Wage-Price Spirals*

Guido Lorenzoni and Iván Werning

November 2023

We interpret recent inflation experience through the lens of a new Keynesian model with price and wage rigidities and non-labor inputs in inelastic supply. The model provides a natural interpretation of some features of the recent episode: an initial surge of non-core inflation, followed by a lagged response of core inflation, and a further lagged, persistent response of wage inflation. The model also provides a natural way of discussing the role and the strength of wage-price spiral dynamics in price setting models.

The model interprets recent developments as symptoms of underlying supply constraints, which can be triggered by both demand and supply shocks. The immediate manifestation of these constraints is in the relative price of scarce, inelastic non-labor inputs (including energy). The secondary effects arise because they produce a gap between lowered real wage aspirations of firms—who try to make up for higher non-labor costs—and increased real wage aspirations of workers—caused by increased labor demand. The gap produces a wage-price spiral, which continues as long as the initial relative scarcity of non-labor inputs persists, even though input prices are falling.

In this view, the fact that nominal wage growth is currently exceeding price inflation can be given a benign interpretation, as a sign of real wages going back to trend, and not necessarily as a concern of an ongoing spiral.

The recent inflation surge in the US and in the rest of the world has reignited debates about its origins and propagation mechanisms. In particular, it has brought to the forefront the separate roles and interaction of prices, wages, and profits, and indeed it has done so at two key junctures.

Early on, at the first juncture, many worried that inflation would emanate from a tight labor market, stimulated by expansionary fiscal and monetary policies, causing wage inflation that would then produce price inflation.[1] This is not how inflation played out, though. Instead, price inflation and profit margins soared, while wage growth picked up later and more gradually, implying an initial fall in real wages, as shown in Figure 1.[2]

*We thank for comments and suggestions Olivier Blanchard, Jason Cummins, Jordi Gali, Bob Rowthorn, Aisegul Sahin, and Jon Steinsson. Pedro Bruera and Valeria Vazquez provided excellent research assistance.

[1]Economists who sent prescient, early warnings on inflation risk, like Blanchard (2021), focused on this transmission mechanism.

[2]In the figure, along with CPI inflation, we show two measures of wage inflation, both of which avoid including compositional effects, the BLS Employment Cost Index (all civilian workers, 12-month change) and the Atlanta Fed's Wage Growth Tracker (overall, 12-month change).

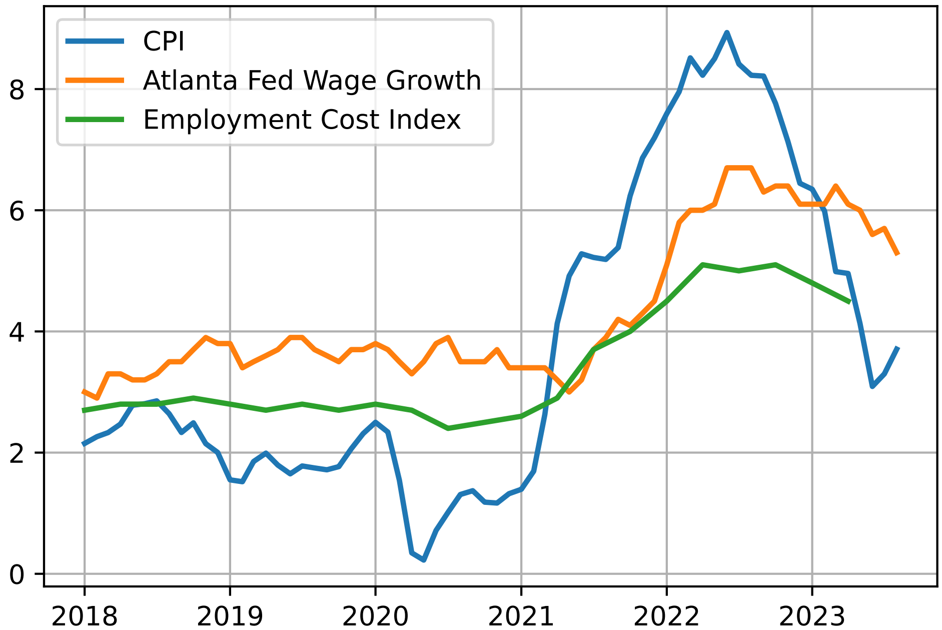

Figure 1: Post-pandemic price and wage inflation in the U.S.

More recently, as price inflation started falling, wage growth rose, surpassing inflation and leading to a rise in real wages. At this second juncture, the concern is that higher wage growth would prevent inflation from going back to target, or even set off an out-of-control wage-price spiral.

This paper aspires to simultaneously improve our understanding of these recent events, while sharpening underlying economic concepts and intuitions surrounding inflation. To this end, we lay out a simple macroeconomic model. We show that this simple model is capable of capturing some key features of the recent episode. Our conceptual analysis dissects the role of prices and wages, isolating their interaction to provide a working definition of wage-price spiral and to understand the dynamics of the real wage.

Our model is relatively close to standard models, but with two essential features not always present in the most basic new Keynesian setups. One important feature of our analysis is the inclusion of a scarce non-labor input with low substitutability in production (lower than Cobb-Douglas). We do not have in mind general forms of capital but rather, inputs like energy, other primary commodities, or intermediate inputs that may be subject to shortages or in relatively fixed supply in the short run, e.g. lumber or microchips. These non-labor inputs provide both a potential supply shock or a supply constraint for demand shocks. This feature of our modeling is motivated by the 2020-23 Covid crises and post-Covid recovery.

The other important feature of our model is that we include both nominal price and wage rigidities, as in many medium-scale models, but unlike the simplest new Keynesian models with only one form of nominal rigidity.

In a model with these features, supply constraints play a crucial role in inflation dynamics, and, when these supply constraints are active, both demand and supply disturbances can set in motion price and wage dynamics that resemble the ones observed.

Namely, the model can produce a three-phase pattern of adjustment in nominal prices. First, there is a bout of very high inflation in the price of the inelastic non-labor inputs, followed by a prolonged gradual fall in the price of these inputs. Second, there is a more persistent period of high general good price inflation. Third, there is a smaller, but even more persistent increase in wage inflation.

The pattern described follows from our assumptions on the role of the inelastic input, which more directly affects price setting firms, and on the relative degree of price stickiness, with the input price being perfectly flexible, and with good prices being more flexible than wages. This pattern implies that at some point wage inflation crosses price inflation, so a period in which real wages fall is followed by a period in which they recover.

Data is always interpreted with a theoretical lens. At one end of the spectrum, commentators and Fed governors' speeches often employ standard macroeconomic concepts, such as a Phillips curve, in their simplest incarnations, to fix ideas or make back of the envelope calculations. On the other end of the spectrum, several papers have contributed by calibrating sophisticated multi-sectoral models. Our paper lives in the gap between these extremes—our model is simpler than medium scale calibrated models allowing us to develop several important concepts, yet goes beyond textbook tools used in day-to-day policy debates.

Turning to the more conceptual points of our paper, one may ask, what do we mean by a wage-price spiral? While there may not be universal agreement, in this paper we use the expression to describe a feedback mechanism where wages and prices compete adjusting upwards: wage earners try to keep up with rising prices; price setters try to keep up with rising wages. This mechanism amplifies and perpetuate the effects of certain inflationary shocks.

Our perspective is that this feedback mechanism is present in virtually all models—including standard new Keynesian varieties. The purpose of this paper is to elucidate and explore this mechanism in detail and focus on the shape of price and wage responses to both supply and demand shocks.

At heart, the economic logic of the wage-price spiral mechanism is that workers and firms disagree on the relative price of goods and labor, that is, on the real wage W/P. When firms adjust nominal prices they do so with some goal for W/P. But workers may adjust nominal wages trying to reach a different, higher ratio for W/P. If they do, the outcome of this disagreement is nominal escalation, with inflation in both prices and wages.

Our interpretation of the concept of a wage-price spiral, highlighting disagreement or conflict as a proximate cause of inflation, is an idea that we explore in greater generality in Lorenzoni and Werning (2022). The present paper studies how this conflict plays out in particular variants of the new Keynesian model and places attention on the path of real wages in response to demand and supply shocks.

Beyond providing an interpretation of recent inflation dynamics, we also use our model to derive a number of general positive and normative results.

First, we derive a general condition for the direction of adjustment of the real wage in response to demand shocks. We show that whether the real wage increases or decreases following a demand shock depends on how strong are the forces set in motion on the price-setting side of the model and on the wage-setting side.

A demand shock acts on the price side by producing an endogenous increase in the price of non-labor inputs. If there is low degree of substitutability between labor and non-labor inputs we get both a large price response of non-labor inputs and a large reduction in the marginal product of labor when non-labor inputs are relatively scarce. The first force will show up in non-core inflation measures, the second will contribute to a the distributional tension between workers and firms that materializes in a wage-price spiral.

A demand shock also acts on the wage side directly. Our model does not feature unemployment and search directly, but the labor supply side of our model captures the basic idea that a overheated labor market will directly affect nominal wage demands, by increasing the rate at which workers are willing to exchange labor for consumption goods. Therefore, this piece of the model captures the basic logic of a wage Phillips curve. Through this channel, excess demand will also produce higher real-wage aspirations for workers and contribute to the wage-price spiral.

However, our general point is that excess demand operates, and contributes to a wage-price spiral, on *both* sides. However, for the movement in the real wage, what matters is the relative strength on the two sides. In our low-elasticity-of-substitution calibration, the effect is stronger on the price side and thus produce, overall, lower real wages.[3]

[3]Incidentally, our analytical result can be taken as a contribution to the classic debate on the cyclicality

An additional observation that comes from our analysis is that both the demand and the supply shocks analyzed create a situation of excess demand. In the demand shock case, natural output is unchanged, but demand temporally expands. In the supply shock case, the "natural" level of output is lower, but demand is unchanged. This excess demand leads to a tension between the level of the real wage that firms and workers aspire to, resulting in a wage-price spiral that produces inflation in both wages and prices. Excess demand is not a sufficient statistic, however. In the supply shock case, real wages always fall, whereas in the demand shock case the real wage may fall depending on parameters. Only under some conditions, the effects on wages and prices is similar for both shocks.

Excess demand is zero when there is a zero output gap. A result that applies in our model is that with a zero output gap there can never be both price and wage inflation, that is, price and wage inflation always have opposite sign. Furthermore, our definition of conflict inflation (from Lorenzoni and Werning, 2022), which we use to capture the wage-price spiral force, is closely related to the size of the output gap in the new Keynesian model here. This connects us immediately to the notion of "divine-coincidence inflation" introduced by Rubbo (2020), which in the model here coincides with conflict inflation.

The result just stated can be rephrased as saying that if the central bank successfully pursues a zero output gap, the central bank can always prevent a wage-price spiral (i.e., achieve zero conflict inflation). But it does not imply that a zero output gap policy is the optimal policy. In Section 4, we study optimal policy and ask two questions. First, could it be part of optimal policy to "run the economy hot", that is, allow for a positive output gap despite high inflation? Second, could it be part of optimal policy to go further and allow for inflation in both prices and wages?

Our answer to the first question is affirmative: if the economy needs a lower real wage, it may be more efficient to reach the adjustment with the help of higher price inflation and moderate wage deflation, rather than though lower price inflation and deeper wage deflation. A positive output gap helps shift the adjustment in the direction of price inflation, so is socially beneficial in this manner.

The answer to the second question is also affirmative. We construct examples in which, at some point, along the adjustment path, the output gap is positive and price and wage inflation are both positive. The economic intuition is that this aspect of policy is a form of "forward guidance": by promising to heat up the economy in the future, we

of the real wage that has spurred a large literature (including Christiano and Eichenbaum 1992; Rotemberg and Woodford 1992). However, our aim here is not to discuss the general cyclical property of real wages, but rather to discuss how potentially sizable real wage movements can be set in motion in special circumstances, like the recent post-pandemic recovery.

speed up the adjustment of the real wage today. Underlying this result is the assumption of forward-looking price- and wage-setting behavior and the commitment of policy. In contrast, when policy has full discretion the equilibrium outcome never features both price and wage inflation.

There is a large and growing literature analyzing the post-pandemic surge in inflation in the U.S. and globally. Our paper is part of a group of papers that emphasizes the role of supply disruptions and supply constraints as playing a crucial role in the recent inflation surge, a group that includes Ball et al. (2022), Amiti et al. (2023), Bernanke and Blanchard (2023), Comin et al. (2023), Gagliardone and Gertler (2023), Kabaca and Tuzcuoglu (2023). We do it here by pointing out the explanatory power of this interpretation for the joint dynamics of prices and wages.

The way in which supply constraints play out here is closely related to the approach in Comin et al. (2023), who develop a quantitative model with an explicit treatment on non-linearities in the supply of non-labor inputs and take an explicit open-economy approach. We believe the virtue of this way of interpreting the facts, is that it shows that a state of global excess demand can causes endogenously sharp input price adjustments, which cannot be taken merely as exogenous price shocks.

Our model emphasizes the role of the real wage as a state variable. This plays an important role in our interpretation of recent events. In particular, we see the recent increase in the real wage as fundamentally driven by a desire of wage setters to make up for the accumulated losses in purchasing power during the early stage of the episode. In other words, we interpret recently high wage inflation as driven by some form of catch-up. The empirical analysis in Bernanke and Blanchard (2023) provides an empirical challenge to this view, as they attempt to measure this catch-up mechanism in the data and fail to find it significant. However, it is not easy to identify structurally this channel of catch-up, and, in general, findings of wage inflation responding to past price inflation, can be taken as supportive of a lag effect, leading to a lag recovery of real wages.[4]

In terms of the broader idea of wage-price spiral, our paper is connected to a vast literature and we'll only make a few close references here. Blanchard (1986) is the seminal paper connecting that idea to new Keynesian models of staggered price setting. The model has nominal prices and wages that are fixed for two periods, with prices reset in even periods and wages in odd periods. The main result in the paper is that the alternating wage and price setting leads to a slow adjustment of the price level in response to a permanent money supply shock and that the adjustment features dampening oscillations

[4]See for example the regressions in Barlevy and Hu (2023)and the literature cited there.

in the real wage. Our paper instead builds on the canonical new Keynesian setting with sticky-price and sticky-wages of the Calvo variety as developed by Erceg et al. (2000). Relative to Blanchard, price and wage setting occur in a staggered fashion without the predictable alternation between wages and prices, so our model is not prone to the same type of oscillations. We also do not focus on a permanent money shock or study monetary policy in terms of money supply. Instead, we focus on supply and demand shocks under different policy responses. Finally, we investigate optimal monetary policy.

Our analysis of wage-price spirals in Section (2) builds on the idea of inflation as the result of distributional conflict, something we explore in more detail in Lorenzoni and Werning (2022). A seminal contribution on this conflict perspective of inflation is Rowthorn (1977). That paper provides a model where, each period, wages are first set by workers and then prices are set by firms. Inflation is shown to be increasing in the conflict or "aspirational gap". Because of the assumed sequential timing of price and wage setting, conflict and inflation must not be fully anticipated by workers. Indeed, no rational expectations equilibrium exists with conflict. In contrast our model features staggered wages and prices that ensure that there is an equilibrium with finite conflict and inflation, even under rational expectations.

Our modeling of non-labor inputs and their connection to price and wage determination connects our analysis to the large literature on models of energy shocks. For example, Blanchard and Gali (2007a),[5] An important modeling difference is that we focus on nominal wage rigidities, while they study a form of real-wage rigidity.

