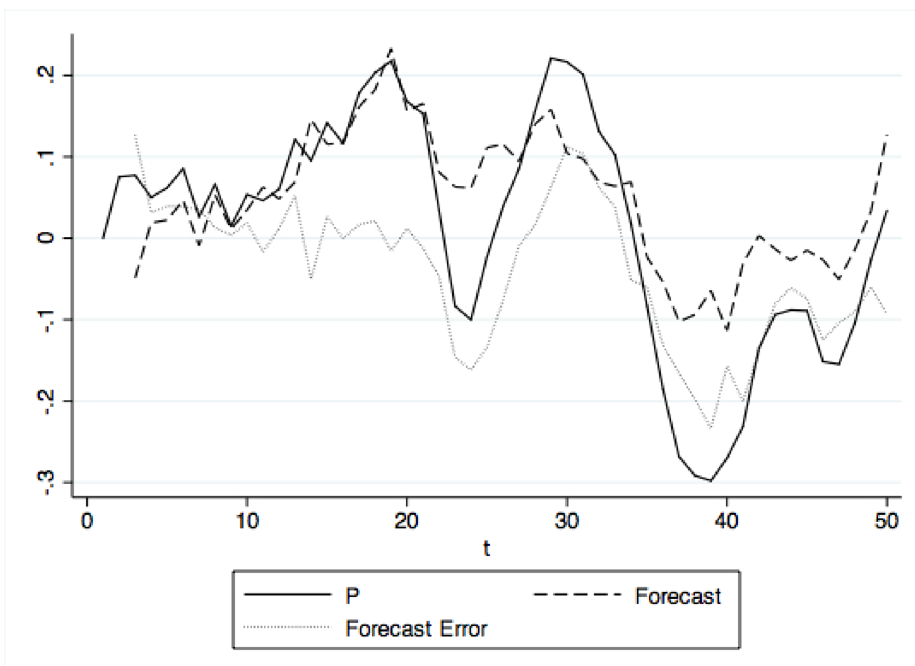

Figure II.4: Simulated One-Step Ahead Forecasted Trend with Lagged Factors

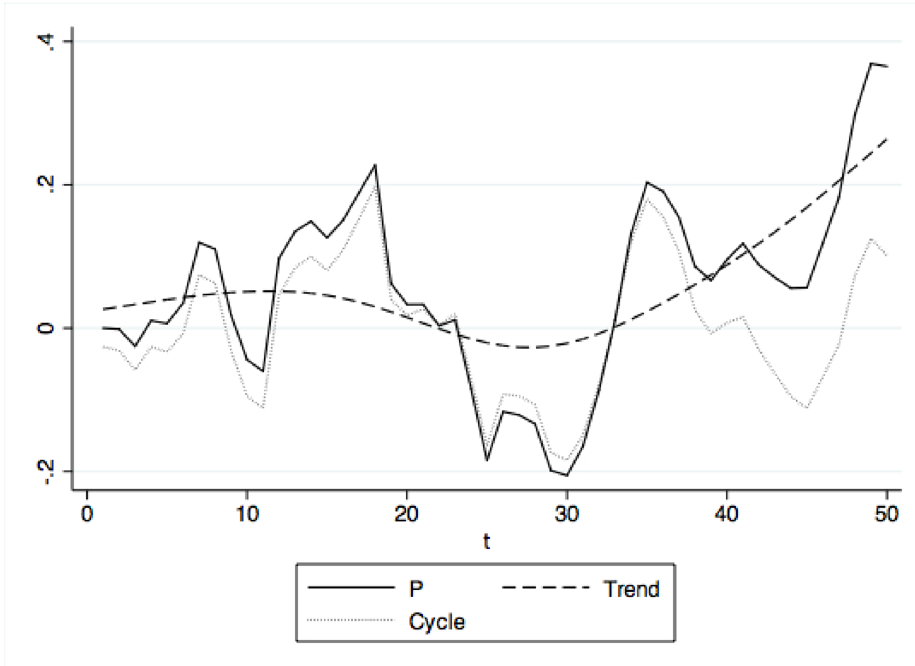

Figure II.5: Simulated HP Trend and Cycle with P Containing Outside Information

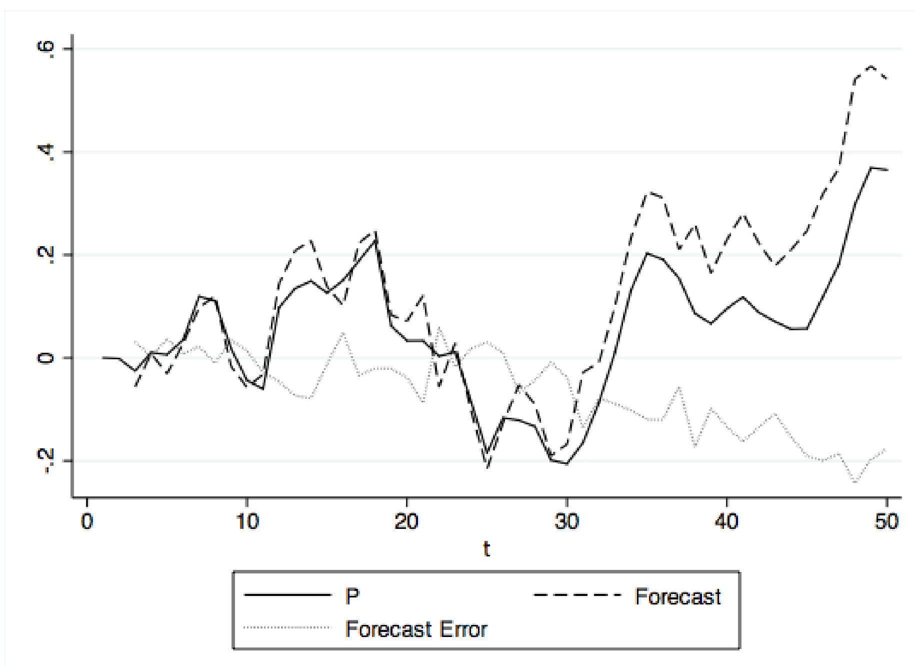

Figure II.6: Simulated One-Step Ahead Forecasted Trend with P Containing Outside Information

