

Abstract

How should monetary policy respond to large fluctuations in world food prices? We study this question in an open economy model in which imported food has a larger weight in domestic consumption than abroad and international risk sharing can be imperfect. A key novelty is that the real exchange rate and the terms of trade can move in opposite directions in response to world food price shocks. This exacerbates the policy trade-off between stabilizing output prices *vis a vis* the real exchange rate, to an extent that depends on risk sharing and the price elasticity of exports. We characterize implications for dynamics, optimal monetary policy, and the relative performance of practical monetary rules. While CPI targeting and expected CPI targeting can dominate PPI targeting if international risk sharing is perfect, even seemingly mild departures from the latter make PPI targeting a winner.

Keywords: Monetary Policy, Small Open Economy, Real Exchange Rates, Terms of Trade

JEL classification: F4,E5

1. Introduction

Inflationary bursts worldwide have been long associated with spiralling food prices. Granger-causality tests on post-1970 data corroborate this old-standing regularity, as global food commodity prices tend to lead rather than lag global CPI changes.¹ That food price shocks greatly matter for aggregate inflation has become particularly important to many inflation targeting countries over the past decade: a burst of food commodities inflation in 2007-08 led to widespread overshooting of inflation targets; this was followed by considerable undershooting of the targets once food prices receded. This evidence may not seem surprising, since food weighs heavily in most consumption baskets and is not easily substitutable by other goods. The surprise is that the monetary policy literature has given little attention to its implications.

To help filling the gap, this paper addresses the related questions of how far monetary policy should accommodate food price shocks and which policy rules, among those that are practically implementable, are best suited to shore up welfare. We model a small open economy that is a net food importer and where food weighs more heavily in domestic consumption than in world consumption. Faced with unexpectedly high world food prices, this economy experiences a terms of trade deterioration, higher CPI inflation, and a real exchange rate appreciation. This combination poses particularly stark policy trade-offs between domestic and external stabilization objectives. We characterize the transmission dynamics of exogenous shocks underlying these trade-offs under various degrees of international financial integration, and examine the implications for welfare and monetary policy choices.

In doing so, this paper relates to a rapidly growing literature on monetary policy in open economies, usefully surveyed by Corsetti, Dedola, and Leduc (2010). As emphasized by these authors, recent dynamic New Keynesian open economy models can yield different monetary policy prescriptions from their closed economy counterparts. In the latter, as summarized by Woodford (2003) and Gali (2008), optimal monetary policy is typically geared towards replicating a flexible price or natural outcome, suitably attainable through the stabilization of producer prices. Further, in the absence of mark-up/cost-push shocks and/or real wage rigidities, these models also imply that PPI stabilization is conducive to output stabilization. In contrast, consumption and production openness introduce additional policy trade-offs. In particular, small open economies can gain from steering the real exchange rate and the terms of trade. This "terms of trade externality" (Corsetti and Pesenti 2001) implies that the flexible price equilibrium is not generally optimal, hence raising the question of whether PPI stabilization remains the best policy. Several studies of the model developed by Gali and Monacelli (2005) have provided a basically affirmative answer (e.g. Faia and Monacelli 2008, and Di Paoli 2009). However, these studies have placed severe restrictions on the model environment, especially perfect international risk sharing.

*lavcatao@gmail.com

†Corresponding Author. Email: chang@econ.rutgers.edu

‡We thank the editors, an anonymous referee, Olivier Blanchard, Rodrigo Caputo, Pietro Cova, José de Gregorio, Philip Lane, Akito Matsumoto, Gian Maria Milesi-Ferretti, Carlos Végh as well as seminar participants at the Bank of Italy, Central Bank of Chile, European University Institute, Fundação Getúlio Vargas, IADB, IMF, PUC-RJ, Swiss National Bank, Trinity College Dublin, and the 2010 World Econometric Society Congress for discussions and comments on earlier drafts. We are particularly grateful to Marola Castillo for outstanding research assistance and to Roberto Rigobón for sharing his food expenditure database with us. The opinions expressed here are, however, the authors' alone and do not necessarily reflect those of the IMF or any of the above colleagues or institutions.

¹This is so even after controlling for oil prices. These claims readily follow from regressing changes in the GDP weighed world CPI on changes in the log of the IMF price indices of global food and of oil commodities over the period 1970-2011. The F-statistic on the exclusion of lagged food inflation is significant at 1%. In contrast, the significance of Granger-causality F-statistics on oil prices is generally weaker and of varying significance across sub-periods.

This paper extends the environment of these previous studies in several ways. We model "food" as a key import, traded in flex-price competitive markets, and which enters as a distinct commodity in the home consumption basket with possibly a very low elasticity of substitution *vis a vis* other goods. Other extensions include: (i) global food prices can vary widely relative to the world price index; (ii) food expenditure shares at home and the rest of the world can differ significantly; (iii) the export price elasticity of the world demand for home exports can differ from the intratemporal elasticity of substitution of home and imported goods in consumption; (iv) international risk sharing can be incomplete.

Extensions i and ii allow our model to capture a much overlooked empirical regularity, already mentioned but still worth emphasizing: with shocks to the world relative price of food, the terms of trade and the real exchange rate can move in opposite directions. Such a negative covariance is ruled out by previous models. Empirically however, and as shown in Figure 1, a negative covariance is often found in economies that are net food importers and that export sticky price high elasticity goods (as allowed by extension iii). Extension iv, that international risk sharing can be imperfect, hardly needs justification. But departing from the assumption of perfect risk sharing introduces several technical difficulties, which may explain why the assumption is ubiquitous. In this paper we follow Schulhofer-Wohl (2011) in assuming complete financial markets but also a costly wedge in the transferring of resources in and out of the domestic household. This formulation implies that domestic consumption is a combination of the polar cases of perfect risk sharing and portfolio autarky, leading to a specification that is parsimonious, intuitive, and relatively easy to calibrate. As such, it is of independent interest.

[PLEASE INSERT FIGURE 1 APPROXIMATELY HERE]

In the resulting setting, we provide a complete characterization of Ramsey and natural allocations, as well as of optimal feasible optimal policy, extending the analysis of Faia and Monacelli (2008) and Di Paoli (2009). We combine both analytical approaches with extensive numerical calibrations to flesh out the role of the degree of international risk sharing and of structural elasticities for optimal policy and welfare-based comparisons of policy rules.

Main findings include: First, in the presence of shocks to world food prices, the relative desirability of home inflation versus output gap stabilization varies significantly depending on the extent of risk sharing and on the export price elasticity. In particular, if the latter is sufficiently but also realistically high (that is, as the economy is "smaller" in export markets), less international risk sharing implies that optimal policy places a heavier weight on domestic price stabilization. Second, when the variance of imported food price shocks is calibrated to be as large as in the data, international risk sharing is perfect, and the home economy's export price elasticity is not too low, CPI targeting can deliver higher welfare than PPI targeting. But targeting "expected" or forecast CPI is even superior. The reason is that expected CPI targeting exploits more heavily the terms of trade externality, resulting in more stable real exchange rate and consumption; in doing so, it delivers a better approximation to the optimal allocation than the competing rules. Third, the welfare-superiority of PPI targeting is easily restored if international risk sharing is less than complete: for a wide range of the other parameters, even small values of the financial transfer cost wedge imply that PPI dominates other rules. In this sense, the conditions for PPI stabilization to be the optimal policy are broader than highlighted in previous work, which relied on perfect risk sharing and the domestic good substitution elasticity being the same as the export good elasticity. Fourth, an optimal feasible policy can be characterized by a "flexible targeting rule", a linear combination of domestic inflation and deviations of output from a target. The output target is a function of exogenous shocks, with coefficients that depend on elasticities of demand and the degree of risk sharing.

The remainder of the paper proceeds as follows. Section 2 lays out the basic framework. Section 3 discusses the model's linearized representation and dynamic responses to world food price shocks. Section 3 characterizes optimal policy. Numerical calibrations and welfare ranking of policy rules is presented in section 4. Section 5 concludes. To preserve space, a Technical Appendix (available from ScienceDirect) gathers several formal expressions and derivations.

2. Model

We study a small open economy populated by identical agents that consume a domestic good and imported food. The domestic good is an aggregate of intermediate varieties produced with domestic labor. The intermediates sector is characterized by monopolistic competition and nominal price rigidities.

The share of food is larger in the domestic consumption basket than in the world basket, so PPP does not hold. Further, the world price of food in terms of world consumption is exogenous. One consequence is that the real exchange rate appreciates when the world relative price of food rises, and domestic consumption fluctuates with world food prices. Also, and in contrast with previous work, our model implies that the terms of trade and the real exchange rate can move in opposite directions.

Another novelty is that international risk sharing is allowed to be imperfect because domestic households may face costs of transferring resources, as in Schulhofer-Wohl (2011).

2.1. Households

The economy has a representative household that maximizes the expected value of $\sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\sigma}}{(1-\sigma)} - \varsigma \int_0^1 \frac{N_t(j)^{1+\varphi}}{1+\varphi} dj \right]$, where $0 < \beta < 1$, σ, φ , and ς are parameters, C_t denotes consumption, and $N_t(j)$ is the supply of labor employed by a firm belonging to industry $j \in [0, 1]$. As in Woodford (2003), there is a continuum of industries, each employing a different type of labor. Labor types are imperfect substitutes if $\varphi > 0$.

Consumption is a C.E.S. aggregate of a home final good C_h and imported food C_f :

$$C_t = \left[(1-\alpha)^{1/\eta} C_{ht}^{(\eta-1)/\eta} + \alpha^{1/\eta} C_{ft}^{(\eta-1)/\eta} \right]^{\eta/(\eta-1)}$$

where η is the elasticity of substitution between home and foreign goods, and α is a parameter that measures the degree of openness. The associated price index, or CPI, expressed in domestic currency, is then

$$P_t = \left[(1-\alpha) P_{ht}^{1-\eta} + \alpha P_{ft}^{1-\eta} \right]^{1/(1-\eta)} \quad (1)$$

where P_{ht} and P_{ft} are the domestic currency prices of the home good and imports. Also, given total consumption C_t and prices P_{ht} and P_{ft} , optimal demands for home goods and foreign goods are given by

$$C_{ht} = (1-\alpha) (P_{ht}/P_t)^{-\eta} C_t \quad (2)$$

and $C_{ft} = \alpha \left(\frac{P_{ft}}{P_t} \right)^{-\eta} C_t$, where P_{ht}/P_t is the price of home output in terms of home consumption (the "real" price of home output).²

The household owns domestic firms and receives their profits. It chooses consumption and labor effort taking prices and wages as given. With respect to trade in assets, we depart from Gali and Monacelli (2005), Di Paoli (2009) and others in allowing for financial frictions and imperfect risk sharing across countries. Specifically, we borrow Schulhofer-Wohl's (2011) closed-economy assumption that the typical household incurs deadweight costs if it transfers resources in or out of the household. Denoting the household's current nonfinancial income by H_t , the assumption is that the household has to pay an extra cost of $\varpi \Phi(C_t, H_t)$ units of consumption, where $\Phi(C, H) = \frac{\varsigma}{2} \left(\log \left(\frac{C}{H} \right) \right)^2$ and ϖ is a parameter controlling the severity of this friction. This formulation implies that optimal risk sharing is given by

$$C_t^\sigma [1 + \varpi \Phi_{Ct}] = \kappa X_t (C_t^*)^\sigma \quad (3)$$

where κ is a positive constant, C_t^* is an index of world consumption, X_t is the *real exchange rate* (the ratio of the price of world consumption to the domestic CPI, both measured in a common currency), and $\Phi_{Ct} = \Phi_C(C_t, H_t)$ is the partial derivative of Φ with respect to C evaluated at (C_t, H_t) .³ If $\varpi = 0$, the preceding expression reduces to the usual perfect international risk sharing condition: marginal utilities of consumption at home and abroad are proportional up to a real exchange rate correction. For nonzero ϖ , optimal risk sharing takes into account that each consumption unit transferred to domestic households involves the extra transfer cost $\varpi \Phi_c$, explaining the appearance of this term in the left hand side. Financial autarky corresponds to ϖ going to infinity: in that case, and using Y_t to denote domestic output, $C_t = H_t = (P_{ht}/P_t) Y_t$ in equilibrium, so the trade balance is zero in all periods. Schulhofer-Wohl's assumptions thus capture market incompleteness in a way that is attractive in its simplicity, encompassing perfect risk sharing and portfolio autarky as special cases, and (as found below) retaining tractability.⁴

Next, if $W_t(j)$ is the domestic wage for labor of type j , optimal labor supply is given by the equality of the marginal

²Home bias corresponds to the case $\alpha < 1/2$. We have assumed $\eta \neq 1$. If $\eta = 1$, C_t and P_t are Cobb Douglas.