On the normative side, our paper is connected to the welfare analysis of alternative policy rules in models where both prices and wages are rigid, going back to the original paper of Erceg et al. (2000) and to the real-rigidity model of Blanchard and Gali (2007b). The starting observation in the literature is that the presence of both price and wage rigidities breaks "divine coincidence" and introduces potentially interesting trade-offs in the response of monetary policy to supply shocks. We offer a complete characterization of optimal policy and explore conditions for the optimum to have a positive output gap in combination with high inflation, as well as cases where it is optimal to have both wage and price inflation.

[5] In turn, this connects us to the enormous literature on the effects of oil shocks, going back to Bruno and Sachs (1985).

1 Model

We build our arguments in a standard new Keynesian model with nominal price and wage rigidities. To capture supply shocks, an important ingredient we include is a scarce non-labor input X, which is used alongside labor for production. We assume this input has a flexible price, and we allow the production function to have elasticity of substitution different from one.[6] An important example is energy inputs, but we interpret X more broadly to also capture shortages, bottlenecks and capacity constraints in the supply of intermediates like microchips or lumber, which have been in the spotlight during the post-pandemic recovery.

We focus on a closed economy in which the supply of X is given while the price of X adjusts endogenously in equilibrium. The analysis can be easily expanded to the case of an open economy in which the good X is imported, and, in particular, to the limit case of a small open economy that takes the world price of X as given. In that case, a supply shock would take the form of a shock to the world price instead of a shock to the endowment.

1.1 Setup

Time is continuous and infinite. The representative household has preferences

$$\int_0^\infty e^{-\rho t}\left(\frac{1}{1-\sigma}C_t^{1-\sigma}-\frac{\Phi_t}{1+\eta}{N_t}^{1+\eta}\right)dt,$$

where C_t is an aggregate of a continuum of varieties of goods $C_t=\left(\int_0^1 C_{jt}^{1-1/\varepsilon_C}dj\right)^{\frac{1}{1-1/\varepsilon_C}}$, N_t is labor supply, and Φ_t is a labor supply shock. Each good variety j is supplied by a monopolistic firm with production function

$$Y_{jt}=F\left(L_{jt},X_{jt}\right)\equiv\left(a_L L_{jt}^{\frac{\epsilon-1}{\epsilon}}+a_X X_{jt}^{\frac{\epsilon-1}{\epsilon}}\right)^{\frac{\epsilon}{\epsilon-1}},$$

where L_{jt} is the labor input and X_{jt} is the non-labor input. The labor input L_{jt} of each firm j is an aggregate of a continuum of labor varieties $L_{jt}=\left(\int_0^1 L_{jkt}^{1-1/\varepsilon_L}dk\right)^{\frac{1}{1-1/\varepsilon_L}}$. Each labor variety k is supplied by a monopolistic union that employs labor from households and turns it, one for one, into specialized labor services of type k. Integrating over firms, total employment of labor variety k is $N_{kt}=\int_0^1 L_{jkt}dj$. Integrating over unions, total labor supply is $N_t=\int_0^1 N_{kt}dk$. The representative household owns an exogenous endowment X_t of

[6] This is formally equivalent to having labor and capital, with capital rented at a flexible price, although the interpretation is different. Erceg et al. (2000) have labor and capital. Closer to the interpretation here, Blanchard and Gali (2007a) have an energy input.

the non-labor input X and sells it to the monopolistic goods producers on a competitive market, at the price P_{Xt}.

Monopolistic firms set the nominal price at which they are willing to sell their variety and then supply the amount chosen by consumers. Similarly, monopolistic unions set the nominal wage and supply the amount chosen by firms. Firms and unions are only allowed to reset their price and their wage rate occasionally. Namely, at each point in time, firms are selected randomly to reset their price with Poisson arrival λ_p, and unions are selected with arrival λ_w.

When the exogenous variables X_t and Φ_t are constant, the model has a steady state in which quantities are constant, nominal prices are constant (zero inflation), all good varieties have the same price, and all labor varieties have the same wage. We will consider an economy in steady state and analyze its response to one time, unexpected shocks, either due to changes (transitory or permanent) to X_t or Φ_t or to changes in monetary policy leading to transitory deviations of C_t and N_t from the path consistent with zero inflation.

1.2 Price and wage setting

Let P_t^* and W_t^* denote the price and wage set by the firms and unions that can reset at time t, while P_t and W_t denote the price indexes for the good and labor aggregates.

The nominal marginal cost of producing good j is

$$\frac{W_t}{F_L\left(L_{jt}, X_{jt}\right)} = \frac{W_t}{a_L Y_{jt}^{\frac{1}{\epsilon}} L_{jt}^{-\frac{1}{\epsilon}}}.$$

Using lowercase variables to denote log-linear deviations from steady state and taking a first-order approximation, nominal marginal costs can then be expressed as

$$w_t - mpl_{jt}, \tag{1}$$

where

$$mpl_{jt} = \frac{1}{\epsilon}\left(y_{jt} - l_{jt}\right)$$

is the marginal product of labor. The production function of firm j in log-linear approximation is

$$y_{jt} = s_L l_{jt} + s_X x_{jt}, \tag{2}$$

where s_L and s_X are the steady state shares of the labor and non-labor inputs, with $s_L + s_X = 1$. All firms being price takers in the input market, they all employ inputs in the same ratio L_{jt}/X_{jt}, so in log-linear approximation

$$l_{jt} - x_{jt} = n_t - x_t$$

where n_t and x_t are the aggregate supplies of the two inputs. Combining these results, the marginal product of labor is

$$mpl_t = \frac{s_X}{\epsilon}\left(x_t - n_t\right). \tag{3}$$

Following standard steps, optimal price setting requires that firms set their price at time t equal to an average of future nominal marginal costs, conditional on not resetting. This gives the following optimality condition for P_t^* in log-linear approximation

$$p_t^* = \left(\rho + \lambda_p\right)\int_t^\infty e^{-\left(\rho+\lambda_p\right)(\tau-t)}\left(w_\tau - mpl_\tau\right)d\tau. \tag{4}$$

Following similar steps, we can derive the wage setting equation

$$w_t^* = \left(\rho + \lambda_w\right)\int_t^\infty e^{-(\rho+\lambda_w)(\tau-t)}\left(p_\tau + mrs_\tau\right)d\tau \tag{5}$$

where

$$mrs_t = \phi_t + \sigma y_t + \eta n_t \tag{6}$$

is the marginal rate of substitution between consumption and leisure of the representative consumer.

The presence of the w_τ's on the right-hand side of equation (4) and of the p_τ's on the right-hand side of equation (5) capture the logic of a wage-price spiral in our model. Firms aim to get prices to be a constant markup over nominal marginal costs, and since marginal costs depend on nominal wages, they set nominal prices to catch up with current and anticipated future nominal wages. Symmetrically, wage setters aim to achieve a real wage that reflects their willingness to substitute leisure with consumption goods, so, they set nominal wages to catch up with current and anticipated future nominal good prices.

The optimality condition for the input-ratio of firms can be written as follows

$$p_{Xt} = w_t - \frac{1}{\epsilon}\left(x_t - n_t\right). \tag{7}$$

This condition will be used to derive the equilibrium input price p_{Xt}.

1.3 Inflation equations

To go from equations (4) and (5) to wage and price inflation, combine them with the differential equations for p_t and w_t:

$$\dot{p}_t = \lambda_p \left(p_t^* - p_t\right), \tag{8}$$

$$\dot{w}_t = \lambda_w \left(w_t^* - w_t\right). \tag{9}$$

As shown in the Appendix, we then obtain the following expressions

$$\rho\pi_t = \Lambda_p \left(\omega_t - mpl_t\right) + \dot{\pi}_t, \tag{10}$$

$$\rho\pi_t^w = \Lambda_w \left(mrs_t - \omega_t\right) + \dot{\pi}_t^w, \tag{11}$$

where we use the notation $\pi_t \equiv \dot{p}_t$ and $\pi_t^w \equiv \dot{w}_t$ for price and wage inflation and $\omega_t \equiv w_t - p_t$ for the real wage, and the coefficients Λ_p and Λ_w are

$$\Lambda_p = \lambda_p \left(\rho + \lambda_p\right), \quad \Lambda_w = \lambda_w \left(\rho + \lambda_w\right).$$

Real wage dynamics are given by

$$\dot{\omega}_t = \pi_t^w - \pi_t. \tag{12}$$

Equations (10) and (11) can be interpreted in terms of a conflict between the real wage aspirations of workers and firms, an interpretation we develop in Lorenzoni and Werning (2022). In the context of the new Keynesian model, the workers' aspiration is given by the marginal rate of substitution mrs_t at which the representative worker is willing to exchange labor for goods, the firms' aspiration is the marginal product of labor mpl_t.[7] As in Lorenzoni and Werning (2022), a discrepancy between the aspirations mpl_t and mrs_t is the proximate cause of inflation.

Equations (10) and (11) can also be expressed as traditional Phillips curves because the expressions $\omega_t - mpl_t$ and $mrs_t - \omega_t$ can be written in terms of gaps between equilibrium objects and their "natural" level. Focusing on the wage equation, we can write[8]

$$\begin{aligned} mrs_t - \omega_t = mrs_t - mrs_t^* - \left(\omega_t - \omega_t^*\right) = \\ = \left(\sigma s_L + \eta\right)\left(n_t - n_t^*\right) - \left(\omega_t - \omega_t^*\right) \end{aligned}$$

[7]The variable ϕ_t in the notation of Lorenzoni and Werning (2022) corresponds to mpl_t here and the variable γ_t is corresponds to mrs_t.

[8]This derivation applies because at the natural allocation the real wage is equalized to the workers' *mrs*. The detailed derivations are in Appendix ...

where ω_t^* is the flexible-price wage rate and n_t^* is the natural level of employment. Substituting this expression in (11) we obtain a wage Phillips curve that connects wage inflation to the employment gap $n_t - n_t^*$. An analogous derivation can be done for the price equation. The crucial observation here is that in both Phillips curves there is an additional term, given by the deviation between the real wage and its flexible-price level ω_t^*. Notice that ω_t is a state variable of our system because both w_t and p_t only move gradually, due to stickiness, so, at a given moment in time ω_t is given by the history of past shocks.

Given an initial condition ω_0 and given paths for mpl_t and mrs_t for $t \geq 0$, the three equations (10)-(12) give unique paths for price and wage inflation.

Our approach in the rest of the paper is to split the analysis in two steps:

1. From the paths for fundamental shocks and aggregate real activity derive the paths of mpl_t and mrs_t;

2. From the paths of mpl_t and mrs_t derive inflation.

In general, in a full-blown general equilibrium model the paths of mpl_t and mrs_t are endogenous and this way of splitting the analysis is somewhat artificial. However, a central point of this paper is to show that this decomposition helps understand the mechanisms underlying inflation in equilibrium.

The next section focuses on step 2. We then go back to step 1 in the following section.

2 From Aspirations to Inflation, With and Without A Spiral

In general, shocks to the economy translate into endogenous changes in the variables mpl and mrs, which, as argued above, reflect the real wage aspirations of firms and workers. In this section, we take the paths of mpl and mrs as given and focus on deriving inflation as a function of them. This part of the analysis isolates how staggered price setting produces inflation for given aspirations, and allows us to identify the wage-price spiral mechanism. The next section shows how shocks and policies determine mpl and mrs and thus completes the analysis. A reader mostly interested in our interpretation of the post-pandemic high inflation episode, can skip this section without loss.

Through the paper, we mostly focus on exponentially decaying paths of mpl and mrs that take the following form.[9] Before $t = 0$ the economy is in steady state: all variables expressed in log deviations from the steady state are equal to zero. At $t = 0$ there is an

[9]In the appendix, we provide a general analytical characterization of the relation between the paths $\{mpl_t, mrs_t\}$ and price and wage inflation.

unexpected shock and mpl_0 and mrs_0 jump discretely to values different from zero (at least for one of them). From then on, they both converge back to the original steady state at constant speed δ, so

$$mpl_t = mpl_0 e^{-\delta t},$$
$$mrs_t = mrs_0 e^{-\delta t}.$$

The demand and supply shocks analyzed in the next section produce paths with this shape, so the analysis here will immediately apply.

To derive price and wage inflation from equations (10) and (11) requires solving first for the endogenous path of the real wage ω_t. In other words, as mentioned earlier, the real wage is a necessary state variable in our inflation equations. The solution for the real wage in terms of mpl and mrs comes from solving a second order ODE and the details are provided in the Appendix. Once we have ω_t, equations (10) and (11) can be solved forward to get

$$\pi_t = \Lambda_p \int_0^\infty e^{-\rho s} \left(\omega_s - mpl_s\right) ds, \tag{13}$$
$$\pi_t^w = \Lambda_w \int_0^\infty e^{-\rho s} \left(mrs_s - \omega_s\right) ds. \tag{14}$$

Price and wage inflation are driven by current and anticipated gaps between the real wage and firms' and workers' aspirations. These two equations are used to provide intuition in this section.

2.1 Two Examples

Consider the two numerical examples plotted in Figure 2.[10]

In the first, mpl and mrs fall by the same amount at date 0, that is, $mrs_0 = mpl_0 < 0$. On impact, the reduction in mpl increases firms' marginal costs, leading firms to increase nominal prices, while the reduction in mrs lowers workers' aspirations and workers reduce nominal wages. In the top left panel of Figure 2, we see that this leads to $\pi_0 > 0 > \pi_0^w$. The real wage starts falling, as shown in the lower left panel. As time goes by, the force of the initial shock goes away while, at the same time, the real wage is lower. Both forces reduce $\omega - mpl$ in the price inflation equation and increase $mrs - \omega$ in the wage inflation

[10] The parameters for the examples are

$$\lambda_p = 2, \lambda_w = 1, \rho = 0.04, \delta = 0.5.$$

equation: the gap between aspirations and the real wage fall for both. After some date, when mpl and mrs are small enough and the real wage has fallen enough, both inflation rates π_t and π_t^w flip sign and we have $\pi_t < 0 < \pi_t^w$. From then on, the real wage starts growing and converges back to its initial level.

In this example, even though wage setters and price setters respond to each other's prices (current and anticipated), this does not produce generalized inflation or deflation, because the two parties are aiming to achieve the same relative price adjustment, so their actions tend to dampen each other. The fact that firms increase prices tends to remove the deflationary impulse on the workers' side. The fact that workers lower their wages tends to remove the inflationary impulse on the firms' side. In this case a wage-price spiral is not present.

In the second example, only the aspirations of firms change, with $mpl_0 < 0$, but mrs_0 is unchanged at zero. In this case there is a positive gap $mrs_0 - mpl_0$. This case is illustrated in the two panels on the right of Figure 2.

On impact, the reduction in mpl increases firms' marginal costs, as in the first example. Now there is no direct effect of mrs on the workers' side, workers anticipate a future reduction in real wages, and react at time 0 by raising their nominal wage demand.[11] Therefore, we get both wage and price inflation, $\pi_0 > \pi_0^w > 0$. In general, in every case in which there is a unilateral change in mpl, with no change in mrs, it is possible to show that price inflation is larger than wage inflation at $t = 0$, given that the price equation is affected directly by the change in mpl, while the wage equation is only affected indirectly through the future equilibrium adjustment in ω.[12]

Notice the back and forth between price and wage inflation, that amplifies the initial shock. The shock originates in the inflation equation but produces an undesirable relative price adjustment for workers, creating a positive gap between workers' aspirations and the real wage path, inducing wage setters to respond. This causes price inflation to spill over into wage inflation. The wage setters' response in turn dampens the adjustment in the real wage, relative to what happens in our first example: comparing the two lower panels in Figure 2, the real wage ω_t falls less in the panel on the right. Therefore, the presence of wage inflation, slowing the fall in real wages, reinforces the price inflation response, as firms, anticipating a weaker reduction in real wages, keep price inflation higher.[13]

[11] In equation (14), $mrs_s = 0$ and $\omega_s < 0$ for all s. Why the real wage falls in this example is explained below.