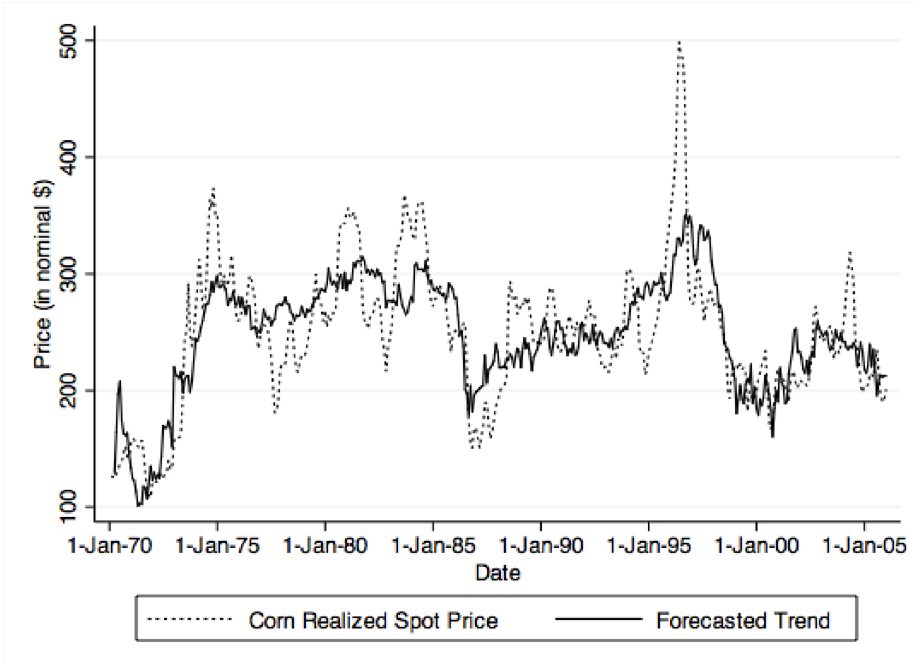

Figure II.7: One Month Ahead Forecasted Trend in Corn

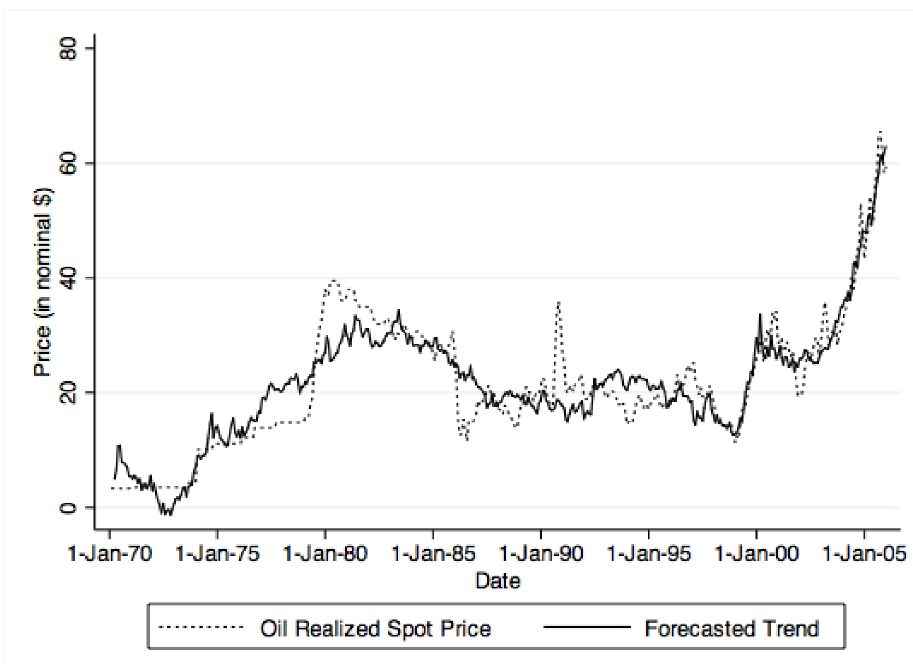

Figure II.8: One Month Ahead Forecasted Trend in Crude Oil

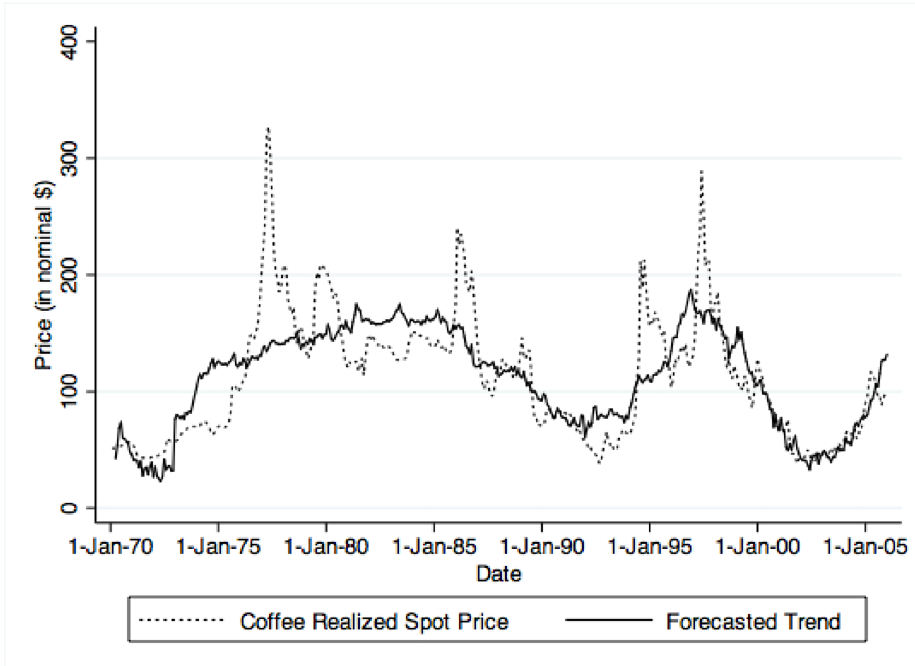

Figure II.9: One Month Ahead Forecasted Trend in Coffee

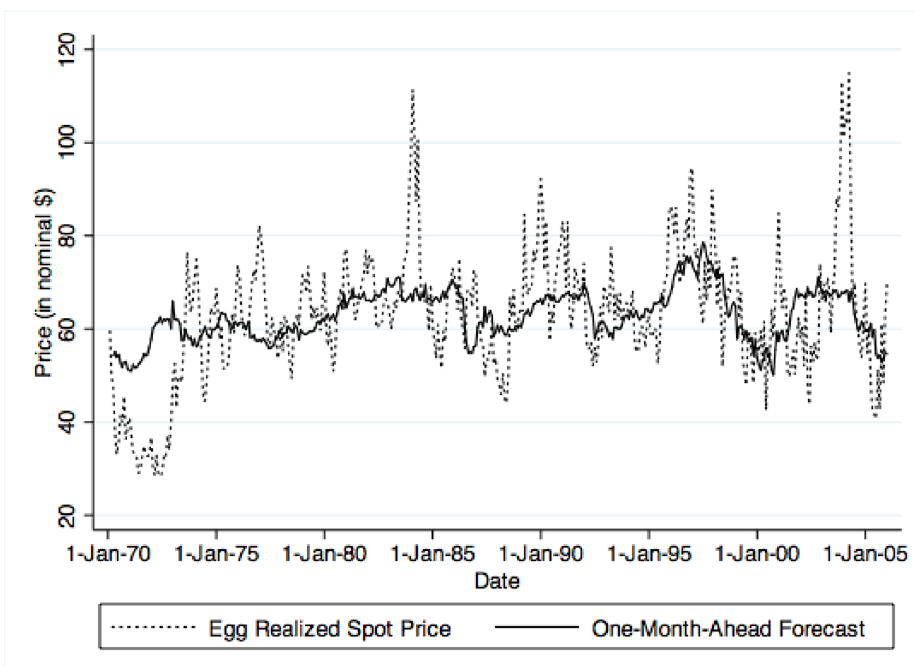

Figure II.10: One Month Ahead Forecasted Trend in Eggs

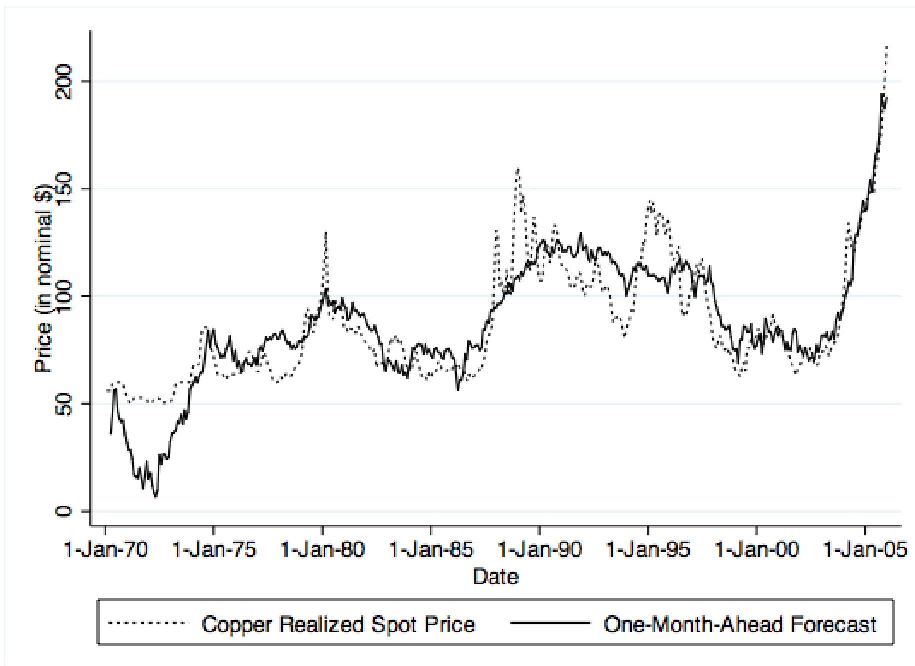

Figure II.11: One Month Ahead Forecasted Trend in Copper

CHAPTER III

Preference for Non-stop Flights and Frequent Flyer Programs: New Empirical Evidence

III.1 Introduction

The US aviation industry has experienced tremendous changes in the past decades. On one hand, legacy carriers have gone through financial distress: all major airlines have filed for bankruptcy since 2000. At the same time, multiple M&A deals are done for the sake of competition reduction and cost saving. As a result, only three legacy carriers still exist as of 2015: American Airlines (merged with US Airways in 2015), Delta Airlines (acquired Northwest Airlines in 2008) and United Airlines (merged with Continental Airlines in 2010). After those major restructuring and consolidations, the remaining carriers were able to come out of bankruptcy protection and started to post profit in 2013. On the other hand, we have observed rapid growth of those so-called low cost carriers (LCCs) during the same period of time. Take Southwest Airlines as an example: it has grown from a small regional airline in the 80's into the biggest airline in the domestic market by passenger volume. Other LCCs like JetBlue and Frontier have also expanded into many new markets. The end result of such industry reshaping is a market consisting of only a few very large legacy carriers and smaller LCCs.

We are interested in the competition of legacy carriers vs LCCs because of the following reasons. Firstly, the route maps look very different between the two groups: All legacy carriers have adopted the "hub-and-spoke" system in which there are usually a few giant hub airports where all non-hub origin/destination passengers have to connect through. Those hubs are usually located close to major metro areas so that customers from those cities don't have to connect. LCCs on the other hand offer non-stop flights to more secondary markets since they don't have the hub system. Due to the consolidation of legacy carriers, the number of hub airports has been shrinking. Therefore, there are less people living in the hub cities and more people in the now secondary markets where they often have to choose between non-stop flights from the LCCs and one-stop flights from the legacy carriers. How do customers in those particular markets choose from those differentiated products has not been studied to my best of knowledge.

Secondly, there is one significant feature of legacy carriers that is not shared by LCCs: a comprehensive Frequent Flyer Program (FFP) with elite status. A FFP consists of two currencies that flyers can earn by flying: the elite qualifying miles and the redeemable miles. The elite qualifying miles are used to determine the status of each passenger: for all legacy carriers, you need to fly at least 25000 miles on paid tickets to qualify for their basic elite status and much more (50000, 75000, 100000 miles etc) to reach the higher tiers. Being an "elite" flyer gives perks like free upgrade to first class and earning redeemable miles at a faster speed. The redeemable miles can be used for free flights around the world at preset prices (e.g. US to Asia requires 35000 miles on AA in the economy cabin, regardless of cash prices). While LCCs do have FFPs, they are much simpler in nature and to the point that they are essentially derivatives of the cash-back system. For example, Southwest gives passengers back 6% of their ticket value in the form of "miles" and those miles can be redeemed at around 1.4 cent per mile for future tickets. Ultimately, it is a 8.4% rebate for everyone. Elite status on LCCs also comes with less perks. The difference is particularly important when analyzing the secondary markets, as frequent flyers having status with legacy carriers need to forgo their privileges to switch to non-stop flights offered by LCCs. The resulting reluctance may be more obvious when LCCs try to penetrate into new markets where legacy carriers are the only incumbents for a long time.