³Let $\Omega_{t,t+1}$ denote the domestic currency price at t of a security that pays a unit of domestic currency at $t+1$ conditional on some state of nature s' being realized at that time. Then optimal consumption requires $\Omega_{t,t+1} = \beta \frac{P_t}{P_{t+1}} \left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \frac{1 + \varpi \Phi_c(C_t, H_t)}{1 + \varpi \Phi_c(C_{t+1}, H_{t+1})}$, reflecting that the effective cost of an extra unit of consumption at t is not P_t but $P_t(1 + \varpi \Phi_c(C_t, H_t))$. For the rest of the world, we assume that there is no transferring cost, so the corresponding FOC is $\Omega_{t,t+1} = \beta \frac{S_t P_t^*}{S_{t+1} P_{t+1}^*} \left(\frac{C_{t+1}^*}{C_t^*} \right)^{-\sigma}$. The usual derivation for the complete markets case can then be amended to yield (3).

⁴One might, of course, object that ϖ may be time-varying and not readily mapped onto observables. But similar objections could be equally raised to the obvious alternative, which is a bond economy. As noted in Schulhofer-Wohl (2011), assuming risk is shared imperfectly via noncontingent bond contracts also amounts to a reduced form specification. In practice, risk sharing takes place through a variety of other financial instruments, both formal and informal, official and private.

disutility of labor with the marginal utility of the real wage, corrected by marginal transfer costs:

$$\varsigma C_t^\sigma N(j)_t^\varphi = \frac{W_t(j)[1 - \varpi \Phi_{Ht}]}{P_t [1 + \varpi \Phi_{Ct}]} \quad (4)$$

Finally, the domestic safe interest rate is given by

$$\frac{1}{1 + i_t} = \beta E_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \frac{P_t (1 + \varpi \Phi_{Ct})}{P_{t+1} (1 + \varpi \Phi_{Ct+1})} \right] \equiv E_t M_{t,t+1} \quad (5)$$

where we have defined $M_{t,t+j}$ as the period t pricing kernel applicable to nominal payoffs in period $t + j$. This extends the familiar expression of the frictionless asset trade case.

2.2. Prices

For simplicity, we assume that all food is imported and that the world price of food is exogenously given in terms of a world currency. Using asterisks to denote prices denominated in world currency, the domestic currency price of food is then $P_{ft} = S_t P_{ft}^*$, where S_t is the nominal exchange rate (domestic currency per unit of foreign currency). So, there is full pass through from world to domestic food prices.

Likewise, we assume that the world currency price of the world consumption index is exogenous.⁵ Denoting it by P_t^* , the real exchange rate is then $X_t = S_t P_t^* / P_t$. It is useful also to define the *terms of trade* by $Q_t = P_{ft} / P_{ht} = S_t P_{ft}^* / P_{ht}$.

As in other models, the terms of trade and the real price of home output are essentially the same, since (1) implies that $(P_{ht} / P_t)^{-(1-\eta)} = (1 - \alpha) + \alpha Q_t^{1-\eta}$. But, in contrast with those other models, the real exchange rate and the terms of trade are not proportional to each other, reflecting fluctuations in the world price of food relative to the world CPI. The previous definitions imply the following relation between the real exchange rate and the real price of home output:

$$P_{ht} / P_t = \left[\frac{1 - \alpha X_t^{1-\eta} Z_t^{1-\eta}}{1 - \alpha} \right]^{1/(1-\eta)} \quad (6)$$

where $Z_t = P_{ft}^* / P_t^*$ is the *world's relative price of food*, which we take as exogenous.

An improvement in the terms of trade (a fall in Q_t) implies an increase in the real price of output (P_{ht} / P_t). If Z_t is held fixed, (6) then implies that X_t must fall (a real appreciation). But this logic does not apply when Z_t moves: X_t and Q_t can then move in opposite directions.

Since this aspect of our model is relatively novel, it deserves further elaboration. Other models have typically assumed that home agents consume a domestic aggregate and a foreign aggregate (such as C^* in our model), and that there is some home bias, so that PPP does not hold. In contrast, we assume that home agents do not consume the foreign aggregate but instead a different good (food). This would not make a difference if the relative price of food were fixed in terms of the foreign aggregate (e.g. if Z were constant). So the main differences between our model and previous ones emerge because Z is allowed to fluctuate.

In particular, the standard specification implies, as just discussed, a very tight link between the terms of trade and the real exchange rate: with constant Z , X_t and Q_t must *always* move in the same direction. Using hatted lowercase variables for log deviations of variables from steady state, it turns out that $\hat{x}_t = (1 - \alpha)\hat{q}_t$ to a first order approximation, so that (to second order) $Var(\hat{x}_t) = (1 - \alpha)^2 Var(\hat{q}_t)$: if $\alpha < 1$, the variance of the real exchange rate is proportional to and strictly less than the variance of the terms of trade. These implications seem quite restrictive.

In our model, in contrast, $\hat{x}_t = (1 - \alpha)\hat{q}_t - \hat{z}_t$ to first order. We then see that fluctuations in the relative price of food mean that \hat{x}_t and \hat{q}_t can move in opposite directions (in response of shocks to \hat{z}_t). This is more likely to be the case if the economy is very open (i.e. if α is large) or if food prices are very volatile (i.e. if the typical size of \hat{z}_t is large). We also see that the variance of \hat{x} can be larger than the variance of \hat{q} , depending on the volatility of \hat{z} .

As we will see, the volatility of food prices is also a crucial factor in the analysis of monetary policy rules. So this model suggests that a negative correlation between the real exchange rate and the terms of trade goes hand in hand with drastic changes in policy evaluation. This is natural in our model, as both aspects of the analysis reflect the impact of food price shocks.

⁵ When calibrating the model, we make the stronger assumption that shocks to C^* and Z are independent. To justify this, one can assume that food has a negligible share in the world consumer basket, in contrast with the domestic basket. This is a defensible assumption since the share of food in the CPI is substantially higher in small emerging economies than in advanced countries.

2.3. Domestic Production

Domestic production follows Gali and Monacelli (2005), Gali (2008) and others, so we refer to those sources for brevity. The home final good Y_t is a Dixit-Stiglitz aggregate of intermediate goods varieties. Cost minimization then implies that the demand for each variety $j \in [0, 1]$ is given by $Y_t(j) = \left(\frac{P_t(j)}{P_{ht}}\right)^{-\varepsilon} Y_t$, where ε is the elasticity of substitution between domestic varieties, $P_t(j)$ is the price of variety j and P_{ht} is the relevant price index (the PPI). Each variety j is produced with only labor of type j according to the production function $Y_t(j) = A_t N_t(j)$, where $N_t(j)$ is the input of type j labor and A_t is a productivity shock, common to all firms in the economy.

Firms take wages as given. We allow for the existence of a subsidy to employment at constant rate v . Hence nominal marginal cost is given by $\Psi_{jt} = (1 - v)W_t(j)/A_t$, where $W_t(j)$ is the wage rate for type j labor.

Variety producers are monopolistic competitors and set prices in domestic currency following a Calvo protocol: each individual producer is allowed to change nominal prices with probability $(1 - \theta)$. All producers with the opportunity to reset prices in period t choose the same price, say \bar{P}_t , reflecting desired current and future markups over marginal costs (see Gali 2008 for a discussion). The price of the home final good is then given by $P_{ht} = \left[(1 - \theta)\bar{P}_t^{1-\varepsilon} + \theta P_{h,t-1}^{1-\varepsilon}\right]^{1/(1-\varepsilon)}$.

2.4. Market Clearing

We assume that the foreign demand for the domestic aggregate is given by a function of its price relative to P_t^* and the index C_t^* of world consumption. Hence market clearing for the home aggregate requires:

$$Y_t = (1 - \alpha) (P_{ht}/P_t)^{-\eta} [C_t + \varpi \Phi(C_t, (P_{ht}/P_t) Y_t)] + \varkappa \left(\frac{P_{ht}}{S_t P_t^*}\right)^{-\gamma} C_t^* \quad (7)$$

where \varkappa is a constant and γ is the price elasticity of the foreign demand for home exports, which is allowed to differ from the domestic elasticity for the home goods, η . The first term in the right hand side is the domestic demand, inclusive of financial transfer costs, for the domestic aggregate; it uses the fact that, in equilibrium, nonfinancial household income equals the value of domestic production, that is, $H_t = (P_{ht}/P_t) Y_{ht}$.

Once a rule for monetary policy is specified, the model can be solved for the equilibrium home output, consumption, and relative prices. We discuss implications in turn.

3. Linear Approximation and Implications

A log linear approximation of the model around its nonstochastic steady state is described in Table 1. As mentioned, we use hatted lowercase variables to denote log deviations from steady state. Equation (L1) is the approximation of the risk sharing condition (3), with $\psi = \sigma/(\varpi + \sigma)$ indicating the degree of risk sharing ($\psi = 1$ denotes perfect risk sharing and $\psi = 0$ portfolio autarky). Equations (L2) and (L3) are linearized versions of (6) and (7). Finally, (L4) summarizes linear versions of Calvo pricing equations: it is the now familiar Phillips Curve relationship between PPI inflation ($\pi_{ht} = \log P_{ht} - \log P_{ht-1}$), its expected future value, and marginal costs (the term in brackets).

[PLEASE INSERT TABLE 1 HERE]

Table 1 also displays the corresponding equations under flexible prices, whose solutions yield the "natural" values, identified by an n superscript. The natural values are linear transformations of the exogenous shocks. Finally, the table shows the equation system in terms of "gaps" or log deviations from natural values, identified with tildes.

The equations in Table 1 yield the solution of the model, up to a linear approximation, once monetary policy is specified. They also allow us to identify how conventional analysis can be modified to accommodate the particular features of our specification.

3.1. Aggregate Supply and Demand

Key to the model's transmission mechanism is the relationship between the output gap and international relative prices, as summarized by the terms of trade q_t or the real exchange rate x_t . Inserting (G1) and (G2) into (G3) yields:

$$\tilde{x}_t = (1 - \alpha)\tilde{q} = \Theta \tilde{y}_t \quad (8)$$

where

$$\Theta = \frac{(1 - \alpha)[1 - \omega(1 - \psi)]}{\omega[\alpha(\eta - \psi/\sigma) - (\gamma - \psi/\sigma)] + \alpha\omega(1 - \psi) + \gamma} \quad (9)$$

and ω is the steady state ratio of domestic expenditure on home goods to home output. To interpret, consider the case of perfect risk sharing ($\psi = 1$) and of equal price elasticities of demand at home and abroad, $\eta = \gamma$. Then $\Theta = (1 - \alpha)/[\eta - \omega(1 - \alpha)(\eta - 1/\sigma)]$; if, further, $\eta = 1/\sigma = 1$, as emphasized in the literature, $\Theta = (1 - \alpha)$. These expressions resemble those studied in previous work, especially in highlighting the importance of the difference between η and $1/\sigma$ for the output response to international relative price shocks (as discussed by Corsetti et al. 2010 and others).