[12] See Proposition 5 in the Appendix.

[13] If nominal wages were perfectly sticky this amplification would not be present and price inflation would be lower throughout. We go back to the relation between stickiness and amplification at the end of

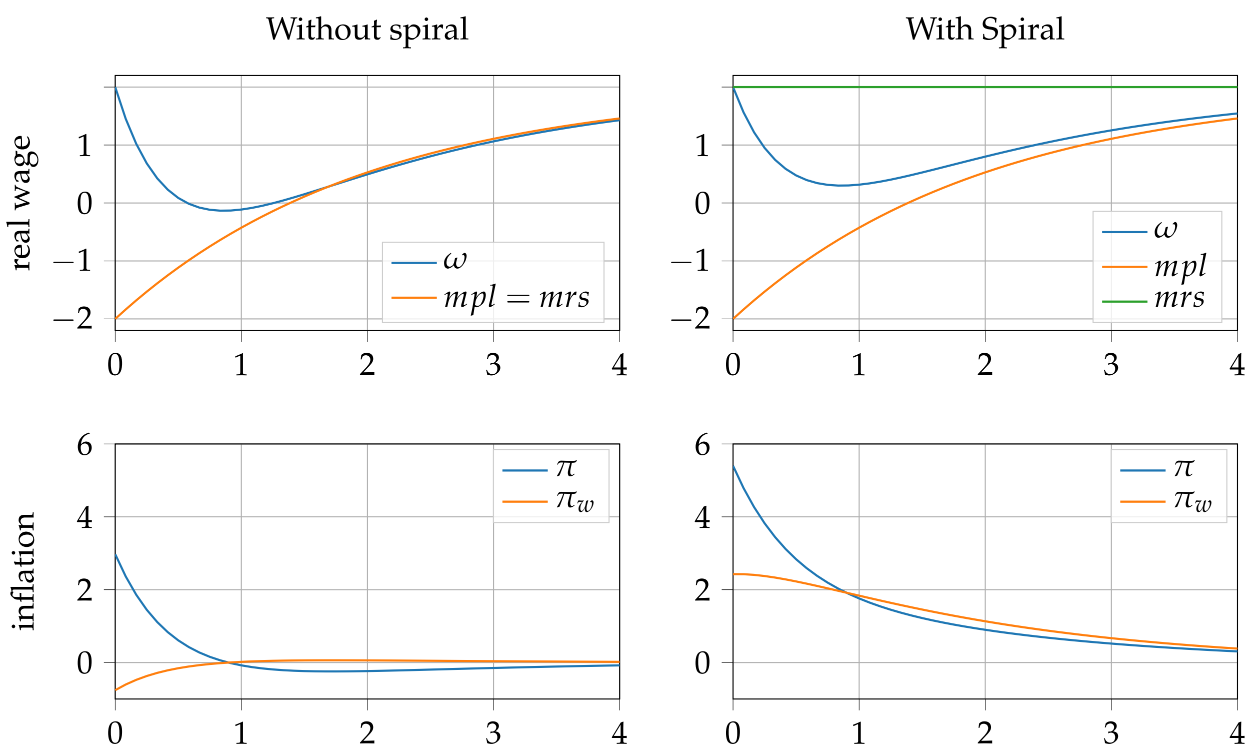

Figure 2: Aspirations and inflation, with and without a spiral

The expression "wage-price spiral" is used to describe these mutually reinforcing dynamics between price and wage inflation. In the first example there is no wage-price spiral, in the second there is.

2.2 Spiral Dynamics and Conflict Inflation

In the two examples above, we just argued that the first shows no spiral while the second does. But how can we distinguish more formally the spiral force in the second example from the relative-price-adjustment mechanism that drives nominal prices and wages in the first?

The crucial difference is that in the second example the attempt of each side to move the relative price in its direction leads to a protracted period of high inflation in both prices and wages. Let us measure the spiral effect in terms of the cumulated effect on price and wage inflation over the entire episode. Since the real wage always mean reverts to 0, cumulated price and wage inflation are the same and we can define

$$\Pi^{Spiral} \equiv \int_0^\infty \pi_t dt = \int_0^\infty \pi_t^w dt.$$

this section.

In the appendix, we prove that

$$\Pi^{Spiral} = \frac{\Lambda_p \Lambda_w}{\Lambda_p + \Lambda_w} \frac{1}{\delta\left(\rho + \delta\right)} \left(mrs_0 - mpl_0\right).$$

Notice the symmetric role of Λ_p and Λ_w in this expression: for the spiral effect to be present we need prices and wages to respond to each other. If one side has fixed nominal prices, for example $\lambda_w = 0$, then the spiral is completely absent. On the other hand, if we vary λ_p and λ_w and hold fixed the total degree of nominal rigidtiy in the economy $\lambda_w + \lambda_p$, then the maximum power of the spiral arises when $\lambda_p = \lambda_w$, that is, when each side responds to the other with equal speed.

The spiral measure just introduced immediately connects spiral dynamics to the notion of conflict inflation proposed in Lorenzoni and Werning (2022), which is defined as follows

$$\Pi_t^{Conflict} \equiv \frac{\Lambda_p \Lambda_w}{\Lambda_p + \Lambda_w} \int_0^\infty e^{-\rho s} \left(mrs_{t+s} - mpl_{t+s}\right) ds,$$

and, with exponentially decaying shocks, yields

$$\Pi_0^{Conflict} = \frac{\Lambda_p \Lambda_w}{\Lambda_p + \Lambda_w} \frac{1}{\rho + \delta} \left(mrs_0 - mpl_0\right).$$

We then conclude that

$$\Pi^{Spiral} = \frac{1}{\delta} \Pi_0^{Conflict},$$

which means that conflict inflation at date 0 fully captures the underlying forces that lead to a protracted period of joint price and wage inflation.

Notice that from (13) and (14), we get

$$\Pi_0^{Conflict} = \alpha \pi_0 + \left(1 - \alpha\right) \pi_0^w, \tag{15}$$

where α is a coefficient of relative stickiness, defined as

$$\alpha \equiv \frac{1/\Lambda_p}{1/\Lambda_p + 1/\Lambda_w}.$$

We then have a "forecasting" interpretation of the result above. Consider an econometrician who does not observe the underlying shocks mrs_0 and mpl_0 at $t = 0$ but only the current inflation rates π_0 and π_0^w. Conflict inflation is the linear combination of π_0 and π_0^w that provides the best estimate of the cumulated future effect of the underlying shocks on inflation.[14]

[14]This result relies on the simple joint AR1 structure of the shocks to mrs_0 and mpl_0. It is an open important question how to extend the connection between conflict inflation and inflation forecasting to

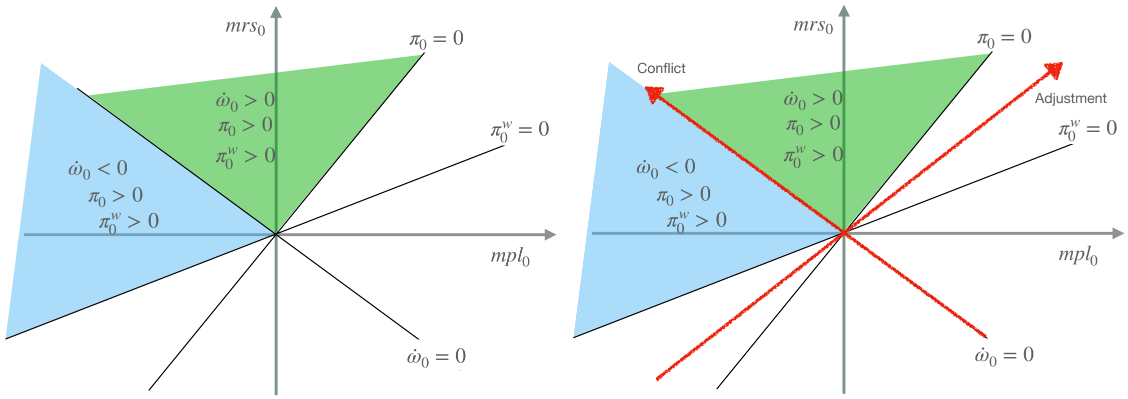

Figure 3: Regions for mpl_0 and mrs_0

From (15) and $\dot\omega_0 = \pi_0^w - \pi_0$, we get the decomposition

$$\pi_0 = \Pi_0^{Conflict} - (1-\alpha)\,\dot\omega_0,$$
$$\pi_0^w = \Pi_0^{Conflict} + \alpha\dot\omega_0.$$

Conflict inflation captures the underlying common component of price and wage inflation due to the gap between the aspirations on the two sides of the market (mpl_0 and mrs_0). The presence of the gap is crucial to set in motion mutually reinforcing responses on the two sides. When there is no gap, there can be no generalized inflation, π_t and π_t^w have opposite sign and the mutual responses tend to dampen the initial shock, consistently with our first example.

Notice that in the new Keynesian model considered here conflict inflation Π_t is proportional to the output gap, as we shall see in the next section. This implies that conflict inflation coincides with the notion of "divine-coincidence" inflation in Rubbo (2020) and with the composite inflation index in the optimal inflation-targeting rule of Giannoni and Woodford (2004).[15]

A graphical representation

A graphical representation can help interpret the decomposition above.

In the left panel of Figure 3 we divide the space (mpl_0, mrs_0) in six regions, depending on the sign of the three variables $\pi_0, \pi_0^w, \dot\omega_t$.

The next proposition shows that the configuration in Figure 3 is general and indepen-

richer structures.

[15]See Chapter 6.4 of Gali (2015) for a textbook discussion.

dent of parameters, given exponentially decaying shocks. The proposition gives conditions in terms of the coefficient ψ, which is a function of the parameters $\Lambda_p, \Lambda_w, \rho$ and δ and is defined in the Appendix.

Proposition 1. *Given exponentially decaying paths for mpl and mrs, at date $t = 0$ price and wage inflation satisfy*

$$\pi_0 > 0 \textit{ iff } (1-\alpha)\,\psi \cdot mrs_0 > (1-\alpha\psi) \cdot mpl_0,$$
$$\pi_0^w > 0 \textit{ iff } (1-(1-\alpha)\,\psi) \cdot mrs_0 > \alpha\psi \cdot mpl_0,$$

and

$$\pi_0^w - \pi_0 = \dot{\omega}_0 > 0 \textit{ iff } \alpha mpl_0 + (1-\alpha)\, mrs_0 > 0.$$

The slope of the boundary of the $\pi_0 > 0$ region is always steeper than that of the $\pi_0^w > 0$ region.

The green and blue regions in Figure 3 are those in which the economy features positive price and wage inflation. Both $mrs_0 > 0$ and $mpl_0 < 0$ are inflationary forces, and produce inflation as long as one of them is present and strong enough.

A positive value for mrs_0 acts directly on wage inflation, a negative mpl_0 acts directly on price inflation. Both also act indirectly through their effects on ω_t. A high mrs_0, by pushing future real wages up tends to increase expected marginal costs and price inflation at $t = 0$. A low mpl_0, by pushing future real wages down, tends to increase wage demands and wage inflation at $t = 0$. The fact that mrs acts directly on wages, while mpl acts directly on prices gives some intuition for why the slope of the $\pi_0 = 0$ line is steeper than that of the $\pi_0^w = 0$ line.

The difference between the green region and the blue region is that in the blue region the real wage falls at $t = 0$ while it increases in the green region. The reason for the difference is the relative strength of the pressure on price setters and wage setters.

The right panel of Figure 3 is identical to the left but adds two axes that represent the conflict and adjustment components of inflation.

The adjustment axis is simply given by the 45 degree line, $mrs_0 = mpl_0$, given that along that line conflict inflation is zero.

The conflict axis is the boundary between the green and blue regions: it is the locus where the power of a wage-price spiral is stronger, because the aspirations of workers and firms are opposite and of equal force, once we adjust for the frequency of price adjustment, that is, where

$$(1-\alpha)\, mrs_0 = -\alpha mpl_0.$$

Along that locus there is zero adjustment inflation: the opposite efforts of workers and firms produce no movement in the real wage and only socially wasteful price dispersion.[16]

To clarify the connection between the figure and the analysis above, it is useful to remember that the figure only shows the impact effect on π_0 and π_0^w. As time goes by and ω_t changes, the same figure applies but with the origin of the conflict and adjustment axes (and of the $\pi_t = 0$ and $\pi_t^w = 0$ loci) shifted along the 45 degree line. So, for example, we can have a shock in the upper-right quadrant that initially produces $\pi_0 < 0$ and $\pi_0^w > 0$, but also gives positive conflict inflation $\Pi_0^{Conflict} > 0$. As time goes by, we will have $\omega_t > 0$ and the origin will shift to the right along the 45 degree line while, at the same time, mrs_t and mpl_t move linearly towards the $(0,0)$ origin. This will at some point produce a combination $\pi_t > 0$ and $\pi_t^w > 0$, consistently with the fact that the shock will eventually produce positive cumulated inflation in both prices and wages.[17]

2.3 Stickiness and Amplification

Consider now a different exercise: fix the size of two initial shocks $mrs_0 > 0$ and $mpl_0 < 0$ and change the economy's parameters λ_w and λ_w to vary the degree by which the shocks get amplified through the wage-price responses.

As we increase the speed at which either prices or wages are reset, the wage-price spiral mechanism gets stronger. This is shown in Figure 4, where we plot level curves for π and π_w. The relatively steeper curves (in absolute value) correspond to π, the flatter ones to π_w. A higher frequency of price adjustment λ_p increases both π and π_w, but has a stronger effect on the former. The reverse holds for λ_w. For ease of illustration, we consider an economy hit by a symmetric shock $mrs_0 = -mpl_0$. This implies that when $\lambda_p = \lambda_w$ Proposition 1 gives $\dot{\omega}_0 = 0$ and $\pi_0 = \pi_0^w$. In the figure, the contour levels corresponding to equal price and wage inflation meet on the 45 degree line.

[16]Projecting any point (mpl_0, mrs_0) on the two axes, the conflict coordinate gives conflict inflation Π_0, while the adjustment coordinate gives $\dot{\omega}_0$. The two coordinates measure adjustment and conflict inflation if we scale the axes as follows: on the adjustment axis the unit vector is

$$\begin{pmatrix} mpl_0 \\ mrs_0 \end{pmatrix} = \frac{r_2 + \delta}{\Lambda_p + \Lambda_w} \begin{pmatrix} 1 \\ 1 \end{pmatrix},$$

(where r_2 is the positive eigenvalue of the real wage ODE, as defined in the Appendix) and on the conflict axis the unit vector is

$$\begin{pmatrix} mpl_0 \\ mrs_0 \end{pmatrix} = \frac{\Lambda_p + \Lambda_w}{\Lambda_p \Lambda_w} (\rho + \delta) \begin{pmatrix} -(1-\alpha) \\ \alpha \end{pmatrix}.$$

[17]Notice also that there is a t—the t at which $\dot{\omega}_t = 0$—where $\pi_t = \pi_t^w = \Pi_t^{Conflict} > 0$.