The main contribution of this chapter is to reveal the choices of customers in those aforementioned mid-sized US markets with particular emphasis on carrier (legacy vs LCC) and route combinations (connecting vs non-stop). We also augment the standard demographic data with market-level elite passenger statistics from legacy carriers. It has been documented that in large metropolitan cities, customers have developed preference for non-stop flights (Berry and Jia 2009). However since large cities are likely to be hubs of at least one legacy carriers, results may not be representative for customers who have been very used to one-stop itineraries. Indeed, our result shows that in those secondary markets, while customers still prefer non-stop flights, the likelihood of switching to non-stop flights offered by LCCs is much lower than previously estimated. The FFPs play a critical role here, as the model indicates that frequent flyers are less likely to switch than the non-elite passengers. This is not true for hub passengers as switching to non-stop flights does not necessarily mean switching carriers.

The rest of the paper is structured as follows: Section 2 provides a literature survey with special focus on the existing popular methodologies, as technical innovations are not the goal in this chapter. Section 3 discusses the new empirical evidences coming from our new data and identification. Section 4 concludes.

III.2 Literature Review

One of the earlier studies on airline competitive pricing is from Morrison and Winston (1990). In this short paper the authors first show us that the de-regulation in 1978 actually had decreased the yields to airlines on average. De-regulation means that airlines were free to choose their own fares instead of having to follow fares set by the Civil Aeronautics Board. In addition competitive requirements for airports were eliminated so individual airlines could control more departure/arrival slots in any single airport than before. This change of policy greatly altered the shape of the industry at that time as all legacy carriers started to build up their route systems under the hub-and-spoke system. Such de-regulation obviously would raise concern about air travel becoming more expensive. The authors construct counterfactual fares by using the Standard Industry Fare Level and show that on average de-regulated fares were 18% lower than regulated fares would have been and saved travelers six billion 1988 dollars annually. The fare regression shows that the number of competitors is negatively correlated with fares on that specific route. Since the average number of competitors has risen from 1.52 to 1.9 at the route level, increased competition on the "average" route can explain the decrease in the average fare. Finally a simple probit model is applied for entry-and-exit analysis for American Airlines and the authors find out that existing activities in origin/destination airports, together with relative fares, are all important factors for exit-entry decisions.

While this paper serves as a fine starting point for looking at the industry at an earlier stage, the model part is not accurate enough for a current study: First of all, the saving figure is obtained with the assumption that demand is perfectly inelastic. Such an assumption is unrealistic and hence has greatly exaggerated the possible saving. Secondly, while overall route level competition may have indeed increased, for hub airports where one airline could control a majority of activities, fares can be substantially higher than in the regulated case as discussed in the hub premium literature. Since hubs are mostly within densely populated metro areas, the total fares paid by all customers as a whole may actually increase rather than decrease as claimed in the paper. Finally, the probit model has overly simplified the real world case and is subject to common endogeneity problems.

