Global Price Shocks and International Monetary Coordination*

Veronica Guerrieri Guido Lorenzoni Iván Werning

June 2025

Individual central banks respond to global supply shocks that transmit inflationary pressures—such as oil prices, shipping costs, and bottlenecks in global supply chains—taking these conditions as given. However, their combined global response determines global demand and, thus, the resulting global price pressure. This paper builds a simple monetary open economy model to explore the economic implications of this channel. We show that, following a negative world supply shock, uncoordinated monetary policy may be excessively loose. Our mechanism for this "expansionary bias" applies to an aggregate shock in a symmetric world economy of small open economies having no individual control over their terms of trade. In these ways, it is distinct from asymmetric shocks and terms-of-trade manipulation motives emphasized in the monetary coordination literature.

1 Introduction

The inflation surge of 2021–2022 had two visible features: it was global in nature, playing out similarly across many advanced and developing economies, and it was strongly associated with scarcity in the global supply of tradable inputs—apparent in various measures of supply chain pressures, increases in shipping costs, and spikes in energy prices.[1]

These two observations have led many economists to view this inflation episode as largely driven by global forces, mostly outside the control of individual central banks. However, global scarcity is relative—a matter of supply versus demand—and global aggregate demand is ultimately shaped by the combined responses of individual countries.[2] Each country may take world commodity prices and bottlenecks in the global supply of tradable intermediates as given, yet their collective decisions shape the tightness of these supply constraints and the resulting prices.

*We thank Javier Bianchi, Tommaso Monacelli, Dimitri Mukhin, Federica Romei, and seminar participants for useful discussions and comments. Bruno Moscarini provided outstanding research assistance.

[1] The Global Supply Chain Pressure Index constructed by the New York Fed spiked to over four standard deviations above its historical average by the end of 2021. The FBX index of container shipping rates increased fivefold from pre-pandemic levels in the same period. Energy prices experienced significant hikes, with, for example, the WTI oil price index more than doubling from early 2021 to mid 2022.

[2] Indeed, Miranda-Pinto et al. (2023) study US monetary policy and find a sizable transmission channel to inflation, both in the US and in other countries through its effects on commodity markets.

These world price linkages and interactions suggest that uncoordinated policy responses may produce a suboptimal world outcome. A widely held view, dating back to the classic work of Canzoneri and Henderson (1991), is that after a global negative supply shock (modeled as a TFP shock), uncoordinated policy choices can lead to a contractionary bias—that is, monetary policy that is excessively tight. More recently, Obstfeld (2022) raised various contractionary-bias concerns surrounding the response of central banks to the late 2022 inflation surge. The general contractionary-bias argument is that each central bank has an incentive to tighten, as this leads to a domestic currency appreciation, which lowers the domestic-currency price of imports and helps reduce inflation. But in response to a global shock, all currencies cannot appreciate simultaneously; the result is a zero-sum game that leads to excessive tightening.

In this paper, we revisit this conventional wisdom and show that, once the role of global commodity markets and other globally determined input prices is explicitly modeled (distinct from TFP shocks), this conclusion is reversed: lack of coordination can lead to excessively loose policy at the global level—an expansionary bias.

Our model is built on a standard open-economy New Keynesian model augmented with two features. The first crucial ingredient is the inclusion of a global commodity market in which the price of the scarce input is determined. As we argued, this is central to recent inflationary episode. For ease of exposition, we focus on a model with only two inputs: a local input, labor, and a global input, which we label "oil" for concreteness but interpret more broadly as a stand-in for traded commodities and intermediate goods. Labor is immobile, while oil is traded in a world market with a perfectly flexible price.

The second crucial ingredient is the presence of nominal wage rigidities. When combined with standard price frictions, wage rigidities introduce a trade-off between inflation stabilization and output gap stabilization following a supply shock, breaking the so-called divine coincidence.

With both ingredients in place, our main exercise explores a global oil supply shock that hits all countries symmetrically. We first solve for the general equilibrium in which each country unilaterally responds to the shock, optimally setting its own monetary policy while taking the policies of others and, crucially, the price of oil as given. We then study a coordinated response at the world level, jointly setting monetary policy to internalize the effect on the endogenous oil price. Comparing these two outcomes, we find that uncoordinated monetary policy can be excessively loose relative to the ideal coordinated benchmark.

To develop economic intuition for our result, we compare Phillips curves at the country and world levels. Each local central bank faces a country-level Phillips curve, relating domestic inflation to output, and takes the world oil price as given. The world-level Phillips curve, in contrast, describes the aggregate trade-off between global inflation and global output, incorporating the effect of activity on the endogenous oil price. We show that under empirically plausible parameter choices the world-level Phillips curve is steeper than the country-level curve. Intuitively, higher global activity raises the equilibrium oil price, fueling additional inflation. Local central banks do not internalize this effect, and thus perceive flatter trade-offs.

These contrasting Phillips curves suggest that a coordinated world planner attempting to stabilize both inflation and output will recognize the less favorable tradeoff and optimally chose lower inflation and lower output than individual countries would. In a nutshell, countries impose an inflationary externality on each other, through the price of oil. The coordinated solution internalizes this externality, leading to lower inflation.

Although the above intuition is an important part of the explanation, optimal policy is more subtle. There are two loose ends that our analysis addresses. First, we establish under what conditions the world Phillips curve is indeed steeper than the country one. We show that this is true whenever the equilibrium oil price is increasing in the level of global economic activity, all else equal. This seems the empirically relevant case. In the model, this property depends on parameters. In particular, it holds when the elasticity of substitution between labor and oil is sufficiently low, which is empirically plausible.

While this insight helps explain the inflationary externality, welfare maximization requires us to go beyond reduced-form tradeoffs. In particular, it is not obvious that local central banks and the world planner share the same objective function, or that these objectives can be reduced to inflation and output only, as is common in closed economy models.

We explicitly study the welfare optimum of the country and of the world central bank, and we show that their respective choices can be decomposed in three elements. The first two are standard: an output gap measure and a measure of inflation distortion. The third captures a distortion in the use of imported oil in the economy, which arises due to wage and price rigidities. However, while this distortion affects an individual central bank choice, it does not actually matter at the world level, due to the fact that oil trade must be balanced in the world equilibrium. Overall, the conditions for an expansionary bias depend not only on the slopes of the Phillips curves, but also on the country-level perceived cost of oil imports.

We provide conditions for an expansionary bias in the uncoordinated equilibrium. As it turns out, a low elasticity of substitution between labor and oil again emerges as the key parameter generating an expansionary bias in response to a negative supply shock.

The economic logic of our result is likely robust to variations of the basic model. Our modeling choices seek to transparently isolate the mechanisms at play. In particular, we assume all countries are symmetric and subject to a common shock. As a result, in the symmetric equilibrium, countries simply consume their own oil and do not run trade surpluses or deficits in oil. We also assume that countries are price takers in the world markets for traded goods. Overall, countries neither care to nor are able to affect world oil prices. Thus, unlike much of the literature on monetary policy coordination, our results are not driven by terms-of-trade manipulation. Section 6 considers a variant of our model where terms-of-trade manipulation motives are present.

The literature on international monetary policy coordination has a long tradition, dating back to Hamada (1976). Early work was indeed motivated by the responses to the high inflation of the 70s, including Oudiz and Sachs (1984) and Canzoneri and Henderson (1991), which we mentioned above. Microfounded New Keynesian models introduced a welfare-based treatment, with key contributions by Corsetti and Pesenti (2001), Obstfeld and Rogoff (2002), Benigno and Benigno (2006), and Engel (2011). This literature focused on the coordination problem arising from terms of trade manipulation motives and the presence of idiosyncratic asymmetric shocks.

Closer to our approach are Fornaro and Romei (2022) and Bianchi and Coulibaly (2024). They study the optimal response to global, symmetric shocks in models without terms of trade manipulation motives. In particular, Fornaro and Romei (2022) study a demand shift toward traded goods and away from non-traded goods, and find a contractionary bias in uncoordinated monetary policy. In their model, higher inflation has the beneficial side effect of promoting reallocation of labor towards the goods in high demand (the traded goods) as in Guerrieri, Lorenzoni, Straub and Werning (2021). In an open economy, however, a country perceives its own expansionary inflationary policy as less effective because domestic households save part of the increase in traded production abroad. This international spillover leads to a contractionary bias. Bianchi and Coulibaly (2024) expand on this mechanism in a more general framework that accommodates a broader set of shocks. They derive conditions in terms of endogenous sufficient statistics for the direction of the bias.[3] In both these papers the spillover works through the world interest rate. This is distinct from our paper that focuses on global markets for commodities and intermediate inputs.

Our paper is also related to the large literature on open economy models with com-

[3] Another strand of research confronts the different sources of inefficiencies by going beyond monetary policy and considering other policy instruments, such as capital controls, insurance, tariffs or fiscal policy, sometimes achieving the first-best (e.g. Farhi and Werning, 2014, Itskhoki and Mukhin, 2023). Here, we focus on monetary policy away from divine coincidence, due to combination of wage and price rigidity.

modity markets. More recently, Auclert, Rognlie and Straub (2024) explore the effects of energy price shocks in a rich open economy model with financial constrained consumers. They focus on oil importing countries, but also emphasize the fact that oil prices are globally determined. However, they do not pursue a normative analysis of optimal policy, neither uncoordinated nor coordinated. Drechsel et al. (2025) study optimal monetary policy in response to commodity price shocks and examine how it varies depending on whether a country is a net exporter or importer of the commodity. However, they take the price shocks as exogenous, so they cannot characterize uncoordinated or coordinated world equilibria.

2 The World Economy

Time is discrete $t = 0, 1, \ldots$ The world comprises a continuum of identical countries. Our focus is on symmetric equilibria so it is natural to omit a subscript for each individual country.

There are two consumption goods, traded and non-traded goods, and two inputs, labor and a traded commodity input. To be concrete, we will refer to this commodity as "oil", but we have in mind a broader set of global inputs into domestic production.

To simplify our analysis, we focus on an unanticipated temporary shock to the oil endowment at $t = 0$.

Preferences. In each country, a representative household consumes tradable goods C_T, consumes non-tradable goods C_N, and supplies labor, L_t. Preferences are represented by the utility function

$$\sum_{t=0}^{\infty} \beta^t U(C_{Tt}, C_{Nt}, L_t),$$

with

$$U(C_T, C_N, L) = C_T + u(C_N) - v(L).$$

This quasilinear-linear utility in tradable consumption is adopted to streamline the analysis but is not critical. Indeed, we later briefly discuss the case with concave utility over C_T.

Non-tradable goods come in a continuum of varieties indexed by $i \in [0, 1]$ that are aggregated according to

$$C_{Nt} = \left(\int_0^1 (C_{Nt}(i))^{1-1/\varepsilon} di \right)^{\frac{1}{1-1/\varepsilon}},$$

a constant elasticity of substitution (CES) function with elasticity $\varepsilon > 1$.

Technology. Each variety i is produced by a monopolistic firm with production function

$$Y_t(i) = F(L_t(i), X_t(i)),$$

where $L_t(i)$ is the labor input and $X_t(i)$ the oil input.[4] The production function F is assumed to have constant returns to scale. In the general analysis, we do not make specific functional form assumptions on F, while in examples we will use a CES production function.

The government sets a subsidy on production to offset the monopolistic distortion, financed by a lump-sum tax. This standard assumption ensures that the steady state is efficient, which allows focusing on stabilization policies.

The representative household in each country receives a constant endowment $\bar{Y}_T$ of tradable goods and an endowment $\bar{X}_t$ of oil in each period. Our focus is on a shock to the endowment sequence $\{\bar{X}_t\}$. In particular, a purely temporary shock in the initial period: $\bar{X}_0 \neq \bar{X}_1 = \bar{X}_2 = \cdots = \bar{X}$.