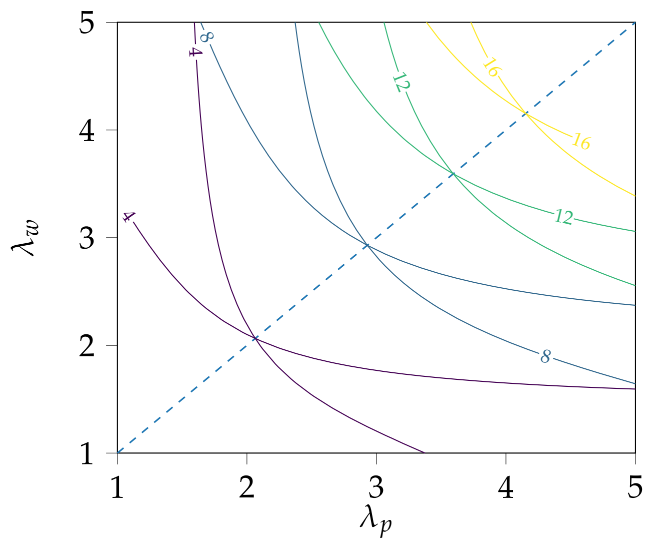

Figure 4: Price and wage inflation contours for different degrees of stickiness

Increasing either price or wage flexibility increases *both* price and wage inflation. This is the total force of the wage-price mechanism. At the same time, what happens to the real wage depends on the relative force on the two sides. Increasing λ_p tends to move us to the region below the 45 degree line, where real wages fall. Increasing λ_w has the opposite effect.

3 Demand and Supply Shocks

We now go back to the full model and trace back price and wage inflation to the general equilibrium effect of two shocks: a demand shock and a supply shock.

We show that if the economy is in an initial state that is "sensitive to supply constraints", in a sense to be made precise, a positive demand shock and a negative supply shock have qualitatively similar implications on inflation. Namely, there will be a dynamic response in three phases, first a fast increase in non-core inflation, captured here by the price of the scarce input X, then a period of sustained general inflation in prices and wages, with price inflation stronger than wage inflation and real wages falling. And finally a period of persistent wage inflation, with price inflation lower than wage inflation and real wages growing back. As argued in the introduction, these dynamics seem to capture well the recent post-pandemic inflationary experience.

3.1 A Demand Shock

Consider an expansionary demand shock, driven by easy monetary policy. In particular, suppose the shock is such that real spending increases to $y_0 > 0$ at date $t = 0$, and, after that, the shock decays exponentially at rate δ, so

$$y_t = y_0 e^{-\delta t}.$$

We have not explicitly modeled monetary policy, which could be done by solving the consumers' intertemporal optimization problem and adding an interest rule to the model. However, it can be shown that the shock above translates immediately into a shock that reduces temporarily the real interest rate below its natural level (here ρ), hence stimulating consumer spending. A demand shock coming from a fiscal impulse or from consumer sentiment would also have similar implications.

3.2 An Inequality for Supply-Constrained Demand Shocks

The responses of the aspirations mpl_t and mrs_t are easily derived from (3) and (6):

$$mpl_t = -\frac{s_X}{\epsilon} e^{-\delta t} \frac{1}{s_L} y_0 < 0, \quad mrs_t = (\sigma s_L + \eta) e^{-\delta t} \frac{1}{s_L} y_0 > 0.$$

The response of the relative price of the X input (expressed in terms of labor) also follows immediately from (7):

$$p_{Xt} - w_t = \frac{1}{\epsilon} e^{-\delta t} n_0 > 0.$$

Giving the sign of these responses, Proposition 1 immediately tell us that both price and wage inflation are positive following this shock. Firms would like to pay lower real wages, given that the marginal product of labor has fallen. Consumers would like to be paid higher real wages, because they are spending more and working more, so income and substitution effect both push for a higher real marginal compensation of labor. These opposing forces produce spiral inflation, i.e., conflict inflation, as discussed in the previous section.

What happens to the real wage is in general ambiguous, but Proposition 1 gives us an easy condition to check, to establish the sign of its response. The next proposition provides this condition.

Proposition 2. *In response to a monetary shock leading to a transitory, exponentially decaying, increase in real output, price and wage inflation are both positive. Price inflation is higher than wage inflation, and consequently real wages fall, at $t = 0$, if and only if the following condition is* *satisfied*

$$\Lambda_p \frac{s_X}{\epsilon} > \Lambda_w \left(\sigma s_L + \eta \right). \tag{16}$$

When an economy satisfies inequality (16), we'll say that it is supply-constrained or sensitive to supply constraints, because, as we shall see, the relative scarcity of the X input, driven by the ratio N_t/X_t plays a central role for price and wage inflation dynamics.

The intuition for inequality (16) is as follows.

Consider first the expression on the left-hand side, $\Lambda_p \frac{s_X}{\epsilon}$. The ratio $\frac{s_X}{\epsilon}$ captures the effect of an increase in employment on the marginal product of labor. To increase output, the economy must increase the labor input, with a fixed supply of the input X. The ratio $\frac{N_t}{X_t}$ goes up, making the X factor relatively scarcer and labor relatively abundant. How much this lowers the marginal product of labor depends on how important is the input X in the production of the final good—the share s_X—and how elastically labor can substitute for X—the elasticity ϵ. If s_X is high and ϵ low, we get a large effect. Finally, the coefficient Λ_p captures how quickly firms can respond to lower marginal productivity, that is, to higher marginal costs, by raising nominal prices.

The expression on the right-hand side $\Lambda_w \left(\sigma s_L + \eta \right)$ comes instead from the workers' side. In particular, the expression $\sigma s_L + \eta$ captures how income and substitution effects change how much workers would like to be compensated on the margin. While Λ_w captures how quickly a higher *mrs* leads to increasing nominal wages.

As we discussed in the previous section, both impulses, to *mpl* on the firms' side and to *mrs* on the workers' side, lead to mutual reactions, that is, to indirect effects: an impulse on firms' marginal costs also leads to increasing nominal wages, and an impulse on workers' marginal rate of substitution also leads to nominal price inflation. However, Proposition 1 shows that the indirect effects are always weaker than the direct effects and that the presence of indirect effects does not change the relative size of the effects on the two sides. Therefore, focusing on the relative strength of the direct effects, we can safely conclude that price inflation will be higher in equilibrium than wage inflation if and only if the direct impulse on prices—the left-hand side of (16)—is stronger than the direct impulse on wages—the right-hand side.

3.3 An Example

Having unpacked analytically the effect of the shock at date $t = 0$, let us turn to a numerical example to look at the full dynamics and get a sense of the magnitudes involved. We focus on an example that satisfies inequality (16).

In Figure 5, we plot the response to a temporary expansionary shock that increases y

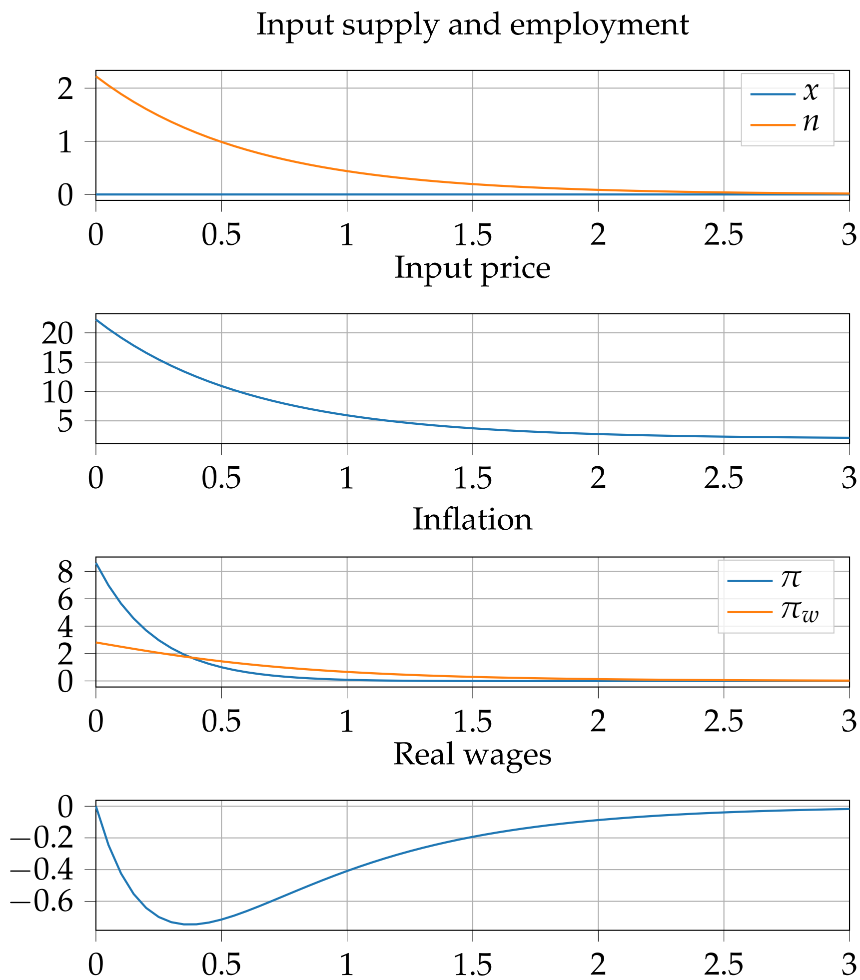

Figure 5: A supply-constrained demand shock

above potential by 2% on impact and converges back to potential at the rate $\delta = 1$. The parameters used are in Table 1.[18]

The first panel shows the path of employment n, which is proportional to output, and the path of x, which, by assumption is constant at 0. The remaining panels show the responses of different prices.

The input price is flexible, so it jumps on impact and then gradually goes back to its initial level, as the shock goes away. This is shown in the second panel of the figure. Notice that this panel shows the level of the input price, not its inflation rate. Inflation for that price is infinite at $t = 0$ and negative afterwards. Due to perfect flexibility P_X jumps by 20% at $t = 0$. This large increase is due to our assumption of a low elasticity of substitution between labor and the input X ($\epsilon = 0.1$), so when the employment is growing

[18] All plots show log deviations from steady state times 100, or, approximately, percentage deviations from steady state.

Preferences	$\sigma = 1$	$\eta = 1/2$	$\rho = 0.04$
Technology	$s_X = 0.1$	$\epsilon = 0.1$,	
Stickiness	$\lambda_p = 4$	$\lambda_w = 1$	

Table 1: Parameters

too fast relative to the supply of X, the price of X reacts strongly.

The effect of the increase in the input price is to increase firms' marginal costs. The impact effect on the nominal marginal cost $w_0 - mpl_0$ is $+2\%$, as the input represents 10% of the cost in steady state, $s_X = 0.1$, and the elasticity if also $\epsilon = 0.1$, so the ratio $s_X/\epsilon = 1$. As we see from the third panel of Figure 5, this increase in marginal costs translates into fast inflation on impact: 10% above its steady state level (so 12% inflation if we assume the central bank is keeping inflation at 2% in steady state).[19] This large response to a relatively small increase in marginal costs is due to our assumption of relatively flexible prices ($\lambda_p = 4$, i.e, prices reset on average every quarter), to the firms' having rational expectations and a long horizon (captured by the discount rate ρ), and, of course, to the wage response, that is, to the presence of a wage-price spiral.

On the wage side the direct, impact effect on the mrs is $(\sigma s_L + \eta) \times 2\% = 2.8\%$, and is close in magnitude to the effect on the marginal cost of goods, both are 2%. However, wages are more sticky ($\lambda_w = 1$), so the effect on wage inflation is weaker. Wage inflation is also plotted in the third panel of Figure 5.

The real wage falls on impact, as shown in the fourth panel. However, as time goes by, the lower level of the real wage pushes workers to ask nominal wage increases larger than price inflation. Wage growth eventually reverse sign and the real wage converges back to trend.

Figure 5 illustrates the three phases of adjustment mentioned in the Introduction. First, very fast inflation in the sector where the supply constraints are binding, here the market for input X. Second, a phase in which price inflation is faster than wage inflation. Third, at some point wage inflation crosses price inflation and we enter the third phase in which real wages recover.

We'll discuss more in depth the connection between this example and current developments at the end of this section. But first, let us look at a supply shock.

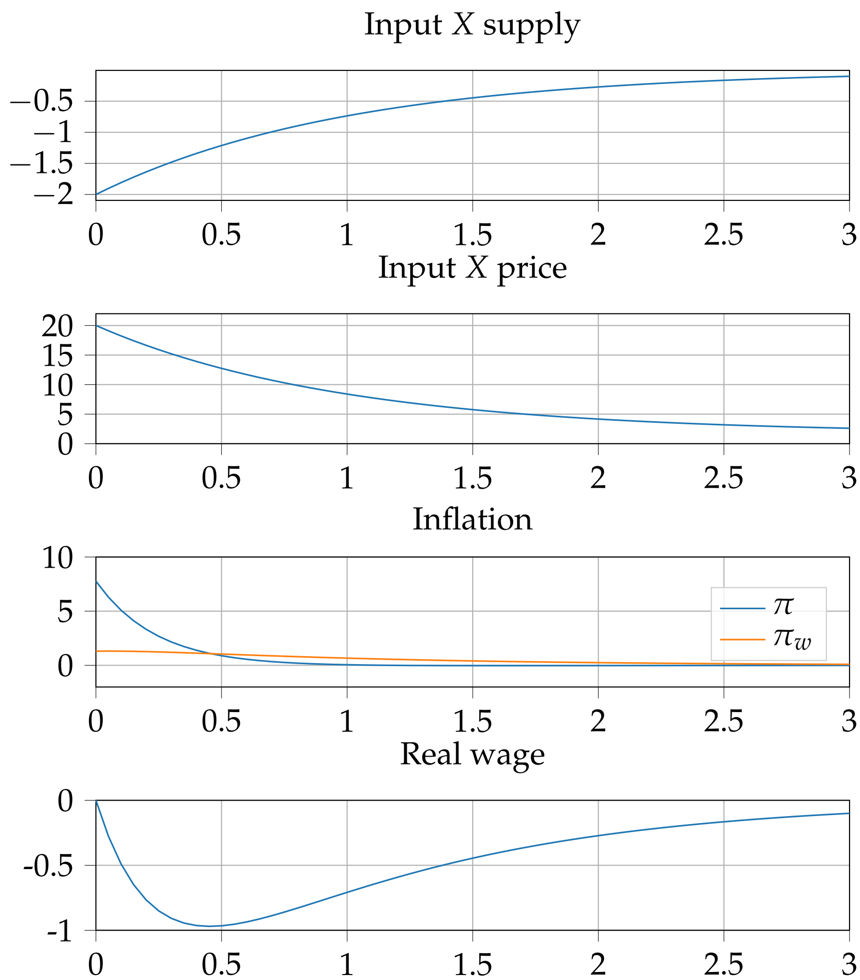

Figure 6: A supply shock

3.4 A Supply Shock

Consider the same economy's response to a temporary reduction in the endowment of input X. Suppose, for now, that the central bank responds in such a way as to keep employment constant at its initial steady state level, $n_t = 0$.

Again, the reaction of monetary policy is left implicit in the path of quantities. Since X falls, constant employment corresponds to a reduction in real output. It can be shown that this means that the central bank is increasing the real interest rate. However, as we shall see, the real rate increase that produces $n_t = 0$ is not large enough to achieve the natural allocation, given our chosen parameters.

[19]Notice that π_t is an instantaneous rate of inflation, expressed in annual terms. Since inflation falls relatively quickly in our example, measured quarterly inflation in the first quarter after the shock is lower than 12%.

