Commodity Dependence and Aggregate Risk

Timothy C. Johnson*

August 2011

Abstract

This paper uses asset pricing methodology to quantify the consequences of the economy's exposure to commodity risk. A simple model of an economy that uses raw materials for both production and consumption yields a rich depiction of the joint behavior of asset returns, dividends, and commodity prices. Financial market data then identifies key elasticity parameters. The fitted model is used to study the ways in which welfare and marginal utility are affected by the evolution of relative commodity supply or the commodity expenditure share. Although the model contains no frictions, there are strong real effects. As the expenditure share rises the increasing volatility of the effective consumption basket can dramatically erode welfare. However until very high levels of expenditure a stronger effect is a decline in volatility of non-commodities output. Locally, higher commodity prices are good, not bad, news and agents would not pay to insure against them. Still, welfare may be significantly enhanced by the ability to avoid extremely high prices, via lowering output volatility, increasing flexibility in production, or forgoing growth in other factors of production.

Keywords: commodity prices, asset prices, welfare, risk.

JEL CLASSIFICATIONS: E23, G12, Q31, Q41

1 Introduction

The global economic expansion prior to 2008 touched off an unprecedented explosion in raw commodities prices. Copper prices soared by a factor of 5 from May 2003 to

*University of Illinois at Urbana-Champaign. I am grateful to Don Fullerton, Dirk Kreuger, Pietro Veronesi, Yilin Zhang and two anonymous referees for thoughtful comments. I also thank seminar participants at Columbia University, the Federal Reserve Bank of Boston, Nanyang Business School, Singapore Management University, and the University of Illinois.

May 2006. Agricultural commodities, according to The Economist index, rose 200 percent in the 20 months to April 2008, triggering food riots and export controls in many developing countries. Oil prices doubled over the year ending in June 2008 even with the developed world's economy already sliding into recession. This commodity-specific inflation – which occurred despite record output of many primary products – raised very real concerns about the sustainability of growth in the face of limited material resources. These concerns were reflected at the time in the stock market, where commodity price increases were increasingly viewed as bad news, and commodity-induced volatility became an important component of market risk. While the contraction of 2009-2010 relieved some of the pressures, signs of renewed commodity inflation already abound.

Motivated by this recent experience, this study seeks to understand the real risks to the economy due to changes in the availability of, and need for, raw inputs. Clearly, commodity risk can become macroeconomically important. Less clear are the implications of this exposure. Commodity price rises are obviously endogenous and reflect both good and bad news. What might happen, then, following another doubling or tripling of real prices? How would the evolving dynamics of output and consumption affect growth, discount rates, and risk premia? How does the possibility of these outcomes influence current hedging demands and welfare? To what extent (and how) should resources be committed to avoiding high commodity price states?

This paper presents an asset-pricing perspective on these questions. Using a frictionless representative agent model of of an economy that needs raw commodities for both consumption and production, I derive the joint dynamics of stock prices and commodity prices. These dynamics – particularly second moments – are then used to back out the key structural parameters that determine how commodities affect agents' welfare. I also show that the intertemporal component of welfare can be represented as the "price-dividend ratio" of a claim to the composite consumption basket. The theory then allows a quantitative assessment of the effects of future scarcity.

The paper complements a large macroeconomic literature modeling the effects of commodity supply on growth. Researchers have suggested many channels through which frictions may imply real effects of commodity price inflation. The model presented here, while excluding important realistic features, offers the benefit of tractability that enables the exploitation of asset pricing information. Moreover, it shows that very rich dynamic patterns – both in financial markets and the real economy – can arise without frictions.

Turning to the results, the estimation yields an elasticity of substitution between commodity and non-commodity inputs that is much lower in consumption than it is in production. This implies that as commodities become more scarce, they become more volatile at the same time as they take up an increasing fraction of expenditures. Agent's effective consumption basket thus experiences rising risk. Combined with high risk aversion, this leads to very large welfare losses as commodity prices rise above a (typically very high) critical level. Below that level, commodity price rises are good news for intertemporal welfare, though. Output of non-commodity goods becomes less volatile as its input mix becomes more diversified. Locally, commodity prices are effectively procyclical.

Armed with the estimation results, the paper can then address the implications of commodity dependence from a number of angles. The findings can be summarized briefly as follows.

- Lowering commodity supply volatility can be extremely welfare enhancing. But specifically targeting lower *price* volatility can actually harm welfare by raising output volatility.

- Increasing flexibility of the production sector by raising its elasticity of substitution can produce sizeable benefits in both high and low price states.

- Reducing productivity growth in order to lower relative scarcity can be an optimal policy at very high levels of commodity expenditure.

- When technological evolution is endogenized, the main benefits come from lowering price volatility. This can involve raising, not lowering, commodity usage in production.

It is important to stress that the conclusions here are not qualitatively dependent on precise assumptions about the future evolution of commodity supply. (Although the precise value today of any particular long-run policy of course depends on state probabilities.) The paper does not attempt to forecast *whether* high commodity prices will occur in the future, but rather to delineate the real mechanisms whereby such an occurrence would affect investors and consumers.

The outline of the paper is as follows. The next section describes the model and explains the determination of aggregate quantities and prices. In Section 3 the technology and preference parameters are estimated and the dynamics of the fitted model are examined in order to understand the identification. Section 4 contains the main results. Using the estimated parameters, I assess the implications of commodity dependence for aggregate risk, marginal utility, and intertemporal wealth. The final section summarizes the paper's contribution.

2 The Economy

This section introduces a simple model of an economy in which an endowment good (the "commodity") plays a role in both production and consumption. The immediate goal is to delineate how the primitive technology and preference parameters that govern the need for the commodity translate into price dynamics. Using price dynamics to estimate the parameters then enables an assessment of the role of commodity dependence in marginal utility and welfare.

2.1 Technology and Preferences

The setting is a two-good generalization of the Lucas (1978) endowment economy set in continuous time. There is a single infinitely-lived representative agent who maximizes lifetime expected utility under rational expectations.

The economy possesses an endowment asset which yields a continuous exogenous flow, Q, of a raw input good. The agent in the economy may use the quantity Q either for direct consumption or for production of finished goods, which comprise a second component of consumption. The amounts of Q devoted to consumption and production are denoted L and N, respectively. Since the commodity is not storable, we have $Q = L + N$.[1]

The quantity of finished goods generated by the productive sector is denoted C. Intuitively, C represents all output of goods and services aside from raw materials, and hence should normally make up most of the consumption basket. For this reason, C is taken to be the numeraire for the economy. The C-producing sector also uses other factors of production summarized by a process, Z, which represents a second exogenous source of uncertainty. This variable can be interpreted as an amalgam of all non-commodity inputs, including capital and labor. Alternatively, as in a standard business cycle model, it can be interpreted as a stochastic level of productivity.

Given realizations of the two state variables, the technology of the economy is described by a production function for finished goods of the constant-elasticity-of-substitution (CES) class:

$$C = F(Z, N) = [\eta Z^{\nu} + (1 - \eta)N^{\nu}]^{\frac{1}{\nu}} . \tag{1}$$

The parameter $\eta \in (0, 1)$ serves to determine the relative weight of the two inputs to production. The parameter ν governs the degree of substitutability of the factors. Formally, the elasticity of substitution is defined as $1/(1-\nu)$. An important special case, which will be used in much of the analysis, is $\nu = 0$, or unit elasticity of substitution. In

[1] The model is implicitly viewing Q as the quantity available for current use after any management of extractive capacity and/or inventory management. Processing or refining of the good for direct consumption is implicitly viewed as a costless, linear technology.

that case, output is given by the Cobb-Douglas production function

$$F(Z,N) = Z^{\eta} \quad N^{1-\eta}. \tag{2}$$

The representative agent in the economy has instantaneous utility, U, defined over his stream of finished goods consumption, C, and commodities consumption, L. Utility is time-separable, with subjective discount rate denoted ϕ. The utility function is also assumed to be of the CES form:

$$U(C,L) = \frac{1}{1-\gamma}\left[\beta C^{\rho} + (1-\beta)L^{\rho}\right]^{\frac{1-\gamma}{\rho}} \tag{3}$$

As in the production function, the relative scale of commodities in consumption is governed by the weight β. The parameter γ is the coefficient of relative risk aversion (and also the inverse of the elasticity or intertemporal substitution). The elasticity of intratemporal substitution between consumption goods is $1/(1-\rho)$.

Combining the production and utility functions, the consumer may be viewed as maximizing a two-level CES utility function $U(f(Z,Q-L),L)$.[2] The model has been used to study asset prices by Basak (1999), who examines the effect of endogenous labor supply on the equity premium. While it inherits the well-known limitations of the constant relative risk aversion (CRRA) class, it clearly delineates a number of important economic mechanisms. In addition, I show that it is capable of providing a realistic description of commodity prices, output, and asset returns.

[2] Two-level CES production functions have been extensively employed in the empirical literature to model differing complementarities between inputs. Polgreen and Silos (2009) estimate a production function involving capital, labor, and energy (oil) inputs.

2.2 Equilibrium

Equilibrium in the model is fully determined by the consumer's allocation decision for Q. The allocation is not subject to any frictions, and the first order condition is simply

$$F_N(Z, N) = \frac{U_L}{U_C}. \tag{4}$$

Given Q and Z, this is an implicit algebraic equation for L (or N). Basak (1999) proves the existence and uniqueness of a solution to this equation for all parameter values. While the equation cannot be explicitly solved, it is straightforward to show that it is of the form $L = Q\ \ell(\zeta)$ where $\zeta \equiv Z/Q$. The latter ratio will be referred to as the relative scarcity of the commodity. In the Cobb-Douglas case, (4) yields

$$\zeta = z_0\ \ell^e\ (1-\ell)^{1-e} \tag{5}$$

where z_0 is a constant and $e = \frac{1}{\eta}(1 - \frac{1}{\rho})$.

The intuition behind the allocation decision is straightforward: when faced with an increase in relative scarcity, more is allocated to the less elastic usage. That is, if $\rho < \nu$, the economy is better able to substitute other inputs for commodities in production than it is in consumption, hence the efficient use of the diminished supply is to use less in production. So $N/Q = 1 - \ell$ will decline. Analytically, the fact that the sign of ℓ' equals the sign of $\nu - \rho$ can be shown by implicit differentiation of the first order condition.[3] The paper will argue that $\rho < \nu$ is the empirically relevant case. This implies that as relative scarcity grows large the allocation fraction ℓ goes to unity.