Markets. The budget constraint of the household is

$$\begin{aligned} P_{Tt}C_{Tt} + P_{Nt}C_{Nt} = P_{Tt}\bar{Y}_T + W_tL_t + P_{Xt}\bar{X}_t + T_t + \Pi_t, \\ + (1+i_t)B_t - B_{t+1} + ((1+i_t^*)B_t^* - B_{t+1})E_t \end{aligned}$$

where T_t is transfers from the government, Π_t aggregate profits from domestic firms, B_t domestic bonds, i_t the domestic nominal interest rate, B_t^* foreign bonds, and i_t^* the foreign nominal interest rate and E_t is the domestic exchange rate. Each country has its own currency and prices P_{Tt}, P_{Nt}, P_{Xt} and the wage W_t are expressed in domestic currency. Here E_t represents the nominal exchange rate of the local currency against a world unit of account, taken to be a equally weighted basket of all currencies. Likewise, the world interest rate i_t^* is the interest rate in this unit of account.

Prices for the tradable good and for oil are perfectly flexible and satisfy the law of one

[4]The fact that oil enters the production of non-tradaded goods, but does not affect the production of traded goods is a simplification due to the fact that we assume a tradable endowment. However, this is not essential for our results.

price

$$P_{Tt} = E_t P^*_{Tt}, \tag{1}$$

$$P_{Xt} = E_t P^*_{Xt}, \tag{2}$$

where P^*_{Tt} and P^*_{Xt} are world prices denominated in a world unit of account.

Nominal Rigidities: Prices and Wages. We introduce nominal rigidities in both prices and wages in a simple and tractable way. To simplify, we only focus on rigidities in $t = 0$, in line with our focus on a one-time, unanticipated temporary shock at $t = 0$, but it is simple to extend the analysis beyond this.

Non-tradable goods prices are subject to the following nominal rigidity at $t = 0$. A fraction λ of firms in the non tradable sector set their price at the beginning of the period, before the aggregate shock is realized, while $1 - \lambda$ set their price after the aggregate shock is realized. The wage is fixed at $t = 0$ and normalized to unity $W_0 = 1$ and flexible in all other periods. We later contrast this to the case of flexible wages.

Nominal wage rigidity creates a trade-off between inflation and real activity following a supply shock, or, in the common jargon, it breaks the "divine coincidence". As we shall see the presence of this trade-off plays a crucial role in evaluating the gains from monetary coordination.

Symmetric Equilibria. We focus on symmetric equilibria after a symmetric shock: all countries experience the same shock and the same outcome. Symmetry has two immediate implications that simplify the possible outcomes.

First, in equilibrium the trade balance must be zero. Indeed, we can focus on equilibria where both oil and the traded consumption good are equal to their endowments

$$C_{Tt} = \bar{Y}_T$$

$$X_t = \bar{X}_t$$

so that net trade is zero in both goods for all $t = 0, 1, \ldots$

Second, the exchange rates are equalized implying an exchange rate of unity $E_t = 1$ relative to the world basket for all $t = 0, 1, \ldots$.

3 Equilibrium and Policy

This section first develops the equilibrium conditions and then defines the two policy problems, uncoordinated and coordinated.

3.1 Equilibrium Conditions

We now develop the equilibrium conditions. Although similar conditions hold for all $t = 0, 1, \ldots$ we focus on $t = 0$ and, in that context, drop time subscripts to ease the notation. The $t = 0$ conditions turn out to be sufficient to develop our analyses.

Price setting. A firm in the non-tradable sector that can reset its price will choose

$$\tilde{P}_N = \mathcal{M}(W, P_X), \tag{3}$$

with the marginal cost function $\mathcal{M}(W, P_X) \equiv \min_{L,X} \{WL + P_X X : F(L, X) = 1\}$. Recall that we normalize the wage to unity, $W = 1$. We also normalize the price of firms that cannot adjust to 1. Aggregating the prices of adjusters and non-adjusters, the price index of non-tradable goods is

$$P_N = (\lambda + (1 - \lambda)\tilde{P}_N^{1-\varepsilon})^{\frac{1}{1-\varepsilon}}. \tag{4}$$

Production. Since they face the same prices and wage, all firms choose the same oil-to-labor ratio, satisfying the cost minimization condition

$$\frac{F_X(L/X, 1)}{F_L(L/X, 1)} = \frac{P_X}{W}. \tag{5}$$

Production for firms that adjust prices is generally different than that of firms that do not adjust their price. Total output equals consumption

$$Y = C_N, \tag{6}$$

and satisfies

$$\Delta \cdot Y = F(L, X), \tag{7}$$

where

$$\Delta \equiv \lambda(1/P_N)^{-\varepsilon} + (1 - \lambda)(\tilde{P}_N/P_N)^{-\varepsilon}$$

captures the distortionary effects of inflation, due to price dispersion.[5] When $\tilde{P}_N = P_N = 1$ inflation is zero and $\Delta = 1$. Whenever, $\tilde{P}_N \neq 1$ there are relative price distortions so that $\Delta > 1$ due to the misallocated across symmetric firms with different prices.

Demand and Trade. The consumer demand for non-tradable goods is given by the optimality condition

$$u'(C_N) = \frac{P_N}{P_T}. \tag{8}$$

The country consolidated budget constraint is

$$C_T = \bar{Y} - Q \cdot (X - \bar{X}), \tag{9}$$

where Q is the world relative price of oil, expressed in tradable goods

$$Q \equiv \frac{P_X^*}{P_T^*}.$$

Equation (9) assumes that whenever a country buys oil in excess of its endowment, $X - \bar{X} > 0$, it pays for it by reducing the consumption of tradable goods in the same period. This is without loss of generality. In general a country could borrow and pay partly in the future, but given quasi-linear preferences the value is the same.

By (1)–(2), we must have

$$\frac{P_X}{P_T} = Q, \tag{10}$$

the domestic relative price of oil to tradable goods equals the world relative price.

Intertemporal Conditions. In equilibrium, for all $t = 0, 1, \ldots$ interest rate parity requires

$$(1 + i_{t+1})E_t / E_{t+1} = 1 + i^*_{t+1} \tag{11}$$

Furthermore, in equilibrium

$$(1 + i^*_{t+1})P^*_{Tt} / P^*_{Tt+1} = R^* = 1/\beta \tag{12}$$

so that the world real interest rate for traded goods R^* is pinned down by the discount rate. This equation is simply the Euler equation for traded goods. We have assumed

[5] The demand for sticky goods is $(1/P_N)^{-\varepsilon} C_N$ and the demand for flexible goods is $(\tilde{P}_N/P_N)^{-\varepsilon} C_N$, so total demand is

$$\left[\lambda (\frac{1}{P_N})^{-\varepsilon} + (1 - \lambda)(\frac{\tilde{P}_N}{P_N})^{-\varepsilon} \right] C_N.$$

quasilinear utility. However, if utility is additively separable and the endowment of the traded good is constant, then the world real interest rate for traded goods is still given exogenously by $R^* = 1/\beta$.[6,7]

Equilibrium. In any given period, we have eight equilibrium conditions (3)–(10) and nine endogenous variables

$$(P_T, \tilde{P}_N, P_N, P_X, C_N, C_T, Y, L, X).$$

Thus, there is one degree of freedom. In what follows, we shall describe policy as a choice over output Y, with equations (3)–(10) determining the rest of the variables and welfare.

In addition, we have the intertemporal conditions (11) and (12). However, the intertemporal conditions do not play a central role and can be omitted from the analysis. Indeed, we can simply note that the parity condition (11) determines the interest rate given a choice for E_t, while (12) constrains the world interest rate. These observations are relevant for implementation purposes, but not required to set up the planning problem.

3.2 Policy Problems

We compare two regimes for international monetary policy coordination. In both regimes, the planner is benevolent and maximizes the utility of the representative consumer.

Definition 1. *(No coordination equilibrium)* A symmetric Nash equilibrium is a vector

$$(Q, P_T, \tilde{P}_N, P_N, P_X, C_N, C_T, Y, L, X)$$

that satisfies two conditions:

1. *Country's optimization.* It satisfies conditions (3)-(10) and, among all the vectors that satisfy these conditions, it maximizes the utility of the representative household, taking Q as given.

2. *Market clearing in the world oil market.* The price Q is such that $X = \bar{X}$.

In the *no-coordination regime*, each central banks acts independently choosing output Y, taking as given relative price of oil. In considering deviations, countries entertain the

[6]This is a key distinction with the work of Fornaro and Romei (2022); Bianchi and Coulibaly (2024).

[7]Note that combining (11) and (12) with the arbitrage conditions and the intratemporal condition gives an intertemporal Euler equation for non-traded goods $U'(C_{Nt}) = \beta(1 + i_{t+1})U'(C_{Nt+1})P_{Nt}/P_{Nt+1}$.

possibility of oil consumption above or below their endowment. However, in the resulting symmetric equilibrium oil consumption equals the endowment. We can interpret the situation as a game, where each country chooses Y taking as given world output, which determines the world price; we study the symmetric Nash equilibrium of this game.

In the *coordination regime*, the planner internalizes the market clearing condition in the world oil market.

Definition 2. *(Coordination equilibrium)* An equilibrium with coordination is a vector

$$(Q, P_T, \tilde{P}_N, P_N, P_X, C_N, C_T, Y, L, X)$$

that maximizes the utility of the representative household, taking as given conditions (3)-(10), plus the condition $X = \bar{X}$.

Since the constraints on the problem are identical, welfare cannot be lower in the coordination regime.

4 Phillips Curves: Country versus World

This section derives Phillips curves: equilibrium relationships between real activity and inflation, at the country level and at the world level. We compare the two and identify conditions under which the world Phillips curve is steeper. This condition plays a crucial role in our subsequent normative analysis.

4.1 Country-Level Phillips Curve and Oil Demand

At the country level, the international relative price of oil Q is taken as given. There are then eight equilibrium conditions and nine variables. This offers one degree of freedom, which we take to be output Y and then solve for inflation of non-traded goods and oil consumption in terms of Y and Q; we briefly discuss the exchange rate.[8]

Country Phillips Curve. Log linearizing around a zero inflation steady state with constant endowments of tradable goods and oil, the appendix derives the country-level Phillips curve

$$\pi_N = \frac{1 - \hat{\lambda}}{\hat{\lambda}} (\sigma y + q), \tag{13}$$

[8]Output is the natural measure of activity since tradable goods are an endowment. Inflation in traded good prices does not generate distortions, so we focus on inflation in non-traded goods.

where lowercase variables denote log-linear deviations from the steady state values, π_N denotes non-tradable inflation, and $\hat{\lambda} \equiv 1 - (1-\lambda)s_X$ where s_X is the oil share in total costs. The value of $\hat{\lambda}$ is a share-adjusted measure of stickiness; it satisfies $\hat{\lambda} \in (s_L, 1)$ (where $s_L = 1 - s_X$) and is increasing in λ.[9] Here σ is the local curvature of the utility function for non-tradable goods.[10,11]

From the perspective of an individual country, an oil price shock ($q > 0$) shifts the Phillips curve (13) upward. The central bank can offset the inflationary effect of the oil shock by accepting a contraction in domestic activity ($y < 0$). This creates a non-trivial trade off between inflation and output.

Why does a contraction in output help contain inflation? Since nominal wages are fixed, the labor market channel is completely shut down. However, monetary policy works through the exchange rate channel. Contractionary policy leads to an appreciation—a lower E. This makes oil cheaper in domestic currency as $P_X = EP_X^*$, which reduces marginal costs and cools non-tradable inflation. The price of tradable goods, P_T, also falls in proportion to the nominal appreciation, due to $P_T = EP_T^*$. Because of wage rigidities and price stickiness, non-tradable prices fall by less than tradable prices, so non-tradable goods become relatively more expensive. In short, the nominal appreciation causes a real appreciation. This relative price effect reduces non-tradable consumption and output. As a result, the currency appreciation achieves lower inflation at the expense of real activity.[12]

Oil Demand. The country demand for oil x is given by

$$x = (1 - \sigma\gamma)y - \gamma q \tag{14}$$

[9]Intuitively, since wages are rigid we have $\hat{\lambda} \geq s_L$, with equality when prices are not sticky $\lambda = 0$. At the other extreme, $\hat{\lambda} = 1$ if $s_X = 0$ or $\lambda = 1$.

[10]Denoting with the subscript SS the steady state allocation they are equal to

$$\sigma = -u''(Y_{SS})Y_{SS}/u'(Y_S),$$

$$s_X = F_X(L_{SS}, X_{SS})X_{SS}/F(L_{SS}, X_{SS}).$$

[11]To derive the CPI Phillips curve, just combine (13) and the following relation between CPI inflation π and non-tradable inflation π_N:

$$\pi = \pi_N + \omega\sigma y,$$

where ω is the steady state share of spending in tradable goods.