The responses of mpl and mrs are now

$$mpl_t = \frac{s_X}{\epsilon} e^{-\delta t} x_0 < 0, \quad mrs_t = \sigma s_X e^{-\delta t} x_0 < 0,$$

while the response of the X good price is

$$p_{Xt} - w_t = \frac{1}{\epsilon} e^{-\delta t} n_0 > 0.$$

The main difference is that now the reduction in output reduces workers' mrs, via an income effect. This weakens real wage demands. Given the parameter choices in Table (1), the inflationary forces on the firms' side are still strong enough that we obtain positive wage and price inflation. In the representation of Figure 3 we are in the portion of the blue region that intersects the lower left quadrant. From Proposition 1, we also know that $mpl_0 < 0$ and $mrs_0 < 0$ implies that the real wage falls on impact for any parameter configuration.

The responses are illustrated in Figure 6. For ease of comparison, we pick a negative shock to x_0 that produces the same increase in the input price as the positive y_0 shock in Figure 5.

While nominal wages are growing less and the real wage drop is larger than in Figure 5, the overall shapes and magnitudes are not very different from the demand shock above. The crucial observation here is that if we scale shocks so that the input price response is the same, we are pinning down the change in the labor-to-X ratio, as

$$p_{X0} - w_0 = \frac{1}{\epsilon} (n_0 - x_0).$$

and the same ratio $n_0 - x_0$ determines

$$mpl_0 = \frac{s_X}{\epsilon} (n_0 - x_0).$$

Once we choose the quantitative size of the fall in $n_0 - x_0$ we have pinned down the inflationary impulse on the firms' side.

The main difference is that in this case the wage-price spiral mechanism is weaker, as workers' aspirations fall instead of increasing in the case of a supply shock. This explains why both price and wage inflation are lower in this case.

3.5 Supply Shocks and the Monetary Response

The response to the supply shock depend on how monetary policy adjusts. So far, we assumed a policy that keeps the employment path unchanged at $n_t = 0$. However, the natural level of employment depends in general on x_t. In particular, keeping employment and output at their the natural level requires $mrs_t = mpl_t$, and n_t^* can be derived from the condition

$$\sigma\left(s_L n_t^* + s_X x_t\right) + \eta n_t^* = \frac{s_X}{\epsilon}\left(x_t - n_t^*\right).$$

The responses of price and wage inflation when

$$n_t = n_t^* = \frac{\frac{1}{\epsilon} - \sigma}{\sigma s_L + \frac{s_X}{\epsilon} + \eta} s_X x_t$$

are plotted in Figure 7. Since our parametrization features a low degree of substitutability between labor and the input X, we have $\frac{1}{\epsilon} - \sigma > 0$ and a reduction in x_t lowers the natural level of employment, as shown in the first panel. The natural level of output $y_t^* = s_X x_t + s_L n_t^*$ is then lower for two reason, the direct effect of a lower x_t and the lower level of natural employment. There is a clear difference in the inflation paths when quantities are at their natural levels: we see positive price inflation, but negative wage inflation. This goes on as long as the real wage falls, once the real wage starts growing again, the signs of price and wage inflation flip. In other words, real wage adjustments always take place with nominal prices and wages moving in opposite directions.

This is not just an outcome of our choice of parameters. When quantities are at their natural level we have $mrs_t = mpl_t$ and both are equal, by definition, to the natural real wage ω_t^*. The inflation equations then become

$$\pi_t = \Lambda_p \int_t^\infty e^{-\rho(s-t)}\left(\omega_s - \omega_s^*\right) ds,$$
$$\pi_t^w = \Lambda_w \int_t^\infty e^{-\rho(s-t)}\left(\omega_s^* - \omega_s\right) ds.$$

The following general result follows immediately.

Proposition 3. *If quantities are at their natural level, price and wage inflation π_t and π_t^w are either both zero or have opposite sign.*

This result can be visualized in the diagram of Figure 3, by noticing that the regions where π and π^w have the same sign are either entirely above or entirely below the 45 degree line, where $mrs = mpl$.

Using the concepts introduced in Section 2, we can then say that if the output gap is always zero, conflict inflation is zero, i.e., a wage-price spiral is not present.[20]

[20] This result also explains why conflict inflation in this model is equal to the divine coincidence inflation of Rubbo (2020).

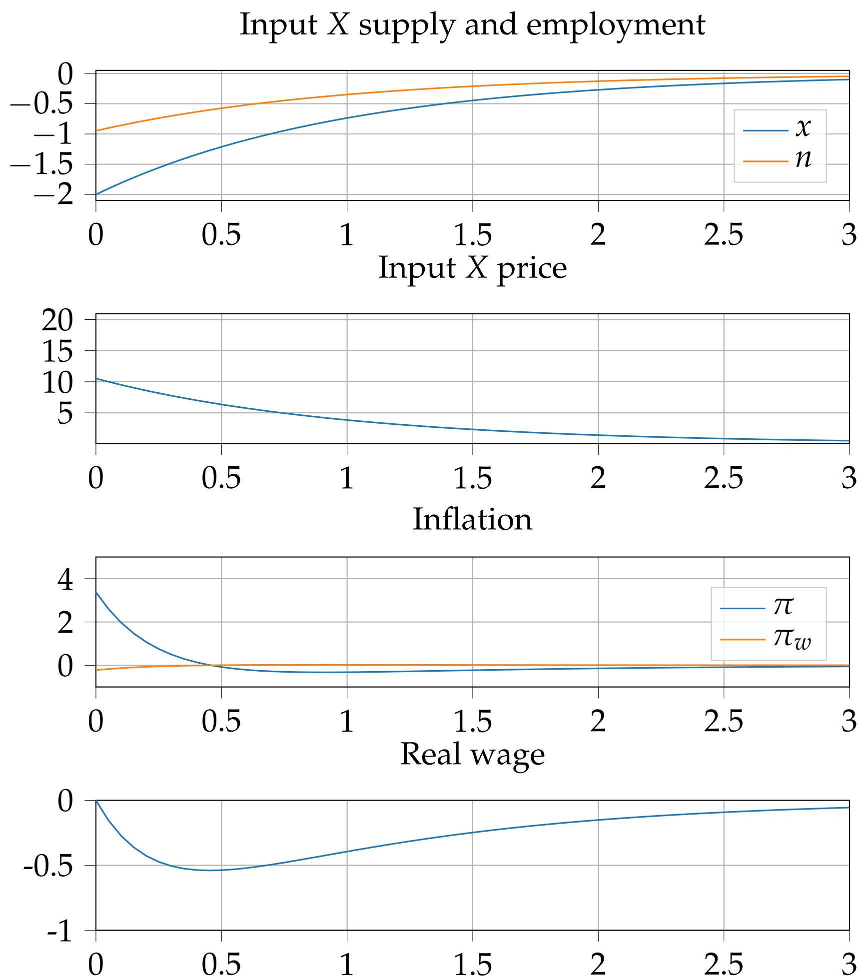

Figure 7: A supply shock with quantities on their natural path

Behind the similar adjustment patterns illustrated in Figures 5 and 6, there is a similar problem of excess demand producing positive conflict inflation. Excess demand can be caused either by a positive demand shock or by a negative supply shock coupled with an insufficient monetary policy response.

However, notice also that, as it's well known, an economy with both price and wage rigidities does not feature "divine coincidence," so a policy of keeping the output gap at zero, that is, of keeping quantities at their flexible price levels, is not necessarily optimal in our environment. We analyze optimal policy in the next section.

Comparing Figures 6 and 7 also shows that while employment falls more at the natural allocation, real wages fall less. This may seem surprising, but it is due to the fact the dynamics of the real wage are more strongly affected by mpl than by mrs, and mpl is higher along the path with lower employment. A different intuition for the same phenomenon is that lower employment reduces the pressure on the market for the scarce input, as seen in the second panel, weakening good inflation due to the high X price and increasing the real wage. Yet another intuition is that due to the fact that prices of goods and non-labor inputs are relatively more flexible than wages, the relation between real wages and employment is dominated by the labor demand side, so higher employment levels push down real wages.

3.6 Interpretation and Connections

This adjustment pattern shows both price and wage inflation, with price inflation stronger early on and wage inflation catching up later. If the central bank keeps always the economy at its flexible price allocation this pattern is not present, as price and wage inflation have opposite sign.

The examples presented are clearly just numerical simulations with parameters chosen mostly for clarity of exposition. Nonetheless, we believe there are some useful lessons and some interesting connections with recent experience.

Demand shocks and wage inflation. Our model helps to clarify that excess demand does not need necessarily to show up primarily through a tight labor market and high wage inflation. A commonly-held view is that excessive demand works its way from a tight labor market, to higher wages through the wage Phillips curve and, eventually, to higher prices. A demand shock then should produce increasing real wages. As we just showed, this is not, generally, the case. In the model, price and wage rigidities interact with general equilibrium forces on both goods and labor markets, and the direction of adjustment of the real wage is, in general, ambiguous. At a general level, the notion that real wages can potentially fall is obvious and commonly noted in the extreme case where nominal wages are fully rigid: in that case, the real wage must fall whenever inflation is positive.[21] Our analysis gives an easy to interpret condition for real wages to fall or rise, clarifying the economic forces at play.

An intuitive way of making our point here is to observe that inflation is in general caused by some form of scarcity on the supply side, relative to existing demand pressures. But there are multiple inputs on the supply side, labor inputs and non-labor inputs. Depending on the episode, scarcity can manifest itself more strongly in labor inputs or in non-labor inputs. When non-labor input scarcity dominates, price inflation will be faster than wage inflation.

[21] See, for example, Figure 6.3 in Gali (2015).

Small and large economies. Many papers measure supply shocks directly in terms of changes in input prices.[22] In this paper, we emphasize the general equilibrium nature of the price shock, by making the price p_X fully endogenous.

It is important to remark that the degree to which p_X should be treated as endogenous or endogenous depends on the size of the economy relative to the world economy. For a small open economy that trades X frictionlessly with the rest of the world (a reasonable approximation for some energy inputs), it make sense to redo the analysis by taking p_{Xt} as given and deriving x_t endogenously instead of shocking x_t and deriving p_{Xt} endogenously. The results for a supply shock would be similar. However, the effects of a demand shock that is completely idiosyncratic to the small open economy (that is, not correlated with a global demand shock) would be very different, as the relative scarcity of X in the world at large would not be affected by a localized shock to demand. On the other hand, a demand expansion in a large country would transmit to smaller economies as a supply shock, via the price p_X.

Pass-through from non-core to core inflation. We can identify the first phase of our three-phase responses as an initial period of high non-core inflation. Technically, the price p_X in our model does not appear directly in the consumer price index, because X is only used as an input, not as a final good. Therefore, there is no distinction between core and non-core inflation in the model. However, it is easy to modify the model to allow for direct consumption of X, or for multiple sectors, some of which use X more intensively than others, and make the distinction between core and non-core more explicit. The fact that the response of p_t lags the response of p_{Xt} shows that our model features a clear mechanism for pass-through from non-core inflation to core inflation. Recent work by Ball et al. (2022) shows empirically that this pass-through has been high in the post-pandemic period.

A related observation is that the fact that p_{Xt} is falling after jumping at $t = 0$ is not in contradiction with the fact that supply constraints are crucial for the inflation episode. It is the level of p_{Xt}, not its rate of change, that reflects the underlying scarcity in the economy, i.e., the a high labor to non-labor inputs ratio $n_t - x_t$, and this scarcity is a crucial driver of the high inflation rate in goods through its effects on mpl_t.

Non-linear Phillips curves. Many economists have pointed out the potentially important role of a non-linear Phillips curve in explaining recent experience, see, e.g., Benigno

[22] For example, this is the strategy in the model used in Bernanke and Blanchard (2023).

and Eggertsson (2023). Our model is linearized, but is linearized around a steady state that captures the economy's state at the moment when the shock hits. Therefore, we can easily see the effect on non-linearities through the parameter s_X in the linearized model. That parameter is not a model's constant, but depends on initial conditions. In particular, s_X is higher if the initial steady state features a relatively high initial ratio N_t/X_t. In other words, if the X input is already relatively scarce when the shock hits, the effects of the shock on inflation will be magnified. It would be interesting to explore model extensions in which the elasticity ϵ is also endogenous and depends on the state of the economy.

Notice that the non-linearity we are pointing out here is not non-linearity in the wage Phillips curve, which is the one that has received more attention, but rather non-linearity in the response of non-labor input prices which affects the price Phillips curve.[23]

Profits. A possible interpretation of the scarce input X is not as a market-supplied input, but rather as capturing fixed production capacity and other bottlenecks at the firm level. The formal analysis is slightly different when the input is fixed at the firm level instead of being fixed economy-wide and frictionlessly traded.[24] But the qualitative responses are similar.

There is, however, a marked difference in interpretation between a model with a market-supplied input X and a model with fixed capacity. In the first model, observed profit margins at the firm level fall in response to the shocks analyzed, because nominal prices increase less than marginal costs, due to stickiness. In the second model, instead observed profit margins increase because firm profits include the shadow price of the scarce input X which increases sharply in all our examples.

The role of ϵ. In our examples, we have used a low elasticity $\epsilon = 0.1$. This low elasticity plays two roles: it magnifies the response of p_{Xt}, explaining the initial jump in non-core inflation, and it magnifies the response of mpl_t, explaining the prolonged inflation episode. To see the central role of this parameter, consider an example with all the same assumptions of our demand shock in Figure 5, but assume a Cobb-Douglas production function, with $\epsilon = 1$. The responses are plotted in Figure 8.

Two differences stand out with our baseline parametrization. First, there is a smaller response of the relative price of the X input in the second panel. With lower elasticity the

[23]Comin et al. (2023) use occasional binding constraints to study a model with a similar non-linearity in the price Phillips curve.

[24]In particular, a model with firm-specific, non-traded X, is a model with decreasing returns to labor at the firm level, which produces strategic complementarity in pricing that is absent in our model with constant returns.

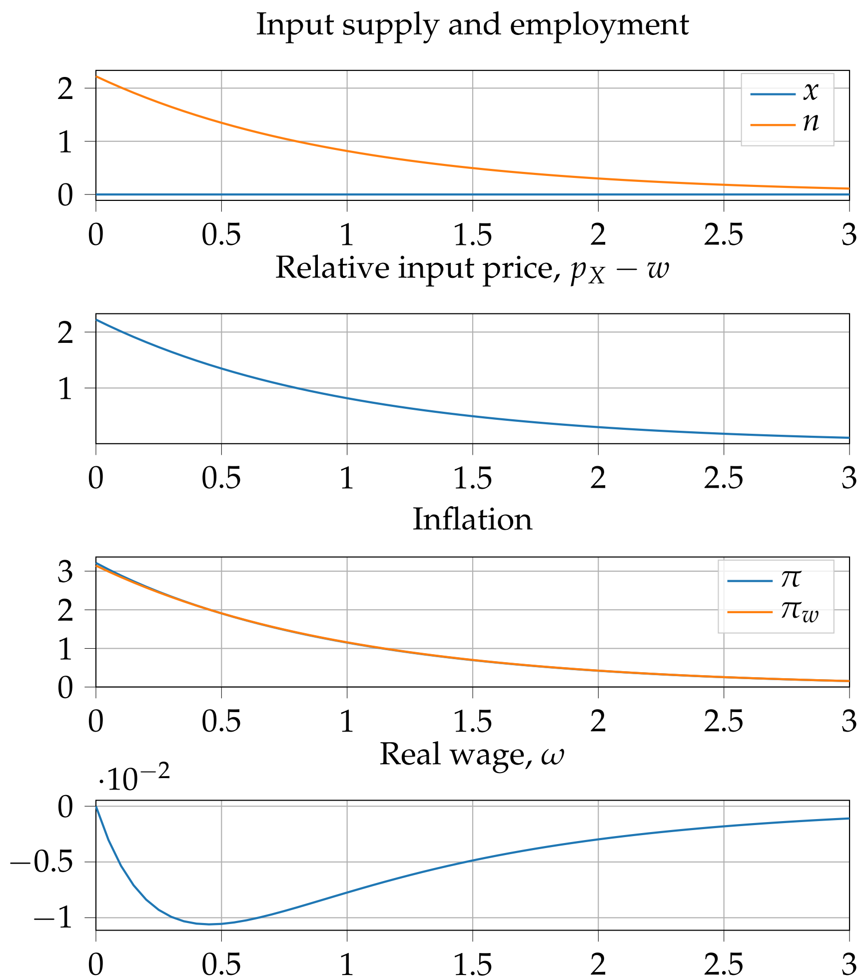

Figure 8: A demand shock with higher elasticity of substitution

relative scarcity of X has a smaller price effect (the effect is proportional to $1/\epsilon$, so it falls by a factor of 10). This implies a smaller overall inflation response. Second, the responses of wage and price inflation are almost indistinguishable and, consequently, the real wage is not affected. This is because the response of mpl is weaker while the response of mrs is unchanged (as we keep the value of s_L unchanged in the two examples).

