Essays on International Macroeconomics

by
Daniel Morgan Rees
B.A., B.Com, University of Sydney (2005)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2013

Author ....
Department of Economics
January 31, 2013

Certified by ....
Ivan Werning
Professor of Economics
Thesis Supervisor

Certified by ....
Guido Lorenzoni
Associate Professor of Economics
Thesis Supervisor

Accepted by ....
Michael Greenstone
3M Professor of Environmental Economics
Chairman, Departmental Committee on Graduate Studies

Essays on International Macroeconomics

by

Daniel Morgan Rees

Submitted to the Department of Economics
on January 31, 2013, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This thesis examines the impact of terms of trade shocks on commodity-exporting small, open economies. The first chapter examines whether households, firms and policymakers in these economies can distinguish between temporary and permanent commodity price shocks. I find that they are largely unable to do so. In fact, my model suggests that the expected future path of commodity prices following a temporary price shock is almost identical to the expected future path of commodity prices following a permanent price shock. However, I also find that these information frictions reduce the magnitude of business cycle fluctuations, contrary to popular belief.

In the second chapter I describe optimal monetary policy in an environment where agents cannot directly observe whether commodity price shocks are temporary or permanent and where an economy's non-commodity sector features a learning-by-doing externality. I find that under optimal monetary policy the non-commodity sector contracts by more during a transitory commodity price boom under incomplete information than it does under full information, but by less during a permanent boom. I also examine the performance of simple monetary policy rules. A policy of responding strongly to deviations of home-produced goods inflation from target with a modest response to changes in the nominal exchange rate comes close to replicating the welfare outcomes of optimal policy. In contrast, an exchange rate peg generally produces large welfare losses.

The third chapter, co-authored with my classmate Patricia Gomez-Gonzales, examines the consequences of changes in the volatility of commodity price shocks on commodity exporters. We first demonstrate the existence of time-varying volatility in the terms of trade of a selection of commodity-exporting small open economies. We then show empirically that increases in terms of trade volatility trigger a contraction in domestic consumption and investment and an improvement in the trade balance in these economies. Finally, we construct a theoretical model and demonstrate that it can replicate our empirical results.

Thesis supervisor: Ivan Werning
Title: Professor of Economics

Thesis supervisor: Guido Lorenzoni
Title: Associate Professor of Economics

Acknowledgement

I would like to thank my advisors Guido Lorenzoni and Ivan Werning. Guido and Ivan provided me with a wealth of guidance and insight over the past few years. At every stage of my dissertation, they managed to find just the right mix of encouragement and prodding, always forcing me to justify my assumptions and challenging me to look at problems in new and interesting ways. Their lessons made an enormous contribution to the papers in this thesis, and will continue to influence the way that I think about economics long after I leave M.I.T.

I would also like to thank the participants at the macroeconomics and international lunches at M.I.T. for their valuable feedback and constructive criticism. In particular, this thesis has benefitted from the comments and suggestions of Daron Acemoglu, Marios Angeletos, Arnaud Costinot, Anna Mikusheva and Athanasios Orphanides.

I am enormously grateful to the classmates with whom I have shared my PhD experience. In particular, I would like to thank Brad Shapiro and Maria Polyakova for helping me to survive my first year in Boston. Isaiah Andrews, Patricia Gomez-Gonzales, Juan Passadore and Su Wang were always willing to take time out of their busy schedules to run over my projects and answer my questions. And Adam, Annalisa, Ben, Dan, Greg, Henry, Mike, Miikka, Sarah, Sally, Stefanie and Xiao were always happy to chat about anything, so long as it had nothing to do with macroeconomics. The past four years wouldn't have been the same without you guys.

My graduate research would not have been possible without the financial support of the Reserve Bank of Australia's Postgraduate Study Allowance scholarship. I would also like to thank my RBA co-authors Dan Andrews, Jeremy Lawson and, especially, Mariano Kulish, as well as my University of Sydney advisors Tony Aspromourgos and Graham White for giving me such a fantastic grounding in how to conduct academic research.

I must also thank my parents, David and Helena and my sister Josephine for their love and support. They have always been there to listen to my complaints and disappointments, to sympathize, and to encourage me to keep my chin up. I would not have had the determination to work my way through a PhD without your example.

Few people would consider an offer to leave a good job and comfortable life in order to move half-way around the world and work long hours to support an inattentive and often grumpy partner who spends his days and nights staring at pages of Matlab code a tempting one. My wife Li did, however, and for that I will forever be thankful.

Contents

Chapter 1

Terms of Trade Shocks and Incomplete Information in Small Open Economies

> "Commodity prices are typically ... volatile ... and many of the price movements last just long enough to convince investors and governments that "this time it is different". And there is always a chance that some day it will be different. In the intervening period, long-range investments may have been set in train, new facilities built, and workers relocated. ... If prices stay high (or low) for a sufficiently long time, these reallocations of capital and labour could well be warranted and yield valuable returns.... The trouble is that businesses, households, and policy-makers often get caught out. ... The inherent difficulty associated with predicting how long a boom (or bust) might last, and how high (or low) prices might go, makes the process extremely risky. Critics worry that a commodity-based economy will constantly find itself in motion, never quite settling down."
>
> - John Murray, Deputy Governor of the Bank of Canada, 6 May 2010.

1.1 Introduction

The terms of trade of many commodity-producing small open economies appear to be characterized by a succession of slow-moving long-run trends, augmented by high-frequency transitory fluctuations. As an illustration, Figure 1.1 shows the terms of trade - defined as the ratio of export prices to import prices - for six small open economies between 1961 and 2011. For each country, the solid line represents the level of the terms of trade in logs, while the dashed line shows the HP-filtered trend.[1] Although the exact patterns differ between countries, each economy has experienced periods in which the trend terms of trade persistently decreased as well as periods in which the trend terms of trade persistently increased. Changes in the trend terms of trade are often large. For example, the trend terms of trade decreased by around 50 per cent in Mexico during the 1980s and increased by over 50 per cent in Australia during the 2000s. Deviations from the trend are also substantial. During the early 1970s, New Zealand's terms of trade was at times 30 per cent above its trend level.

The presence of both persistent and transitory movements in the terms of trade matters because the optimal response to a terms of trade shock depends upon the persistence of the shock. A positive terms of trade shock is similar to a positive income shock in that it allows an economy to increase absorption without a corresponding increase in production. A simple permanent income model would suggest that consumption smoothing households will respond to a temporary increase in the terms of trade by saving some of the windfall and increasing consumption by the annuity value of the shock. In contrast, a permanent increase in the terms of trade will induce a larger immediate consumption response and a smaller increase in saving.[2]

[1] Although the data is quarterly, the trend was calculated using a smoothing parameter of 64,000 rather than the usual 1,600. This reflects the fact that commodity price cycles - which drive the terms of trade in these economies - are typically longer than business cycles.

[2] The response to temporary and permanent terms of trade shocks may differ from this simple permanent income example depending, for example, on consumers' willingness to substitute intertemporally and between tradeable and non-tradeable goods. Nonetheless, the key point that the optimal responses to transitory and permanent shocks differ is generally true.

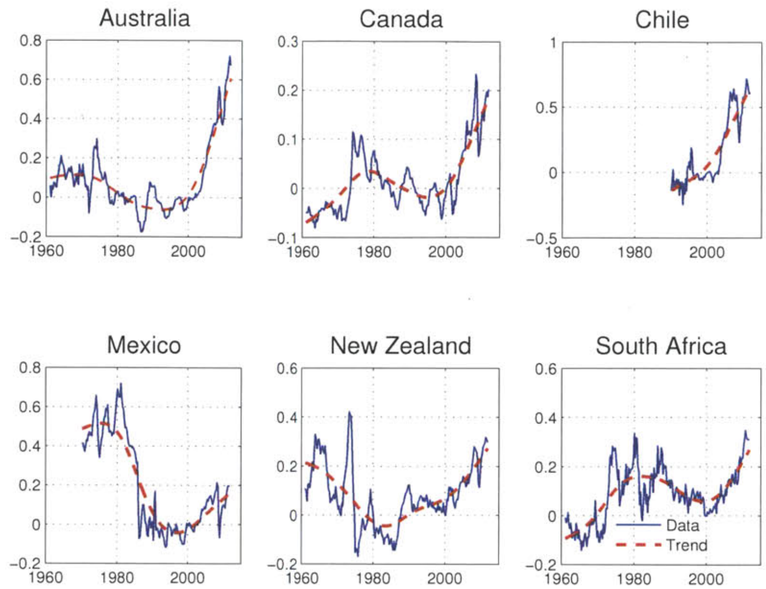

Figure 1.1: Terms of Trade for Selected Economies - 2003Q1 = 0

Notes: The data are quarterly. Trend calculated using an HP filter with a smoothing parameter of 64,000. The sources are national statistical agencies.

But in order for households and firms to respond in this manner to terms of trade shocks, they must first be able to identify which shocks are permanent and which are transitory. There is some reason to believe that they can do so. Unlike many other drivers of macroeconomic fluctuations - such as productivity or consumption preference shocks - the terms of trade are observable. Moreover, for many countries, changes in the terms of trade reflect broad global economic developments. For example, the increases in the terms of trade during the 2000s for the economies shown in Figure 1.1 were largely due to rising commodity prices driven by strong economic growth in countries such as China and India (Kearns and Lowe 2011, Plumb et al. 2012, Kilian and Hicks 2012). To the extent that agents recognize the underlying causes of changes in the terms of trade, it seems plausible to think that they are able to forecast the persistence of these changes accurately.

And yet there is also evidence which suggests that identifying the persistence of terms of trade shocks is difficult. Consider Figure 1.2. This shows the evolution of the terms of trade in Australia during the 2000s, as well as successive forecasts of the the terms of trade published by the Reserve Bank of Australia. It is striking how consistently the forecasts underestimated the persistence of increases in the terms of trade despite the fact that the

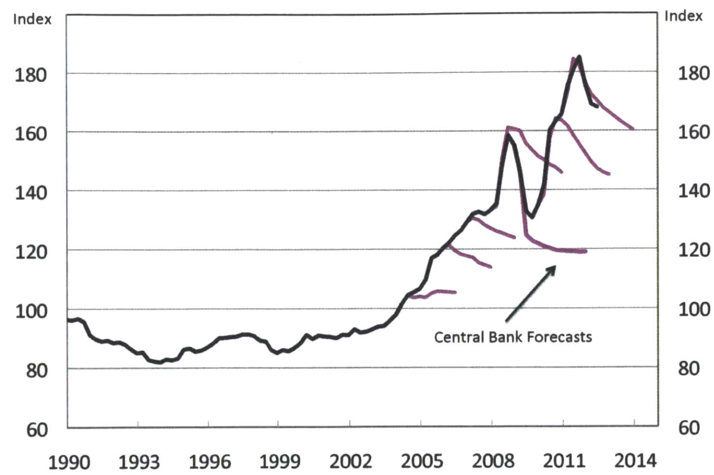

Figure 1.2: Forecasts of Australian Terms of Trade: 2003-04 = 100
Source: Plumb et al. 2012

underlying drivers of these increases were readily observable. Of course, a number of interpretations of Figure 1.2 are possible. It may be that central banks are bad at forecasting. Or this economy may have experienced a succession of positive, but temporary, shocks. The results of this paper, however, suggest that the patterns of Figure 1.2 reflect the fact that it is difficult to determine the persistence of terms of trade shocks.

To reach this conclusion, I augment an otherwise standard small open economy model to include incomplete information about the persistence of terms of trade shocks. I then estimate the model using Bayesian methods on Australian data. The results suggest that agents face considerable difficulties in untangling the persistence of terms of trade shocks. In fact, agents' beliefs about the future path of the terms of trade are largely independent of the type of terms of trade shock that hits the economy. Consequently, it should come as no surprise if the response of the economy to terms of trade shocks differs substantially from that implied by models in which agents are perfectly informed about the nature of these shocks.

As the quotation at the beginning of this paper illustrates, it is often argued that an inability to forecast accurately the persistence of commodity price shocks exacerbates macroeconomic volatility in small open economies. I demonstrate that, at least in the model used in this paper, this is not the case. That is because, while incomplete information about the persistence of terms of trade shocks increases the volatility of investment, it makes consumption, the trade balance and output less volatile than they would be if agents had full information.

This paper is related to several strands of literature. Most directly, it adds to the literature examining the effects of incomplete information about the composition of of structural shocks, as in Blanchard et al. (2012) and Angeletos and L'ao (2010). An application of this methodology to international macroeconomics is found in Boz et al. (2011), who estimate open economy real business cycle models for Canada and Mexico including uncertainty about the persistence of productivity shocks. This paper contributes to this literature in two ways. First, it provides empirical evidence of the existence of incomplete information about the persistence of an economically meaningful shock that has not previously been examined. Beyond this modest goal, the paper may also shed light on the pervasiveness of informational frictions about other shocks. Because terms of trade shocks are observable and can be rationalized in terms of broader economic developments, it seems plausible that households and firms have more information about these shocks than they do about other, unobserved, shocks. Consequently, estimates of the extent of uncertainty regarding the persistence of terms of trade shocks may well represent a lower bound of the uncertainty regarding other shocks.

The paper also contributes to the literature examining the determinants of business cycles in small open economies. Aguiar and Gopinath (2007) demonstrate that a small open economy business cycle model can better match the moments of macroeconomic variables in developing economies if augmented with persistent growth shocks, which accumulate over time, to accompany standard transitory mean-reverting productivity shocks. Boz et al. (2011) demonstrate that a similar result can be obtained with smaller structural shocks if one assumes that agents have incomplete information about whether shocks are temporary or permanent. An open question in both of these papers is why some economies should experience more persistent, or less observable, shocks than others. This paper provides a potential answer to this question by highlighting the difficulty of identifying the persistence of commodity price shocks. If developing economies are more exposed to commodity price movements than developed economies, then commodity price shocks could provide one explanation for the why the shock structure of developing and developed economies appears to differ.

This paper is also related to the literature on the response of small open economies to terms of trade shocks. The key theoretical papers in this literature are Harberger (1950) and Laursen and Metzler (1950), who use a simple Keynesian approach, and Sachs (1981), Obstfeld (1982) and Svensson and Razin (1983), who examine the response to a terms of trade shock in an intertemporal optimization setting. A number of papers have examined these relationships empirically. Otto (2003) constructs structural VAR models for a number of small-open economies to examine the effect of transitory terms of trade shocks on the trade balance. He concludes that a positive terms of trade shock generally leads to an improvement in the trade balance, consistent with a basic consumption-smoothing model of the current account in a model with only transitory shocks. Kent and Cashin (2003) separate countries into those whose terms of trade shocks are typically permanent and those whose terms of trade shocks are typically transitory. They find that a positive terms of trade shock leads to a deterioration in the current account in the former countries and an improvement in the latter. They argue that their results are also consistent with standard intertemporal approaches to the current account in which agents smooth their consumption in response to transitory shocks and adjust consumption and investment fully in response to persistent shocks.

Other papers have examined the importance of terms of trade shocks as a source of macroeconomic fluctuations. The empirical results here are mixed. Based on structural VARs estimated for a number of developing countries, Broda (2004) concludes that terms of trade shocks typically explain less than 10 per cent of output volatility in developing countries. In contrast, using a simulated real business cycle model, Mendoza (1995) finds that terms of trade disturbances explain 56 per cent of output fluctuations in developing countries and 33 per cent of output fluctuations in developed economies.

To some extent, the results of this paper reinforce those of the previous empirical literature. For example, I find that transitory positive terms of trade shocks cause an improvement in the trade balance while permanent positive terms of trade shocks cause a deterioration of the trade balance. However, as outlined in Blanchard et al. (2012), if agents have incomplete information about the persistence of shocks then it is not possible for an econometrician to identify permanent and transitory shocks in the data. Consequently, the finding that agents are largely unable to differentiate between permanent and transitory terms of trade shocks raises questions about the identification of these shocks in other papers.[3]

The paper proceeds as follows. Section 2 outlines the model and clarifies the information structure. Section 3 describes the estimation and summarizes the key results. Section 4 discusses the implications of the empirical results for the response of the economy to terms of trade shocks. Section 5 reports a series of robustness checks and Section 6 presents conclusions.

1.2 A Small Open Economy Model

The basic setup is a standard small open economy model with incomplete markets, similar to those in Mendoza (1995) and Aguiar and Gopinath (2007). I augment the model by assuming that agents are imperfectly informed about the contribution of permanent and transitory shocks to the observed terms of trade, requiring agents to solve a signal extraction problem.

In the model, households choose consumption, saving and labor supply to maximize lifetime utility. Households consume two goods - domestically produced goods and importables. The relative price of the two goods is the terms of trade. Households can invest in two assets - physical capital and a one-period non-contingent bond traded in international capital markets. The price of the bond is set exogenously, except for a small risk-premium included to ensure that the economy's net foreign debt is stationary. On the firm side, the model features production with endogenous capital and labor. I augment the model with

[3]This issue may be less of a concern for Kent and Cashin (2003) as they do not identify individual transitory or permanent shocks. Their approach can be viewed as implicitly assuming that agents have no information about the persistence of individual terms of trade shocks and merely expect the persistence of the average shock. It turns out that this assumption about the information structure is not a bad approximation to the results of this paper.

permanent and transitory productivity shocks as well as preference shocks and include capital adjustment costs. These features help the model to fit the data, but play little role in the analysis.

1.2.1 The Environment

1.2.1.1 Firms

The economy features a single perfectly competitive firm that produces a tradeable good using a Cobb-Douglas production technology of the form:

$$Y_t = A_t K_t^{\alpha} \left(X_t N_t\right)^{1-\alpha} \tag{1.1}$$

where Y_t denotes output in period t, K_t denotes capital \ N_t denotes hours worked and A_t and X_t represent productivity shocks. The first productivity shock, A_t, is a stationary shock that follows an autoregressive process in logs. In what follows, I use lower-case letters to represent log deviations from a variable's steady state, so that $a_t = \ln A_t$. The evolution of A_t then follows:

$$a_t = \rho_a a_{t-1} + \varepsilon_t^a;\ \varepsilon_t^a \sim N\left(0, \sigma_a^2\right) \tag{1.2}$$

The second productivity shock, X_t, is nonstationary. Let

$$M_t \equiv \frac{X_t}{X_{t-1}} \tag{1.3}$$

I assume that the logarithm of M_t follows a first-order autoregressive process of the form:

$$m_t = \left(1 - \rho_m\right) \mu + \rho_m m_{t-1} + \varepsilon_t^m;\ \varepsilon_t^m \sim N\left(0, \sigma_m^2\right) \tag{1.4}$$

The parameter μ measures the deterministic growth rate of the productivity factor X_t. The parameters $\rho_a, \rho_m \in [0,1)$ govern the persistence of a_t and m_t. With some abuse of notation, I refer to a_t and m_t as transitory and permanent productivity shocks, respectively.

Profit maximization by the firm ensures that factor prices reflect marginal products:

$$W_t = (1-\alpha) P_t^H \frac{Y_t}{N_t} \tag{1.5}$$

$$R_t^k = \alpha P_t^H \frac{Y_t}{K_t} \tag{1.6}$$

where W_t is the nominal wage, P_t^H is the price of the home-produced good and R_t^k is the rate of return to capital.

1.2.1.2 Households

Households maximize expected lifetime utility given by:

$$\sum_{t=1}^{\infty} \beta^t \left(V_t \ln C_t - A_L \frac{N_t^{1+\varphi}}{1+\varphi} \right) \tag{1.7}$$

where C_t is consumption and V_t is an exogenous preference shock that follows a first order autoregressive process in logs,

$$v_t = \rho_v v_{t-1} + \varepsilon_t^v; \ \varepsilon_t^v \sim N\left(0, \sigma_v^2\right) \tag{1.8}$$

The household's consumption bundle is a Cobb-Douglas aggregate of home- and foreign-produced goods,

$$C_t = \frac{\left(C_t^H\right)^{1-\eta} \left(C_t^F\right)^{\eta}}{(1-\eta)^{1-\eta} \eta^{\eta}} \tag{1.9}$$

where C_t^H are home-produced goods and C_t^F are foreign-produced goods. The parameter $\eta \in (0,1)$ governs the relative weights of home- and foreign-produced goods in the household's consumption bundle. Let P_t be the consumer price index corresponding to C_t. Then,

$$P_t = \left(P_t^H\right)^{1-\eta}\left(P_t^F\right)^{\eta} \tag{1.10}$$

where P_t^H is the price of the home-produced good and P_t^F is the price of the foreign produced good. Household optimality ensures that the demand for home- and foreign-produced goods is given by:

$$C_t^H = (1-\eta)\left(\frac{P_t}{P_t^H}\right)C_t \tag{1.11}$$

$$C_t^F = \eta\left(\frac{P_t}{P_t^F}\right)C_t \tag{1.12}$$

Households have access to two assets: domestic capital and a single-period, risk-free bond, denominated in the foreign good. The household's period-by-period budget constraint is:

$$Q_tB_{t+1} + P_tC_t + I_t + \frac{\phi}{2}\left(\frac{K_{t+1}}{K_t} - \mu\right)^2 K_t \leq W_tN_t + R_t^kK_t + B_t \tag{1.13}$$

where Q_t denotes the price of one-period risk-free bonds, B_{t+1} denotes the stock of bonds acquired in period t, I_t denotes gross investment, and ϕ is a parameter that controls the cost of adjusting the size of the capital stock. The capital stock evolves according to the law of motion:

$$K_{t+1} = (1-\delta)K_t + I_t \tag{1.14}$$

where $\delta \in [0,1)$ denotes the depreciation rate of capital.

To ensure that the solution to the model is stationary, I assume that the country faces a debt-elastic interest-rate premium as in Schmitt-Grohe and Uribe (2003). Specifically,

$$\frac{1}{Q_t} = 1 + r^* - \psi \left(e^{B_{t+1}/X_t (P_t^H)^{\frac{1}{1-\alpha}} - \bar{B}} - 1 \right) \tag{1.15}$$

where r^* is the exogenous foreign rate of interest on a risk-free bond and $\bar{B}$ is the steady-state foreign asset level.

Household utility maximization implies the following first order conditions:

$$\begin{aligned}
\frac{V_t}{C_t} &= \lambda_t P_t && (1.16) \\
N_t^{\varphi} &= \lambda_t W_t && (1.17) \\
\beta \mathbb{E}_t \left\{ \lambda_{t+1} \left[1 - \delta + R_{t+1}^K \right] \right\} &= \lambda_t \left(1 + \phi \left(\frac{K_{t+1}}{K_t} - \mu \right) \right) && (1.18) \\
&\quad + \beta \mathbb{E}_t \left\{ \lambda_{t+1} \left[\frac{\phi}{2} \left(\frac{K_{t+2}}{K_{t+1}} - \mu \right)^2 - \phi \frac{K_{t+2}}{K_{t+1}} \left(\frac{K_{t+2}}{K_{t+1}} - \mu \right) \right] \right\} \\
Q_t \lambda_t &= \beta \mathbb{E}_t \{ \lambda_{t+1} \} && (1.19)
\end{aligned}$$

where λ_t is the Lagrange multiplier on the household's budget constraint.[4]

1.2.1.3 Relative Prices

I take the price of the foreign good, P_t^F, as the numeraire and normalize it to 1. Define the terms of trade, S_t, as the relative price of home-produced goods in terms of foreign-produced goods. It follows from the definition of the consumer price index that:

$$S_t = P_t^H; \; P_t = S_t^{1-\eta}$$

The model economy is assumed to be small in the sense that it is a price-taker on world markets. Consequently, changes in its terms of trade are exogenous to domestic variables. The terms of trade are assumed to follow the process,

[4]Note that in taking first order conditions with respect to the foreign debt level, I have assumed that agents take the interest rate on foreign assets as given - that is, they do not internalize the effect of their decisions on their borrowing costs. This is standard in the literature (for example, Schmitt-Grohe and Uribe 2003). For a discussion of the implications of internalization of the risk premium, see Lubik 2007.

$$S_t = Z_t \Gamma_t \tag{1.20}$$

The first component, Z_t, represents a transitory shock to the terms of trade, which is assumed to follow a first-order autoregressive process in logs. That is,

$$z_t = \rho_z z_{t-1} + \varepsilon_t^z;\ \varepsilon_t^z \sim N\left(0, \sigma_z^2\right) \tag{1.21}$$

The second component, Γ_t is a permanent terms of trade shock. Let,

$$G_t \equiv \frac{\Gamma_t}{\Gamma_{t-1}} \tag{1.22}$$

I assume that the logarithm of G_t follows a first-order autoregressive process of the form:

$$g_t = \rho_g g_{t-1} + \varepsilon_t^g;\ \varepsilon_t^g \sim N\left(0, \sigma_g^2\right) \tag{1.23}$$

The decomposition of the terms of trade outlined in equations (1.20) - (1.23) is extremely flexible and encompasses many of the assumptions about the evolution of the terms of trade used in other papers. For example, if $\sigma_g^2 = 0$, the terms of trade is subject to purely transitory shocks, while if $\sigma_z^2 = 0$ and $\rho_g = 0$ then the terms of trade follows a random walk.[5]

1.2.1.4 Market clearing

Market clearing requires that the quantity of goods produced in the Home economy equals the consumption of these goods at home and abroad. This is ensured by the current account condition:

$$Q_t B_{t+1} + P_t C_t + I_t + \frac{\phi}{2}\left(\frac{K_{t+1}}{K_t} - \mu\right)^2 K_t \leq P_t^H Y_t + B_t \tag{1.24}$$

[5]The terms of trade will also follow a random walk if $\rho_g = \rho_z = \rho$ and $\rho \sigma_g^2 = (1-\rho)^2 \sigma_z^2$.

1.2.1.5 Equilibrium

An equilibrium is a sequence of quantities $\{C_t, N_t, I_t, Y_t, K_{t+1}, B_{t+1}\}_{t=0}^{\infty}$, prices $\{W_t, R_t^K, Q_t, P_t, P_t^H, S_t\}_{t=0}^{\infty}$ and exogenous processes $\{A_t, X_t, V_t, Z_t, \Gamma_t\}_{t=0}^{\infty}$ such that (i) firms maximize profits, given by Equations (1.5) and (1.6), (ii) households maximize utility, given by Equations (1.16)-(1.19), and (iii) markets clear, given by Equation (1.24), subject to the technological and resource constraints in Equations (1.1), (1.14), (1.15) and (1.20) and the exogenous processes (1.2), (1.4), (1.8), (1.21) and (1.23).

1.2.2 Information Structure

I assume that agents have complete information about all aspects of the economy other than the components of the terms of trade, about which they are imperfectly informed. In particular, I assume that agents can observe the levels of the terms of trade, but cannot observe Z_t or Γ_t directly. Reflecting the fact that agents are likely to have some information about the persistence of these shocks, I assume that agents receive a noisy signal regarding the permanent component of terms of trade growth. I refer to this signal as h_t, where $h_t = g_t + \varepsilon_t^h$ and ε_t^h is both independently and identically distributed with mean zero and variance σ_h^2. The agents' information set as of time t includes the entire history of terms of trade shocks and signals; $I_t \equiv \{S_t, h_t, S_{t-1}, h_{t-1}, \dots\}$.

In the model, agents form expectations about the decomposition of the terms of trade using the Kalman filter. To implement this, I represent the agent's filtering problem in state space form using the decomposition in Boz et al. (2011). First, I define the growth rate of the terms of trade as:

$$\begin{aligned} \Delta s_t &\equiv \ln s_t - \ln s_{t-1} \\ &= z_t - z_{t-1} + g_t \end{aligned}$$

The measurement equation includes a reformulation of this definition and the definition of h_t, as well as the observed exogenous processes:

$$\underbrace{\begin{bmatrix} a_t \\ m_t \\ v_t \\ \Delta s_t \\ h_t \end{bmatrix}}_{\mathcal{S}_t} = \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}}_{C} \underbrace{\begin{bmatrix} a_t \\ m_t \\ v_t \\ z_t \\ z_{t-1} \\ g_t \end{bmatrix}}_{\mathcal{X}_t} + \underbrace{\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}}_{D} \underbrace{\begin{bmatrix} \varepsilon_t^a \\ \varepsilon_t^m \\ \varepsilon_t^v \\ \varepsilon_t^z \\ \varepsilon_t^g \\ \varepsilon_t^h \end{bmatrix}}_{U_t} \tag{1.25}$$

The transition equation summarizes the evolution of the unobserved variables:

$$\underbrace{\begin{bmatrix} a_t \\ m_t \\ v_t \\ z_t \\ z_{t-1} \\ g_t \end{bmatrix}}_{\mathcal{X}_t} = \underbrace{\begin{bmatrix} \rho_a & 0 & 0 & 0 & 0 & 0 \\ 0 & \rho_m & 0 & 0 & 0 & 0 \\ 0 & 0 & \rho_g & 0 & 0 & 0 \\ 0 & 0 & 0 & \rho_z & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \rho_g \end{bmatrix}}_{A} \underbrace{\begin{bmatrix} a_{t-1} \\ m_{t-1} \\ v_{t-1} \\ z_{t-1} \\ z_{t-2} \\ g_{t-1} \end{bmatrix}}_{\mathcal{X}_{t-1}} + \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}}_{B} \underbrace{\begin{bmatrix} \varepsilon_t^a \\ \varepsilon_t^m \\ \varepsilon_t^v \\ \varepsilon_t^z \\ \varepsilon_t^g \\ \varepsilon_t^h \end{bmatrix}}_{U_t} \tag{1.26}$$

where $U_t \sim N(0,Q)$ and $Q \equiv \begin{bmatrix} \sigma_a^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & \sigma_m^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_v^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_z^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma_g^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma_h^2 \end{bmatrix}$.

The Kalman filter can be used to express the consumers' expectations in recursive form as:

$$\mathscr{X}_{t|t} = (I - KC)A\mathscr{X}_{t-1|t-1} + K\mathscr{S}_t \tag{1.27}$$

where I is an identity matrix of size the matrix of size 6×6 and K is the Kalman gain, calculated as:

$$K = PC' \left(CPC' + DQD'\right)^{-1} \tag{1.28}$$

and P is the steady-state error covariance matrix, calculated as the solution to:

$$P = APA' - APC' \left(CPC'\right)^{-1} CPA' + BQB' \tag{1.29}$$

Equations (1.27)-(1.29) fully characterize learning.

1.2.3 Steady State

Before deriving the non-stochastic steady state of the model, I normalize the following variables to be in a form that is stationary: $\tilde{Y}_t = Y_t / \left(X_{t-1} S_{t-1}^{\frac{\alpha}{1-\alpha}}\right)$, $\tilde{C}_t = C_t / \left(X_{t-1} S_{t-1}^{\frac{\alpha+\eta(1-\alpha)}{1-\alpha}}\right)$, $\tilde{K}_t = K_t / \left(X_{t-1} S_{t-1}^{\frac{1}{1-\alpha}}\right)$, $\tilde{B}_t = B_t / \left(X_{t-1} S_{t-1}^{\frac{1}{1-\alpha}}\right)$ and $\tilde{I}_t = I_t / \left(X_{t-1} S_{t-1}^{\frac{1}{1-\alpha}}\right)$, where a $\tilde{}$ denotes a stationary variable. Using these normalizations, the non-stochastic steady state of the model is given by:

$$\tilde{Y} = \tilde{K}^{\alpha} (\mu N)^{1-\alpha} \tag{1.30}$$

$$(\mu + \delta - 1)\tilde{K} = \tilde{I} \tag{1.31}$$

$$\frac{1}{Q} = 1 + r^* \tag{1.32}$$

$$\tilde{C} A_L N^{1+\varphi} = (1-\alpha)\tilde{Y} \tag{1.33}$$

$$\alpha \frac{\tilde{Y}}{\tilde{K}} = \mu + \delta - 1 \tag{1.34}$$

$$Q = \beta \tag{1.35}$$

$$(Q\mu - 1)\tilde{B} + \tilde{C} + \tilde{I} = \tilde{Y} \tag{1.36}$$

where I have replaced the wage rate and the rate of return on capital in the solution to the consumers' problem using the firm's profit maximization conditions, given by Equations (1.5) and (1.6).

1.2.4 Log-linearized equilibrium conditions

To solve the model, I log-linearize the model around the steady-state derived in the previous section. The log-linearized equilibrium conditions are:

Production:

$$\tilde{y}_t = a_t + \alpha \tilde{k}_t + (1-\alpha)(m_t + n_t) \tag{1.37}$$

Intratemporal optimization:

$$(1+\varphi) n_t + \tilde{c}_t - v_t = \tilde{y}_t + \eta \Delta s_t \tag{1.38}$$

Bond market Euler equation:

$$\begin{aligned} \tilde{c}_t - v_t - q_t &= E_t \{\tilde{c}_{t+1} - v_{t+1}\} + \omega \Delta s_t \\ &\quad + (1-\eta) E_t \{\Delta s_{t+1}\} + m_t \end{aligned} \tag{1.39}$$

Capital Euler equation:

$$\begin{aligned} \tilde{c}_t - v_t + \phi \mu \tilde{k}_t &= E_t \{\tilde{c}_{t+1} - v_{t+1}\} + (1+\phi\mu) m_t - \left(1 - \eta - \frac{1+\phi\mu}{1-\alpha}\right) \Delta s_t \\ &\quad - \frac{\beta \alpha \hat{Y}}{\mu \hat{K}} E_t \{\tilde{y}_{t+1}\} + \left(1 - \eta - \frac{\beta \phi \mu}{1-\alpha} - \frac{\beta \alpha \bar{Y}}{\mu \bar{K}}\right) E_t \{\Delta s_{t+1}\} \\ &\quad + \left(\phi \mu (1+\beta) + \frac{\beta \alpha \bar{Y}}{\mu \bar{K}}\right) \tilde{k}_{t+1} - \beta \phi \mu E_t \{\tilde{k}_{t+2}\} \\ &\quad - \beta \phi \mu E_t \{m_{t+1}\} \end{aligned} \tag{1.40}$$

Capital accumulation:

$$\mu\tilde{k}_{t+1} + \mu m_t + \frac{\mu \Delta s_t}{1-\alpha} = (1-\delta)\tilde{k}_t + \frac{\tilde{I}}{\tilde{K}}\tilde{i}_t \tag{1.41}$$

Risk-free rate of return:

$$q_t = \psi \tilde{b}_{t+1} \tag{1.42}$$

Current account:

$$\tilde{C}\left(\tilde{c}_t + (1-\eta)\Delta s_t\right) + \tilde{I}\tilde{i}_t + Q\tilde{B}\mu\left(q_t + \tilde{b}_{t+1} + m_t + \frac{\Delta s_t}{1-\alpha}\right) = \tilde{Y}\left(\tilde{y}_t + \Delta s_t\right) + \tilde{B}\tilde{b}_t \tag{1.43}$$

where a ~ denotes agents' beliefs about the value of an imperfectly observed variable and lower case letters denote a variable's log-deviation from its steady state, that is $d_t = (D_t - D)/D$.

1.2.5 Model Solution

The solution of the model follows Uhlig (1999) and Blanchard et al. (2012). Let $\mathscr{Y}_t$ denote the endogenous variables controlled by the agent. The economic model can be represented as the stochastic difference equation:

$$FE_t\{\mathscr{Y}_{t+1}\} + G\mathscr{Y}_t + H\mathscr{Y}_{t-1} + M\mathscr{S}_t + NE_t\{\mathscr{S}_{t+1}\} = 0 \tag{1.44}$$

where F, G, H, M and N are matrices of parameters and $\mathscr{S}_t$ is the vector of observable variables described in section 1.2.2. The unique stable solution of the model is:

$$\mathscr{Y}_t = P\mathscr{Y}_{t-1} + Q\mathscr{S}_t + R\mathscr{X}_{t|t} \tag{1.45}$$

where $\mathscr{X}_{t|t}$ represents the agents' expectation of the unobserved states described in section 1.2.2. The matrices P, Q, R can be found by solving the three matrix equations:

$$\begin{aligned} FP^2 + GP + H &= 0 \\ (FP+G)Q + M &= 0 \\ (FP+G)R + F(QC+R)A &= 0 \end{aligned}$$

and the matrices A and C are as defined in section 1.2.2. The matrix P can be recovered using the techniques discussed in Uhlig (1999).

1.3 Estimation

I estimate the model using Bayesian methods. This section outlines the estimation strategy, including the choice of priors, and explains how the variables of the theoretical model map into observable time series.

1.3.1 Measurement

The initial stage of the estimation is to map the model's variables, which are generally unobservable, into observable variables that can be used to estimate the model's parameters. To do this, I first express the log-linear equilibrium conditions, derived in the previous section, in state-space form as:

$$\tilde{\xi}_t = F\tilde{\xi}_{t-1} + v_t \tag{1.46}$$

$$\tilde{Y}_t = G + H\tilde{\xi}_t + \zeta_t \tag{1.47}$$

$$\begin{bmatrix} v_t \\ \zeta_t \end{bmatrix} \sim N\left(0, \begin{bmatrix} Q & 0 \\ 0 & R \end{bmatrix}\right) \tag{1.48}$$

where the theoretical variables are collected in the state vector ξ_t, and the observable variables are collected in the vector $\tilde{Y}_t$. Equation 1.46 governs the transition of the state variables, while Equation 1.47 maps the state into observable variables. The matrices F, G, H and Q are functions of the parameters of the model.

The observable variables I use to estimate the model are the growth rates of real GDP, consumption, gross fixed capital formation and the terms of trade as well as the level of the trade-balance-to-GDP ratio. That is:

$$\tilde{Y}_t = \begin{bmatrix} \Delta \ln Y_t & \Delta \ln C_t & \dots \\ \Delta \ln I_t & \Delta \ln S_t & \frac{NX_t}{Y_t} \end{bmatrix} \times 100 \tag{1.49}$$

All variables are expressed in per capita terms and are seasonally adjusted. I estimate the model using Australian data over the period 1973Q1 - 2012Q2. The starting point reflects the first quarter for which per-capita national accounts data are available for Australia. This is somewhat earlier than the starting date for most Australian DSGE models, which typically use data spanning the period after the adoption of a floating exchange rate in 1983 or inflation targeting in 1993. A later starting date is appropriate for models containing nominal interest rates or inflation, whose behavior is likely to be affected by changes in the conduct of monetary policy. In contrast, the model in this paper contains no nominal variables. And, given the presence of long-lived trends in the terms of trade, it seems preferrable to use a longer time series to estimate the model.

Following Jaaskela and Nimark (2011), the covariance matrix, R, of the vector of measurement errors, ζ_t, in Equation 1.47 are set to $E_t\left[\tilde{Y}_t\tilde{Y}_t'\right] \times 0.1$ so that ten per cent of the variance of the data series is assumed to be owing to measurement errors.

1.3.2 Bayesian Estimation

I estimate the parameters of the model using Bayesian methods that combine prior information with information from the data. The estimation works in the following way. Denote the vector of parameters to be estimated as Θ and the log of the prior probability of observing a given vector of parameters $\mathscr{L}(\Theta)$. The function $\mathscr{L}(\Theta)$ summarizes what is known about the parameters prior to estimation. The log likelihood of observing the dataset $\tilde{Y}_t$ for a given parameter vector Θ is denoted $\mathscr{L}\left(\tilde{Y}_t|\Theta\right)$, and is computed as:

$$\mathscr{L}\left(\tilde{Y}_t|\Theta\right) = -0.5\sum_{t=0}^{T}\left[p\ln(2\pi)+\ln|\Omega|+u_t'\Omega^{-1}u_t\right] \tag{1.50}$$

where p is the dimension of $\tilde{Y}_t$ and

$$\Omega = H'PH+R \tag{1.51}$$

is the covariance of the one-step ahead forecast errors u_t. The one-step ahead forecast errors, and the estimate of the model states, $\tilde{\xi}_t$, are computed recursively using the Kalman filter as follows:

$$u_t = \tilde{Y}_t - G - H\tilde{\xi}_t \tag{1.52}$$

$$\tilde{\xi}_{t+1} = F\tilde{\xi}_t + Ku_t \tag{1.53}$$

where K is the Kalman gain matrix, defined as:

$$K = FPH'\left(H'PH+R\right)^{-1} \tag{1.54}$$

and P is the steady state error covariance matrix, calculated as:

$$P = FPF' - KH'PF' + Q \tag{1.55}$$

The numerical procedure that I use to estimate the posterior distribution follows the methodology outlined in An and Schorfheide (2007). The algorithm is as follows:

Algorithm 1. *Random Walk Metropolis Algorithm*

1. Use a numerical maximization routing to calculate the posterior mode of the parameter vector, that is: $\hat{\Theta} = \arg\max \left[\mathscr{L}(\Theta) + \mathscr{L}\left(\tilde{Y}|\Theta\right)\right]$.[6] *Let . Let* $\tilde{\Sigma}$ *denote the inverse of the Hessian computed at the posterior mode.*

2. Draw an initial vector of parameters Θ^0 *from* $\mathscr{N}\left(\hat{\Theta}, \tilde{\Sigma}\right)$.

3. For $s = 1, \ldots, n_{sim}$ *draw* ϑ *from the proposal distribution* $\mathscr{N}\left(\Theta^{s-1}, c^2\tilde{\Sigma}\right)$*, where* c *is a scaling parameter included to induce the desired probability of accepting an individual draw. Accept this draw with probability* $\min\left\{1, r\left(\Theta^{s01}, \vartheta|Y\right)\right\}$*, where:*

$$r\left(\Theta^{s01}, \vartheta|Y\right) = \frac{\mathscr{L}(\vartheta|Y)\,p(\vartheta)}{\mathscr{L}(\Theta^{s-1}|Y)\,p(\Theta^{s-1})} \tag{1.56}$$

If I accept the draw, I set $\Theta^s = \vartheta$*. If I reject the draw, I set* $\Theta^s = \Theta^{s-1}$.