[12]The exchange rate elasticity of P_T is 1, the exchange rate elasticity of P_N is $(1-\lambda)s_X = 1 - \hat{\lambda} < 1$, so the elasticity of P_N/P_T is $(1-\lambda)s_X - 1 = -\hat{\lambda} < 0$. Given that $1/\sigma$ is the elasticity of non-tradable consumption to P_N/P_T, the response of non-tradable consumption is $(1/\sigma)\hat{\lambda} > 0$. A 1% nominal appreciation gives a $(1-\hat{\lambda}) \cdot 1\%$ reduction in non-tradable inflation at the cost of a $(1/\sigma)\hat{\lambda} \cdot 1\%$ reduction in non-tradable consumption. This gives an exact intuition for the slope of the Phillips curve.

where $\gamma = \eta s_L / \hat{\lambda} > 0$ and η is the local elasticity of substitution between oil and labor in the production function.[13] An increase in price q for fixed output y induces substitution towards oil, away from labor.

An increase in output has less than a one for one effect on oil, $1 - \sigma\gamma < 1$. To see why, note that (5)–7 imply, to a first order, that if P_X/W were fixed, then an increase in Y induces a one-for-one increase in X and L, with X/L fixed.[14] The equilibrium response is less than this because Q is fixed, but P_X/W is not fixed. Indeed, as output rises the exchange rate depreciates (by an amount that depends on σ), E rises, raising the nominal price P_X, while W is fixed. This induces a shift away from oil towards labor (in an amount determined by η). This explains why the effect of y on x is less than one.

Note that output increases the use of oil whenever $1 - \sigma\gamma > 0$, or equivalently when

$$\eta s_L < \frac{\hat{\lambda}}{\sigma}. \tag{15}$$

Economically, there are two opposing forces. On the one hand, by reducing domestic output, contractionary policy lowers oil demand. On the other hand, by making oil cheaper in domestic currency, relative to the fixed wage, it increases the quantity of oil demanded per unit of output.[15]

The empirically plausible case has oil consumption rising with output, which holds within our model when the elasticity of substitution η is low enough.

4.2 World-Level Phillips Curve

Turning to the world Phillips curve, we no longer take Q as given, but instead impose the symmetric general equilibrium condition that $X = \bar{X}$. This then determine Q as a function of oil supply $\bar{X}$ and output Y. In log-linear terms, this requires

$$x = \bar{x}.$$

Suppose all countries choose the same output level $\bar{y}$. Then, equating world oil demand (14) with oil supply $\bar{x}$ and solving for q gives the equilibrium oil price

$$\bar{q} = \frac{1}{\gamma}(1 - \sigma\gamma)\bar{y} - \frac{1}{\gamma}\bar{x}. \tag{16}$$

[13] That is, $\eta \equiv F_L F_X / (F_{XL} F)$ evaluated at (L_{SS}, X_{SS}).

[14] Note that $\Delta = 1$ to a first order.

[15] Following up on the intuition in footnote 12, the reduction in consumption caused by a 1% appreciation is $(1/\sigma)\hat{\lambda} \cdot 1\%$, and causes a proportional shift in the oil demand function. The price effect of the appreciation, due to substitution between oil and labor, is $\eta s_L \cdot 1\%$. Therefore, the two terms on the right-hand side of the inequality capture exactly the two effects described in the text.

Notice that condition $1 - \sigma\gamma > 0$ implies that the equilibrium oil price is increasing in world output y.

Substituting $q = \bar{q}$ into (13) gives the world Phillips curve

$$\bar{\pi}_N = \frac{1-\hat{\lambda}}{\hat{\lambda}}\frac{1}{\gamma}(\bar{y} - \bar{x}). \tag{17}$$

Comparing the slopes of our two Phillips curves leads to the following result.

Proposition 1. *The world-level Phillips curve* (17) *is steeper than the country-level Phillips curve if and only if any of the following equivalent conditions hold: (i) the oil demand of the individual country is increasing in the country's output; (ii) the world oil price is increasing in world output; (iii) inequality* (15) *holds.*

Recall that the condition that oil consumption rises with output, at the country level, is empirically plausible and holds if the elasticity of substitution η is low enough. The economic intuition for a steeper world Phillips curve is as follows. When world output y expands it raises oil demand (assuming (15) holds) and pushes the relative price of oil q up; this then causes higher world inflation.

Suppose a deviating individual country chooses $y \neq \bar{y}$ when all other countries in the world have $\bar{y}$.[16] Then the oil price is pinned down by $\bar{y}$ and given by $q = \bar{q}$. Substituting into (13) now gives

$$\pi_N = \frac{1-\hat{\lambda}}{\hat{\lambda}}\sigma(y - \bar{y}) + \frac{1-\hat{\lambda}}{\hat{\lambda}}\frac{1}{\gamma}(\bar{y} - \bar{x}). \tag{18}$$

The second term is the world Phillips curve, relating inflation to $\bar{y}$ and $\bar{x}$. The first term captures the inflation effects of deviations $y \neq \bar{y}$ at the country-level Phillips curve slope $(1-\hat{\lambda})/\hat{\lambda}\sigma$. In a symmetric equilibrium, $y = \bar{y}$ so the first term vanishes and we recover the world Phillips curve (17). Although the world Phillips curve acts as a constraint on equilibria, an individual country behavior is incentivized by the first term, among other things. The resulting symmetric equilibrium lies on the world Phillips curve, but individual countries act as if it did not have to.

These results are illustrated in Figure 1, where we plot country-level and world-level Phillips curves in an example where (15) holds. The figure shows two country-level Phillips curves: one (in blue, solid) corresponding to a higher level of world demand, and thus a higher oil price; and one (in blue, dashed) corresponding to a lower level of world demand and a lower oil price. A higher oil price appears as an exogenous shift in the country-level curve. In any symmetric world equilibrium, however, the oil price

[16] Alternatively, we can interpret $\bar{y}$ as the average choice of y in the world $\bar{y} = \int y^i di$.

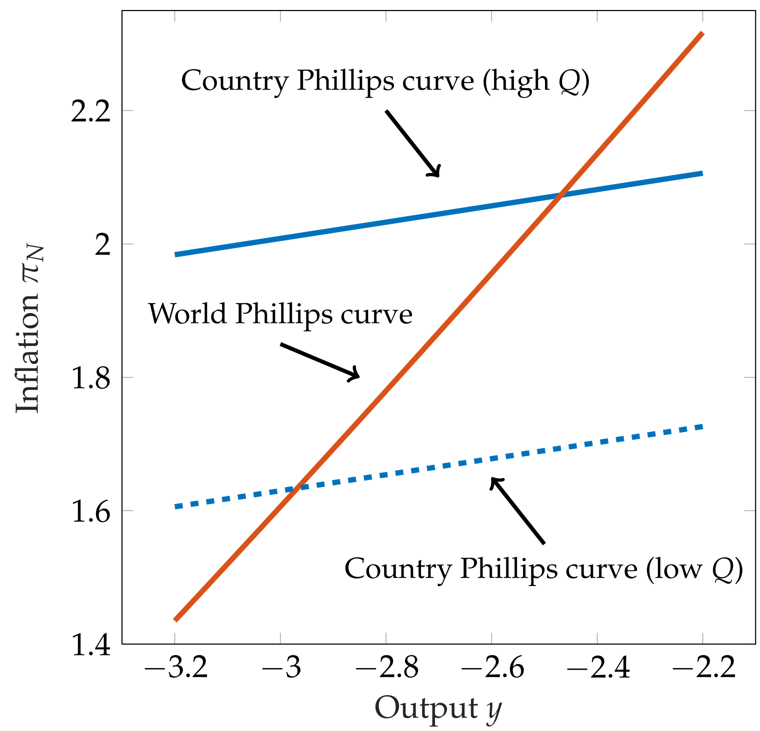

Figure 1: Phillips curves

is endogenous and the economy moves along the steeper world Phillips curve shown in red.[17]

Exchange Rate Role. It is useful to clarify the role of the exchange rate in the comparison between local and global Phillips curves. The exchange rate channel was crucial in our intuition for the slope of the local Phillips curve. Since all countries cannot appreciate at once, one could ask: why the absence of the exchange rate channel at the world level is not making the global Phillips curve flatter than the local one? Recall that the exchange rate channel was captured by the term σy in equation (13). Why is that term still present in equation (17)? The answer is that the term σy in equation (17) captures a different but equivalent effect: it captures the fact that a global planner controls the nominal world prices P_T^* and P_X^*. These prices are taken as given by the individual country, but not by a global planner. When we look at the price of oil in domestic currency

$$P_X = E \cdot P_X^*,$$

the country planner thinks it can change it by moving E, taking P_X^* as given, the world planner thinks it can change it by levering P_X^*, with $E = 1$. Given that E and P_X^* enter

[17]The example uses the same parameters used for Figure 3 below.

perfectly symmetrically in this equation, the two effects are identical. When we go from the local Phillips curve to the global Phillips curve the term σy that was capturing the exchange rate channel, now captures the nominal world oil price channel.

4.3 Ad Hoc Dual Mandate

Before turning to our normative welfare analysis, we discuss an ad hoc version to the policy analysis which provides some insight, despite its shortcomings.

Suppose central banks have some given preferences over inflation and output. For example, consider a quadratic loss function

$$(y - y^*)^2 + \lambda(\pi_N - \pi^*)^2$$

for some targets y^*, π^* which may potentially respond to the $\bar{x}$ shock.

Suppose individual countries and the world planner have the same loss function. In the uncoordinated equilibrium, each country chooses y to minimize the loss function subject to (18), taking world output $\bar{y}$ as given. The symmetric equilibrium then occurs when $y = \bar{y}$. The coordinated equilibrium minimizes the loss function subject to the same constraint but optimizing over $\bar{y}$ and taking into account $y = \bar{y}$. In other words, it considers the world Phillips curve (17). Then whenever (15) holds, so that the world Phillips curve is steeper, the uncoordinated equilibrium will feature an expansionary bias, with higher output and inflation, as compared to the coordinated equilibrium.

The loss function may be thought of as capturing a dual mandate in a reduced-form manner. In simple closed-economy models quadratic loss functions can be justified as proper second order approximations to welfare. However, this is not the case in our model, so we must go beyond this reduced form analysis.

5 Expansionary Bias and Gains from Coordination

The previous section showed that disinflationary benefits of a contraction in output are smaller for an individual country than for the world as a whole. This section examines under what conditions this leads to an expansionary bias, with higher output and inflation as compared to the coordinated outcome.

5.1 No Coordination

Using the equilibrium conditions (3)–(10), we can express the levels of domestic equilibrium variables as functions of the level of local activity Y and of the world oil price Q. Let

$$X = \mathcal{X}(Y, Q), \qquad L = \mathcal{L}(Y, Q), \qquad P_N = \mathcal{P}(Y, Q)$$

denote the equilibrium oil consumption, employment, and the price of non-tradable goods.

In a world with no coordination, the domestic central bank's best response, that is, the level of activity Y chosen for a given world oil price Q, is characterized compactly as the solution to the following problem:

$$\mathcal{Y}(Q) = \arg\max_{Y} \mathcal{U}(Y, Q). \tag{19}$$

where

$$\mathcal{U}(Y, Q) \equiv U(\bar{Y}_T - Q \cdot (\mathcal{X}(Y, Q) - \bar{X}), Y, \mathcal{L}(Y, Q)).$$

Substituting this best response in the oil demand function, and imposing market clearing in the world oil market, gives

$$\mathcal{X}(\mathcal{Y}(Q), Q) = \bar{X}. \tag{20}$$

The two conditions (19)-(20) fully characterize a symmetric Nash equilibrium.

To further analyze the tradeoff faced by the individual central bank, let us derive the first order condition of the maximization problem in (19). Omitting the functions' arguments for ease of notation, and, using subscripts to denote partial derivatives, we have

$$U_{C_N} + U_L \cdot \mathcal{L}_Y - Q \cdot U_{C_T} \cdot \mathcal{X}_Y = 0.$$

To interpret this expression and similar expressions to come, it is useful to recall that $U_L < 0$, capturing the marginal disutility of labor.