Our derivation of Θ in (9) extends the previous intuition in two directions: it highlights that frictions to full risk sharing ($\psi < 1$) introduce a wedge in the "gap" between η and $1/\sigma$; and it shows that the home response of output to relative prices changes change with the export elasticity γ .

Using (8) and combining (G1)-(G4) one then obtains a New Keynesian Phillips Curve:

$$\pi_{ht} = \chi \tilde{y}_t + \beta E_t \pi_{ht+1},$$

with slope

$$\chi = \lambda \left[\varphi + \left(\psi + \frac{\alpha}{1 - \alpha} \right) \Theta + \sigma(1 - \psi) \left(1 - \frac{\alpha}{1 - \alpha} \Theta \right) \right]$$

While the form of the resulting Phillips Curve is standard, our analysis shows that its slope depends on the degree of international capital mobility (as parameterized by ϖ , which affects χ both directly and through Θ), as well as on the elasticities η and γ (which affect Θ and hence χ). Food price shocks, in turn, affect natural output and hence the output gap \tilde{y} , as implied by the equations in Table 1.

The aggregate demand side can be summarized by a dynamic IS curve: log linearizing the Euler equation and noting that the definition of the CPI implies that $\Delta p_{ht} = \log(P_{ht}/P_{ht-1}) - \log(P_t/P_{t-1}) = \pi_{ht} - \pi_t = -\frac{\alpha}{1-\alpha} \Delta(\hat{x}_t + \hat{z}_t^*)$, one obtains:

$$\tilde{y}_t = -\Lambda [i_t - E_t \pi_{ht+1} - r_t^n] + E_t \tilde{y}_{t+1} \quad (10)$$

where $\Lambda = \frac{1-\alpha}{\Theta}$, and the natural interest rate is:

$$r_t^n = E_t [\sigma \Delta \hat{c}_{t+1}^* + \frac{1}{1-\alpha} (\alpha \Delta \hat{z}_{t+1} + \Delta x_{t+1}^n)] = E_t [\sigma \Delta \hat{c}_{t+1}^* + \Delta q_{t+1}^n - \Delta \hat{z}_{t+1}] \quad (11)$$

the last step having made use of $x_t^n = (1 - \alpha)q_t^n - z_t$.

Two features of the resulting IS curve are worth highlighting. One is that financial market frictions affect the coefficient Λ , which depends on Θ and hence on the market incompleteness parameter ψ . The other is that shocks to food prices affect output through the real natural interest rate – in particular, the latter will change with expected changes in the natural terms of trade (q_t^n) and in the world terms of trade z_t .

3.2. Impulse Responses

To shed light on the dynamic responses of this economy to world food price shocks, abstract from other shocks by setting $\hat{c}_t^* = a_t = 0$, and assume that \hat{z}_t follows an AR(1) process with persistence parameter $|\rho_z| < 1$.

For concreteness, here we assume a standard Taylor rule on PPI inflation⁶:

$$i_t = \phi_y \tilde{y}_t + \phi_\pi \pi_{ht} \quad (12)$$

This rule, together with the Phillips Curve and the dynamic IS curve, can then be solved for the output gap, domestic inflation, and the interest rate as functions of the natural interest rate. In particular⁷, $\pi_{ht} = \chi \Lambda_v r_t^n$ and $\tilde{y}_t = (1 - \beta \rho_z) \Lambda_v r_t^n$, where

$$\Lambda_v = \frac{1}{[\frac{1-\rho_z}{\Lambda} + \phi_y](1 - \beta \rho_z) + \chi(\phi_\pi - 1)}$$

Here we have used the fact that the natural interest rate can be written as a linear function of the food price shock: $r_t^n = \phi E_t \Delta \hat{z}_{t+1} = -\phi(1 - \rho_z) \hat{z}_t$ for a constant ϕ described below, so that r_t^n has the same autocorrelation ρ_z as z . The

⁶To keep the notation compact, we omit the constant in the policy rule (which equals the real interest rate in steady state). This is consistent with assuming zero inflation and zero world interest rate in steady state. Since we are abstracting from other than food price shocks, we also abstract from stochastic shocks to policy.

⁷This is a straightforward exercise in undetermined coefficients. See e.g. Gali (2008), subsection 3.4.1.

1 same observation allows us to rewrite the solutions for inflation and the output gap as functions of z_t :

$$\begin{aligned}\pi_{ht} &= -\chi\Lambda_v\phi(1-\rho_z)\hat{z}_t \\ \tilde{y}_t &= -(1-\beta\rho_z)\Lambda_v\phi(1-\rho_z)\hat{z}_t\end{aligned}\tag{13}$$

2 This is a closed form solution that fully characterizes the responses of the model to food price shocks. Notably, since
3 $\chi > 0$, home inflation and the output gap will always move in the same direction in response to z_t . This can be up or
4 down, however, depending on elasticities. To establish the direction and respective magnitudes, we need to compute ϕ ,
5 that is, how the natural rate of interest responds to z . From the natural system of equations one obtains:

$$\phi = -\frac{\alpha\{(\sigma-1)(1-\psi) + \omega\tau\}}{\alpha\{(\sigma-1)(1-\psi) + \omega\tau\} - \gamma(1-\omega)(\varphi + \sigma(1-\psi)) - \psi(1 + \omega\varphi/\sigma)}$$

where

$$\tau = [(1-\psi)(\varphi + 1 - \eta\sigma) - \varphi(\eta - \psi/\sigma)]$$

6 This shows that ϕ can be positive or negative depending on parameter values. But the expression is complex and
7 difficult to interpret directly. We can gain intuition, however, by examining the extremes of perfect risk sharing and
8 portfolio autarky. Given (L1), any intermediate case is a convex combination of those two.

9 With perfect risk sharing: $\psi = 1$ and ϕ simplifies greatly:

$$\phi = \frac{-\alpha\omega\varphi(\eta - 1/\sigma)}{1 + \varphi[(1-\omega)\gamma + \omega(\frac{1-\alpha}{\sigma} + \alpha\eta)]} \quad \text{if } \psi = 1$$

10 It is now easy to spot the critical role of the relative values of η and σ in shaping the economy's response to z shocks.
11 If $\eta = 1/\sigma$, a parameterization that is not too unrealistic in our model, $\phi = 0$: the natural interest rate does not move
12 at all with the z shock. Likewise, (13) reveals that the output gap and domestic inflation do not move either. The
13 assumed PPI Taylor rule then prescribes that the nominal interest rate does not change.

14 The terms of trade and the real exchange rate do change and, notably, in opposite directions. The terms of trade
15 depreciate in proportion to z , as can be readily inferred from (11) and the fact that r^n does not react to z when $\phi = 0$.
16 From $\hat{x}_t = (1-\alpha)\hat{q}_t - \hat{z}_t$, the real exchange rate appreciates by $\alpha\hat{z}_t$. Under perfect international risk sharing and given
17 world consumption, domestic consumption declines *pari passu* with the real appreciation. These responses are illustrated
18 by the dotted green line in Figure 2, which has $\eta = 1/\sigma = 0.5$ (and other parameter values set as in subsection 5.1.
19 below).

20 **[PLEASE INSERT FIGURE 2 HERE]**

21 Consider now the case of $\eta > 1/\sigma$, still maintaining complete risk sharing ($\psi = 1$). Now ϕ is negative and r_t^n increases
22 with a positive z shock. All else constant, the right hand side of (10) is higher, reflecting that the real interest rate
23 falls below the natural interest rate. This induces an increase in aggregate demand, leading to an increase in domestic
24 inflation and the output gap, as given by (13).

25 The rationale is that domestic and imported goods are substitutes when $\eta > 1/\sigma$. Hence demand for home goods
26 at unchanged relative prices increases with higher imported food prices. Under PPI inflation targeting, the nominal
27 interest rate must go up and, since $\phi_\pi > 1$ (as required by the Taylor Principle), the real interest rate also rises. This in
28 turn dampens output and home inflation. The gradual decline of r_t^n entailed by its AR(1) dynamics, coupled with the
29 forward-looking behavior of the output gap and inflation implied by the Phillips curve and the dynamic IS, determine
30 that convergence in output and home inflation to pre-shock levels is gradual.

31 Since r^n increases, (11) implies that that q^n must undershoot Δz , so the natural terms of trade deteriorate by less
32 than if $\eta = 1/\sigma$. But the rise in the output gap and home inflation are not enough to fully offset the terms of trade
33 deterioration on impact. Meanwhile, the real exchange rate still appreciates. So, again, the terms of trade and the real
34 exchange rate move in opposite directions. Given full risk sharing, consumption falls by more than with $\eta = 1/\sigma$. These
35 responses are depicted by the dotted lines of Figure 2. The reasoning is the opposite if $\eta < 1/\sigma$ (bold line in Figure 2).

36 Now consider the opposite case of portfolio autarky. Then $\psi = 0$ and the response of the natural interest rate to z
37 is given by:

$$\phi = -\frac{(\sigma-1) + \omega[\varphi(1-\eta) - \sigma(\eta - 1/\sigma)]}{(\sigma-1) + \omega[\varphi(1-\eta) - \sigma(\eta - 1/\sigma)] - \gamma(\varphi + \sigma)(1-\omega)/\alpha} \quad \text{if } \psi = 0$$

38 This is more complex than under perfect risk sharing, but still yields useful insights. Importantly, ϕ does not become

zero even if $\eta = 1/\sigma$, as long as these elasticities are not exactly equal to one.⁸ In particular, if $\eta \leq 1/\sigma$ and $\sigma > 1$, and γ is not too large, ϕ becomes strictly negative; this gives a strictly positive response of output and home inflation to the shock even when $\eta = 1/\sigma$. Further, the lower η , the stronger the response. This is illustrated in Figure 3.

[PLEASE INSERT FIGURE 3 HERE]

The intuition is that, under significant frictions to international risk sharing, lower substitutability between the domestic and foreign goods implies that, in response to a positive z_t shock, domestic agents must produce and export more (in quantity terms) of the domestic good to maintain pre-shock consumption levels. Given that foreign demand for the home good is not perfectly elastic, this causes a deterioration of the terms of trade, so q_t overshoots z_t , and y_t^n and $(y_t - y_t^n)$ both rise. In contrast, with perfect risk sharing, the economy receives an insurance payment from abroad in response to the shock. This payment is intended to stabilize the marginal utility of domestic consumption and, hence, is bigger the smaller η : complete financial markets effectively allow the trade balance to turn negative as the world terms of trade turn against the small open economy – and sharply so if η is very small.⁹ As domestic demand for the home good is stabilized, its supply in world markets does not increase by as much, shoring up the world price of home exports and thus preventing further terms of trade deterioration.

A comparison between Figures 2 and 3 reveals that consumption drops by more than under perfect risk sharing. The sharper drop in consumption and the rise in output then imply that the fall in the ratio of consumption to labor effort and of welfare are also steeper.

Also in contrast with the perfect risk sharing case, the terms of trade and the real exchange rate can now display some positive covariance when $\eta < 1$. This follows from the fact that $\hat{x}_t = (1 - \alpha)\hat{q}_t - \hat{z}_t$ to first order, and that the terms of trade response to \hat{z} is larger than under perfect risk sharing. This reemphasizes that the model can deliver various covariance patterns between those two variables, depending on parameterization. Finally, with less than perfect risk sharing, the export elasticity parameter γ is also important for the magnitude and sign of the responses of output and inflation to the z_t shock.