Another descriptive paper by Borenstein (2005) has nice updates from more recent data on airport-specific fares. From their data it is observed that airline prices overall have been declining but at the same time the hub premium is confirmed: there are hubs where the average prices are significantly higher than the national average, though some of those premiums are also declining largely due to aggressive entries of LCCs during this period. Yet for hubs without such competition, fares remain high. The reduction in air fares in general is believed to be generated from lower cost as modern regional fleet entering service and the use of computers and online booking tools. Strategic bankruptcies combined with mergers have lowered labor/fixed cost as well. But since technological shocks are industry wide and all legacy carriers have gone through corporate re-structuring, relative competitiveness should not be affected as much. Yet we can see that market shares vary greatly in 2004 from a decade ago. Such dynamics in market shares cannot be explained by simple regression models.

To better explain the hub and some particular airport premium as illustrated in the above paper, Berry (1990) suggests that simple models may not be conclusive, as both simple cost-reducing and naive market-power stories are inappropriate for the airline industry. He suspects that product differentiation in those dominated airports is the reason why hub premiums exist. To support his argument, he constructs a structural model in an oligopolistic differentiated products framework. He introduces unobservable components into consumer preference and market products, hence making the demand heterogeneous. Then the demand function in this discrete choice model can be expressed in the simplest form as:

$$U_{ij} = v_{ij} - \alpha p_j \tag{III.1}$$

where v_{ij} is assumed to be exponential with mean μ_j and $\mu_j = exp(x\beta + \xi_j)$ and ξ_j is the market unobservable variable. Together with the price assumption that

$$p_j = mc_j + \mu_j/\alpha \tag{III.2}$$

where mc_j is the marginal cost which is assumed to be constant for simplicity. If we assume for ξ_j some density function, we can compute out this discrete choice model. The author does this by assuming a normal distribution. His results show that whether airlines choose an airport as a hub depends on the demand characteristics and dominance may come from both supply side (cost reduction) and demand side (high demand for premium fare). This is somewhat counter intuitive, as the model suggests that customers in those markets are willing to pay a premium if any airline can dominate in their airport and airlines are willing to dominate in those markets since by doing so will bring down the cost. I suspect that such results are due to the less than ideally identified demand function, as many critical factors are not included. By specifying the demand function more accurately in a parsimonious fashion we can increase model accuracy but at the cost of lengthy computing time and the possibility of no analytical solution of market share if parameters are not carefully chosen.

Berry (1992) does follow up with a seminal paper on estimating a specific model of entry. In this paper the author, following the preliminary results as above, investigates the importance of airport presence in determining the profits of operating in a given city pair. The analysis framework is a simple two-stage oligopolistic game in each market i. In the first stage, firms choose to be "in" or "out" of the market and in the second stage, the "in" firms play some unspecified game that determines post-entry profits. Such profits are not the main concern of this paper. If there are no unobserved factors for firms, a Nash equilibrium strategy vector s^* shall satisfy:

$$s_k^* \pi_{ik}(s^*) \geq 0 \tag{III.3}$$

and

$$(1 - s_k^*)\pi_{ik}(s^{*+k}) \leq 0. \tag{III.4}$$

It is necessary to impose restrictions on profit function $\pi_{ik}()$ so to have a unique pure strategy NE while acknowledging the heterogeneity of firms. The author proves that if the profit function takes the form

$$\pi_{ik}(s) = v_i(N(s)) + \phi_{ik} \tag{III.5}$$

where $v(N)$ is strictly decreasing, then there is a unique market equilibrium with N^* entering firms. The expression ϕ_{ik} is used to capture firm heterogeneity and can be treated as firm specific fixed cost. For estimation purposes, market common characteristics and firm specific characteristics are specified in this model as:

$$v_i(N) = X_i\beta - \delta ln(N) + \rho u_{io} \tag{III.6}$$

$$\phi_{ik} = Z_{ik}\alpha + \sigma u_{ik} \tag{III.7}$$

where X_i is vector of common market characteristics and Z_{ik} is firm specific characteristics. In order for the unobserved component to have a variance of 1, we also have to restrict errors with $\sigma = \sqrt{1-\rho^2}$. Hence the model is complete and fully specified. The author then proceeds to identify the distribution of the equilibrium number of firms. However since entry decisions are dependent on both the kind of firms that are incumbent at the time and those that are trying to enter, additional assumptions are needed. To get over this issue author first puts the model in special cases (the most profitable firms move first, for example) and also uses simulation estimators. For simulation estimators, the idea is straightforward. First, we can always define a "prediction error":

$$v_{io}(N_i^*, W_i, \theta) = N_i^* - E[N^*/W_i, \theta] \tag{III.8}$$

The term W_i is the exogenous market data and θ is the vector of parameters to be estimated. So by construction, when θ is at its true value θ^*:

$$E[v_{io}(N_i^*, W_i, \theta)/W_i, \theta = \theta^*] = 0 \tag{III.9}$$

the number of firms can be calculated as

$$N_i^* = \hat{N}(W_i, \theta, \hat{\mu}_i) + \hat{v_{io}} \tag{III.10}$$

where

$$\hat{N}(W_i, \theta, \hat{\mu}_i) = \frac{1}{T}\sum_{t=1}^{T} \hat{n}(W_i, \theta, \hat{\mu}_i) \tag{III.11}$$

and $\hat{n}$ are estimators from profit functions from various markets.

Finally the author makes use of this estimating technology and contrasts with simple estimators (OLS) and MLE estimators (Probit) under special cases. The general consensus among different streams of models is that it is better to serve more cities than focusing only on a few while there is still a substantial effect of airport presence, especially in the full structural model. Increasing airport share will increase profits from directional markets both originating from and arriving at the city. The full structural model is arguably better since it can capture policy-relevant effects (like controlling one airport brings bureaucratic benefits to incumbents), which simple models cannot explain.

Berry, Carnall and Spiller (1996) on the other hand have targeted the hub premium as revealed in data. They extend the story in Berry (1990) by looking specifically into cost, markups and implied customer heterogeneity. Their model is also a nice extension of previous simulated discrete choice models. One major contribution in this updated model is that they have been able to explicitly solve for market shares from discrete choices of individual customers.

In the BCS setup, the utility function now takes the form

$$\mu_{ij} = x_j\beta_i - \alpha_i p_j + \xi_j + v_i(\lambda) + \lambda\varepsilon_{ij} \tag{III.12}$$

where x_j is a vector of observed product characteristics, p_j is the price, ξ_j captures the unobserved product characteristics and the additive error $v_i(\lambda) + \lambda \varepsilon_{ij}$ gives the utility a "nested logit" feature. The distribution of this sum is parameterized by λ and can be estimated together with other coefficients. The idea behind this model is that a certain consumer type i may prefer a certain group of similar products, hence the correlations of certain markets. If λ goes to 1, $v_i(\lambda)$ goes to 0, then there is no such correlation. Within any group, such added error would cancel out since it does not vary among products and for consumer type 1, the within market share is expressed in logit term:

$$s^1_{j1g} = \frac{exp[(x_j\beta_1 - \alpha_1 p_j + \xi_j)/\lambda]}{\sum_j exp[(x_j\beta_1 - \alpha_1 p_j + \xi_j)/\lambda]} \tag{III.13}$$

and the group share is of logit form as well:

$$\bar{s_1} = \frac{\{\sum_j exp[(x_j\beta_1 - \alpha_1 p_j + \xi_j)/\lambda]\}^\lambda}{1 + \{\sum_j exp[(x_j\beta_1 - \alpha_1 p_j + \xi_j)/\lambda]\}^\lambda}. \tag{III.14}$$

So assuming only two types of customers, the market share of product j is just the weighted sum:

$$s_j = \gamma s^1_j(p,x,\xi) + (1-\gamma)s^2_j(p,x,\xi) \tag{III.15}$$

The term γ may or may not be revealed by data.