If we substitute for $\mathcal{L}_Y$ in the expression above, we obtain the following decomposition that is easier to interpret (the necessary steps are in the appendix):

$$\underbrace{(U_{C_N} + \frac{\Delta}{F_L} U_L)}_{\text{Labor wedge}} + \underbrace{U_L \frac{\Delta'(P_N) Y}{F_L}}_{\text{Inflation distortion}} \cdot \mathcal{P}_Y + \underbrace{(-U_L \frac{F_X}{F_L} - Q \cdot U_{C_T})}_{\text{Oil wedge}} \cdot \mathcal{X}_Y = 0. \tag{21}$$

The first two terms on the left-hand side are standard.

The first captures a notion of output gap. When the expression is positive, there is a negative output gap: increasing output is welfare improving as the marginal utility of consumption exceeds the marginal labor cost.

The second term captures the welfare cost of inflation, due to the distortion Δ. Notice that Δ can be written solely in terms of P_N, by defining the function,[18]

$$\Delta(P_N) \equiv P_N^{\varepsilon}\left[\lambda + (1-\lambda)^{\frac{1}{1-\varepsilon}}(P_N^{1-\varepsilon} - \lambda)^{-\frac{\varepsilon}{1-\varepsilon}}\right]. \tag{22}$$

The derivative of this function appears in (21). When $P_N > 1$ and inflation is positive we have $\Delta'(P_N) > 0$. Increasing output increases the inflation distortion, given that the partial derivative $\mathcal{P}_Y$ is positive. This lowers welfare because $U_L < 0$.

The third term in (21) is novel and captures the gap between the social benefits of higher oil consumption and the cost of importing oil. Suppose inequality (15) holds, which implies $\mathcal{X}_Y > 0$. Oil imports are socially beneficial, as more oil reduces the need for labor in production in proportion to the marginal rate of substitution between oil and labor F_X/F_L. The cost of oil is given by the loss of tradable consumption, as the country needs to pay for the oil. The difference between these two expressions captures the net social benefit of increasing oil demand via expansionary policy. We call this term the "oil wedge". The sign of this wedge can in general be positive or negative. In our examples below it is typically negative, so the central bank perceives an incentive to reduce domestic oil consumption.

Wage rigidity is important in determining the presence of an oil wedge. With flexible wages the oil wedge is always zero. With flexible wages the real wage in tradable goods W/P_T is equalized to the ratio $-U_L/U_{C_T}$ by workers' optimality. Since the ratio of the factor prices QP_T/W is equal to the factors' marginal rate of substitution, by firms' optimality, we get

$$-\frac{U_L}{U_{C_T}} = \frac{W}{P_T} = \frac{F_L}{F_X}Q,$$

which implies a zero oil wedge.

By a similar line of reasoning, when wages are rigid, the oil wedge always has the opposite sign of the expression:

$$\frac{W}{P_T} + \frac{U_L}{U_{C_T}},$$

which is the gap between the real wage and its' flexible price counterpart. When this expression is positive firms perceive an over-priced labor input and tend to use too much oil from a social perspective. This makes the oil wedge negative.

[18]Inverting condition (4) gives

$$\tilde{P}_N = \left(\frac{P_N^{1-\varepsilon} - \lambda}{1-\lambda}\right)^{\frac{1}{1-\varepsilon}}.$$

Substituting in the expression for Δ gives the function $\Delta(P_N)$.

5.2 Coordination and Inefficiency

We now turn to the case of coordinated macroeconomic policies.

The world social planner solves the same problem of an individual country's central bank, but with a crucial difference. In an uncoordinated equilibrium, the market clearing condition (20) is imposed after countries optimize, as a fixed point condition. In a coordinated equilibrium, the same condition becomes a constraint on the world planner's problem.

The planner's problem can be written as

$$\max_{Y,Q}\{\mathcal{U}(Y,Q) \text{ s.t. } Q=\mathcal{Q}(Y,\bar{X})\}.$$

The objective function is identical to the one in the country problem (19), but Q is now a choice variable and the world oil market clearing condition is an additional constraint—implicit in the equilibrium mapping $\mathcal{Q}$. Namely, $\mathcal{Q}(Y,\bar{X})$ gives the equilibrium oil price and is defined as the Q that solves the market clearing condition $\mathcal{X}(Y,Q)=\bar{X}$. It is the nonlinear analog of (16).

Notice that the feasibility conditions for the world and country planner are the same. The only difference is that the world planner internalizes the oil market clearing condition. As noticed above, this implies that social welfare under coordination is necessarily higher or equal than in the uncoordinated equilibrium. When it is higher we say that there are gains from coordination.

Proceeding as in the case of the individual country, we can derive the optimality conditions for the world planner and, after substituting and rearranging, obtain the single optimality condition

$$\underbrace{(U_{C_N}+\frac{\Delta}{F_L}U_L)}_{\text{Labor wedge}}+\underbrace{U_L\frac{\Delta'(P_N)Y}{F_L}}_{\text{Inflation distortion}}\cdot(\mathcal{P}_Y+\mathcal{P}_Q\mathcal{Q}_Y)+$$
$$+\underbrace{(-U_L\frac{F_X}{F_L}-Q\cdot U_{C_T})}_{\text{Oil wedge}}\cdot(\mathcal{X}_Y+\mathcal{X}_Q\mathcal{Q}_Y)=0. \tag{23}$$

The optimality conditions (21) and (23) can now be easily compared. The labor wedge appears in the same way in both. However, the world planner's optimality contains two additional elements, in the second and the third component, both including the expression $\mathcal{Q}_Y$. This is intuitive: the world planner internalizes the welfare effects caused by the endogenous response of the oil price, captured by $\mathcal{Q}_Y$.

Let us now discuss in more detail these two additional terms and their welfare interpretation.

First, individual countries fail to internalize the benefit of reducing world inflation through lower oil prices, as they ignore the expression $\mathcal{P}_Q \mathcal{Q}_Y$. This expression captures exactly the difference between the slopes of the local and global Phillips curves, so the discussion in the previous section is immediately relevant. As shown in Proposition 1, condition (15) implies $\mathcal{P}_Q \mathcal{Q}_Y > 0$. If that condition holds, countries underestimate the social benefit of world-wide restrictive policy through their effects on the oil price and inflation.

However, there is a second difference between the world and the country optimality conditions, captured in the third term in (21) and (23): the oil wedge term. The definition of $\mathcal{Q}$ implies $\mathcal{X}_Y + \mathcal{X}_Q \mathcal{Q}_Y = 0$, which implies that the oil wedge gets a zero weight in the world planner's marginal calculation. The explanation is intuitive. The world planner internalizes the fact that the world as a whole cannot change its oil consumption. On the other hand, $\mathcal{X}_Y$ is in general non-zero and the oil wedge gets a non-zero weight in the country planner's calculation.

In the examples below, the oil wedge is negative in equilibrium and $\mathcal{X}_Y > 0$. All else equal, this pushes the country central bank to be relatively more contractionary in reaction to the oil shock.

Whether the uncoordinated equilibrium displays excessively expansionary policy or the opposite, depends on the sign and magnitude of the two differences just discussed, as summarized in the following proposition.

Proposition 2. *(Expansionary bias) In the no coordination equilibrium real output and inflation are inefficiently high iff*

$$\left[U_L \frac{\Delta'(P_N) Y}{F_L} \cdot \mathcal{P}_Q + \left(-U_L \frac{F_X}{F_L} - Q \cdot U_{C_T}\right) \cdot \mathcal{X}_Q \right] \mathcal{Q}_Y < 0. \tag{24}$$

Small shocks. The result above is stated in terms of equilibrium allocations and holds generally for any size of the oil shock $\bar{x}$. We now focus on the case of a small $\bar{x}$, to derive conditions on the model primitives that imply condition (24).

Proposition 3. *(Expansionary bias, local conditions) Assume prices are neither fully rigid or fully flexible $\lambda \in (0, 1)$ and $s_X > 0$. The non-coordinated response to a small negative supply shock $\bar{x} < 0$ displays an expansionary bias if and only if:*

$$[\varepsilon(1 - \lambda) - \eta][\sigma(1 - (1 - \lambda)s_X) - \eta s_L] > 0. \tag{25}$$

The proof is in the appendix and is based on linearizing the world planner's optimality conditions near the steady state and checking the sign of inequality (24). The two expressions in square brackets on the left-hand side of (25) control the signs of two factors on the left-hand side of (24): the first expression controls the sign of the long expression in square brackets in (25), the second expression simply controls the sign of $\mathcal{Q}_Y$. Notice, that this second expression in (25) is the same as our familiar condition (15), which controls the relative slope of the country and world Phillips curves.

An expansionary bias can emerge under two different parameter configurations: one with a low elasticity η, in which both factors in (25) are positive, and one with a high elasticity η, in which they are both negative. Therefore, Proposition (3) implies that the condition on the Phillips curves' relative slopes is neither necessary nor sufficient for an expansionary bias.

However, as argued above, we believe the empirically relevant case to be the one in which condition (15) holds. Therefore, let us focus our interpretation on the case in which an expansionary bias arises from inequality (15) and

$$\varepsilon(1-\lambda) - \eta > 0.$$

To capture the intuition for this inequality, remember from the discussion above that the inflation distortion tends to make the world planner more contractionary. The inflation distortion is larger when the goods are more substitutable and when prices are less sticky, which explains why the inequality is more easily satisfied when ε is large and λ is small. Moreover, when η is low, a given expansion in world output has a larger effect on the marginal product of labor and thus on marginal costs, for a fixed oil supply. This magnifies the inflation effect not internalized by the country—as the country always perceives a perfectly elastic oil supply. Summing up, condition $\varepsilon(1-\lambda) - \eta > 0$ ensures that the inflation distortion term dominates the oil wedge term.

Concave Utility in Tradable Goods. We assumed linear utility in C_T to simplify the derivations, but the analysis could be equally developed in the case in which the per-period preferences took the form $u(C_T) + u(C_N) - v(L)$. The main difference is in the treatment of the trade surplus or deficit due to oil trade in period 0. With concave preferences the single country will consider the possibility of borrowing to finance an oil deficit $Q \cdot (X - \bar{X})$ and the problem of the individual country will be to choose Y to maximize

$$u(Y) - v(\mathcal{L}(Y, Q)) + V(Q \cdot (X - \bar{X}))$$

where V is a value function that captures the optimal choice of the path $C_{T0}, C_{T1}, \ldots$ given the trade deficit $Q \cdot (X - \bar{X})$.[19] The only other difference is that the function $\mathcal{L}$ is different from the quasi-linear case, due to a different optimality condition for the relative demand of traded and non-traded consumption goods. But analog equations to the first-order conditions 21 and 23 can be derived and would contain the same forces at play.

5.3 Flexible Prices or Flexible Wages: Divine Coincidence

The combination of price and wage rigidities is crucial to produce a need for international coordination. If we allow for either flexible prices or flexible wages the following proposition shows that in equilibrium condition (24) becomes an equality and the Nash equilibrium is efficient.

Proposition 4 (Divine coincidence, no need for coordination.)**.** *With either flexible prices or flexible wages, the Nash equilibrium yields the first best allocation and coincides with the world planner optimum.*

This result is independent of the slopes of the world and country Phillips curves. When divine coincidence holds, both Phillips curves intersect at the natural allocation with zero inflation, so the slope of the trade off is irrelevant.[20]

5.4 Examples

We now parametrize preferences and technology to construct examples. The utility function is

$$U(C_T, C_N, L) = C_T + \ln C_N - \frac{\Psi}{1+\phi} L^{1+\phi},$$

and the production function is

$$F(L, X) = \left(a_L^{\frac{1}{\eta}} L^{1-\frac{1}{\eta}} + a_X^{\frac{1}{\eta}} X^{1-\frac{1}{\eta}}\right)^{\frac{1}{1-\frac{1}{\eta}}}.$$

[19]Given that the world interest rate in tradable goods is $1/\beta$ the function V can be solved in closed form and is

$$V(Q \cdot (X - \bar{X})) = \frac{1}{1-\beta} u(\bar{Y}_T - \frac{1-\beta}{\beta} Q \cdot (X - \bar{X})).$$

[20]Notice also that this result is only meaningful if some nominal rigidity is present in the model, either on the price or wage side. If both prices and wages are perfectly flexible, monetary policy no longer has the power to affect the real allocation, so the problem of coordination is trivial as both planners' choice set is a point.