4. Welfare and Policy Trade-Offs

Our model economy is not completely "small", since it produces a differentiated good facing a downward sloping world demand. Nominal rigidities imply that monetary policy can affect the price of the home aggregate in terms of world consumption. Hence policy choices must take into account not only the domestic distortions caused by inflation but also international relative price effects (known as *terms of trade externalities*), as known since Corsetti and Pesenti (2001).

This section discusses how this trade-off plays out in our model. Following Faia and Monacelli (2008), the first subsection compares the solution of the social planner's or *Ramsey* problem against the *natural* allocations that would emerge in a flexible price competitive equilibrium (and hence PPI targeting). The second subsection follows Benigno and Woodford (2006) and Benigno and Benigno (2006) in deriving a purely quadratic approximation of the representative agent's welfare. The latter can then be exploited to characterize optimal policy as a targeting rule involving inflation and output.

4.1. The Ramsey Solution versus the Natural Outcome

The market clearing condition for home goods, (7), can be written as:

$$\begin{aligned} Y_t &= A_t N_t = (1 - \alpha) (P_{ht}/P_t)^{-\eta} [C_t + \varpi \Phi(C_t, (P_{ht}/P_t) Y_t)] + \varkappa X_t^\gamma (P_{ht}/P_t)^{-\gamma} C_t^* \\ &\equiv \Theta(Y_t, C_t, (P_{ht}/P_t), X_t) \end{aligned} \quad (14)$$

since $H_t = (P_{ht}/P_t) Y_t$. Equation (14) must hold at all times and is a key constraint for the planner's choices of consumption, leisure, and the real exchange rate.

Another key constraint is the international risk sharing equation (3), which amounts to

$$C_t^\sigma [1 + \varpi \Phi(C_t, H_t)] = \kappa X_t (C_t^*)^\sigma \quad (15)$$

⁸This is consistent with our earlier discussion in that risk sharing imperfections become irrelevant in this model if elasticities are unitary, up to first order. That remains true up to second order under an appropriate tax subsidy that balances out the monopolistic competition and the terms of trade externality distortions in steady state.

⁹The trade balance dynamics is not plotted to save on space and preserve the readability of the figures but is available from the working paper version.

Finally, recall that (6) gives P_{ht}/P_t , the real price of home output, a function of the real exchange rate and the food price shock, say $g(X_t, Z_t)$. Given this, the Ramsey problem is then to maximize $u(C_t) - v(N_t)$ subject to (14) and (15). Remarkably, the Ramsey problem is static: it has the same form at each date, in each state. The Technical Appendix shows that the Ramsey optimality condition can be written as:

$$\frac{C_t u'(C_t)}{N_t v'(N_t)} = \frac{\sigma + s_t \epsilon_{C_t}^{\Phi_C} + \vartheta_t \epsilon_{C_t}^{\Theta}}{\vartheta_t [1 - \epsilon_{Y_t}^{\Theta}] - s_t \epsilon_{H_t}^{\Phi_C}} \quad (16)$$

where $s_t = \varpi \Phi_{C_t} / (1 + \varpi \Phi_{C_t})$ and $\vartheta_t = \left(1 - s_t \epsilon_{H_t}^{\Phi_C} \cdot \epsilon_{X_t}^g\right) / (\epsilon_{X_t}^{\Theta} + \epsilon_{p_{ht}}^{\Theta} \cdot \epsilon_{X_t}^g)$, with $\epsilon_{X_t}^{\Theta}$ denoting the partial elasticity of Θ_t with respect to X_t , $\epsilon_{p_{ht}}^{\Theta}$ the partial elasticity of Θ_t with respect to P_{ht}/P_t , etc.¹⁰

To interpret the preceding expression, consider the perfect risk sharing case: $\varpi = 0$. Then the optimality condition reduces to:

$$C_t u'(C_t) = N_t v'(N_t) \{ [\epsilon_{X_t}^{\Theta} + \epsilon_{p_{ht}}^{\Theta} \cdot \epsilon_{X_t}^g] \sigma + \epsilon_{C_t}^{\Theta} \}$$

This just says that the Ramsey planner equates the utility benefit of a one percent increase in consumption (the LHS of the preceding equation) to its cost in terms of increased labor effort (the RHS). To understand the latter, notice that a one percent increase in consumption increases directly the demand for home output by $\epsilon_{C_t}^{\Theta}$. But perfect risk sharing requires that a one percent increase in consumption be coupled with a real depreciation of σ percent. This, in turn, is associated with an increase in demand of $[\epsilon_{X_t}^{\Theta} + \epsilon_{p_{ht}}^{\Theta} \cdot \epsilon_{X_t}^g] \sigma$ percent.

If $\varpi > 0$, the intuition remains the same, except that, if there is to be a one percent increase in consumption, the risk sharing condition (15) can now be satisfied by changing output instead of the exchange rate. The planner then chooses an optimal mix of output and exchange rate adjustment, taking into account the impact on world demand and labor effort. This accounts for the extra terms in the optimality condition.

The Ramsey allocation is then given by (14), (15), and (16). This system depends on exogenous shocks, so the Ramsey solution (C_t, N_t, X_t) is stochastic and generally time varying. In particular, the different elasticities in the RHS of the optimality condition (16) are generally time varying and, more crucially, summarize the role of the different elasticities of demand and substitution in the model. These elasticities, in fact, determine the incentives for the planner to exploit the "terms of trade externality". To see this in the case of perfect risk sharing, note that in that case a one percent depreciation always increases consumption by $1/\sigma$ percent, but the size of the associated increase in labor effort, with the resulting cost, is smaller or larger depending on $[\epsilon_{X_t}^{\Theta} + \epsilon_{p_{ht}}^{\Theta} \cdot \epsilon_{X_t}^g]$. This implies, *a fortiori*, that the relative desirability of different policy rules will depend on the interplay between elasticities and how much each rule attempts to exploit the terms of trade externality.

Our characterization of the Ramsey solution is most useful to evaluate the optimality of the *natural* allocation, and hence of perfect PPI stabilization. In any flexible price market equilibrium, prices are set as a markup over marginal cost, $P_{ht} = \mu MC_t = \mu(1 - v)W_t/A_t = \mu(1 - v)(W_t/P_t)/g(X_t, Z_t)A_t$. And since $W_t/P_t = v'(N_t)/u'(C_t)$ we get:

$$\frac{C_t u'(C_t)}{N_t v'(N_t)} = \frac{\mu(1 - v)C_t}{g(X_t, Z_t)A_t N_t} \quad (17)$$

The natural allocation is therefore pinned down by (14), (15), and (17). Hence the Ramsey allocation and the natural allocation will, in general, differ because (and only because) the Ramsey optimality condition (16) and the markup condition (17) are not the same. The basic difference is, in fact, the terms of trade externality: the Ramsey planner takes into account the impact of its policies on the real exchange rate, while the natural allocation ignores that impact. To see this, assume that $\mu(1 - v) = 1$. Then the preceding expression for the natural allocation reduces to $v'(N_t) = u'(C_t)A_t g(X_t, Z_t)$, which is easily seen to be the optimal labor choice condition for a planner that takes the relative price $g(X_t, Z_t) = P_{ht}/P_t$ as given.

Hence the optimality of PPI, a mainstay of the literature, hinges on how far apart the Ramsey and natural allocations can be. Again, this will depend on the parameters underlying the different elasticities in (16) and (17). This perspective clarifies many of the results in the literature. For example, under perfect risk sharing, if $\eta = \gamma = 1/\sigma$, (16) and (17) coincide exactly provided that $\mu(1 - v) = 1 + \varkappa/(1 - \alpha)$. Under the additional assumption $\alpha = \varkappa$, this is Gali and Monacelli's (2005) condition for PPI stabilization to be optimal. Clearly, however, this is a very special case.

¹⁰Explicit expressions for the elasticities are given in the Technical Appendix.

4.2. Optimal Feasible Outcomes

A Ramsey planner is constrained only by the market clearing condition (14) and the risk sharing condition (15). The Ramsey solution, therefore, may not be available to a central bank that must also take nominal rigidities as given. In other words, optimal feasible allocations are constrained also by the Calvo pricing equations.

To tackle this problem, we derive a quadratic approximation to the representative household's utility function and then maximize the resulting objective subject to the linearized equilibrium equations described in Table 1. As discussed in Woodford (2003), for that procedure to be correct the quadratic approximation to utility must contain no linear terms. This can be achieved by deriving a second order approximation to some structural equations of the model, and using them to eliminate any linear term in the second order approximation to utility.

In our model, the Technical Appendix adapts the techniques of Benigno and Woodford (2006) and Benigno and Benigno (2006) to derive a second order approximation to the utility of the representative agent of the form

$$-E \sum_{t=0}^{\infty} \beta^t \left[\frac{1}{2} v_t' L_v v_t + v_t' L_e e_t + \frac{1}{2} l_\pi \pi_{ht}^2 \right]$$

plus a term independent of policy, where $v_t = (\hat{y}_t, \hat{c}_t, \hat{p}_{ht}, \hat{x}_t)'$ is a vector of endogenous variables, $e_t = (\hat{a}_t, \hat{c}_t^*, \hat{z}_t^*)'$ is the vector of shocks, L_v and L_e are matrices, and l_π is a scalar. The entries of L_v , L_e and the value of l_π are functions of the underlying parameters of the model. The optimal policy problem is then to maximize the previous objective subject to the linearized equations given in Table 1. One can gain further insight, however, by recognizing that one can use three of the linearized equations to express the vector v_t in terms of only one of its components and the vector of shocks. In particular, the linearized equations imply that v_t can be written as

$$v_t = N \hat{y}_t + N_e e_t \quad (18)$$

for some straightforward matrices N and N_e . Inserting in the utility function, rearranging, and ignoring terms independent of policy, the objective can then be rewritten as:

$$-E \sum_{t=0}^{\infty} \beta^t \left[\frac{1}{2} l_y (\hat{y}_t - \tilde{e}_t)^2 + \frac{1}{2} l_\pi \pi_{ht}^2 \right] \quad (19)$$

where $l_y = N' L_v N$ is a scalar and $\tilde{e}_t = -(1/l_y) (N' L_v N_e + N' L_e) e_t$ is a linear combination of shocks. \tilde{e}_t can then be seen as an appropriate *output target*. The weights l_y and l_π measure the appropriate tradeoff between inflation and output stabilization.

The relevant constraint for the policy problem is then obtained by inserting the previous representation of v_t into the linearized New Keynesian Phillips Curve to obtain

$$\pi_{ht} = \lambda_y \hat{y}_t + \lambda_e e_t + \beta E_t \pi_{ht+1} \quad (20)$$

where λ_y and λ_e are functions of the parameters of the model, as described in the Technical Appendix.

In order to define optimal policy, we need to take a stand about the commitment possibilities open to the central bank in order to deal with the possibility of time inconsistency. We focus on the optimal policy under the "timeless perspective" advocated by Woodford (2003). This reduces, in our problem, to the maximization of (19) subject to (20), taking π_0 as given (which is a form of limited commitment).¹¹

The resulting first order condition for optimality can be written as:

$$\pi_{ht} + \mathcal{W}_y (\Delta \hat{y}_t - \Delta \tilde{e}_t) = 0 \quad (21)$$

where $\mathcal{W}_y = l_y / \lambda_y l_\pi$ is the weight on the output term and Δ is the difference operator. This can be interpreted as a version of "flexible inflation targeting": it combines targeting zero inflation and the change in output around the target $\Delta \tilde{e}_t = \tilde{e}_t - \tilde{e}_{t-1}$.

This formulation shows how optimal monetary policy can be analyzed in basically the same way as in the closed economy of, say, Benigno and Woodford (2006). The open economy aspects of the model affect the rule in at least two ways: through the weight \mathcal{W}_y assigned to output targeting and also via the definition of the output target \tilde{e}_t .