4. Approximate the posterior expected value of a function $h(\Theta)$ *by* $\frac{1}{n_{sim}}\sum_{s=1}^{n_{sim}} h(\Theta^s)$.

In computing the posterior distribution, I set the number of Metropolis-Hastings draws, n_{sim}, equal to 250,000, and select these after discarding an intial 250,000 burn-in draws. I calibrate the parameter c to achieve an acceptance rate of approximately 25 per cent.

1.3.2.1 Priors

For the AR(1) parameters of the exogenous processes, I assign beta priors with a mean of 0.5 and standard deviation of 0.15. The distribution of these priors ensures that the estimated parameters lie between 0 and 1, consistent with economic theory. I assign inverse gamma priors with a mean of 1 and a standard deviation of 1 to the standard deviations of the exogenous processes. Finally, for the deterministic growth rate of the economy, μ, and the investment adjustment cost parameter, ϕ, I assign truncated normal priors, with means 1.004 and 7.5 and standard deviations 0.1 and 1.50 respectively.[7] The mean of the prior for μ is consistent with the mean growth rate of per capita real GDP over the sample period.

[6]For this procedure, I use Chris Sims' csminwel program suite, available at www.princeton.edu/~sims/.

[7]The truncation ensures that μ is greater than 1 and ϕ is greater than 0. In the estimation, the bulk of the posterior distribution of both parameters lie far away from these truncation points.

The theoretical model, of course, contains a number of additional parameters. Many of these are likely to be poorly identified using only the observed data series included in the model but have been estimated many times previously. Rather than rely on imprecise estimates of these parameters, I calibrate them using values determined by previous research or economic theory. In a Bayesian framework, calibration can be thought of as a very tight prior. Table 1.1 outlines the calibrated parameters.

Table 1.1: Calibrated Parameters

Parameter		Description	Calibrated from
β	0.99	Discount factor	Standard value in literature.
δ	0.02	Depreciation rate	Match I/Y.
A_L	7.5	Utility cost of labor	Agents spend 1/3 of time in work.
$\tilde{B}$	1.10	Steady state foreign assets	Match TB/Y.
α	0.29	Capital share of income	Average compensation to capital as share of GDP.
η	0.20	Share of foreign goods in consumption	Average share of imports in consumption basket.
φ	1.00	Inverse Frisch elasticity	As in Jaaskela and Nimark (2011).
ψ	0.001	Debt-elastic interest rate premium	Estimated value in Jaaskela and Nimark (2011)

1.3.3 Posterior Distribution

Table 1.2 shows the main results of estimation. The transitory terms of trade shock is reasonably persistent, with a posterior median of the AR(1) coefficient ρ_z equal to 0.83. The persistence of the permanent shock is similar, with ρ_g equal to 0.85. Nonetheless, a shock to ε_g ultimately has a much larger and more lasting impact on the terms of trade. A positive shock to ε_z causes a once-off increase in the terms of trade, which then diminishes. The high value of ρ_z implies that it takes some time for the terms of trade to return to its initial level following the shock. In particular, the half-life of this shock is around five quarters, and the terms of trade does not return to its trend level for around five years. In contrast, a positive shock to ε_g increases the terms of trade on impact and then continues to increase the terms of trade further, albeit at a diminishing rate, over time. The accumulation continues over several quarters, and the terms of trade ultimately settles at slightly over twice the level of the initial impulse.

In terms of the magnitude of the shocks, the standard deviation of transitory terms of trade shocks, ε_z, is quite large at 2 per cent, while the standard deviation of the permanent terms of trade shock, ε_g, is much smaller at just 0.3 per cent. The standard deviation of the noise shocks is also large, at 7.8 per cent. This suggests that agents receive a fairly weak signal about the persistence of terms of trade shocks.

Table 1.2: Prior and Posterior Distributions - Structural Parameters - Partial Information Model

		Prior		Posterior			
Parameter	Distribution	Mean	SD	Mode	Median	5%	95%
Exogenous Processes - AR(1) coefficients							
ρ_a	Beta	0.500	0.150	0.853	0.834	0.648	0.961
ρ_m	Beta	0.500	0.150	0.161	0.177	0.050	0.358
ρ_v	Beta	0.500	0.150	0.495	0.548	0.246	0.827
ρ_z	Beta	0.500	0.150	0.833	0.823	0.768	0.874
ρ_g	Beta	0.500	0.150	0.848	0.831	0.742	0.899
Exogenous Processes - standard deviations							
σ_a	Inv. Gamma	1.000	1.000	0.517	0.521	0.400	0.656
σ_m	Inv. Gamma	1.000	1.000	0.601	0.599	0.375	0.808
σ_v	Inv. Gamma	1.000	1.000	0.460	0.488	0.314	0.647
σ_z	Inv. Gamma	1.000	1.000	2.028	2.030	1.804	2.280
σ_g	Inv. Gamma	1.000	1.000	0.241	0.269	0.169	0.402
σ_h	Inv. Gamma	1.000	1.000	5.675	7.819	2.315	48.796
Other parameters							
μ	Trunc. normal	1.005	0.001	1.004	1.004	1.003	1.005
ϕ	Trunc. normal	7.50	1.500	12.580	12.883	10.158	15.798
Log marginal density						-1475.9	

Although the remaining parameter estimates are not the focus of this paper, it is comforting to note that the results seem plausible and are broadly consistent with other empirical estimates. The persistence of the transitory productivity shocks is broadly in line with those typically found in the literature, although slightly larger than estimates from other small-scale Australian DSGE models, including Jaaskela and Nimark (2011) and and Kulish and Rees (2011). The magnitude of the permanent TFP shocks are similar to that estimated for Canada in Aguiar and Gopinath (2007). The estimate of the deterministic trend per-capita TFP growth rate of 0.4 per cent per quarter is only marginally lower than the average GDP per capita growth rate over the sample. Finally, the results imply large capital adjustment costs. This is a common finding in the open economy literature. In the absence of these adjustment costs, the ability of agents in the model to borrow and lend at an exogenous risk-free interest rate would imply implausibly large investment volatility.

Appendix (1. B) contains graphs of the prior and posterior distributions of the model's parameters. The posterior distributions appear to be well behaved and do not seem to be unduly influenced by the shape of the priors.

1.4 Response of the Economy to Terms of Trade Shocks

In this section I first show how incomplete information affects the response of the economy to terms of trade shocks and then discuss its implications for aggregate macroeconomic volatility.

1.4.1 Dynamic responses to terms of trade shocks

1.4.1.1 Transitory terms of trade shocks

Figure 1.3 shows the response of the economy to a one standard deviation transitory positive terms of trade shock. I focus first on the solid red lines, which show the response of the economy in the baseline case in which agents have incomplete information. The shock increases the terms of trade by around 2 per cent on impact. After the initial impulse, the

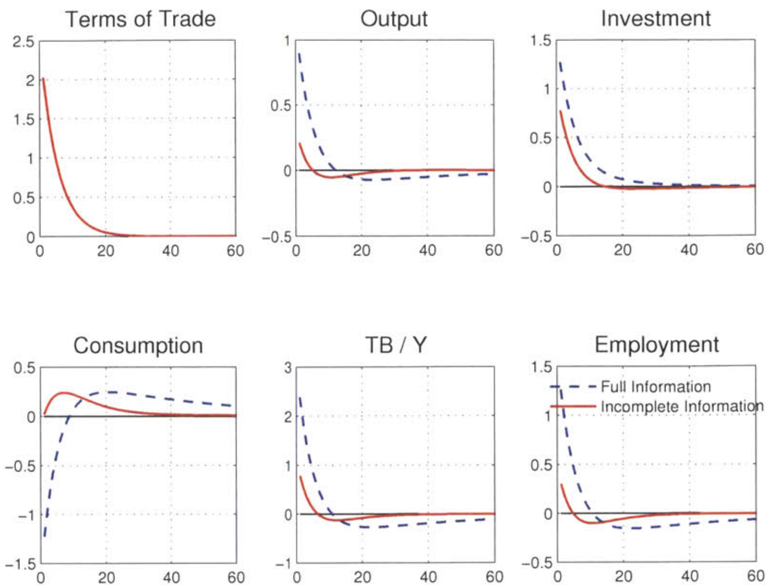

Figure 1.3: Impulse response to a one standard deviation positive transitory terms of trade shock

terms of trade decrease and after around two years stabilize at their original level. The shock increases the price of the economy's output relative to the price of consumption and investment goods. This induces households to work and invest more, leading to an increase in production. The boom in employment is short-lived, however, and within four quarters employment falls below its trend level. The investment boom is somewhat more persistent and it takes four years for investment to return to trend. Output increases by around 0.2 per cent on impact although it declines in subsequent periods, reflecting the contraction in employment. Consumption also increases in the first few quarters after the shock before slowly reverting to trend. The impact on the trade balance is quite small, with an initial improvement, followed a few quarters later by a small deterioration.

It is instructive to compare the response of the economy under partial information to its response under full information, shown by the dashed blue lines.[8] In the full-information case, the initial responses of investment, employment and output to the shock are substantially larger than in the partial information case. This is because, with full information, agents realize that the shock is entirely transitory and bring forward production to take

[8]To calculate the response under the full information, I use the parameter estimates from Table 1.2 but set the standard deviation of the terms of trade noise shocks, σ_h, equal to zero. In Section (1.5) I re-estimate the model parameters under the assumption of full information.

advantage of temporarily high relative export prices. Consumption initially decreases by around one per cent. It then increases to be marginally above trend after eight quarters and remains at an elevated level for a considerable period after that. The initial decrease in consumption occurs because agents expect the price of the consumption good in terms of importables to fall.[9] As the bond price in this model is largely unresponsive to the domestic economy and denominated in terms of importables, the expected decrease in the CPI causes an increase in the real interest rate. This induces households to postpone consumption.[10] A larger increase in output and decrease in consumption implies a larger initial improvement in the trade balance in the full information case compared to the partial information case. After four years or so, households start to draw down on the foreign assets that they accumulate through the improved trade balance, and use this to fund additional consumption.

The response of the economy under full information reflects a standard consumption smoothing response to a temporary income shock. Agents produce more when the relative price of output is high and save part of the windfall to fund higher consumption in the future. To understand the response of agents under partial information, it is necessary to examine their beliefs. These are illustrated in Figure 1.4. The left panel shows how agents' beliefs about the two components of the terms of trade shock, z_t and g_t, evolve following a transitory shock. Agents have some success in identifying the shock. They attribute over half of the 2 per cent increase in the terms of trade to the transitory shock and only a small proportion to the permanent shock.[11] Agents are less successful in inferring the evolution of z_t and g_t in future periods. But they still correctly attribute most of the evolution in the terms of trade to transitory shocks.

[9] In log-linear terms, $\tilde{p}_t = (1-\eta)\tilde{s}_t$, so that the CPI is proportional to the terms of trade in this model. The transitory terms of trade shock increases the CPI on impact, but agents expect deflation in subequent periods as the terms of trade decline.

[10] It should be noted that while this mechanism is general, the sign and magnitude of the consumption response are sensitive to the parameterization of the utility function. For example, in a similar model, Mendoza (1995) assumes an intertemporal elasticity of substition of 2.6, which causes consumption to increase following a transitory terms of trade shock.

[11] Letting the symbol ^ refer to agents' beliefs about the components of the terms of trade, the sum of $\hat{z}_t$ and $\hat{g}_t$ does not equal the change in the terms of trade because agents also adjust their beliefs regarding $\hat{z}_{t-1}$ following the shock. Specifically, on impact, $\hat{z}_{t-1}$ is equal to around -0.8, so that $\hat{z}_t - \hat{z}_{t-1} + \hat{g}_t = \Delta s_t$.

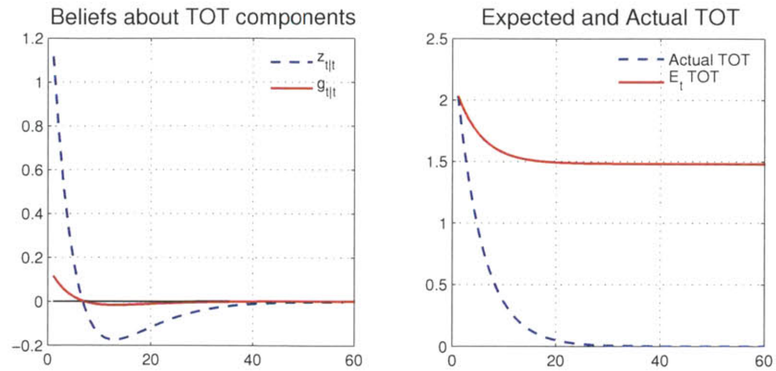

Figure 1.4: Agents' beliefs following a transitory terms of trade shock. The left panel shows the evolution of agents beliefts about the level of the transitory (dashed line) and permanent (solid line) components of the terms of trade. The right panel shows agents' expectations about the future evolution of the terms of trade (solid line) in the period in which the shock hits as well as the actual evolution of the terms of trade (dashed line).

Given that agents correctly identify transitory shocks as the main cause of the observed changes in the terms of trade, why do their reactions differ so much between the full information and partial information cases? The key to understanding this is to recall that the permanent shock increases the terms of trade in future periods as well as on impact. Hence, even a small initial increase in agents' beliefs about g_t can translate into a large increase in the expected long-run level of the terms of trade. To illustrate this, the right panel of Figure 1.4 shows the actual path of the terms of trade as well as agents' expectations about the evolution of the terms of trade calculated *in the period when the shock hits.* Although agents initially attribute only a small portion of the shock to the permanent component, this small permanent shock is ultimately expected to leave the terms of trade 1.5 per cent above its initial level. That is, agents expect that most of the increase in the terms of trade will be permanent. This explains why, in the partial information case, agents in the economy feel less urgency to work and invest more in the near term to take advantage of the high terms of trade than they do in the full-information case. It also explains why agents consume more and accumulate fewer foreign assets when they have incomplete information.

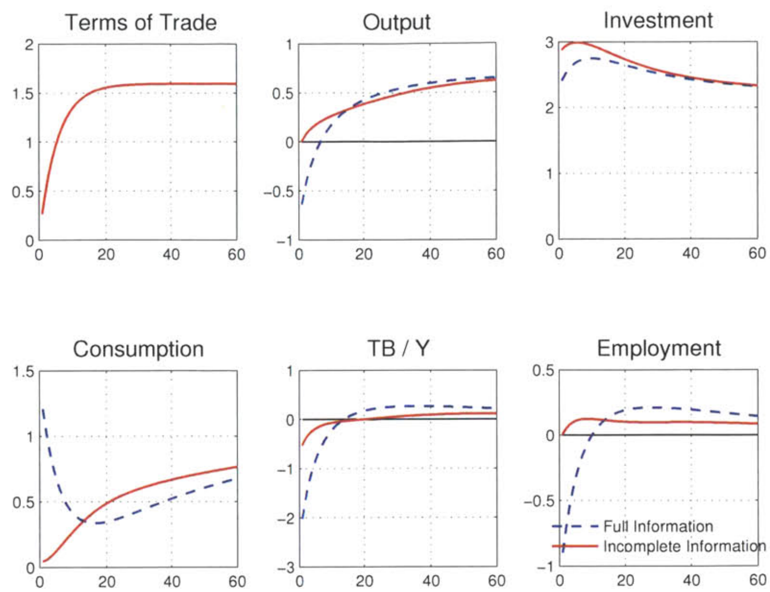

Figure 1.5: Impulse response to a one standard deviation permanent terms of trade shock

1.4.1.2 Permanent terms of trade shocks

Turning to the permanent shock, Figure 1.5 shows the economy's response to a one standard deviation shock to ε_g. The shock increases the terms of trade by 0.3 per cent on impact, and accumulates over time so that the terms of trade ultimately settles at around 1.6 per cent above its initial level after four years. Output and investment both increase on impact. The increase in output is initially small and accumulates over time. In contrast, the response of investment is quite large, as agents seek to build up the capital stock rapidly, and then diminishes. While output and investment increase permanently following the shock, employment eventually returns to trend. It takes a long time to do so, however, and 15 years after the shock employment remains above trend. Consumption initially responds little to the shock, but then increases over time. The investment boom overwhelms the increase in revenue from the higher terms of trade. Hence, the economy's trade balance deteriorates for around four quarters following the shock, although the economy ultimately runs trade later on.

Once again, it is informative to examine the economy's response to the shock under full information. In this case, output and employment both decrease following the shock and

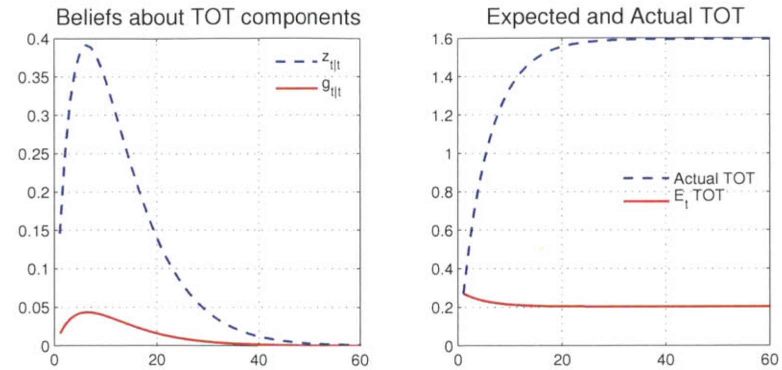

Figure 1.6: Agents' beliefs following a permanent terms of trade shock

only return to their initial level after around two years. Following that, however, the response of output in the full and partial information cases are broadly similar. Investment increases following the shock, but by less than in the partial information case. Agents also increase consumption by around 1.2 per cent when the shock hits. A stronger response of consumption and investment, and weaker response of output, translates into a larger initial deterioration in the trade balance. This is offset by a stronger improvement in the trade balance in future years.

Figure 1.6 shows agents' beliefs about the composition of the permanent shock. Agents make substantial errors in interpreting this shock. Initially, they attribute most of the shock to changes in the transitory component of the terms of trade, rather than the permanent component. And even as the terms of trade continues to increase in future periods, agents continue to attribute most of the increases to changes in the transitory components of the terms of trade. Indeed, agents' expectations about the future evolution of the terms of trade immediately following a permanent shock, shown in the right-hand panel of Figure 1.6, aren't substantially different from their beliefs following a transitory shock. Agents expect that most, but not all, of the initial increase in the terms of trade will be permanent. Or, put another way, the high estimated standard deviation of the noise shock means that agents struggle to distinguish between permanent and transitory terms of trade shocks in real time.

1.4.1.3 Noise shocks

As a final exercise, Figure 1.7 shows the response to a noise shock. Although it is difficult to place a structural interpretation to this shock, it can be thought of as a signal that the terms of trade will increase permanently that ultimately proves to be unfounded. While this signal is noisy, it contains sufficient information value that agents respond to the shock. In particular, agents consume more in anticipation of future increases in income. They also work, invest and produce less, preferring to perform these activities in the future, when they expect that a higher terms of trade will make production more lucrative. Agents fund their additional consumption by borrowing from abroad. This translates into a deterioration in the trade balance. When it becomes clear that the signal was misleading and that the terms of trade will not increase, agents are forced to draw back on consumption, and to work and produce more to repay their accumulated foreign borrowing.

Although one must be cautious about interpreting this shock, it is interesting that the behavior of the economy is similar to that of an economy running a "bad" current account deficit, of the type described by Blanchard and Milesi-Ferretti (2009), in which private saving decreases in anticipation of an income boom that does not occur. Examining the behavior of household and firm expectations in the lead up to balance of payments crises to see whether this mechanism is empirically relevant would be a useful avenue for further research.

1.4.2 Implications for volatility

The previous section demonstrated how incomplete information alters the macroeconomic effects of permanent and transitory terms of trade shocks. In light of these results, one might wonder whether incomplete information makes the economy more sensitive to terms of trade shocks. Specifically, if agents have incomplete information about the persistence of terms of trade shocks, is the variance of macroeconomic variables greater and do terms of trade shocks contribute more to macroeconomic volatility?

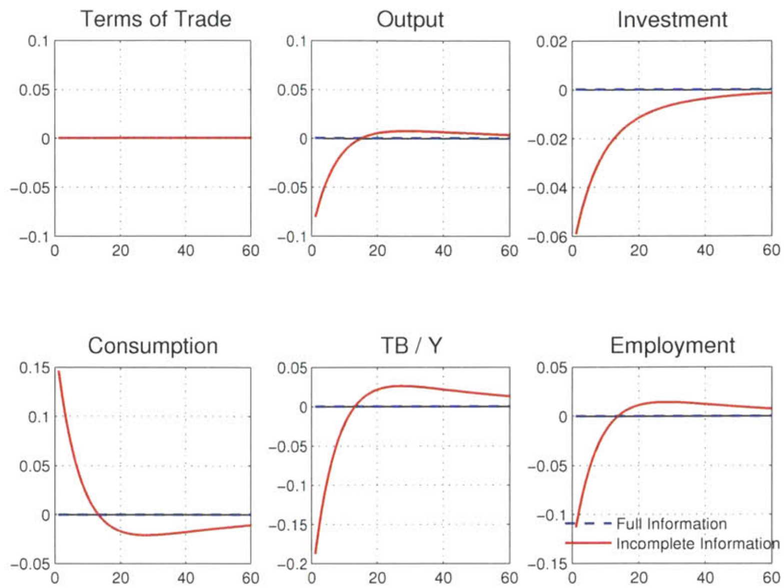

Figure 1.7: Impulse response to a terms of trade noise shock

One can imagine why this might be the case. With incomplete information, agents' forecasts of future terms of trade movements are likely to be less accurate than if they had full information. Agents make consumption and investment decisions based, in part, upon their expectation of future terms of trade movements. If the evolution of the terms of trade differs from expectations, agents may need to revise their consumption plans and expand or contract their investment projects. So, in an environment in which agents make larger forecasting errors about the evolution of the terms of trade, we might expect them also to make larger revisions to their consumption and investment plans, leading to greater macroeconomic volatility.

It turns out, however, that this is not the case. In fact, the model suggests that incomplete information actually reduces macroeconomic volatility and diminishes the importance of terms of trade shocks as a source of macroeconomic fluctuations. We can see this in Table (1.3). This shows the standard deviation of the growth rates of output, consumption and investment as well as the trade balance both in the data and in the model under partial and full information. In the baseline case of partial information, the model suggests a degree of macroeconomic volatility broadly comparable to that which we see in the data. Under full information, the standard deviation of output growth increases by around a 50 per cent, while the volatility of consumption growth and the trade balance more than double.

The intuition for this result comes from the fact that, under full information, transitory terms of trade shocks cause large change in the timing of production, investment and consumption across time. If agents expect that an increase in the terms of trade will be temporary, they will work more and invest to maximize production while prices are high. In contrast, permanent shifts in the terms of trade induce smaller intertemporal changes in production and consumption.[12] When agents are unable to observe the persistence of shocks, they react more cautiously to temporary terms of trade shocks. As the measured variance of these transitory shocks is high, this caution reduces macroeconomic volatility.

Table 1.3: Variance Decomposition

	Standard Deviation			Per cent of variance explained by:				
Variable	Data	Partial Information	Full Information	Partial Information			Full Information	
				ε_z	ε_g	ε_h	ε_z	ε_g
$\Delta \ln Y_t$	0.94	0.92	1.47	6.9	1.4	0.8	42.0	22.8
$\Delta \ln C_t$	0.75	0.81	1.97	2.3	2.4	3.6	44.8	39.9
$\Delta \ln I_t$	2.91	3.20	2.98	6.6	81.1	0.0	19.9	65.9
NX_t/Y_t	1.70	2.12	5.79	27.7	44.7	2.9	48.2	47.1

The table also shows the proportion of the variance of each of the variables explained by the three terms of trade shocks. Under partial information, terms of trade shocks explain a relatively modest proportion of the variance of output, consumption and investment, but a large proportion of the variance of the trade balance. Terms of trade noise shocks account for a modest portion of the variance of consumption and the trade balance, but have little effect on output or investment. In the absence of noise shocks, the terms of trade becomes a much more important driver of the variances of output growth and consumption.[13] The contribution of terms of trade shocks to the variance of the trade balance and investment

[12]Permanent shifts may, however, cause changes in the sectoral composition of output, for example between the tradeables and non-tradeables sectors or amongs tradeable industries. Examining these changes is not the focus of this paper and is left for future research.

[13]When $\sigma_h = 0$ noise shocks do not contribute to the variance of the model variables.

also increases somewhat, from a high level. There is also a change in the relative contribution of transitory and permanent terms of trade shocks, with the relative contribution of transitory terms of trade shocks increasing.

These results help to reconcile the conflicting results of Broda (2004) and Mendoza (1995) regarding the importance of terms of trade disturbances as a source of macroeconomic fluctuations. The baseline results correspond to Broda's finding that the terms of trade shocks explain a relatively modest proportion of output volatility. In contrast, the results under the assumption of full information are much closer to those in Mendoza (1995), who also assumed that agents have full information about the persistence of terms of trade shocks.

1.5 Robustness Checks

In this section, I show that the results are robust to alternative plausible estimates of the magnitude of the terms of trade noise shock, demonstrate the credibility of the model's forecasts for the evolution of the terms of trade and compare the baseline results to those from a model estimated under the assumption that agents have full information.

1.5.1 Alternative Magnitude of Noise Shock

In light of the importance of noise shocks for the conclusions of this paper, one might wonder how sensitive the results are to other plausible assumptions about the magnitude of σ_h. To illustrate this, Figure 1.8 shows the initial response of output, consumption and investment to temporary and permanent terms of trade shocks when σ_h takes alternative values from 0 to 50.[14] When σ_h equals zero, agents receive a perfect signal about the persistence of terms of trade shocks. Consequently, the responses of output, consumption and investment correspond to the initial responses in the full-information cases in Figures 1.3 and 1.5. As the amount of noise increases, agents find it harder to identify whether

[14]Recall that the posterior median value of this parameter was 8.9 and the 90 per cent probability interval spanned 2.3 - 48.8. All other parameters were set to their posterior median values in this exercise.

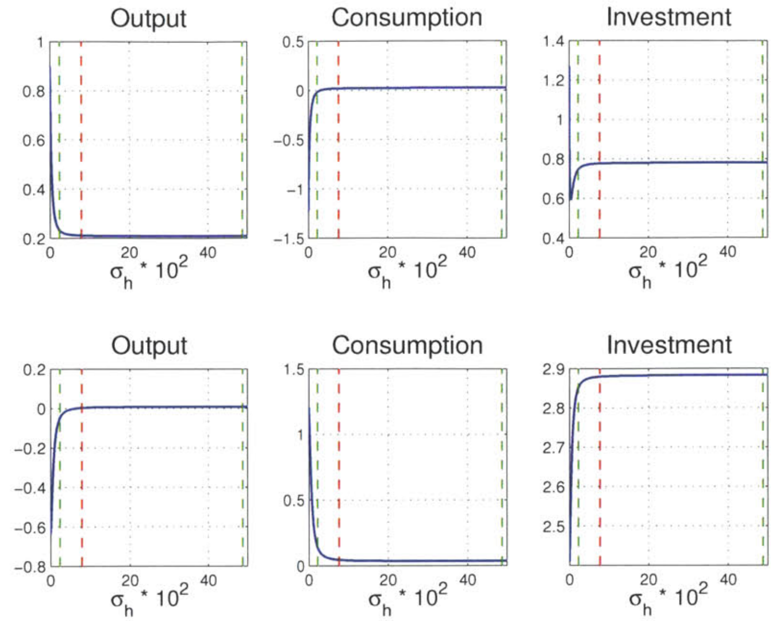

Figure 1.8: Impulse Response to a one standard deviation temporary (top row) and permanent (bottom row) terms of trade shock at t=1 for alternative values of σ_h. Dashed lines show posterior mean estimate of σ_h and 5 and 95 per cent confidence bands.

observed changes in the terms of trade are the result of permanent or transitory shocks and their initial responses to these shocks changes.However, once the standard deviation of the noise shocks increases beyond 3 or so, agents' responses to terms of trade shocks become largely unresponsive to changes in σ_h. Intuitively, at this point the information content of the signal that agents receive about the persistence of terms of trade shocks, h_t, is so weak that additional noise has little effect on agents' beliefs. The mass of the posterior distribution of σ_h lies in this region. That is, the model strongly suggests that agents receive little information about the persistence of terms of trade shocks. However, given the relatively small changes in the responsiveness of model variables in this region of the parameter space for σ_h, there is likely to be considerable uncertainty about exactly how much noise agents face.

1.5.2 Model Terms of Trade Forecasts

The forecasting process for the terms of trade in the model is extremely simple. One might be concerned that the large noise shocks that the model implies are the result of excluding other sources of information that agents might use to forecast the terms of trade. To examine this, for each draw from the posterior distribution, I use the Kalman Filter to back out estimates of agents' beliefs about the values of the permanent and transitory shocks to the terms of trade given the data included in the model. Given these beliefs, I then reconstruct the forecasts of the evolution of the terms of trade that agents in the model would have made, given their beliefs. Figure 1.9 shows the median forecasts at various times over the past decade. Comparing this to Figure 1.2, it appears that the model's forecasts are reasonably close to those produced by actual forecasters. This suggests that the model's forecasting process is a good approximation to that used by real-world agents and that the results of the paper are not driven by artificially constraining the information sets of agents in the model.

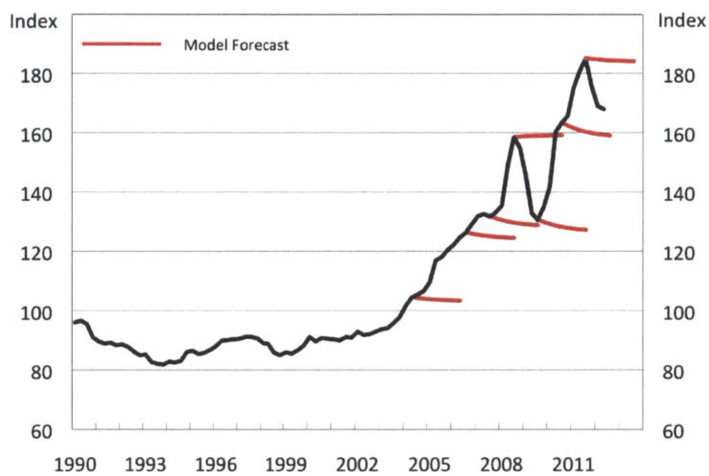

Figure 1.9: Australia terms of trade. The black line shows the actual evolution of the Australian terms of trade. The red lines show the forecasts produced by the model at various times.

1.5.3 Comparison to Full Information Model

As a final exercise, I compare the results of the incomplete information model presented in this paper to results from a model estimated assuming that agents have full information about the persistince of terms of trade shocks. For this exercise, I re-estimated the model in Section 1.2, using the same priors, but restricting the standard deviation of terms of trade noise shocks, σ_h, to equal zero. The estimation results are presented in Appendix 1. C.

The parameter estimates for the full information model differ from the incomplete information model in several respects. The persistence of shocks to productivity and consumer preferences in the full information model are larger than in the partial information model, while the persistence of the two terms of trade shocks are smaller. In addition, the size of transitory productivity and preference and persistent terms of trade shocks shocks are considerably larger in the full information model. The log marginal density for this model, computed using the Geweke (1999) modified harmonic mean estimator, is lower than in the partial information model, suggesting that the latter has better relative model fit.

As a further test of the relative merits of the partial- and full information models, I examine their ability to replicate the dynamic interactions between the terms of trade and domestic macroeconomic variables that we see in the data. To do this, I first estimate a vector autoregression using HP-filtered Australian data over the period 1973Q1 - 2012Q2 of the form:

$$AY_t = v + B(L)Y_t + u_t \tag{1.57}$$

where $Y_t' = (tot, y, c, i, nx/y)$ is a vector of stationary endogenous variables, v is a vector of constants, $u_t' = \left(u^{tot}, u^y, u^c, u^i, u^{nx/y}\right)$ is an error vector, A is a matrix, $B(L)$ is a matrix polynomial in the lag operator and $var(u_t) = \Omega$. I restrict the matrices A and $B(L)$ so that the domestic variables do not affect the terms of trade, either contemporaneously or with a lag. This is consistent with the assumption throughout this paper that the terms of trade is exogenous with respect to domestic economic developments. Having estimated the VAR, I calculate the impulse response of the domestic variables to a one standard deviation shock to the terms of trade. As discussed in Section 1.1, under the null-hypothesis of partial information about the persistence of terms of trade shocks, it is not possible to give a structural interpretation to this shock. However, the exercise still provides a useful summary of the dynamic reduced form relationships between empirical variables.

To compare the empirical responses to the model, I simulate 155 observations of synthetic data (equivalent to the sample size of the empirical data) for the terms of trade, output, consumption, investment and the trade-balance-GDP ratio using the partial- and

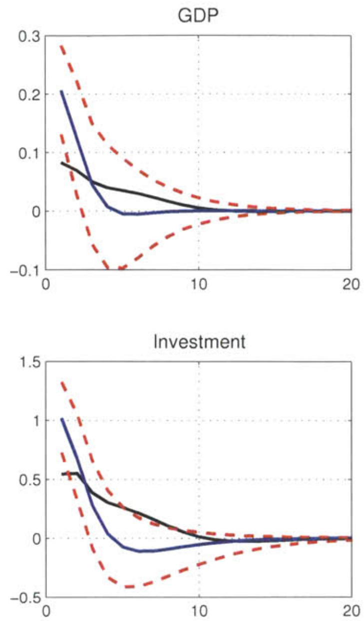

Figure 1.10: Impulse Response to a Reduced Form Terms of Trade Shock. The solid black line shows the empirical response. The blue line shows the median response from 10,000 simulations of model-generated data using parameter draws from the posterior of the model estimated assuming incomplete information about the persistence of terms of trade shocks. The dashed lines show the 95 per cent confidence bands of the impulse responses from these simulations.

full-information models, in each case taking a random draw from the posterior parameter distribution to simulate the model. For each model, I estimate a VAR as in Equation (1.57) using the synthetic data and calculate impulse responses to a terms of trade shock. I repeat this process 10,000 times for each model.

Figure 1.10 shows the empirical responses to a one standard deviation terms of trade shock, as well as the mean and 95 per cent confidence bands of the theoretical responses generated by the partial information model. The pattern and magnitude of the responses are broadly similar for all of the variables, although the initial responses of GDP, investment and the trade-balance-GDP ratio to the shock are somewhat smaller in the data than in the model, while the response of consumption is larger.[15] Figure 1.11 shows the responses for the full-information model. Here, there are substantial differences between the model and data responses. In particular, the initial responses of output, investment and trade-balance-GDP ratio are between three and five times larger in the model than they are in the data.

[15] I also repeated this exercise instrumenting the terms of trade with a commodity price index. The results were almost identical.

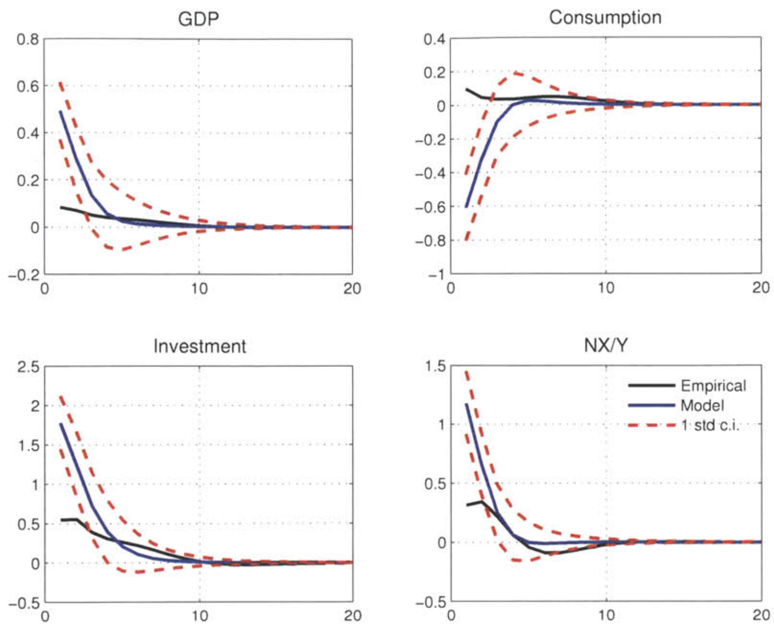

Figure 1.11: Impulse Response to a Reduced Form Terms of Trade Shock. The solid black line shows the empirical response. The blue line shows the median response from 10,000 simulations of model-generated data using parameter draws from the posterior of the model estimated assuming incomplete information about the persistence of terms of trade shocks. The dashed lines show the 95 per cent confidence bands of the impulse responses from these simulations.

Meanwhile, the model predicts that the terms of trade shock should trigger a sharp decrease in consumption, while the data suggests a small increase. In sum, these results suggest that the partial information model comes far closer to reproducing the dynamic relationships between the terms of trade and other macroeconomic variables that we see in the data than does the full-information model.

1.6 Conclusion

This paper has examined the extent to which agents are uncertain about the persistence of terms of trade shocks and described the effects of these shocks when agents have incomplete information. The results suggest that agents find it difficult to identify whether terms of trade disturbances are permanent or transitory in real time. In fact, agent's expectations about the evolution of the terms of trade are largely invariant to the type of shock that hits the economy. A corrollary of this result is that we should not expect households or firms to respond to terms of trade shocks in a first-best manner. Instead, in response to a temporary shock, agents will consume more and produce less than in a full information environment. And, in response to a permanent shock, agents will consume less and produce more than they would if they had full information. But despite the fact that agents make mistakes in identifying the source of terms of trade shocks, incomplete information about these shocks does not increase macroeconomic volatility.

A number of extensions to this work deserve consideration. First, it may be worthwhile to replicate the estimation for other small open-economies, including those featured in Figure 1.1. In particular, it would be interesting to learn whether the extent of incomplete information, and the effect of terms of trade noise shocks, differs between a developed country like Australia and a developing small open economy like Brazil or Chile. It may also be interesting to extend the model to include a non-traded sector and to incorporate nominal rigidities, including sticky wages and prices into the model. The former extension might reveal whether incomplete information about relative prices in the tradeable sector could have spillover effects to the rest of the economy. In particular, it may be interesting to see whether incomplete information about terms of trade movements could cause "Dutch Disease" type effects on the non-traded or non-commodity sectors. The inclusion of nominal rigidities could allow one to examine how uncertainty about the persistence of terms of trade shocks affects monetary policy. Many of the economies included in Figure 1.1 were early adopters of inflation targetting. Is there something about this monetary policy framework that makes it particularly desirable for countries facing incomplete information about terms of trade movements?

Finally, a word on policy more generally. Agents in the model optimize given the information available to them. Moreover, the model contains no market failures. Hence, the results of this paper provide no explicit reason for policy intervention in response to terms of trade shocks. Nonetheless, the differences between the estimated responses to terms of trade shocks under incomplete and complete information might cause one to wonder about the merits of policy interventions along the lines suggested by Caballero and Lorenzoni (2009). If policy makers have information about the persistence of terms of trade shocks not available to other agents in the economy then such interventions may well be warranted. If however, policy makers also operate under the constraints of incomplete information, as strongly suggested by Figure 1.2, then the results of this paper demonstrate how difficult such interventions are likely to be.