The profit of firm f in each market m that produces in the set $J(f,m)$ is given as

$$\pi_f = \sum_m \sum_j q_{jm}(p)p_{jm} - \sum_s C(Q_s, w_s) - \sum_m \sum_j \omega_{jm} q_{jm}(p) - FC_f. \tag{III.16}$$

The observed cost term $C(Q_s, w_s)$ is a function of spoke density Q_s and ω is some random error observable upon realization. The first order condition can be derived as

$$\frac{\partial \pi_f}{\partial p_{jm}} = q_{jm} + \sum_k \frac{\partial q_k}{\partial p_{jm}}(p_k - mc_k) = 0 \tag{III.17}$$

which completes the model. The estimation strategy is similar to Berry's algorithm. The results show that the nested logit demand function is a better fit as compared to the standard logit model. The hub airline charges higher prices for travelers originated from hub airports. Such a strategy is particularly effective in consumer types with higher demand for non-stop flights and less travel time. Marginal cost corresponds to flight density in a non-linear, hump-shaped fashion and it is related with route distances as well. Exit and entry decisions are based on the relative marginal cost that the firm faces.

We can see from the evolution of the models that researchers have benefited from the development of simulation estimators. The main topics can be categorized into two strands: the first general topic is about pricing and trying to understand why in some airport/airline combinations the average fares are higher than the rest. Usually most attention is on hubs since they are where most high fares are from. The second strand focuses on exit and entry decisions of airlines in particular markets. Earlier studies tend to treat the two aspects separately. Such analyses are valid only under rather restrictive conditions, as many crucial factors must be set as exogenous by not considering pricing and exit/entry decision simultaneously. The structural models, on the other hand, can relax many restrictive assumptions and are therefore more realistic. However, the caveat is that if the algorithm is not carefully written, estimating structural models can be very time consuming, even with advanced computational power. Also, the distributions of unobserved errors and characteristic terms need to behave nicely to ensure closed form solutions to key variables. As a consequence, assumptions on distributions are not necessarily based on real world observations. But such idealization is minor in comparison with classic models.

There are some recent updates on this topic as well. Berry and Jia (2009) employ the same methodology to analyze recent data and find out that customers' preference has shifted toward non-stop flights while being more price sensitive than before for both business and leisure travelers. On the opposite side, Gayle (2004) suspects that price may not be a crucial factor in customers' choices. Aguirregabiria and Ho (2012) in the mean time use a time series state space model to check on the exit and entry oligopoly game and concludes that the adoption of hub-and-spoke networks is because the sunk cost of entry in a route declines with the number of cities that airlines connect from the origin and destination airports of that route. This conclusion somewhat coincides with Berry's earlier finding of cost saving in hubs. Alderighi (2010) instead looks at fare dispersions and finds that in the duopoly, the max fare dispersion is rather limited regardless of models adopted. Such dispersion is a direct measure of price differentiation so it is less needed for airlines to differentiate customers when competition is not keen.

III.3 New Empirical Evidence

In this section we are going to build upon existing findings as discussed above and focus on the mid-sized markets. Methodology innovation is not the concern in this paper. Instead, I try to construct a new market level consumer-side dataset and target on the question that has not been studied but is obviously important: How will LCCs' decisions to offer new direct flights affect legacy carriers' competitive margin in those secondary markets? Will such entries be as disastrous to legacy carriers as some models have predicted? Berry and Jia (2008) predict that up to 80% of the market share of connecting flights would be shifted to the LCC offering non-stop flights. While the number certainly sounds substantial, it may be too big in our case: In their dataset they have only selected airports in densely populated cities. Such airports are more likely to be hubs for legacy carriers. Therefore, while a new non-stop flight is more convenient, most of the passengers that will switch are likely from connecting flights offered by the same carrier, since it is always easier to change within the carrier. For legacy carriers, such product cannibalization is less interesting, as the added revenue would be from only the premium of non-stop flights instead of coming from new passengers. It is not the case in non-hub city pairs, since legacy carriers are much less likely to offer non-stop products and those products are more costly to operate (Aguirregabiria and Ho 2012).

Another important customer characteristic that can affect purchase decisions are the “loyalty” to specific airlines created by FFP. The FFPs have proven to be a successful business innovation. In addition to the miles flyers earn that can be redeemed for award flights, if customers fly frequent enough with a specific carrier each year, they will be categorized as “elite” flyers and will be entitled to certain perks. Currently all US legacy carriers have this feature in their FFPs. The yearly re-qualification mechanism can hold frequent flyers back from flying other carriers as perks cannot be transferred to other airlines and “occasional” switching makes it harder to meet the re-qualification requirements. Therefore, an elite flyer affiliated with Delta Airlines can behave very differently from a passenger who always purchases cheapest fares. If elite passengers are indeed very loyal to the airline, a new non-stop flight from, say, Southwest, may not lure away that many passengers as in the scenario of no FFPs.

Time of booking on the other hand, measures the profitability of each customer in that specific itinerary. Since it is generally true that the more advance that a booking is made, it is more likely to be booked into deeply discounted fares. However such customers are worth less to airlines than those who travel on higher fares. So if premium passengers react differently to the offering of non-stop flights than discounted passengers, airlines should react more to those who buy more expensive tickets. Even though it is not directly observable in the data, we will try to partially identify it with instruments that are available.

III.3.1 The Model

As this study is to apply existing econometric methods to the new question, my full model will follow common structural setup as in Berry, Levinsohn and Pakes (1995) and Petrin (2002). But before simulating the full model, we can have a preliminary look at the data by estimating the consistent fixed-effect panel model:

$$log(S_{ijt}) - log(S_{0t}) = x_{ijt}\beta - \alpha p_{ijt} + \xi_{it} + \Delta\varepsilon_{ijt} \tag{III.18}$$

where S_{ijt} stands for the share of product i in market j in time t. S_{0t} is the share of "outside good". Such share is constructed in the fashion of "proportional expansion": It is the additional capacity invented for that route if the departure airport's average load factor were 1 while IIA assumption among shares of departure slots of different routes holds. By this construction we have a good approximation to the total market size as required for identification. Vector x_{ijt} here is the product characteristic vector which contains information like miles flown, number of stops, fare class, etc. Fare paid is denoted by p_{ijt}. The airline-specific unobserved effect is captured by ξ_{it}. The advantage of this fixed-effect estimation is that the computation is easy and we can have instant preliminary estimates on distribution of different types of customers. Disadvantages, though, are that since market/airline effects are absorbed we will not be able to see it directly. In addition, the substitution elasticities implied are implausible, similar to those in the standard multinomial logit models. Nonetheless it is a useful starting point.