The homogeneity of the allocation solution carries over to that of finished goods output and consumption, which can be written $F = Q\ f(\zeta)$. The key feature of output dynamics is that the exposure of output to each type of shock shifts from one to the

[3] This can be seen immediately for $\nu = 0$ by implicitly differentiating (5) in logs, which gives the sensitivity of ℓ to ζ as $\ell'\zeta = \ell(1-\ell)/(e-\ell)$. If $\rho > 0$ then $e < 0$ and the denominator is negative. If $\rho < 0$ then $e > 1$ and it is positive.

other with relative scarcity. Formally, writing the stochastic process for dC as,

$$\frac{dC}{C} = \mu_C(Q,Z)\ dt + \delta_1(Q,Z)\ \sigma_Z\ dB^Z + \delta_2(Q,Z)\ \sigma_Q\ dB^Q$$

homogeneity immediately gives $\delta_1 = \delta_1(\zeta) = 1 - \delta_2(\zeta) = f'\zeta/f$. The coefficients δ_1 and δ_2 are thus mirror images of each other. When $\rho < \nu$ the exposure of (non-commodity) consumption to commodity supply shocks increases with relative scarcity, whereas the exposure to the productivity factor Z diminishes and may even become negative.

2.3 Commodity Prices

Equation (4) determines the price, P, of a unit of the commodity (in C units) because market clearing in the primary goods market dictates that $P = F_N$. Note also, from the form of the utility function, that

$$P = \frac{U_L}{U_C} = \frac{1-\beta}{\beta}\ \left(\frac{L}{C}\right)^{\rho-1}. \tag{6}$$

This expression yields the convenient relationship for the value of commodity consumption relative to finished good consumption,

$$X \equiv PL/C = p_0\ P^{\rho/(\rho-1)} \tag{7}$$

where $p_0 \equiv (\beta/(1-\beta))^{1/(\rho-1)}$. For future reference, I also define the expenditure fractions as $x \equiv X/(1+X)$ and $y \equiv 1 - x = 1/(1+X)$. From the homogeneity of L and C, we immediately have that all these quantities are solely functions of ζ. In particular, under very general conditions, the commodity price is an increasing function of relative scarcity.

As with C, the exact dynamics of P can be obtained from Itô's lemma up to expressions in $\ell(\zeta)$, via implicit differentiation. Section 3 will examine several specific features of implied higher moments in particular.

There is a large literature on commodity price dynamics that documents the theoretical and empirical importance of numerous considerations left out of the present model.[4] I augment this literature by illustrating (in the next section) that the simple equilibrium mechanisms captured in this model can also induce some complex and empirically relevant dynamic properties.

2.4 Asset Prices

Given the equilibrium quantities in the economy, we can compute the prices of long-lived claims to components of output. As usual, the value of a claim to any payment stream is the expected sum of the product of future dividends with future marginal utility. Since C is the numeraire, the pricing kernel process, $M \equiv U_C$, can be written

$$M_t = e^{-\phi t}\, C_t^{-\gamma}\, y_t^{-\psi}. \tag{8}$$

where $\psi \equiv (1-\gamma-\rho)/\rho$ and y, defined above, is the fraction of C in total consumption. The pricing kernel describes a two-factor model in terms of the (endogenous) variables C and y. Or, since y is a monotonic function of P, we can equivalently view marginal utility as determined by numeraire consumption and commodity prices.

The kernel is related to some recent contributions in asset pricing (Piazzesi, Schneider, and Tuzel (2007), Yogo (2006), Lochstoer (2009)) utilizing "two-good, two-tree" endowment models with non-separability in utility. As in those papers, the relative composition of the consumption basket may matter directly to investors: if $\rho \neq 1-\gamma$, agents care about commodity price risk for its own sake, not just because of its contribution to C risk. Unlike pure endowment models, however, here the dynamics of both factors affecting marginal utility – as well as those of the dividend streams to be valued and

[4]An incomplete list includes industry competition (Stiglitz (1976), Deaton and Laroque (1996)); resource uncertainty (Pindyck (1980), Sundaresan (1984)); inventory conditions (Pindyck (1994)); and real investment frictions (Casassus, Collin-Dufresne, and Routledge (2005), Carlson, Khokher, and Titman (2009)).

the spot market price of the second good – are determined endogenously by the production/consumption trade-off. The endogeneity of the price moments is what permits one to recover technology and preference parameters from the data.

The model also nests single-good, two tree models such as Santos and Veronesi (2006) and Cochrane, Longstaff, and Santa Clara (2008) if C and L are perfect substitutes and there is no usage of commodities in production ($\eta = 1$). Two insights from those models carry over to the present economy. First, aggregate consumption may become less risky as the contributions of a dominant component becomes smaller. Second, the risk premium on an individual asset will fall with its dividend's weight in composite consumption as the covariance between the two falls.

To my knowledge, there are no extant general equilibrium models that feature endogenous determination of both commodity prices and asset prices.[5] Wei (2003) models a production economy with putty-clay technology and and a single exogenous energy price shock. In her setting, labor, not capital, suffer losses from a price increase, and the stock market response is minimal. The macroeconomic literature has recognized the dual consumption-production role of oil, in particular. (See Bodenstein, Erceg, and Guerrieri (2011).) But the implications for the dynamics of discount rates and risk premia have not been studied.

There are several claims one could consider pricing in this economy. The analysis below will be primarily concerned with the "net profits" of the production sector of the economy. That is, a representative firm endowed with the production technology has output flow C and costs given by PN. Thus I define dividends as

$$D \equiv C - PN = F(Q, Z) - N(Q, Z)F_N(Q, Z)$$

which is just ηC in the Cobb-Douglas case. I view a claim to this stream as "the stock market" and denote its value S^D.

[5] The lack of such models is noted in Kilian and Park (2009), footnote 7.

The representative agent in the economy also owns the commodity output stream Q. So one can define S^Q as the value (in C units) of a claim to the stream PQ, representing the profits of the primary goods sector. The sum of the profits of the two sectors is

$$C - P(Q - L) + PQ = C + PL \equiv \hat{C}$$

which is the total consumption bundle. The value of a claim to it will be denoted $S^{\hat{C}}$.

One can also consider a variety of composite stock markets having more or less ownership of the primary goods sector. For a given $\lambda \in [0, 1]$, the value of a claim to the stream of finished goods plus a fraction λ of the commodity goods will be denoted S^λ. For example, the major global stock market indices, which represent a large fraction of global non-commodity output, but which exclude the (mainly sovereign) claims to most of the world's natural resources, could be represented via $\lambda = 0.05$, whereas an economy with a large commodity producing sector might have $\lambda = 0.5$. Note that we can write $D^\lambda = D + \lambda PQ = D + \lambda(\hat{C} - D)$. Hence $S^\lambda = \lambda \; S^{\hat{C}} + (1 - \lambda) \; S^D$.

One way of assessing the risks posed by the need for commodities is to examine the prices of assets whose cash flows are tied to P. The analysis in Section 4.1 will consider how much agents are willing to pay for insurance against high-P states. Options or forward contracts on P can be readily valued under the model as their risk-neutral discounted expected payoff, given the process for M.

Since the dynamics of C and P and L and N are not obtainable in fully closed form, the same is true of both the pricing kernel and the payoff streams of interest, and hence of the expectations determining asset prices. However computing these prices is straightforward numerically, via Monte Carlo integration.

2.5 Welfare

The measure of how well off agents are, under any circumstances, is the magnitude of their lifetime discounted expected utility, i.e., the value function. For the utility function (3), we have $U_C = (1-\gamma)\, U/\hat{C}$ (recall $\hat{C} \equiv C + PL$). Using this, the value function can be written

$$J_t = \mathrm{E_t}\left[\int_t^\infty e^{-\phi(s-t)}\; U(C_s, L_c)\; ds\right] = \frac{1}{1-\gamma}\; \mathrm{E_t}\left[\int_t^\infty e^{-\phi(s-t)}\; U_C(C_s, L_c)\; \hat{C}_s\; ds\right]$$

$$= \frac{1}{1-\gamma}\; U_C(C_t, L_t)\; \mathrm{E_t}\left[\int_t^\infty e^{-\phi(s-t)}\; M_s\; \hat{C}_s\; ds\right] = U(C_t, L_t)\; \left(\frac{S_t^{\hat{C}}}{\hat{C}_t}\right) \equiv U\; V. \tag{9}$$

The right-most expression here represents a decomposition between the value of current consumption, represented by the utility function, U, and the value of future consumption, represented by the "price-dividend ratio" of a claim to the stream of the composite consumption basket. For ease of reference, define this component of the value function to be $V \equiv S^{\hat{C}}/\hat{C}$. It is straightforward to show that this ratio is also only a function of relative scarcity. When $\gamma > 1$, bear in mind that decreases in V and decreases in $|U|$ are improvements in welfare, since utility is negative.[6]

Section 4 will assess the welfare consequences of alternative parameterizations of the model. In doing so, the focus will be on the intertemporal component, $V(\zeta)$.

2.6 Stochastic Setting

The two exogenous state variables will be assumed to obey Itô processes, that is,

$$\frac{dZ}{Z} = \mu_Z(\zeta)\; dt + \sigma_Z(\zeta)\; dB^Z, \qquad \frac{dQ}{Q} = \mu_Q(\zeta)\; dt + \sigma_Q(\zeta)\; dB^Q \tag{10}$$

[6]It seems counterintuitive that agents prefer *lower* asset values. But for these preferences, the value of the future composite consumption stream reflects future marginal utility more than future cash flows. So, holding current consumption fixed, a high price-dividend ratio indicates a greater likelihood of being worse off in the future.

where dB^Z and dB^Q are standard Brownian motions with instantaneous correlation θ. None of the parameters of these processes enter into the allocation decision. Thus the fundamental properties of the equilibrium are all robust to state dependent growth rates or volatilities.

On the other hand, long-run properties of the economy clearly will depend on the likelihood of future scarcity. The paper seeks to identify implications of commodity dependence that do not depend on this likelihood, which can be altered by changing the relative growth rates or the correlation between the innovations. While the model does not include any endogenous mechanism whereby the supply of either factor increases in response to its relative scarcity, such a response could always be captured in reduced form in (10) by making the processes cointegrated. In the baseline case estimated in the next section, the two variables will be assumed to grow at the same rate.

2.7 Technological evolution

Since the paper aims to deduce implications about the economy's well-being in the face of possible shortage of resources, it may be important to consider the robustness of the analysis to the assumption that technology is fixed. For this reason, Section 4 will examine a generalization of the model which includes the possibility of costly adaptation. Specifically, retaining the Cobb-Douglas output function (2), I write the variable Z as the product of a stock of a productive resource H with a technology shock ϵ_t and assume that the agent can draw down the level of H to alter the weight $1-\eta$ of commodity inputs to production. If one thinks of H as representing technology-specific human capital ("know how"), then the intuition is that adopting new production technologies imposes a cost in accumulated knowledge. Effectively, changing the drift rate of H is the same as altering the parameter μ_Z.

Denoting the log of the usage ratio $\eta/(1-\eta)$ as $\hat{\eta}$, then the ability to alter the technology is captured by the drift of this process. Agents may choose a rate of alteration i such that $d\hat{\eta} = bi\ dt$. In doing so, the stock of H is assumed to decay at rate $a(i; Q, Z, \eta)\ H\ dt$. Here the function a is chosen so that the impact of the alteration on the growth rate of output, dC/C, is proportional to the (absolute value) of the policy. Specifically, setting

$$a(i) = -[\frac{\partial \log f}{\partial \log Z}]^{-1}\ [\frac{d \log f}{d\hat{\eta}}\, i + k\ |i|]$$

the loss in instantaneous production is $k\ |i|$, The constant k is thus the effective elasticity of C with respect to the usage ratio. There are no constraints on the choice of i.