¹¹See Woodford (2003) or Benigno and Woodford (2005) for a discussion of optimality from a timeless perspective. The optimal policy under the timeless perspective then depends on π_0 , and becomes the optimal policy under commitment if π_0 is, in turn, chosen optimally. For a full argument in a related model, see Chang (1998).

To illustrate, the upper panel of Table 2 gives the relative output weight \mathcal{W}_y under different assumptions on parameters and capital market imperfections. International financial markets are parameterized by the values of $\psi = 1$ (perfect risk sharing) down to $\psi = 0$ (portfolio autarky), displayed in the first column of the panel. The second column corresponds to a case in which all relevant elasticities are equal to one. In that case, emphasized by Gali and Monacelli (2005) and others, $\mathcal{W}_y = 1/5$ regardless of the degree of international capital mobility.

[PLEASE INSERT TABLE 2 HERE]

The third column assumes unit elasticities of demand, except for $\eta = 0.25$. As discussed in more detail in subsection 5.1. below, this is a good assumption for items such as food, which are not easily substitutable with other goods. We see that, in that case, the emphasis on output relative to inflation in the optimal targeting rule is somewhat larger than in the unit elasticity case. In addition, the output weight increases as international capital markets become less perfect, except close to the limit case of portfolio autarky. Anticipating some of our results in the next section, this suggests that an interest rate rule that depends on output in addition to inflation will become more valuable in this case if risk sharing is less perfect.

The last column in the panel assumes not only that $\eta = 0.25$ but also that $\gamma = 5$. In other words, it looks at the consequences of assuming that the price elasticity of world demand for exports is large. In this case, we see that the optimal targeting rule places much more emphasis on output than if $\gamma = 1$, under perfect financial markets ($\psi = 1$): \mathcal{W}_y becomes almost two. In this case, the result suggest, that targeting domestic inflation may be not as desirable as targeting other variables. However, the value of \mathcal{W}_y falls substantially as ψ decreases. Under portfolio autarky, in fact, the value of \mathcal{W}_y is very close to $1/5$, its value in the unit elasticity case.

To illustrate the influence of food price shocks on the targeting policy, the lower panel of Table 2 displays the values of the coefficient of the food price shock z_t^* in the definition of the output target \tilde{e}_t . In the case of unit elasticities, the coefficient is always zero. This in fact confirms the results of Gali and Monacelli (2005) and Gali (2008), as it says that optimal policy does not need to react to food price shocks at all. But this is, of course, a special result. In the case of low η , the coefficient is negative, indicating that an increase in the world relative price of food reduces the output target. The target rule then prescribes that it is optimal for policy to respond with less domestic inflation and lower actual output. The reaction is stronger as ψ falls, moving away from perfect risk sharing.

In the case of $\eta = 0.25$ and $\gamma = 5$, the coefficient of z_t^* is again negative. Its absolute value is large under perfect risk sharing, and falls quite substantially as ψ approaches zero. Under perfect risk sharing the output target is more sensitive to food price shocks if the price elasticity of the world demand for exports is larger than one, but the opposite inference can be easily drawn if risk sharing is imperfect.

Before closing this section, it is worth mentioning that the optimal targeting policy just discussed is one of many available. In particular, the vector v_t was summarized by (18) in terms of output and exogenous shocks, but we could have equally expressed v_t in terms of any of its components, say the real exchange rate or consumption. This would have led to an optimal targeting rule in terms of inflation and that other variable. In other words, the policy prescriptions discussed here are optimal but not unique.¹² It may be the case that alternative representations may be advantageous in terms of other grounds. For example, some may argue that an optimal target rule that includes the real exchange rate in addition to domestic inflation and output is superior in terms of intuition and transparency. That argument, however, remains to be spelled out and is beyond our paper.

5. Welfare Properties of Simple Rules

Besides its implications for optimal policy, as characterized in the previous section, our model is useful to examine the welfare properties of monetary policy rules often seen in practice. In this section we calibrate the model and compare the PPI Taylor rule (12), which in the context of our model amounts to the so called *core* inflation targeting rule, against three well known alternatives: a *headline* CPI-based inflation targeting rule (same as (12), but with π_t replacing π_{ht}), a *CPI forecast* rule (with $E_t\pi_{t+1}$ instead of π_{ht} in (12)), and an exchange rate peg.

5.1. Calibration

We calibrate the model to a quarterly frequency and assume that all shocks follow AR(1) processes. From regressions of the IMF global index of food commodity prices relative to the US WPI between 2000Q1 and 2011Q4, we set the

¹²This explains, in particular, the differences between our targeting discussion and that of Di Paoli (2009). Di Paoli expresses v_t in terms of output and the real exchange rate. This naturally leads her to conclude that it is optimal to target both variables in addition to inflation. This is correct but is not the only optimal targeting procedure.

standard deviation of the z shock to five percent and the AR persistence coefficient to 0.85. For productivity shocks, we set the standard deviation at 1.2 percent and the persistence parameter at 0.7. Both are consistent with estimates for Chile, a typical small open economy for which suitably long data series exists, and also with Gali and Monacelli's (2005) estimates for Canada, once differences in output volatility between Chile and Canada are adjusted for. For interest rate shocks, we set a standard deviation of 0.62 and a persistence parameter of 0.6, based on our own estimates of the Taylor rule for Chile on 1991-2008 data.

The transfer cost parameter ϖ is calibrated from $\psi = \sigma/(\sigma + \varpi)$, so that $\psi \in [0, 1]$. Besides the two extremes, we also consider $\psi = 0.9$ which is consistent with Schulhofer-Wohl's (2011) closed economy estimates. We also explored lower values but found that the critical range is in the interval $[0.9, 1]$. This suggests that even small departures from full risk sharing can have significant implications for welfare-based policy comparisons.

Regarding intra-temporal elasticities, while previous studies have assumed that $\gamma = \eta$, there is no compelling reason to impose the equality in a small open economy context, specially given the differences between imported goods (food) and exported goods (manufacturing/services) that motivate our model. Hence we allow γ to differ from η . Our baseline value of $\gamma = 5$ is consistent with estimates for manufacturing elasticities. The labor supply parameter φ is set to one in the baseline¹³. The ratio of home good consumption to income in steady-state (ω) is set to 0.66, consistent with food expenditure shares in GDP of around thirty percent.

We set $Y^* = C^* = 1$ as an obvious choice of numeraire. We set $\varepsilon = 6$, in line with the literature. We set ν so that $\varepsilon(1 - \nu)/(\varepsilon - 1) = 1 + \varkappa/(1 - \alpha)$. As discussed in subsection 4.1, this implies that when $\eta = 1/\sigma = \gamma$ the nonstochastic steady state is efficient.

We assumed that in steady state the representative household allocates about two thirds of time to leisure, trade is balanced, and $A = 1$. These assumptions imply that $C = Y = N = 0.33$ in steady state and that all relative prices are one. From the risk sharing condition, one can then obtain κ . The economy-wide market clearing equation then yields \varkappa , and the household first order condition for labor gives ς .

The coefficients of Taylor rules are calibrated as follows. Sticking to the baseline values of $\gamma = 5$ and $\sigma = 2$, we compute discounted utility values over a grid spanning from 1.25 to 3.05 (with 0.2 increments) for the coefficient on inflation (ϕ_π), and from 0 to 1.0 (with 0.125 increments) for the coefficient on the output gap (ϕ_y). This is done for each η and ψ under consideration. In each case, the pair (ϕ_π, ϕ_y) that maximizes discounted utility for each rule is then picked.

Finally, for the peg rule we need to specify the stochastic process for the world consumer price index. We assumed an AR(1) with considerable persistence ($\rho = 0.99$) and standard deviation of 1.3 percent, as obtained from a quarterly regression of an unweighted average of advanced countries (G-8) producer price indices during the 1990-2008 period.¹⁴

5.2. Welfare Results

In order to gain intuition on the relative performance of the simple policy rules, Figure 4 plots their responses to a food price shock. For comparison, the response associated with the optimal policy is also plotted. The figure assumes complete financial markets and $\eta = 1/\sigma = 0.5$: the low value of η is motivated by the price elasticity of food, while setting $\eta = 1/\sigma$ is a natural starting point since, as discussed in Section 3.2, intra- and inter-temporal substitution effects cancel each other out. In that case, optimal policy delivers a zero coefficient on the z component of target output and, as can be gleaned from equations (20) and (21), this implies a zero response of output and home inflation.

[PLEASE INSERT FIGURE 4 HERE]

Figure 4 emphasizes that the PPI rule coincides with the optimal policy in delivering a zero response of output and domestic inflation to the z shock. But the PPI rule also implies a larger appreciation of the real exchange rate than optimal and, by perfect risk sharing, a suboptimally large drop in consumption.

The CPI rule allows for output and home inflation to react to the food price shock. However, CPI stabilization requires, on impact, a large real appreciation and a drastic drop in output, the latter to mitigate the increase in the PPI. The appreciation translates into a large consumption drop on impact. The intuition is that stabilizing CPI inflation in the first period amounts to stabilizing the *level* of the CPI in that period. If the PPI were not to move, the only way to prevent an increase in the CPI is to engineer an exchange rate appreciation. The CPI rule limits the appreciation at the expense of letting the PPI increase some in the period of the shock.

Expected CPI targeting performs much better. Intuitively, targeting expected CPI inflation does not prevent an increase in the PPI on impact, and hence does not call for a real appreciation as large and under the headline CPI rule.

¹³We also considered $\varphi = 0$ and $\varphi = 3$, values usually found in the literature, but these choices did not alter the thrust of our results.

¹⁴Restricting estimation to the pre-2008 mitigates potential small sample biases due to the deflationary effects of the 2009-10 financial crisis but either way, our results are not critically affected by the choice of this estimation window.

As a consequence, output and domestic inflation rise on impact, but only trivially (0.08% per quarter). The benefit is a consumption response that is closest to the optimal rule relative to the other rules.

In short, in the case of Figure 4, PPI targeting replicates the optimal behavior of output and domestic inflation but delivers a larger real appreciation and consumption drop than optimal; expected CPI targeting does the opposite. Which rule delivers higher welfare will depend on the relative weights of consumption and leisure in utility. We show below, with all shocks in place, that expected CPI typically has an edge, which grows with η , i.e. as home goods and imports become more Edgeworth substitutable.

Figure 4 also indicates that the terms of trade and the real exchange rate co-vary negatively under the optimal policy. But we find that the covariance can be positive and, in general, depends on the policy rule.¹⁵ To understand why, recall that $\hat{x}_t = (1 - \alpha)\hat{q}_t - \hat{z}_t$ to first order. Hence the covariance is $E\hat{x}_t\hat{q}_t = (1 - \alpha)E\hat{q}_t^2 - E\hat{q}_t\hat{z}_t$. In the case of Figure 3, $E\hat{q}_t\hat{z}_t$ is positive, and dominates the term $(1 - \alpha)E\hat{q}_t^2$. But the sizes of $E\hat{q}_t\hat{z}_t$ and $E\hat{q}_t^2$ clearly depend on the policy rule.

Our earlier discussion emphasized that policy analysis depends on the severity of financial frictions, as given by the degree of international risk sharing. To examine this point, Figure 5 displays impulse responses for the same case as Figure 4, except that financial autarky is assumed. The expected CPI rule now delivers a consumption response quite far away from that of the optimal policy. The optimal policy no longer entails a flat response of home output and inflation to the z shock, but rather a response in between those of the PPI rule and the expected CPI rule. This suggests that the expected CPI rule is unlikely to be superior to the PPI rule under financial autarky.