This suggests that at the aggregate level, to capture episodes in which the relative scarcity of non-labor inputs trigger an inflationary episode, with a lagged response of wage inflation, a low degree of elasticity at the aggregate level is a needed ingredient.

4 Optimal Policy

In the previous section, we looked at economies in which the central bank unnecessarily stimulates the economy (demand shock) or in which the central bank responds weakly to a supply shock, so as to allow for both price and wage inflation (the supply shock with $n_t = 0$). The first example is a policy mistake, by construction. Of course, due to imperfect information and lags in the effects of monetary policy, similar mistakes can happen. However, in this section, we focus on the second shock, a supply shock, and ask what is the optimal response. Throughout, we assume monetary policy has perfect information on the underlying shocks and instantaneous control on the level of real activity.

The questions we address in this section are two: is it possible that following a supply shock the optimal response is to let the economy overheat, that is, to choose a positive output gap $y_t - y_t^* > 0$? Is it possible that the optimal response entails both positive price and wage inflation?

It is well known that divine coincidence fails in our environment. But that is really just a statement about feasibility: an outcome with no inflationary distortions, $\pi_t = \pi_t^w = 0$, and a zero output gap, $y_t = y_t^*$, are simply not feasible in our economy. The real wage needs to move in the flexible price equilibrium and that is incompatible with zero nominal inflation in p_t and w_t. Our contribution here is to characterize the signs of the deviations of π_t, π_t^w and $y_t - y_t^*$ from zero, under optimal policy.

In particular, Proposition 5 in the previous section tells us that if the central bank chooses $y_t = y_t^*$, then the signs of π_t and π_t^w will always be opposite. In other words, with a zero output gap the adjustment in the real wage never requires *both* price and wage inflation. Therefore, one could conjecture that generalized inflation, that is, inflation in both prices and wages is never optimal. However, a zero output gap is not necessarily optimal so that conjecture is not generally correct.

4.1 Optimal policy problem

Following standard steps, the objective function of the central bank can be derived as a quadratic approximation to the social welfare function:

$$\int_0^\infty e^{-\rho t} \frac{1}{2} \left[-\left(y_t - y_t^*\right)^2 - \Phi_p \pi_t^2 - \Phi_w \left(\pi_t^w\right)^2 \right] dt. \tag{17}$$

Deviations from first-best welfare come from two type of distortions: output deviations from its natural level, that is, from the level that equalizes the marginal benefit of producing goods with its marginal cost in terms of labor effort; and inflation in prices and wages that causes inefficient dispersion in relative prices of different varieties. The terms in (17) reflect these distortions. The value of the coefficients Φ_p and Φ_w depend on the model parameters and are derived and reported in the appendix.

The natural level of the real wage following a supply shock is

$$\omega_t^* = \frac{s_X}{\epsilon} \frac{\sigma + \eta}{\sigma s_L + \frac{s_X}{\epsilon} + \eta} x_t.$$

We can then express mpl and mrs in terms of the natural real wage and deviations of employment from its natural path

$$mpl_t = \omega_t^* - \frac{s_X}{\epsilon} \left(n_t - n_t^* \right), \tag{18}$$

$$mrs_t = \omega_t^* + \left(\sigma s_L + \eta \right) \left(n_t - n_t^* \right). \tag{19}$$

The optimal policy problem is to maximize (17), subject to the constraints coming from price setting (10) and (11), condition

$$\dot{\omega}_t = \pi_t^w - \pi_t,$$

and the aggregate production function

$$y_t = s_L n_t + s_X x_t.$$

The optimality conditions that characterize an optimal policy are derived in the Appendix.

4.2 Examples

We now consider examples that illustrate a variety of possible outcomes.

It helps the interpretation of the policy trade-offs to focus on the simple case of a permanent shock to x_t. With this shock, in all our examples, in the long run, the real wage is permanently lower and so are mpl and mrs, so that the economy eventually reaches a new steady state with zero inflation and zero output gap. To reach that new steady state requires ω_t to fall. This can be achieved by many combinations of price and wage inflation or deflation, as long as price inflation is larger than wage inflation. The question is what is the optimal way to get there.

Example 1. A symmetric case

Our first example is an economy with parameters that have the following properties:[25]

[25] The parameters are as follows:

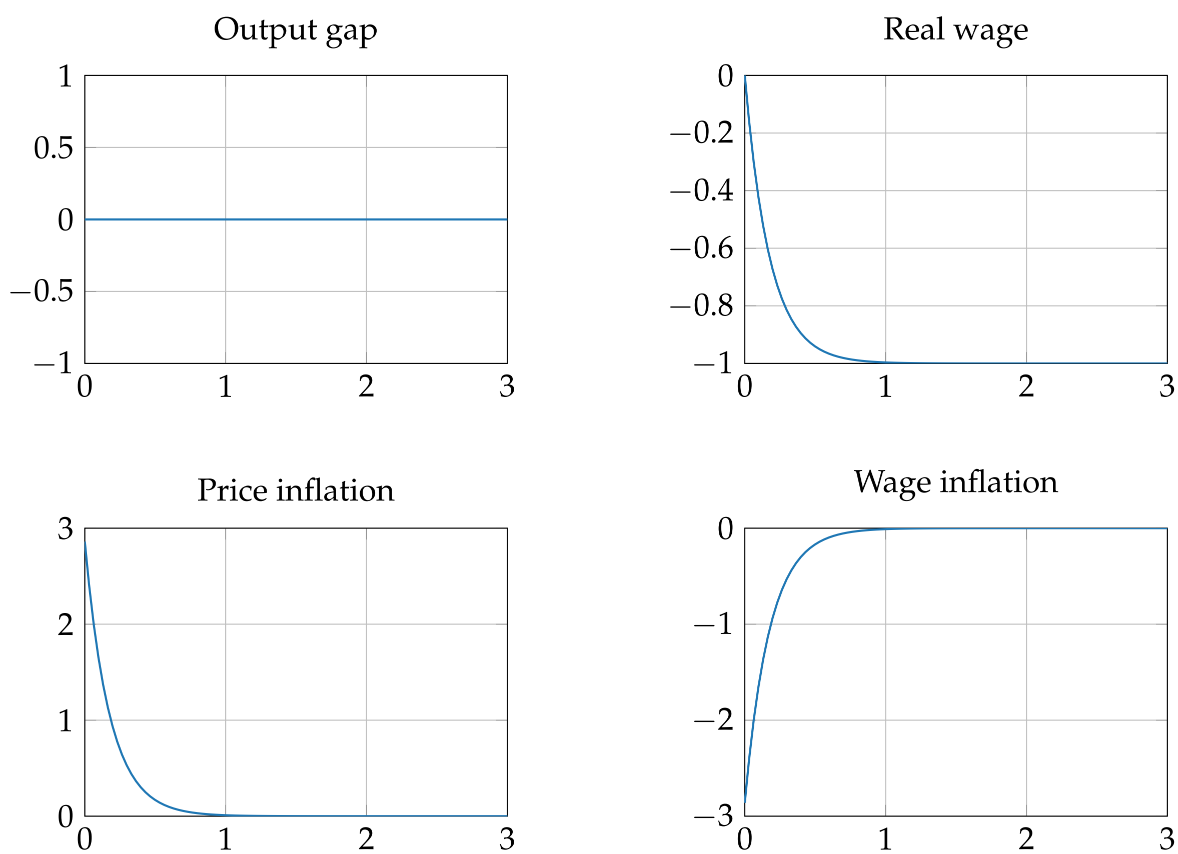

Figure 9: A symmetric example

- the welfare costs of wage and price inflation enter symmetrically the objective function, $\Phi_p = \Phi_w$;
- wages and prices are equally sticky, $\Lambda_p = \Lambda_w$;
- the output gap has symmetric effects on mpl and mrs.[26]

Figure 9 illustrates optimal policy outcomes in this example. Given the symmetry of the problem, the reduction in real wages is achieved by spreading the adjustment equally between nominal wage deflation and nominal price inflation. The output gap is kept exactly at zero. This example is clearly a knife edge case and relies on the symmetry of the parameters. As soon as we abandon this symmetry things get more interesting.

$\sigma = 1$	$\eta = 0$	$\rho = 0.05$	
$s_X = 1/2$	$\epsilon = 1,$	$\varepsilon_C = 1.5$	$\varepsilon_L = 3$
$\lambda_p = 4$	$\lambda_w = 4$		

[26]Given the expressions above this requires $\frac{s_X}{\epsilon} = \sigma s_L + \eta$.

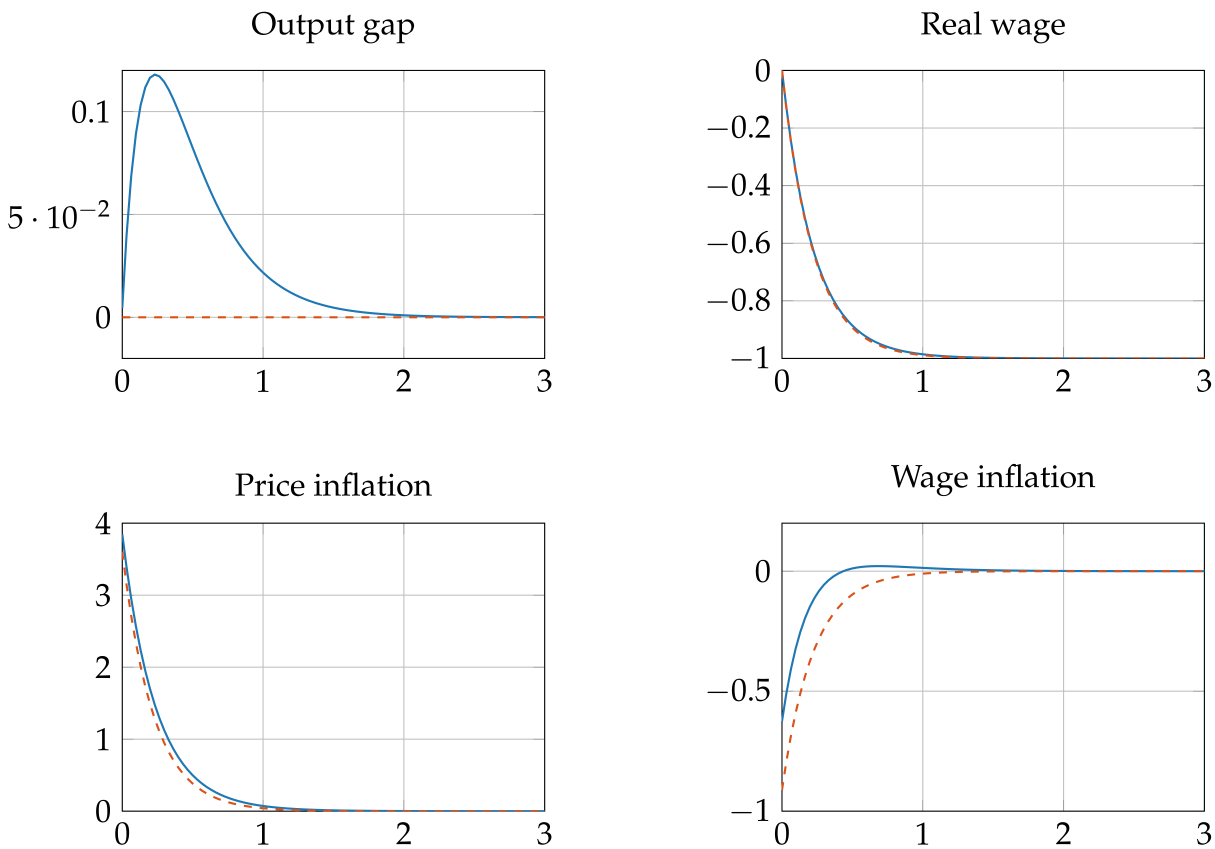

Figure 10: An optimal hot economy

Example 2. A hot economy

In the second example, the parameters chosen imply that:[27]

- the welfare cost of wage inflation is larger than that of price inflation, $\Phi_p < \Phi_w$;
- wages are more sticky than prices, $\Lambda_p > \Lambda_w$.

We still have a set of parameters that imply roughly symmetric effects of the output gap on *mpl* and *mrs*, but the differences above are sufficient to obtain a quite different result. Figure 10 illustrates optimal policy outcomes in this case. For comparison, in the figure we also plot outcomes under a zero output gap policy (red dashed lines).

In this second example, it is optimal to have a positive output gap throughout the transition. To get some intuition for this result it is useful to recall from equations (10)-

[27] The parameters are as follows:

$\sigma = 1$	$\eta = 0$	$\rho = 0.05$	
$s_X = 0.1$	$\epsilon = 1,$	$\varepsilon_C = 1.5$	$\varepsilon_L = 4$
$\lambda_p = 4$	$\lambda_w = 2$		

(11) and (18)-(19) that increasing the output gap has two direct effects. By decreasing mpl it leads to higher price inflation, by increasing mrs it leads to higher wage inflation. If we start at a zero-output-gap policy, with positive price inflation and negative wage inflation, the effect can be welfare improving because the welfare cost of price inflation is smaller than the welfare cost of wage deflation.

The role of $\Lambda_p > \Lambda_w$ is subtler and has to do with dynamics. With $\Lambda_p > \Lambda_w$ and $\xi_p \approx \xi_w$ a higher output gap also implies a faster declining real wage. Since a lower real wage in the future requires less adjustment, lowering the real wage today is welfare improving from a dynamic point of view. Therefore, a parametrization with $\Lambda_p > \Lambda_w$ makes it easier to obtain examples with a welfare improving positive output gap.[28]

By choosing parameters that yield the opposite inequality, $\Phi_p > \Phi_w$, in the welfare coefficients it is possible to construct examples of the opposite: economies in which it is optimal to run a negative output gap in the transition.

Example 3. Generalized inflation and a hot economy

Our third example is a variant on the second example, with an even larger welfare cost associated to wage dispersion (a larger Φ_w), a larger distance between price and wage stickiness, and with a smaller value of the elasticity of substitution between labor and the X input, ϵ, which implies that running a hot economy has larger benefits in terms of lowering the real wage by having a larger effect on firms' marginal costs and thus on price inflation.[29]

The parametric choices above amplify the forces we saw in example 2 and they imply that there is an interval during the transition in which the optimal policy yields both a hot economy ($y_t > y_t^*$) and generalized price and wage inflation ($\pi_t > 0$ and $\pi_t^w > 0$).[30]

This result is surprising from a static point of view. Given the welfare function (17), at any point in time in which $y_t > y_t^*, \pi_t > 0$ and $\pi_t^w > 0$ it is welfare improving, from a static point of view, to reduce y_t, as it unambiguously lowers π_t and π_t^w and leads to an increase in the current payoff. However, from a dynamic perspectives there is an additional argument. Increasing y_t at time t has the effect of increasing π_s and π_s^w in all

[28] The discussion of Figure X in the Appendix expands on this argument.