This extension adds two interesting features to the economy. It allows the long-run elasticity of substitution (in production) to differ from the short-run elasticity.[7] It also incorporates endogenous real costs of reallocation as output growth may be sacrificed to alleviate the need for Q (or, in the case of a resource glut, to exploit its abundance).

While solving the extended model involves dynamic programming, the formulation preserves the instantaneous commodity allocation solution of the basic model. Since the costs of reallocation come out of the growth of the factor H and not out of the instantaneous amounts of Q or Z, and since the choice of L and N is not subject to any additional friction, that choice is still governed by the first-order condition (4) and the spot price is still given by (6). Moreover, the economy is still characterized by just three state variables, e.g., Q and $\zeta = (H\epsilon)/Q$ and η. Homogeneity ensures that the optimal policy and the intertemporal component of welfare V will only be a function of the latter two.

3 Estimation and Implied Dynamics

The model described above can encompass a wide range of possible behavior depending on the parameter values. I now consider what asset market data tell us about the

[7] To be careful, as technology evolves the instantaneous elasticity is always 1.0 because the production functions stays Cobb-Douglas. But the long-run output response to a given change in relative supply of inputs can differ from the instantaneous response.

most likely specifications. This empirical analysis complements a large macroeconomic literature on the technology and preference parameters governing commodity usage. The fitted model actually provides an equilibrium account of a number of salient features of prices and returns. This is, however, an intermediate step in the paper: the goal (taken up in Section 4) is to examine the implications of the fitted model for scenarios beyond those experienced historically.

3.1 Estimation

I estimate the model using the historical time-series of the Goldman Sachs Commodity Index (GSCI). In the model, "commodities" represent all natural resources required in both production and consumption. As such, the most appropriate empirical counterpart is of a broad, quantity-weighted basket such as the GSCI.[8] As of the end of 2010, energy products comprised 67% of the index, with agricultural products and metals making up 21% and 12% respectively. Daily data start in January 1970 and run through December 2010. In the model, non-commodity goods are the numeraire. So I deflate the GSCI series using (daily interpolated) values of the Bureau of Labor Statistics consumer price index excluding food and energy ("core CPI").

The estimation also uses three financial series. The CRSP value-weighted index is taken as the proxy for the claim to the stream D. Its one-year trailing real dividend flow is used as D itself. The real interest rate is computed as 3-month Treasury bill yield minus the contemporaneous change in the core CPI.

For purposes of this section, I focus on the sub-case of unit elasticity of substitution in production ($\nu = 0$). Cobb-Douglas production functions are a common specification in macroeconomic modeling. And, while the choice here is primarily motivated by considerations of tractability, it is not manifestly counterfactual. There is an empirical literature

[8]The index is currently computed by Standard and Poors. Weights are updated annually based on the most recent five years' global output of each component. Prices are those of nearest dated futures contracts. I treat them as spot prices in terms of the model.

that attempts to estimate the degree of substitutability between energy and capital in production, which is similar to the quantity captured by $1/(1-\nu)$. There is no consensus in the estimates on whether this elasticity is greater or less than one. (See Berndt and Wood (1979) and Solow (1987).)

Another maintained assumption of the estimation is that the exogenous variables Z and Q follow (possibly correlated) geometric random walks with equal mean. This assumption merits further comment. It is not an aim of the analysis to draw conclusions about the long-run path of relative scarcity. Hence I do not attempt to identify the drift of commodity output from the historical growth rate of the commodity price index. In fact, this in-sample growth rate is close to zero. The assumption of constant and equal growth rates is consistent with this realization while not imposing that agents believed *a priori* in stationary commodity prices.

I estimate the model parameters by the simulated method of moments (SMM). The model implies that essentially all dynamic quantities are functions of the level of commodity prices (or the expenditure ratio). I thus focus on picking up the extent to which the prices and moments covary with p, the log real level of the commodity index. Specifically, I attempt to match the mean product of dividend growth, market returns, market volatility, commodity volatility, and the price-dividend ratio with p. I also include commodity volatility and the covariance of market returns with commodity returns. Finally, the set of moments also includes the real rate, the mean and volatility of dividend growth and equity returns, and the log price-dividend ratio. I weight the moment errors using an asymptotically optimal matrix.

The data include no actual quantity information. One result of this is that the parameter β (the weight of non-commodity inputs in utility) is not identified, as it has no effect on price dynamics. In addition, the level of the commodity price itself is not identified. (The units of the GSCI are arbitrary.) The estimation only requires an initial value to identify the level of the p path, Expenditure on energy goods and services was

6.4% of consumption in the fourth quarter of 1969, suggesting a plausible lower bound. Expenditure on food and energy goods was 20% of consumption, probably an upper bound. The baseline results assume $X_0 = 0.10$; results are also shown for $X_0 = 0.05$ and $X_0 = 0.20$.

Under the random walk assumption for the exogenous variables, the economy is non-stationary and the moments to be matched do not exist. To address this, I compute the *conditional* moments given that a non-divergent path has been observed. That is, model moments are computed in a simulated sample in which the path of p is constrained to stay within the range that was historically observed.[9] Of course, computing the moments of the data also implicitly conditions upon the stationarity of the generating process.

Asset prices in the model are not available in closed-form. To make the estimation computationally feasible, a fast approximation method for the price-dividend ratio is developed whose accuracy has been verified under a wide range of parameter values.[10]

Results of the estimation are shown in Table 1. All cases show evidence of low elasticities of substitution between commodities and non-commodity components of consumption. The t statistics in the table are misleading due to the nonlinearity of the parameter restrictions. In particular, the estimation of ρ is only imprecise in one direction. The log elasticity is very significantly different from zero. To show this, Figure 1 computes the Newey and West (1987) distance statistic (and associated p value) for restricted models estimated with varying elasticities. Based on this statistic, a 99% confidence interval for the elasticity does not include 0.75, while the hypothesis that the elasticity is effectively zero (e.g., $\rho < -100$) is not rejectable. In fact, asset market and commodity price data together are highly informative about ρ's upper bound.

---

[9] More precisely, the state variable ζ is drawn from an Ornstien-Uhlenbeck process whose ratio of instantaneous variance to unconditional variance is the same as that computed for the log GSCI index.

[10] The methodology is based on a weighted average of the exact price-dividend values that would obtain if the (risk-neutral) growth rate and volatility of dividends were constant. Since these moments vary over the range of x, the approximation weighs the static price dividend ratios by the occupation time of the process in each region of x. The method can be shown to be exact as the partition of x grows finer, and is exact at the endpoints of the range. Alternative log-linearizations perform much more poorly. Further details are available upon request.

Table 1: *Parameter Estimates*

X_0	ρ	η	γ	μ_Z	μ_Q	σ_Z	σ_Q	θ
0.10	-9.5636	0.9488	19.3836	0.0379	0.0379	0.0631	0.0625	0.5776
	(0.43)	(14.76)	(5.36)	(3.87)	(3.87)	(28.43)	(3.34)	(2.69)
0.20	-18.3888	0.9329	21.1073	0.0375	0.0375	0.0602	0.0595	0.7222
	(0.21)	(9.47)	(5.26)	(3.62)	(3.62)	(25.52)	(6.11)	(4.47)
0.05	-4.1919	0.9673	21.7209	0.0422	0.0422	0.0630	0.0617	0.3946
	(0.35)	(14.59)	(5.30)	(3.97)	(3.97)	(33.54)	(2.30)	(1.622)

Simulated method of moments estimates are shown for the model of Section 2 for the case of Cobb-Douglas technology and geometric random walk commodity supply and productivity. Results are shown for three values of X_0, the initial commodity expenditure ratio. The estimation imposes $\mu_Z = \mu_Q$. The parameter ϕ is estimated to be zero, its lower bound. The parameter β is not identified. Asymptotic t statistics are shown in parentheses.

The finding here is consistent with strong results in the empirical literature documenting very low price elasticities of demand for oil and energy commodities. Hamilton (2009) summarizes results from numerous estimations which find aggregate price elasticities in the range of 0.10. Using vector autoregressions with time-varying parameters, Baumeister and Peersman (2009) report strongly declining elasticities over time, with values under 0.2 for the last 20 years. Note from equation (7) that in the model $1 - \epsilon = \rho/(\rho - 1)$ is identified from the slope of the log expenditure share with respect to the log real price. A simple regression (in log differences) of the NIPA consumption share of gasoline and fuel on the CPI-energy commodities series since 1959 yields a slope of 0.85, corresponding to $\rho \approx -6$.

Table 2 shows the primary moments from the sample and from the estimated model. While commodities risk cannot resolve all the empirical deficiencies of consumption-based

Figure 1: *Inference on Elasticity of Substitution in Consumption*

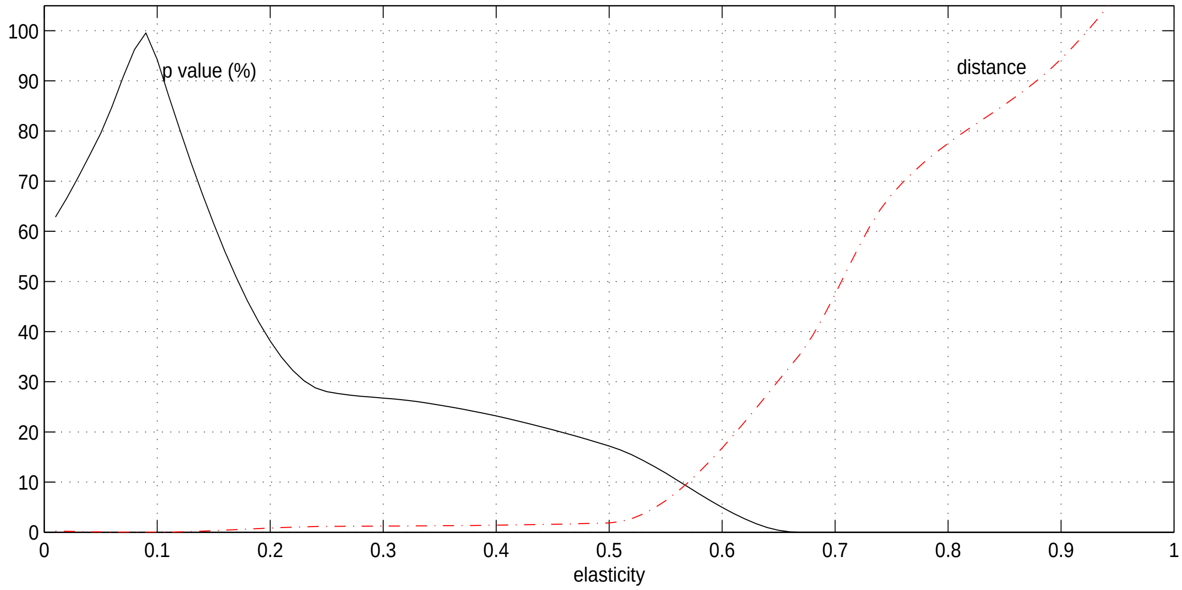

The figure shows the distance in maximized objective function of restricted SMM estimation from unrestricted estimation, where the restriction fixes the elasticity of substitution in utility $1/(1-\rho)$. The asymptotic p value (in percent) is also plotted.

models with constant relative risk aversion, the point estimates do deliver a plausible economy. The model is able to do reasonably well on the main financial dimensions, although it has trouble generating enough volatility.[11]

[11]The fit is actually better than the table would indicate. The estimation uses mean absolute values, which are less noisy than mean squares, for volatilities. Even a perfect fit to mean absolute values would lead to a divergence in standard deviations because the model (simulated) innovations are Gaussian and the daily data has substantial leptokurtosis.