1. A Data

The dataset spans the quarters 1973:2 to 2012:2. The start date is chosen because quarterly estimates of the Australian population, published by the Australian Bureau of Statistics, is not available before that date.

Consumption

Household final consumption expenditure, expressed in chain volume terms and seasonally adjusted divided by population (Australian Bureau of Statistics Cat No. 5206.0).

Gross Domestic Product

Real gross domestic product, expressed in chain volume terms and seasonally adjusted divided by population (Australian Bureau of Statistics Cat No. 5206.0).

Investment

Private gross fixed capital formation, expressed in chain volume terms and are seasonally adjusted divided by population (Australian Bureau of Statistics Cat No. 5206.0).

Population

Total resident population (Australian Bureau of Statistics Cat No. 3101.0).

Terms of Trade

Australia: Terms of trade index, seasonally adjusted (Australian Bureau of Statistics Cat No. 5206.0).

Trade balance - to - GDP ratio

Ratio of nominal net exports to nominal gross domestic product (Australian Bureau of Statistics Cat No. 5206.0).

1. B Prior and Posterior Distributions

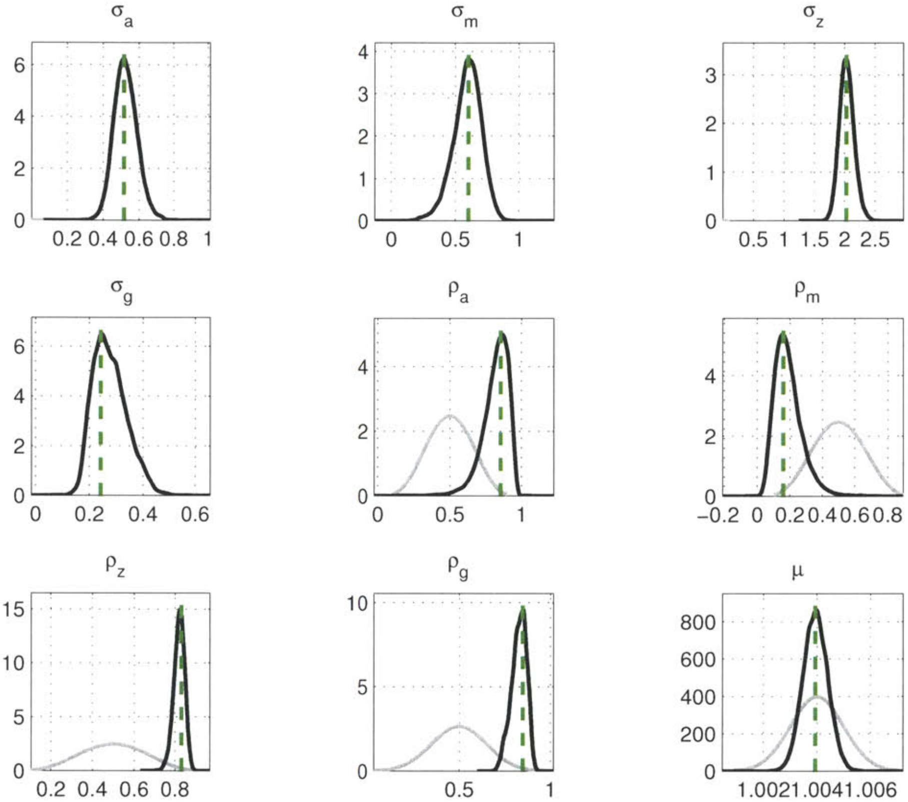

Figure 1.12: Prior and Posterior Distributions: solid grey line indicates the prior distribution; solid black line indicates the posterior distribution; dashed line indicates the posterior mode.

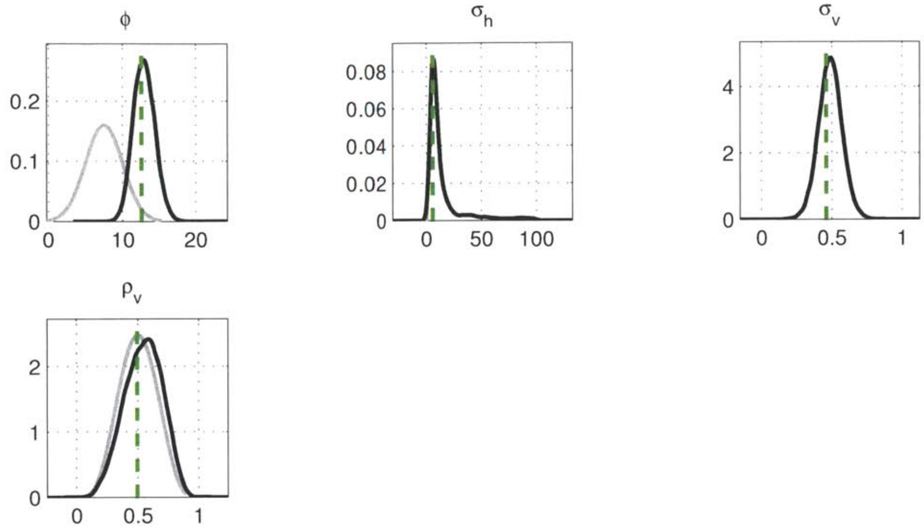

Figure 1.13: Prior and Posterior Distributions: solid grey line indicates the prior distribution; solid black line indicates the posterior distribution; dashed line indicates the posterior mode.

1. C Full Information Model Estimation

Table 1.4: Prior and Posterior Distributions - Structural Parameters - No TOT Noise Shocks

		Prior		Posterior			
Parameter	Distribution	Mean	SD	Mode	Median	5%	95%
Exogenous Processes - AR(1) coefficients							
ρ_a	Beta	0.500	0.150	0.977	0.976	0.959	0.988
ρ_m	Beta	0.500	0.150	0.375	0.353	0.141	0.587
ρ_v	Beta	0.500	0.150	0.715	0.712	0.595	0.823
ρ_z	Beta	0.500	0.150	0.659	0.656	0.563	0.739
ρ_g	Beta	0.500	0.150	0.065	0.091	0.000	0.227
Exogenous Processes - standard deviations							
σ_a	Inv. Gamma	1.000	1.000	0.848	0.844	0.714	0.981
σ_m	Inv. Gamma	1.000	1.000	0.249	0.289	0.124	0.545
σ_v	Inv. Gamma	1.000	1.000	1.555	1.580	1.360	1.831
σ_z	Inv. Gamma	1.000	1.000	1.329	1.347	1.142	1.567
σ_g	Inv. Gamma	1.000	1.000	0.958	0.935	0.706	1.150
Other parameters							
μ	Trunc. normal	1.004	0.001	1.004	1.004	1.003	1.005
ϕ	Trunc. normal	7.500	2.500	8.293	8.479	6.492	10.648
Log marginal density						-1658.0	

Chapter 2

Incomplete Information, Dutch Disease and Monetary Policy

> The main thing we don't know is how long the boom will last. This matters a great deal. If the rise in income is only temporary, then ... the economic restructuring that would reduce the size of other sectors that would be viable at 'normal' relative prices and a 'normal' exchange rate would be wasteful if significant costs are associated with that change only to find that further large costs are incurred to change back after the resources boom ends. If, on the other hand, the change is going to be quite long-lived ... a great deal of structural economic adjustment is bound to occur. In fact it almost certainly could not really be stopped. It would not be sensible to try to stop it. ... How then should we respond to our knowledge, and to the limits of our knowledge?
>
> - Glenn Stevens, Reserve Bank of Australia Governor, 23 February 2011

2.1 Introduction

Commodity-producing small open economies often experience large fluctuations in their terms of trade. In addition to their direct impact on national income, these shocks may also affect the allocation of resources between different sectors of the economy. A leading example of this occurs when a large commodity price boom, which is typically associated with an appreciation of the real exchange rate, triggers a contraction in the non-commodity tradeable sector, a phenomenon commonly known as the 'Dutch Disease'. Aside from distributional concerns, this type of sectoral reallocation is not generally considered to be a reason for policy intervention to support the contracting sector in a neoclassical competitive economy. But a policy response may be warranted if a shock is temporary and the contraction of the non-commodity sector generates negative externalities that firms, workers and consumers in that sector do not internalize.

But what if, as is suggested by the quotation above, policymakers are unsure about the persistence of a commodity price boom? Does this strengthen the case for policymakers to lean pre-emptively against contractions in the non-commodity sector? And, if so, how much weight should they put on this objective relative to other concerns, such as the costs of inflation? I examine these issues in the context of monetary policy using a New Keynesian model of a small open economy. The model economy exports non-commodity goods and a commodity endowment. Commodity prices are exogenously determined and subject to stochastic shocks that may be either brief or long-lived. Agents in the model, including the central bank, understand the commodity price process. But they cannot directly observe the persistence of commodity price shocks. Consequently, they must act based upon their best forecasts of future commodity price movements, and alter their plans if these forecasts turn out to be incorrect. Non-commodity firms are subject to a dynamic learning-by-doing externality through which the current level of aggregate output affects firm-level productivity in the future. Individual firms rationally ignore the effects of their production decisions on aggregate output. This creates an additional reason for policy intervention in the economy beyond the desire to maintain price stability.

A commodity price boom alters the value of a learning-by-doing externality in two ways. First, higher commodity prices provides agents with additional income that makes future revenue from the non-commodity sector less valuable. All else equal, this reduces the value of the learning-by-doing externality and implies that optimal policy will provide less support to the non-commodity sector during a commodity price boom. However, if accumulating learning takes time then a temporary contraction in the non-commodity sector that must later be reversed generates costs that individual firms may not internalize. This provides a rationale for policymakers to subsidize the non-commodity sector to prevent the loss of accumulated learning. The relative importance of these two considerations depends upon a number of factors, including the persistence of commodity price shocks, the process governing the evolution of learning-by-doing and the nature of information acquisition. In general, however, the value of learning and the optimal output of the non-commodity sector will be greater at the onset of a transitory commodity boom than it is at the beginning of a persistent boom.

When policymakers have incomplete information they must trade off the costs of letting the non-commodity sector contract too much during a temporary boom against the costs of maintaining a larger-than-optimal non-commodity sector during a prolonged boom. I find that the optimal response sees the non-commodity sector initially contract by more during temporary price booms under incomplete information than it does under full information, but by less during long-lived booms. This reflects the fact that, during a temporary price boom, imperfectly informed agents believe that there is some chance that the boom will be persistent. Consequently, the perceived value of learning is lower than it is when agents have full information, as is the optimal subsidy to the non-commodity sector. The opposite occurs during persistent booms.

As prices evolve and policymakers receive more information about the persistence of the boom, the optimal subsidy adjusts to reflect this new information. If a boom is transitory, it may be optimal to support the recovery of the non-commodity sector after prices have fallen. In contrast, if a boom turns out to be persistent, policymakers may wish to accomodate a prolonged contraction in the non-commodity sector to encourage resources to move to more productive uses. However, the evolution of the optimal subsidy will depend upon the commodity price process as well as the nature of the information frictions. Moreover, across a range of parameter values, I find that optimal policy largely stabilizes the prices of domestically-produced non-commodity goods, suggesting that the externality has only a small impact on policy settings relative to other considerations.

Because it may be difficult to implement optimal policy, I also examine the performance of simple policy rules. It turns out that a policy of targeting the rate of inflation of domestically-produced goods comes close to matching the welfare outcomes of optimal policy. Policies that feature a strong response to domestic inflation as well as a modest response to changes in the nominal exchange rate can produce marginal improvements in welfare over strict domestic inflation targeting and reduce the volatility of output and employment. Consumer price index inflation targeting produces somewhat worse welfare outcomes, although the quantiative cost of this policy is estimated to be small. An exchange rate peg generates worse welfare outcomes than other simple policy rules for most reasonable parameter values. However, its performance improves when the small economy's non-commodity goods are close substitutes for foreign goods and imports account for a large share of the domestic consumption basket.

2.1.1 Related Literature

Economists have long recognised that a commodity windfall can have an adverse effect on a small open economy's non-commodity tradeable sector. Much of the early work examining this issue, including Gregory (1976) and Corden (1982), featured models in which these windfalls were unambiguosly welfare improving for the economy as a whole, albeit with possibly adverse distributional consequences. Later research noted the potential for temporary commodity booms to reduce welfare if they trigger adverse long-term shifts in the pattern of comparative advantage, for example due to the presence of learning-by-doing externalities in the tradeable sector (van Wignbergen 1984, Krugman 1987 and Torvik 2001). When this occurs, some type of subsidy to protect the non-commodity sector is often optimal, although this conclusion is sensitive to assumptions about the persistence of commodity price shocks, the extent of financial market incompleteness and whether learning-by-doing externalities are limited only to tradeable sectors of the economy. Moreover, even if policy intervention is desirable, it may be optimal to provide it after commodity prices have fallen, rather than during a boom, as highlighted by Caballero and Lorenzoni (2009).

In most models, the contraction of the non-commodity tradeable sector during a commodity boom, which occurs despite an increase in domestic consumption, results from an appreciation of the real exchange rate that reduces the competitiveness of domestically-produced tradeables (Corden and Neary 1982). In a model with sticky prices, monetary policy can limit these appreciations, at least in the short-run, through its influence on the nominal exchange rate. It is therefore a natural candidate to act as a policy instrument to mitigate the adverse effects of transitory commodity price fluctuations. Indeed, this possibility has recently been examined in a closely related paper by Lama and Medina (2012). They compare the performance of simple monetary policy rules in a small open economy with a manufacturing sector that is subject to a learning-by-doing externality and a commodity sector that faces exogenous price shocks. They find that for most reasonable parameter values optimal policy responds primarily to inflation in home-produced manufactured goods and that the learning-by-doing externality is not a reason to stabilize the nominal exchange rate. Unlike this paper, Lama and Medina assume that agents and the central bank have full information about the state of the economy. The contribution of this paper is to examine this issue in a setting in which agents have incomplete information about the persistence of commodity shocks. Rees (2012) finds that such frictions are pervasive. I also provide a more rigorous analysis of the effect of learning-by-doing externalities, and demonstrate that the particular externality used in this paper and by Lama and Medina actually suggests that policymakers may want to hasten the contraction of the non-commodity sector during the early stages of a commodity price boom.

A number of other papers have examined the conduct of monetary policy in commodity exporting economies more generally. Recent examples in this literature include Lartey (2008), Dib (2008) and Hevia and Nicolini (2012). Although exact specifications differ across models, these papers tend to conclude that some form of home-produced goods inflation targeting is optimal. This largely reflects the fact that, in these models, the major departure from economic efficiency occurs due to staggered price setting. This means that relative prices are distorted when the inflation rate of home-produced goods is non-zero. A policy of home-produced goods inflation targeting mitigates this distortion. An exception is Catao and Chang (2012), who demonstrate that, when the elasticity of substitution between domestic and foreign tradeable goods is high, an exchange rate peg can produce superior welfare outcomes to inflation targeting rules.

This paper also contributes to work examining optimal monetary policy when agents have imperfect information about the state of the economy. The starting points for this literature were Pearlman (1992) and Svensson and Woodford (2003). In an environment in which policymakers and private sector agents receive the same noisy signal about the state of the economy, they show that optimal monetary displays certainty equivalence. That is, the optimal response to the best estimate of the state of the economy under incomplete information is the same as it would be if the state of the economy was perfectly observable.[1] But, while optimal policy displays certainty equivalence, this is not generally true for the simple policy rules, including Taylor rules, that are often thought to provide a more realistic representation of actual central bank decision making. A number of papers have characterised optimal simple rules under imperfect information. For example, Aoki (2003) and Orphanides (2003) consider simple policy rules when agents and policymakers receive a noisy signal about potential output and underlying inflation. They find that the optimal Taylor rule implies a degree of policy cautiousness that manifests itself in a stronger autoregressive component to policy than is optimally the case under full information. And, using calibrated business cycle models, Rudebusch (2001), Smets (2002) and Ehrmann and Smets (2003) argue that it is optimal for policymakers to respond less to variables that are measured with error. This paper applies some of the tools used in the existing literature to a previously unexplored question. And, to the best of my knowledge, this is the first paper to examine optimal monetary policy under incomplete information in a small open economy model.

[1]This result occurs because the selection of optimal policy and the estimation of the current state of the economy can be treated as separate problems. When policymakers and the private sector have different information sets, this separation principle does not hold and certainty equivalence may fail, as discussed in Svensson and Woodford (2004).

In what follows, I begin by setting out a flexible price version of the model to demonstrate how the value of a learning-by-doing externality evolves following imperfectly-observed commodity price shocks. I then outline the sticky-price version of the model, explain how to derive an approximation to welfare and construct optimal monetary policy. Having outlined the optimal monetary policy response to commodity price shocks, I then examine whether alternative simple rules can replicate the welfare outcomes of optimal policy. The following section discusses robustness checks and extensions to the basic model. A final section draws conclusions and presents suggestions for future research.

2.2 Income Shocks, Learning-by-Doing and Optimal Policy under Incomplete Information: A Simple Example

Before considering optimal monetary policy responses to commodity price shocks, I describe the optimal response to a commodity price boom under incomplete information in a simpler model with flexible prices. In the model, households choose consumption, labor supply and savings to maximize expected lifetime utility. In addition to wages and returns on financial assets, households also receive a commodity endowment, whose price is unaffected by developments in the domestic economy. The economy features a large number of perfectly competitive firms who produce identical goods using labor and and a second factor of production, which I refer to as learning. As in Cooper and Johri (1997), the evolution of learning depends on aggregate output. This creates an externality as each firm takes aggregate output as given when making its own production decision, and provides a rationale for policy intervention in the economy.

2.2.1 The Market Economy

I begin by describing equilibrium in the market economy. The economy features a representative household with the utility function:

$$U_t = E_0 \sum_{t=0}^{\infty} \beta^t \left[\ln(c_t) - \frac{l_t^{1+\varphi}}{1+\varphi} \right] \tag{2.1}$$

where c is consumption and l is labor. I treat the consumption good as the numeraire and normalize its price to one. The household receives income from three sources. The first is labor income. The second is a constant commodity endowment (normalized to one) that the household receives each period and whose relative price is exogenous and equal to s. The third is income from bond holdings, b, which accrues at an exogenous real interest rate r. All bonds are denominated in units of the consumption good, are risk free, and have a maturity of one period. The household's budget constraint is:

$$b_{t+1} + c_t \leq w_t l_t + s_t + b_t (1+r) \tag{2.2}$$

where w is the real wage. On the production side, a large number of identical firms produce consumption goods using labor and a factor of production called learning, lrn. Each firm's production function is Cobb-Douglas:

$$y_t(j) = lrn(j)_t^{\alpha} l_t(j)^{1-\alpha} \tag{2.3}$$

where $y(j)$, $lrn(j)$ and $l(j)$ are the production, stock of learning and labor employed by firm j. A firm's stock of learning evolves according to:

$$lrn(j)_{t+1} = lrn(j)_t^{\chi} y_t^{1-\chi} \tag{2.4}$$

where $y = \int_0^1 y(j)dj$ is aggregate output, which firms take as given when making their production decisions. I assume that all firms initially have the same level of learning. Given Equation 2.4, this ensures that they have the same stock of learning in subsequent periods as well.

The factor of production lrn represents a lagged production complementarity that is external to the firm. It may be interpreted in several ways. It could, for example, reflect the contribution of past economic activity to labor productivity in a model where workers accumulate human capital through on-the-job learning, as in Drazen (1985). This interpretation motivates the naming of the externality. But it could also represent external economies of scale, possibly through the integration and streamlining of supply chains (as documented by Bartelsman et al. 1994 and Holmes 1999), or knowledge spillovers through the transfer of information about efficient production techniques or innovations (as discussed in Jaffe et al. 1993).

2.2.2 Competitive Equilibrium

The conditions for household maximization are:

$$\begin{aligned} c_t l_t^{\varphi} &= w_t && (2.5) \\ 1 &= \beta (1+r) E_t \left\{ \frac{c_t}{c_{t+1}} \right\} && (2.6) \end{aligned}$$

I assume throughout that $\beta (1+r) = 1$, which implies that the Euler equation becomes $1 = E_t \{c_t / c_{t+1}\}$.

Firms equate the cost of labor to its marginal value product, which generates the optimality condition:

$$w_t = (1-\alpha) \left[\frac{lrn_t}{l_t} \right]^{\alpha} \quad (2.7)$$

where $lrn = \int_0^1 lrn(j) dj$. Here, I use the fact that all firms are identical to write the firm's first order condition in terms of aggregates.

A competitive equilibrium is a sequence of allocations $\{c_t, l_t, y_t, lrn_{t+1}, b_{t+1}\}_{t=0}^{\infty}$ and wages $\{w_t\}_{t=0}^{\infty}$, given initial conditions $\{lrn_0, b_0\}$ and commodity prices $\{s_t\}_{t=0}^{\infty}$ such that households and firms optimize (Equations 2.5-2.7) and the budget constraints and production functions (Equations 2.2-2.4) hold.

2.2.3 The Social Planner's Solution

I now consider the problem from a social planner's perspective. The planner's problem is to choose a sequence of allocations $\{c_t, l_t, lrn_{t+1}, b_{t+1}\}_{t=0}^{\infty}$ to maximize household utility subject to the technical conditions given by Equations 2.2-2.4.

The solution to the planner's problem is:

$$
\begin{aligned}
1 &= \beta(1+r)E_t\left\{\frac{c_t}{c_{t+1}}\right\} && (2.8)\\
c_t l_t^{\varphi} &= (1-\alpha)\left[\frac{lrn_t}{l_t}\right]^{\alpha} && (2.9)\\
&\quad +(1-\chi)(1-\alpha)\mu_t\left[\frac{lrn_t}{l_t}\right]^{\alpha+\chi(1-\alpha)}\\
\mu_t &= \beta E_t\left\{\frac{c_t}{c_{t+1}}\left[\alpha\left[\frac{lrn_{t+1}}{l_{t+1}}\right]^{\alpha-1}\right.\right. && (2.10)\\
&\quad \left.\left.+\mu_{t+1}(\chi+(1-\chi)\alpha)\left[\frac{lrn_{t+1}}{l_{t+1}}\right]^{-(1-\chi)(1-\alpha)}\right]\right\}
\end{aligned}
$$

where μ reflects the shadow value of learning. Note the difference between Equation 2.9 in the planner's problem and the equivalent equilibrium conditions in the market economy. This difference reflects the externality associated with learning-by-doing in this model. An individual firm's output today affects the productivity of all firms in the future. In the competitive equilibrium, individual firms disregard the external effects of their production decisions. In contrast, the social planner does take account of these effects. And, the wedge between the private and social value of labor is equal to the shadow value of learning, μ, multiplied by the marginal effect of labor on the learning stock.

2.2.4 Implementing the Planner's Solution

The market economy can replicate the planner's solution through a time-varying wage subsidy funded by lump sum taxation. With this, the first order conditions for the competitive firm become:

$$\frac{w_t}{\tau_t} = (1-\alpha)\left[\frac{lrn_t}{l_t}\right]^{\alpha} \tag{2.11}$$

The competitive equilibrium with the wage subsidy will replicate the planner's solution if:

$$\tau_t = 1 + (1-\chi)\,\mu_t \left[\frac{lrn_t}{l_t}\right]^{\chi(1-\alpha)} \tag{2.12}$$

where μ_t is as defined in Equation 2.10.

If the economy is in a steady state where all variables are constant, then the optimum wage subsidy is:

$$\tau = 1 + \frac{\alpha\beta\,(1-\chi)}{[1-\beta\,(\chi+(1-\chi)\,\alpha)]}$$

This is independent of the scale of production as well as of the household's bond holdings or endowment prices. If $\alpha = 0$ or $\chi = 1$, which correspond to situations in which learning either does not contribute to production or is unaffected by aggregate output, then $\tau = 1$. In these cases, the competitive equilibrium replicates the social planner's solution even if there is no wage subsidy.

To gain more intuition for the factors influencing the value of the optimal subsidy, let $\chi = 0$. In this case, there is no persistence in the evolution of learning as $lrn_{t+1} = y_t = lrn_t^{\alpha} l_t^{1-\alpha}$. The optimal subsidy is:

$$\tau_t = 1 + \mu_t \tag{2.13}$$

And, μ_t evolves according to the process:

$$\mu_t = \beta E_t \left[\frac{\alpha c_t}{c_{t+1}} \left(\frac{lrn_{t+1}}{l_{t+1}}\right)^{\alpha-1} (1+\mu_{t+1})\right]$$

From the equation governing the evolution of learning and the production function, we know that:

$$\frac{lrn_{t+1}}{lrn_t} = \left(\frac{lrn_t}{l_t}\right)^{\alpha-1}$$

Substituting these last two results into the optimal subsidy formula, given in Equation 2.13, and iterating forward gives:

$$\tau_t = 1 + \sum_{j=1}^{\infty} (\alpha\beta)^j E_t \left\{ \frac{lrn_{t+j+1}}{lrn_{t+1}} \right\}$$

In this special case, the optimal subsidy depends only on the expected evolution of the learning stock. Intuitively, while the learning stock is increasing the learning-labor ratio is below its steady-state level. A low learning-labor ratio corresponds to a high marginal product of learning and, hence, a high value of μ_t. Consequently, a shock that is expected to increase non-commodity output in the future also increases the social value of learning. This implies a higher optimal subsidy to encourage increased non-commodity sector production. Conversely, a shock that is expected to reduce production in the non-commodity sector will lower μ_t as well as the optimal subsidy.

When $\chi > 0$, an additional factor besides the value of learning also influences the optimal subsidy. This is reflected in the term $[lrn_t/l_t]^{\chi(1-\alpha)}$ in Equation 2.12. When $\chi > 0$ the marginal product of labor in the production of learning exceeds the marginal product of labor in the production of goods. A decrease in labor supply increases the value of this term (recall that learning responds to changes in production with a lag) and, hence, raises the optimal subsidy. This term plays a role similar to fixed costs of entry in a model with nonconvexity in production. When an economy has a large stock of learning, setting production at a low level is especially costly because some of the accumulated learning stock will be lost. In contrast, when an economy has only a small stock of learning, setting production at a low level is less costly.

2.2.5 Numerical Exercises

I now provide some numerical examples to illustrate how the optimal subsidy evolves following changes in commodity prices. In what follows, I make the following parametric assumptions: $\varphi = 1$, $\alpha = 0.33$, $\chi = 0.3$, $\beta = 0.99$, although the results are not sensitive to these choices. In each exercise, I assume that the economy begins in a steady state with no net foreign assets and that agents initially expect that the price of the commodity endowment will remain constant. The market economy achieves the socially optimal allocation through use of the subsidy scheme described above.

I first consider the effect of a permanent rise in the relative price of the commodity endowment. Agents observe the price rise and correctly anticipate that it will be permanent. Figure 2.1 shows the responses of consumption, hours worked, learning, the optimal subsidy, bond holdings and output to the shock.

Higher commodity prices allow households to consume more while working less. The decrease in hours worked reduces output and, with a lag, the stock of learning. While learning is decreasing the optimal subsidy falls below its steady state level. That is, in this example, the impact of lower expected marginal products of learning on the value of μ_t overwhelms the costs of denuding the economy's learning stock.

Figure 2.1 also shows dynamic responses to a transitory change in the price of the commodity. In this case, the commodity price boom lasts for two periods, after which the price returns to its steady-state level. Once again, I assume that agents are aware of the duration of the commodity boom when it occurs. As was the case following the permanent shock, consumption increases, while hours, learning and output all decrease. But the magnitude of these changes are much smaller than in the previous case, reflecting the smaller impact of a transitory shock on permanent income. The optimal subsidy also decreases, although by less than when the shock is permanent.

As a second exercise, I consider another rise in the price of the commodity endowment. But, in this scenario, I assume that agents are unsure of the persistence of the boom. Instead they believe that there is a 50 per cent chance that the shock will turn out to be temporary,

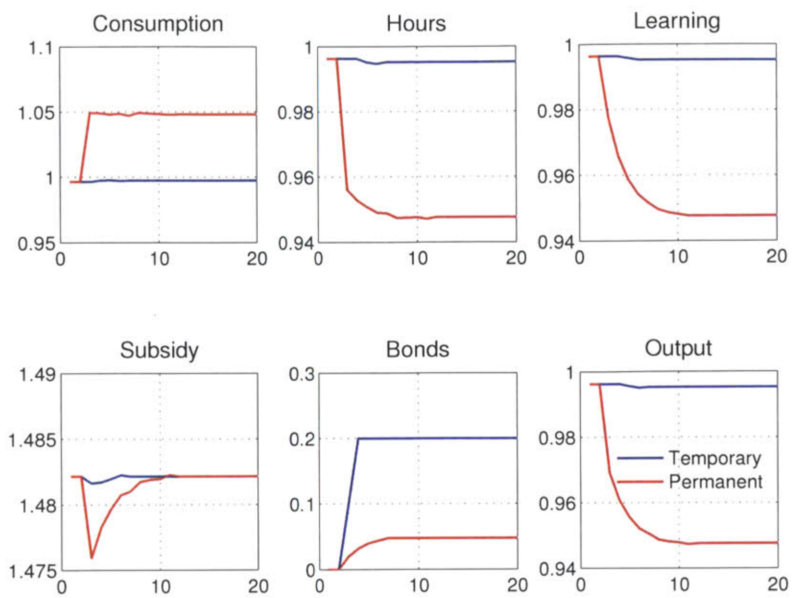

Figure 2.1: Commodity price shock: Full information. Response of variables to a ten per cent increase in the relative price of the economy's commodity endowment in period three. The red line shows the response when the increase in the commodity price is permanent. The blue line shows the response when the increase in the commodity price lasts for two periods, after which it reverts to its original value.

meaning that commodity prices will return to their original level after one period. If this does not occur, then prices remain permanently higher.

In the period in which the shock occurs - when agents are uncertain of its duration - the economy responds in a qualitatively similar way to the permanent shock in the first exercise. But, in terms of magnitudes, the increase in consumption and contraction in output are both smaller than they are when agents know that the shock is permanent. As a consequence, the initial decrease in learning is also smaller, while agents accumulate substantially more foreign assets. The decrease in the subsidy is also smaller, although the optimal subsidy is still below its steady state level.

If the shock turns out to be temporary, consumption, hours worked and learning rapidly return to near their original levels. The optimal subsidy increases above its steady state level, which supports the recovery in output. If the shock is permanent, consumption increases, hours worked and output decrease, and the subsidy decreases further to hasten the contraction in output.

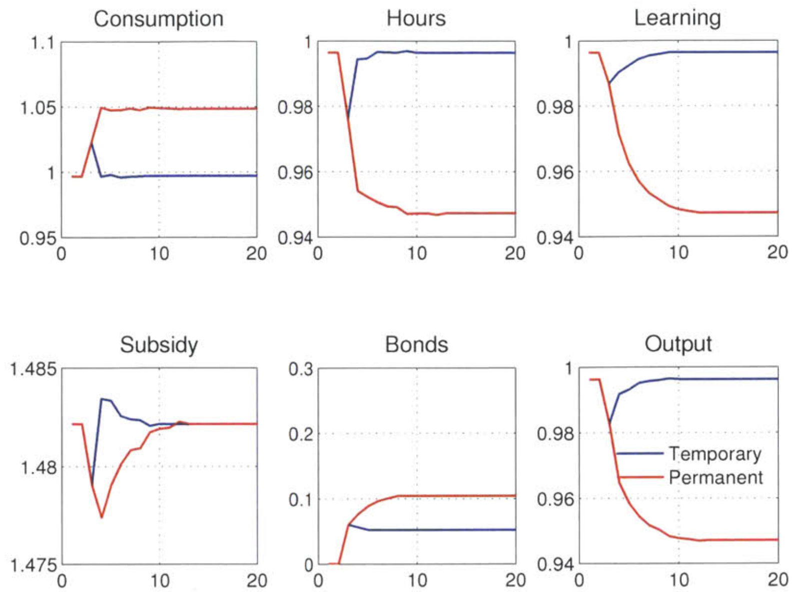

Figure 2.2: Commodity price shock: Incomplete information. Response of variables to a ten per cent increase in the relative price of the economy's commodity endowment in period three. The red line shows the response when the increase in the commodity price is permanent. The blue line shows the response when the increase in the commodity price lasts for two periods, after which it reverts to its original value. Agents initially place a 50 per cent probability on the shock being permanent.

These results demonstrate that, while learning-by-doing externalities provide a reason to subsidize domestic production, they do not imply that it is optimal to provide additional support during a commodity price boom that contracts the non-commodity sector. In fact, to the extent that a prolonged increase in commodity prices leads to a permanently smaller non-commodity sector, support for this sector may optimally decrease following a boom in commodity prices. When agents (including the social planner) are unsure whether a shock will be temporary or permanent, it is optimal to restrict the contraction of the non-commodity sector relative to the case where agents know that the shock is permanent. But, rather than providing an increased subsidy for the the non-commodity sector, this may take the form of a smaller reduction in the production subsidy, followed by an increase in the subsidy if it becomes clear that the price shock is temporary.

2.3 The Monetary Model

Having demonstrated the basic mechanisms in a flexible price environment, I now present the full model that I use to examine optimal monetary policy. I focus on a single small open economy, which I refer to as Home, that interacts with the rest of the world, which I refer to as Foreign. I modify the New Keynesian small open economy framework of Gali and Monacelli (2005) in three respects. First, I assume that financial markets are incomplete. Specifically, agents in the model may trade only a single-period risk-free bond, denominated in units of the Foreign-produced good. Second, I provide the home economy with an endowment of a tradeable commodity each period. The price of the commodity in units of the Foreign good is exogenous to the Home economy. Third, I introduce learning-by-doing to the non-commodity sector in the Home economy.

2.3.1 Households

The economy is populated by an infinitely-lived representative household that maximizes its expected discounted utility. This is a function of consumption, C_t, and labor, L_t.

$$U_{t_0} \equiv E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\sigma}}{1-\sigma} - A_L \frac{L_t^{1+\varphi}}{1+\varphi} \right] \tag{2.14}$$

where β is the household's discount factor, $1/\sigma$ is the intertemporal elasticity of substitution and φ is the Frisch labor supply elasticity. The aggregate consumption bundle is a constant elasticity of substitution (CES) aggregate of home, $C_{H,t}$, and foreign, $C_{F,t}$, goods:

$$C_t \equiv \left[(1-\eta)^{1/\gamma} (C_{H,t})^{\frac{\gamma-1}{\gamma}} + \eta^{1/\gamma} (C_{F,t})^{\frac{\gamma-1}{\gamma}} \right]^{\frac{\gamma}{\gamma-1}} \tag{2.15}$$

The parameter η indexes the degree of home bias, which is a measure of openness. The home and foreign goods are themselves CES aggregates of individual varieties:

$$C_{H,t} \equiv \left(\int_0^1 C_{H,t}(j)^{\frac{\zeta-1}{\zeta}} dj\right)^{\frac{\zeta}{\zeta-1}} \tag{2.16}$$

$$C_{F,t} \equiv \left(\int_0^1 C_{F,t}(j)^{\frac{\zeta-1}{\zeta}} dj\right)^{\frac{\zeta}{\zeta-1}} \tag{2.17}$$

where γ is the elasticity between goods produced in the home and foreign economies and ζ is the elasticity between varieties produced within a single country.

The household maximizes (2.14) subject to a sequence of budget constraints of the form:

$$\int_0^1 P_{H,t}(j)C_{H,t}(j)dj + \int_0^1 P_{F,t}(j)C_{F,t}(j)dj + \frac{B_{H,t+1}}{1+r_t} + \frac{\mathscr{E}_t B_{F,t+1}}{(1+r^*)\Psi(b_{F,t+1})} \tag{2.18}$$
$$\leq W_t L_t + \Upsilon_t + B_{H,t} + \mathscr{E}_t B_{F,t} + T_t + \mathscr{E}_t S_t Y^X$$

for $t = 0, 1, 2, \cdots$. In the inequality, $P_{H,t}(j)$ and $P_{F,t}(j)$ are the prices of individual good varieties expressed in domestic currency. W_t is the nominal wage, Υ_t are nominal profits and T_t represent nominal lump sum transfers, all expressed in domestic currency. The assets of the households comprise $B_{H,t+1}$, which are holdings of domestic bonds, and $B_{F,t+1}$, which are holdings of foreign bonds. Domestic bonds are denominated in domestic currency, while foreign bonds are denominated in foreign currency. The rate of return on domestic and foreign bonds are r_t and r^*. $\mathscr{E}_t$ is the bilateral nominal exchange rate between the Home and Foreign currency (that is, the price of the foreign country's currency in terms of domestic currency). The function $\Psi(b_{t+1}) = \frac{\psi}{2}\left(\exp(-B_{F,t+1}\mathscr{E}_t/P_t)\right)$ is a portfolio adjustment cost term that penalizes the domestic economy when it has a net foreign asset position different from zero. As discussed in Schmitt-Grohe and Uribe (2003), including this term ensures that the foreign asset level is stationary. The final element of the inequality relates to the commodity sector. Y^X is the home economy's constant commodity endowment and S_t is the relative price of the commodity in terms of foreign currency. The commodity is entirely exported and so does not enter the household's consumption basket.

The optimal allocation of expenditure within each category of goods yields the demand functions:

$$C_{H,t}(j) = \left(\frac{P_{H,t}(j)}{P_{H,t}}\right)^{-\zeta} C_{H,t}; \quad C_{F,t}(j) = \left(\frac{P_{F,t}(j)}{P_{F,t}}\right)^{-\zeta} C_{F,t} \tag{2.19}$$

for all $j \in [0,1]$, where $P_{H,t} \equiv \left(\int_0^1 P_{H,t}(j)^{1-\zeta} dj\right)^{\frac{1}{1-\zeta}}$ is the home-produced goods price index (PPI) and $P_{F,t} \equiv \left(\int_0^1 P_{F,t}(j)^{1-\zeta} dj\right)^{\frac{1}{1-\zeta}}$.

The optimal allocation of expenditures between domestic and imported goods is given by:

$$C_{H,t} = \left(\frac{P_{H,t}}{P_t}\right)^{-\gamma} C_t; \quad C_{F,t} = \left(\frac{P_{F,t}}{P_t}\right)^{-\gamma} C_t \tag{2.20}$$

where $P_t \equiv \left[(1-\eta)(P_{H,t})^{1-\gamma} + \eta (P_{F,t})^{1-\gamma}\right]^{\frac{1}{1-\gamma}}$ is the consumer price index (CPI). In the standard fashion, these results imply that total consumption expenditure by the representative household is equal to $P_t C_t$.

The optimality conditions for the household's problem are:

$$A_L C_t^{\sigma} L_t^{\varphi} = \frac{W_t}{P_t} \tag{2.21}$$

which is the standard intratemporal optimality condition,

$$1 = \beta E_t \left\{ \left(\frac{C_{t+1}}{C_t}\right)^{-\sigma} \frac{1+r_t}{\Pi_{t+1}} \right\} \tag{2.22}$$

where $\Pi_t = P_t / P_{t-1}$ is the rate of CPI inflation in the Home country, which is the household's Euler equation and,

$$1 + r_t = (1 + r^*) E_t \left\{ \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} \Psi(b_{t+1}) \right\} \tag{2.23}$$

which is the uncovered interest rate parity (UIP) condition.

2.3.2 International Relative Prices: Some Identities

In what follows, it will be useful to define some indices of international relative prices. First, the bilateral real exchange rate between the Home and Foreign economies, Q_t, is:

$$Q_t \equiv \frac{\mathscr{E}_t P_{F,t}}{P_t} \tag{2.24}$$

Second, the relative price of Home goods, $\hat{P}_{H,t}$, is:

$$\hat{P}_{H,t} \equiv \frac{P_{H,t}}{P_t} \tag{2.25}$$

Rearranging this gives an expression for PPI inflation in terms of CPI inflation and the rate of change of $\hat{P}_{H,t}$:

$$\Pi_{H,t} = \frac{\hat{P}_{H,t}}{\hat{P}_{H,t-1}} \Pi_t \tag{2.26}$$

The law of one price holds for individual goods at all times. This implies that $P_{Ft}(j) = \mathscr{E}_t P^*_{F,t}(j)$ and $P_{H,t}(j) = \mathscr{E}_t P^*_{H,t}(j)$ for all $j \in [0,1]$.