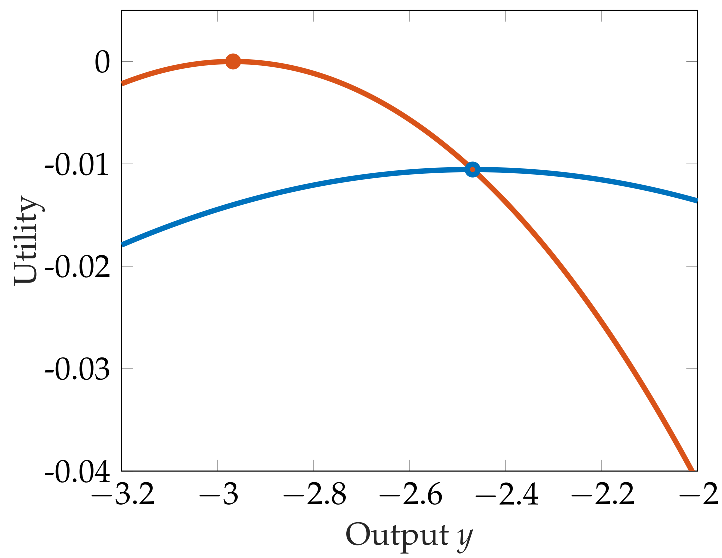

Figure 2: Optimal policy: the country planner at the Nash equilibrium and the world planner

Parameters are chosen to normalize all steady state prices to unity.[21]

Consider an example with the following parameters:

$$a_X = 0.2, \quad \eta = 1/6, \quad \phi = 2, \quad \lambda = 1/2, \quad \varepsilon = 3.$$

In Figure 2 we plot the objective function of the domestic central bank, given Q at its Nash equilibrium level (blue line), and social welfare at a symmetric allocation (red line), as functions of the output level in deviation from the steady state, y. The point where the two curves intersect gives the Nash equilibrium payoff. This example satisfies condition (25) in Proposition 3, and, as the figure shows, social welfare is indeed maximized at an output level below the Nash equilibrium.

To investigate the forces at work, in Figure 3 we plot the first derivative of the objective function of the country planner and of the world planner, computed at a symmetric allocation in which all countries choose the same y. These expressions correspond to the expressions on the right-hand side of equations (21) and (23). We also report the different components labeled in those equations. In the left panel, the Nash equilibrium corresponds to the point where the total first derivative (the blue line) is equal to zero. At that point, the country's central bank balances the marginal cost of keeping output below its natural level (a positive labor wedge, in red), with the marginal benefit of reducing inflation (in yellow), and of reducing an inefficiently high level of oil consumption (in purple).

[21]In particular, we set $\Psi = {a_L}^{-\phi}$ and $\bar{X}_{SS} = a_X$.

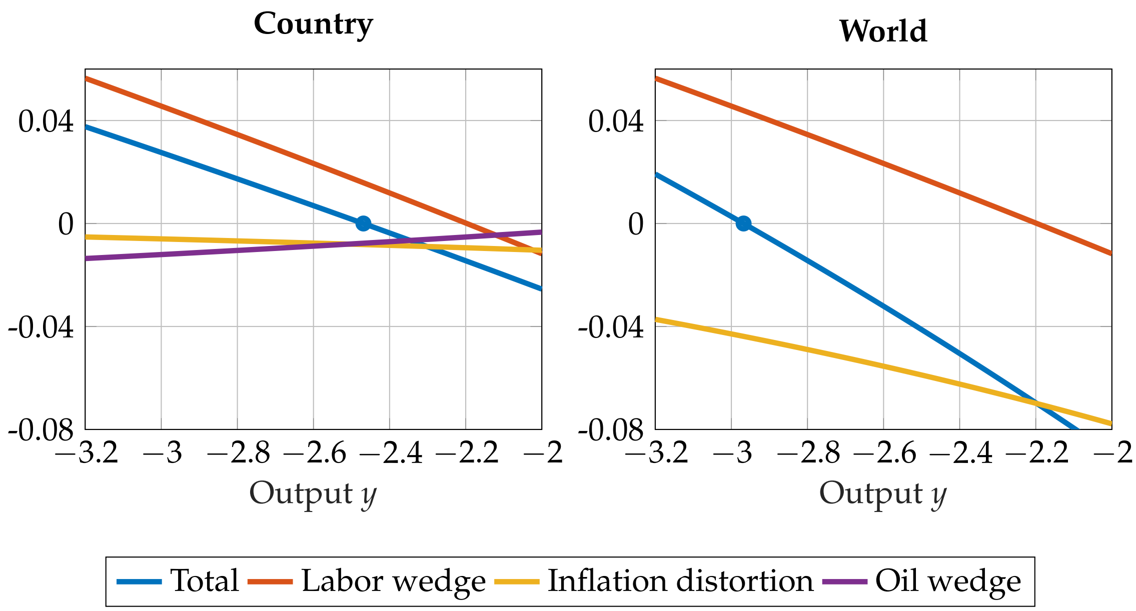

Figure 3: Decomposition of the net marginal benefits of increasing output, for the country and the world planner

In the right panel, we see that the world planner chooses instead the allocation where the total first derivative is zero, with lower output y. While the labor wedge is the same in the two panels, the crucial difference is that the world planner internalizes a much larger cost of inflation distortions (in yellow). Even though in the world planner's problem the oil wedge is absent, the larger cost of inflation distortions is sufficient to make optimal output lower.

In the examples below, the oil wedge is negative in equilibrium and $\mathcal{X}_Y > 0$. All else equal, this pushes the domestic central bank to be relatively more contractionary in reaction to the oil shock. The logic of this negative wedge was discussed above: the real wage is above its natural level, so firms tend to import too much oil from a social welfare point of view, to substitute for labor. Contracting the economy is a way for the domestic central bank to partly correct for this bias, reducing oil imports. Therefore, in those examples there are two opposing forces: the inflation distortion makes the world planner relatively more contractionary, while the oil wedge makes the country planner relatively more contractionary.

The outcomes in terms of inflation can be read from the Phillips curves in Figure 1, which is drawn for the same numerical example analyzed here. The two intersections of the world Phillips curve with the country Phillips curves correspond to the Nash equilibrium (on the right) and to the world planner optimum (on the left). Therefore, in this example, world coordination delivers lower world inflation.

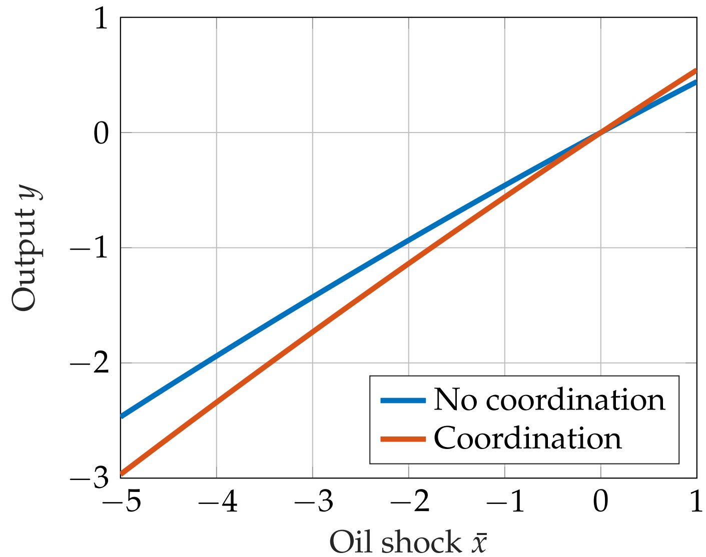

Figure 4: Comparative statics

The example above is for a given oil shock $\bar{x}$. Figure 4 shows comparative statics for different values of $\bar{x}$. With no shock ($\bar{x} = 0$) there is no inefficiency, that is, no gain from coordination. As the size of the shock grows (in absolute value) the distance between the two equilibria grows as well. The figure only looks at negative shocks, but the model behavior is symmetric, so following a positive oil shock we get a contractionary bias, i.e., countries expand too little and there is a too large deflation.

We conclude this section by showing by counterexample that the second condition in Proposition 3 is indeed needed. When the condition is not satisfied, even if the economy displays a steeper world Phillips curve, the planner may prefer a more expansionary stance. In Figure 5, we show an example where the oil wedge is large enough that countries, in Nash equilibrium, have a strong incentive to contract, in fact, too strong from the point of view of a global planner, who would select a larger level of output. The parameters needed to produce a result like this are equal to the ones above except for

$$\eta = 1.2 \qquad s_X = 2/3 \qquad \varepsilon = 2.$$

We view this example as enlightening but empirically implausible: it assumes unrealistically large oil share s_X and a modestly large elasticity of substitution $\eta > 1$. These conditions are fine tuned to obtain a steeper world Phillips curve, yet obtain a contractionary bias.[22]

[22]This does not mean that other relevant forces may lead overall to a contractionary bias in a quantified model. Our statement here is just that this specific way of obtaining a contractionary bias seems quantitatively far fetched.

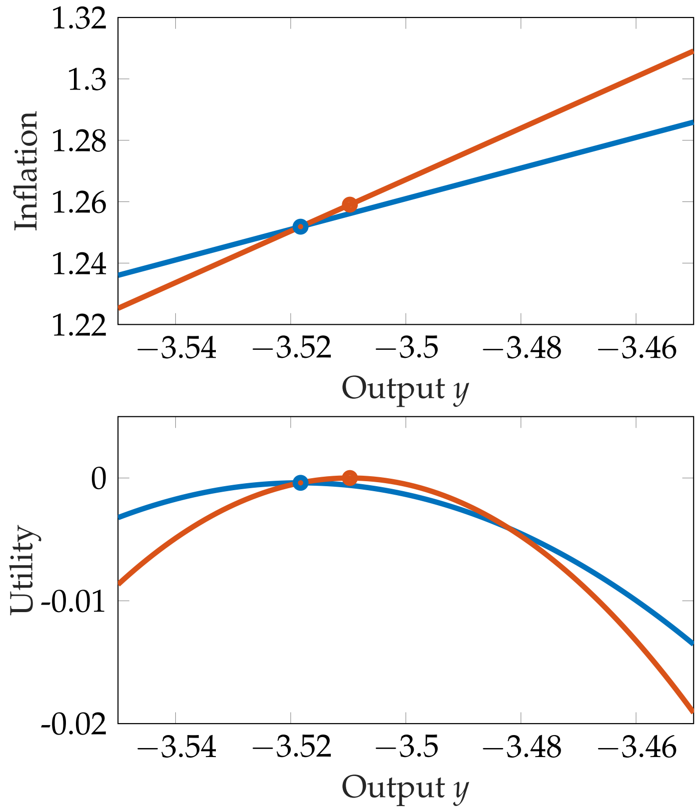

Figure 5: Nash equilibrium and world optimum in an example with a relatively large oil wedge

6 A Model with Terms of Trade Manipulation

In this section, we study a variant of the model in which we have countries producing differentiated tradable goods. The model in this section is similar to the canonical model of Gali and Monacelli (2005) and introduces term-of-trade effects of monetary policy. For convenience, we will make assumptions of unit elasticity of substitution in most places, so this economy has many convenient properties of a Cole and Obstfeld economy.

Term-of-trade effects can in general impart a contractionary bias to monetary policy, as reducing domestic production has the beneficial side effect of making domestically produced goods more expensive in the world. However, as we will show, this is a bias that is present independently of shocks. The question we are asking in this paper is not whether on average monetary policy is too contractionary or too expansionary, but how supply shocks affect this bias. We will show that by allowing for a sufficient set of policy instruments, we can separate these two questions and derive results that are parallel to the ones derived in the economy of Section 2. Namely, we show that it is possible to have an expansionary bias in response to a world supply shock, which comes from the same mechanism identified in the model of the previous sections: a failure to internalize the effects of world demand on the demand for scarce traded inputs.