Table 2: *Moment Comparison*

	$\text{E}[\Delta d]$	σ_d	$\text{E}[p-d]$	$\text{E}[r]$	$\text{E}[r^m]$	σ_{r^m}	σ_p
Sample:	1.39	9.05	3.59	1.34	4.00	16.70	19.52
Model:	3.78	6.35	3.42	1.49	5.59	8.98	12.81

The table shows estimated moments from the data and from the benchmark model shown in Table 1. The first two columns give the mean and standard deviation of log real dividends. The third column is the mean log price-dividend ratio. The fourth column is the real interest rate. The fifth and sixth columns are the mean and standard deviation of excess market returns. The seventh column is the standard deviation of log changes in the real commodity price series. Model moments are conditional upon the realization of a stationary sample path of the scarcity ratio Z/Q. All quantities except for the price-dividend ratio are in annualized percentage.

3.2 Model Properties

In the fitted model commodity prices influence asset prices in a number of ways. This section explains how the estimation identifies properties of technology and preferences, and illustrates the implied dynamic relationships.

The low elasticity of substitution in production is a robust property of estimated specifications. This is primarily caused by the difficulty of generating sufficient volatility when the elasticity is closer to one. As discussed in Section 2, the need to reallocate scarce raw materials from production to consumption rises with the difference in substitutibilites in the two uses. This reallocation provides an amplification mechanism to the volatility of commodity prices. Moreover, in the data commodity volatility itself increases with the level of commodities prices. (See Ferderer (1996) for empirical evidence of this property in crude oil.) In the model, commodity risk, σ_P, rises with scarcity for all parameter values. But the effect becomes negligible as ρ approaches zero. See Figure 2.

The volatility of commodity prices also adds an important component to the volatility

Figure 2: *Commodity Volatility*

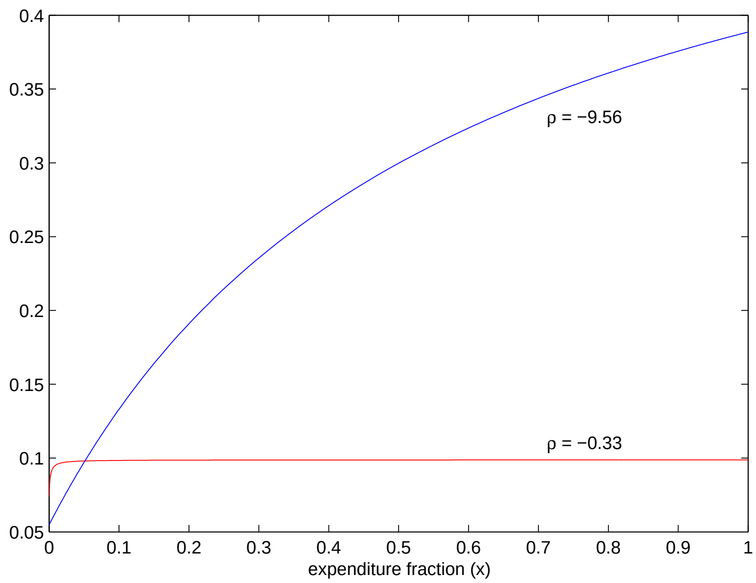

The figure shows the instantaneous volatility of log changes in the commodity price index for the benchmark case whose parameters are given in Table 1 (top line) and for the best-fit specification having $\rho = -1/3$ (lower line). The horizontal axis is the commodity expenditure fraction $x = PL/(PL + C)$.

of equity returns. The price-dividend ratio in the model is driven solely by changes in the expenditure ratio. So reallocation again amplifies fundamental volatility in stock prices. Specifications with less negative values of ρ are strongly penalized in the estimation because of their inability to do as well on this dimension. At the same time, there is a limit to how much commodity-induced volatility can increase stock return volatility because the correlation between commodity changes and stock returns is close to zero in the data.

Another feature of the data favoring low elasticity in production is a negative relationship between the level of commodity prices and the market price-dividend ratio. The correlation between the two series in the data is 11 percent.[12] With elasticities closer to one (ρ closer to zero) this relationship flips in the model because the risk-free rate falls with ζ (or with x or P) which raises asset values. The decline in r comes mostly

[12]The estimation takes into account the sampling uncertainty in the relationship through the moment weighting matrix.

through a precautionary savings channel. Commodity risk enters the expression for r with a coefficient equal to $-\psi = -(1-\gamma-\rho)/\rho$. When $1-\gamma < \rho < 0$, this is positive and it grows larger as ρ approaches zero from below. Since a large γ is necessary to match the equity premium, a relatively large negative ρ is necessary to keep the riskless rate from dropping in response to rises in commodity prices.

Figure 3 shows the fitted model's price-dividend ratio and riskless rate as a function of the commodity expenditure fraction, x. For this specification r always rises with x because dividend volatility σ_C falls with scarcity, which lowers the usual precautionary savings effect. For very high values of x, the precautionary effect with respect to σ_P (described above) can reverse the rise of r and the fall in the price-dividend ratio, but the σ_C effect prevails in the historically relevant region (e.g., $x < 0.25$) for statistically supported specifications.

Figure 3: *Riskless Rate and Price-Dividend Ratio*

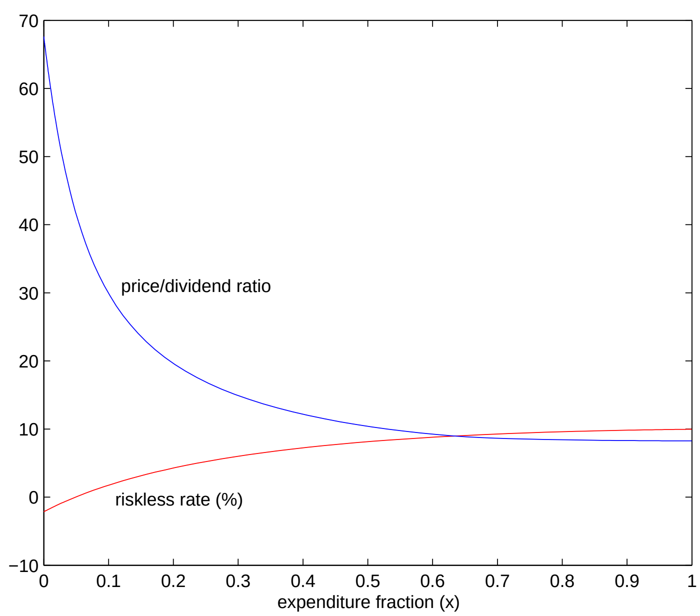

The figure shows the instantaneous riskless interest rate and the price-dividend ratio for the dividend claim for the benchmark case whose parameters are given in Table 1. The horizontal axis is the commodity expenditure fraction $x = PL/(PL + C)$.

The decline in volatility of dividends (and, equivalently, of finished goods output

C) with x is a diversification effect. The production function is mainly exposed to productivity shocks, Z, rather than commodity supply shocks Q, because the share of commodity inputs to production $1-\eta$ is low. As described in Section 2, the exposure to Z diminishes and the exposure to Q rises as scarcity increases. Since the exposures are imperfectly correlated, the result is a drop in overall risk.

The declining volatility of C with P leads to a corresponding decline in the volatility of marginal utility (despite the fact that P-volatility itself is rising) lowering risk premia.[13] In addition, the falling price-dividend ratio means that negative productivity shocks can actually raise stock prices, meaning the dividend claim can act as a hedge. Together with falling dividend risk, these effects cause the equity premium to decline with P for the empirically relevant range. See Figure 4. (For higher values of x the risk premium rebounds somewhat because the hedging effect vanishes when the price-dividend ratio stabilizes.)

The model thus implies an economically significant return predictability effect, with expected returns falling by around 400 basis points as x rises from 0 to 0.25. This is consistent with evidence presented in Driesprong, Jacobsen, and Maat (2008) and Jacobsen, Marshall, and Visaltanachoti (2009) who find that rises in commodities prices are followed by predictably low stock returns. Moreover, the effect is weaker in countries with higher oil reserves (e.g., Norway and the U.K.). This is also a prediction of the model. Figure 4 also plots the risk premium when the underlying asset is a claim, not just to profits of the production sector (the lowest curve), but also to some component of the profits of the commodity supplying sector. As in the data, increasing the share of commodities ownership in the equity claim results in less decline in the risk premium with x, and hence less predictive power of changes in P for future returns.

[13] Using equation (8) the return premium for any asset whose price is S can be written as

$$\gamma \ \mathrm{cov}[dS/S, dC/C] + \psi \ \mathrm{cov}[dS/S, dy/y]. \tag{11}$$

For most parameter values, the second term is smaller than the first.

Figure 4: *Expected Stock Returns*

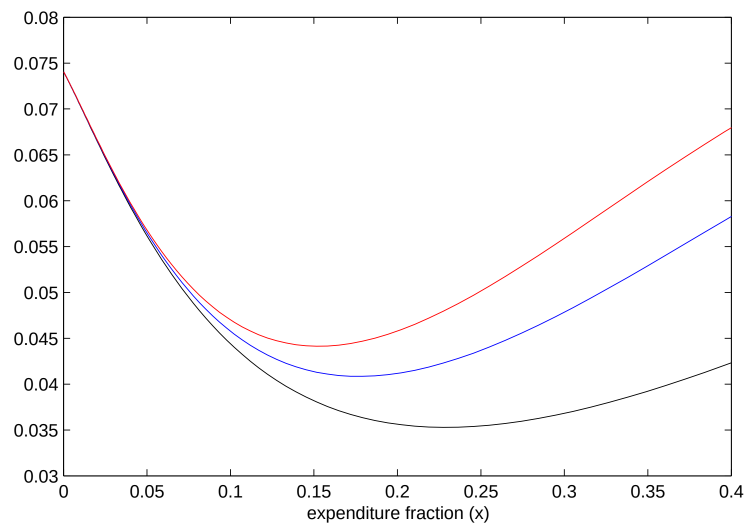

The lowest line is the instantaneous expected excess return for the dividend claim for the benchmark parameters given in Table 1. The higher lines are the expected excess returns for claims to dividends and to some fraction of the profits from the commodity producing sector. For the middle line the fraction is 0.5; for the top line it is 1.0. The horizontal axis is the commodity expenditure fraction $x = PL/(PL + C)$.