[PLEASE INSERT FIGURE 5 HERE]

For further insight, Table 3 reports first and second moments of key observables over 10,000 random shock realizations for each of the three shocks in the model (the imported price shock, the productivity shock, and the monetary policy shock). This is done across policy rules and market structures, and also by adding the exchange rate peg rule to the menu of policy options.

[PLEASE INSERT TABLE 3 HERE]

Under perfect capital mobility, the expected CPI rule gets closer than the PPI rule to the optimal policy in terms of the variance of consumption. This is consistent with Figure 4. But the expected CPI rule also displays output and home inflation variability closer to optimal than the PPI rule. This reflects that Table 3 includes all shocks, while Figure 4 focuses on the response to only food price shocks.

The table then says that expected CPI targeting should dominate PPI targeting in terms of welfare, since the latter is determined by consumption, output/employment, and home inflation variability: for all of these observables, the outcomes of the expected CPI rule are closer to optimal than PPI targeting. A similar analysis implies that expected CPI targeting dominates headline CPI targeting. Finally, an exchange rate generates home inflation variability that is closest to the optimal rule, but loses badly to the other rules on other dimensions.

Table 3 also displays the corresponding statistics for $\psi = 0.9$ and $\psi = 0$ (portfolio autarky). It shows that, as international risk sharing becomes less perfect, the PPI rule progressively delivers outcomes closer to the optimal policy than other rules with regard to consumption, output and home inflation, becoming most clearly dominant under financial autarky.

To round up our discussion, Figure 6 reports the welfare rankings of the different rules, and examine its sensitivity to the "food-like" low elasticity assumption, letting η vary from the (very) low bound of 0.25 through 5. The other elasticities are kept at their baseline values. Each line in the figure reports the welfare difference, in terms of steady state consumption, between CPI targeting, expected CPI targeting, or an exchange rate peg, respectively, and PPI targeting (a positive value means that PPI is beaten by the competing rule)¹⁶.

[PLEASE INSERT FIGURE 6 HERE]

The upper panel of Figure 6 assumes complete markets. Confirming our discussion of impulse responses and moments, expected CPI targeting dominates PPI targeting across most of the relevant values of η . This panel also shows that

¹⁵In particular, this is the case for CPI targeting in the first two quarters following the shock. In that case, we have found that the resulting currency appreciation, combined with home price stickiness, can raise the foreign price of the domestic composite good, compensating for the negative impact of rising import prices on the terms of trade.

¹⁶Here we follow Schmitt-Grohe and Uribe (2007) and others, and report welfare measures conditional on starting at the nonstochastic steady state. Computing welfare values amounts to a simple addition of a control variable V_t to our system of non-linear equations, where V_t evolves according to the law of motion $V_t - \beta V_t = U(C_t, N_t)$.

headline CPI targeting and the exchange rate peg dominate PPI targeting as well if η is above two, and that the welfare advantages grow with η .

The middle panel of Figure 6 sticks to complete market and the same parameterization as the first panel except that the z shocks are assumed to be negligible. This is for comparability with previous work and to emphasize how our policy analysis depends on the properties of imported food prices. Notably, welfare gaps become noticeably narrower, thus suggesting that the smaller welfare gaps found in previous studies is partly due to the absence of z shocks in those models. More importantly for our discussion, the middle panel of Figure 6 confirms that expected CPI targeting is now welfare-dominated by PPI targeting across the η spectrum; it also says, however, that both rules are dominated by the headline CPI rule and the exchange rate peg rules once the intra-temporal elasticity is sufficiently large. The latter results are consistent with the theoretical claims of Faia and Monacelli (2008) as well as calibration results in Cova and Sondergaard (2004) and Di Paoli (2009). That the CPI rule or the peg can beat the PPI rule provides some rationale for the usual central bank practice of targeting headline CPI rather than the PPI. What is new here is that expected CPI targeting, a rule that was not considered in previous studies, appears superior to all others if food price shocks are realistically volatile.

Finally, the lower panel of Figure 6 reinstates the full menu of shocks but assumes that risk sharing is imperfect, even if seemingly slightly so ($\psi = 0.9$). The figure corroborates the results of Table 3 in restoring the welfare supremacy of PPI targeting, except again when η is (unrealistically for food) high. This proviso disappears, however, as ψ becomes smaller. Under portfolio autarky, PPI targeting easily beats the other rules.¹⁷

We have explored many other parameterizations, which we do not report to save space. The role of the size of the parameter ψ bears mentioning, however. While the analysis is starkest under the polar assumptions of perfect risk sharing and portfolio autarky, we have found that it changes quickly as ψ drops just below one. This suggests that only mild frictions to the financial technology can be quite significant, although a further examination is warranted. Finally, we have explored the roles of export price elasticity γ , the elasticity of labor supply ($1/\varphi$), and the price-stickiness parameter θ . Provided that the latter is kept within sensible ranges (above 0.3) so a New Keynesian setup remains meaningful, the export price elasticity γ has the greatest effect on welfare rankings. Yet, once risk sharing is incomplete ($\psi \leq 0.9$), the dominance of flexible PPI targeting is only marginally dented for alternative calibrations of these parameters and generally becomes stronger with lower γ .

6. Final Remarks

While our analysis places fewer restrictions on the economic environment than previous studies, it by no means includes all cases of interest. It may be necessary to ask, for example, what are the implications of assuming that food can be produced at home. While this extension is beyond the scope of the present paper, our analysis suggests the resulting policy prescriptions are likely to depend on export price elasticities and degree of capital mobility, both of which can be highly country specific.

One obvious avenue for future research is to estimate versions of the model discussed here, which would yield lessons for real world policy formulation. In this regard, our formulation of international risk sharing should prove especially useful, because of its tractability and parsimony.

References

- Benigno, G., Benigno, P., 2006. Designing Targeting Rules for International Monetary Cooperation. *Journal of Monetary Economics* 53, 473-506.
- Benigno, P., Woodford, P., 2006. Linear Quadratic Approximation of Optimal Policy Problems. NBER Working Paper 12672.
- Bodenstein, M., Erceg, C., Guerrieri, L., 2008. Optimal Monetary Policy with Distinct Core and Headline Inflation Rates. *Journal of Monetary Economics* 55, 518-533.
- Chang, R., 1998. Credible Monetary Policy in an Infinite Horizon Model: Recursive Approaches. *Journal of Economic Theory* 81, 431-61.
- Corsetti, G., Pesenti, P., 2001. Welfare and Macroeconomic Independence. *Quarterly Journal of Economics*, 116, 421-445.

¹⁷The corresponding figure looks very much like the lower panel of Figure 6 and is omitted here to save space.

- 1 Corsetti, G., Dedola, L., Leduc, S., 2010. Optimal Monetary Policy in Open Economies. In: Friedman B.M, Woodford,
2 M. (Eds.), *Handbook of Monetary Economics*, Vol. III, Elsevier, Amsterdam, 861-933.
- 3 Cova, P., Sondergaard, J., 2004. When should monetary policy target the exchange rate? Royal Economic Society
4 Annual Conference, Paper Number 51.
- 5 Di Paoli, B., 2009. Monetary policy and welfare in a small open economy. *Journal of International Economics* 77, 11–22.
- 6 Faia, E., Monacelli, T., 2008. Optimal Monetary Policy in a Small Open Economy with Home Bias. *Journal of Money,*
7 *Credit and Banking* 40, 721-750.
- 8 Galí, J., 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework.*
9 The MIT Press, Cambridge.
- 10 Galí, J., Monacelli, T., 2005. Monetary Policy and Exchange Rate Volatility in a Small Open Economy. *Review of*
11 *Economic Studies* 72, 707-734.
- 12 Kollmann, R., 2002. Monetary Policy Rules in the Open Economy: Effects on Welfare and Business Cycles. *Journal of*
13 *Monetary Economics* 49, 989-1015
- 14 Schmitt-Grohe, S., Uribe, M., 2007. Optimal Simple and Implementable Monetary and Fiscal Rules. *Journal of Monetary*
15 *Economics* 54, 1702-1725.
- 16 Schulhofer-Wohl, S., 2011. Heterogeneity and Tests of Risk Sharing. *Journal of Political Economy* 119, 925-958.
- 17 Sutherland, A., 2005. Incomplete Pass-through and welfare effects of exchange rate variability. *Journal of International*
18 *Economics* 65, 375-399.
- 19 Woodford, M., 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton University Press,
20 Princeton.

Table 1 : Linearized Model Equations

Log linear versions of structural equations:

$$\hat{c}_t = \psi \left[\frac{1}{\sigma} \hat{x}_t + \hat{c}_t^* \right] + (1 - \psi) [\hat{p}_{ht} + \hat{y}_{ht}] \quad (\text{L1})$$

$$0 = (1 - \alpha) \hat{p}_{ht} + \alpha(\hat{x}_t + \hat{z}_t) \quad (\text{L2})$$

$$\hat{y}_{ht} = -[\eta\omega + \gamma(1 - \omega)] \hat{p}_{ht} + \omega \hat{c}_t + \gamma(1 - \omega) \hat{x}_t + (1 - \omega) \hat{c}_t^* \quad (\text{L3})$$

$$\pi_{ht} = \lambda[\sigma \hat{c}_t + \varphi \hat{y}_{ht} - \hat{p}_{ht} - (1 + \varphi) \hat{a}_t] + \beta E_t \pi_{ht+1} \quad (\text{L4})$$

where $\psi = \frac{\sigma}{\sigma + \varpi}$ and $\lambda = \frac{1 - \theta}{\theta} \frac{1 - \beta\theta}{1 + \varphi\varepsilon}$.

Flexible Price ("Natural") Variables:

$$\hat{c}_t^n = \psi \left[\frac{1}{\sigma} \hat{x}_t^n + \hat{c}_t^{*n} \right] + (1 - \psi) [\hat{p}_{ht}^n + \hat{y}_{ht}^n] \quad (\text{N1})$$

$$0 = (1 - \alpha) \hat{p}_{ht}^n + \alpha(\hat{x}_t^n + \hat{z}_t) \quad (\text{N2})$$

$$\hat{y}_{ht}^n = -[\eta\omega + \gamma(1 - \omega)] \hat{p}_{ht}^n + \omega \hat{c}_t^n + \gamma(1 - \omega) \hat{x}_t^n + (1 - \omega) \hat{c}_t^{*n} \quad (\text{N3})$$

$$0 = \sigma \hat{c}_t^n + \varphi \hat{y}_{ht}^n - \hat{p}_{ht}^n - (1 + \varphi) \hat{a}_t^n \quad (\text{N4})$$

Deviations from Natural Variables ("Gaps "):

$$\tilde{c}_t = \psi \left[\frac{1}{\sigma} \tilde{x}_t \right] + (1 - \psi) [\tilde{p}_{ht} + \tilde{y}_{ht}] \quad (\text{G1})$$

$$0 = (1 - \alpha) \tilde{p}_{ht} + \alpha \tilde{x}_t \quad (\text{G2})$$

$$\tilde{y}_{ht} = -[\eta\omega + \gamma(1 - \omega)] \tilde{p}_{ht} + \omega \tilde{c}_t + \gamma(1 - \omega) \tilde{x}_t \quad (\text{G3})$$

$$\pi_{ht} = \lambda[\sigma \tilde{c}_t + \varphi \tilde{y}_{ht} - \tilde{p}_{ht}] + \beta E_t \pi_{ht+1} \quad (\text{G4})$$

Note: This table collects the linearized equations of the model.