[29] The parameters are as follows:

$\sigma = 1$	$\eta = 0$	$\rho = 0.05$	
$s_X = 0.1$	$\epsilon = 0.1,$	$\varepsilon_C = 1.5$	$\varepsilon_L = 8$
$\lambda_p = 4$	$\lambda_w = 1$		

[30] Notice, that these qualitative features can actually be seen in example 2 too, but it is useful to choose an example where they are more clearly visible.

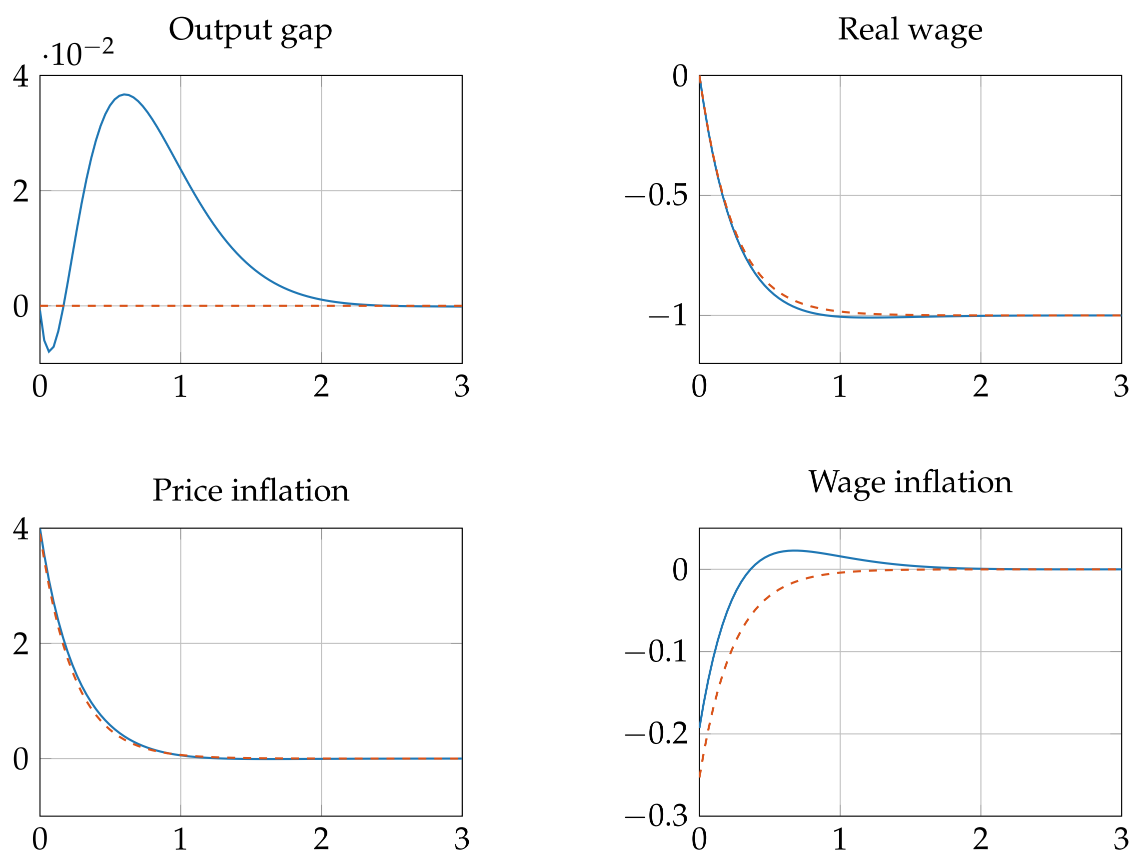

Figure 11: An example with generalized inflation and a hot economy

previous periods, due to the forward looking element in price setting. This entails welfare gains in early periods in the transition in which $\pi^w_s < 0$. Through this forward looking force a positive output gap later in the transition can be beneficial even if, at that point $\pi^w_t > 0$.

Now, while this example is theoretically interesting, it does have the flavor of a overly sophisticated form of forward guidance. Therefore, we do not think it provides a strong argument in favor of policies that deliver $y_t > y^*_t, \pi_t > 0$ and $\pi^w_t > 0$ at the same time. In the context of the present model, given the distortions it captures, it is hard to make a compelling practical case that the combination of a hot economy with positive wage and price inflation are a desirable outcome, even in response to a supply shock and even in presence of inelastic supply constraints.[31]

[31] This does not mean that such a case could not maybe be made in richer models, which capture, just to make an example, the benefits of labor reallocation. But that is clearly outside the scope of this paper.

5 Adaptive Expectations and Real Rigidities

The model with rational expectations analyzed so far has two embedded features: the effect of any shock tends to be front-loaded, as agents perfectly anticipate its future effects on prices, and there is no room for persistent deviations of inflation expectations from target, as agents anticipate the economy will go back to its initial steady state. We now explore variants of the model that deviate from rational expectations and allow for more inertial responses by introducing two ingredients: adaptive expectations on expected inflation and a gradual adjustment of price-setters' and wage-setters' relative price objectives. For this second ingredient we use the label "real rigidities."

The objective of this sections is twofold. First, by allowing for inertial responses we allow the feedback between price and wages to play out more explicitly over time: shocks that produce high prices in the goods market only gradually lead to higher wage demands in the labor market. In other words, the wage-price spiral instead of playing out in the "virtual time" of best responses, plays out in the observed dynamics of prices and wages. Second, by allowing for deviations of inflation expectations from target we capture the common concern of central bankers that prolonged episodes of high inflation may lead to de-anchoring of inflation expectations.

From an empirical perspective, we show that adaptive expectations and inertia reinforce the main prediction of the baseline model in Section 3: there is a lagged and persistent increase in wage inflation following a large increase in price inflation. However, the medium term implications are different depending on the sources of inertia: if inertia is mostly due to de-anchoring, inflation can take a long time to go back to target, absent a recession, if instead inertia is mostly due to real rigidities, then a path of immaculate disinflation is possible.

Let us begin by rewriting the price setting conditions making explicit agents' expectations. Letting E_t^f and E_t^w denote firms' and workers' expectations, we can write

$$
\begin{aligned}
p_t^* &= \left(\rho+\lambda_p\right) E_t^f \int_t^\infty e^{-\left(\rho+\lambda_p\right)(\tau-t)} \left(w_\tau + s_X \left(p_{X\tau} - w_\tau\right)\right) d\tau = \\
&= w_t + \left(\rho+\lambda_p\right) E_t^f \int_t^\infty e^{-\left(\rho+\lambda_p\right)(\tau-t)} s_X \left(p_{X\tau} - w_\tau\right) d\tau + E_t^f \int_t^\infty e^{-\left(\rho+\lambda_p\right)(\tau-t)} \dot{w}_\tau d\tau.
\end{aligned}
$$

Reset prices are decomposed in three components: the current nominal wage, the expected path of the relative price of input X vs labor, the expected path of future wage inflation.

We assume that agents expect a constant inflation rate over the future horizon

$$E_t^f \dot{w}_t = \pi_t^{w,e},$$

and expected inflation is driven by the simple adaptive, constant-gain rule

$$\dot{\pi}_t^{w,e} = \gamma \left(\dot{w}_t - \pi_t^{w,e} \right). \tag{20}$$

Moreover, we assume that agents perfectly anticipate the path of real variables n_t, x_t, y_t and can deduce the path of the relative price $p_{Xt} - w_t$ from the equilibrium condition in factor markets

$$x_t - n_t = -\epsilon \left(p_{Xt} - w_t \right).$$

Combining these assumptions with exponentially decaying, one time shocks at date 0, as in Section 3, we can substitute in the expression above for p_t^*, substitute in the inflation equation (8), and obtain the following

$$\dot{p}_t = \lambda_p \left[\frac{s_X}{\epsilon} \frac{\rho + \lambda_p}{\rho + \lambda_p + \delta} \left(n_t - x_t \right) - \left(p_t - w_t \right) \right] + \frac{\lambda_p}{\rho + \lambda_p} \pi_t^{w,e}. \tag{21}$$

Similar steps on the wage setting side of the model lead to

$$\dot{w}_t = \lambda_w \left[\frac{\rho + \lambda_w}{\rho + \lambda_w + \delta} \left(\sigma y_t + \eta n_t \right) - \left(w_t - p_t \right) \right] + \frac{\lambda_w}{\rho + \lambda_w} \pi_t^e, \tag{22}$$

where price inflation follows the adaptive rule

$$\dot{\pi}_t^e = \gamma \left(\dot{p}_t - \pi_t^e \right). \tag{23}$$

Equations (20)-(23) can be solved forward for any given initial condition w_0, p_0.

An example of de-anchoring

Figure 12 shows the response of inflation to a supply shock in a numerical example analogous to the one shown in Figure 6, except for the assumption of adaptive expectations. The parameters are the same as in Table 1 and we set $\gamma = 1$. There are two main differences from the case of rational expectations. First, wage inflation is weaker on impact and only picks up gradually, as initially workers do not anticipate higher prices and so do not start trying to catch up until their purchasing power has actually been eroded by

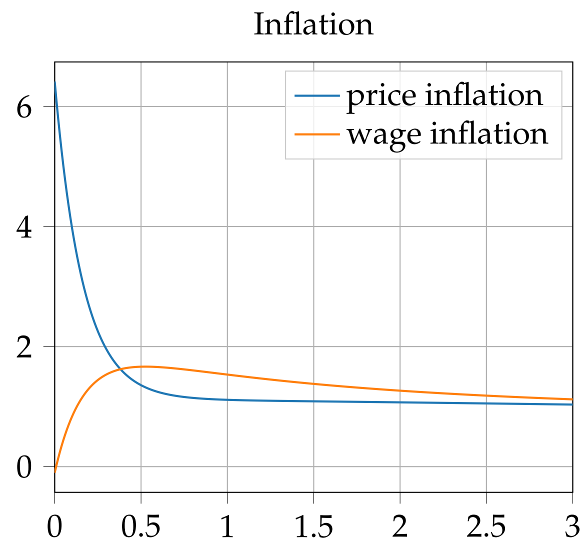

Figure 12: A supply shock with adaptive expectations

past inflation.[32] Second, there is a very persistent effect on inflation, due to the learning dynamics. Since ρ is small, the coefficients on the expected inflation terms on the right-hand side of equations (21)-(22) are close to 1. This implies that even though all quantities and all relative price targets for workers and firms have gone back to steady state, we can have a prolonger period of self-sustaining inflation. This is a case of de-anchoring, in which the only way to go back to target inflation faster is for the central bank to keep activity low for some time.

The wage-price spiral is active in the self-sustaining phase of prolonged inflation, but it is exactly balanced on the two sides, so real wages remain constant.

An example with real rigidities

We now consider a different source of inertia, due to a gradual adjustment of the relative price targets of price and wage setters. In particular, we assume that changes in real marginal costs and in the marginal rate of substitution between consumption and leisure only gradually change the behavior of price and wage setters. We replace the inflation

[32] Notice that given that n is kept on its pre-shock path ($n = 0$) and that output falls due to the supply shock ($y_0 = s_X x_0 < 0$), there is an income effect that depresses the real wage demands of workers on impact, causing a very small initial nominal wage deflation, which is barely visible in the figure.

dynamics above, (21)-(22) with the following equations

$$\dot{p}_t = \lambda_p \left[a_t^p - (p_t - w_t)\right] + \frac{\lambda_p}{\rho + \lambda_p}\pi_t^{w,e},$$
$$\dot{w}_t = \lambda_w \left[a_t^w - (w_t - p_t)\right] + \frac{\lambda_w}{\rho + \lambda_w}\pi_t^{e}.$$

The real aspirations of price setters and wage setters, a_t^p and a_t^w, follow the adjustment equations

$$\dot{a}_t^p = \xi_p \left[\frac{s_X}{\epsilon}\frac{\rho + \lambda_p}{\rho + \lambda_p + \delta}(n_t - x_t) - a_t^p\right],$$

and

$$\dot{a}_t^w = \xi_w \left[\frac{\rho + \lambda_w}{\rho + \lambda_w + \delta}(\sigma y_t + \eta n_t) - a_t^w\right].$$

Aspirations are driven by the same forces that drive them in the baseline model, which, in the case of firms are anticipated real input prices, captured by the term $\frac{s_X}{\epsilon}\frac{\rho+\lambda_p}{\rho+\lambda_p+\delta}(n_t - x_t)$, and in the case of workers are anticipated marginal rates of substitution between consumption and leisure, captured by $\frac{\rho+\lambda_w}{\rho+\lambda_w+\delta}(\sigma y_t + \eta n_t)$. However, these forces only gradually modify the aspirations of firms in terms of the desired margins ($p_t - w_t$ for the firms and $w_t - p_t$ for the workers).

We assume that the inflation expectations $\pi_t^{w,e}$ and π_t^e still follow the learning processes (20) and (23), so this version of the model includes both inertia caused by slow adjustment of inflation expectations and inertia caused by real rigidities. The choice to combine the two is because an interpretation of the real rigidities here is also some form of bounded rationality in processing observed changes in input prices and changes in labor market conditions, and combining that with perfect foresight on future price paths seems less natural. However, to focus on the role of real rigidities we choose a parametrization with a lower $\gamma = 0.1$, relative to the parametrization used for Figure 12, so inflation expectations play a more limited role. For the parameters ξ_p and ξ_w we experiment with values equal to 4 and 1, so the degree of real rigidity in the goods and labor market mirror the degree of nominal rigidity (capture by λ_p and λ_w). The inflation responses to the same supply shock used above are reported in Figure 13.

In this economy, both price and wage inflation display hump-shaped responses and the wage response is more delayed and more persistent than in the rational expectations baseline. The delay in the wage response is essentially due to the same reason as in model with only adaptive inflation expectations: wage setters only start to demand higher nominal wages when price inflation has been going on for a while and has moved real wages away from their aspirations. The additional delay here is due to the fact that prices

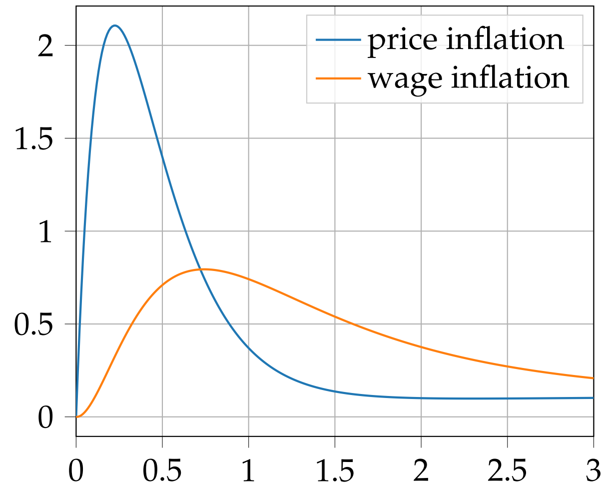

Figure 13: A supply shock with adaptive expectations

also take longer to respond, due to the real rigidity in price setting.[33]

The example in Figure 13 comes closest to capture an immaculate inflation-disinflation scenario. The shock causes persistent responses of prices and wages. The persistence is purely due to the fact that price setters takes some time to respond and wage inflation follows with further delay because wage setters only start responding after price setters have increased the price level enough to lower $w - p$. The persistence of wage inflation in this scenario is not a symptom of persistent overheating in the labor market, but of a gradual return to pre-shock trends for the real wage.