The full model, which will improve upon the FE model, starts with a latent utility function:

$$u_{ijt} = \beta_{it}x_{jt} - \alpha_{it}p_{jt} + \xi_{jt} + \varepsilon_{ijt} \tag{III.19}$$

where u_{ijt} is the utility of consumer type i purchasing product j in market t. The vector x_{jt} contains the product characters of product j in market t and ξ_{jt} is the unobserved product characteristics. Usually, ξ_{jt} will be further decomposed into market and brand specific effects, but since my goal is to specifically inspect the airline effects for frequent and premium travelers, I will put the airline dummies into x_{jt} to interact with demographics once my full set of demographic data is constructed. The taste coefficients, $(\beta_{it}, \alpha_{it})$ is modeled as:

$$\begin{pmatrix} \alpha_i \\ \beta_i \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} + \Pi D_i + \Sigma v_i \tag{III.20}$$

where $\binom{\alpha}{\beta}$ is the average taste that corresponds to the FE model estimation. Observed demographic characteristics are captured in D_i and so its density distribution, $\widehat{P_D}(D)$, either parametric or nonparametric, is known. However, the density of v_i, i.e. the unobserved demographics, is unknown and need to be parametrically specified before solving the share of product j in market t:

$$s_{jt} = \int_{A_{jt}} dP_\varepsilon(\varepsilon) dP_v(v) d\widehat{P_D}(D) \tag{III.21}$$

where A_{jt} is the mass of consumers that choose product j in market t. Such shares have no analytical solution in general but can be estimated by simulating ε, v and D.

Following BLP (1995) and Berry and Jia (2008), the markups are computed from demand data and parameters and are denoted as $b(s_t, x_t, p_t, \theta_d)$. The marginal cost of product j in market t is simplified as:

$$mc_{jt} = w_{jt}\psi + \omega_{jt} = p_{jt} - b_{jt} \tag{III.22}$$

where w_{jt} is considered as cost-shifters. The unobserved part, ω_{jt} shall be instrumented by exogenous supply and demand side shifters. Equation (22) closes the model.

III.3.2 Data

There are four major pieces of data used in this study: The first part is the Airline Origin and Destination Survey (DB1B) from Department of Transportation. It is a 10% random sample of domestic airline tickets by US carriers. The original dataset are split into three parts: coupon, market and ticket. The coupon section contains information on origin, destination, distance and fare class of each segment on itineraries. The market data consists of market fare, miles flown and nonstop miles for each complete one-way trip. I employed the itinerary id and market id for each reporting carrier that are consistent across the two dataset to merge with exact match. After merging, I have, in addition to the original market data, connecting city (if any), fare class and distance of each segment. Such construction allows us to get maximum information on product characteristics of all products offered on each city pair.

The supply side characteristics are computed from Air Carrier Statistics published by DOT as well which is a monthly data on aircraft type, available capacity and load factor for each segment. Based on origin airports the information are summarized to measure average operating cost of each carrier in each airport. The average operating cost per passenger should be decreasing with the efficiency of aircraft as well as the load factor and possibly have a non-linear relationship with total available capacity.

The demographic data are from the publicly available March Current Population Survey and a partial summary of the proprietary Delta Air Lines Origin-Destination Statistics. While March CPS provides general average demographic information for each metro area that is served by at least one airport, data directly from the carrier provides detailed passenger information on each segment like their frequent flyer status, fare class booked and fare class flown. Though it does not have the complete itinerary information, it provides a great way of distinguishing customers at the market-level. For example, 14% of the passengers that have flown out of Nashville on Delta are at least their silver tier frequent fliers in 2012. By comparing to the 21% of Boston, the Delta data greatly enhances our ability to identify behavior heterogeneities among passengers with and without status. As of now, I have not been able to get similar data from other legacy carriers.

It is worth noting that the proprietary data is not required for this study and an alternative method of collecting for a similar dataset is also used: Since Delta gives unlimited upgrades to first class to all the elite flyers (the "medallions"), the number of elites is usually much larger than the number of available first class seats. In order to determine the priority of such upgrades, Delta lists all elite passengers and clears upgrades from top to bottom. Such lists are displayed both online and at departure gates. As we do not need to peg the status of each passenger to their tickets, the numbers collected are equally useful in our study. We acknowledge that the numbers collected this way do not coincide exactly with the official report, partly due to confirmed upgrades in advance as well as last minute changes. But since the correlation of the two dataset is very high, we believe that the survey data is a valid instrument for FFP status when the true values are not revealed.

III.3.3 New Empirical Results

Table 1 shows descriptive statistics from market level data and compares prices and shares of legacy and LCCs in selected representative routes (Boston to Nashville and Boston to San Jose). Due to the constraint on time and budget, the origin and destination airports are selected under the criteria that: 1. They are not hub airports for any of the legacy carriers and 2. They are within major metropolitan areas with top 50 total population in the United States. Based on those criteria, markets between nine qualified airports (Austin, Boston, Nashville, Las Vegas, Kansas City, Portland, San Diego, San Jose and St. Louis) are considered and there are a total of 22 markets that are served by newly offered non-stop flights from LCCs during 2010-2012. For example, JetBlue started the Boston to San Jose non-stop in May 2010 and Southwest started Boston to Nashville in early 2011. Both routes were only served by connecting flights before. If Berry and Jia's claim is applicable to those markets, we should observe huge increases in both shares and fares of the LCCs who offer the non-stop products. Numbers in Table 1 show otherwise. While we do see increase in shares for Southwest and JetBlue, such increase is quite moderate. The non-stop flights only give around 10% market shares to the firms and for the case of Southwest, it is at the cost of nearly 7% decrease in its connecting flight market share (product cannibalization). In addition, both non-stop flights offer among the lowest fares in their specific routes, implying that market shares may be gained by heavily discounted prices. What we see here offers an alternative story in those mid-sized markets that non-stop products are not as competitive as in bigger markets that Berry and Jia have focused on. The reasons may be that passengers are reluctant to switch airlines and in bigger airports where it is more likely that the same route is served by both non-stop and connecting flights from one single carrier. Thus switching from connecting to non-stop flights may not involve a carrier change. Such reluctance may come from FFPs. Also business travelers may have their preferred airline of travel from corporate policies. In general, it seems true that in those mid-sized airports, non-stop flights from LCCs do not have the pricing power and the ability to gain passengers as shown in legacy carriers' products from hub cities (Berry 1992).

Table 2 shows the classic OLS and FE estimations over the period 2007-2012 on a quarterly basis. The dependent variable, $log(S_{ijt}) - log(S_{0t})$, is restricted to connecting flights only and the non-stop fares are the average non-stop fares offered by LCCs on the same routes. The fixed effects are the market-specific effects and airline dummies are used to identify any differences in airlines. Results show that: Firstly, the prices for connecting flights alone are not that important in determining market shares. In fact the coefficients are significant in none of these specifications. But such shares are very sensitive to the non-stop fares in the same market. An increase in prices of non-stop flights causes the shares of all connecting flights to rise. That may explain why in these markets it is harder to charge non-stop premiums for LCCs. This is the opposite of the hub premium from legacy carriers where they can charge much higher prices for direct flights and passengers accept them for the convenience. Secondly, passengers do show preferences to non-stop flights if travel distances are long enough. They care less about connecting though, when they travel short-haul. Thirdly, the airline dummies play rather limited roles in this case without further demographic information, so as an average passenger it looks like she or he is indifferent among carriers.