2.3.3 Home Firms

Production: Home firms produce differentiated goods using the technology:

$$Y_t(j) = LRN_t(j)^{\alpha} L_t(j)^{1-\alpha} \tag{2.27}$$

where $Y_t(j)$ is output of firm j. As in the simple example described above, the term $LRN_t(j)$ reflects a learning-by-doing productivity shifter. It evolves according to:

$$LRN_{t+1}(j) = LRN_t(j)^{\chi} Y_t^{1-\chi} \tag{2.28}$$

where once again the evolution of learning depends upon aggregate output, which each firm takes as given when making its pricing and production decisions.

Real marginal cost for firm j is given by:

$$MC_t(j) = (1-\tau)\frac{W_t}{P_{H,t}}\left[(1-\alpha)\frac{Y_t(j)}{L_t(j)}\right]^{-1} \tag{2.29}$$

where τ is a production subsidy, funded by lump-sum taxation. I set the value of this subsidy to eliminate distortions caused by monopolistic competition and the learning-by-doing externality in the steady state. Its derivation is presented in Appendix 2. B.1.

Price setting: Producers set prices in a staggered fashion, as in Calvo (1983). Each period a fraction, $1-\theta$, of randomly selected firms are able to reset their prices, with an individual firm's probability of re-optimizing in any given period being independent of the time elapsed since it last reset its price.

The problem facing a firm that sets its price at time t is:

$$\max_{\bar{P}_{H,t}(j)} \sum_{k=0}^{\infty} \theta^k E_t\left\{\mathscr{Q}_{t,t+k}\left[Y_{t+k}(j)\left(\bar{P}_{H,t}(j) - MC_{t+k}(j)P_{H,t+k}\right)\right]\right\} \tag{2.30}$$

subject to the demand constraint that $Y_{t+k}(j) = \left(\bar{P}_{H,t}(j)/P_{H,t+k}\right)^{-\zeta} Y_t$. $\bar{P}_t^H$ is the price chosen by a firm that resets its price at time t and $\mathscr{Q}_{t,t+k}$ is the household's stochastic discount factor between period t and $t+k$.

The optimality condition for the firm's pricing problem is:

$$E_t\left[\sum_{k=0}^{\infty} \theta^k \mathscr{Q}_{t,t+k} Y_{t+k}(j)\left[\frac{\bar{P}_{H,t}}{P_{H,t+k}} - \frac{\zeta}{\zeta-1}MC_{t+k}(j)\right]\right] = 0 \tag{2.31}$$

As all firms that reset their price at time t will choose the same price, this expression can be rewritten in terms of the difference between the reset price and the average price in the home goods sector:

$$\frac{\bar{P}_{H,t}}{P_{H,t}} \equiv \Xi_t = \left(\frac{F_t}{J_t}\right)^{\frac{1-\alpha}{1-\alpha+\alpha\zeta}} \tag{2.32}$$

where:[2]

$$F_t = \frac{\zeta}{\zeta - 1} C_t^{-\sigma} Y_t MC_t \hat{P}_{H,t} + \beta\theta E_t \left[(\Pi_{H,t+1})^{\frac{\zeta}{1-\alpha}} F_{t+1} \right] \tag{2.33}$$

and

$$J_t = C_t^{-\sigma} Y_t \hat{P}_{H,t} + \beta\theta E_t \left[(\Pi_{H,t+1})^{\zeta-1} J_{t+1} \right] \tag{2.34}$$

Under the assumed pricing arrangements, the domestic price index is:

$$P_{H,t} = \left[\theta P_{H,t-1}^{1-\zeta} + (1-\theta) \bar{P}_{H,t}^{1-\zeta} \right]^{\frac{1}{1-\zeta}}$$

By maniputing this expression, we can obtain an expression for the inflation rate of home-produced goods in terms of Ξ_t:

$$1 = \theta (\Pi_{H,t})^{\zeta-1} + (1-\theta)(\Xi_t)^{1-\zeta} \tag{2.35}$$

Finally, define price dispersion, Δ_t as:

$$\Delta_t = \int_0^1 \left(\frac{P_{H,t}(j)}{P_{H,t}} \right)^{-\zeta/(1-\alpha)} dj$$

By the law of large numbers, it can be shown that this is equivalent to:

$$\Delta_t = \theta \Delta_{t-1} (\Pi_{H,t})^{\zeta/(1-\alpha)} + (1-\theta)(\Xi_t)^{-\zeta/(1-\alpha)} \tag{2.36}$$

[2]Note that this expression contains average marginal cost, $MC_t = (1-\tau)\frac{W_t/P_t^H}{(1-\alpha)Y_t/L_t}$, rather than firm-specific marginal cost, $MC_t(j)$. Appendix 2. A shows the derivation of the firm's pricing problem.

2.3.4 Commodity Sector

The output of the commodity sector is constant and equal to Y^X. The relative price between commodities and foreign goods is unaffected by developments in the Home economy and is denoted by S_t. Hence, revenue from the commodity sector in domestic currency is $\mathcal{E}_t S_t Y^X$.

I assume that the relative price of commodities is a function of two shocks:

$$S_t = Z_t G_t \tag{2.37}$$

The first shock, Z_t follows an AR(1) process in logs:

$$z_t = \rho_z z_{t-1} + \varepsilon_t^z \tag{2.38}$$

where $z_t = \log Z_t$ and $\varepsilon_t^z \sim \mathcal{N}\left(0, \sigma_z^2\right)$.

The second shock, G_t, follows an AR(2) process that captures a growth rate component and imposes an error correction term to ensure stationarity in levels. Specifically, defining $g_t = \log G_t$:

$$\Delta g_t = \rho_{g1} \Delta g_{t-1} - \rho_{g2} g_{t-1} + \varepsilon_t^g \tag{2.39}$$

where $\varepsilon_t^g \sim \mathcal{N}\left(0, \sigma_g^2\right)$ and ρ_{g2} is a small value.[3]

Although the two shocks both directly affect the level of commodity prices, they differ in their persistence. A positive impulse to ε_t^z provides a one-off boost to commodity prices that decays over time. I refer to this as the transitory shock. In contrast, a positive impulse to ε_t^g increases commodity prices on impact, and then continues to increase them further in subsequent periods. That is, a shock to ε_t^g has a considerably more persistent impact on the level of commodity prices. I refer to this as the persistent shock.

[3]Note that this process can be re-written as an AR(2) process as $g_t = (1 + \rho_{g1} - \rho_{g2}) g_{t-1} - \rho_{g1} g_{t-2} + \varepsilon_t^g$. If $\rho_{g1} = \rho_{g2} = 0$ the process follows a random walk.

2.3.5 Foreign Economy

I abstract from all foreign shocks, except for the commodity price shock. I also assume that preferences in the foreign economy are as in the home economy. Finally, I normalize the price of foreign goods in foreign currency to one. Hence, the price of foreign goods in terms of domestic currency, $P_{F,t}$, is equal to $\mathscr{E}_t$.

2.3.6 Monetary Policy

Under optimal monetary policy, the nominal interest rate will be a complex function of the model's parameters and shocks. I also consider the performance of simple policy rules. In this case, monetary policy in the domestic economy is given by a Taylor-type rule of the form:

$$\left(\frac{1+r_t}{1+r^*}\right) = \left(\frac{1+r_t}{1+r^*}\right)^{\psi_r} \left[\left(\frac{\Pi_t}{\Pi}\right)^{\psi_\pi} \left(\frac{\Pi_{H,t}}{\Pi_H}\right)^{\psi_{\pi H}} \left(\frac{Y_t}{Y}\right)^{\psi_y} \left(\frac{\mathscr{E}_t}{\mathscr{E}_{t-1}}\right)^{\psi_e}\right]^{1-\psi_r} \tag{2.40}$$

2.3.7 Aggregation and Market Clearing:

Aggregate GDP is defined analogously to aggregate home consumption as: $Y_t \equiv \left[\int_0^1 Y_t(j)^{\frac{\zeta-1}{\zeta}} dj\right]^{\frac{1}{1-\zeta}}$. The relationship between aggregate labor supply and the labor demand of firms is given by:

$$L_t = \left[\frac{Y_t}{LRN_t^{\alpha}}\right]^{\frac{1}{1-\alpha}} \Delta_t \tag{2.41}$$

Market clearing for home goods is:

$$Y_t = (1-\eta)\left(\hat{P}_{H,t}\right)^{-\gamma} C_t + \eta \left(\frac{\hat{P}_{H,t}}{Q_t}\right)^{-\gamma} C^* \tag{2.42}$$

where the right-hand-side of the expression uses the fact that the foreign price index is normalized to one and the foreign economy faces no shocks.

Finally, I assume that home-denominated bonds are traded only domestically and are in zero net supply. Then, the current account equation is:

$$C_t + \frac{b_{t+1}}{(1+r^*)\Psi(b_{t+1})} = \hat{P}_{H,t}Y_t + Q_tS_tY^X + \frac{Q_t}{Q_{t-1}}b_t \tag{2.43}$$

2.3.8 Equilibrium

An equilibrium is a sequence of quantities $\{C_t, Y_t, L_t, LRN_{t+1}, B_{t+1}\}_{t=1}^{\infty}$, prices $\{W_t, \Pi_t, \Pi_{H,t}, \hat{P}_{H,t}, Q_t, r_t, \Xi_t, F_t, J_t, \Delta_t\}_{t=1}^{\infty}$ and exogenous processes $\{S_t, Z_t, G_t\}_{t=0}^{\infty}$ such that households maximize utility (Equations 2.21-2.23), firms maximize profits (Equations 2.32-2.36) and markets clear (Equations 2.41-2.43), subject to the technical conditions (Equations 2.24, 2.26-2.28), the exogenous processes (Equations 2.37-2.39) and a monetary policy rule.

2.3.9 Information

The key innovation of this paper is to examine the optimal conduct of monetary policy in an environment in which agents cannot observe the persistence of a commodity price shock directly. Specifically, I assume that agents can observe commodity prices, S_t, but cannot observe Z_t and G_t. I do, however, allow agents to receive a noisy signal about the value of the persistent shock. I refer to this signal as h, and assume that it evolves according to the process:

$$h_t = g_t + \varepsilon_t^h \tag{2.44}$$

where $\varepsilon_t^h \sim \mathcal{N}\left(0, \sigma_h^2\right)$. The agent's information set as of time t includes the entire history of commodity price shocks and signals; $I_t \equiv \{s_t, h_t, s_{t-1}, h_{t-1}, \ldots\}$, where $s_t = \log S_t$.[4]

Agents form expectations about the decomposition of commodity prices using the Kalman filter. The measurement equation for the filter includes the definitions of the observable commodity price variables, s_t and h_t:

$$\underbrace{\begin{bmatrix} s_t \\ h_t \end{bmatrix}}_{\Gamma_t} = \underbrace{\begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \end{bmatrix}}_{C} \underbrace{\begin{bmatrix} z_t \\ g_t \\ g_{t-1} \end{bmatrix}}_{X_t} + \underbrace{\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{D} \underbrace{\begin{bmatrix} \varepsilon_t^z \\ \varepsilon_t^g \\ \varepsilon_t^h \end{bmatrix}}_{V_t} \tag{2.45}$$

The transition equation summarizes the evolution of the unobserved variables, and is given by:

$$\underbrace{\begin{bmatrix} z_t \\ g_t \\ g_{t-1} \end{bmatrix}}_{X_t} = \underbrace{\begin{bmatrix} \rho_z & 0 & 0 \\ 0 & 1+\rho_{g1}-\rho_{g2} & -\rho_{g1} \\ 0 & 1 & 0 \end{bmatrix}}_{A} \underbrace{\begin{bmatrix} z_{t-1} \\ g_{t-1} \\ g_{t-2} \end{bmatrix}}_{X_{t-1}} + \underbrace{\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{B} \underbrace{\begin{bmatrix} \varepsilon_t^z \\ \varepsilon_t^g \\ \varepsilon_t^h \end{bmatrix}}_{V_t} \tag{2.46}$$

where $V_t \sim N(0, Q)$ and $Q \equiv \begin{bmatrix} \sigma_z^2 & 0 & 0 \\ 0 & \sigma_g^2 & 0 \\ 0 & 0 & \sigma_h^2 \end{bmatrix}$.

The Kalman filter can be used to express the agents' beliefs about the components of commodity prices as:

$$X_{t|t} = (I - KC) A X_{t-1|t-1} + K\Gamma_t \tag{2.47}$$

where $X_{t|t} \equiv E_t\left[X_t | I_t\right]$ are the households' expectations of the unobservable variables given their information set at time t. K is the Kalman gain, calculated as:

[4] I also assume that agents are able to observe the all of the non-commodity variables in the model. However, observing these variables does not help agents to identify z_t or g_t.

$$K = PC' \left(CPC' + DQD'\right)^{-1} \tag{2.48}$$

where P is the steady-state error covariance matrix, calculated as the solution to:

$$P = APA' - APC' \left(CPC'\right)^{-1} CPA' + BQB' \tag{2.49}$$

If $\sigma_h^2 = 0$, then h_t provides a perfect signal of the persistent shock. Using this signal and their observation of S_t, agents are able to infer the persistence of a shock exactly. I refer to this as the full information case. If $\sigma_h^2 > 0$, agents must form beliefs about Z_t and G_t based on their observations of commodity prices and the noisy signal. I refer to this environment as the incomplete information case.

2.3.10 Log-linearized Equations

I solve the model by log-linearizing around the steady state of the non-linear model.[5] In what follows, lower case letters refer to log-deviations from a variable's steady state. That is, $x_t = \left(X_t - \bar{X}\right)/\bar{X}$. The log-linear optimality conditions for the households are:[6]

$$\begin{aligned} c_t &= E_t\{c_{t+1}\} - \frac{1}{\sigma}\left(r_t - E_t\{\pi_{t+1}\}\right) \\ r_t &= E_t\{\Delta q_{t+1} + \pi_{t+1}\} - \psi b_{t+1} \end{aligned}$$

The optimality conditions for domestic firms are:

[5]The steady-state of the model is described in Appendix 2. B.1.

[6]In the UIP equation, note that $\mathscr{E}_t = Q_t P_t$ and recall that r^* is constant.

$$
\begin{aligned}
\xi_t &= \frac{1-\alpha}{1-\alpha+\alpha\zeta}(f_t - j_t) \\
f_t &= (1-\beta\theta)(1+\varphi)l_t + \beta\theta E_t\left\{\frac{\zeta}{1-\alpha}\pi_{H,t+1} + f_{t+1}\right\} \\
j_t &= (1-\beta\theta)(y_t - \sigma c_t + \hat{p}_{H,t}) + \beta\theta E_t\{(\zeta-1)\pi_{H,t+1} + j_{t+1}\} \\
\pi_{H,t} &= \frac{1-\theta}{\theta}\xi_t
\end{aligned}
$$

The aggregators imply that:

$$
\begin{aligned}
\pi_{H,t} &= (\hat{p}_{H,t} - \hat{p}_{H,t-1}) + \pi_t \\
\hat{p}_{H,t} &= -\frac{\eta}{1-\eta}q_t
\end{aligned}
$$

The market clearing conditions imply that:

$$
\begin{aligned}
y_t &= \alpha lrn_t + (1-\alpha)l_t \\
lrn_{t+1} &= \chi lrn_t + (1-\chi)y_t \\
y_t &= -\gamma\hat{p}_{H,t} + \frac{(1-\eta)\bar{C}}{\bar{Y}}c_t + \frac{\bar{Y}-(1-\eta)\bar{C}}{\bar{Y}}\gamma q_t \\
c_t + \beta b_{t+1} &= \frac{\bar{Y}}{\bar{C}}(y_t + \hat{p}_{H,t}) + \frac{\bar{C}-\bar{Y}}{\bar{C}}(s_t + q_t) + b_t
\end{aligned}
$$

Finally, the price of the commodity endowment and its components evolve according to:

$$
\begin{aligned}
s_t &= z_t + g_t \\
z_t &= \rho_z z_{t+1} + \varepsilon_t^z \\
g_t &= (1+\rho_{g1} - \rho_{g2})g_{t-1} - \rho_{g2} + \varepsilon_t^g
\end{aligned}
$$

Note that the optimality conditions for the firm imply that:

$$\pi_{H,t} = \beta E_t\left\{\pi_{t+1}^H\right\} + \kappa\left(\varphi l_t + \sigma c_t - \hat{p}_t^H + \frac{\alpha}{1-\alpha}\left(y_t - lrn_t\right)\right)$$

where $\kappa \equiv (1-\beta)(1-\beta\theta)(1-\alpha)/(\theta(1-\alpha+\alpha\zeta))$, which is the standard open economy New Keynesian Phillips Curve.

2.4 Model Solution

2.4.1 Solution

The solution of the model follows Uhlig (1999) and Blanchard et al. (2012). Let $\mathscr{Y}_t$ denote the endogenous variables controlled by the agent. The economic model can be represented as the stochastic difference equation:

$$FE_t\left\{\mathscr{Y}_{t+1}\right\} + G\mathscr{Y}_t + H\mathscr{Y}_{t-1} + M\Gamma_t = 0 \tag{2.50}$$

where F, G, H, M are matrices of parameters. The unique stable solution of the model is:

$$\mathscr{Y}_t = P\mathscr{Y}_{t-1} + Q\Gamma_t + RX_{t|t} \tag{2.51}$$

where the matrices P, Q, R can be found by solving the three matrix equations:

$$\begin{aligned} FP^2 + GP + H &= 0 \\ (FP+G)Q + M &= 0 \\ (FP+G)R + F(QC+R)A &= 0 \end{aligned}$$

and the matrices A and C are as defined in Section 2.3.9. The matrix P can be recovered using the techniques discussed in Uhlig (1999).

2.4.2 Approximating Welfare

To evaluate welfare, I follow the linear-quadratic approach of Sutherland (2002) and Benigno and Woodford (2012). This approximates welfare by taking a second-order Taylor expansion of the representative household's utility function and then using second-order approximations of the model's equilibrium conditions to eliminate linear terms from the welfare approximation. This leaves a purely quadratic approximation to the objective function that can be analyzed using a first-order approximation to the model. More formally, the model presented above is of the form:

$$\begin{aligned} \max_{\{\mathcal{Y}_t\}} E_0 \sum_{t=0}^{\infty} \beta^t U\left(\mathcal{Y}_{t-1}, \mathcal{Y}_t, V_t\right) & \qquad (2.52) \\ s.t.\ E_t\left\{g\left(\mathcal{Y}_{t-1}, \mathcal{Y}_t, \mathcal{Y}_{t+1}, V_t, V_{t+1}\right)\right\} = 0,\ t \geq 0 & \end{aligned}$$

where $\mathcal{Y}_t$ is a vector of endogenous variables, V_t is a vector of exogenous processes, U is a standard utility function, g is a vector of constraints and E_t denotes rational expectations conditional on information available at time t. Denoting the lagrange multipliers associated with the constraints at time t by λ_t, the first order conditions for $t > 0$ are:

$$U_{2,t} + \beta U_{1,t+1} + \beta \lambda_{t+1} g_{1,t+1} + \lambda g_{2,t} + \beta^{-1} \lambda_{t-1} g_{3,t-1} = 0 \qquad (2.53)$$

$$E_t\left\{g\left(\mathcal{Y}_{t-1}, \mathcal{Y}_t, \mathcal{Y}_{t+1}, V_t, V_{t+1}\right)\right\} = 0 \qquad (2.54)$$

where $a_{j,t}$ denotes the first derivative of the function $a(\cdot)$ with respect to the jth argument. In a steady state for $t > 0$, it must be that:

$$U_2 + \beta U_1 + \beta \lambda g_1 + \lambda g_2 + \beta^{-1} \lambda g_3 = 0 \qquad (2.55)$$

$$g = 0 \qquad (2.56)$$

The second order Taylor approximation to utility around the steady state is:

$$\sum_{t=0}^{\infty} \beta^t U(\mathscr{Y}_{t-1}, \mathscr{Y}_t, V_t) \approx \sum_{t=0}^{\infty} \beta^t \left(U_2(\mathscr{Y}_t - \mathscr{Y}) + \beta U_1(\mathscr{Y}_t - \mathscr{Y})\right) + s.o.t. + t.i.p. \tag{2.57}$$

where $s.o.t.$ stands for second order terms and $t.i.p.$ stands for terms independent of policy. It is well known that maximization of this approximation to utility using a linearized approximation to the constraints, g, may not deliver a correct local linear approximation to true optimal policy.[7]

To resolve this problem, one first takes a second order approximation of the constraints in Equation 2.54 and sums over time, delivering:[8]

$$\begin{aligned} \sum \beta^t g(\mathscr{Y}_{t-1}, \mathscr{Y}_t, \mathscr{Y}_{t+1}, V_t, V_{t+1}) \;\approx\; & \sum_{t=0}^{\infty} \beta^t \Big(\beta g_1(\mathscr{Y}_t - \mathscr{Y}) + g_2(\mathscr{Y}_t - \mathscr{Y}) \\ & + \beta^{-1} g_3(\mathscr{Y}_t - \mathscr{Y})\Big) + s.o.t. + t.i.p. \end{aligned} \tag{2.58}$$

Multiplying this expression by λ, one can then use the steady state relationship given in Equation 2.55 to eliminate the linear terms in Equation 2.57. This leaves a purely quadrative objective function, which can be maximized subject to linear constraints to deliver the solution to the optimal policy problem (Benigno and Woodford 2012, Debortoli and Nunes 2006).

Given the large dimension of the optimal policy problem in this paper, I calculate a Matlab-coded version of the linear-quadratic welfare approximation.

[7] For an example of where this naiive LQ approximation fails, see Kim and Kim (2003)

[8] In taking this approximation, I assume that the constraint in Equation 2.54 is also present at $t = -1$. This is consistent with analyzing welfare from a timeless perspective - in which the policymaker chooses policy according to a self-consistent rule that applies to $t = 0$ as well as subsequent periods.

2.4.3 Calibration

Table 2.1 shows the baseline parameter values that I use. Most parameters are based on standard values in the literature or, where available, are taken from papers that have estimated these parameters for commodity-producing small open economies of the type that are likely to be affected by the issues discussed in this paper. In Section 2.7 I discuss the sensitivity of the results to alternative parametric assumptions.

For the commodity price processes, I calibrate the parameters using the results of Rees (2012), who estimated the parameters for a similar process for the Australian terms of trade. I set the persistence of the long-lived shock, ρ_{g1}, equal to 0.86. Recalling that this is a shock to the growth rate of commodity prices, this parameter implies that a small shock to g cumulates over time and ultimately has a large impact on the level of commodity prices. I set the error correction term, ρ_{g2}, equal to 0.001. This small value implies that the error correction term has little impact on the dynamics of commodity prices following a persistent shock while ensuring that commodity prices remain stationary. I set the persistence of the transitory shock equal to 0.84, implying a half life of five quarters. Finally, I set the standard deviations of the two shocks, σ_g and σ_z, equal to 0.002 and 0.02 and the standard deviation of the noise shock, σ_h equal to 0.06.

2.5 Optimal Monetary Policy

In this section, I describe the model economy's response to transitory and persistent commodity price shocks when agents have incomplete information and monetary policy is set optimally. For comparison, I contrast the responses to those when agents have full information about the persistence of commodity price shocks.

2.5.1 Transitory Shock

Figure 2.3 plots the economy's response to a transitory commodity price boom. The shock increases commodity prices by around two per cent on impact. Commodity prices fall in

Table 2.1: Baseline Parameter Values

Parameter		Description	Source
β	0.99	Discount factor	Standard value in literature.
σ	1	Intertemporal elasticity of substitution	Specifying a log utility function.
φ	1	Frisch labor supply elasticity	Standard value in literature.
ψ	5×10^{-3}	Portfolio adjustment cost	Small value to ensure stationarity.
θ	0.75	Calvo parameter	Gali and Monacelli (2005)
α	0.33	Learning share of income	Lama and Medina (2012)
χ	0.63	Rate of depreciation rate of learning	Lama and Medina (2012)
η	0.30	Share of imports in consumption basket	This implies a 30 per cent import share of consumption.
γ	1.3	Elasticity of substitution between domestic and foreign goods	Jaaskela and Nimark (2011)
ζ	6	Elasticity of substitution between varieties	Gali and Monacelli (2005)
Y^x/C	0.20	Commodity endowment	This implies that commodities account for two-thirds of export values in steady state.
		Commodity Price Process	
ρ_{g1}	0.86	AR(1) coefficient in persistent commodity price process.	
ρ_{g2}	1×10^{-3}	Error correction coefficient in persistent commodity price process.	
ρ_z	0.84	AR(1) coefficient in transitory commodity price process.	
σ_g	2×10^{-3}	Standard deviation of persistent commodity price shock.	
σ_z	2×10^{-2}	Standard deviation of transitory commodity price shock.	
σ_h	6×10^{-2}	Standard deviation of shock to persistent commodity price signal.	

subsequent quarters, and after six years have returned to their initial level. Under both informational assumptions, the shock leads to appreciations of the nominal and real exchange rates and a fall in consumer prices. Higher commodity prices increase the wealth of domestic households. This triggers a consumption boom and a fall in hours worked.[9]

[9]This aspect of the model might seem unrealistic. In a model with a non-traded sector, some of the decrease in hours worked in the tradeable sector would reflect a shift in labor supply to the non-tradeable sector. In contrast, in the model in this paper, it entirely reflects increased leisure. However, the implications for the tradeable sector - which is the focus of this paper - are broadly similar.

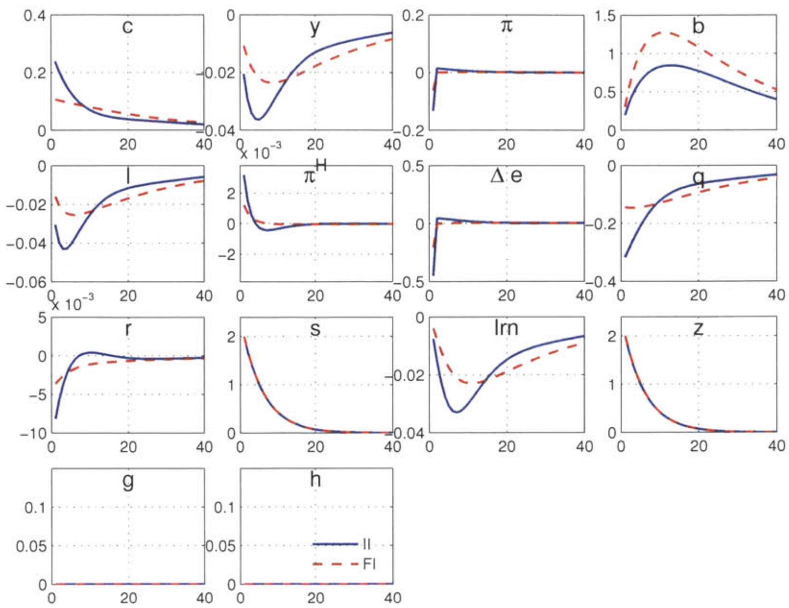

Figure 2.3: Impulse Responses to a one standard deviation shock to the transitory component of commodity prices. Solid lines show the responses under incomplete information. Dashed lines show the responses under full information.

The latter corresponds to a decrease in domestic goods production and a contraction in the stock of learning. The initial real exchange rate appreciation lowers the relative price of imported goods. This offsets some of the effect of increased consumption on the demand for home-produced tradeables. Nonetheless, the rate of inflation of home-produced tradeables is initially positive, although close to zero.

Comparing the two scenarios, it is clear that incomplete information magnifies the initial response of the economy to the shock. In particular, for the first eight quarters, consumption is stronger under incomplete information, while the contractions in hours worked, output and learning are larger. This translates into less foreign asset accumulation. In later quarters, this situation reverses and consumption is weaker under incomplete information, while output is stronger. Having accumulated fewer foreign assets while commodity prices were high, agents in the incomplete information scenario must consume less once the boom has passed. The recovery in output is supported by a persistent depreciation in the real exchange rate, which occurs far more rapidly under incomplete information than under full information.

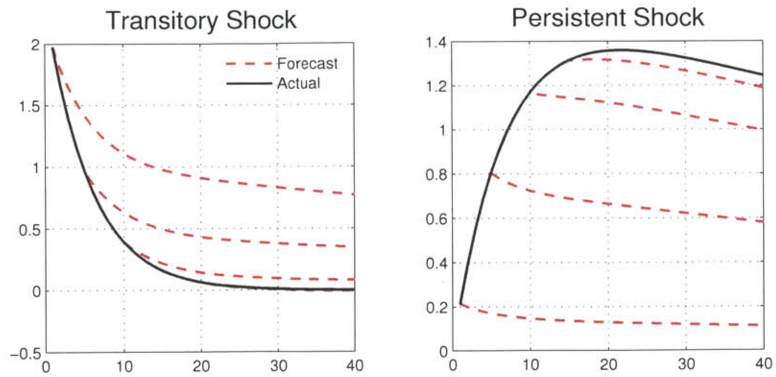

Figure 2.4: Commodity price forecasts after transitory and persistent shocks. The left panel shows forecasts following a transitory commodity price shock. The solid line shows the actual evolution of commodity prices. The dashed lines show forecasts made after one, six, eleven and sixteen quarters. The right panel shows the same following a shock to the persistent component of commodity prices.

To gain an intuition for these responses, it is useful to consider agents' beliefs about the source of changes in commodity prices. To illustrate this, the left-hand panel of Figure 2.4 shows the evolution of commodity prices as well as successive forecasts that agents in the model make based on the level of commodity prices and the noisy signal, h_t. I include forecasts for the period when the shock hits, as well as after six, eleven and sixteen quarters.

At the assumed parameter values, the signal that agents receive about the persistence of the commodity price shock is weak. Agents initially believe that there is a good chance that some of the rise in commodity prices reflects an increase in the value of the persistent shock. They therefore expect that commodity prices will remain above their initial level for some time, as illustrated by the first dashed line. A highly persistent shock implies a substantial increase in the wealth of domestic residents. This explains the large initial responses of consumption and labor supply in Figure 2.3. If they realized that the boom was transitory, policymakers may prefer to limit the growth of consumption and moderate the contraction of the non-commodity sector. However, in the model all agents, including policymakers, have the same information and so share the same optimistic beliefs about the persistence of the commodity price shock.

How does the learning-by-doing externality affect these results? To guage this, I calculate the value of the externality, defined as the marginal product of labor excluding its direct effect on output, in each scenario as follows:

$$ext_t = \mu_t + \chi (1-\alpha)(lrn_t - l_t) \tag{2.59}$$

where μ_t denotes the log deviation of the shadow value of learning from its steady-state level, calculated as:[10]

$$\begin{aligned} \mu_t &= \sigma c_t - \hat{p}_t^H - E_t \left\{ \sigma c_{t+1} - \hat{p}_{t+1}^H \right\} + \frac{\alpha\beta}{\mathscr{M}} (\alpha - 1)(lrn_{t+1} - l_{t+1}) \\ &\quad + \beta (\chi + \alpha (1-\chi)) (\mu_{t+1} - (1-\chi)(1-\alpha)(lrn_{t+1} - l_{t+1})) \end{aligned} \tag{2.60}$$

The left panel of Figure 2.5 shows how ext_t evolves following a transitory commodity price shock. Under both informational assumptions, the shock prompts an expected near-term contraction in the stock of learning. As discussed in the context of the simple model, this lowers the shadow value of learning, μ_t. However, because the shock reduces hours worked on impact while the stock of learning responds only with a lag, it also increases the marginal response of learning to additional labor, reflected in the term $\chi (1-\alpha)(lrn_t - l_t)$. In the full information case the latter effect overwhelms the former and the value of the externality increases when the shock hits. That is, firms undervalue the benefits of additional labor. In the incomplete information case, the value of the externality initially is marginally below its steady state level. While it becomes positive in subsequent quarters, it remains below its value under full information.

This last result might seem puzzling in light of the fact that the initial contraction in the economy's stock of learning is much sharper under incomplete information than it is under full information. This reflects two factors. First, information frictions in the model are

[10] This is the log-linearized form $\mathscr{M}_t = \beta E_t \left\{ \left(\frac{C_{t+1}}{C_t}\right)^{-\sigma} \frac{\hat{P}_{t+1}^H}{\hat{P}_t^H} \left[\alpha \left(\frac{LRN_{t+1}}{L_{t+1}}\right)^{\alpha-1} + (\chi + \alpha(1-\chi)) \mathscr{M}_{t+1} \left(\frac{LRN_{t+1}}{L_{t+1}}\right)^{-(1-\chi)(1-\alpha)} \right] \right\}$, where $\mu_t = \frac{\mathscr{M}_t - \mathscr{M}}{\mathscr{M}}$ and $\mathscr{M} = \frac{\alpha\beta}{1-\beta(\chi+\alpha(1-\chi))}$. Appendix 2. B.1 shows the derivation of these expressions.

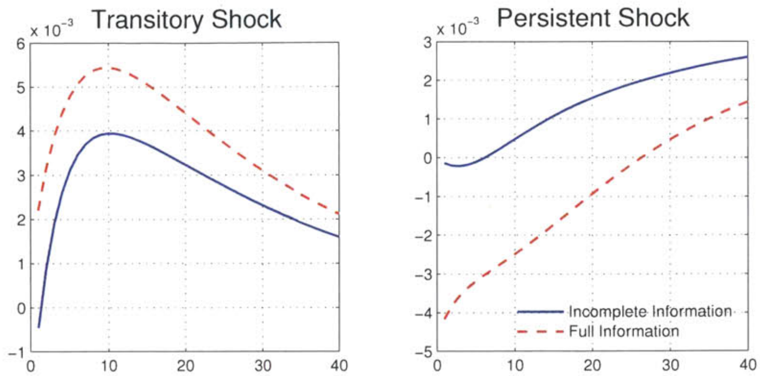

Figure 2.5: Perceived value of learning-by-doing externality following a transitory (left panel) and persistent (right panel) commodity price shock. Solid lines show values when agents have incomplete information. Dashed lines show values when agents have full information.

highly persistent. This is illustrated in Figure 2.4 where even eleven quarters after the transitory shock, when more than half of the initial increase in commodity prices has disipated, agents still expect that some of the boom will prove to be long-lived. These optimistic commodity price forecasts reduce the perceived value of additional learning. Second, the sharp decrease in learning under incomplete information reduces the marginal impact of additional labor supply on the learning stock. This also depresses the value of the externality. In the sensitivity analysis section of the paper, I discuss how the value of the externality would evolve under alternative assumptions about information acquisition.

In sum, under optimal monetary policy a transitory commodity price boom leads to a quantitatively small increase in PPI inflation, an increase in consumption and a contraction in the non-commodity tradeable sector. The contraction is larger under incomplete information than under full information, for two reasons. First, under incomplete information the expected impact of the shock on permanent income is larger, reflecting the fact that agents believe that the shock may be persistent. Second, the value of the externality is lower under incomplete information than it is under full information. As well as these effects, monetary policymakers must also weigh up other considerations, including the costs of departures from price stability and the benefits of strategic terms of trade manipulation of the relative price of non-commodity goods (Corsetti and Pesenti 2001). At the assumed parameter values, the results point to the optimality of a high degree of Home good inflation stabilization.

2.5.2 Persistent Shock

Figure 2.6 shows the response to a persistent increase in commodity prices. The shock increases commodity prices by 0.2 per cent on impact and cumulates over time, with prices ultimately peaking around 1.2 per cent above their steady state level. The initial responses of the economy to the persistent shock are qualitatively similar to the transitory shock. Consumption increases, the real and nominal exchange rates appreciate, consumer prices fall and output, hours worked and learning all contract. But, unlike the previous example, the initial response of the variables under incomplete information are muted compared to their responses under full information. Take consumption. In the full information case, consumption immediately jumps and then remains broadly constant, reflecting the desire of agents to smooth marginal utility across time. In contrast, in the incomplete information case, the increase in consumption is initially modest, but grows over time. This is because, under incomplete information, households initially attribute much of the boom to a shock to the transitory component of commodity prices, and so expect commodity prices to fall in the future (see Figure 2.4). Consequently, they react cautiously and accumulate foreign assets to provide the income necessary to smooth their expected consumption stream over time. When it turns out that the commodity price boom is long-lived, households are able to use these assets to support additional consumption and leisure.

As in the transitory commodity price boom example, monetary policy accomodates a small initial increase in PPI inflation. Under incomplete information, the initial appreciation of the nominal and real exchange rates are far smaller than under full information. However, when it turns out that the commodity price boom is highly persistent, the nominal exchange rate continues to appreciate. In contrast, in the full information case, there is a one-time appreciation in the nominal exchange rate. The smaller initial exchange rate appreciation limits the decrease in the relative price of imported goods as well as the contraction of the non-commodity sector, guarding against the possibility that the shock will prove to be short-lived.

What about learning-by-doing in this case? The right-hand panel of Figure 2.5 shows the value of the externality under incomplete and full information. In the full information case,

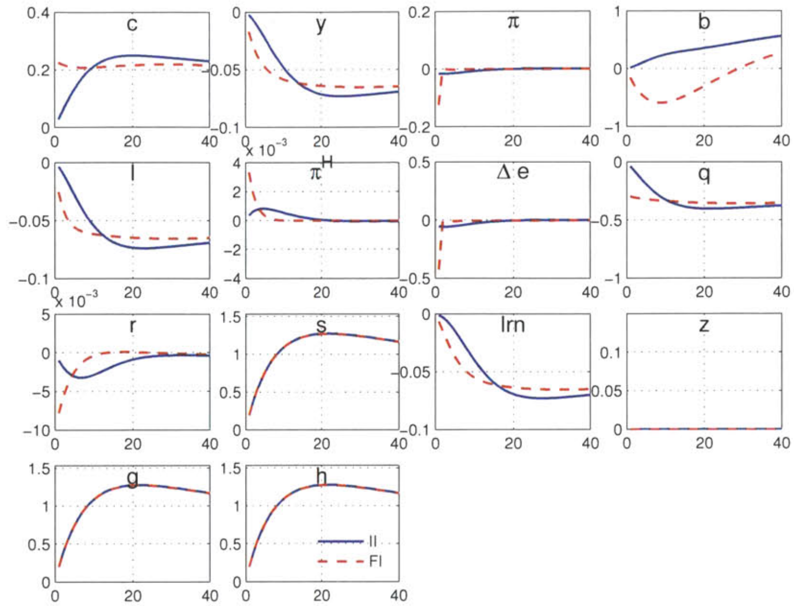

Figure 2.6: Impulse responses to a one standard deviation shock to the persistent component of commodity prices. Solid lines show responses under incomplete information. Dashed lines show responses under full information.

the shock initially causes a sharp decrease in the value of the externality. This reflects the fact that a persistent increase in commodity prices greatly reduces the shadow value of learning. The initial decrease is much less pronounced under incomplete information. This is because agents believe that the shock may be short-lived, in which case the value of additional information will be higher than it is if the boom is long-lived. As the boom persists, the value of additional information continues to decline for a while and eventually follows a path broadly similar to the full information case, albeit at a higher level.

In sum, the results in this section demonstrate that under optimal policy, the non-commodity sector will contract by less following a persistent commodity price shock under incomplete information than it does when agents have full information. This reflects the fact that agents believe that the shock may be transitory and that the value of learning is higher following a transitory shock than it is following a long-lived shock. To the extent that optimal monetary policy responds to the learning-by-doing externality, it will provide greater support to the non-commodity sector under incomplete information than it does under full information.

2.6 Simple Policy Rules

Although the policy responses described in the previous section were optimal, they may be hard to implement because they include variables that are difficult to monitor in real time, or because their weights are complex functions of the model's structural parameters. Previous work in other monetary policy settings has shown that simple policy rules can provide a close approximation to optimal policy (Schmitt-Grohe and Uribe 2007). In light of this, in this section I examine the model economy's response to commodity price shocks under simple policy rules.