The model presented in this section clearly provides only a limited depiction of the many factors influencing asset prices. Never the less, the estimation shows that financial data are informative enough about the crucial elasticities in the model to reveal the key mechanisms determining the effect of commodities risk on the real economy.

4 Commodity Risk and Welfare

This section addresses the primary topic of the paper: how the need for commodities affect agents' well-being. To start, I illustrate the behavior of real quantities in the benchmark specification. I then ask how alternative specifications of the economy – having more or less risk of high scarcity states – affect intertemporal welfare and marginal utility.

4.1 The Benchmark Case

Are high commodity prices bad? In the model economy, there are no frictions and hence no loss of real resources when commodity prices inflate. A first important conceptual goal, then, is to understand whether and how the real economy changes as ζ (or P or x) rises. and how this affects the representative agent.

Figure 5: *Consumption Moments*

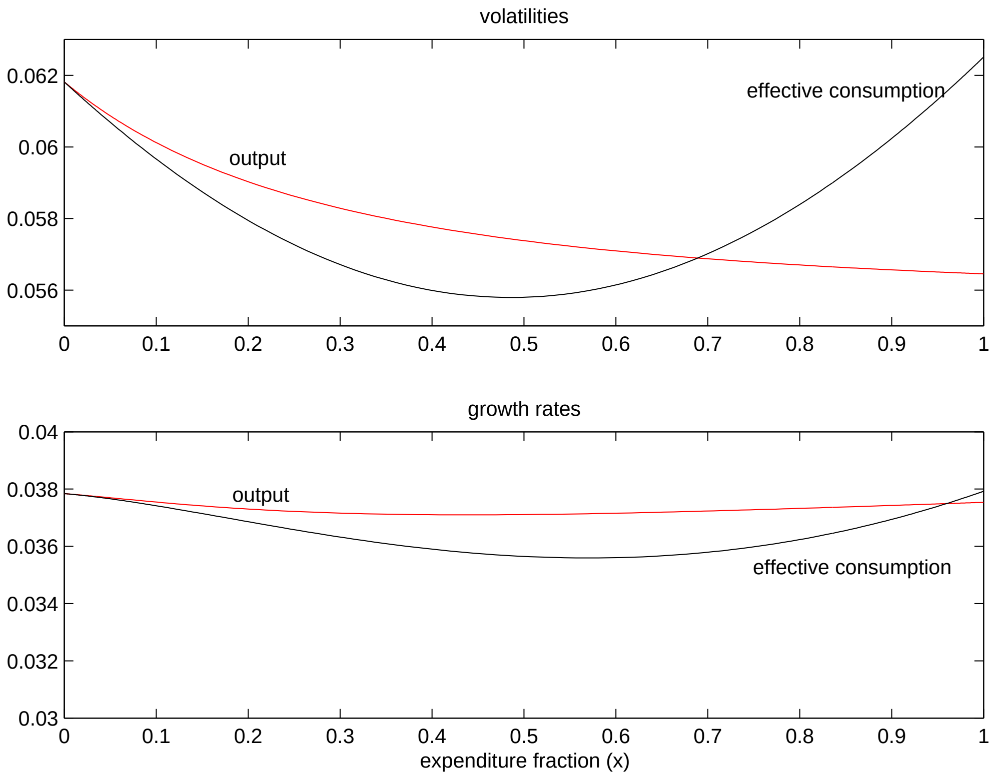

The figure shows the instantaneous first and second moments of non-commodity consumption, C, and effective consumption. $\tilde{C} = (\beta C^{\rho} + (1-\beta)L^{\rho})^{1/\rho}$, as a function of the commodity expenditure fraction, x for the benchmark parameterization of Section 3.

Figure 5 shows how the instantaneous growth rate and volatility of output, C, and the effective consumption basket, $\tilde{C}$, vary with x. In terms of $\tilde{C}$, utility is just of the power form: i.e. $\tilde{C}^{1-\gamma}/(1-\gamma)$. The benchmark specification has large γ, hence agents are very averse to $\tilde{C}$ risk. The top panel of the figure suggests then that *both* states of high and of low x will be disliked. While the U-shape in $\sigma_{\tilde{C}}$ is not large numerically, it turns out to be economically large. (By contrast, the growth rate effects in the bottom panel are economically insignificant.) The U-shape in effective consumption risk is due to the simultaneous decline of σ_C and the rise in σ_P, both discussed in the last section. Recall that $\tilde{C} = C\,(1+X)^{1/\rho}$ and that σ_X and σ_P are monotonically related.

Figure 6 shows the intertemporal component of the representative agent's welfare, V (equation (9)). The plot affirms that both high and low x states can be significantly worse in terms of this measure than mid-range ones and that the effect is extreme as x approaches one.[14] With very high curvature in utility, increases in risk have extremely persistent effects on expected utility growth It is worth emphasizing that these large changes in V with x are *independent* of how the change in x occurs. That is, positive productivity shocks have the exact same effect as negative commodity output shocks.

Figure 6: *Intertemporal Welfare*

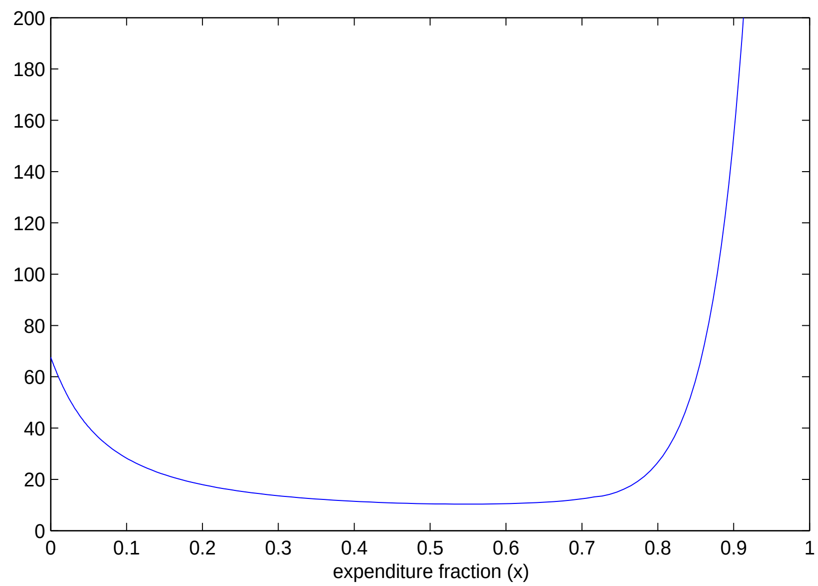

The figure shows the intertemporal component of the representative agent's value function as a function of the commodity expenditure fraction, x for the benchmark parameterization.

While Figure 6 does show that extreme increases in commodity scarcity can be very bad for welfare, it also shows that the variation in V is *negative* over the historically relevant range of x, i.e., $x < 0.25$. Here the risk gradient shown in Figure 5 is downwards.

[14]The function does stay finite. The plot is truncated for clarity.

This raises the question of whether high P-states are bad in the sense that agents would pay to insure against them.

The answer to this is shown in Figure 7. The figure plots the ratio of risk-neutral probability to true probability for deciles of the commodity price distribution at a range of time-horizons, T. These ratios correspond to the (forward) price agents would pay per unit probability for a contract which pays 1.0 if P_T is in the indicated decile of its (physical) distribution. For example, at the 10-year horizon, there is a 5 percent probability that $P_T > 2.3\,P_0$. But, as indicated by the rightmost point on the graph, agents would pay almost nothing (4 percent of expected value) for a bet on this outcome. By contrast, contracts paying off in low price states are very valuable. Agents in the benchmark economy are actually far more concerned with *falling* commodity prices than rising ones.

Figure 7: *Commodity Price Insurance*

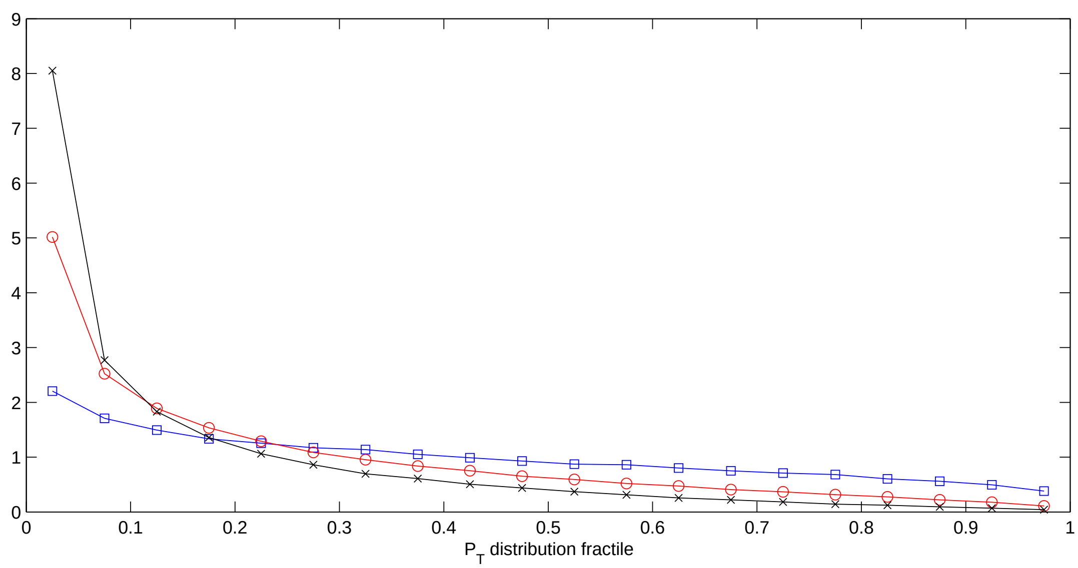

The figure shows the ratio of risk-neutral probability to true probability for five-percent fractiles of the commodity price distribution for the benchmark specification estimated in Section 3. Lines plotted as squares are for 1-year horizon, T; circles are T = 5-years; crosses are T = 10-years.