Table 2. Calibrated Weights in the Optimal Policy Rule

a) Relative Weight of Output

ψ	Unit Elasticities	$\eta = 0.25$	$\eta = 0.25$ and $\gamma=5$
1	0.2	0.217	1.968
0.8	0.2	0.225	0.352
0.6	0.2	0.238	0.262
0.4	0.2	0.258	0.24
0.2	0.2	0.300	0.246
0	0.2	0.296	0.201

b) Relative Weight of Food Price Shock in Target Output

ψ	Unit Elasticities	$\eta = 0.25$	$\eta = 0.25$ and $\gamma=5$
1	0	-0.173	-0.372
0.8	0	-0.199	-0.152
0.6	0	-0.231	-0.078
0.4	0	-0.271	-0.043
0.2	0	-0.321	-0.02
0	0	-0.324	-0.022

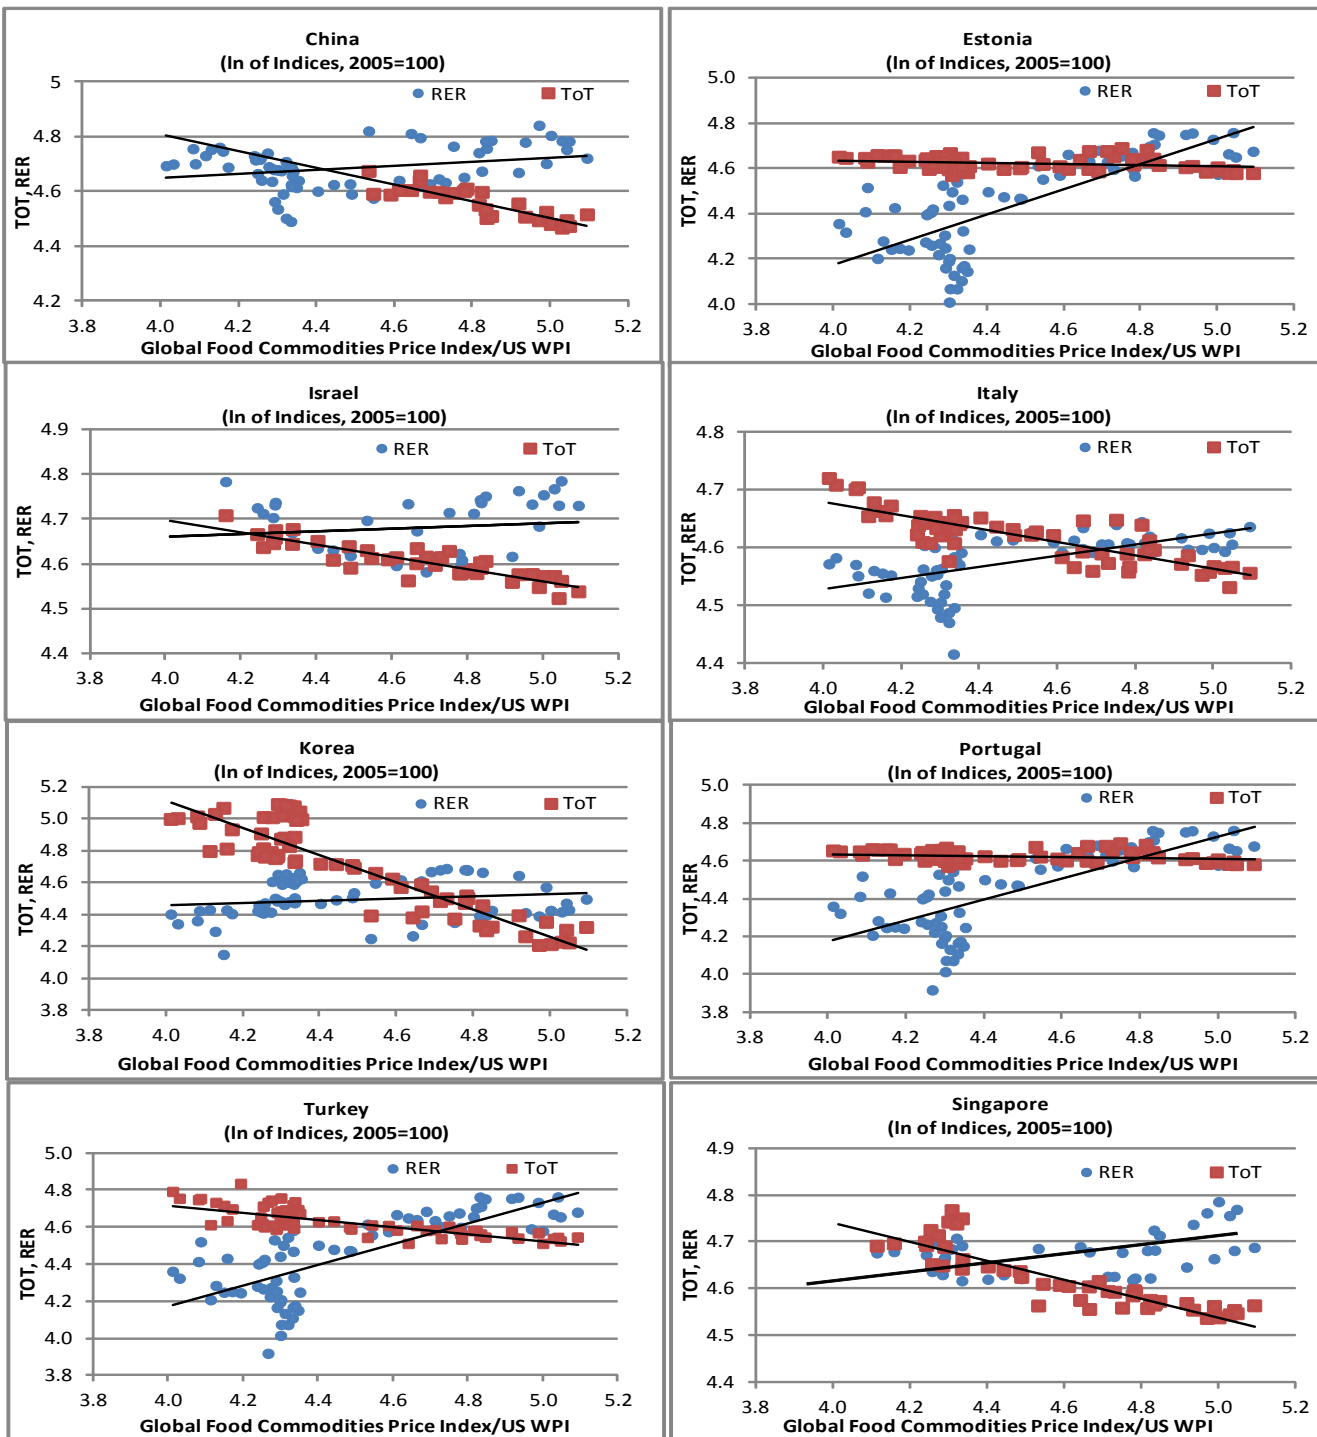
Note: The first panel reports the weight of output relative to inflation in an optimal policy rule. The second panel gives the weight of the food price shock in target output under the optimal policy rule.

Table 3. Model Statistics Under Simulated Random Shocks

	Complete Markets ($\psi=1.0$)					Imperfect Risk Sharing ($\psi=0.9$)					Financial Autarky ($\psi=0$)				
	Optimal Policy	PPI Rule	CPI Rule	EXP(CPI) Rule	PEG Rule	Optimal Policy	PPI Rule	CPI Rule	EXP(CPI) Rule	PEG Rule	Optimal Policy	PPI Rule	CPI Rule	EXP(CPI) Rule	PEG Rule
Standard deviations (in %)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Domestic Output	0.990	0.668	0.780	0.752	1.047	0.628	0.642	0.748	0.711	1.053	0.492	0.502	0.568	0.540	0.540
Consumption	0.408	0.529	0.547	0.499	0.546	0.434	0.751	0.598	0.534	0.612	0.677	0.572	0.598	0.629	1.756
Real Exchange Rate	2.475	3.206	3.315	3.021	3.310	2.353	3.112	3.219	2.948	3.256	1.985	2.504	2.493	2.429	2.764
CPI inflation	1.615	1.777	1.519	1.839	1.844	1.464	1.754	1.538	1.846	1.841	1.503	1.898	1.803	1.857	8.786
Domestic Inflation	0.361	0.234	0.129	0.152	0.473	0.279	0.049	0.136	0.144	0.486	0.007	0.091	0.290	0.301	0.713
Means in % of SS deviation															
Domestic Output	0.021	-0.038	0.003	0.000	-0.132	0.009	0.009	0.003	0.002	-0.131	0.007	0.006	-0.008	-0.011	-0.121
Consumption	-0.003	-0.013	-0.007	-0.008	-0.032	-0.002	-0.005	-0.007	-0.008	-0.041	0.001	-0.007	-0.022	-0.029	-0.113
Real Exchange Rate	-0.001	-0.053	-0.012	-0.027	-0.168	0.003	-0.005	-0.006	-0.007	-0.006	-0.011	-0.007	-0.013	-0.025	-0.161
CPI inflation	0.012	0.001	0.073	0.038	0.017	0.010	0.025	0.083	0.044	0.017	0.011	0.025	0.083	0.044	0.017
Domestic Inflation	0.000	-0.015	0.061	0.022	0.001	0.000	0.009	0.071	0.027	0.001	0.000	0.002	0.153	0.173	0.003

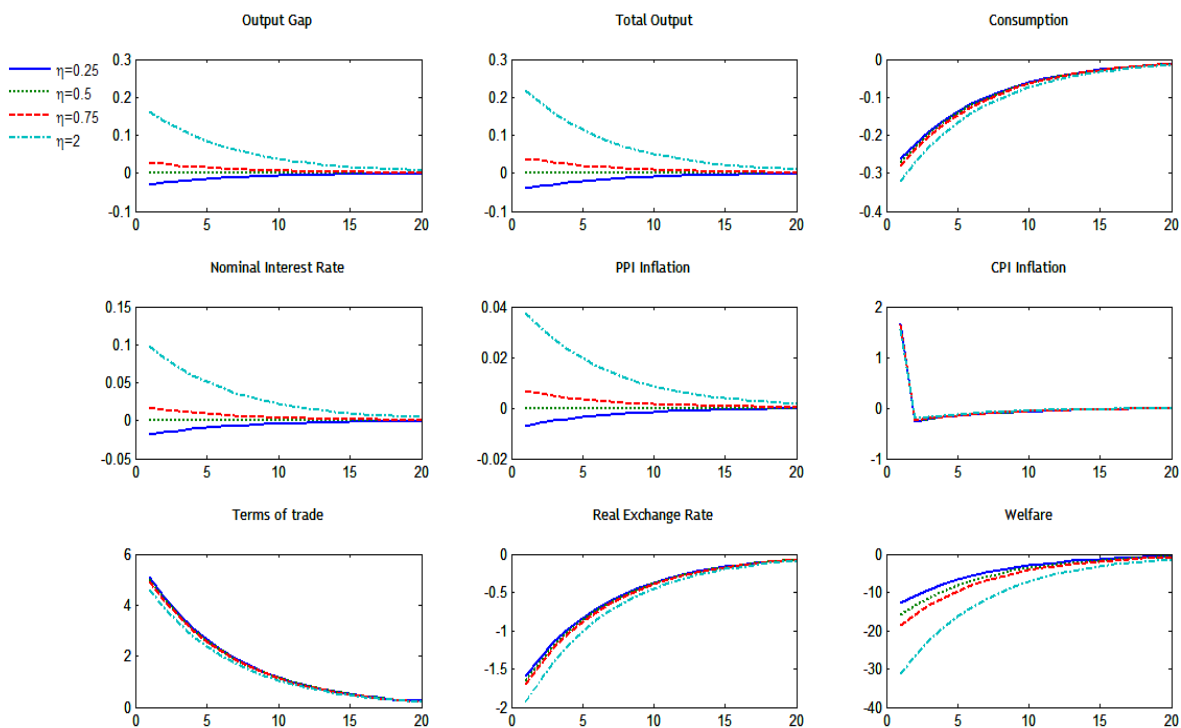
Note: This table reports standard deviations and mean deviations from non-stochastic steady state for main aggregates. They are computed from simulating 10,000 random shocks to the processes for import food price index, home productivity, and monetary policy, for the case $\eta=1/\sigma=0.5$, $\gamma=5$, and baseline values for the other parameters. Results are displayed for optimal monetary policy, PPI inflation targeting (IT), *headline* CPI IT, *expected* CPI IT, and an exchange rate peg. The three panels correspond to perfect risk sharing, imperfect risk sharing, and financial autarky.