2.6.1 Optimal Simple Policy Rules

As a first exercise, I calculate optimal simple policy rules. To do this I search in the grid of parameters $\left\{\psi_{\pi_H}, \psi_y, \psi_e, \psi_r\right\}$ for the rule that delivers the highest level of welfare, which I define as the optimal simple policy rule.[11] The resulting coefficients on the optimal simple rule under incomplete information are given by $\psi_r = 0.77$, $\psi_{\pi_H} = 3.00$, $\psi_e = 0.07$ and $\psi_y = 0.02$. The optimal coefficients under full information are similar, but feature marginally smaller responses to the lagged nominal interest rate, output and movements in the nominal exchange rate: $\psi_r = 0.73$, $\psi_{\pi_H} = 3.00$, $\psi_e = 0.06$ and $\psi_y = 0.01$. Under both informational assumptions, the optimal simple rule features a strong reaction to departures of PPI inflation from its target value and only a small response to deviations of output from its steady-state level or changes in the nominal exchange rate.[12]

Figures 2.7 and 2.8 compare the responses to persistent and transitory commodity price shocks under the optimal simple rule with the optimal policy responses, in both cases when agents have incomplete information. Focussing first on the transitory shock, the optimal simple rule features a larger appreciation of the real exchange rate when the shock hits.

[11] I search over the range $[1.5,3]$ for ψ_{π_H}, $[-0.2,0.2]$ for ψ_y, $[-0.2,0.2]$ for ψ_e, $[0,0.99]$ for ψ_r. With the exception of the upper bound on ψ_{π_H}, the bounds on the other policy rule parameters were never binding.

[12] Experiments with a wider grid for ψ_{π_H} revealed that larger values for this parameter could deliver further marginal increases in welfare. However, I chose to restrict the maximum value for this parameter, reflecting the fact that extremely large responsiveness to inflation is inconsistent with estimated Taylor rules for most economies.

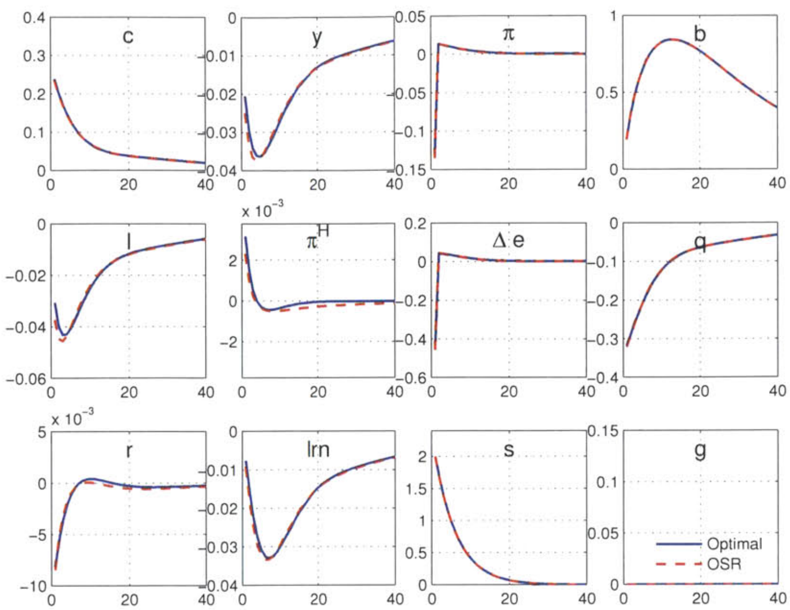

Figure 2.7: Impulse response to a one standard deviation transitory commodity price shock under optimal policy (solid line) and the optimal simple policy rule (dashed line).

This leads to a larger initial contraction in output, hours worked and learning that continues for around four quarters. Thereafter, the responses of these variables are similar to those under optimal policy. Consumption increases by marginally less under the optimal simple rule than it does under optimal policy, although its overall path is similar in the two cases. PPI inflation initially increases by less than it does under optimal policy, although it takes longer to return to its target level.

The behavior of the macroeconomic variables following a persistent shock is similar. The initial real exchange rate appreciation is smaller under the optimal simple rule than it is under optimal policy, while the decreases in output and hours worked are larger. The increase in PPI inflation is quite prolonged in this scenario although, once again, is quantitatively small. Overall, the optimal simple rule comes extremely close to replicating the dynamic responses of the economy under optimal policy.

2.6.2 Other Simple Rules

Figure 2.9 compares the economy's response to a transitory commodity price shock under PPI targeting, CPI targeting and an exchange rate peg to its response under optimal policy,

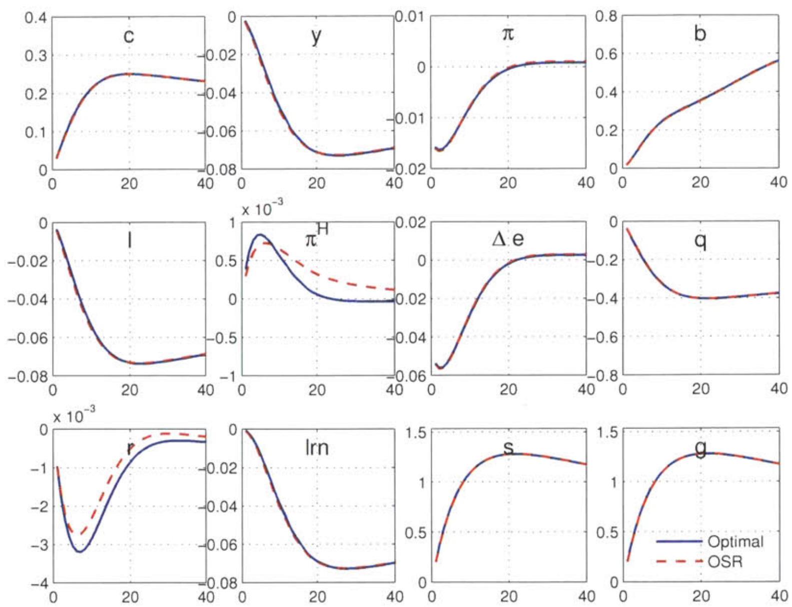

Figure 2.8: Impulse response to a one standard deviation persistent commodity price shock under optimal policy (solid line) and the optimal simple rule (dashed line).

in each case under incomplete information.[13] Compared to optimal policy, PPI targeting delivers a larger initial decrease in output, hours worked and learning and smaller initial increase in consumption. The initial appreciation of the nominal exchange rate and fall in CPI inflation are also larger under this policy rule. Like optimal policy, PPI targeting produces very small absolute movements in PPI inflation. In sum, the responses of the economy under PPI inflation targeting is broadly similar to that under optimal policy.

Under CPI inflation targeting, the central bank responds to the shock with a large decrease in nominal interest rates. This limits the initial appreciation of the nominal exchange rate, which helps to stabilize CPI inflation. Because domestic prices are sticky, the smaller appreciation of the nominal exchange rate reduces the initial change in relative prices between domestic and foreign produced tradeable goods. As a consequence, more of the additional expenditure feeds through into demand for home-produced tradeables. Output initially increases following the shock, although it decreases below its steady-state value after a few

[13]Specifically, under PPI targeting I set $\psi_{\pi_H} = 1.5$, while under CPI targeting I set $\psi_\pi = 1.5$. Under an exchange rate peg, the central bank sets interest rates consistent with $\Delta e_t = 0$. See Benigno et al. (2007) for a discussion of how policymakers can implement this policy in a manner consistent with a determinant and unique rational expectations equilibrium.

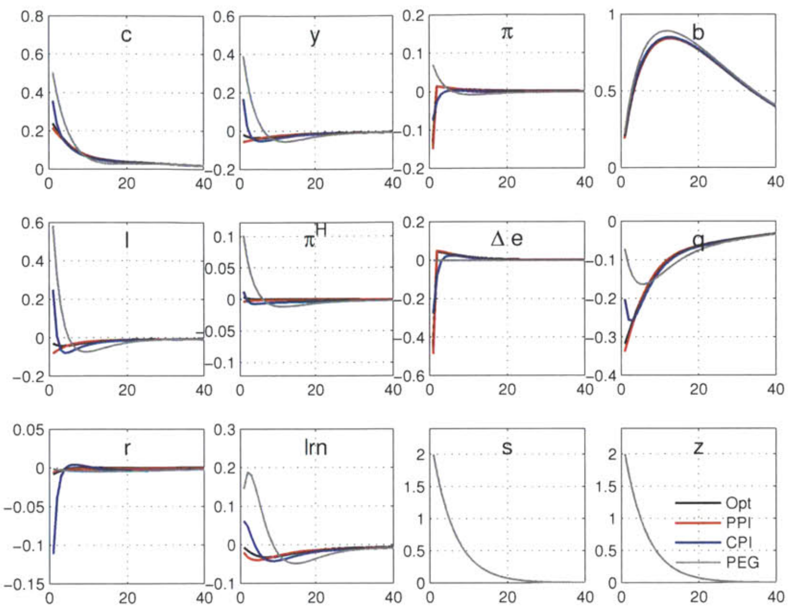

Figure 2.9: Impulse response to a one standard deviation transitory commodity price shock under optimal policy and alternative simple policy rules.

quarters. Likewise, the initial increase in consumption is larger under CPI targeting than it is under optimal policy.

Even more than CPI targeting, an exchange rate peg limits the response of relative prices to the commodity price shock. Because of this, much of the additional consumption generated by the commodity price boom flows into the domestic tradeable sector. Consequently, there is a large increase in the production of these goods and a prolonged increase in the economy's stock of learning. PPI inflation increases following the shock and, as there is no offsetting effect from a change in the nominal exchange rate, CPI inflation increases as well. The appreciation of the real exchange rate is also smaller than it is under optimal policy.

Figure 2.10 compares the response to a persistent commodity price shock with the alternative simple policy rules to the response under optimal policy, once again assuming incomplete information about the persistence of the shock. As was the case for the transitory shock, PPI inflation targeting produces broadly similar responses to optimal policy, albeit with a slightly larger initial decrease in output and a smaller initial increase in consumption. CPI inflation targeting limits the initial appreciation of the nominal exchange rate

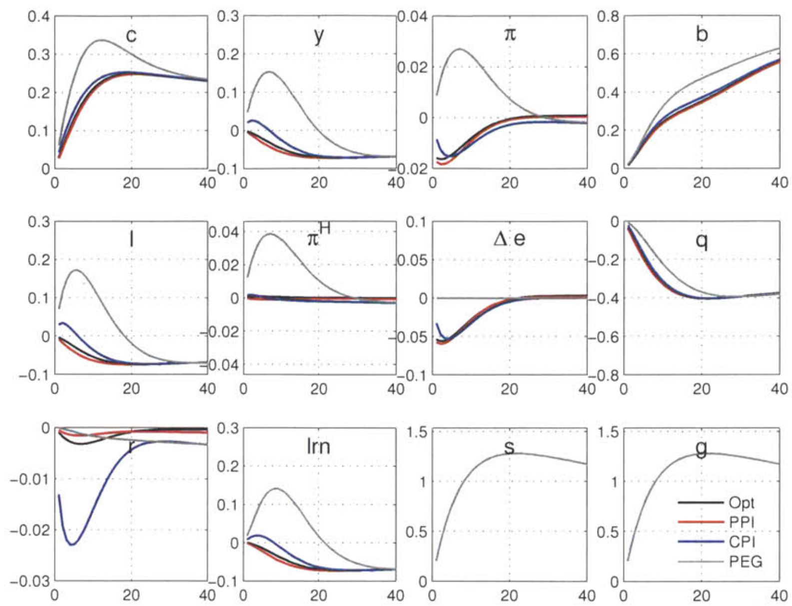

Figure 2.10: Impulse response to a one standard deviation persistent commodity price shock under optimal policy and alternative simple policy rules.

and decrease in CPI inflation, although after three quarters the response of these variables is broadly similar to the optimal policy case. CPI inflation targeting also causes an initial increase in output, reflecting the smaller increase in the relative price of domestically-produced goods. Optimal policy, PPI targeting and CPI targeting all generate relatively small changes in PPI inflation. In contrast, under an exchange rate peg the shock triggers a boom in domestic production that lasts for around five years. The additional income funds a large increase in consumption. However, it also leads to a substantial increase in both PPI and CPI inflation.

2.6.3 What are the consequences of sub-optimal policy?

In this section, I attempt to quantify the consequences of adopting simple policy rules rather than following optimal policy. As a first exercise, Table 2.2 presents the differences in welfare, measured as a percentage shift in steady state consumption, between each policy rule and optimal policy.[14] Under both incomplete and full information the optimal simple

[14]Specifically, for alternative policy rules, I calculate the value of λ such that $\lambda \equiv \left(W^{Pol} - W^{Opt}\right)/U_C\bar{C}$, where W^{Pol} is welfare under a given policy rule, W^{Opt} is welfare under optimal policy and $\bar{C}$ is the steady state consumption level.

rule and PPI targeting deliver welfare outcomes that are largely indistinguishable from optimal policy. The welfare cost of CPI targeting is also rather small, at around 1/2 of a per cent of steady state consumption. The welfare costs of an exchange rate peg are somewhat larger, at around 6 per cent of steady state consumption. The welfare costs of deviating from optimal policy are smaller under full information than under incomplete information.

Table 2.2: Welfare costs of suboptimal policy

	Cost relative to optimal policy	
Policy	Incomplete information	Full information
Optimal simple rule	0.00	0.00
PPI Targeting	0.04	0.03
CPI Targeting	0.50	0.48
PEG	7.14	5.21

To understand these results, it is helpful to consider the implications of the alternative policy rules for macroeconomic volatility. These are shown in Table 2.3. The optimal simple rule and PPI targeting both come extremely close to replicating the moments of optimal policy. The optimal simple rule produces marginally more inflation volatility while PPI inflation targeting produces more inflation, output and real exchange rate volatility, but marginally less consumption volatility. In contrast, the exchange rate peg leads to substantially more output, consumption and PPI inflation volatility, albeit with slightly less real exchange rate volatility. In terms of macroeconomic volatility, CPI inflation targeting lies somewhat between the PPI inflation targeting and exchange rate peg cases.

2.7 Some Sensitivity Analysis

In this section, I discuss the sensitivity of my results to alternative modelling choices and parameter values.

Table 2.3: Theoretical Data Moments: Incomplete Information

	Optimal Policy	Optimal Simple Rule	PPI Targeting	CPI Targeting	PEG
Standard Deviations					
y	0.73	0.73	0.75	0.76	1.03
c	2.48	2.48	2.46	2.51	2.79
π	0.16	0.17	0.19	0.11	0.14
π_H	0.01	0.01	0.01	0.05	0.20
q	4.02	4.02	4.03	4.01	3.90

2.7.1 Alternative Parameter Values

The results so far have been calculated for a particular set of calibrated parameter values. In this sub-section, I examine how the simple policy rules perform under alternative parametric assumptions.

As a first exercise, I examine alternative choices for the elasticity of substitution between domestic and foreign goods, γ. In the baseline examples I set this parameter equal to 1.3, based on estimates for Australia. Other models estimated using data from commodity-producing small open economies have found similar values, including Bache et al. (2010) for Norway and Murchison and Rennison (2006) for Canada. However, empirical estimates for this parameter vary substantially. At the upper end, Broda and Weinstein (2006) report a value between four and six, while Corsetti et al. (2008) argue that a lower value of around 0.85 helps open economy models to better match features of open economy business cycles.[15] The first two rows of Table 2.4 show the differences in welfare between PPI targeting and both CPI targeting and an exchange rate peg when γ takes the values of 0.8 and 6, holding all other parameters at their baseline levels. When $\gamma < 1$, home and foreign goods are complements in utility. When this is the case, the commodity price shock causes a smaller decrease in domestic output. This is because, when their income increases, domestic households have a strong preference to increase their consumption of both home-

[15]See Bodenstein (2010) for a review of the estimates of this parameter in the literature.

and foreign-produced goods. This scenario improves the welfare properties of PPI targeting, which dampens changes in PPI inflation, relative to the other two simple policy rules, which limit the contraction in domestic output at the cost of more volatile inflation. In contrast, when $\gamma > 1$, domestic and foreign goods are substitutes in utility. When $\gamma = 6$, the commodity price boom produces relatively little change in the domestic price level, but a large decline in production of domestically produced goods. This improves the welfare properties of CPI targeting and the exchange rate peg.

I focus next on the intertemporal elasticity of substitution, $1/\sigma$. Although the macroeconomic literature often assumes a value for this parameter of around one, much of the recent empirical literature has concluded that the elasticity of substitution lies well below one (implying a value of $\sigma > 1$), for example see Hall (1988) and Dynan (1993) among others. I consider two alternative values: 2 and 4. A higher value of this parameter raises the costs of economic fluctuations. This, in turn, increases the welfare losses associated with inferior rules, and so improves the performance of PPI targeting relative to the other two simple rules.

Turning next to price stickiness, a typical assumption in calibrated macroeconomic models is that firms reset prices roughly once a year, implying a value of θ of around 0.75, as in the baseline calibration.[16] In contrast, microeconomic studies typically find far less evidence of price rigidity, with an implied duration between price changes of between four and nine months, implying a lower value for θ (Bils and Klenow 2004; Nakamura and Steinsson 2008). More frequent price resetting reduces the costs of domestic inflation, which with staggered price setting largely result from distortions to relative prices. Consequently, lower values of θ reduce the welfare costs of CPI inflation and an exchange rate peg, both of which cause greater variation in home-produced goods prices than a policy of strict PPI targeting.

In the baseline calibration, I considered a country with an import share of consumption of 30 per cent. Increasing the openness of the economy, holding the commodity endowment fixed, improves the welfare properties of an exchange rate peg relative to PPI inflation.

[16] Empirical evidence for this estimate is provided in Blinder et al. (1998).

This is partly explained by the fact that, in a more open economy, domestic goods make up a smaller share of the consumption basket. Hence, as η increases, the welfare benefit of reducing price distortions among Home-produced goods decreases.

Table 2.4 also examines the implications for the performance of policy rules of higher values of α, the parameter that governs the importance of learning-by-doing in the production function. A higher value of α improves the welfare benefits of PPI targeting relative to CPI targeting or an exchange rate peg. Finally, Table 2.4 also considers the implications of alternative values of φ, the inverse Frisch elasticity of labor supply. A higher value of φ (that is, less elastic labor supply) improves the performance of PPI inflation targeting relative to the other simple policy rules.

Table 2.4: Welfare: PPI Targeting vs CPI and Exchange Rate Peg

		CPI	PEG
Elasticity of substitution between home and foreign goods	$\gamma = 0.8$	-0.008	-0.148
	$\gamma = 6$	-0.001	-0.006
Intertemporal elasticity of substitution	$\sigma = 2$	-0.006	-0.081
	$\sigma = 4$	-0.008	-0.102
Price Stickiness	$\theta = 0.25$	-0.001	-0.011
	$\theta = 0.50$	-0.002	-0.028
Openness	$\eta = 0.5$	-0.003	-0.014
	$\eta = 0.7$	-0.002	-0.003
Importance of learning	$\alpha = 0.1$	-0.002	-0.034
	$\alpha = 0.5$	-0.011	-0.125
Elasticity of labor supply	$\varphi = 0.5$	-0.004	-0.062
	$\varphi = 5$	-0.010	-0.123

It is also noteable that in all of the exercises in Table 2.4, PPI inflation targeting produced superior welfare outcomes to either CPI inflation targeting or an exchange rate peg. While there are parameter combinations for which these policy rules outperform PPI inflation targeting, the optimality of the latter appears robust across a broad range of common parameter estimates.

2.7.2 HP-Filtered Moments

All of the variables in the model are stationary. However, several of them, including consumption and output, experience prolonged departures from their steady state values following a persistent commodity price shock, and in simulations appear close to a random walk. Consequently, it could be argued that the welfare comparisons of alternative policy rules should be calculated using moments of detrended or filtered data. To examine this, I repeated the exercises above using HP-Filtered moments. To calculate these moments, for each policy rule I simulated 100 draws of 5,000 observations from the model, dropping the first 500 observations in each simulation. I then calculate the moments based on the average covariances across the 100 draws.

As a first exercise, I calculated the optimal policy rule for the incomplete information model using HP-filtered moments. This produced the parameters $\psi_r = 0.85$, $\psi_{\pi_H} = 3.00$, $\psi_e = 0.15$ and $\psi_y = -0.08$. That is, the optimal simple policy rule features greater persistence in the response to commodity price shocks and a larger response to changes in the nominal exchange rate, although a strong response to deviations in PPI inflation remains optimal.

As a second exercise, Table 2.5 shows HP-filtered moments and welfare costs relative to optimal policy for each of the simple policy rules. The HP-filtered moments feature much less volatility in output, consumption and the real exchange rate, all of which experience prolonged deviations from their steady state following a persistent commodity price shock. The decreases in output and consumption volatility are greatest under the optimal policy regimes and PPI targeting. The decrease in real exchange rate volatility is greatest under an exchange rate peg. However, the relative welfare rankings of the alternative policy rules at the baseline parameter values were not affected by the use of HP-filtered moments.

2.7.3 When is learning valuable?

The baseline results were calculated for a particular parameterization of the commodity price process. In this sub-section, I examine the robustness of the results regarding the impact of the learning-by-doing externality to alternative commodity price processes.

Table 2.5: HP Filtered Moments and Welfare Costs: Incomplete Information

Standard Deviations	Optimal Policy	Optimal Simple Rule	PPI Targeting	CPI Targeting	PEG
y	0.07	0.06	0.09	0.17	0.60
c	0.37	0.37	0.35	0.48	0.76
π	0.16	0.16	0.18	0.10	0.10
π_H	0.00	0.00	0.00	0.03	0.15
q	0.53	0.53	0.55	0.45	0.29
Welfare cost*	-	0.00	0.01	0.21	3.85

* Relative to optimal policy

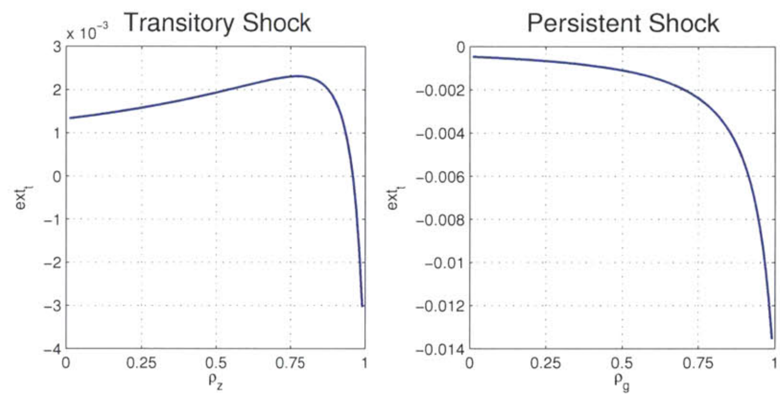

Figure 2.11: Value of externality in the first period of a commodity price boom for alternative values of ρ_z (left panel) and ρ_g (right panel).

Figure 2.11 shows the value of ext_t in the period in which a commodity price shock hits when I vary the persistence of the commodity price processes, holding all other parameters constant. I calculate these figures using the full information version of the model - the incomplete information results are a weighted average that depends on the persistence of the two shocks as well as the size of the noise shock. In the case of the transitory shock, shown in the left panel, the value of the externality increases with the persistence of the shock until ρ_z reaches a value of around 0.8, after which it falls rapidly. That is, support for the non-commodity sector is most valuable following a highly persistent, but not permanent price shock that potentially leads to a large contraction in the stock of learning but does not produce a permanent rise in commodity income. In contrast, for the persistent shock, the value of the externality is monotonically decreasing in ρ_g. A persistent shock produces a large increase in wealth that reduces the value of production externalities in the non-commodity sector. And, the larger and more persistent the price shock is, the less valuable are the non-commodity externalities.

2.7.4 Alternative Informational Setup

In the simple model presented in Section 2.2, the discovery that the commodity boom is transitory leads to a sharp increase in the optimal subsidy to support the recovery in the non-commodity sector. In contrast, in the monetary model, under incomplete information the transitory shock is followed by a gradual increase in the value of the externality (which is equivalent to the optimal subsidy in the simple model). In this section, I show that the difference in the value of the externality between the two exercises follows largely from different assumptions about information acquisition and the nature of the shock process, rather than from other features of the monetary model, including the fact that all variables in that model are stationary.

To illustrate this, I modify the monetary model so that the signal that agents receive about the persistent shock takes the form:

$$h_t = g_{t-1} \tag{2.61}$$

That is, agents receive a perfect signal about the persistent component of the commodity price shock, but they receive this signal only after one period. I also set $\rho_z = 0.1$ and $\rho_g = 0.01$, so that the commodity shock either disappears almost entirely after one period or else is close to a random walk. Finally, I assume that $\sigma_z = \sigma_g = 0.01$, so that when a shock hits agents assign equal probability to it being permanent and transitory. With these assumptions, the commodity price process in the monetary model broadly matches that presented in the simple example of Section 2.2.

Figure 2.12 shows the value of the externality following transitory and persistent commodity price shocks. Under incomplete information, the initial value of the externality falls

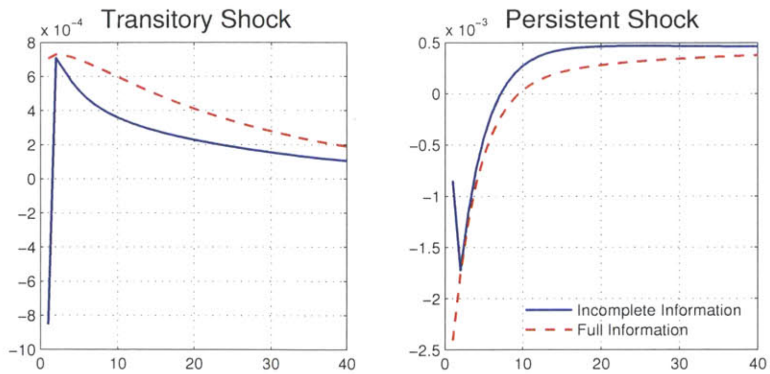

Figure 2.12: Perceived value of externality during a transitory (left panel) and persistent (right panel) commodity price boom. The solid line shows the results under incomplete information while the dashed line shows the results under full information.

below its steady state value and is the same following each shock. This reflects the fact that agents initially believe that each shock is equally likely. In the second period after the shock, the signal reveals whether the shock is transitory or persistent. In the former case, the value of the externality immediately increases to above its steady state value, before gradually declining. In contrast, if agents realize that the shock is persistent, the value of the externality sharply declines, before recovering in subsequent periods.

2.8 Conclusions

In this paper, I have examined the optimal monetary policy response to commodity price shocks when agents, including the central bank, cannot directly identify the persistence of the commodity price shock and the non-commodity sector features a learning-by-doing externality.

The key conclusion of this paper is that neither incomplete information about the persistence of commodity price shocks, nor the existince of learning-by-doing externalities in the non-commodity sector necessarily imply that monetary policy should support the non-commodity sector during a commodity price boom. Rather, optimal policy will generally allow the non-commodity sector to contract by more during a short-lived commodity price boom under incomplete information than it does when agents have full information. If policymakers come to realize that the boom is transitory, however, they will support the recovery of the non-commodity sector. And, the opposite will occur if a boom is more persistent than initially expected. These results reflect the fact that the value of a learning-by-doing externality is initially lower at the start of a persistent commodity price boom than it is at the start of a transitory boom and that, under incomplete information, the policy response will resemble a weighted average of the optimal responses to transitory and persistent booms under full information.

In a calibrated version of the model, I also demonstrate that optimal policy largely stabilizes the rate of inflation of home-produced non-tradeable goods, and allows movements in the nominal exchange rate to adjust the relative price of Home and Foreign goods. A policy of reacting strongly to changes in home-goods price inflation and a modest response to changes in the nominal exchange rate comes close to matching the welfare properties of optimal policy. In contrast, an exchange rate peg generally leads to large welfare losses unless unless home- and foreign-produced goods are very close substitutes.

A number of extensions to this paper would be interesting to pursue. For example, the model in this paper considered a particular type of externality - learning-by-doing. An extension to this work could examine whether the results are robust to other types of externalities, for example pecuniary externalities associated with collateral constraints that may bind more tightly during commodity booms. It may also be useful to consider a model in which the costs associated with changing the scale of the non-commodity sector are non-convex. An example of this would be if there were fixed costs associated with changing the scale of firms in the non-commodities sector. Finally, additional work to estimate the nature and magnitude of learning-by-doing externalities in the import-competing sectors of small open economies would help to improve the plausibility of the paper's numerical results.

2. A Derivation of the Firm's Pricing Problem

The first order condition for the firm's pricing problem is:

$$E_t \left[\sum_{k=0}^{\infty} \theta^k \mathscr{Q}_{t,t+k} Y^H_{t+k}(j) \left[\bar{P}^H_t - \frac{\zeta}{\zeta - 1} MC_{t+k|t} P^H_{t+k} \right] \right] = 0$$

where $MC_{t+k|t}$ is the marginal cost at time $t+k$ of a firm that last reset its price at time t.

Using the fact that $\mathscr{Q}_{t,t+k} = \beta^k \left(C_{t+k}/C_t \right)^{-\sigma} \left(P_t/P_{t+k} \right)$ and $Y_{t+k}(j) = Y_{t+k} \left(P^H_t(j)/P^H_t \right)^{-\zeta}$, we can re-write this as:

$$E_t \left[\sum_{k=0}^{\infty} (\beta\theta)^k \frac{C^{-\sigma}_{t+k} Y_{t+k}}{\Pi_{t,t+k}} \left(\frac{\bar{P}^H_t}{P^H_{t+k}} \right)^{-\zeta} \left[\bar{P}^H_t - \frac{\xi}{\xi - 1} MC_{t+k|t} \frac{P^H_{t+k}}{P_{t+k}} P_{t+k} \right] \right] = 0$$

where $\Pi_{t,t+k}$ is CPI inflation between period t and $t+k$.

Re-arranging this expression gives:

$$\begin{aligned} & \frac{\bar{P}^H_t}{P^H_t} E_t \left[\sum_{k=0}^{\infty} (\beta\theta)^k C^{-\sigma}_{t+k} Y_{t+k} \left(\frac{P^H_{t+k}}{P^H_t} \right)^{\zeta} \frac{P^H_t}{P^H_{t+k}} \frac{P^H_{t+k}}{P_{t+k}} P_t \right] \\ = & \frac{\zeta}{\zeta - 1} E_t \left[\sum_{k=0}^{\infty} (\beta\theta)^k C^{-\sigma}_{t+k} Y_{t+k} MC_{t+k|t} \left(\frac{P^H_{t+k}}{P^H_t} \right)^{\zeta} \frac{P^H_{t+k}}{P_{t+k}} P_t \right] \end{aligned}$$

We can replace the firm specific marginal cost with average marginal cost as follows. First, recall that:

$$MC_{t+k|t} = \frac{W_{t+k}/P^H_{t+k}}{(1-\alpha)\left(Y_{t+k}(j)/L_{t+k}(j)\right)}$$

Defining the average marginal cost at time $t+k$, MC_{t+k}, as:

$$MC_{t+k} = \frac{W_{t+k}/P^H_{t+k}}{(1-\alpha)\left(Y_{t+k}/L_{t+k}\right)}$$

and using the firm-specific demand functions and production function, we can express $MC_{t+k|t}$ as:

$$MC_{t+k|t} = MC_{t+k}\left(\frac{\bar{P}_t^H}{P_{t+k}^H}\right)^{-\frac{\alpha\zeta}{1-\alpha}}$$

The pricing equation becomes:

$$\left(\frac{\bar{P}_t^H}{P_t^H}\right)^{\frac{1-\alpha+\alpha\zeta}{1-\alpha}} E_t\left[\sum_{k=0}^{\infty}(\beta\theta)^k C_{t+k}^{-\sigma}Y_{t+k}\left(\frac{P_{t+k}^H}{P_t^H}\right)^{\zeta}\frac{P_t^H}{P_{t+k}^H}\frac{P_{t+k}^H}{P_{t+k}}P_t\right]$$
$$= \frac{\zeta}{\zeta-1}E_t\left[\sum_{k=0}^{\infty}(\beta\theta)^k C_{t+k}^{-\sigma}Y_{t+k}MC_{t+k}\left(\frac{P_{t+k}^H}{P_t^H}\right)^{\frac{\zeta}{1-\alpha}}\frac{P_{t+k}^H}{P_{t+k}}P_t\right]$$

Now:

$$E_t\left[\sum_{k=0}^{\infty}(\beta\theta)^k C_{t+k}^{-\sigma}Y_{t+k}\left(\frac{P_{t+k}^H}{P_t^H}\right)^{\zeta-1}\frac{P_{t+k}^H}{P_{t+k}}\right]$$
$$= C_t^{-\sigma}Y_t\hat{P}_t^H + \beta\theta E_t\left[\sum_{k=0}^{\infty}(\beta\theta)^k C_{t+k+1}^{-\sigma}Y_{t+k+1}\left(\frac{P_{t+k+1}^H}{P_{t+1}^H}\frac{P_{t+1}^H}{P_t^H}\right)^{\zeta-1}\hat{P}_{t+k+1}^H\right]$$

And so, defining $J_t = E_t\left[\sum_{k=0}^{\infty}(\beta\theta)^k C_{t+k}^{-\sigma}Y_{t+k}\left(\frac{P_{t+k}^H}{P_t^H}\right)^{\zeta-1}\frac{P_{t+k}^H}{P_{t+k}}\right]$, we have that:

$$J_t = C_t^{-\sigma}Y_t\hat{P}_t^H + \beta\theta E_t\left[J_{t+1}\left(\Pi_{t+1}^H\right)^{\zeta-1}\right]$$

Similarly:

$$
\begin{aligned}
E_t\left[\sum_{k=0}^{\infty}(\beta\theta)^k C_{t+k}^{-\sigma}Y_{t+k}MC_{t+k}\left(\frac{P_{t+k}^H}{P_t^H}\right)^{\frac{\zeta}{1-\alpha}}\frac{P_{t+k}^H}{P_{t+k}}\right] \quad = \quad & \frac{\zeta}{\zeta-1}C_t^{-\sigma}Y_tMC_t\hat{P}_t^H \\
& +\beta\theta E_t\left[\frac{\xi}{\xi-1}\sum_{k=0}^{\infty}(\beta\theta)^k C_{t+k+1}^{-\sigma}Y_{t+k+1}\times\right. \\
& \left. MC_{t+k+1}\left(\frac{P_{t+k+1}^H}{P_{t+1}^H}\frac{P_{t+1}^H}{P_t^H}\right)^{\frac{\zeta}{1-\alpha}}\hat{P}_{t+k+1}^H\right]
\end{aligned}
$$

Defining $F_t = E_t\left[\frac{\xi}{\xi-1}\sum_{k=0}^{\infty}(\beta\theta)^k C_{t+k}^{-\sigma}Y_{t+k}MC_{t+k}\left(\frac{P_{t+k}^H}{P_t^H}\right)^{\frac{\zeta}{1-\alpha}}\hat{P}_{t+k}^H\right]$, we have that:

$$
F_t = \frac{\xi}{\xi-1}C_t^{-\sigma}Y_tMC_t\hat{P}_t^H + \beta\theta\left[F_{t+1}\left(\Pi_{t+1}^H\right)^{\frac{\zeta}{1-\alpha}}\right]
$$

Finally, defining $\Xi_t = \bar{P}_t^H / P_t^H$, we have that:

$$
\Xi_t = \left(\frac{F_t}{J_t}\right)^{\frac{1-\alpha}{1-\alpha+\alpha\zeta}}
$$

which is Equation (2.32) in the text.

2. A.1 Forming the Standard New Keynesian Phillips Curve from the non-linear firm equations

The log-linearized firm pricing equations are (dropping expectations to ease notation):

$$
\begin{aligned}
\chi_t &= \frac{1-\alpha}{1-\alpha+\alpha\zeta}\left(f_t - j_t\right) \\
f_t &= \left(1-\beta\theta\right)\left(1+\varphi\right) l_t + \beta\theta\left(f_{t+1} + \frac{\zeta}{1-\alpha}\pi_{t+1}^H\right) \\
j_t &= \left(1-\beta\theta\right)\left(y_t - \sigma c_t + \hat{p}_t^H\right) + \beta\theta\left(j_{t+1} + \left(\zeta-1\right)\pi_{t+1}^H\right) \\
\pi_t^H &= \frac{1-\theta}{\theta}\chi_t \\
\delta_t &= \theta\delta_{t-1} + \frac{\zeta}{1-\alpha}\left[\theta\pi_t^H - \left(1-\theta\right)\chi_t\right]
\end{aligned}
$$

Note that the final equation implies that $\delta = 0$. Now, substituting the second and third equation into the first, we have that:

$$
\begin{aligned}
\chi_t &= \frac{\left(1-\beta\theta\right)\left(1-\alpha\right)}{\left(1-\alpha+\alpha\eta\right)} mc_t + \beta\theta\left(\chi_{t+1} + \pi_{t+1}^H\right) \\
\Longrightarrow \pi_t^H &= \frac{\left(1-\beta\theta\right)\left(1-\theta\right)\left(1-\alpha\right)}{\theta\left(1-\alpha+\alpha\eta\right)} mc_t + \beta\pi_{t+1}^H
\end{aligned}
$$

which is the standard New Keynesian Phillips curve.

2. B Steady State and Subsidies

2. B.1 Steady state

The non-stochastic long-run equilibrium is characterized by constant real variables and nominal variables growing at a constant rate. The equilibrium conditions reduce to:

$$
\begin{aligned}
1 &= \beta\left[\frac{1+r}{\Pi}\right] \\
1+r &= (1+r^*)\,\Pi\Delta Q \\
1 &= (1-\eta)\left(\hat{P}_H\right)^{1-\gamma}+\eta Q^{1-\gamma} \\
\Pi_H &= \Pi \\
LRN &= Y \\
MC &= (1-\tau)\frac{A_L C^{\sigma} L^{\varphi}}{\hat{P}_H}\left[(1-\alpha)\frac{Y}{L}\right]^{-1} \\
\Xi &= \frac{F}{J} \\
F &= \frac{\zeta}{\zeta-1}\frac{C^{-\sigma} Y \hat{P}_H MC}{1-\beta\theta\left(\Pi_H\right)^{\zeta}} \\
J &= \frac{C^{-\sigma} Y \hat{P}_H}{1-\beta\theta\left(\Pi_H\right)^{1-\zeta}} \\
1 &= \theta\left(\Pi_H\right)^{\zeta-1}+(1-\theta)\,\Xi^{1-\zeta} \\
\Delta &= \frac{(1-\theta)\,\Xi^{-\zeta/(1-\alpha)}}{1-\theta\Pi_H^{\zeta/(1-\alpha)}} \\
S &= 1 \\
L &= Y\Delta \\
Y &= (1-\eta)\left(\hat{P}_H\right)^{-\gamma}C+\eta\left(\frac{\bar{\hat{P}}_H}{Q}\right)^{-\gamma}C^* \\
C &= \bar{\hat{P}}_H Y+QY^X
\end{aligned}
$$

I normalize $C=1$ and set the values of the parameters $\sigma,\varphi,\gamma,\zeta,\alpha,Y^X$ and η as described in Table (2.1). I assume that $r^*=\beta^{-1}-1$ and that CPI inflation is zero in the steady state, that is $\Pi=1$. The steady state of the consumers' Euler equation implies that $r=\beta^{-1}-1$. Substituting this result into he uncovered interest rate partity condition confirms that $\Delta Q=1$. And, the equation giving the relationship between PPI and CPI inflation ensures that $\Pi_H=1$. As $\Pi=\Pi_H=\Delta Q$, the ratio of $\hat{P}_H$ to Q is constant over time. Without loss of generality, I normalize both of these variables to equal one in the steady state.

Given Y^X and C, I can then determine $Y = C - Y^X$. I then use this result to set C^* so that $Y - (1-\eta)C = \eta C^*$.

Zero steady state inflation implies that all firms choose the same price. As a consequence, $\Delta = 1$, which implies that $\Xi = 1$, $L = Y$ and $F = J$. The latter condition requires that $MC = (\zeta - 1)/\zeta$.