To understand this result, consider how marginal utility adjusting affects the growth rate of the exogenous process for scarcity. Figure 8 shows both the physical and risk-neutral drifts μ_ζ and μ_ζ^{RN}. The former is almost zero by design in this specification. The latter is negative for $x < 0.72$. This is because marginal utility is initially more (negatively) exposed to Z shocks than it is (positively) exposed to Q shocks when commodity expenditure is low. The risk-adjustment to μ_P works the same way. Even though commodity prices are expected to rise, the risk-neutral expected rise is lower (and actually initially negative). In a nutshell, commodity prices are procyclical until a very high level of scarcity obtains.[15]

Figure 8: *Physical and Risk-Neutral Growth Rates*

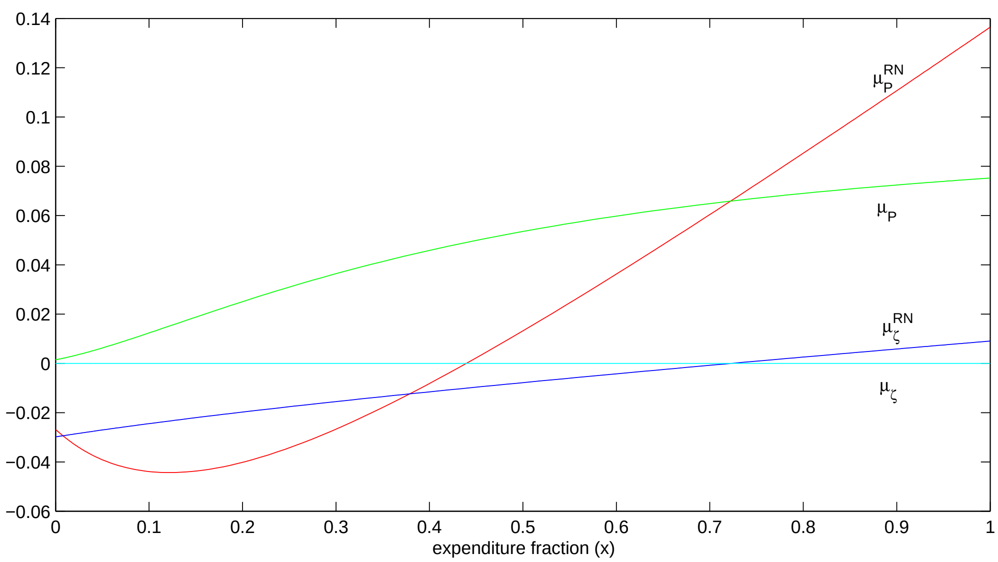

The figure shows the true and risk-neutral expected growth rates of commodity prices and of relative scarcity $\zeta = Z/Q$ for the benchmark specification estimated in Section 3. The horizontal axis is the commodity expenditure fraction $x = PL/(PL + C)$.

Returning to insurance prices, agents would indeed want to hedge against rising commodity prices *conditional* on reaching such a level. However the risk-neutral probability of

[15]The conclusion here is equivalent to the assertion that forward commodity prices should be below expected spot prices. This is consistent with empirical evidence of unconditional backwardation for (nominal) oil futures presented in Litzenberger and Rabinowitz (1995). Recently, strong positive commodity risk premia have been documented by Gorton and Rouwenhorst (2006) and Hong and Yogo (2009).

prices rising that high is low except at very long horizons, and even then it is significantly less than the true probability of doing so. Unconditionally, then, the Arrow-Debreu price of these states is still below their expected value. This conclusion is robust to variation in the *true* growth rates of the Z and Q. That is, even if the true probability of attaining states of high scarcity is raised (e.g. by putting $\mu_Q < \mu_Z$) the *difference* between true and risk-neutral growth rates is is not changed, nor is the ratio of true to risk-neutral densities.[16]

Summarizing, the benchmark case implies a very different nature of commodity risk in the range of extreme scarcity than in other times. It is important to note that the predictions for the high-x range are robust in a statistical sense even though there is no data in this range. The U-shape risk profile of the effective consumption bundle $\tilde{C}$ is common to all specifications having low ρ and high γ. As shown in the last section, the statistical evidence in asset market data for both is strong. That said, the normative conclusion is that the states in which rising prices are harmful are not worth insuring against.

4.2 The Real Cost of Commodity Risk

Following Lucas (1987), another natural way to quantify the economy's exposure to commodity risk is to consider the welfare effect of varying the degree of that risk. The real cost of "commodities cycles" is exactly analogous to the real cost of "business cycles", which is typically assessed by varying the volatility of productivity shocks.

In the present economy, commodity risk is driven by σ_ζ, the volatility of the scarcity ratio. Commodity price volatility scales directly with σ_ζ. If ζ were constant, then the allocation ℓ would be as well, and the model would reduce to a one-factor endowment economy.

The volatility σ_ζ declines if *either* the correlation, θ, between Z and Q shocks is raised

[16] An earlier draft of this paper gives example calculations under a variety of assumptions, including cointegration, about the two drift rates.

or if the volatility of commodity output, σ_Q, is moved toward the value $\theta\sigma_Z$ (when θ is held fixed). Interestingly, the comparative statics of these two alternatives yield different conclusions.

Consider two variations on the benchmark model: case **A** raises θ to 0.6634 holding σ_Q fixed; and case **B** lowers σ_Q from 0.0625 to $\theta\sigma_Z = 0.0365$ where σ_ζ attains its minimum when $\theta = 0.5776$. Both cases imply a 10 percent decline in σ_ζ from the benchmark. Varying the stochastic parameters does not change the allocation function $\ell(\zeta)$ or $x(\zeta)$. Thus the commodity volatility function $\sigma_P(x)$ in both cases is identical to the benchmark – but 10 percent lower – for all values of x.

Figure 9 shows the intertemporal component of welfare for these two cases. Case **B** has everywhere strongly improved welfare, yet case **A** has worsened it. The reason for this is that the former case has lowered the volatility of effective consumption, $\tilde{C}$, and the latter has raised it. As we have seen, welfare is extremely sensitive to this risk.

Figure 9: *Welfare and Commodity Risk*

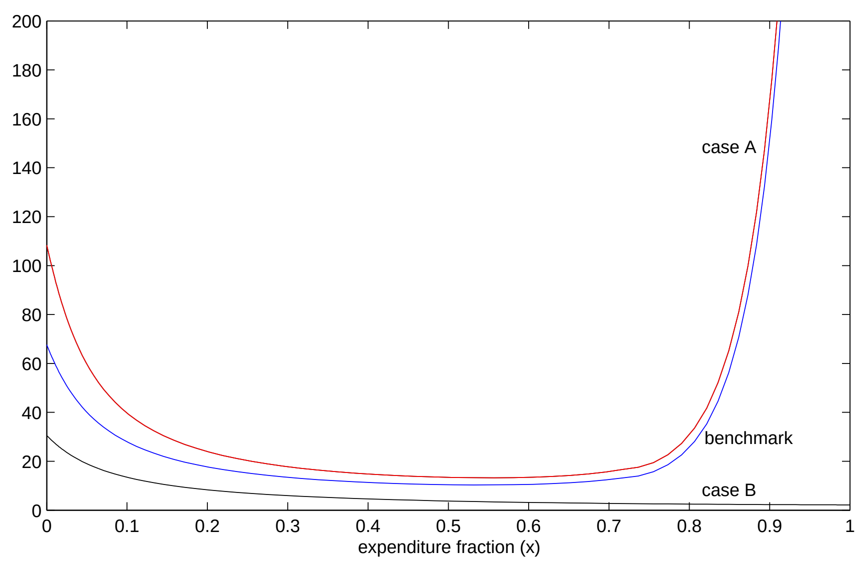

The figure shows the intertemporal component of welfare for two cases in which commodity volatility is 10 percent lower than in the benchmark estimation. Case A raises the correlation between the exogenous factors Q and Z. Case B lowers σ_Q. The horizontal axis is the commodity expenditure fraction $x = PL/(PL + C)$.

To understand the effect of the parameters on effective consumption risk, consider first just the effect on σ_C. Recall that the diffusion terms for C are of the form

$$\delta_1 \sigma_Z \; dB^Z + (1 - \delta_1) \sigma_Q \; dB^Q$$

where $\delta_1 = \delta_1(\zeta)$ does not depend on the stochastic parameters. Hence the variance σ_C^2 is

$$\delta_1^2 \sigma_Z^2 + (1 - \delta_1)^2 \sigma_Q^2 + 2\theta \delta_1 (1 - \delta_1) \sigma_Z \sigma_Q.$$

For the cases under consideration, $0 < \delta_1 < 1$ which implies that the partial derivatives of σ_C^2 with respect to *both* θ and σ_Q are positive. Output (of non-commodity goods) is always less volatile when commodity output is. However, as discussed above, production of C-goods benefits from a diversification effect which is larger when its inputs are *less* correlated. To take the extreme case, if Z and Q were perfectly correlated, there would be almost no commodity price risk, but C risk would actually be higher than when $\theta < 1$.

Effective consumption has a second component, $(1+X)^{1/\rho}$, multiplying C. This term becomes unambiguously less volatile when σ_ζ declines. So, when σ_Q is lowered both of the components of $\sigma_{\tilde{C}}$ shrink. When instead θ is raised, this term mitigates the increase coming from σ_C. However, it does not reverse it. See Figure 10.

This analysis highlights an important lesson regarding commodity risk, namely, that the direct effect of commodity price uncertainty on welfare is secondary to its indirect effect on production. Indeed it may be the case that $\sigma_{\tilde{C}}$ and σ_P move in opposite directions.[17] The normative implication is that, although alleviating supply shocks makes sense as a policy objective, "smoothing" commodity prices by e.g. managing inventories to off-set demand shocks may be misguided.

[17]As illustrated above, this happens in response to changes in θ. In addition, when σ_Q is lowered below $\theta \sigma_Z \; \sigma_\zeta$ – and hence σ_P – starts to increase. Yet effective consumption volatility continues to decline and welfare increases monotonically.

Figure 10: *Effective Consumption Risk*

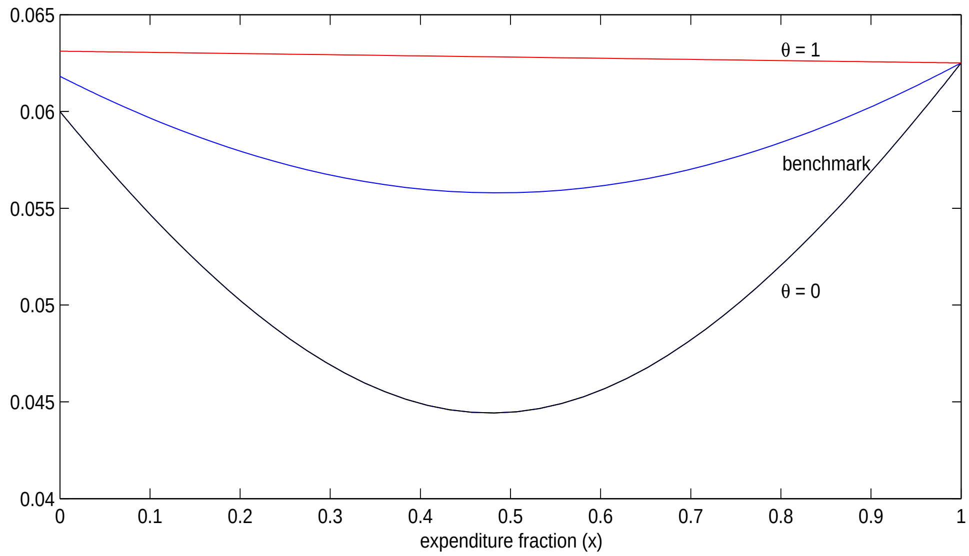

The figure shows effective consumption volatility as a function of the correlation between the exogenous factors Q and Z. Other parameters are as per the benchmark specification. The horizontal axis is the commodity expenditure fraction $x = PL/(PL + C)$.