Table 3 moves to the full structural model both without and with the survey demographic data in 2012. The estimation algorithm follows BLP (1995) and Petrin (2004) and the major steps are: 1) Assuming error terms and standard deviations of random terms, simulate market shares; 2) Assuming the standard deviations, calculate the errors that match observed market shares using contraction mapping theorem; 3) calculate the unobserved fixed effects by subtracting observed components from the errors in step two; 4) minimize the value function using the set of instruments. The set of instruments are challenging to find for our current dataset. In this study we use miles flown and total seats available as such supply side instruments.

The first two columns show the results without the elite survey data and the last two columns are computed with it augmented. The elite ratio is defined as the proportion of departing passengers that have at least Delta Silver Elite Medallion status. This ratio varies with time and departing airports. The connection factor is defined as market miles flown divided by non-stop miles which is always equal or bigger than one and has an effect on the coefficient of the connection factor. A larger number represents a less convenient connection. The load factor enters the demand function for the purpose of trying to capture the relative ease of schedule changes. A higher average factor load for an airport/carrier combination means that it is less likely to switch to other flights when necessary. Such convenience may be valued by certain types of customers. In order to save the time of computing and hunting for valid instruments, I assume that demographic characteristic variation is only from elite ratios across airports and it only has effect on coefficient of connection factor. Results have confirmed our suspicion that those mid-sized markets are more sensitive to prices than to convenience. Switching from offering a non-stop flight to a connecting flight with double the mileage will only cost about 8% of the market share and even less so when the route is elite-heavy. The load factor does not seem to be important in determining market shares in both cases so excess capacity should be avoided. The shares are, however, extremely volatile to prices: a small percentage increase in average fare can cause a big decrease in market shares.

III.4 Conclusion and Further Extensions

This chapter first reviews existing literature on airline pricing under oligopolistic market setup and then examines the determining factors of market shares in mid-sized air travel markets. My results show that the convenience of non-stop flights and their associated pricing power may be over-stated in those markets as in previous studies. When direct flights are offered by LCCs, they are not able to charge higher prices due to the fact that loyal customers of other legacy carriers are reluctant to switch and remaining passengers are very price sensitive. So even in airports that LCCs' presence is significant, prices remain more competitive than in hub airports. The market power associated with FFPs helps legacy carriers keep their most profitable customers with less convenient connecting products.

Further extensions include collecting passenger status data from other legacy airlines, as only the information of Delta passengers is analyzed in this chapter. Also better instruments can improve results and computation efficiency. In addition, if data availability allows, a dynamic panel model can be a more appealing choice of modeling as optimization across periods is not considered in this study.

Table III.1: Descriptive Statistics

	BOS-BNA				BOS-SJC			
	2010q1		2012q1		2010q1		2012q1	
	Share (%)	Price ($)	Share (%)	Price ($)	Share (%)	Price ($)	Share (%)	Price ($)
American	6.64	190.25	0.99	444.56	29.24	213.54	27.66	260.42
Delta	18.58	227.24	22.66	251.95	15.79	190.78	18.14	243.50
United	1.55	202.10	3.69	198.45	25.44	233.26	10.88	253.02
US Air	10.40	248.33	16.50	275.14	8.19	235.99	13.38	273.65
Jetblue (c)					2.34	247.95	4.26	303.29
Southwest (c)	56.16	208.12	49.12	242.77	8.19	186.74	15.42	258.77
Jetblue (d)							10.03	241.31
Southwest (d)			13.27	228.32				

Table III.2: Descriptive Estimations

	log(share of connecting flights) - log(share of outside good)					
	OLS	FE	OLS	FE	OLS	FE
Fare	-0.0001	0.00004	-0.0001	-0.0002	0.0001	0.0001
	(0.0006)	(0.00002)	(0.0006)	(0.0003)	(0.0004)	(0.0004)
Non-stop Fare			0.0015	0.0014**	0.0008**	0.0010**
			(0.0021)	(0.0006)	(0.0004)	(0.0004)
Miles Flown	-0.0001	-0.0001	-0.0001	-0.00007	-0.0001	-0.0001
	(0.0002)	(0.0001)	(-0.0002)	(0.0001)	(0.0001)	(0.0001)
Non-stop Miles	-0.0625**	-0.0579**	-0.0121	-0.0114	-0.0303**	-0.0322
	(0.0305)	(0.0131)	(0.0778)	(0.026)	(0.0155)	(0.0183)
Airline Dummies	No	No	No	No	Yes	Yes
R-square	0.1642	0.1616	0.1789	0.1762	0.1929	0.1951

Table III.3: The Full Structural Models

		Without Demographic Data		With Demographic Data	
	Variable	Estimate	Std. Errors	Estimate	Std. Errors
Mean Utility	constant	112.42	45.92	106.93	50.39
	connection factor	-7.64	5.41	-11.25	7.10
	elite on connection			5.19	3.18
	origin load factor	-2.52	1.81	1.59	3.03
	log(fare)	-18.66	7.43	-8.25	8.02
Std. Dev.	connection factor	3.68	4.88	3.88	4.93
	origin load factor	2.08	1.02	2.25	1.33
	log(fare)	0.17	1.16	0.01	3.29

BIBLIOGRAPHY

[1] Aguiar, M. and G. Gopinath (2007) "Emerging Market Business Cycles: The Cycle Is the Trend," *Journal of Political Economy 115*, 69-102.

[2] Aguirregabiria V. and C. Ho (2012):"A dynamic oligopoly game of the US airline industry: estimation and policy experiments," *Journal of Econometrics 168*, 156-173.

[3] Alderighi M. (2010):"Fare dispersions in airline markets: a quantitative assessment of theoretical explanations," *Journal of Air Transport Management 16*, 144-150.

[4] Alquist, R., L. Kilian and R. J. Vigfusson (2011) "Forecasting the Price of Oil," *International Finance Discussion Paper No. 1022*, Board of Governors of the Federal Reserve System.

[5] Backus, D. K. and M. J. Crucini (2000) "Oil Prices and the Terms of Trade," *Journal of International Economics 50*, 185-213.

[6] Bai, J. and S. Ng (2002) "Determining the Number of Factors in Approximate Factor Models," *Econometrica 70*, 191-221.

[7] Bai, J. (2004) "Estimating Cross-Section Common Stochastic Trends in Nonstationary Panel Data," *Journal of Econometrics 122*, 137-183.

[8] Bai, J. and S. Ng (2008) "Forecasting Economic Time Series Using Targeted Predictors," *Journal of Econometrics 146*, 304-317.

[9] Bansal, R. and I. Shaliastovich (2012) "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," *NBER Working Papers 18357*.

[10] Bansal, R., and A. Yaron (2005) "Risk for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance 59(4)*, 1481-1509.

[11] Bansal, R., D. Kiku and A. Yaron (2010) "Long Run Risks, the Macroeconomy, and Asset Prices," *American Economic Review 100(2)*, 542-546.

[12] Baumeister, C. and L. Kilian (2015) "Forecasting the Real Price of Oil in a Changing World: A Forecast Combination Approach," *Journal of Business and Economic Statistics*, forthcoming.

[13] Baumeister, C., P. Guerin and L. Kilian (2015) "Do High-Frequency Financial Data Help Forecast Oil Prices? The MIDAS Touch at Work," *International Journal of Forecasting 31(2)*, 238-252.