2. B.2 Deriving the employment subsidy

To set subsidies, I adopt the following approach. First, I consider a social planner's problem in an economy with flexible prices. This amounts to selecting a sequence of allocations $\{C_t, L_t, P_t^H, Q_t, LRN_{t+i}, B_{t+i}\}_{t=0}^{\infty}$ to maximize:

$$\max_{\{C_t, L_t, P_t^H, Q_t, LRN_{t+1}, B_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\sigma}}{1-\sigma} - A_L \frac{L_t^{1+\varphi}}{1+\varphi} \right] \tag{2.62}$$

subject to:

$$\begin{aligned}
B_t + Y^X Q_t + P_t^H LRN_t^{\alpha} L_t^{1-\alpha} &\geq C_t + \frac{B_{t+1}}{1+r} && (2.63)\\
LRN_t^{\alpha} L_t^{1-\alpha} &\geq (1-\eta) C_t \left(P_t^H\right)^{-\gamma} + \eta C^* Q_t^{\gamma} P_t^{H-\gamma} && (2.64)\\
(LRN_t)^{\chi + \alpha(1-\chi)} L_t^{(1-\chi)(1-\alpha)} &\geq LRN_{t+1} && (2.65)\\
1 &\geq (1-\eta)\left(P^H\right)^{1-\gamma} + \eta Q^{1-\gamma} && (2.66)
\end{aligned}$$

The optimal choices imply the following conditions:

$$
\begin{aligned}
C^{-\sigma} &= \lambda_{1,t} + (1-\eta)\lambda_{2,t}\left(P_t^H\right)^{-1\gamma} && (2.67)\\
A_L L_t^{\varphi} &= (1-\alpha)\left[\lambda_{1,t}P_t^H + \lambda_{2,t}\right]\left(\frac{LRN_t}{L_t}\right)^{\alpha} + (1-\alpha)(1-\chi)\lambda_{3,t}\left(\frac{LRN_t}{L_t}\right)^{\chi+\alpha(1-\chi)} && (2.68)\\
0 &= \lambda_{1,t}LRN_t^{\alpha}L_t^{1-\alpha} + \lambda_{2,t}\left[(1-\eta)\gamma C_t\left(P_t^H\right)^{-\gamma-1} + \eta\gamma C^* Q^{\gamma}\left(P_t^H\right)^{-\gamma-1}\right]\\
&\quad -(1-\eta)(1-\gamma)\lambda_{4,t}\left(P^H\right)^{-\gamma} && (2.69)\\
\lambda_{1,t} &= \beta(1+r)E_t\left\{\lambda_{1,t+1}\right\} && (2.70)\\
\lambda_{3,t} &= \beta E_t\left\{\alpha\left[\lambda_{1,t+1}P_{t+1}^H + \lambda_{2,t+1}\right]\left(\frac{LRN_{t+1}}{L_{t+1}}\right)^{\alpha-1}\right.\\
&\quad \left. + (\chi+\alpha(1-\chi))\lambda_{3,t+1}\left(\frac{LRN_{t+1}}{L_{t+1}}\right)^{-(1-\chi)(1-\alpha)}\right\} && (2.71)\\
0 &= \lambda_{1,t}Y^X - \lambda_{2,t}\eta\gamma C^* Q_t^{\gamma-1}\left(P_t^H\right)^{-\gamma} - \lambda_{4,t}\eta(1-\gamma)Q_t^{-\gamma} && (2.72)
\end{aligned}
$$

In a steady state where $Q = P^H = 1$ and $LRN = L$:

$$
\begin{aligned}
\lambda_1 + \gamma\lambda_2 &= C^{-\sigma} - (1-\eta-\gamma)\lambda_2 && (2.73)\\
0 &= \lambda_1 L + \gamma\lambda_2\left[(1-\eta)C + \eta C^*\right] - (1-\eta)(1-\gamma)\lambda_4 && (2.74)\\
\lambda_4 &= \frac{\lambda_1 Y^X - \eta\gamma\lambda_2 C^*}{\eta(1-\gamma)} && (2.75)
\end{aligned}
$$

Substituting 2.75 into 2.74 and using the steady state current account and goods market clearing conditions, it follows that:

$$
\lambda_1 + \gamma\lambda_2 = -\frac{(1-\eta)\gamma\lambda_2 C}{C^*} \qquad (2.76)
$$

Equating 2.73 and 2.76 gives an expression for λ_2:

$$
\lambda_2 = \left(C^* C^{-\sigma}\right)\frac{1}{C^*(1-\eta-\gamma) - \gamma(1-\eta)C} \qquad (2.77)
$$

Note that if $\gamma = \sigma = 1$:

$$\lambda_2 = -\frac{C^*}{CL} \tag{2.78}$$

Using the result above, we can solve for λ_1:

$$\lambda_1 = C^{-\sigma}\left[1 - \frac{(1-\eta)\,C^*}{C^*\,(1-\eta-\gamma) - \gamma(1-\eta)\,C}\right] \tag{2.79}$$

From which, it follows that:

$$\lambda_1 + \lambda_2 = C^{-\sigma}\left[1 + \frac{\eta C^*}{C^*\,(1-\eta-\gamma) - \gamma(1-\eta)\,C}\right] \tag{2.80}$$

In steady state:

$$\lambda_3 = \frac{\alpha\beta\,(\lambda_1 + \lambda_2)}{1 - \beta\,(\chi + \alpha\,(1-\chi))} \tag{2.81}$$

$$= \frac{\alpha\beta C^{-\sigma}\left[1 + \frac{\eta C^*}{C^*(1-\eta-\gamma)-\gamma(1-\eta)C}\right]}{[1 - \beta\,(\chi + \alpha\,(1-\chi))]} \tag{2.82}$$

Therefore:

$$A_L L^{\varphi} C^{\sigma} = (1-\alpha)\left[1 + \frac{\eta C^*}{C^*\,(1-\eta-\gamma) - \gamma(1-\eta)\,C}\right] \times \left[1 + \frac{(1-\chi)\,\alpha\beta}{1 - \beta\,(\chi + \alpha\,(1-\chi))}\right] \tag{2.83}$$

Once again, note that in the Cole-Obstfeld case have that:

$$A_L L^{\varphi} C^{\sigma} = (1-\alpha)\,(1 - \eta C^*/L)\left[1 + \frac{(1-\chi)\,\alpha\beta}{1 - \beta\,(\chi + \alpha\,(1-\chi))}\right] \tag{2.84}$$

And, if $\chi = 1$, and $C^* = L$, we have the familiar condition that:

$$A_L L^{\varphi} C^{\sigma} = (1-\alpha)(1-\eta) \tag{2.85}$$

Normalizing $C = 1$ and noting that $L = C - Y^X$, we can derive the value of A_L:

$$\begin{aligned} A_L \quad = \quad & \frac{(1-\alpha)}{(1-Y^X)^{\varphi}} \left[1 + \frac{\eta C^*}{C^*(1-\eta-\gamma) - \gamma(1-\eta)}\right] \times \\ & \left[1 + \frac{(1-\chi)\alpha\beta}{1-\beta(\chi+\alpha(1-\chi))}\right] \end{aligned} \tag{2.86}$$

Finally, I set τ so that:

$$1 - \tau = \left(\frac{\zeta - 1}{\zeta}\right) \frac{1-\alpha}{A_L L^{\varphi} C^{\sigma}} \tag{2.87}$$

where $A_L C^{\sigma} L^{\varphi}$ is as defined in 2.86.

2. B.3 Calculating the value of the externality

Consider the problem of a social planner that chooses labor demand to maximize firm profits, taking prices as given. The problem is:

$$\max_{\{L_t, LRN_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \mathscr{Q}_t \left[\frac{P_t^H LRN_t^{\alpha} L_t^{1-\alpha}}{P_t} - \frac{W_t}{P_t} L_t\right] \tag{2.88}$$

subject to:

$$LRN_t^{\chi+\alpha(1-\chi)} L_t^{(1-\chi)(1-\alpha)} \geq LRN_{t+1} \tag{2.89}$$

The optimality conditions are:

$$\frac{W_t}{P_t} = (1-\alpha)\hat{P}_t^H\left(\frac{LRN_t}{L_t}\right)^{\alpha}\left(1+(1-\chi)\mu_t\left(\frac{LRN_t}{L_t}\right)^{\chi(1-\alpha)}\right) \tag{2.90}$$

$$\mu_t = \beta E_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\sigma}\frac{\hat{P}_{t+1}^H}{\hat{P}_t^H}\left[\alpha\left(\frac{LRN_{t+1}}{L_{t+1}}\right)^{\alpha-1}\right.\right. \tag{2.91}$$

$$\left.\left.+(\chi+\alpha(1-\chi))\mu_{t+1}\left(\frac{LRN_{t+1}}{L_{t+1}}\right)^{(\alpha-1)(1-\chi)}\right]\right]$$

Log-linearization of this final term produces Equation 2.60 in the main text.

Chapter 3

Stochastic Terms of Trade Volatility in Small Open Economies[1]

3.1 Introduction

The terms of trade of many commodity-producing small open economies are subject to large shocks that can be an important source of economic fluctuations. Alongside times of high volatility, however, these economies also experience periods in which their terms of trade are comparatively stable. The effect of shocks to the level of the terms of trade has been widely studied. But little is known about the impact of changes in the volatility of terms of trade shocks. We study the macroeconomic effects of these shocks and quantify their importance as a source of business cycle fluctuations.

Figure 3.1 shows the growth rate of the terms of trade for a selection of commodity-producing small open economies. At various times, each economy has experienced an increase or decrease in its terms of trade of more than 10 per cent in a quarter, while fluctuations of five per cent or more are common. The existence of these large shocks has motivated a substantial literature examining the impact of changes in the level of the terms of trade on these economies (Mendoza (1995), Kose and Riezman (2001), Broda (2004)).

[1]This paper is co-authored with my classmate Patricia Gómez-González.

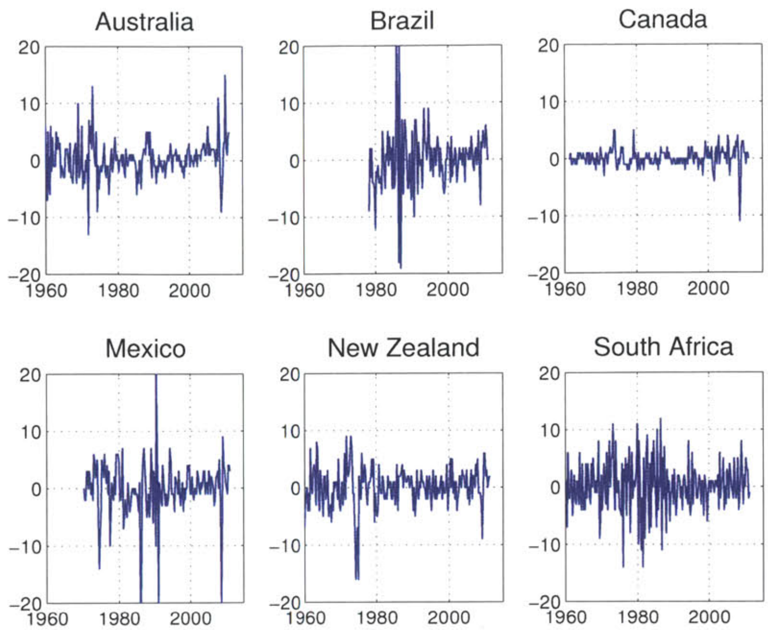

Figure 3.1: Terms of Trade Growth: Selected Countries

Source: See Appendix 3. A.

In addition to these large shocks, Figure 3.1 also suggests that the terms of trade volatility that small open economies experience varies over time. Each economy in the figure has undergone episodes of extremely high terms of trade volatility, including during the 1970s for Australia and New Zealand, 1980s for Brazil, Mexico and South Africa, and 2000s for Canada. But these economies have also experienced sustained periods in which their terms of trade were comparatively stable, such as the late 1990s for Australia, New Zealand and Mexico. This paper examines the economic relevance of these changes in terms of trade volatility.

To address this question, we first estimate the empirical process of the terms of trade for the six economies featured in Figure 3.1. We use the estimated time series of terms of trade volatility produced in this exercise to identify the effect of volatility shocks on output, consumption, investment, the current account and prices in a vector autoregression (VAR). We then set up and augment a small open economy real business cycle model to incorporate stochastic terms of trade volatility. We test whether this model can replicate the empirical responses produced by the VAR and use it to explore the theoretical causes and sectoral impacts of these responses. Finally, we compute variance decompositions to quantify the importance of terms of trade volatility shocks as a source of macroeconomic fluctuations.

Our empirical results suggest that an increase in terms of trade volatility depresses domestic demand and leads to an improvement in the current account, leaving the response of aggregate output ambiguous. Our model successfully replicates these patterns. It also suggests increased terms of trade volatility causes a shift in the composition of output from non-treadeables to tradeables and a substitution in factor inputs from capital to labor.

The effects of terms of trade volatility shocks are generally small. But, interacted with shocks to the level of the terms of trade, variance decompositions suggest that they have an economically meaningful impact. For a typical small open economy we find that shocks to volatility account for around one quarter of the total impact of terms of trade shocks on the standard deviations of output, consumption and investment.

3.1.1 Literature review

Our paper is related to several strands of literature. Most directly, it complements other papers that have studied the real effects of uncertainty and time-varying volatility on macroeconomic aggregates. Examples here include **?** who explores the short-run fluctuations of output, employment and productivity growth after a shock to macroeconomic uncertainty, Justiniano and Primiceri (2008) who shed light on the sources of changes in US macroeconomic volatility in the postwar period using a structural model and **?** who use a financial accelerator model to study the effects of an idyosincratic risk shock to entrepeneurs' productivity. Closely related to our paper is Fernandez-Villaverde et al. (2011), who examine shocks to the volatility of sovereign debt interest rates and Fernandez-Villaverde et al. (2012), who study how changes in uncertainty about future fiscal policy affects aggregate economic activity. Our main contribution to this literature is empirical. We document time-varying volatility in a variable, the terms of trade, that has not previously been studied and explore the effects of changes in this volatility.

The paper also builds on the literature examining the macroeconomic effects of terms of trade shocks. Many papers in this literature have examined terms of trade shocks using calibrated business cycle models. These typically conclude that terms of trade shocks are an important source of small open economy business cycles. For example, Mendoza (1995) concludes that terms of trade shocks account for around half of the fluctuations in GDP in developing countries and slightly less in developed economies. In a model calibrated to match features of a standard developing economy, Kose and Riezman (2001) find that terms of trade shocks account for 45 per cent of output volatility and 86 per cent of investment volatility. And, in a model calibrated for Canada, Macklem (1993) finds that a 10 per cent temporary deterioration in the terms of trade - a large but not unprecedented shock for the economies in Figure (3.1) - reduces real GDP by almost 10 per cent and investment by almost 20 per cent.

Other papers in this literature have adopted a more reduced form approach and examined the effects of terms of trade shocks in vector autoregression models. These papers typically find smaller effects of terms of trade shocks than those that rely on structural business cycle models. For example, using a panel VAR covering 75 developing countries, Broda (2004) concludes that a 10 per cent permanent deterioration in the terms of trade reduces the level of GDP by around one per cent, and that the terms of trade shocks explain between 10 - 30 per cent of GDP growth. Similarly, Collier and Goderis (2012) find that a 10 per cent rise in commodity prices increases the level of GDP by around one percentage point after two years for a typical developing country. Our contribution to this literature is to illustrate another channel through which the terms of trade can have macroeconomic effects. In particular, we show how changes in terms of trade volatility can have an impact even if the level of the terms of trade remains constant.

Alongside the literature examining the dynamic effect of shocks to the level of the terms of trade, another empirical literature documents a negative link between terms of trade volatility and long-run economic growth. Using a panel of 35 developed and developing economies over the period 1870 to 1939, Blattman et al. (2007) conclude that, for commodity producers, a one standard deviation increase in terms of trade volatility (in their sample, from 8 per cent to 13 per cent per year) causes a 0.4 percentage point reduction in annual per capita GDP growth. In related work, Williamson (2008) attributes much of the gap in economic performance in the early 19th century between economies in Western Europe and those in Eastern Europe, the Middle East and East Asia to fact that the latter groups experienced more terms of trade volatility. Focussing on more contemporary patterns, Bleaney and Greenaway (2001) estimate a cross-country panel regression using data from 14 sub-Saharan African countries over the period 1980-95 and also conclude that terms of trade volatility, measured as the residuals from a GARCH model of the terms of trade, reduces GDP growth.

As well as long-run growth, papers in this literature have examined links between terms of trade volatility and the volatility of other macroeconomic variables. For example, using a panel of countries **?** show that times of high terms of trade volatility tends to be correlated with times of more volatile GDP growth, while Andrews and Rees (2009) also establish a link with consumption and inflation volatility. The theoretical grounding for these results was established in Mendoza (1997). Using a stochastic growth model, he demonstrates that terms of trade volatility can affect growth through its effects on households' incentives to save, but that the direction of the effect depends on the degree of households' risk aversion. He also shows that, regardless of its impact on growth, an increase in terms of trade volatility reduces welfare. Our paper complements this literature by illustrating the links between terms of trade volatility and macroeconomic outcomes in a fully-specified macroeconomic model and by tracing out the dynamic effects of changes in terms of trade volatility on output, external accounts and prices.

3.2 Estimating the Law of Motion for the Terms of Trade

In this section, we estimate the empirical process for the terms of trade for six small open economies: Australia, Brazil, Canada, Mexico, New Zealand and South Africa. We selected these countries based on two criteria. First, we focussed on commodity-producing small open economies whose terms of trade are both volatile and plausibly exogenous to domestic economic developments. Second, we required countries to have reasonably long time series data for the terms of trade and other macroeconomic variables.

Table 3.1: Summary Statistics: Merchandise Exports, 2010

	Share of world merchandise exports	Export Composition		
		Food Items and Agricultural Raw Materials	Fuels, Ores and Metals	Manufactured Goods
Australia	1.4	13.1	69.2	12.8
Brazil	1.3	34.8	29.5	35.8
Canada	2.5	13.5	35.6	47.8
Mexico	2.0	6.3	18.7	74.5
New Zealand	0.2	63.3	10.1	22.9
South Africa	0.6	10.5	47.4	39.2

Source: UNCTAD Handbook of Statistics 2011

Our claim that the terms of trade for these countries are exogenous may be controversial. To support our contention, Table 3.1 provides descriptive statistics about the size and export composition of each economy. The six economies each account for a small share of world GDP and merchandise trade. This suggests that economic developments within these countries are unlikely to have a substantial effect on world economic activity. Moreover, the exports of these countries are geared towards agriculture, fuels and mining - that is, commodities - with these goods accounting for more than 50 per cent of merchandise export values for each country, except Mexico.[2] Commodities tend to be less differentiated, and more substitutable, than manufactured goods and commodity producers generally have less pricing power on world markets.[3] Further evidence to support our contention comes from the numerous studies that have used statistical techniques to examine the exogeneity of the terms of trade for small open economies. For example, using Granger causality tests, Mendoza (1995) and Broda (2004) conclude that the terms of trade is exogenous for a large sample of small open economies, including Brazil, Mexico and Canada.

[2]Even for Mexico, petroleum is the largest single export good at the three digit SITC 3 level, accounting for almost 12 per cent of total exports in 2010. Moreover, commodities accounted for the bulk of Mexico's exports in the early part of our sample, before the expansion of manufacturing exports that accompanied Mexico's trade liberalization in 1986 and entry into NAFTA in 1994 (?).

[3]While this is not strictly true for all commodity producers, such as large oil producers for example, it seems reasonable for the countries in our sample.

3.2.1 Estimation

For each country, we specify that the terms of trade, q_t, follow an AR(1) process described by:

$$q_t = \rho_q q_{t-1} + e^{\sigma_{q,t}} u_{q,t} \tag{3.1}$$

where $u_{q,t}$, are normally distributed shocks with mean zero and unit variance. The log of the standard deviation of the terms of trade shocks, $\sigma_{q,t}$, varies over time, according to an AR(1) process:

$$\sigma_{q,t} = (1-\rho_\sigma)\sigma_q + \rho_\sigma \sigma_{q,t-1} + \eta_q u_{\sigma,t} \tag{3.2}$$

where $u_{\sigma,t}$, are normally distributed shocks with mean zero and unit variance. To emphasize, innovations to $u_{q,t}$ alter the level of the terms of trade, while innovations to $u_{\sigma,t}$ alter the magnitude of shocks to the terms of trade, with no direct effect on its level. The parameter σ_q is the log of the mean standard deviation of terms of trade shocks, while η_q is the standard deviation of shocks to the volatility of the terms of trade. The parameter ρ_σ controls the persistance of terms of trade volatility shocks. Throughout, we assume that $u_{q,t}$ and $u_{\sigma,t}$ are independent of each other.

Equations (3.1)-(3.2) represent a standard stochastic volatility model. Inference in these models is challenging because of the presence of two innovations, to the level of the terms of trade and to its volatility, that enter the model in a nonlinear manner. To overcome this issue, we follow Fernandez-Villaverde et al. (2011) and use a sequential Markov Chain Monte Carlo filter, also known as a particle filter, that allows us to evaluate the likelihood of the model using simulation methods. We estimate the model using a Bayesian approach that combines prior information with information that can be extracted from the data.

Denote the vector of parameters to be estimated as as $\Psi = \{\rho_q, \rho_\sigma, \sigma_q, \eta_q\}$ and the log of the prior probability of observing a given vector of parameters $\mathscr{L}(\Psi)$. The function $\mathscr{L}(\Psi)$

summarizes what is known about the parameters prior to estimation. The log-likelihood of observing the dataset $q^T \equiv \{q_1, \ldots, q_T\}$ for a given parameter vector is denoted $\mathscr{L}\left(q^T|\Psi\right)$. The likelihood of the data given the parameters factorizes to:

$$\exp\left(\mathscr{L}\left(q^T|\Psi\right)\right) = p\left(q^T|\Psi\right) = \prod_{t=1}^{T} p\left(q_t|q^{t-1};\Psi\right)$$

The final term in this expression expands as follows:

$$\prod_{t=1}^{T} p\left(q_t|q^{t-1};\Psi\right) = \prod_{t=1}^{T} \int p\left(q_t|q_{t-1},\sigma_{q,t};\Psi\right) p\left(\sigma_{q,t}|q^{t-1};\Psi\right) d\sigma_{q,t} \tag{3.3}$$

Computing this expression is difficult because the sequence of conditional densities $\left\{p\left(\sigma_{q,t}|q^{t-1};\Psi\right)\right\}_{t=1}^{T}$ has no analytical characterization. A standard procedure, which we follow, is to substitute the density $p\left(\sigma_{q,t}|q^{t-1};\Psi\right)$ with an empirical draw from it. To obtain these draws, we follow Algorithm 1, which we borrow from Fernandez-Villaverde et al. (2011).

Algorithm 2. *Particle Filter*

Step 0, Initialization:

Sample N particles, $\left\{\sigma^i_{q,0|0}\right\}_{i=1}^{N}$ *from the initial distribution* $p\left(\sigma_{q,0}|\Psi\right)$.

Step 1, Prediction:

Sample N one-step ahead forecasted particles $\left\{\sigma^i_{q,t|t-1}\right\}_{i=1}^{N}$ *using* $\left\{\sigma^i_{q,t-1|t-1}\right\}_{i=1}^{N}$, *the law of motion for the states (Equation 3.2) and the distribution of shocks* $u_{\sigma_q,t}$.

Step 2, Filtering:

Assign each draw $\left(\sigma^i_{q,t|t-1}\right)$ *the weight* ω^i_t, *where:*

$$\omega^i_t = \frac{p\left(q_t|q_{t-1},\sigma^i_{q,t|t-1};\Psi\right)}{\sum_{i=1}^{N} p\left(q_t|q_{t-1},\sigma^i_{q,t|t-1};\Psi\right)} \tag{3.4}$$

Step 3, Resampling:

Generate a new set of particles by sampling N times with replacement from $\left\{\sigma^i_{q,t|t-1}\right\}_{i=1}^N$ *using the probabilities* $\left\{\omega^i_t\right\}_{i=1}^N$*. Call the draw* $\left\{\sigma^i_{q,t|t}\right\}_{i=1}^N$*. In effect, this step builds the draws* $\left\{\sigma^i_{q,t|t}\right\}_{i=1}^N$ *recursively from* $\left\{\sigma^i_{q,t|t-1}\right\}_{i=1}^N$ *using the information on* q_t*. If* $t < T$*, set* $t = t+1$ *and return to step 1. Otherwise stop.*

Using the law of motion for the terms of trade in Equation 3.1, we can evaluate $p\left(q_t|q_{t-1},\sigma^i_{q,t|t-1};\Psi\right)$ for any $\sigma^i_{q,t|t-1}$. Moreover, from the Law of Large numbers we know that:

$$\int p\left(q_t|q_{t-1},\sigma_{q,t};\Psi\right) p\left(\sigma_{q,t}|q^{t-1};\Psi\right) d\sigma_{q,t} \approx \frac{1}{N}\sum_{i=1}^{N} p\left(q_t|q_{t-1},\sigma^i_{q,t|t-1};\Psi\right)$$

Algorithm 1 provides a sequences of $\left\{\sigma^i_{q,t|t-1}\right\}_{i=1}^N$ for all t. Consequently, the algorithm gives us the information needed to evaluate Equation (3.3).

To calculate the posterior distribution of the parameters, we repeat this procedure 25,000 times. At each iteration, we update our parameter draw using a random walk Metropolis-Hastings procedure, scaling the proposal density to induce an acceptance ratio of around 25 per cent. We discard the initial 5,000 draws and conduct our posterior inference on the remaining draws. For each evaluation of the likelihood we use 2,000 particles.

Other methods of modelling time-varying volatility processes, including Markov switching models and GARCH models, also exist. Although these methods have advantages in other contexts, we do not believe that they provide a satisfactory description of terms of trade volatility. For example, a GARCH model does not sharply distinguish between innovations to the terms of trade and its volatility. High levels of volatility are triggered only by large innovations to the terms of trade. In contrast, our methodology allows changes in the volatility of the terms of trade to occur independently of innovations to the level of the terms of trade. A Markov switching model would require us to restrict the number of potential realizations of terms of trade volatility in a way that seems inconsistent with the patterns in Figure 3.1.

3.2.1.1 Data

The terms of trade for each country are defined as the ratio of the export price deflator to the import price deflator and sourced from national statistical agencies.[4] As we wish to estimate changes over time in the variance of the terms of trade, we require our data to be stationary. The stationarity of commodity prices (which drive the terms of trade for the countries in our sample) is a source of contention. Previous studies by Cashin et al. (2000), Powell (1991) and Lee et al. (2006), among others, have concluded that commodity prices are stationary. Others, including Kim et al. (2003), Newbold et al. (2005) and Maslyuk and Smyth (2008) have found that they are not.

In light of the disagreement in the literature, we adopt a compromise approach and detrend our data using a bandpass filter that excludes cycles of longer than 30 years. This preserves all but the lowest frequency movements in the terms of trade for each country while ensuring that the data is stationary.[5]

3.2.1.2 Priors

Table 3.2 reports our priors for the parameters of the terms of trade process. For the persistence parameters, ρ_q and ρ_σ, we impose a Beta prior with mean 0.9 and standard deviation of 0.1. The shape of this prior restricts the value of these parameters to lie between 0 and 1, consistent with economic theory. For the log of the mean standard deviation of terms of trade shocks, σ_q, we impose a Normal prior. For each country, we set the mean of this prior equal to the OLS estimate of this parameter calculated assuming an AR(1) process for the terms of trade without stochastic volatility. For the standard deviation of terms of trade volatility shocks, η_q, we use a Truncated Normal prior thus ensuring that this parameter is positive. We experimented with alternative priors and found that these had very little impact on our results.

[4] Appendix 3. A includes a full list of data sources and descriptions.

[5] We also estimated the models with HP filtered data (see Appendix 3. B for the results). The choice of detrending method has some effect on the estimated persistence of terms of trade shocks, but relatively little impact on the estimated magnitude of shocks to the terms of trade or its volatility.

Table 3.2: Prior Distribution of Parameters

Parameter	ρ_q	σ_q	ρ_σ	η_q
Prior	$\beta(0.9,0.1)$	$\mathcal{N}(\hat{\sigma}_{OLS},0.4)$	$\beta(0.9,0.1)$	$\mathcal{N}^+(0.5,0.3)$

Notes: 1) β, $\mathcal{N}$ and $\mathcal{N}^+$ stand for Beta, Normal and truncated normal distributions.

3.2.1.3 Posterior estimates

Table 3.3 reports the posterior medians of the parameter estimates and associated confidence bands. The first row shows the posterior estimates of ρ_q, the persistence of the terms of trade processes. This parameter lies above 0.9 for all countries except for South Africa, indicating that shocks to the terms of trade for these countries tend to be highly persistent. The parameter estimates for σ_q reveal substantial differences in the average size of shocks to the terms of trade between countries. Converting the parameters into standard deviations, the results suggest that the magnitude of the average terms of trade shock varies from around 1.2 per cent for Canada to 4.0 per cent for Brazil.[6] The estimates for ρ_σ indicate that shocks to the volatility of the terms of trade are highly persistent for Australia, Brazil, New Zealand and South Africa, but somewhat less so for Canada and Mexico. The final row of the table confirms that the magnitude of shocks to the volatility of the terms of trade differs between countries. Of the countries in our sample, Mexico has tended to experience the largest volatility shocks, while New Zealand and South Africa have experienced the smallest. To put these numbers in context, a one standard deviation shock to $u_{\sigma,t}$ increases the standard deviation of terms of trade shocks in Mexico from 3.2 per cent to 4.7 per cent and in South Africa from 3.7 per cent to 4.1 per cent.

To give a clearer insight into what our results imply for the time-varying terms of trade volatility, Figure (3.2) shows the model's estimates of the evolution of the standard deviations of terms of trade shocks for each country. The average level of these series is higher for Brazil and Mexico than for the other countries in the sample, reflecting the fact that these countries have typically experienced larger terms of trade shocks. The changes in the level are also greatest for Mexico, as that country has experienced the largest shocks to

[6] Recall, that the standard deviation of shocks to the terms of trade is equal to $\exp(\sigma_q)$.

Table 3.3: Posterior Medians: Bandpass Filtered

	Australia	Brazil	Canada	Mexico	New Zealand	South Africa
ρ_q	0.93	0.96	0.90	0.92	0.96	0.81
	[0.88, 0.98]	[0.90, 0.99]	[0.85, 0.95]	[0.87, 0.97]	[0.89, 0.97]	[0.73, 0.86]
σ_q	-3.65	-3.22	-4.40	-3.43	-3.48	-3.30
	[-4.16, -3.03]	[-3.74, -2.57]	[-4.74, -4.04]	[-3.78, -3.03]	[-3.93, -2.90]	[-3.81, -2.76]
ρ_σ	0.94	0.92	0.85	0.82	0.93	0.97
	[0.83, 0.99]	[0.79, 1.00]	[0.57, 0.98]	[0.64, 0.94]	[0.84, 1.00]	[0.85, 1.00]
η_q	0.21	0.21	0.26	0.38	0.13	0.10
	[0.12, 0.33]	[0.11, 0.36]	[0.13, 0.48]	[0.25, 0.54]	[0.07, 0.23]	[0.05, 0.23]

Notes: 95 per cent set in brackets.

the volatility of its terms of trade. In contrast, the standard deviation of shocks to Canada's terms of trade have typically been small and stable over time, at least compared to those experienced by other commodity exporters. The experiences of Australia, New Zealand and South Africa lie somewhere in between those of Canada and the Latin American countries. They have typically experienced relatively large terms of trade shocks, with an average standard deviation of around 3 per cent. They have also experienced periods of hightened volatility, although not to the same extent as Brazil and Mexico.

The patterns of volatility suggested by Figure (3.2) broadly conform to our understanding of macroeconomic developments over the sample. For example, the average magnitudes of terms of trade shocks increased in most countries during the mid 1970s, mid 1980s and late 2000s, while the 1990s was generally a period of low terms of trade volatility.

In sum, our results indicate that the volatility of shocks to the terms of trade for small open economy commodity producers varies over time. Historically, the variation has been largest for Latin American countries such as Brazil and Mexico, where the standard deviation of terms of trade shocks has at times increased from an average level of around three per cent to over 10 per cent. But countries like Australia, New Zealand and South Africa have also

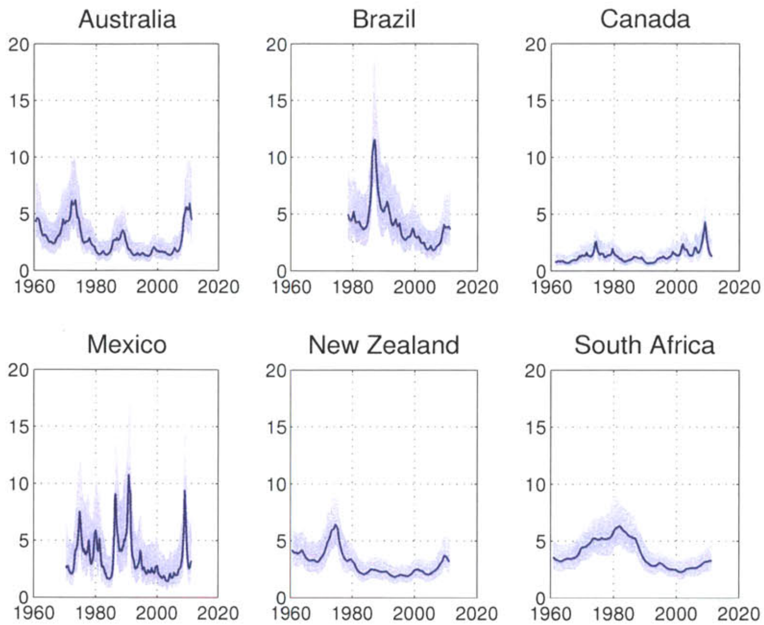

Figure 3.2: Standard Deviation of Terms of Trade Shocks

Notes: Darker lines are median values of $\exp(\tilde{\sigma}_{qt|T})$ where $\tilde{\sigma}_{q,t|T}$ are the estimated volatility conditional on all information calculated using the particle smoother. The shaded area represents 5 and 95 per cent confidence bands.

experienced shocks that have increased the standard deviation of their terms of trade shocks from around three per cent to around six per cent.

3.3 The Impact of Volatility Shocks: Empirics

3.3.1 Panel VAR

This section models the responses of real GDP (y), consumption (c), investment (i), the current account (ca) and the GDP deflator (p) to the terms of trade volatility shocks identified in the previous section. Because each economy in our sample has experienced only a relatively small number of sizeable volatility shocks, we pool the data for all six countries. The model can be expressed as a panel vector autoregression (VAR):

$$Y_{it} = v + A(L)Y_{it} + B(L)X_{it} + u_{it} \tag{3.5}$$

where $Y'_{it} = (y, c, i, ca, p)$ is a vector of stationary endogenous variables, v is a vector of constants, $X'_{it} = (q_{it}, \sigma_{q,it})$ is a vector containing the level of the terms of trade as well as its volatility, $u'_{it} = (u^{y}_{it}, u^{c}_{it}, u^{i}_{it}, u^{ca}_{it}, u^{p}_{it})$ is an error vector, $A(L)$ and $B(L)$ are matrix polynomials in the lag operator and $var(u_{it}) = \Omega$.[7] Note that although the variables in Y_{it} respond to the terms of trade and its volatility, we do not include terms of trade variables as endogenous variables in the VAR. This is consistent with our assumption in Section 3.2 that the terms of trade is exogenous with respect to domestic economic developments for the small open economies in our sample.

The empirical model described in Equation (3.5) can be thought of as a simplified reduced form version of a DSGE model with stochastic volatility, like the one described in Section 3.4 below. Of course, the empirical model cannot fully capture the nonlinear relationships implied by a theoretical model. However, we argue that it nonetheless provides a meaningful indication of the relationships between exogenous terms of trade volatility shocks and macroeconomic variables that appear in the data and serves as a useful benchmark against which to compare the results of our theoretical model. In Appendix (3. C) we provide evidence to support this contention.

3.3.2 Results

To illustrate the consequences of a terms of trade volatility shock, we report the dynamic effects of an innovation to σ_q of 0.22, roughly equivalent to the average of η_q across the countries estimated in Section 3.2. After the initial shock, we allow σ_q to decay by 10 per cent per quarter, again broadly consistent with the estimates for Section 3.2.

Figure 3.3 shows the dynamic response of y_{it}, c_{it}, i_{it}, ca_{it} and p_{it} to an increase in the volatility of terms of trade shocks. Solid lines are the point estimates of the impulse response functions and dashed lines represent one standard deviation (16th and 84th percentile) of the empirical distribution of responses.

[7] In the results below, we include four lags of the endogenous variables and the contemporaneous value and one lag of the terms of trade variables. Experiments with alternative lage structures produced broadly similar results.

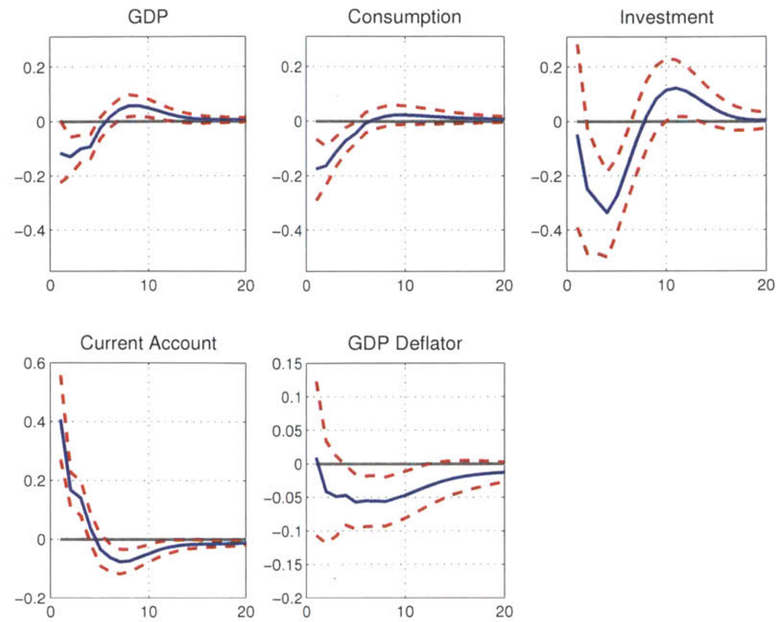

Figure 3.3: Impulse Responses to a Terms of Trade Volatility Shock: Pooled Sample

The volatility shock reduces both consumption and investment on impact. Although the investment response is not initially significant, it becomes so in later quarters, eventually troughing two quarters after the shock. Consumption and investment both return to their original levels six - eight quarters after the shock, although there is some evidence of a small boom in domestic demand in later quarters. Aggregate output also decreases in the periods after the shock. The size of its response is smaller than the responses of consumption and investment, however, suggesting an offsetting response of net exports. This shows up in the current account-GDP ratio, which increases in the quarter in which the shock hits. It remains above its trend level for two subsequent quarters, before declining as domestic demand recovers. There is also a persistent decrease in the GDP deflator. As we have held the terms of trade constant in this exercise, this implies a fall in the price of non-tradeable goods relative to tradeables.

A possible criticism of our empirical approach is that pooling data conceals cross-country heterogeneity in the impact of volatility shocks. In particular, one might wonder whether economies in which households and firms are less able to hedge the risks associated with terms of trade volatility are more responsive to these types of shocks. As a first step to answering this question, Figures 3.4-3.5show responses to volatility shocks when we separate

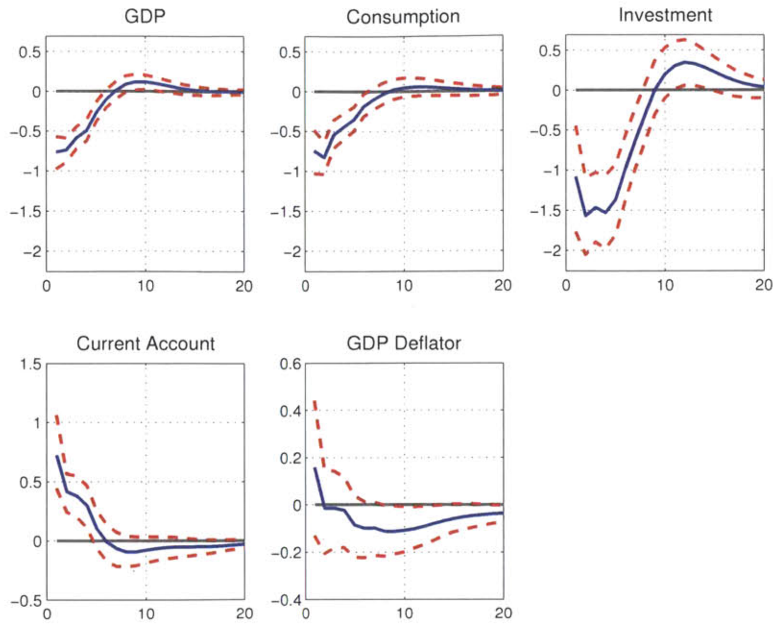

Figure 3.4: Terms of Trade Volatility Shock: Emerging Economies

our sample into emerging economies - Brazil, Mexico and South Africa - and developed economies - Australia, Canada and New Zealand.