4.3 Flexibility in Production

The benchmark model has unit elasticity of substitution between commodity and non-commodity inputs to production. In the general case given by (1) one can characterize commodity dependence in production via the substitution ratio

$$\left(\frac{NF_N}{ZF_Z}\right) = \frac{1-\eta}{\eta}\ \left(\frac{Z}{N}\right)^{-\nu}.$$

The left side is the percentage increase required in Z to offset a one percent drop in commodity inputs. The right side identifies two distinct components of this requirement. The first, $(1-\eta)/\eta$, the relative weight in the production function, is static. The second term depends upon the current level of scarcity. When $\nu < 0$ the requirement for commodities increases with scarcity; when $\nu > 0$ it decreases. A natural hypothesis, then, is that the response of the economy to high commodity-price states should depend on ν.

Figure 11 shows the effect on the intertemporal component of welfare when ν is increased from 0.0 to 0.2. (The horizontal axis is now $\log(\zeta)$ instead of x because $x(\zeta)$ changes when ν is changed.) In line with intuition, there are indeed substantial increases in welfare (lower V) for high scarcity states. The mechanism at work is precisely the substitution described above: output becomes increasingly *positively* sensitive to ζ shocks as ζ increases when $\nu > 0$ because the marginal product of Z rises and that of N declines. Since this coincides with rapidly rises commodity expenditure, there is a strong positive effect on the real interest rate coming from the positive covariance of the two components of marginal utility. The higher real rate indicates expected declines in marginal utility, reflected in the lower (negative) utility growth measured by V.

Figure 11: *Welfare and Commodity Risk*

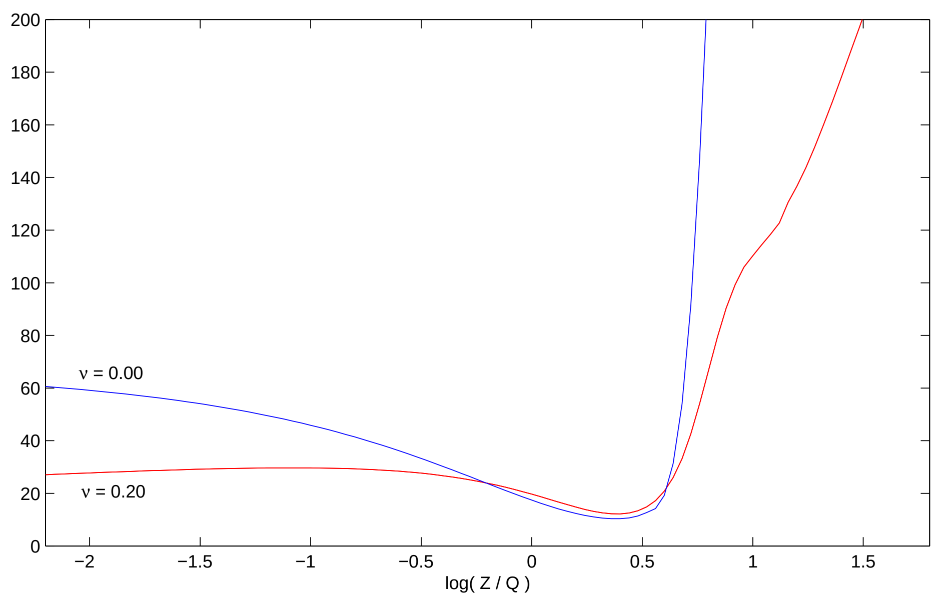

The figure shows the intertemporal component of welfare, V, for both the benchmark case with $\nu = 0.0$ and for the same parameters with $\nu = 0.2$. The horizontal axis is the log ratio of the exogenous state variables.

Interestingly, the figure also indicates significant welfare gains relative to the benchmark model in states of *low* scarcity. These gains too can be attributed to an effect of the substitution dynamics. For $\nu > 0$ the volatility of output, C, declines with ζ because the exposure to productivity shocks, Z, diminishes. The lower volatility induces a higher riskless rate via the precautionary savings channel. The middle of the graph indicates only minor differences in welfare over a range of ζ that corresponds to approximately $0.15 < x < 0.85$. In this range, the dynamics of the effective consumption basket are similar in the two models. The normative implication is that, while expending resources to increase production flexibility may not be welfare enhancing in this region, having the option to do so in either direction might be extremely valuable. This observation motivates consideration of versions of the economy in which the possibility of technological evolution affords such an option.

4.4 Technological Evolution

The model extension described in Section 2.7 allows for the possibility that the economy may be willing to sacrifice productivity growth in order to alter the input usage ratio $\eta/(1-\eta)$. If indeed it is worthwhile to alleviate the need for commodities, then the model offers a quantification in terms of the degree of growth sacrificed.

In this extension, η becomes an additional state variable. Figure 12 shows V as a function of both x and η for the benchmark case before allowing η to change.[18] Note the bowl-shape: previous sections have discussed the U-shape in the x direction, however the U-shape on the η axis is a new feature. This implies that utility growth considerations favor intermediate values of η. This can be counterintuitive: when commodities are relatively scarce (x is high) one might think that higher η (lower commodities share) would always be desirable. This indeed is the case for current utility, U, but not for V, e.g., when η is above one half. The reason is that commodity volatility – and hence effective consumption volatility – decreases with η.

Now consider what happens when the economy has the technology described in in

[18]For this exercise, it is necessary that V be finite for all values of η. The benchmark case as estimated does not satisfy this condition. However it is satisfied if the subjective discount rate ϕ is raised from 0.00 to 0.02 – a change which is not statistically rejectable.

Figure 12: *Welfare as a function of Commodities Usage Rate in Production*

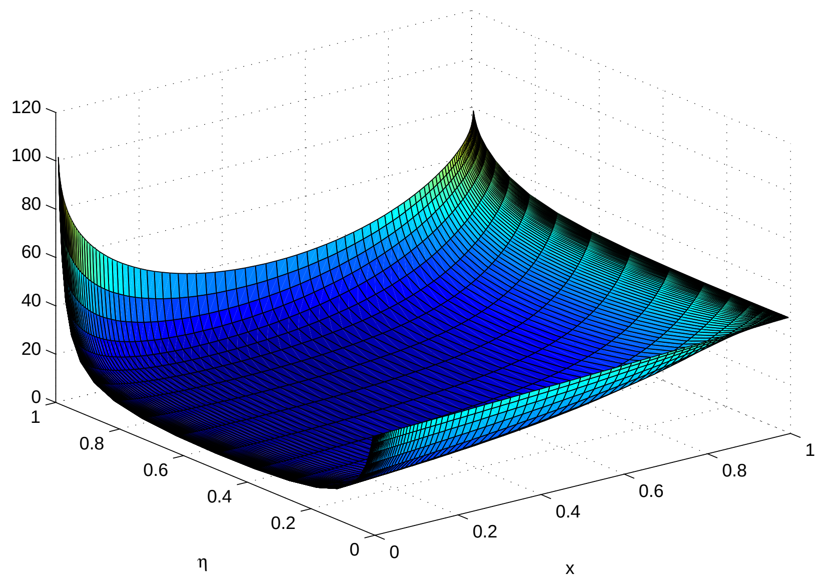

The figure shows the intertemporal component of welfare, V, for the baseline parameter case as a function of the expenditure ratio, x, and the non-commodities share in production, η. The subjective discount rate ϕ is set to 0.02.

Section 2.7. Agents can deplete the productivity-enhancing factor H, which lowers the growth rate of the input Z to the production function, lowering the growth rate of output, C. The model nests the case that agents may choose to lower growth even if it has no effect on the production function. (Put $b = 0$.) Arguably, this is the appropriate benchmark for studying the effect of technological evolution, since presumably any economy has the ability to not utilize exogenous increases in productivity.

The top line in Figure 13 shows the optimal investment policy for the benchmark case, setting $b = 0$, $k = 1$. The latter means that the policy units are the amount by which the growth rate is lowered. Indeed it is the case in this model that it can be optimal to lower Z-growth when relative scarcity is high, even with no technological benefit. This is a direct consequence of the fact that V increases with x for high values of x. The policy quantifies the intertemporal benefit of pushing Z (hence x) down in units of forgone future output. The size of the policies can become extremely large, effectively ensuring

Figure 13: *Technological Evolution: Policies*

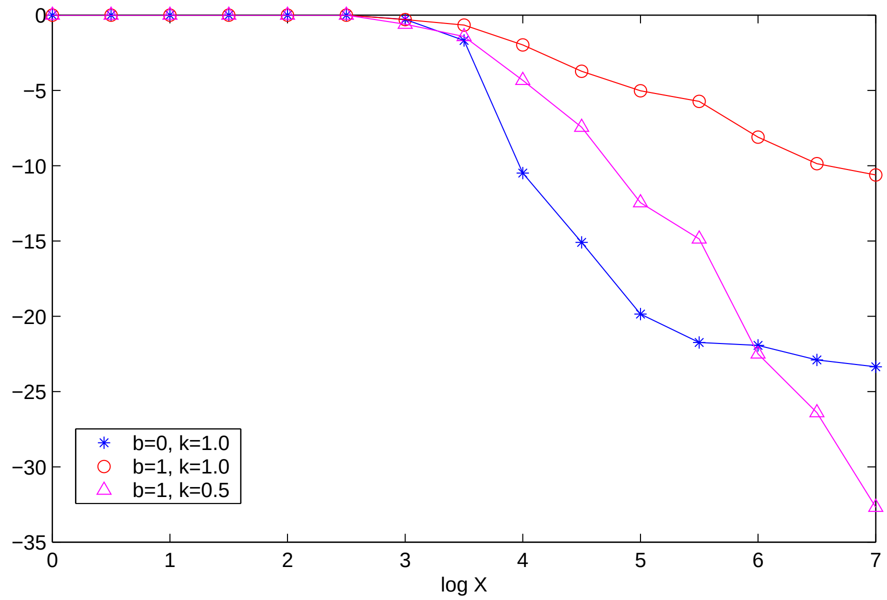

The figure shows the optimally chosen quantity i which alters the drift rate of $log(\eta/(1-\eta)$ by bi and lowers the drift of dC/C by $k\,|i|$. The plot shows three possible technologies described by the pair (b, k). The horizontal axis is the log of the commodities expenditure ratio.

that the high-x region becomes inaccessible.