6.1 Environment

To simplify, we consider a two-period version of the model.

The representative consumer in a typical country has preferences represented by the utility function

$$\ln C - \frac{\Psi}{1+\phi} L^{1+\phi} + Z$$

where C and L are consumption and labor supply in the first period, and Z is consumption of a uniform tradable good in the second period. The consumption index C is a unit-elasticity composite of the home good H and of imported foreign goods,

$$C = (\frac{C_H}{\omega})^{\omega} (\frac{C_F}{1-\omega})^{1-\omega},$$

where C_H is a constant-elasticity index of consumption of a continuum of varieties $j \in [0,1]$ produced in the home economy,

$$C_H = (\int_0^1 \left(C_H(j) \right)^{\frac{\varepsilon-1}{\varepsilon}} dj \right)^{\frac{\varepsilon}{\varepsilon-1}},$$

C_F is a unit-elasticity index of consumption of the goods produced in all other countries $i \in [0,1]$

$$C_F = \exp \int_0^1 \ln C_{Fi} di,$$

and C_{Fi} is a consumption index of the varieties produced in country i imported by the domestic economy

$$C_{Fi} = \left(\int_0^1 C_{Fi}(j)^{1-\frac{1}{\varepsilon}} dj \right)^{\frac{1}{1-\frac{1}{\varepsilon}}}.$$

Each variety j of home good is produced according to the production function

$$Y(j) = F(L(j), X(j))$$

$$F(L, X) \equiv ((1-s_X)^{\frac{1}{\eta}} L^{1-\frac{1}{\eta}} + s_X^{\frac{1}{\eta}} X^{1-\frac{1}{\eta}})^{\frac{1}{1-\frac{1}{\eta}}},$$

Price setting is as in Section 2, with λ firms setting prices before shocks are realized, and $1-\lambda$ after. As in Section 2 we have fully sticky wages: the nominal wage $\bar{W}$ is set before shocks are realized at its steady-state flexible wage level but cannot adjust when shocks are realized.

The following assumptions are all as in Section 2. Each country has a fixed supply of oil, all countries have the same endowment $\bar{X}$,the world oil market is perfectly integrated, and the oil price is perfectly flexible. There are flexible exchange rates, the world unit of account is an equally-weighted basket of all currencies, and the exchange rate is the price of the world unit of account in domestic currency denoted by E.

Let R^* denote the world real interest rate in tradable goods, defined as the price at which one unit of the composite good F (which is the same for all countries) can be converted into a unit of good Z in the second period. The real rate R^* is determined by the collective choices of all central banks, as it depends on nominal interest rates in each currency and by the nominal prices of each local good in terms of the world unit of account. Since each country is small, each country takes R^* as given.

From expenditure minimization, the price index for the composite foreign good C_F, expressed in the world unit of account, is

$$P^* = \exp \int_0^1 \ln(P_i/E_i)di,$$

(just for this expression we index the exchange rate E_i explicitly with the country index i). The price index of good H and the domestic CPI are then respectively

$$P_H = (\int_0^1 P_H\left(j\right)^{1-\varepsilon}dj\right)^{\frac{1}{1-\varepsilon}},$$
$$P = P_H^{\omega}(EP^*)^{1-\omega}.$$

The country's real output is defined as

$$Y = \frac{1}{P_H}\int_0^1 P_H(j)Y(j)dj,$$

and the country-level intertemporal budget constraint is

$$Z = \frac{R^*}{EP^*}\left[P_H Y - PC - EP_X^*(X-\bar{X})\right],$$

where P_X^* is the oil price expressed in the world unit of account and X is the country's oil consumption $X = \int_0^1 X(j)dj$.

6.2 Equilibrium

From intertemporal optimization, we have the Euler equation for consumption of the foreign good

$$\frac{1-\omega}{C_F} = R^*. \tag{26}$$

Optimal relative demand for home and foreign goods gives

$$\frac{C_H}{C_F} = \frac{\omega}{1-\omega}(\frac{P_H}{EP^*})^{-1}. \tag{27}$$

Demand for the home good comes from home and foreign consumers and is equal to

$$Y = C_H + (\frac{P_H}{EP^*})^{-1} C_F^*, \tag{28}$$

where the second term on the right-hand side is derived from the demand of all other countries, and C_F^* denotes the average demand for imported goods in all other countries.

The price setting condition for firms who set their price after the shock is

$$\tilde{P}_H = (1+\mu)((1-s_X)\bar{W}^{1-\eta} + s_X\left(EP_X^*)^{1-\eta}\right)^{\frac{1}{1-\eta}}, \tag{29}$$

where $\bar{W}$ is the steady state wage rate, μ is the markup that satisfies

$$1+\mu = \frac{1}{1+\varsigma}\frac{\varepsilon}{\varepsilon-1},$$

and ς is a proportional subsidy to producers set by the domestic fiscal authority.

The price of good H is then

$$P_H = (\lambda \bar{P}_H^{1-\varepsilon} + (1-\lambda)\tilde{P}_H^{1-\varepsilon})^{\frac{1}{1-\varepsilon}}, \tag{30}$$

where $\bar{P}_H$ is the initial price level, set at its steady state level by firms who expect no shocks.

The distortion function $\Delta(P_H)$ has the same form as (22), except with P_H in place of P_N. The resource constraint is then

$$\Delta(P_H) \cdot Y = F(L, X). \tag{31}$$

Cost minimization by each domestic firm gives the relative demand for oil and labor

$$\frac{X}{L} = \frac{s_X}{1-s_L}(\frac{EP_X^*}{\bar{W}})^{-\eta}. \tag{32}$$

Conditions (26)-(32) give seven equilibrium condition, which determine, up to one degree of freedom, the vector of eight equilibrium variables:

$$(E, \tilde{P}_H, P_H, C_H, C_F, Y, L, X). \tag{33}$$

As in the previous section, there is one degree of freedom and we assume that a benevolent domestic central bank selects the equilibrium vector.

We can now define a Nash equilibrium. The definition is similar to 1, but now four global variables need to be specified: R^*, P^*, P_X^*, C_F^*.

Definition 3. *(No coordination equilibrium)* A symmetric Nash equilibrium is given by a vector

$$(R^*, P^*, P_X^*, C_F^*, E, \tilde{P}_H, P_H, C_H, C_F, Y, L, X)$$

that satisfies two conditions:

1. *Country's optimization.* The vector of local prices and quantities (33) satisfies equilibrium conditions (3)-(10) and, among all vectors of prices and quantities that satisfy those conditions, it maximizes the utility of the representative household in country H, taking (R^*, P^*, P_X^*, C_F^*) as given.

2. *World equilibrium.* The exchange rate satisfies $E = 1$, the price of imports satisfies $P^* = P_H$, foreign demand satisfies $C_F^* = C_F$, and oil consumption satisfies $X = \bar{X}$.

The definition of an equilibrium with coordination is analogous to Definition 2, and we omit it.

Dealing with steady state bias

Take a steady state allocation and suppose that, as we did in Section 3, we set the subsidy ς to achieve a first-best steady-state allocation. Then the following first order condition is satisfied in equilibrium

$$\frac{1}{C} - \frac{1}{F_L}\Psi L^{\phi} = 0. \tag{34}$$

Consider now the marginal benefit of increasing consumption of the home good C_H for a domestic central bank, satisfying the local equilibrium conditions (3)-(10). The marginal benefit for the representative local consumer is now

$$\frac{\omega}{C_H} - \frac{1}{\omega}\frac{1}{F_L}\Psi L^{\phi} = \frac{1}{C} - \frac{1}{\omega}\frac{1}{F_L}\Psi L^{\phi} < 0,$$

where the equality follows because in equilibrium $C = C_H/\omega$, and the inequality follows from (34) and $\omega < 1$. The presence of $1/\omega$ in the last term has a simple interpretation: to stimulate domestic consumption the central bank needs to depreciate the real exchange rate $\frac{P_H}{EP^*}$, but this stimulates foreign spending on domestic goods, so the increase in labor effort needed to satisfy total demand is magnified relative to the increase in domestic consumption. In our simple Cole-Obstfeld economy the magnification is simply given by the factor $1/\omega$.[23] Therefore, if we start at a steady state that is first-best optimal for the global planner, increasing consumption is welfare reducing for a local planner. A domestic central bank will have an incentive to contract the domestic economy.

In other words, the local monetary authority has a deflationary bias in steady state. This bias arises from a terms-of-trade manipulation incentive: reducing domestic activity improves the domestic terms of trade $\frac{P_H}{EP^*}$. This is a well-known source of deflationary bias in open economy new Keynesian models, but it is a bit of a nuisance in the present context, since we are trying to identify the effects of supply shocks, not the steady state inflationary bias of the two regimes.

Fortunately, there is a simple solution to avoid this tension. In the analysis that follows, when we compare the two regimes, no coordination and coordination, we assume that in the uncoordinated case local authorities choose *both* the response of monetary policy to shocks and the level of the fiscal subsidy ς in an uncoordinated way, while in the coordinated case, both are chosen by the world planner. This implies that the steady state is different in the two cases. In the coordination equilibrium the steady state satisfies (34), in the no coordination equilibrium the steady state satisfies

$$\frac{1}{C} - \frac{1}{\omega}\frac{1}{F_L}\Psi L^{\phi} = 0. \tag{35}$$

This condition comes from letting countries choose ς optimally ex ante, taking the ς chosen by all other countries as given, and assigning zero probability to the oil shock.

An immediate implication of this approach is that the benevolent central bank will choose zero inflation in steady state in both no coordination and coordination equilibrium.

[23]Combining equations (27) and (28) we have

$$\frac{Y}{C_H} = \frac{1}{\omega}.$$

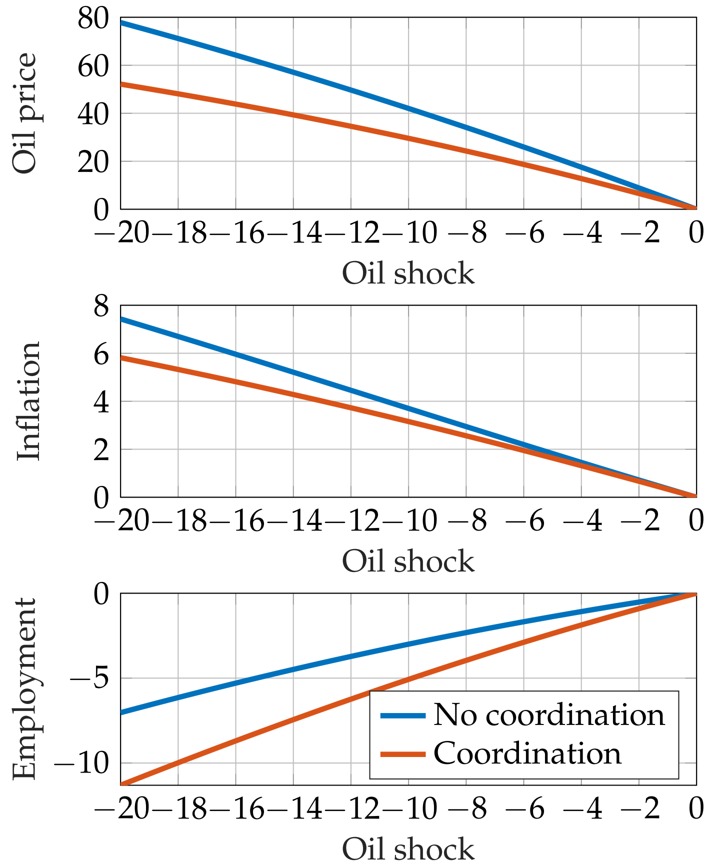

Figure 6: Comparative statics in the model with home and foreign goods

6.3 Supply shocks and expansionary bias

We can now turn to the central question of our paper: is there an expansionary bias following a supply shock?

Let $P_H = \mathcal{P}(Y)$ and $\mathcal{X}(Y)$ denote the equilibrium mappings that give home consumption, the price of the home price good, and oil consumption as functions of home output Y. For economy of notation, we leave implicit the global variables R^*, P^*, P_X^*, C^* and assume they are at their equilibrium level.