As Figure 3.4 shows, the effect of volatility shocks on output and its components is considerably greater when we estimate the model on a sample including only the emerging economies. The responses of these variables is roughly four times as large and the responses of investment and GDP are now significant from the quarter of impact. It also takes an additional quarter or two for these variables to return to trend after the shock. The current account-GDP ratio continues to increase following the shock, consistent with the decrease in domestic demand exceeding the decrease in GDP. The point estimate of the response of the GDP deflator is qualitatively similar to the pooled response, although the response is only marginally statistically significant.

Figure 3.5 reveals a somewhat different response to the shock among the developed economies. The responses of investment and consumption for these economies are not significantly different from zero, while the point estimates suggest that investment may actually increase immediately following the shock. The initial response of aggregate output is also positive, albeit only significant in the period in which the shock hits. In contrast, the response of the GDP deflator remains negative and significant, and is quantitatively

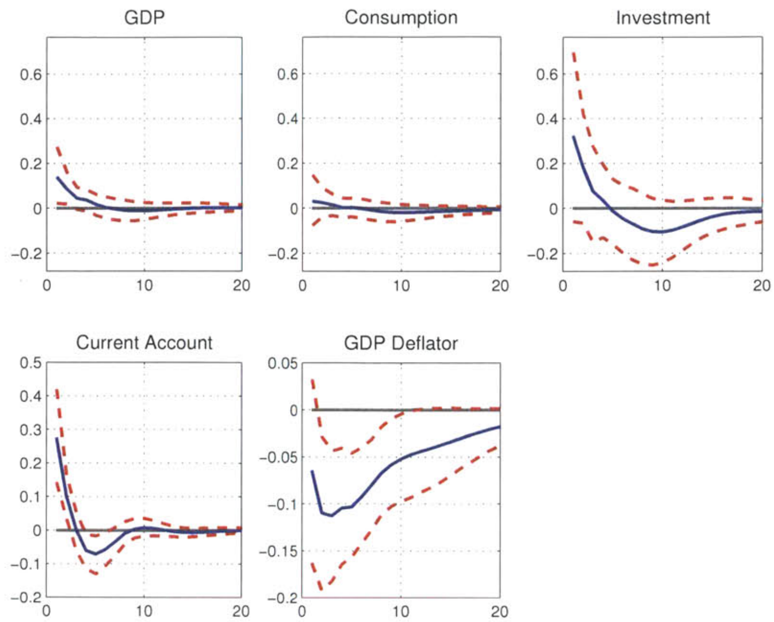

Figure 3.5: Terms of Trade Volatility Shock: Developed Economies

larger relative to the response of the other variables than in Figure 3.3. These economies also continue to experience a substantial improvement in their current account following the volatility shock.

In sum, the empirical results suggest that an increase in terms of trade volatility triggers a slump in domestic demand that is partly offset by an increase in net exports, leading to a relatively small impact on aggregate output. These shocks also cause a decrease in the domestic price level which, given that we have held the level of the terms of trade constant, suggests a relative decrease in the price of domestic non-tradeable goods. There is some evidence that the response of output and its components is larger in developing economies, while the price response is larger in developed economies. However, given the relatively small number of countries in our sample, we are reluctant to place too much weight on this conclusion. In the following section, we show that a standard international real business cycle model, augmented with stochastic terms of trade volatility, is broadly able to replicate these responses. We then use the model to shed light on the theoretical causes and sectoral implications of these responses.

3.4 The Impact of Volatility Shocks: Theory

In this section, we embed stochastic terms of trade volatility in an otherwise standard small open economy real business cycle model with incomplete markets. In the model, households choose consumption, saving and labor supply to maximize expected lifetime utility. Households consume three goods - non-tradables and home and foreign-produced tradeable goods - and can invest in three assets - a one-period risk-free bond traded in international capital markets and physical capital in the two domestic sectors. On the production side, firms produce output using capital, which is industry-specific, and labor, which is mobile across sectors, and aim to maximize profits. As well as terms of trade shocks, we also include productivity shocks in the model. These shocks help the model to match key features of the data, but play little role in the analysis.

3.4.1 Households

The economy features a representative household that maximizes its expected lifetime utility given by:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{c_t^{1-\sigma}}{1-\sigma} - \frac{l_t^{1+\zeta}}{1+\zeta} \right) \tag{3.6}$$

where c is consumption of goods and l is hours of work.

Consumption is a composite of tradable and non-tradable goods:

$$c_t \equiv \left[\omega_T^{\frac{1}{\vartheta}} \left(c_t^T\right)^{\frac{\vartheta-1}{\vartheta}} + (1-\omega_T)^{\frac{1}{\vartheta}} \left(c_t^{NT}\right)^{\frac{\vartheta-1}{\vartheta}} \right]^{\frac{\vartheta}{\vartheta-1}} \tag{3.7}$$

where the elasticity of substitution between tradables and non-tradables is ϑ, the weight of tradables in the consumption basket is ω_T and c^{NT} is the household's consumption of non-tradables. c^T is the household's consumption of tradeable goods, which is itself a composite of home- and foreign-produced tradable goods:

$$c_t^T \equiv \left[\omega_H^{\frac{1}{\eta}}\left(c_t^H\right)^{\frac{\eta-1}{\eta}} + (1-\omega_H)^{\frac{1}{\eta}}\left(c_t^F\right)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} \tag{3.8}$$

where the elasticity of substitution between the two tradeable goods is η, the weight of home-produced goods is ω_H, c^H is the household's consumption of home-produced tradable goods and c^F is the household's consumption of foreign-produced tradable goods.

To smooth consumption across time, households have access to three assets: a one-period risk-free bond, denominated in units of the foreign-produced tradeable good, and physical capital in the non-traded and home-tradeable sectors. Reflecting the fact that the domestic economy is small relative to the rest of the world, we assume that the interest rate faced by the economy on its debt issuance, r, is exogenous. Households face a portfolio adjustment cost from holding foreign debt at a different level than its steady-state level, d. This ensures that the economy's foreign debt level is stationary and prevents precautionary savings diverging to infinity.[8]

Household capital holdings, k^{NT} and k^H, are sector specific. We assume that the price of all capital goods are denominated in units of the foreign-produced tradeable good.

We take the price of the foreign good as numeraire and set it equal to one. With this normalization, the household's budget constraint is given by:

$$\begin{aligned} & c_t^F + e^{q_t} c_t^H + p^{NT} c_t^{NT} + + i_t^{NT} + i_t^H + d_t(1+r) \\ & \leq w_t l_t + r_t^{NT} k_t^{NT} + r_t^H k_t^H + d_{t+1} - \frac{\psi}{2}\left(d_{t+1} - d\right)^2 \end{aligned} \tag{3.9}$$

where e^{q_t} is the price of home-produced tradeable goods in terms of foreign-produced tradeable goods - the terms of trade - and p^{NT} is the relative price of non-tradeable goods. The terms of trade is exogenous in the model, while p^{NT} is determined endogenously.

[8]Portfolio adjustment costs are one of the several ad-hoc methods commonly used to close small open economy models. Others include a debt-elasic interest rate premium or a time preference rate that varies with aggregate consumption. Schmitt-Grohe and Uribe (2003) show that all of these methods deliver almost identical dynamics at business-cycle frequencies. Another way of attaining a stationary asset distribution is to assume that the rate of time preference is smaller than the interest rate as in **?**.

i^{NT} and i^{H} are investment in the non-tradeable and home-tradeable sectors. And w, r^{NT} and r^{H} are the wage rate and return on capital in the non-tradeable and home-tradeable sectors. Note that as labor is mobile between the two sectors, firms in each sector pay the same wage. The final term on the right hand side of the equation represents the portfolio adjustment costs.

The capital stock of each sector evolves according to:

$$k_{t+1}^{j} = i_{t}^{j} + (1-\delta)k_{t}^{j} - \frac{\phi}{2}\frac{\left(k_{t+1}^{j} - k_{t}^{j}\right)^{2}}{k_{t}^{j}} \tag{3.10}$$

for $j = \{NT, H\}$. The parameter δ represents the depreciation rate of capital, while the final term represents adjustment costs associated with changing the size of the capital stock. We include these to prevent the model from delivering unrealistically large movements in investment.

Household optimization implies that the demand for for home- and foreign-produced tradeable goods is:

$$c_{t}^{H} = \omega_{H}\left(\frac{e^{q_{t}}}{p_{t}^{T}}\right)^{-\eta} c_{t}^{T}; c_{t}^{F} = (1-\omega_{H})\left(\frac{1}{p_{t}^{T}}\right)^{-\eta} c_{t}^{T} \tag{3.11}$$

where $p_{t}^{T} \equiv \left[\omega_{H}\left(e^{q_{t}}\right)^{1-\eta} + (1-\omega_{H})\right]^{\frac{1}{1-\eta}}$ is the traded goods price index. The demand for tradeable and non-tradeable goods is:

$$c_{t}^{T} = \omega_{T}\left(\frac{p_{t}^{T}}{p_{t}}\right)^{-\vartheta} c_{t};\ c_{t}^{NT} = (1-\omega_{T})\left(\frac{p_{t}^{NT}}{p_{t}}\right)^{-\vartheta} c_{t} \tag{3.12}$$

where $p_{t} \equiv \left[\omega_{T}\left(p_{t}^{T}\right)^{1-\vartheta} + (1-\omega_{T})\left(p_{t}^{NT}\right)^{1-\vartheta}\right]^{\frac{1}{1-\vartheta}}$ is the consumer price index (CPI).

Using the optimal household decisions over different good types, we can re-write the household's budget constraint as:

$$p_t c_t + i_t^{NT} + i_t^H + d_t(1+r) \tag{3.13}$$
$$\leq w_t l_t + r_t^{NT} k_t^{NT} + r_t^H k_t^H + d_{t+1} - \frac{\psi}{2}(d_{t+1} - d)^2$$

The household's optimal choice over consumption, labor supply and asset holdings implies the following intra- and intertemporal conditions:

$$c_t^{\sigma} l_t^{\zeta} = \frac{w_t}{p_t} \tag{3.14}$$

$$1 = \beta \mathbb{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \frac{(1 + r + \psi(d_{t+1} - d))}{\pi_{t+1}} \right] \tag{3.15}$$

$$\left[1 + \phi \left(\frac{k_{t+1}^H}{k_t^H} - 1 \right) \right] = \beta \mathbb{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \frac{1}{\pi_{t+1}} \left(1 - \delta + r_{t+1}^H + \frac{\phi}{2} \left[\left(\frac{k_{t+2}^H}{k_{t+1}^H} \right)^2 - 1 \right] \right) \right] \tag{3.16}$$

$$\left[1 + \phi \left(\frac{k_{t+1}^{NT}}{k_t^{NT}} - 1 \right) \right] = \beta \mathbb{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \frac{1}{\pi_{t+1}} \left(1 - \delta + r_{t+1}^{NT} + \frac{\phi}{2} \left[\left(\frac{k_{t+2}^{NT}}{k_{t+1}^{NT}} \right)^2 - 1 \right] \right) \right] \tag{3.17}$$

where $\pi_{t+1} = p_{t+1}/p_t$ is the rate of CPI inflation between period t and period $t+1$.

3.4.2 Firms

The home-tradeable and non-tradeable sectors both feature perfectly competitive firms that maximize profits given by:

$$\pi_t^H = e^{q_t} y_t^H - w_t l_t^H - r_t^H k_t^H \tag{3.18}$$

$$\pi_t^{NT} = p_t^{NT} y_t^{NT} - w_t l_t^{NT} - r_t^{NT} k_t^{NT} \tag{3.19}$$

Firms in each sector produce output using a Cobb-Douglas production function:

$$y_t^H = e^{a_t}(k_t^H)^{\alpha}(l_t^H)^{1-\alpha} \tag{3.20}$$

$$y_t^{NT} = e^{a_t}(k_t^{NT})^{\alpha}(l_t^{NT})^{1-\alpha} \tag{3.21}$$

where e^{a_t} is a productivity shifter that is common to both sectors.

Profit maximization by firms implies that factor prices are given by:

$$w_t = (1-\alpha)\,e^{q_t}\frac{y_t^H}{l_t^H} \tag{3.22}$$

$$w_t = (1-\alpha)\,p_t^{NT}\frac{y_t^{NT}}{l_t^{NT}} \tag{3.23}$$

$$r_t^H = \alpha e^{q_t}\frac{y_t^H}{k_t^H} \tag{3.24}$$

$$r_T^{NT} = \alpha p_T^{NT}\frac{y_t^{NT}}{k_t^{NT}} \tag{3.25}$$

3.4.3 Shock Processes

The model features three exogenous processes. First, productivity evolves according to an AR(1) process:

$$a_t = \rho_a a_{t-1} + \varepsilon_t^a \tag{3.26}$$

Second, the terms of trade and its variance evolve according to the processes described in the empirical section and repeated here for convenience:

$$q_t = \rho_q q_{t-1} + e^{\sigma_{q,t}} u_{q,t} \tag{3.27}$$

$$\sigma_{q,t} = (1-\rho_\sigma)\sigma_q + \rho_\sigma \sigma_{q,t-1} + \eta_q u_{\sigma,t} \tag{3.28}$$

The interpretation of the parameters is also given in the empirical section.

3.4.4 Equilibrium definition

The competitive equilibrium is given by an allocation $\{c_t, l_t, l_t^H, l_t^{NT}, k_t^H, k_t^{NT}, i_t^H, i_t^{NT}, d_t\}_{t=0}^{\infty}$ and goods and factor prices $\{w_t, r_t^H, r_t^{NT}, p_t^{NT}, p_t\}_{t=0}^{\infty}$ where (i) consumers' satisfy their optimality conditions (3.14- 3.17) and equations for evolution of capital in both sectors (3.10); (ii) firms' zero-profit conditions given in Equations (3.22) to (3.25) hold; (iii) productivity and the terms of trade, a_t, q_t and $\sigma_{q,t}$, follow the exogenously given processess in Equations (3.26)-(3.28) and (iv) factor and goods markets clear.

Regarding factor market clearing, labor is fully mobile across sectors. Hence, its market clearing condition is given by:

$$l_t = l_t^H + l_t^{NT} \tag{3.29}$$

Capital is by assumption sector-specific. Market clearing is defined similarly: capital supplied by households has to equal capital demanded by firms.

Goods market clearing implies that all production in the tradeable and non-tradeable sectors is consumed:

$$y_t^{NT} = c_t^{NT} \tag{3.30}$$

$$y_t^H = c_t^H + c_t^{H*} \tag{3.31}$$

where c_t^{H*} is consumption of the home-produced tradable good by foreigners. The latter can be expressed in terms of home-variables only. To do so, we use the equation for the evolution of foreign debt $d_{t+1} - d_t = rd_t - nx_t$ where nx_t denotes net exports. Net exports equals nominal exports minus nominal imports (where the latter include capital goods):

$$e^{q_t} c_t^{H^*} = nx_t + c_t^F + i_t^{NT} + i_t^H \tag{3.32}$$

Substituting in the tradable goods market clearing condition and replacing nx_t using the debt evolution equation we obtain the condition for home-produced goods market clearing in terms of home variables only:

$$e^{q_t} y_t^H = (1+r)d_t - d_{t+1} + p_t^T c_t^T + i_t^{NT} + i_t^H + \frac{\psi}{2}(d_{t+1} - d)^2 \tag{3.33}$$

3.4.5 Model solution and calibration

We solve the model using perturbation methods, taking a third-order approximation of the policy functions of the agents and the law of motion of the exogenous variables around the model's steady state. As Fernandez-Villaverde et al. (2011) discuss, in models with stochastic volatility it is necessary to take a third-order approximation of the model to capture the effects of volatility shocks independent of the other innovations in the model.[9]

We fix the value of a number of parameters in all calibrations using values generally found in the literature (Table 3.5). For the households, we set the discount rate, β, equal to 0.99, the inverse of the elasticity of substitution, σ, and the inverse of the Frisch elasticity, ζ, both equal to 2, which are all consistent with values commonly used in the literature. We base the values of ϑ and η on available estimates for the elasticity of substitution between traded and non-traded goods. For the elasticity of substitution between tradeables and non-tradeables, ϑ, we use the estimate by Mendoza (1995), calculated for a sample of industrialized countries, and set that elasticity equal to 0.74.[10] For the elasticity of substitution between home and foreign tradeables, η, we use the estimate of Corsetti et al. (2008) and select a value of 0.85. We set the share of traded goods in the households' consumption basket, ω_T, equal to 0.5, consistent with the estimates of Stockman and Tesar (1995). We also set the share of home goods in the tradable goods basket equal to 0.5. We set the capital share of income, α, equal to $1/3$ for both sectors. We follow Fernandez-Villaverde

[9]Specifically, a first order approximation eliminates all of the effects of volatility shocks as certainty equivalence holds. A second-order approximation captures the effects of volatility shocks only through their interaction with shocks to the level of the terms of trade. It is only in a third-order (or higher) approximation that stochastic volatility shocks enter as independent arguments in the policy functions.

[10]For a sample of developed and developing countries, Stockman and Tesar (1995) estimate a lower elasticity of 0.44. We examine the sensitivity of our results to these elasticities in our robustness checks.

Table 3.5: Parameter Values

Parameter	Description	Value	Comments
β	Discount factor	0.99	Standard value.
σ	Inverse of elasticity of subtitution	2	Standard value.
ζ	Inverse of Frisch elasticity	2	Standard value.
ω_T	Share of tradables in consumption basket	0.5	As in Stockman and Tesar (1995).
ω_H	Share of home goods in tradable consumption basket	0.5	
ϑ	Elasticity of substitution between tradable and non-tradable goods	0.74	As in Mendoza (1995).
η	Elasticity of substitution between home and foreign tradable goods	0.85	As in Corsetti et al. (2008).
α	Capital share of income	$\frac{1}{3}$	Standard value.
ρ_a	Persistence of shock to productivity	0.95	As in Fernandez-Villaverde et al. (2011)

et al. (2011) in setting the persistence of productivity shocks, ρ_a, equal to 0.95. This choice has little effect on our results as we merely use this shock to calibrate the model. We set ψ, the portfolio adjustment cost of foreign debt, equal to 10^{-3} for all the countries. This small value ensures that the foreign debt level is stationary, without significantly affecting the dynamic properties of the model (Schmitt-Grohe and Uribe (2003); Fernandez-Villaverde et al. (2011)).

Conditional on these choices, we pick the last three parameters to match moments of the ergodic distribution generated by simulating the model to moments of the data. The parameters are: (i) σ_a, the standard deviation of productivity shocks; (ii) ϕ, the adjustment cost of capital; and (iii) d, the parameter that controls the average stock of foreign debt. The moments of the data that we match are: (i) output volatility; (ii) the volatility of investment relative to output; and (iii) the ratio of net exports to output. Because the moments are affected by a nonlinear combination of parameters, we choose the parameters to minimize

Table 3.6: Calibrated Parameters

Parameter	Australia	Brazil	Canada	Mexico	New Zealand	South Africa
d	-2.00	1.77	3.96	-3.13	2.71	7.55
ϕ	32.87	54.11	11.32	42.07	19.91	11.78
σ_a	1.00×10^{-2}	1.16×10^{-2}	1.06×10^{-2}	1.86×10^{-2}	1.00×10^{-2}	1.05×10^{-2}

the sum of the quadratic distance between the model moments and the moments from the data.[11] Table 3.6 shows the calibrated parameter values.

3.5 Results

In this section, we analyze the quantitative implications of our model. First, we compare the moments of the model to those of the data. Second, we construct impulse response functions to illustrate how an innovation to the volatility of terms of trade shocks affects the other variables in the model. Third, we use variance decompositions to quantify the contribution of terms of trade volatility shocks to the variance of the key macroeconomic variables in the model.

3.5.1 Moments

Table 3.7 compares the moments of the model to those of the data. The model matches the three calibrated moments - the variance of output, the relative variance of investment and the level of net exports relative to GDP - successfully for all countries. The model comes reasonably close to matching the correlation of investment and output and consumption and output. It is less successful at replicating some of the other moments of the data. In particular, the volatility of consumption relative to output is generally lower in the model than it is in the data. This is a common finding in small open economy real business cycle models and is generally resolved by assuming the absence of wealth effects on labor supply

[11] Specifically, for each economy, we simulate a sample of 200 observations and calculate moments based on these observations. We then repeat this procedure 200 times and calculate the mean of each moment across the 200 draws.

Table 3.7: Empirical Second Moments

	Australia		Brazil		Canada	
	Data	Model	Data	Model	Data	Model
σ_y	1.35	1.36	1.52	1.51	1.38	1.37
σ_c/σ_y	0.81	0.53	1.06	0.51	0.83	0.47
σ_i/σ_y	2.97	2.97	3.67	3.69	2.97	2.98
$\rho_{c,y}$	0.49	0.57	0.62	0.87	0.80	0.81
$\rho_{i,y}$	0.67	0.66	0.87	0.46	0.71	0.83
$\rho_{nx,y}$	-0.22	0.63	-0.26	0.59	0.18	0.62
nx/y	-0.94	-0.94	0.30	0.30	1.41	1.41

	Mexico		New Zealand		South Africa	
	Data	Model	Data	Model	Data	Model
σ_y	2.43	2.43	1.39	1.39	1.60	1.61
σ_c/σ_y	1.16	0.51	1.04	0.53	1.29	0.69
σ_i/σ_y	1.82	1.83	4.46	4.47	3.70	3.70
$\rho_{c,y}$	0.97	0.71	0.76	0.48	0.76	0.00
$\rho_{i,y}$	0.83	0.71	0.83	0.67	0.64	0.86
$\rho_{nx,y}$	-0.07	0.62	-0.18	0.63	-0.52	0.64
nx/y	-1.53	-1.54	0.80	0.80	2.53	2.53

(as in ?) or adding trend growth shocks (as in Aguiar and Gopinath (2007)) to the model.[12] The model also produces too much correlation between net exports and income. The latter result might be due, in part, to our assumption that all investment goods are imported. We examine this issue in the robustness checks section below.

[12]In our model, one can induce greater consumption volatility by increasing the magnitude of the portfolio adjustment cost, ψ. This makes it more costly for households to borrow and lend, which reduces consumption smoothing. However, we found that an extremely high value of ψ - generally in the order of 0.1 - was required for the volatility of consumption in the model to match that found in the data. And, with ψ at such a high level, the effect of portfolio adjustment costs on the model's dynamics cease to be negligible.

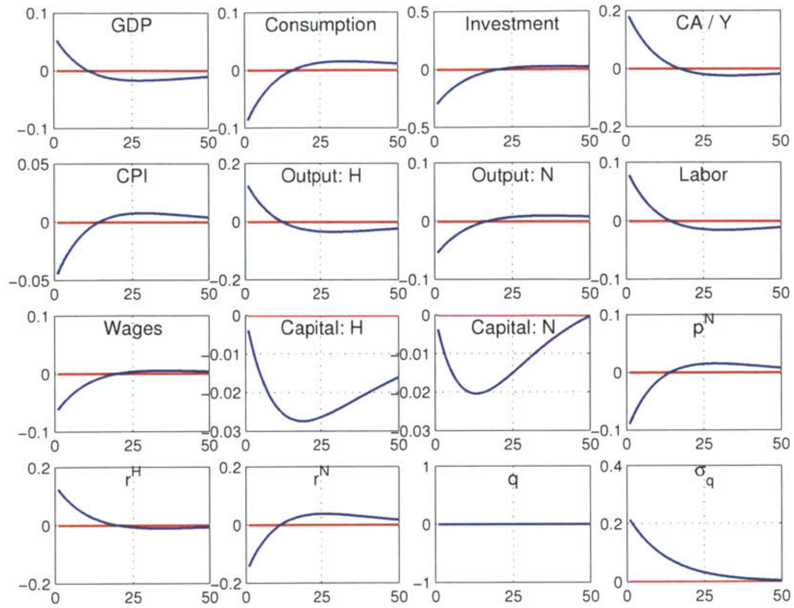

Figure 3.6: Terms of Trade Volatility Shock: Brazil

3.5.2 Impulse response functions

We now turn to the dynamic response of the economies to a shock to terms of trade volatility. We focus and describe in detail the Brazilian case and consign the results for the other economies to Appendix 3. D. Figure (3.6) shows the response of the Brazilian economy to a one standard deviation shock to the volatility of the terms of trade, that is a shock to u_{σ_t}. Note that the shock has no effect on q_t, the level of the terms of trade. Despite this, the shock induces a decrease in consumption of almost 0.1 per cent on impact and a larger decrease in investment, of around 0.3 per cent. The current account-to-GDP ratio also increases by around 0.2 per cent following the shock, while the price level decreases. Because the terms of trade does not change following the shock, this implies a decrease in the GDP deflator. In sum, the model qualitatively matches the responses to a terms of trade volatility shock identified in the panel VAR. The magnitudes of the responses are slightly smaller than the VAR, although in the same ballpark as the ones found in Fernandez-Villaverde et al. (2011) and Fernandez-Villaverde et al. (2012).

The theoretical intuition for these responses comes from the household's optimality conditions. Consider first the household's Euler equation:

$$1 = \beta \mathbb{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\sigma} \frac{(1 + r + \psi(d_{t+1} - d))}{\pi_{t+1}} \right]$$

The shock to terms of trade does not affect the expected level of consumption directly. But it does make agents more uncertain about their future income flows, which increases the expected marginal utility of future consumption, $E_t \left\{ c_{t+1}^{-\sigma} \right\}$. As households prefer to smooth marginal utility across time, they reduce consumption today. This increases the marginal utility of consumption today and, by freeing up more resources for future consumption, reduces the expected future marginal utility of consumption. Moreover, the reduction in consumer demand lowers prices today relative to future prices, which increases $E_t \left\{ \pi_{t+1} \right\}$. Because the terms of trade is exogenous, the adjustment in prices must occur entirely through changes in the relative price of non-tradeable goods.

The decrease in current consumption and the reduction in the price level also affects labor supply through the household's intratemporal optimality decision:

$$c_t^{\sigma} l_t^{\zeta} = \frac{w_t}{p_t}$$

The decrease in current consumption increases the marginal utility of consumption. This brings about an increase in the labor supply because the utility cost of working is now lower. The increase in the labor supply causes the equilibrium wage rate to drop.

On the production side of the economy, the volatility shock brings about a change in the sectoral composition of output away from non-tradeables towards tradeables. The decrease in non-tradeable production follows directly from the decrease in consumption. Although consumption of tradeable goods also decreases, this effect is overwhelmed by an increase in exports, which helps agents to accumulate foreign assets.

The volatility shock also affects factor utilization. The increase in labor supplied by households lowers the capital-labor ratio, which reduces the real wage relative to the return on capital. This change in factor prices encourages firms to adopt more labor-intensive production methods. Combined with the changes in output, this reduces firms' demand for

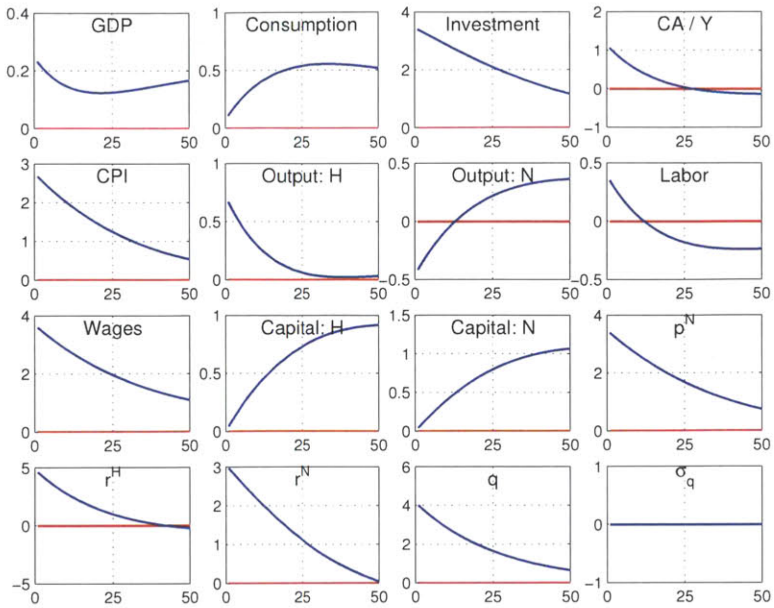

Figure 3.7: Terms of Trade Level Shock: Brazil

capital. The consequent reduction in investment also helps households to reduce domestic absorption, further contributing to their accumulation of foreign assets.

The qualitative behavior of variables in response to the terms of trade volatility shocks is similar for all the countries we study, although the magnitude of the responses are somewhat more modest.[13]

It is also instructive to compare the response of the economy to a volatility shock to its response to a one standard deviation shock to the level of the terms of trade, shown in Figure (3.7). The shock brings about (i) a prolonged increase in consumption; (ii) an investment boom; (iii) an improvement in the current account; (iv) an increase in home-produced goods output and a temporary decrease in non-tradable goods production; and (v) an increase in hours worked and real wages. These results are consistent with the findings in Mendoza (1995).

Although both shocks lead to a boom in domestic tradeable output and a reduction in foreign debt, the terms of trade level shock is far more favorable to domestic agents. This

[13]Appendix 3. D reports the IRFs for Mexico, Australia, New Zealand, South Africa and Canada to a shock to the volatility of terms of trade.

shock encourages firms to invest in order to increase production to take advantage of temporarily high relative goods prices. The resulting increase in the capital-labor ratio drives the increase in real wages, which triggers the expansion in labor supply.

3.5.3 Variance decompositions

In this section we study the contribution to aggregate fluctuations of each of the three shocks in our model. Because of our nonlinear approximation to the policy function, it is not possible to divide total variance among the shocks as in a linear model. Therefore, in this exercise, we set the realizations of one or two of the shocks to zero and measure the volatility of the economy when we simulate the economy with the remaining shocks.

We study four macro-aggregates: output, consumption, investment and net exports and explore four scenarios: (i) all shocks; (ii) terms of trade level shocks only; (iii) terms of trade volatility shocks only; and (iv) terms of trade level and volatility shocks jointly.[14]

Table 3.8 reports the variance decompositions for all six countries. For each of them, productivity shocks are the main contributor for output fluctuations, while shocks to the level and volatility of the terms of trade are key drivers of investment and net exports fluctuations.

By themselves, volatility shocks account for only a very small portion of the standard deviation of output and consumption for all of the countries in our sample. The impact of these shocks for investment and net exports is somewhat greater - with these shocks alone the standard deviation of Brazilian investment is estimated to be 0.36 and net exports 0.22 - although still modest given the high variance of these series.

However, interacted with shocks to the level of the terms of trade, volatility shocks are estimated to have a meaningful impact of macroeconomic outcomes. For example, with only shocks to the level of the terms of trade, the standard deviation of Brazilian investment is estimated to be 4.27 per cent. With shocks to the volatility as well as the level of the terms

[14]Note that in each decomposition agents in the model believe that that the shocks are distributed according to the law of motion specified in the previous section. Consequently, they will respond to volatility shocks even when the realization of shocks to the level of the terms of trade is always zero.

Table 3.8: Variance Decompositions

	All Shocks	Terms of Trade Level Only	Terms of Trade Volatility Only	TOT Level and TOT Volatility
	(i)	(ii)	(iii)	(iv)
		Australia		
Y	1.36	0.39	0.01	0.52
C	0.71	0.30	0.02	0.40
I	4.04	2.88	0.10	3.79
NX	1.90	1.35	0.05	1.81
		Brazil		
Y	1.51	0.32	0.07	0.45
C	0.78	0.23	0.11	0.32
I	5.57	4.27	0.36	5.46
NX	2.19	1.50	0.22	2.08
		Canada		
Y	1.37	0.22	0.00	0.27
C	0.64	0.21	0.00	0.26
I	4.08	2.35	0.02	2.90
NX	0.75	0.54	0.01	0.67
		Mexico		
Y	2.43	0.51	0.01	0.71
C	1.24	0.40	0.02	0.57
I	4.42	2.81	0.10	3.86
NX	2.86	1.89	0.05	2.66
		New Zealand		
Y	1.39	0.47	0.01	0.57
C	0.74	0.37	0.02	0.45
I	6.20	4.92	0.14	5.89
NX	1.71	1.37	0.05	1.65
		South Africa		
Y	1.61	0.78	0.00	0.89
C	1.10	0.80	0.00	0.92
I	5.96	4.56	0.05	5.24
NX	2.71	2.34	0.02	2.69

of trade, the standard deviation of investment is estimated to be 5.46 per cent - 28 per cent greater.

Indeed, for countries like Brazil, Mexico and Australia, the volatility of the key macroeconomic variables is between 20 and 30 per cent higher when there are both volatility and level shocks than it is when terms of trade level shocks operate alone. That is, for these countries between a fifth and a third of the effect of the terms of trade on macroeconomic volatility comes through changes in the volatility of terms of trade shocks. For countries like Canada, New Zealand and South Africa, the contribution of volatility to the overall effects of terms of trade shocks is smaller. However, even for these countries, our results suggest that between 10 and 20 per cent of the impact of the terms of trade on the key macroeconomic variables is due in part to volatility in the terms of trade.

Because of the nonlinear structure of our model, it is difficult to isolate the exact channels through which interactions between the level and volatility of the terms of trade affect the macroeconomy. However, much of the explanation may come from the fact that stochastic volatility increases the variance of the terms of trade and larger shocks to the terms of trade imply greater macroeconomic volatility. To see the impact of stochastic volatility on the variance of the terms of trade, first note that in the absence of stochastic volatility, that is if $\sigma_{q,t} = \sigma_q \forall t$, then the variance of the terms of trade, $var(q_t)$ is:

$$var(q_t) = \frac{\exp(2\sigma_q)}{1-\rho_q^2} \tag{3.34}$$

In contrast, when stochastic volatility is present, the variance of the terms of trade is:

$$var(q_t) = \frac{\exp(2\sigma_q + 2\zeta^2)}{1-\rho_q^2} \tag{3.35}$$

where $\zeta^2 = \frac{\eta_q^2}{1-\rho_\sigma^2}$. For Brazil, the presence of stochastic volatility increases the standard deviation of terms of trade shocks by a third, from 0.14 per cent to 0.19 per cent. Although other effects are likely to exist, this direct impact appears roughly large enough to explain the change in macroeconomic volatility between the scenario with terms of trade level shocks only and the scenario with both terms of trade level and volatility shocks.

3.6 Robustness Checks

In this section we examine the robustness of our results to alternative parameter assumptions and modelling choices.

3.6.1 Alternative Parameter Values

As a first exercise, we test the sensitivity of our model's dynamics to alternative parameter values. In particular we consider: (i) increasing the inverse Frisch elasticity, ζ, from 2 to 100; (ii) increasing the inverse of the elasticity of substitution, σ, from 2 to 10; and (iii) increasing the parameter governing the sensitivity of the risk-free interest rate to the foreign debt level, ψ, from 10^{-3} to 10^{-2}. We examine each of these alternative parameter choices seperately, leaving the other parameters at the same level as in the baseline model presented above. Figure 3.8 shows impulse responses to a one standard deviation terms of trade volatility shock in Brazil under the alternative parameter values.

An increase in ζ makes labor supply less sensitive to changes in other macroeconomic variables. In fact, with $\zeta = 100$, aggregate hours worked are almost unchanged in response to a terms of trade volatility shock. Because of the muted labor supply response, home tradeables output increases by less following the shock than it does in the baseline case, while non-tradeables output decreases by more. The decrease in consumption and increase in the current account are also somewhat larger. The response of the current account reflects the fact that, when labor supply elasticity decreases, households are less able to smooth consumption by adjusting their working hours and are forced to rely more on changes in asset holdings. The decrease in consumption is required to square the smaller increase in tradeables output with a larger decrease in foreign debt.

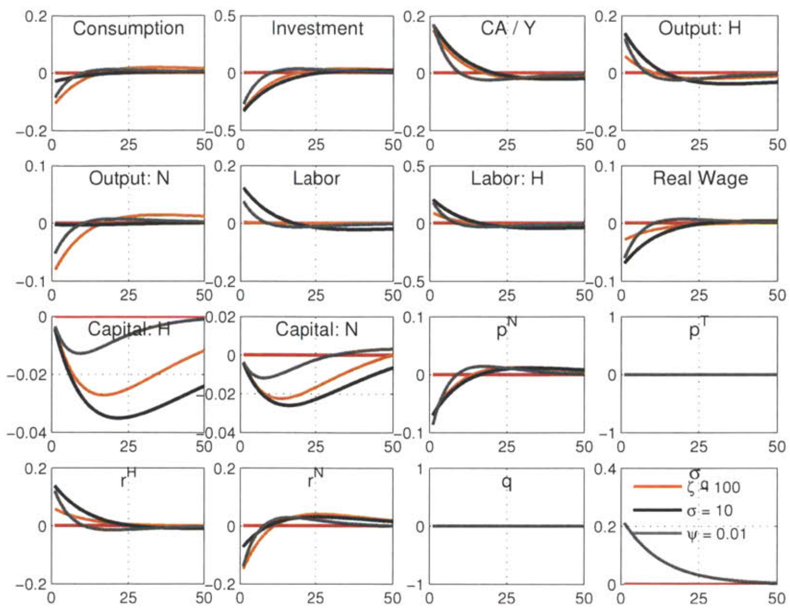

Figure 3.8: Terms of Trade Volatility Shock: Brazil - Alternative Parameters

Although macroeconomic models typically assume a low value for σ, microeconomic studies tend to report higher values (Hall (1988), Dynan (1993)). An increase in σ reduces the willingness of agents to trade consumption intertemporally. As a consequence, the decrease in consumption following the terms of trade volatility shock is considerably smaller in this scenario than it is in the baseline case. The decrease in non-tradeables output is also smaller than in the baseline case. However, the responses of the other variables are broadly similar to the baseline case.

In our baseline results we assumed an extremely low value of ψ in order to minimize the effect of this parameter on the dynamics of the model. Other papers that have estimated the value of this parameter for various countries have tended to find higher values (for example, Fernandez-Villaverde et al. (2011) for a selection of South American economies and Jaaskela and Nimark (2011) for Australia, while Justiniano and Preston (2010) calibrated this parameter to 10^{-2} in models of Australia, Canada and New Zealand). A higher value of ψ penalizes the economy for accumulating foreign assets (or reducing its foreign debt) by reducing the interest rate that agents receive on those assets. Setting this parameter to 10^{-2} reduces the amount of time that it takes for the model to converge to its steady state.

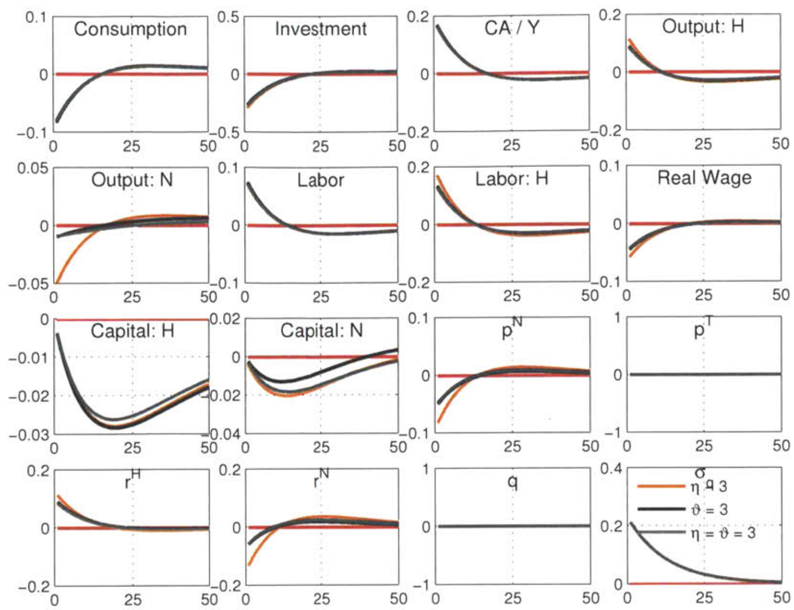

Figure 3.9: Terms of Trade Volatility Shock: Brazil - Alternative Parameters

However, the initial responses of the variables to the terms of trade volatility shock are broadly similar to the baseline model.