The lower two lines in the figure show how the policy changes when the impact on the production function is turned on ($b = 1$). When $k = 1$, the policy is less extreme. This reflects the fact that additional welfare benefits are now achieved by altering the production function, hence less output needs to be foregone to attain a given benefit level. For $k = 0.5$, the output cost has been cut in half, resulting in policies that approximately double. All policies go to zero away from high-scarcity states. (Note that $\log X = 3$ corresponds to an expenditure fraction of about $x = 0.95$.) It is not optimal for agents to lower η when x is low – despite the negative slope of V – because here the slope of V in the x direction is positive. So sacrificing growth in Z (lowering x) is too painful.

The surprise in the picture is the sign of the policies. Negative values indicate that welfare is increased under conditions of scarcity by lowering $\eta/(1-\eta)$, or *raising* the commodity usage, $(1-\eta)$, in production. This reflects the negative slope of the V

surface along the η dimension noted earlier.[19] The natural intuition that agents would use the technology to lower commodity usage is thus belied in this region of the state space. (The policy does turn positive, though, for lower-η values with high scarcity.) The influence of η on commodity price volatility is actually the dominant motivation. This is a direct consequence of commodities being in the utility function.

How does technological evolution affect welfare? Figure 14 shows the log levels of V achieved for each case. The interesting observation here is that – even though the policy actions themselves are confined to states of extremely high scarcity, the benefits are not. With x as low as 0.3 the $b = 0$ policy represents an improvement of about 0.01 (one percent) over the benchmark. And the $b = 1, k = 0.5$ case yields an improvement of about 0.01 over $b = 0$. Viewing these amounts as percentages of global wealth that would be willingly exchanged in order to acquire the technology, they are not negligible, and quickly grow in size as x increases.

Over all, including technological evolution in the model does not change the basic conclusions that applied to the benchmark case. For the cases examined here, the welfare function retains its U-shape in x. As a consequence, in typical times (e.g. $x < 0.25$) the intertemporal component of welfare increases (and marginal utility decreases) with commodity prices.

5 Conclusion

Much has been written about the possibility of a future in which the prices of raw commodities rise ever higher. Analyzing the economic consequences of this possibility is an important step in understanding whether and how resources should be expended to insure against or prevent such an outcome.

This paper has used asset-pricing techniques to assess the nature of the economy's dependence on commodities. I build a representative-agent model in which the con-

[19] The sign of the policy is unidentified in the $b = 0$ case. It is plotted as negative for ease of comparison.

Figure 14: *Technological Evolution and Welfare*

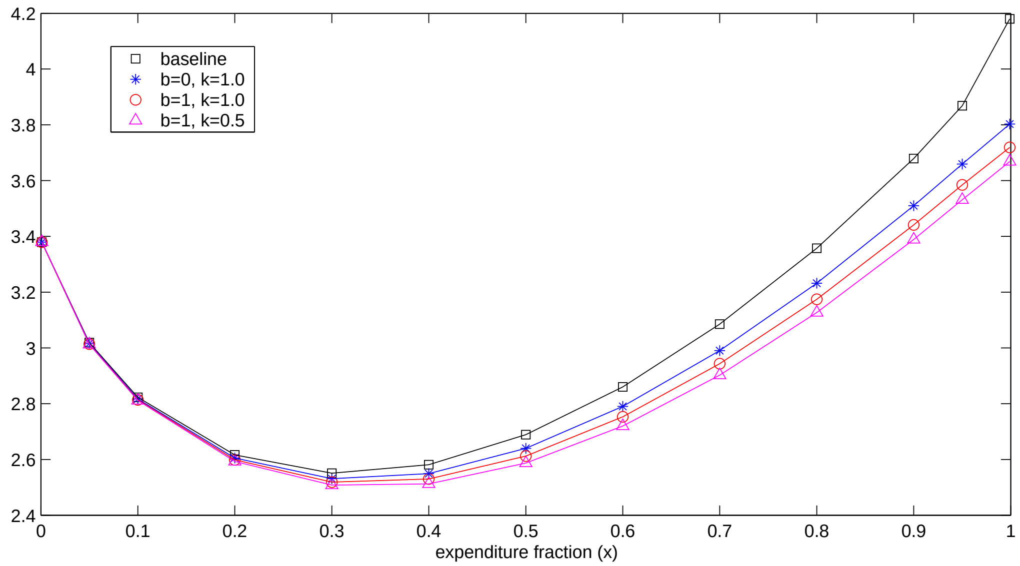

The figure shows the log of the intertemporal component of welfare, V, for the benchmark case and for three cases allowing for policies to alter the production function and the growth rate of output C. The horizontal axis is the commodities expenditure fraction.

sumption basket includes both finished goods (non-commodities) and raw primary goods (commodities). The model yields a rich description of the joint dynamics of commodities prices, equity values, and interest rates. I then use 40 years of data to back out the crucial preference and technology parameters from stock and commodities returns. I show that the intertemporal component of welfare can be represented as the "price-dividend ratio" of a claim to the composite consumption bundle. I compute this quantity under a variety of assumptions about the long-run behavior of the economy.

A basic observation is that there *are* real effects of high commodity prices on utility even though there are no frictions in the model. Prices are driven by the the scarcity of commodities relative to other factors of production. This ratio is the sole determinant of the rate of growth and volatility of the effective consumption basket. So, for the intertemporal component of welfare, the impact of rising scarcity is the same whether it occurs through low growth of commodity supply or high growth of other factors. The model implies that states of high scarcity are indeed bad. As prices rise two things happen: first commodity volatility rises; second the exposure of consumption to this volatility rises.

While all versions of the model imply that these outcomes are extremely bad, they all also imply that they are a long way off. In the baseline case, for example, welfare actually improves with the commodity price level until the commodities share in consumption is over 70 percent – many times its historical level. Locally, the welfare improvements are also driven by second moments: the volatility of (non-commodity) production benefits from a diversification effect as output becomes less reliant on other inputs. Even with a high estimated correlation between the factors and a low commodity input share this effect is large because agents are extremely averse to shocks to this component of consumption. As a result, the risk premium associated with high-price states is actually positive meaning that (except for extremely high levels) agents would not be willing to pay to insure against them.

To assess the role of the different mechanisms in the model, I conduct several comparisons. First, commodity risk, like business-cycle risk, can be quantified via the benefits of mitigating the exogenous exposure. Very large welfare gains do indeed result from lowering the variability of commodity supply. However increasing the correlation of that supply with other factors (e.g., through procyclical management of buffer stocks) can actually do more harm than good by decreasing the diversification effect. Next, I show that increasing the elasticity of substitution in production can raise welfare substantially in states of high or low scarcity. Finally, when the model is extended to allow for technological evolution, the ability to avoid states of extreme scarcity does improve welfare at intermediate ranges even when the policy actions only take place at much higher levels. Interestingly, sacrificing growth to lower prices can be optimal even with no change in technology, and the direction of optimal technological change may actually increase its usage in production.

Overall, the analysis shows that the information in financial markets can provide important perspective on a broadly relevant macroeconomic topic. Future research may extend the basic framework here both to incorporate additional features of commodities themselves and to explain further features of asset dynamics.

References

Basak, Suleyman, 1999, On the fluctuations in consumption and market returns in the presence of labor and human capital: An equilibrium analysis, *Journal of Economic Dynamics and Control* 23, 1029–1064.

Baumeister, Christiane, and Gert Peersman, 2009, Sources of the volatility puzzle in the crude oil market, Ghent University working paper.

Berndt, Ernst R., and David O. Wood, 1979, Engineering and econometric interpretations of energy-capital complementarity, *American Economic Review* 69, 342–354.

Bodenstein, Martin, Christopher J. Erceg, and Luca Guerrieri, 2011, Oil shocks and external adjustment, *Journal of International Economics* 83, 163–184.

Carlson, Murray, Zeigham Khokher, and Sheridan Titman, 2009, Equilibrium exhaustible resource price dynamics, *Journal of Finance* forthcoming.

Casassus, Jaime, Pierre Collin-Dufresne, and Bryan R. Routledge, 2005, Equilibrium commodity prices with irreversible investment and nonlinear technologies, Carnegie Mellon University working paper.

Cochrane, John, Francis Longstaff, and Pedro Santa Clara, 2008, Two trees, *Review of Financial Studies* 21, 347–385.

Deaton, Angus, and Guy Laroque, 1996, Competitive storage and commodity price dynamics, *Journal of Political Economy* 104, 896–923.

Driesprong, Gerben, Ben Jacobsen, and Benjamin Maat, 2008, Striking oil: Another puzzle, *Journal of Financial Economics* 89, 307–327.

Ferderer, J. Peter, 1996, Oil price volatility and the macroeconomy, *Journal of Macroeconomics* 18, 1–26.

Gorton, Gary, and K. Geert Rouwenhorst, 2006, Facts and fantasies about commodity futures, *Financial Analysts Journal* 62, 4768.

Hamilton, James D., 2009, Understanding crude oil prices, *Energy Journal* 30, 179–206.

Hong, Harrison, and Motohiro Yogo, 2009, Digging into commodities, Princeton University working paper.

Jacobsen, Ben, Ben Marshall, and Nuttawat Visaltanachoti, 2009, Return predictability revisited, Massey University working paper.

Kilian, Lutz, and Cheolbeom Park, 2009, The impact of oil price shocks on the u.s. stock market, *International Economic Review* 50, 12671287.

Litzenberger, Robert H., and Nir Rabinowitz, 1995, Backwardation in oil futures markets: Theory and empirical evidence, *Journal of Finance* 50, 1517–1445.

Lochstoer, Lars A., 2009, Expected returns and the business cycle: Heterogeneous goods and time-varying risk aversion, *Review of Financial Studies* 22, 5251–5294.

Lucas, Jr., Robert, 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.

——— , 1987, *Models of Business Cycles* (Basil Blackwell: New York).

Newey, Whitney K., and Kenneth D. West, 1987, Hypothesis testing with efficient method of moments estimation, *International Economic Review* 28, 777–787.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel, 2007, Housing consumption and asset pricing, *Journal of Financial Economics* 83, 531–569.

Pindyck, Robert S., 1980, Uncertainty and exhaustible resource markets, *Journal of Political Economy* 88, 1203–1225.

——— , 1994, Inventories and the short-run dynamics of commodity prices, *RAND Journal of Economics* 25, 141–159.

Polgreen, Linnea, and Pedro Silos, 2009, Crude substitution: The cyclical dynamics of oil prices and the skill premium, *Journal of Monetary Economics* 56, 409–418.

Santos, Tano, and Pietro Veronesi, 2006, Labor income and predictable stock returns, *Review of Financial Studies* 19, 1–44.

Solow, John L., 1987, The energy-capital complementarity debate revisited, *American Economic Review* 77, 605–614.

Stiglitz, Joseph E., 1976, Monopoly and the rate of extraction of exhaustible resources, *American Economic Review* 66, 655–661.

Sundaresan, Suresh M., 1984, Equilibrium valuation of natural resources, *Journal of Business* 57, 493–518.

Wei, Chao, 2003, Energy, the stock market, and the putty-clay investment model, *American Economic Review* 93, 311–323.

Yogo, Motohiro, 2006, A consumption-based explanation of expected stock returns, *Journal of Finance* 61, 539–580.