Figure 1. Covariance of the Terms of Trade and Real Exchange Rate with World Food Prices



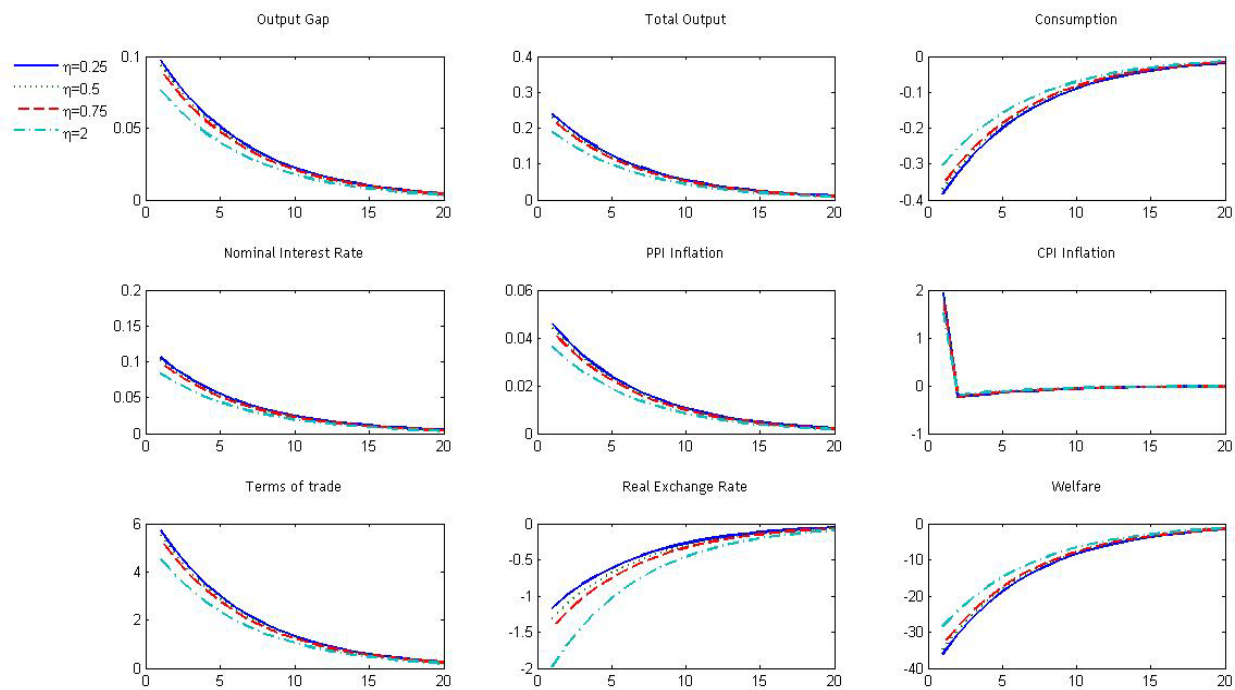
Note: This figure plots the natural logarithm of the terms of trade and real effective exchange rate indices for each country (both normalized so that 2005=100), against the natural logarithm of the IMF index of world food commodity prices deflated by the US Wholesale Price Index (also normalized so that 2005=100) for the period of 1994Q1 to 2011:Q4.

Figure 2. Impulse-responses to a world food price shock under full risk sharing



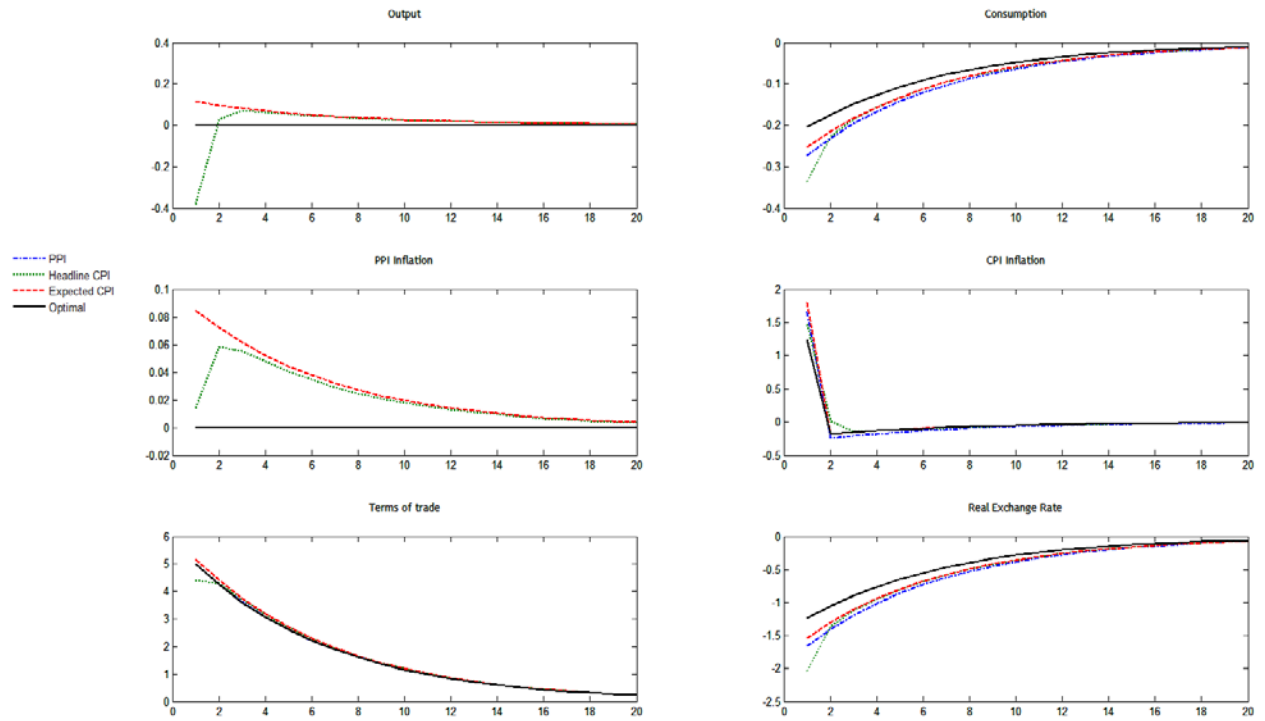
Note: This figure plots the responses (in percentage points) of the main aggregates of the model to a one standard deviation shock (5 percent) to the imported food price index under perfect international risk sharing ($\psi=1$). A flexible PPI Taylor rule is assumed. The inter-temporal substitution elasticity (σ) is equal to 2. The remaining parameters are set to the baseline values discussed in section 5.1 of the main text.

Figure 3. Impulse-responses to a world food price shock under financial autarky



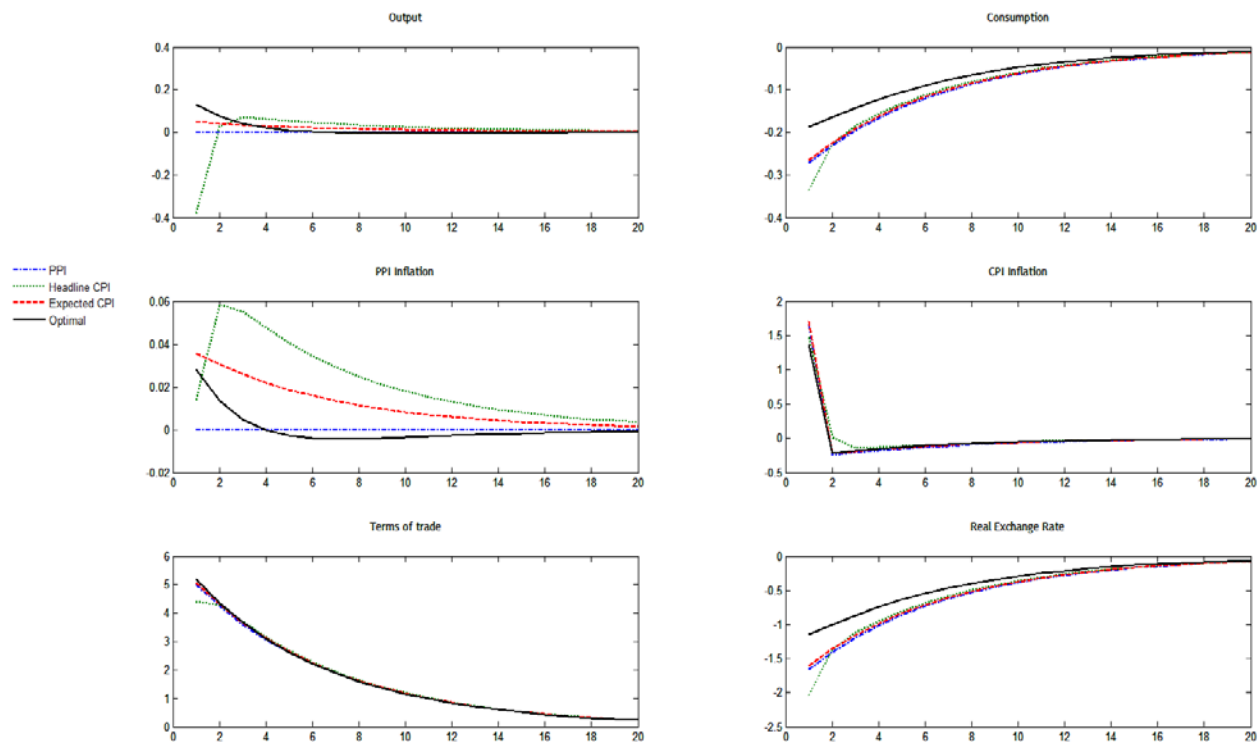
Note: This figure plots the responses (in percentage points) of the main aggregates of the model to a one standard deviation shock (5 percent) to the imported food price index under financial autarky ($\psi=0$). A flexible PPI Taylor rule is assumed. The inter-temporal elasticity (σ) is equal to 2. The remaining parameters are at baseline values.

Figure 4: Optimal Policy vs. Simple IT Rules under Complete Markets: IRs to z^* shock




Note: This figure plots the responses (in percentage points) of the main aggregates of the model to a one standard deviation shock (5 percent) to the imported (food) price index under perfect international risk sharing ($\psi=1$). It is assumed that $\eta=1/\sigma=0.5$, $\gamma=5$, and that other parameters are kept at baseline values.

Figure 5: Optimal Policy vs. Simple IT Rules under Financial Autarky: IRs to z^* shock



Note: This figure plots the responses (in percentage points) of main model aggregates to a one standard deviation shock (5 percent) to the imported food price index under financial autarky ($\psi=0$). It is assumed that $\eta=1/\sigma=0.5$ and $\gamma=5$, and that other parameters are kept at baseline values.

Figure 6: Welfare Differences Across Simple Policy Rules



Note: This figure plots differences in conditional welfare, expressed as percent of non-stochastic steady state consumption, between pairs of competing policy rules for alternative values of η . The upper panel assumes perfect risk sharing ($\psi=1$) and all shocks present. The middle panel assumes perfect risk sharing but z^* shocks are set to zero. The lower panel assumes all shocks are present but international risk sharing is imperfect ($\psi=0.9$).