The optimality condition for the domestic planner can then be written as follows, as shown in the appendix,

$$\underbrace{\left(\frac{\omega}{C_H} - \frac{1}{\omega}\frac{\Delta}{F_L}\Psi L^{\phi}\right)}_{\text{Labor wedge (local)}} - \underbrace{\Psi L^{\phi} \cdot \frac{\Delta'(P_H)Y}{F_L}}_{\text{Inflation distortion}}\mathcal{P}'(Y) + \underbrace{\left(\Psi L^{\phi}\frac{F_X}{F_L} - \frac{R^* P_X^*}{P^*}\right)}_{\text{Oil wedge}}\mathcal{X}'(Y) = 0.$$

Notice the similarity with the decomposition (21) in Section 5.

From the optimality of the world planner we get the optimality conditions

$$\underbrace{\left(\frac{1}{C} - \frac{\Delta}{F_L}\Psi L^{\phi}\right)}_{\text{Labor wedge (world)}} - \underbrace{\Psi L^{\phi} \frac{\Delta'(P)C}{F_L}}_{\text{Inflation distortion}} \cdot \mathcal{P}^{*\prime}(Y) = 0$$

where the mapping $\mathcal{P}^*(Y)$ captures the relation between aggregate activity and the world CPI in a symmetric world equilibrium, and it includes by definition the endogenous effects of oil prices on inflation. Notice again the similarity with 23 in Section 5.

Both the domestic central bank and the world planner are trading off distortions in the aggregate labor allocation against inflation distortions. However, the domestic central bank takes into account the additional effects of its choices on domestic oil imports (see the presence of the oil wedge), while the world central bank internalizes the equilibrium in world oil markets and thus uses a different relation between aggregate activity and inflation (see the different inflationary effects in the second element of our decompositions).

We now present a numerical example based on the following parameters:

$$s_X = 0.2, \quad \eta = 1/6, \quad \omega = 0.8, \quad \phi = 2, \qquad \lambda = 1/2, \quad \varepsilon = 3.$$

All other parameters are set so that prices are unity at the steady state.[24]

Figure 6 illustrates what happens to the oil price P_X^*, to CPI inflation P, and to employment in the two regimes considered, as a function of the oil price shock $\bar{x}$. All variables are expressed in log deviations from the steady state.[25]

By construction, in steady state both the local planner and the world planner choose to implement the zero inflation allocation. However, for negative oil shocks, both the local and world planner face a trade off between reducing employment below its efficient level and containing inflation. In the example, the inflation distortion perceived by the world planner is large enough, due to the endogeneity of the oil price, that the world planner chooses a larger contraction and lower world inflation.

Overall, the example shows a result that is qualitatively similar to the one analyzed in Section 5 for the model with tradable and non-tradable goods. In response to negative oil shocks a world planner would recommend a tighter monetary stance to contain inflation.

[24]The parameter Ψ is set equal to $(1 - s_X)^{-\phi}$ and the oil supply in steady state is set to $\bar{X} = s_X$, so that in the first-best steady state all prices and consumption levels are equal to 1, while $L = 1 - s_X$ and $X = s_X$.

[25]The equilibrium allocations are computed numerically using the equilibrium relations and the planner's optimality conditions in their non-linear form.

7 Conclusions

This paper challenges the conventional wisdom that uncoordinated monetary policy leads to excessive tightening in response to global negative supply shocks. Indeed, we show that when a global supply shock drives inflation up symmetrically in all countries, uncoordinated monetary policy can be too loose. Our result moves the focus away from traditional terms-of-trade considerations, and places emphasis on the endogenous prices of globally traded, non-labor inputs.

References

Auclert, Adrien, Matt Rognlie, and Ludwig Straub, "Managing an Energy Shock: Fiscal and Monetary Policy," in Sofía Bauducco, Andrés Fernández, and Giovanni L. Violante, eds., *Heterogeneity in Macroeconomics: Implications for Monetary Policy,*, Vol. Santiago, Chile,, Banco Central do Chile, 2024, pp. pp. 39–108.

Benigno, Gianluca and Pierpaolo Benigno, "Designing targeting rules for international monetary policy cooperation," *Journal of Monetary Economics*, 2006, *53* (3), 473–506.

Bianchi, Javier and Louphou Coulibaly, "Financial integration and monetary policy coordination," Technical Report, National Bureau of Economic Research 2024.

Canzoneri, Matthew B. and Dale W. Henderson, *Monetary Policy in Interdependent Economies: A Game-Theoretic Approach*, The MIT Press, 1991.

Corsetti, Giancarlo and Paolo Pesenti, "Welfare and macroeconomic interdependence," *The Quarterly Journal of Economics*, 2001, *116* (2), 421–445.

Drechsel, Thomas, Michael McLeay, Silvana Tenreyro, and Enrico D Turri, "Optimal monetary policy and exchange rate regimes in commodity-exposed economies," 2025.

Engel, Charles, "Currency misalignments and optimal monetary policy: a reexamination," *American Economic Review*, 2011, *101* (6), 2796–2822.

Farhi, Emmanuel and Ivan Werning, "Dilemma not trilemma? Capital controls and exchange rates with volatile capital flows," *IMF Economic Review*, 2014, *62* (4), 569–605.

Fornaro, Luca and Federica Romei, "Monetary policy during unbalanced global recoveries," 2022.

Gali, Jordi and Tommaso Monacelli, "Monetary policy and exchange rate volatility in a small open economy," *The Review of Economic Studies*, 2005, *72* (3), 707–734.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Ivan Werning, "Monetary Policy in Times of Structural Reallocation," 2021. ¡em¿Proceedings of the Jackson Hole Symposium¡/em¿, 2021.

Hamada, Koichi, "A strategic analysis of monetary interdependence," *Journal of Political Economy*, 1976, *84* (4, Part 1), 677–700.

Itskhoki, Oleg and Dmitry Mukhin, "Optimal exchange rate policy," Technical Report, National Bureau of Economic Research 2023.

Miranda-Pinto, Jorge, Mr Andrea Pescatori, Ervin Prifti, and Guillermo Verduzco-Bustos, *Monetary policy transmission through commodity prices*, International Monetary Fund, 2023.

Obstfeld, Maurice, "Uncoordinated monetary policies risk a historic global slowdown," Blog, Peterson Institute for International Economics 2022.

— **and Kenneth Rogoff**, "Global implications of self-oriented national monetary rules," *The Quarterly journal of economics*, 2002, *117* (2), 503–535.

Oudiz, Gilles and Jeffrey Sachs, "Macroeconomic policy coordination among the industrial economies," *Brookings Papers on Economic Activity*, 1984, *1984* (1), 1–75.

A Appendix

A.1 Equilibrium

The system of equations (3)-(10) characterizes an equilibrium for given Y and Q. In this section, we derive a first-order approximation of the equilibrium conditions, computed in general at any equilibrium allocation. This approximation is then used to:

1. Derive a linear approximation of the equilibrium near the initial steady state;
2. Derive the derivatives of the mappings $\mathcal{L}(Y, Q)$ and $\mathcal{X}(Y, Q)$ in general, possibly away from the initial steady state.

For a generic variable X, we use the notation

$$x = \ln X - \ln X_{SS},$$

where X_{SS} denotes the steady state level and x log deviations from the steady state.

The labor share and the oil share are

$$s_L = \frac{LF_L(L, X)}{F(L, X)}, \quad s_X = \frac{XF_X(L, X)}{F(L, X)}. \tag{36}$$

Throughout the appendix s_L and s_X are in general to be interpreted as functions of L and X. In the text, they denote their values at the initial steady state.

Combining (3) and (4) gives

$$P_N^{1-\varepsilon} = \lambda + (1 - \lambda)(\mathcal{M}(1, P_X))^{1-\varepsilon}.$$

Differentiating both sides and rearranging, using differential notation, we obtain

$$dp_N = h \cdot dp_X \tag{37}$$

where

$$h = (1 - \lambda)(\frac{\tilde{P}_N}{P_N})^{1-\varepsilon} \frac{P_X \mathcal{M}_{P_X}}{\mathcal{M}}.$$

Rearranging (4) we have

$$(1 - \lambda)(\frac{\tilde{P}_N}{P_N})^{1-\varepsilon} = 1 - \lambda P_N^{\varepsilon-1},$$

and standard properties of the cost function imply

$$\frac{P_X \mathcal{M}_{P_X}}{\mathcal{M}} = s_X.$$

We can then rewrite h as a function of P_N and s_X:

$$h = (1 - \lambda P_N^{\varepsilon-1}) s_X. \tag{38}$$

Differentiating the resource constraint (7) gives

$$Y\Delta'(P_N)dP_N + \Delta(P_N)dY = F_L(L,X)dL + F_X(L,X)dX,$$

dividing, side by side, by (7) and rearranging gives

$$\delta(P_N)dp_N + dy = s_L dl + s_X dx, \tag{39}$$

where δ is the elasticity of the distortion function

$$\delta(P_N) = \frac{P_N \Delta'(P_N)}{\Delta(P_N)}. \tag{40}$$

Differentiating (8)-(10) gives

$$dp_T = dp_N + \sigma dy, \tag{41}$$

$$dx - dl = -\eta \cdot dp_X, \tag{42}$$

$$dy = -\frac{Q(X-\bar{X})}{C_T}dq - \frac{QX}{C_T}dx, \tag{43}$$

$$dp_X = dq + dp_T. \tag{44}$$

Substituting (41) and (44) in (37) gives

$$dp_N = h \cdot (dq + dp_N + \sigma dy)$$

and solving for dp_N

$$dp_N = \frac{h}{1-h} \cdot (\sigma dy + dq). \tag{45}$$

Computing this expression at the steady state gives the country Phillips curve (13) in the text, as, at the steady state, $h = 1 - \hat{\lambda}$.

Substituting back in (41) and (44) we get the other two prices,

$$dp_T = dp_X = \frac{1}{1-h} \cdot (\sigma dy + dq). \tag{46}$$

From (39) we obtain

$$dl = dy + \delta \cdot dp_N - s_X(dx - dl),$$

and from (42) and (46) we obtain

$$dx - dl = -\eta \frac{1}{1-h} \cdot (\sigma dy + dq).$$

Combingin the last two equations and rearranging eventually gives

$$dl = (1 + \sigma \frac{\eta s_X + \delta h}{1-h}) dy + \frac{\eta s_X + \delta h}{1-h} dq, \tag{47}$$

$$dx = (1 - \sigma \frac{\eta s_L - \delta h}{1-h}) dy - \frac{\eta s_L - \delta h}{1-h} dq. \tag{48}$$

The last equation computed at the steady state and again using $h = 1 - \hat{\lambda}$ gives the oil demand equation (14).

From conditions (47)-(48) we obtain the following properties of the mappings $\mathcal{L}$ and $\mathcal{X}$ that are useful in the analysis of optimal policy:

$$\mathcal{L}_Y = (1 + \sigma \frac{\eta s_X + \delta h}{1-h}) \frac{L}{Y},$$

$$\mathcal{L}_Q = \frac{\eta s_X + \delta h}{1-h} \frac{L}{Q}, \tag{49}$$

$$\mathcal{X}_Y = (1 - \sigma \frac{\eta s_L - \delta h}{1-h}) \frac{X}{Y},$$

$$\mathcal{X}_Q = -\frac{\eta s_L - \delta h}{1-h} \frac{X}{Q}. \tag{50}$$

A.2 Derivation of equations (21) and (23)

Totally differentiating the resource constraint

$$\Delta(P_N) Y = F(L, X),$$

yields

$$\Delta'(P_N) Y \mathcal{P}_Y + \Delta = F_L(L, X) \mathcal{L}_Y + F_X \mathcal{X}_Y,$$

which can be rearranged to give

$$\mathcal{L}_Y = \frac{1}{F_L} \Delta'(P_N) Y \mathcal{P}_Y + \frac{1}{F_L} \Delta - \frac{F_X}{F_L} \mathcal{X}_Y.$$

Similar steps, differentiating with respect to Q, give

$$\mathcal{L}_Q = \frac{1}{F_L} \Delta'(P_N) Y \mathcal{P}_Q - \frac{F_X}{F_L} \mathcal{X}_Q.$$

Substituting for $\mathcal{L}_Y$ in the country optimality condition yields

$$U_{C_N} + \frac{1}{F_L} U_L (\Delta'(P_N) Y \mathcal{P}_Y + \Delta - F_X \mathcal{X}_Y) - Q \cdot U_{C_T} \cdot \mathcal{X}_Y = 0$$

which, after rearranging, yields (21) in the text.