As a second parameter exercise, we also examined the effect of different elasticities of substitution between tradeable and non-tradeable goods, ϑ, and between home- and foreign-produced goods, η. In our baseline scenario, we calibrated both ϑ and η to values below one. This meant that these goods were complements in consumption. While this is a common assumption in the literature, other papers have also estimated values for these parameters above one, implying that these goods are substitutes in consumption.[15] To illustrate the sensitivity of our results to alternative values for these parameters, we simulated the model assuming that one, or both, of these parameters equaled three - a reasonable value in the literature. Figure 3.9 shows the results.

Changing the values of these parameters has little effect on aggregate consumption, investment or labor supply. It does, however, affect the composition of consumption between traded and non-traded goods and, consequently, the production of non-traded goods. Take the case where $\vartheta = 3$. In this scenario, traded and non-traded goods are substitutes, while home- and foreign-produced tradeable goods are complements. Relative to the baseline

[15] For a review of empirical estimates of these parameters, see Bodenstein (2010).

case, households now reduce their consumption of imports not only by reducing their overall consumption, but also by substituting towards non-tradeable goods.

Relative to the baseline case, a small change in the relative price of non-tradeable goods now leads to a large change in the demand for those goods. An increase in terms of trade volatility still reduces consumer demand for all goods, including non-traded goods - which reduces their price. However, consumers are now more willing to substitute from foreign-produced traded goods to non-traded goods. As a consequence, the contraction in the output of non-traded goods is smaller than in the baseline case, as is the decrease in the price of those goods. By shifting consumption from imports to non-traded goods, households are able to achieve a larger reduction in foreign debt despite a similar decrease in consumption. In contrast, increasing η to three has a smaller impact on the results relative to the baseline case. This reflects the fact that while the increase in η makes agents more willing to substitute between home- and foreign-produced tradeable goods, the shock to volatility does not change the relative price of these goods. Consequently, this parameter plays a relatively small role in influencing the response of the economy to a terms of trade volatility shock. In sum, altering the elasticities of substitution has relatively little qualitative effect on the results, although it has a quantitative effect on the responses in individual sectors of the economy.

3.6.2 Home-Produced Components of Investment

In our baseline model we assumed that the investment good was priced in units of the foreign tradeable good. This choice was motivated by the stylized fact that prices of investment goods differ less across countries than the prices of consumption goods (see **?**, Figures 4 and 5). In this exercise, we instead assume that the investment good is priced in the same units as the economy's consumption good. That is, we allow the prices of home-produced goods also to affect prices of investment goods. To do this, we assume that the investment good is a CES aggregate of tradeable and non-tradables goods and that the tradable component is itself an aggregate of home and foreign tradeable goods. For simplicity, we set the weights of each good and elasticities of substitution between alternative goods to the same values as they are for the consumption aggregate.[16]

Because we are interested in how including home-produced investment goods affects the fit of the model, as well as the dynamics, we first recalibrate the new model to match the same moments of the data as we did with the baseline model. Table 3.9 shows the resulting model moments for Brazil. Including a home-produced component of investment goods reduces the correlation between consumption and output and increases the correlation between investment and output. This improves the fit of the model for Brazil in this dimension, although it would worsen the fit of the model for other countries. It also reduces the correlation between net exports and output, although by not nearly enough to match the correlation seen in the data.

Table 3.9: Moments: Model with Home Produced Investment

	σ_y	σ_c/σ_y	σ_i/σ_y	$\rho_{c,y}$	$\rho_{i,y}$	$\rho_{nx,y}$	nx/y
Brazil	1.52	0.44	3.70	0.66	0.76	0.45	0.29

Next, we examine the dynamic response of the economy to a terms of trade volatility shock, shown in Figure 3.10. Including a home-produced component of investment has very little effect on the response of the economy to a volatility shock. The response of consumption is slightly more muted, while the investment response is slightly larger. However, the labor supply response is somewhat larger, as is the output response. In sum, in terms of both model fit and dynamics there seems to be little to choose between the two model specifications.

[16]To be precise: $i_t \equiv \left[\omega_T^{\frac{1}{\vartheta}} \left(i_t^T\right)^{\frac{\vartheta-1}{\vartheta}} + (1-\omega_T)^{\frac{1}{\vartheta}} \left(i_t^{NT}\right)^{\frac{\vartheta-1}{\vartheta}}\right]^{\frac{\vartheta}{\vartheta-1}}$ and $i_t^T \equiv \left[\omega_H^{\frac{1}{\eta}} \left(i_t^H\right)^{\frac{\eta-1}{\eta}} + (1-\omega_H)^{\frac{1}{\eta}} \left(i_t^F\right)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}}$.

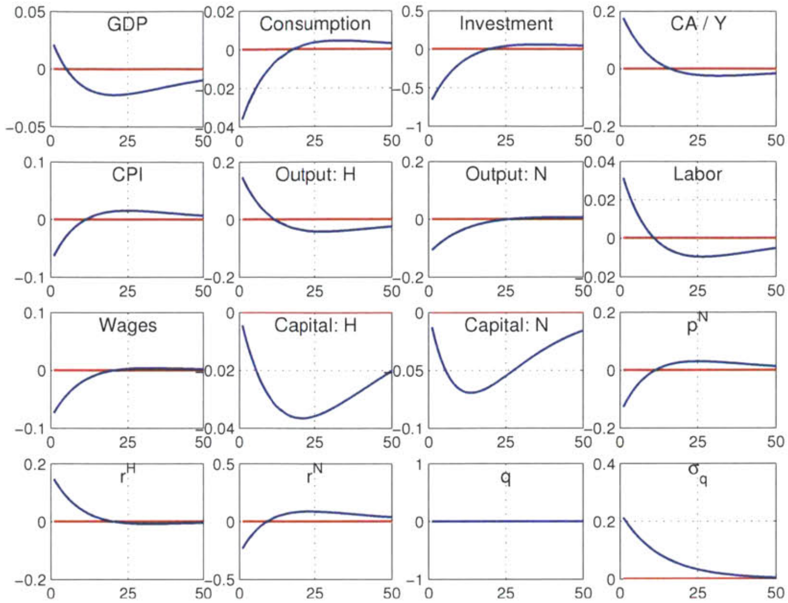

Figure 3.10: Terms of Trade Volatility Shock: Brazil - Home-Produced Investment

3.7 Conclusion

This paper has contributed to the literature examining time-varying volatility in macroeconomics by studying the effects of changes in the volatility of the terms of trade, a plausibly exogenous price index for small open commodity-exporting economies. Our empirical section estimated the empirical process for the terms of trade for six commodity-producing small open economies. We demonstrate that the magnitude of terms of trade shocks varies considerably over time for all the economies in our sample. Using a panel vector autoregression we then demonstrated that a volatility shock reduces both consumption and investment. Aggregate output also decreases following the shock and the current account-GDP ratio increases when the shock hits, and remains above trend before decreasing as domestic demand recovers. There is also a persistent decrease in the price level. Dividing the sample into emerging and developed economies shows some differences between the two categories, with the volatility shock having a larger price effect in developed economies and a larger effect on quantities in emerging economies.

In our theoretical section, we set up a small open economy real business cycle model and demonstrate that it can replicate the responses to the volatility shock generated by the VAR.

We then use the model to further explore the mechanisms behind these responses and also to examine their sectoral impacts. In the model, a shock to terms of trade volatility reduces consumption, causes a boom in the tradable sector at the expense of the non-tradeable sector and triggers a shift in the factor intensity of production away from capital towards labor. The decrease in domestic absorption and the increase in tradeables production leads to an improvement in the trade balance that allows the economy to reduce its foreign borrowing.

Finally, the model allows us to quantify the direct contribution of terms of trade volatility shocks to the fluctuations of macro-aggregates. Although the direct contribution of terms of trade volatility shocks to the variance of key variables is rather small, we find that these shocks have a meaningful economic effect in interaction with shocks to the level of the terms of trade. Our estimates suggest that terms of trade volatility shocks account for between one-fifth and one-third of the total effect of the terms of trade on the volatility of output, consumption, investment and net exports in the countries in our sample.

Our results point to a number of promising avenues for further research. The disaggregated VAR results hint that, for emerging economies, the response to volatility shocks occur mainly through quantities while, for developed economies, the response occurs mainly through prices. More detailed empirical work using a larger sample of economies could shed light on the robustness of this result. And, if it does turn out to be robust, further theoretical work is needed to understand the economic drivers of this observation.

3. A Data Sources and Definitions

Terms of Trade Data

With the exception of Canda, all terms of trade data was sourced from national statistical agencies. We retrieved data for Canada from the OECD. For Australia, Brazil, New Zealand and South Africa, published terms of trade indexes were used. For Canada, we constructed a terms of trade index by dividing the exports of goods and services deflator by the imports of goods and services deflator. For Mexico, we constructed a terms of trade index by dividing the exports price index by the imports price index. The raw data for Australia, Canada, New Zealand and South Africa was quarterly. For Brazil and Mexico, we constructed a quarterly series using quarterly averages of monthly data. Samples and sources for the individual countries are:

Australia: sample 1959-2010Q4. Source: Australian Bureau of Statistics (www.abs.gov.au).

Brazil: Sample 1978Q1 - 2010Q4. Source: IPEA (www.ipeadata.gov.br).

Canada: Sample: 1961Q1 - 2010Q4. Source: OECD (www.oecd.org).

Mexico: Sample: 1970Q1 - 2010Q4. Source: Instituto Nacional de Estadistica Y Geografia (dgcnesyp.inegi.org.mx).

New Zealand: Sample: 1957Q1 - 2010Q4. Source: Statistics New Zealand (www.stats.govt.nz).

South Africa: Sample 1960Q1 - 2010Q4. Source: Reserve Bank of South Africa (www.resbank.co.za).

National Accounts Data

For all countries, data for Gross Domestic Product and its components was sourced from the OECD economic outlook database (www.oecd.org). All national accounts data are HP-filtered using a smoothing parameter of 1600.

3. B Terms of Trade Processes: HP Filtered Data

Table 3.10: Posterior Medians: HP Filtered

	Australia	Brazil	Canada	Mexico	New Zealand	South Africa
ρ_q	0.84	0.77	0.83	0.78	0.83	0.68
	[0.81, 0.89]	[0.65, 0.89]	[0.77, 0.91]	[0.71, 0.89]	[0.77, 0.92]	[0.60, 0.80]
σ_q	-3.69	-3.36	-4.45	-3.46	-3.53	-3.38
	[-4.19, -2.95]	[-3.87, -2.74]	[-4.82, -4.00]	[-3.88, -3.02]	[-3.96, -2.91]	[-3.95, -2.81]
ρ_σ	0.94	0.93	0.89	0.87	0.95	0.97
	[0.78, 0.99]	[0.78, 0.99]	[0.62, 0.99]	[0.66, 0.97]	[0.86, 1.00]	[0.88, 1.00]
η_q	0.20	0.21	0.23	0.31	0.14	0.10
	[0.12, 0.35]	[0.11, 0.38]	[0.11, 0.45]	[0.20, 0.46]	[0.08, 0.24]	[0.05, 0.21]

3. C What does the empirical VAR capture?

In this Appendix, we demonstrate the ability of our empirical vector autoregression exercise to capture the macroeconomic impacts of exogenous shocks to terms of trade volatility. To do this, we compare impulse responses from our empirical VAR estimated using simulated data to the impulse responses to exogenous terms of trade volatility shocks generated by our model. Specifically, we simulate our model for 200 periods setting all parameters at their baseline values for Brazil. We then estimate our empirical VAR using this data and calculate impulse responses to an innovation to the terms of trade volatility variable as in Section (3.3). We repeat this process 50,000 times to characterize the distribution of VAR responses.

Figure (3.11) shows the median, 5 and 95 per cent responses of the simulated VAR for each variable as well as the theoretical responses to a terms of trade volatility shock from the model. Despite its linear structure, the VAR comes extremely close to matching the theoretical model responses. This gives us confidence that our empirical model reflects a response to an exogenous terms of trade volatility shock.

Figure 3.11: Terms of Trade Volatility Shock: Model and VAR

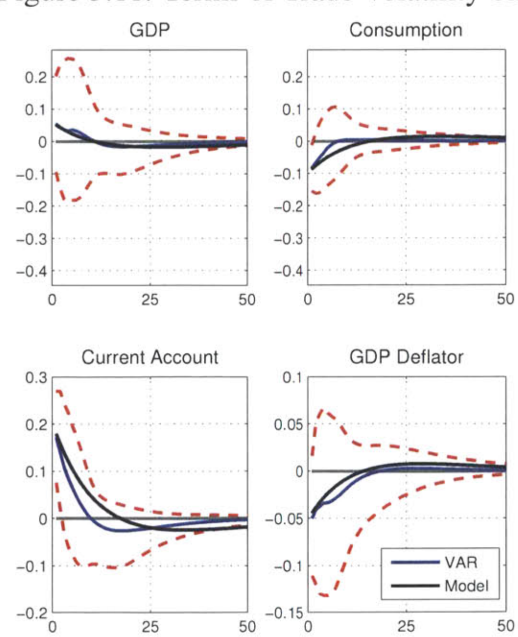

3. D Theoretical Impulse Response Functions: Other Economies

Figure 3.12: Terms of Trade Volatility Shock: Australia

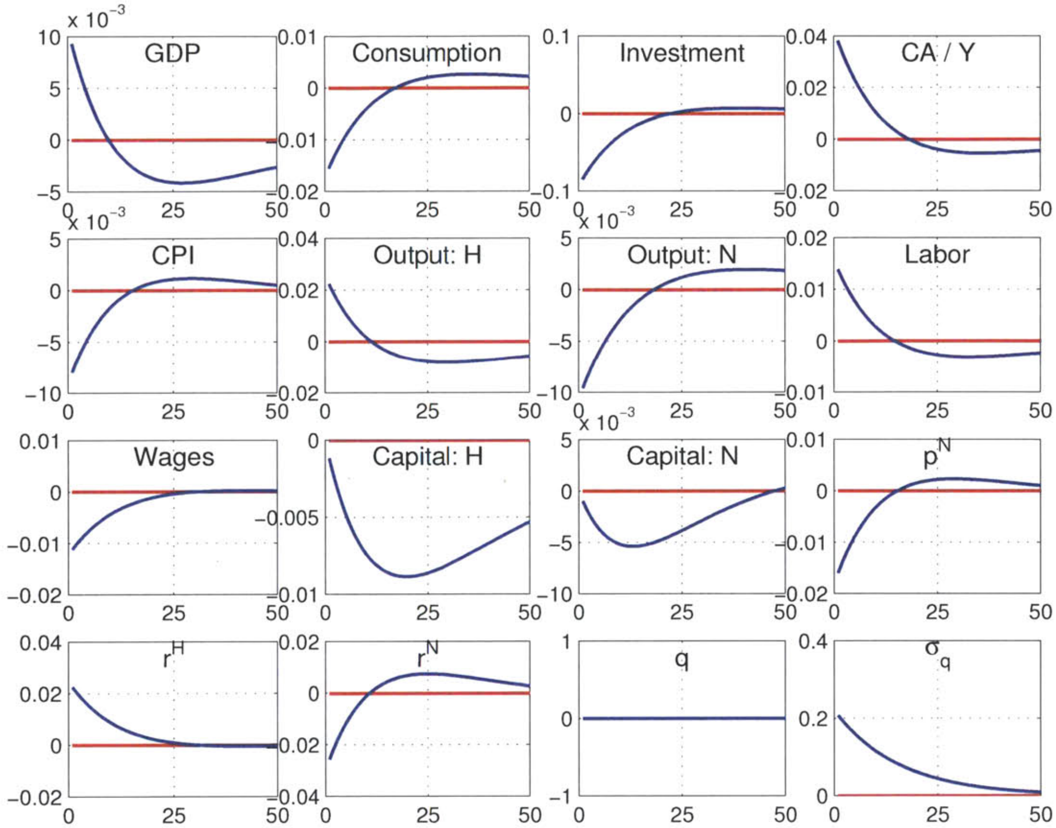

Figure 3.13: Terms of Trade Volatility Shock: Canada

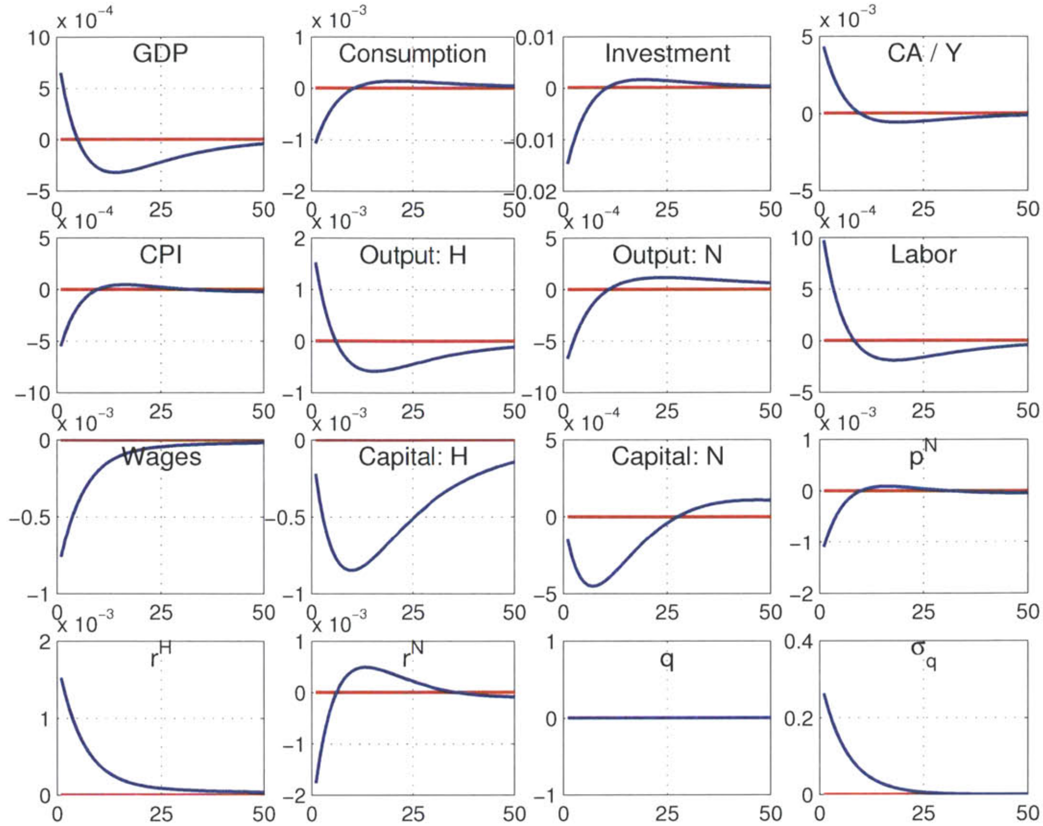

Figure 3.14: Terms of Trade Volatility Shock: Mexico

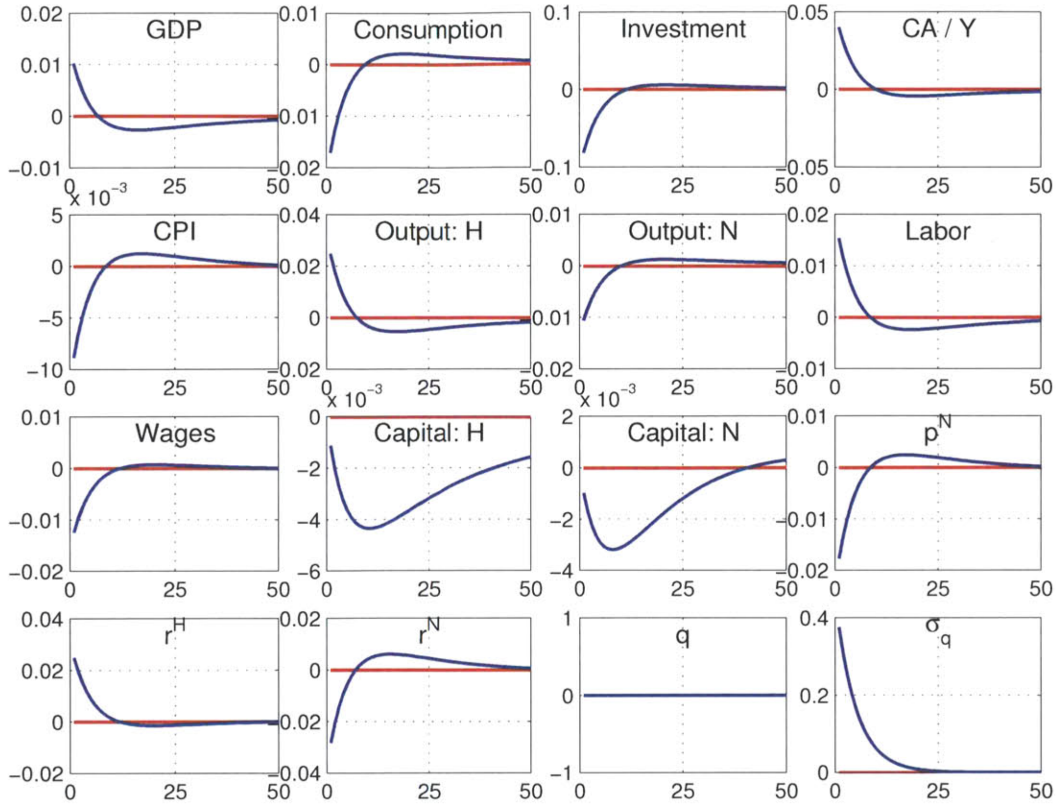

Figure 3.15: Terms of Trade Volatility Shock: New Zealand

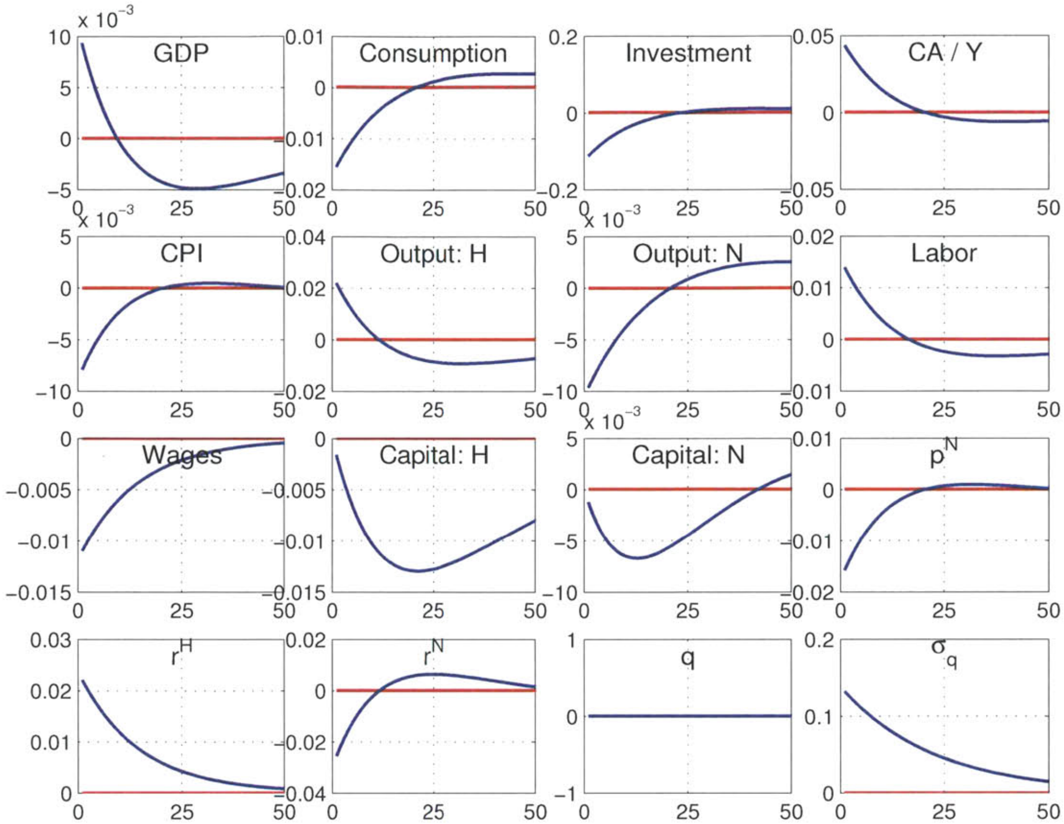

Figure 3.16: Terms of Trade Volatility Shock: South Africa

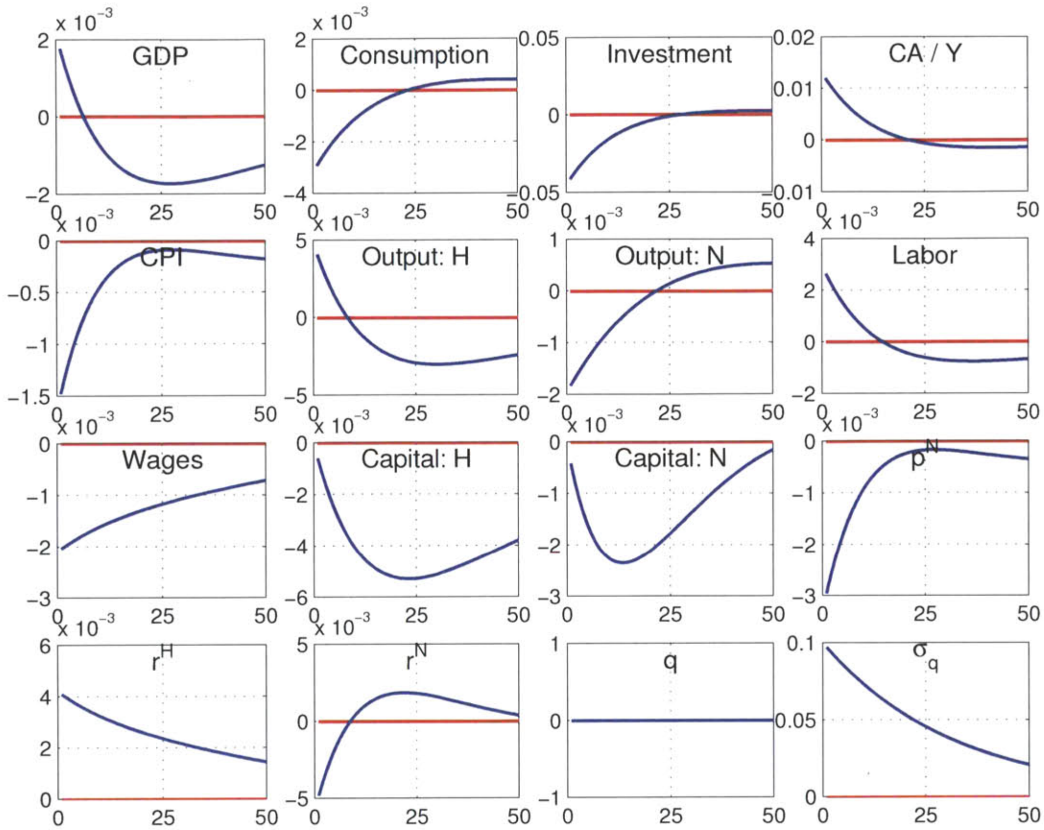

Bibliography

Aguiar, M., Gopinath, G., 2007. Emerging market business cycles: The cycle is the trend. Journal of Political Economy 115 (1), 69–102.

An, S., Schorfheide, F., 2007. Bayesian analysis of DSGE models. Econometric Reviews 26, 113–172.

Andrews, D., Rees, D., 2009. Macroeconomic volatility and terms of trade shocks. Research Discussion Paper 2009-05, Reserve Bank of Australia.

Angeletos, G. M., L'ao, J., 2010. Noisy business cycles. NBER Macroeconomics Annual 24 (1), 319–378.

Aoki, K., 2003. On the optimal monetary policy response to noise indicators. Journal of Monetary Economics 50, 501–523.

Bache, I. W., Brubakk, L., Maih, J., 2010. Simple rules versus optimal policy: What fits? Norges Bank Working Paper 2010-3.

Bartelsman, E. J., Caballero, R. J., Lyons, R. K., 1994. Customer- and supplier-driven externalities. American Economic Review 84 (4), 1075–1084.

Benigno, G., Benigno, P., Ghironi, F., 2007. Interest rate rules for fixed exchange rate regimes. Journal of Economic Dynamics and Control 31 (7), 2196–2211.

Benigno, P., Woodford, M., 2012. Linear-quadratic approximation of optimal policy problems. Journal of Economic Theory 147, 1–42.

Bils, M., Klenow, P., 2004. Some evidence on the importance of sticky prices. Journal of Political Economy 112 (5), 927–985.

Blanchard, O., Milesi-Ferretti, G. M., 2009. Global imbalances: In midstream. Imf staff position note, International Monetary Fund.

Blanchard, O. J., L'Huillier, J.-P., Lorenzoni, G., 2012. News, noise, and fluctuations: An empirical exploration. mimeo.

Blattman, C., Hwang, J., Williamson, J. G., 2007. Winners and losers in the commodity lottery: The impact of terms of trade growth and volatility in the periphery 1870-1939. Journal of Development Economics 82, 156–179.

Bleaney, M., Greenaway, D., 2001. The impact of terms of trade and real exchange rate volatility on investment and growth in sub-saharan africa. Journal of Development Economics 65, 491–500.

Blinder, A. S., Canetti, E. R. D., Lebow, D. E., Rudd, J. B., 1998. Asking About Prices. Russell Sage Foundation, New York, NY.

Bodenstein, M., 2010. Trade elasticity of substitution and equilibrium dynamics. Journal of Economic Theory 145 (3), 1033–1059.

Boz, E., Daude, C., Durdu, C. B., 2011. Emerging market business cycles: Learning about the trend. Journal of Monetary Economics 58, 616–631.

Broda, C., 2004. Terms of trade and exchange rate regimes in developing countries. Journal of International Economics 63 (1), 31–58.

Broda, C., Weinstein, D., 2006. Globalization and gains from variety. Quarterly Journal of Economics 121 (541-585).

Caballero, R. J., Lorenzoni, G., 2009. Persistent appreciations and overshooting: A normative analysis. mimeo, MIT.

Calvo, G., 1983. Staggered prices in a utility-maximizing framework. Journal of Monetary Economics 12 (3), 383–398.

Cashin, P., Liang, H., McDermott, C. J., 2000. How persistent are shocks to world commodity prices? IMF Staff Papers 47 (2), 177–217.

Catao, L., Chang, R., 2012. Monetary rules for commodity traders. Presented at the Policy Responses to Commodity Price Movements conference April.

Collier, P., Goderis, B., 2012. Commodity prices and growth: An empirical investigation. European Economic Review 56, 1241–1260.

Cooper, R., Johri, A., 1997. Dynamic complementarities: A quantitative analysis. Journal of Monetary Economics 40, 97–119.

Corden, W. M., 1982. Exchange rate policy and the resources boom. Economic Record 58 (160), 18–31.

Corden, W. M., Neary, J. P., 1982. Booming section and de-industrialisation in a small open economy. Economic Journal 92 (368), 825–848.

Corsetti, G., Dedola, L., Leduc, S., 2008. International risk sharing and the transmission of productivity shocks. The Review of Economic Studies 75, 443–473.

Corsetti, G., Pesenti, P., 2001. Welfare and macroeconomic interdependence. Quarterly Journal of Economics 116 (2), 421–445.

Debortoli, D., Nunes, R., 2006. On linear-quadratic approximations. Draft, university of california at san diego.

Dib, A., 2008. Welfare effects of commodity price and exchange rate volatilities in a multi-sector small open economy model. Working Paper 2008-8, Bank of Canada.

Drazen, A., 1985. Cyclical determinants of the natural level of economic activity. International Economic Review 26 (2), 387–397.

Dynan, K., 1993. How prudent are consumers? Journal of Political Economy 101, 1104–1113.

Ehrmann, M., Smets, F., 2003. Uncertain potential output: Implications for monetary policy. Journal of Economic Dynamics and Control 27, 1611–1638.

Fernandez-Villaverde, J., Guerron-Quintana, P., Kuester, K., Rubio-Ramirez, J., January 2012. Fiscal volatility shocks and economic activity. mimeo.

Fernandez-Villaverde, J., Guerron-Quintana, P., Rubio-Ramirez, J. F., Uribe, M., 2011. Risk matters: The real effects of volatility shocks. American Economic Review 101, 2530–2561.

Gali, J., Monacelli, T., 2005. Monetary policy and exchange rate volatility in a small open economy. Review of Economic Studies 72, 707–734.

Geweke, J., 1999. Computational experiments and reality. University of Minnesota and Federal Reserve Bank of Minneapolis.

Gregory, R., 1976. Some implications of the growth of the mineral sector. Australian Journal of Agricultural Economics 20 (2), 79–91.

Hall, R. E., 1988. Intertemporal substitution in consumption. Journal of Political Economyl 96 (2), 339–357.

Harberger, A., 1950. Currency depreciation, income and the balance of trade. Journal of Political Economy 58, 47–60.

Hevia, C., Nicolini, J. P., 2012. Optimal devaluations. mimeo.

Holmes, T. J., 1999. Localization of industry and vertical disintegration. Review of Economics and Statistics 81 (2), 314–325.

Jaaskela, J., Nimark, K. P., 2011. A medium-scale open economy model of Australia. Economic Record 86, 11–36.

Jaffe, A. B., Trajtenberg, M., Henderson, R., 1993. Geographic localization of knowledge spillovers and evidenced by patent citations. Quarterly Journal of Economics 108 (3), 577–598.

Justiniano, A., Preston, B., 2010. Can structural small open-economy models account for the influence of foreign disturbances? Journal Of International Economics 81 (1), 61–74.

Justiniano, A., Primiceri, G. E., 2008. The time-varying volatility of macroeconomic fluctuations. American Economic Review 98 (3), 604–641.

Kearns, J., Lowe, P., 2011. Australia's prosperous 2000s: Housing and the mining boom. Reserve Bank of Australia Research Discussion Papers 2011-07.

Kent, C., Cashin, P., 2003. The response of the current account to terms of trade shocks: Persistence matters. Working Paper 143, International Monetary Fund.

Kilian, L., Hicks, B., 2012. Did unexpectedly strong economic growth cause the oil price shock of 2003–2008? Journal of Forecasting Forthcoming.

Kim, J., Kim, S. H., 2003. Spurious welfare reversals in international business cycle models. Journal of International Economics 60 (2), 471–500.

Kim, T.-H., Pfaffenzeller, S., Rayner, T., Newbold, P., 2003. Testing for linear trend with application to relative primary commodity prices. Journal of Time Series Analysis 24 (5), 539–551.

Kose, M. A., Riezman, R., 2001. Trade shocks and macroeconomic fluctuations in africa. Journal of Development Economics 65, 55–80.

Krugman, P., 1987. The narrow moving band, the Dutch disease, and the competitive consequences of Mrs. Thatcher: Notes on trade in the presence of dynamic scale economies. Journal of development Economics 27 (1), 41–55.

Kulish, M., Rees, D., 2011. The yield curve in a small open economy. Journal Of International Economics 85, 268–279.

Lama, R., Medina, J. P., 2012. Is exchange rate stabilization an appropriate cure for the dutch disease. International Journal of Central Banking 8 (1), 5–46.

Lartey, E. K. K., 2008. Capital inflows, Dutch disease effects and monetary policy in a small open economy. Review of International Economics 16 (5), 971–989.

Laursen, S., Metzler, L., 1950. Flexible exchange rates and the theory of employment. Review of Economics and Statistics 32, 281–299.

Lee, J., List, J. A., Strazicich, M. C., 2006. Non-renewable resource prices: Deterministic or stochastic trends? Journal of Environmental Economics and Management 51 (3), 354–370.

Lubik, T. A., 2007. Non-stationarity and instability in small open-economy models even when they are "closed". Federal Reserve Bank of Richmond Economic Quarterly 93 (4), 393–412.

Macklem, R. T., July 1993. Terms-of-trade disturbances and fiscal policy in a small open economy. Economic Journal 103 (419), 916–936.

Maslyuk, S., Smyth, R., 2008. Unit root properties of crude oil spot and future prices. Energy Policy 36 (7), 2591–2600.

Mendoza, E., 1995. The terms of trade, the real exchange rate, and economic fluctuations. International Economic Review 36 (1), 101–137.

Mendoza, E. G., 1997. Terms-of-trade uncertainty and economic growth. Journal of Development Economics 54, 323–356.

Murchison, S., Rennison, A., 2006. ToTEM: The Bank of Canada's new quarterly projection model. Bank of Canada Technical Report 97.

Murray, J., 2010. Is there a commodity curse? Lessons from the past. Presentation to the University of Alberta Institute for Public Economics 6 May 2010.

Nakamura, E., Steinsson, J., 2008. Five facts about prices: A reevaluation of menu cost models. Quarterly Journal of Economics 123 (4), 1415–1464.

Newbold, P., Pfaffenzeller, S., Rayner, T., 2005. How well are long-run commodity prices characterized by trend components? Journal of International Develpment 17, 479–494.

Obstfeld, M., 1982. Aggregate spending and the terms of trade: Is there a Laursen-Metzler effect. Quarterly Journal of Economics 91, 251–270.

Orphanides, A., 2003. Monetary policy evaluation with noisy information. Journal of Monetary Economics 50, 605–631.

Otto, G., 2003. Terms of trade shocks and the balance of trade: there is a Harberger-Laursen-Metzler effect. Journal Of International Money and Finance 22, 155–184.

Pearlman, J., 1992. Reputational and nonreputational policies under partial information. Journal of Economic Dynamics and Control 16, 339–357.

Plumb, M., Kent, C., Bishop, J., 2012. Implications for the Australian economy of strong growth in Asia. Paper presented at the Structural Change and the Rise of Asia Conference, Sydney, September 2012.

Powell, A., 1991. Commodity and developing country terms of trade: What does the long run show? The Economic Journal 101 (409), 1485–1496.

Rees, D. M., 2012. Terms of trade shocks and incomplete information in small open economies. mimeo.

Rudebusch, G. D., 2001. Is the Fed too timid? monetary policy in an uncertain world. The Review of Economics and Statistics 83 (2), 203–217.

Sachs, J., 1981. The current account and macroeconomic adjustment in the 1970s. Brookings Papers on Economic Activity 1, 201–268.

Schmitt-Grohe, S., Uribe, M., 2003. Closing small open economy models. Journal of International Economics 61, 163–185.

Schmitt-Grohe, S., Uribe, M., 2007. Optimal simple and implementable monetary and fiscal rules. Journal of Monetary Economics 54, 1702–1725.

Smets, F., 2002. Output gap uncertainty: Does it matter for the Taylor rule? Empirical Economics 27, 113–129.

Stevens, G., 2011. The resources boom. Reserve Bank of Australia Bulletin March 2011, 67–72.

Stockman, A. C., Tesar, L. L., 1995. Tastes and technology in a two-country model of the business cycle: Explaining international comovements. The American Economic Review 85, 168–185.

Sutherland, A., 2002. A simple second-order solution method for dynamic general equilibrium models. Tech. rep., CEPR Discussion Paper No. 3554.

Svensson, L., Razin, A., 1983. The terms of trade and the current account: The Harberger-Laursen-Metzler effect. Journal of Political Economy 91, 97–125.

Svensson, L., Woodford, M., 2003. Indicator variables for optimal policy. Journal of Monetary Economics 50, 691–720.

Svensson, L., Woodford, M., 2004. Indicator variables for optimal policy under asymmetric information. Journal of Economic Dynamics and Control 28, 661–690.

Torvik, R., 2001. Learning by doing and the Dutch disease. European Economic Review 45 (2), 285–306.

Uhlig, H., 1999. A toolkit for analyzing nonlinear dynamic stochastic models easily. In: Marimom, R., Scott, A. (Eds.), Computational Methods for the Study of Dynamic Economies. Oxford University Press, pp. 30–61.

van Wignbergen, S., 1984. The 'Dutch disease': A disease after all? The Economic Journal 94, 41–55.

Williamson, J. G., 2008. Globalization and the great divergence: terms of trade booms, volatility and the poor periphery, 1782-1913. European Review of Economic History 12, 355–391.