[14] Bernanke, B., J. Boivin and P. S. Eliasz (2005) "Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach," *The Quarterly Journal of Economics 120(1)*, 387-422.

[15] Berka, M., M. J. Crucini and C.W. Wang (2012) "International Risk Sharing and Commodity Prices." *Canadian Journal of Economics, 45(2)*, 417-447.

[16] Berry S. (1990): "Airport presence as product differentiation," *American Economic Review 80*, 394-399.

[17] Berry S. (1992): "Estimation of a model of entry in the airline industry," *Econometrica 60*, 889-917.

[18] Berry S., M. Carnall and P.T. Spiller (1996): "Airline hubs: costs, markups and the implications of customer heterogeneity," *NBER Working Paper Series 5561*.

[19] Berry S., J. Levinsohn and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica 63*, 841-890.

[20] Berry S., P. Jia (2008): "Tracing the woes: an empirical analysis of the airline industry," *NBER Working Paper Series 14503*.

[21] Bhardwaj, G. and A. Dunsby (2013) "The Business Cycle and the Correlation Between Stocks and Commodities," Working Paper.

[22] Bloom, N. (2009) "The Impact of Uncertainty Shocks," *Econometrica 77(3)*, 623-685.

[23] Borenstein S. (2005): "U.S. Domestic Airline Pricing, 1995-2004," *University of California Working Paper No. CPC05-48*.

[24] Breeden, D. T. (1980) "Consumption Risk in Future Markets." *Journal of Finance 35(2)*, 503-520.

[25] Breiman, L. (1996) "Bagging Predictors." *Machine Learning 36*, 105-139.

[26] Buyuksahin, B., M.S. Haigh and M.A. Robe (2008) "Commodities and Equities: A Market of One?" *Working Paper*.

[27] Casassus, J., P. Collin-Dufrene and B. R. Routledge (2005) "Equilibrium Commodity Prices with Irreversible Investment and Non-Linear Technology." *NBER Working Paper 11864*.

[28] Chen, H. (2010) "Macroeconomic Conditions and the Puzzles of Credit Spreads and Capital Structure," *Journal of Finance 65(6)*, 2171-2212.

[29] Chen, Y., K. Rogoff and B. Rossi (2010) "Can Exchange Rates Forecast Commodity Prices?" *The Quarterly Journal of Economics 125(3)*, 1145-1194.

[30] Chiang, I., W. K. Hughen and J. Sagi (2014) "Estimating Oil Risk Factors Using Information from Equity and Derivative Markets", *Journal of Finance,* forthcoming.

[31] Christiano, L., R. Motto and M. Rostagno (2014) "Risk Shock", *American Economic Review 104(1)*, 27-65.

[32] Cochrane, J. H. (1994) "Permanent and Transitory Components of GNP and Stock Prices," *Quarterly Journal of Economics 30,*

[33] Cole, H. L. and M. Obstfeld (1991) "Commodity Trade and International Risk Sharing." *Journal of Monetary Economics 28*, 3-25.

[34] Crucini, M. J. and M. Shintani (2014) "Measuring Business Cycles by Saving for a Rainy Day," *NBER working paper no. 16075*.

[35] de Roon, F. and M. Szymanowska (2010) "The Cross-Section of Commodity Future Returns." *SSRN Working Paper 891073*.

[36] Dusak, K. (1973) "Futures Trading and Investor Returns: An Investigation of Commodity Market Risk Premiums," *Journal of Political Economy 81*, 1387-1406.

[37] Fama, E.F. and K.R. French (1993) "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics 33*, 3-56.

[38] Fama, E.F and J. MacBeth (1973) "Risk, Return and Equilibrium: Empirical Tests," *Journal of Political Economy 81(3)*, 607-636.

[39] Ferson, W., S. Nallareddy and B. Xie (2013) "The Out-of-Sample Performance of Long Run Risk Models," *Journal of Financial Economics 107(3)*, 537-556.

[40] Frankel, J. and A. Rose (2009) "Determination of Agricultural and Mineral Commodity Prices," *Inflation in an Era of Relative Price Shocks*, 9-51.

[41] Gao, Y. and K.X.D Huang (2014) "Risk Shocks and Staggered Production in International Business Cycles," *Unpublished manuscript*.

[42] Gargarno, A. and A. Timmermann (2014) "Predictive Dynamics in Commodity Prices." *International Journal of Forecasting*, forthcoming.

[43] Gayle P.G. (2004):"Does price mater? Price and non-price competition in the airline industry," *Kansas State University Working Paper*.

[44] Ghysels, E., P. Santa-Clara and R.Valkanov (2005) "There is a Risk-Return Trade-Off after all," *Journal of Financial Economics 76(3)*, 509-548.

[45] Gorton, G. and K. G. Rouwenhorst (2006) "Facts and Fantasies about Commodity Futures." *Financial Analyst Journal 62(2)*, 47-68.

[46] Gospodinov, N. and S. Ng (2013) "Commodity Prices, Convenience Yields, and Inflation," *The Review of Economics and Statistics 95(1)*, 206-219.

[47] Hansen, L.P., J.C. Heaton and N. Li (2008) "Consumption Strikes Back? Measuring Long-Run Risk," *Journal of Political Economy 116(2)*, 260-302.

[48] Heston, S. (1993) "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies 6(2)*, 327-343.

[49] Inoue, A. and L. Kilian (2008) "How Useful is Bagging for Forecasting Economic Time Series? A Case Study of U.S. CPI Inflation," *Journal of the American Statistical Association 103*, 511-522.

[50] Johannes, M. and N. Polson (2009) "MCMC methods for financial econometrics," *Handbook of Financial Econometrics vol.2*, 1-72.

[51] Jurado, K., S.C. Ludvigson and S. Ng (2015) "Measuring Uncertainty, " *American Economic Review 105(3)*, 1177-1216.

[52] Kim, C. and C.R. Nelson (1999) "State Space Models with Regime Switching," MIT Press.

[53] Kim, H. and N. Swanson (2014) "Forecasting Financial and Macroeconomic Variables Using Data Reduction Methods: New Empirical Evidence," *Journal of Econometrics 178*, 352-367.

[54] Kose, M.A., C. Otrok and C.H. Whitman (2003) "International Business Cycles: World, Region and Country-Specific Factors," *American Economic Review 93(4),* 1216-1239.

[55] Morrison S.A. and C. Winston (1990): "The dynamics of airline pricing and competition," *American Economic Review 80*, 389-393.

[56] Pastor, L. and R. Stambaugh (2012) "Are Stocks Really Less Volatile in the Long Run?" *Journal of Finance 67(2)*, 431-478.

[57] Petrin A. (2002):" Quantifying the benefits of new products: the case of the minivan," *Journal of Political Economy 110*, 705-729.

[58] Stock J. and M. Watson (2002) "Macroeconomic Forecasting using Diffusion Indexes," *Journal of Business and Economic Statistics 20(2)*, 147-162.

[59] Stock J. and M. Watson (2004) "Combination Forecasts of Output Growth in a Seven-Country Data Set," *Journal of Forecasting 23*, 405-430.

[60] Tang, K. and W. Xiong (2012) "Index Investment and the Financialization of Commodities," *Financial Analyst Journal 68(6)*, 54-74.

[61] Wang, C.W. (2008) "Commodity Price Dynamics: Evidence and Theory," *Vanderbilt University Working Paper.*