The world planner optimality condition is

$$U_{C_N} + U_L \cdot (\mathcal{L}_Y + \mathcal{L}_Q \mathcal{Q}_Y) - Q \cdot U_{C_T} \cdot (\mathcal{X}_Y + \mathcal{X}_Q \mathcal{Q}_Y) = 0.$$

Substituting the expression for $\mathcal{L}_Y$ and $\mathcal{L}_Q$ derived above yields (23).

A.3 Divine coincidence at the first-best allocation (Proposition 4)

We show that if nominal wages are set at their flexible wage level, the coordinated and uncoordinated equilibrium coincide with the first best allocation. This shows that there are no gains of coordination at the steady state, and also provides the argument for Proposition 4.

At the first best efficient allocation the following two conditions are satisfied

$$U_{C_N} F_L + U_L = 0, \tag{51}$$

$$Q = \frac{U_{C_N}}{U_{C_T}} F_X, \tag{52}$$

and the real wage $W/P_N = F_L$ supports this allocation.

Totally differentiating the resource constraint gives

$$Y\Delta'(P_N)dP_N + \Delta(P_N)dY = F_L(L,X)dL + F_X(L,X)dX.$$

Therefore, at an allocation with $P_N = 1$, $\Delta(P_N) = 1$, and $\Delta'(P_N) = 0$ we have

$$1 = F_L(L,X)\mathcal{L}_Y + F_X(L,X)\mathcal{X}_Y \tag{53}$$

and

$$0 = F_L(L,X)\mathcal{L}_Q + F_X(L,X)\mathcal{X}_Q. \tag{54}$$

Using (51) and (52) the optimality condition of the country planner becomes

$$U_{C_N} + U_L \mathcal{L}_Y - Q U_{C_T} \mathcal{X}_Y = U_{C_N} - U_{C_N}(F_L \mathcal{L}_Y + F_X \mathcal{X}_Y) = 0,$$

where the last equality follows from (53). Condition (52) can be used to derive the analogous result for the world planner.

A.4 Proof of Proposition 3

Step 1. The first step is to define two functions ψ and g and derive their properties. Define

$$\psi(l) \equiv \left(\lambda + (1-\lambda)\left(\frac{\bar{W}}{F_L(l,1)}\right)^{1-\varepsilon}\right)^{\frac{1}{1-\varepsilon}},$$

and

$$g\left(l\right) \equiv \frac{F(l,1)}{\Delta(\psi(l))},$$

which give, respectively, the price P_N and output per unit of oil Y/X at a symmetric allocation where $L/X = l$. Differentiating ψ gives

$$\psi'\left(l\right) = -\left(1-\lambda\right)\left(\psi(l)\right)^{\varepsilon}\left(\frac{\bar{W}}{F_L(l,1)}\right)^{1-\varepsilon}\frac{F_{LL}(l,1)}{F_L(l,1)}.$$

From the definition of g we have

$$\ln g(l) = \ln F(l,1) - \ln \Delta(\psi(l)),$$

and differentiating two times gives

$$\frac{g'(l)}{g(l)} = \frac{F_L}{F} - \frac{\Delta'(\psi)}{\Delta(\psi)}\psi'(l),$$

$$\frac{g''\left(l\right)}{g\left(l\right)} - \left(\frac{g'\left(l\right)}{g}\right)^2 = \frac{F_{LL}}{F} - \left(\frac{F_L}{F}\right)^2 - \frac{\Delta''\left(\psi\right)}{\Delta\left(\psi\right)}\left(\psi'\left(l\right)\right)^2 + \left(\frac{\Delta'\left(\psi\right)}{\Delta\left(\psi\right)}\right)^2\psi'\left(l\right) - \frac{\Delta'\left(\psi\right)}{\Delta\left(\psi\right)}\psi''\left(l\right).$$

Recall that at the steady state we have

$$\psi\left(l_{SS}\right) = 1, \quad \Delta(\psi(l_{SS})) = 1, \quad \Delta'(\psi(l_{SS})) = 0, \quad \bar{W} = F_L(l_{SS},1).$$

Using these properties in the expressions above, we obtain

$$\psi'(l_{SS}) = -(1-\lambda)\frac{F_{LL}(l_{SS},1)}{F_L(l_{SS},1)},$$

$$g'(l_{SS}) = F_L(l_{SS},1),$$

$$g''(l_{SS}) = F_{LL}(l_{SS},1) - F(l_{SS},1)\Delta''(1)\left(\psi'(l_{SS})\right)^2.$$

The properties of constant returns to scale functions imply that, at the steady state,

$$\frac{F_{LL}}{F_L} = -\frac{X}{L}s_X\frac{1}{\eta}.$$

Differentiating twice the distortion function (22) gives

$$\Delta'(P_N) = \varepsilon P_N^{\varepsilon-1}\left[\lambda + (1-\lambda)^{\frac{1}{1-\varepsilon}}(P_N^{1-\varepsilon}-\lambda)^{-\frac{\varepsilon}{1-\varepsilon}}\right] +$$
$$-\varepsilon(1-\lambda)^{\frac{1}{1-\varepsilon}}(P_N^{1-\varepsilon}-\lambda)^{-\frac{1}{1-\varepsilon}},$$

and

$$\Delta''(P_N) = \varepsilon(\varepsilon-1)P_N^{\varepsilon-2}\left[\lambda + (1-\lambda)^{\frac{1}{1-\varepsilon}}(P_N^{1-\varepsilon}-\lambda)^{-\frac{\varepsilon}{1-\varepsilon}}\right] + \\ -\varepsilon^2 P_N^{\varepsilon-1}(1-\lambda)^{\frac{1}{1-\varepsilon}}(P_N^{1-\varepsilon}-\lambda)^{-\frac{\varepsilon}{1-\varepsilon}-1}P_N^{-\varepsilon} + \\ +\varepsilon(1-\lambda)^{\frac{1}{1-\varepsilon}}(P_N^{1-\varepsilon}-\lambda)^{-\frac{1}{1-\varepsilon}-1}P_N^{-\varepsilon},$$

which, computed at the steady state, give

$$\Delta'(1) = 0, \quad \Delta''(1) = \varepsilon\frac{\lambda}{1-\lambda}.$$

Substituting in the expressions above, we finally obtain

$$\psi'(l_{SS}) = (1-\lambda)\frac{1}{l_{SS}}s_X\frac{1}{\eta},$$

and

$$g''(l_{SS}) = -F_L(l_{SS},1)\frac{1}{l_{SS}}\frac{s_X}{\eta}\left[1+\varepsilon\lambda(1-\lambda)\frac{1}{\eta}\frac{s_X}{s_L}\right]. \tag{55}$$

Step 2. In this step we derive the slope of the policy function of the world planner expressed in terms of the optimal labor to input ratio l as a function of the oil supply X. Using the function g the problem of the world planner can be stated compactly as choosing the labor to input ratio l to solve

$$\max_l u\left(g(l)X\right) - v\left(lX\right).$$

Define the function

$$H\left(l,X\right) \equiv u'\left(g\left(l\right)X\right)g'\left(l\right) - v'\left(lX\right)$$

so the optimality condition of the world planner is $H\left(l,X\right)=0$. Differentiating H yields

$$H_l = u''\left(g\left(l\right)X\right)\left(g'\left(l\right)\right)^2 X + u'\left(g\left(l\right)X\right)g''\left(l\right) - v''\left(lX\right)X,$$
$$H_X = u''\left(g\left(l\right)X\right)g\left(l\right)g'\left(l\right) - v''\left(lX\right)l.$$

Using the expression for g'' derived in step 1, in equation (55), and defining

$$\phi \equiv \frac{v''(L_{SS})L_{SS}}{v'(L_{SS})},$$

we can compute the expressions for H_l and H_X at the steady state:

$$H_l(l_{SS},X_{SS}) = -v'\left(L_{SS}\right)\frac{1}{l_{SS}}\left(\sigma s_L + \phi + s_X\frac{1}{\eta}\left(1+\varepsilon\lambda\left(1-\lambda\right)\frac{1}{\eta}\frac{s_X}{s_L}\right)\right),$$
$$H_X(l_{SS},X_{SS}) = -\frac{v'(L_{SS})}{X_{SS}}(\sigma+\phi).$$

By the implicit function theorem, the elasticity of the world planner's policy function $l(X)$ at the steady state is

$$\frac{l'(X_{SS})X_{SS}}{l(X_{SS})} = -\frac{\sigma + \phi}{\sigma s_L + \phi + s_X \frac{1}{\eta}\left(1 + \varepsilon\lambda(1-\lambda)\frac{1}{\eta}\frac{s_X}{s_L}\right)}.$$

Step 3. In the last step, we characterize conditions that imply that for small shocks, at the planner's allocation $l(X)$, the inefficiency condition (24) holds. We focus on the case in which both terms in (25) are positive. The other case can be proved along similar lines.

Suppose $\sigma(1-(1-\lambda)s_X) > \eta s_L$ and recall that it implies $\mathcal{Q}_Y > 0$. Therefore, condition (24) is equivalent to

$$v'(L)\frac{\Delta'(P_N)Y}{F_L} \cdot \frac{\mathcal{P}_Q}{\mathcal{X}_Q} + \left(v'(L)\frac{F_X}{F_L} - Q\right) > 0,$$

Using the conditions

$$\frac{W}{QP_T} = \frac{F_L}{F_X}, \quad u'(Y) = \frac{P_N}{P_T},$$

gives

$$Q = \frac{F_X}{F_L}\frac{W}{P_T} = \frac{F_X}{F_L}\frac{u'(Y)}{P_N}W.$$

From (49) and (50) we obtain

$$\frac{\mathcal{P}_Q}{\mathcal{X}_Q} = \frac{(1-\lambda)s_X}{\eta s_L}\frac{P_N}{X}.$$

We can then define the function

$$K(l,X) \equiv v'(lX)\frac{\Delta'(\psi(l))F(l,1)}{F_L(l,1)}\frac{(1-\lambda)s_X(l)}{\eta(l)s_L(l)}\psi(l) + \frac{F_X(l,1)}{F_L(l,1)}\left(v'(lX) - \frac{W}{\psi(l)}u'(g(l)X)\right),$$

which gives the value of the externality term as a function of l and X.

Exploiting the fact that at the steady state

$$\Delta'(1) = 0,$$

and

$$v'(L_{SS}) - Wu'(Y_{SS}) = 0,$$

after some algebra, we obtain the following expressions for the derivatives of K at the steady state

$$K_l(l_{SS}, X_{SS}) = \frac{s_X}{s_L}v'(L_{SS})\left(\sigma s_L + \phi + (1-\lambda)s_X\frac{1}{\eta} + \varepsilon\lambda(1-\lambda)\frac{s_X}{s_L}\frac{1}{\eta^2}\right),$$
$$K_X(l_{SS}, X_{SS}) = \frac{l_{SS}}{X_{SS}}\frac{s_X}{s_L}v'(L_{SS})(\sigma + \phi).$$

We can then differentiate $K(l(X), X)$ with respect to X and compute its value at the steady state. To have $K(l(X), X) > 0$ for a small negative shock $\bar{x} < 0$ the following inequality is necessary and sufficient

$$K_l(l_{SS}, X_{SS}) \cdot l'(X_{SS}) + K_X(l_{SS}, X_{SS}) < 0.$$

Substituting the expressions derived above for l', K_l, K_X gives

$$\frac{\sigma + \phi}{\sigma s_L + \phi + s_X \frac{1}{\eta} \left(1 + \varepsilon\lambda \left(1 - \lambda\right) \frac{1}{\eta} \frac{s_X}{s_L}\right)} > \frac{\sigma + \phi}{\sigma s_L + \phi + \left(1 - \lambda\right) s_X \frac{1}{\eta} + \varepsilon\lambda \left(1 - \lambda\right) \frac{s_X}{s_L} \frac{1}{\eta^2}},$$

which, after some algebra, gives

$$\eta < \varepsilon \left(1 - \lambda\right).$$