6 Conclusions

We explored the wage-price spiral in a canonical model of price and wage setting.

Interpreting inflation as the outcome of inconsistent aspirations for the real wage (or other relative prices) opens the door to many theoretical and empirical questions. We are especially interested in extending our work to explore potential sources of inertia in the inflation process, expanding the models explored in Section 5.

In the model analyzed here there is an instantaneous connection between the output gap and the real wage aspirations of workers' and firms. However, it is plausible that workers' real wage aspirations respond gradually to changes in labor market conditions. Similarly, changes in goods market conditions could affect slowly firms' expected profit

[33]The real rigidity in wage setting does not really play an important role in this simulation, because with a pure supply shock to x the effect on $\sigma y + \eta n$ is very small, so workers' aspirations are essentially constant at 0. In line with this observation, simulations with larger and smaller values of ξ_w produce responses very similar to those in Figure 13. Of course, in the case of other shocks this is no longer the case.

margins. These are sources of inertia in inflation that come from agents' views on relative prices, and so are different from sources of inertia tied to future inflation expectations, on which most research has focused on. Even if inflation expectations are well anchored it is possible for inflation to persist if the disagreement between firms and workers is inertial. On the empirical front, while there is a large literature measuring inflation expectations, there has been limited effort so far at measuring workers' and firms' aspirations for real pay and for real profit margins.

References

Amiti, Mary, Sebastian Heise, Fatih Karahan, and Aysegul Sahin, "Inflation Strikes Back: The Role of Import Competition and the Labor Market," *NBER Macroeconomics Annual*, 2023, *38.*

Ball, Laurence, Daniel Leigh, and Prachi Mishra, "Understanding US Inflation During the COVID Era," *Brooking Papers on Economic Activity*, 2022.

Barlevy, Gadi and Luojia Hu, "Unit Labor Costs and Inflation in the Non-Housing Service Sector," Chicago Fed Letter, Federal Reserve Bank of Chicago 2023.

Benigno, Pierpaolo and Gauti Eggertsson, "It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve," Technical Report, NBER WP 2023.

Bernanke, Ben and Olivier Blanchard, "What caused the US pandemic-era inflation?," *Brooking Papers on Economic Activity*, 2023.

Blanchard, Olivier, "The Wage Price Spiral," *The Quarterly Journal of Economics*, August 1986, *101* (3), 543–566.

— **and Jordi Gali**, "The Macroeconomic Effects of Oil Shocks. Why are the 2000s So Different from the 1970s?," Technical Report 2007.

— **and** — , "Real Wage Rigidities and the New Keynesian Model," *Journal of Money, Credit and Banking*, 2007, *39* (1), 35–65.

Blanchard, Olivier J., "In defense of concerns over the 1.9 trillion relief plan," Technical Report, Peterson Institute for International Economics February 2021.

Bruno, Michael and Jeffrey Sachs, *Economics of Worldwide Stagflation*, Cambridge Massachusetts: Harvard University Press, 1985.

Christiano, Lawrence J and Martin Eichenbaum, "Current real-business-cycle theories and aggregate labor-market fluctuations," *The American Economic Review*, 1992, pp. 430–450.

Comin, Diego, Robert C. Johnson, and Callum Jones, "Supply Chain Constraints and Inflation," Technical Report 31179, NBER 2023.

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin, "Optimal monetary policy with staggered wage and price contracts," *Journal of Monetary Economics*, 2000, *46* (2), 281–313.

Gagliardone, Luca and Mark Gertler, "Oil Prices, Monetary Policy and Inflation Surges," Technical Report, National Bureau of Economic Research 2023.

Gali, Jordi, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications Second edition* number 10495. In 'Economics Books.', Princeton University Press, 2015.

Giannoni, Marc and Michael Woodford, "Optimal inflation-targeting rules," in "The Inflation-Targeting Debate," University of Chicago Press, 2004, pp. 93–172.

Kabaca, Serdar and Kerem Tuzcuoglu, "Supply Drivers of U.S. Inflation Since the COVID-19 Pandemic," Technical Report, Bank of Canada 2023.

Lorenzoni, Guido and Ivan Werning, "Inflation is Conflict," Technical Report, MIT 2022.

Rotemberg, Julio J and Michael Woodford, "Oligopolistic pricing and the effects of aggregate demand on economic activity," *Journal of political Economy*, 1992, *100* (6), 1153–1207.

Rowthorn, Robert J., "Conflict, inflation and money," *Cambridge Journal of Economics*, 1977, *1*, 215–239.

Rubbo, Elisa, "Networks, Phillips Curves, and Monetary Policy," Working Paper 2020.

A Appendix

A.1 Derivations for Section 1

Derivation of equations (10) and (11)

Differentiate both sides of (4) and (8) with respect to time to get

$$\dot{p}_t^* = -\left(\rho + \lambda_p\right)(w_t - mpl_t) + \left(\rho + \lambda_p\right) p_t^*,$$

and

$$\ddot{p}_t = \lambda_p \left(\dot{p}_t^* - \dot{p}_t\right).$$

Substituting $\dot{p}_t^*$ from the first equation on the right-hand side of the second equation and changing notation for inflation, yields

$$\dot{\pi}_t = \lambda_p \left(-\left(\rho + \lambda_p\right)(w_t - p_t - mpl_t) + \left(\rho + \lambda_p\right)(p_t^* - p_t) - \pi_t\right).$$

Using $\lambda_p \left(p_t^* - p_t\right) = \pi_t$ and rearranging gives

$$\dot{\pi}_t = -\lambda_p \left(\rho + \lambda_p\right)(w_t - p_t - mpl_t) + \rho \pi_t,$$

which corresponds to (10). Equation (11) is derived in a similar way.

A.2 Additional Material for Section 2

Real wage dynamics

Combining equations (10)-(12) gives the second order ordinary differential equation

$$\ddot{\omega}_t = \rho \dot{\omega}_t + \Lambda \left(\omega_t - \tilde{\omega}_t\right), \tag{24}$$

where

$$\Lambda = \Lambda_p + \Lambda_w,$$

and where

$$\tilde{\omega}_t = \alpha mpl_t + (1 - \alpha) mrs_t,$$

is the average of the aspirations of workers and firms, weighted by the relative degree of price rigidity

$$\alpha = \frac{\Lambda_p}{\Lambda_p + \Lambda_w}.$$

The next proposition provides the saddle-path stable solution of (24).

Proposition 4. *The real wage satisfies the first order ODE*

$$\dot{\omega}_t = r_1 \omega_t + \Lambda \int_t^{\infty} e^{-r_2(\tau - t)} \tilde{\omega}_\tau d\tau, \tag{25}$$

where $r_1$and r_2 are the roots of the quadratic equation

$$r(r-\rho)=\Lambda,$$

and satisfy $r_1 < 0 < \rho < r_2$. The solution of (25) is

$$\omega_t = e^{r_1 t}\omega_0 + \int_0^\infty H_{s,t}\tilde{\omega}_s ds, \tag{26}$$

where $H_{s,t}$ is defined as

$$H_{s,t} = \frac{\Lambda}{r_2 - r_1}\left(e^{\min\{r_1(t-s),-r_2(s-t)\}} - e^{r_1 t - r_2 s}\right).$$

Proof. Since $\Lambda > 0$ there are two real eigenvalues r_1, r_2 that solve

$$r^2 - \rho r - \Lambda = 0.$$

The ODE can then be written as

$$(\partial - r_1)(\partial - r_2)\omega_t = -\Lambda\tilde{\omega}_t$$

where ∂ is the time-derivative operator. Integrating forward gives

$$(\partial - r_1)\omega_t = -\frac{1}{\partial - r_2}\Lambda\tilde{\omega}_t = \Lambda\int_t^\infty e^{-r_2(\tau-t)}\tilde{\omega}_\tau d\tau,$$

which gives (25). Integrating backward gives

$$\omega_t = e^{r_1 t}\omega_0 + \Lambda\int_0^t e^{r_1(t-s)}\int_s^\infty e^{-r_2(\tau-s)}\tilde{\omega}_\tau d\tau ds.$$

Changing the order of integration, the double integral on the right-hand side becomes

$$\int_0^t\int_0^\tau e^{r_1(t-s)}e^{-r_2(\tau-s)}\tilde{\omega}_\tau ds d\tau + \int_t^\infty\int_0^t e^{r_1(t-s)}e^{-r_2(\tau-s)}\tilde{\omega}_\tau ds d\tau$$

which gives

$$\omega_t = e^{r_1 t}\omega_0 + \Lambda\int_0^t \frac{e^{r_1(t-s)} - e^{r_1 t - r_2 s}}{r_2 - r_1}\tilde{\omega}_\tau ds + \Lambda\int_t^\infty \frac{e^{-r_2(s-t)} - e^{r_1 t}e^{-r_2 s}}{r_2 - r_1}\tilde{\omega}_\tau ds,$$

which can be written compactly as (26). □

The second term in (24) shows that real wage dynamics are driven by a forward-looking expression, capturing the anticipated levels of the average aspiration $\tilde{\omega}_t$.

The first term in (25) shows that the real wage tends to mean revert, since $r_1 < 0$. The intuition for the mean-reversion is that a higher ω_t increases $\omega_t - mpl_t$, i.e., the distance between the real wage and the firms' aspiration mpl_t, pushing up price inflation. It also

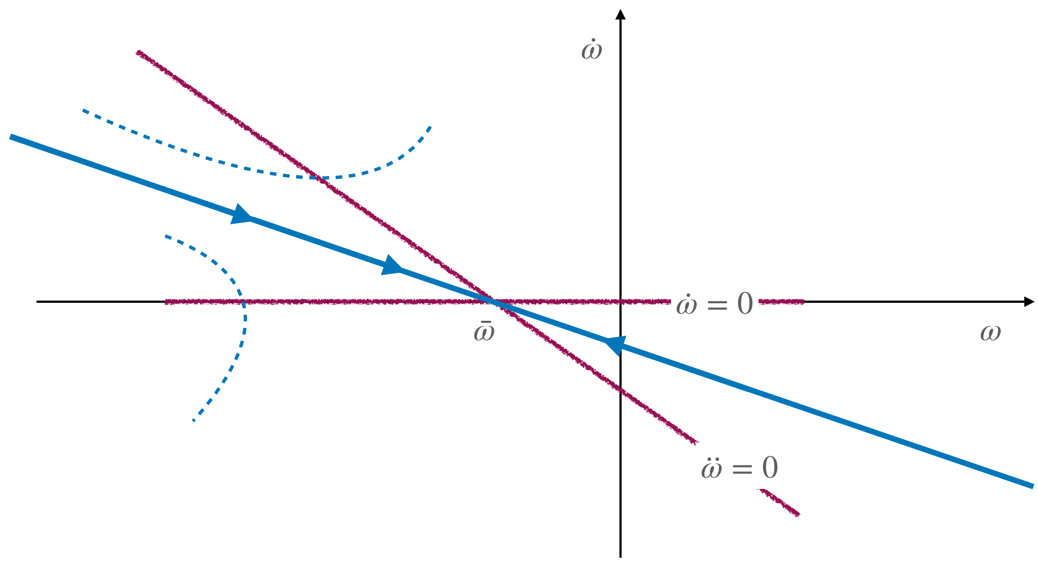

Figure 14: A permanent shock

reduces $mrs_t - \omega_t$, i.e., the distance between the workers' aspiration mrs_t and the real wage, which pushes down wage inflation. Higher price inflation and lower wage inflation reduce the real wage.

Graphical analysis

Suppose the economy is in steady state with all variables equal to 0. At date 0, unexpectedly, there is a one time, permanent reduction in mpl, which goes to $\overline{mpl} < 0$. The level of mrs remains unchanged at 0.

The phase diagram in Figure 14 represents the second order ODE (24). The stationary locus $\dot{\omega} = 0$ coincides with the horizontal axis. The stationary locus $\ddot{\omega} = 0$ is downward sloping. Both are drawn in purple. The saddle path, in blue, is given by the equation

$$\dot{\omega}_t = r_1 \left(\omega_t - \overline{\omega} \right),$$

where

$$\overline{\omega} = \frac{\Lambda_p}{\Lambda_p + \Lambda_w} \overline{mpl}$$

is the constant value of $\tilde{\omega}_t$ after the shock and is also the long-run level of the real wage. The expression for the saddle path comes from 25, using the condition $-r_1 r_2 = \Lambda_p + \Lambda_w$.

The diagram shows that starting at $\omega_0 = 0$, we initially have $\dot{\omega}_t < 0$. Gradually, as the real wage reaches its new long-run level $\overline{\omega}$, this effect goes away.

There are initially two forces pushing up price inflation: a permanently higher conflict component, plus a temporarily positive adjustment component, reflecting the initial fall in the real wage. On the wage inflation side, adjustment inflation has the opposite effect

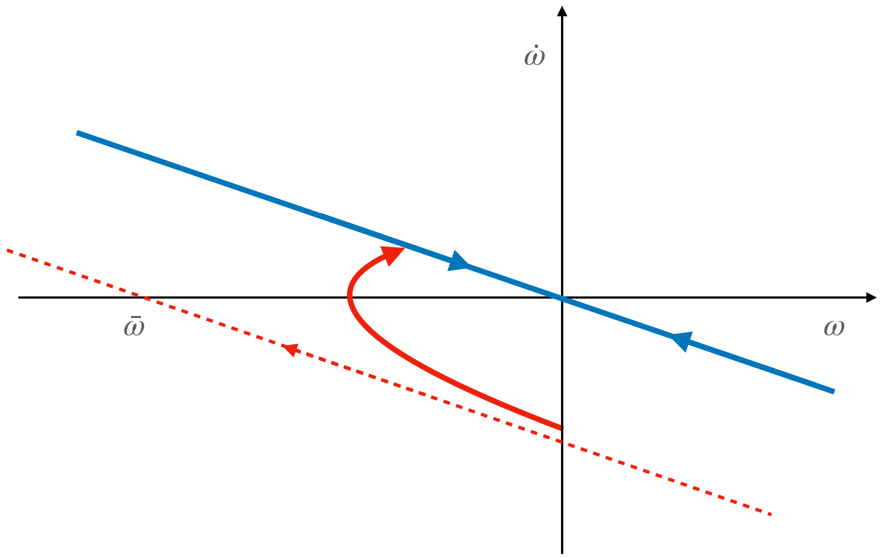

Figure 15: A transitory shock

and initially keeps wage inflation lower than Π^C.[34]

In the long run, the adjustment component goes away, and wage and price inflation converge to the same level, equal to the conflict component.

A similar graphical analysis can be applied to a temporary shock. Consider an economy in steady state with all variables at 0. At $t = 0$, unexpectedly, firms realize that for a finite time interval $[0, T]$ they will face $\overline{mpl} < 0$. At T, mpl goes back to zero. The value of mrs remains at zero throughout. Figure 15 illustrates the phase diagram in this case, with the real wage first moving towards $\overline{\omega}$ and then gradually reverting towards zero.

Proposition 5 in the appendix provides formal derivations for a general class of experiments like the two just analyzed, in which only one side of the labor market is affected, that is, where only mpl or only mrs deviate from zero.

General Result for Unilateral Changes in mrs and mpl

Proposition 5. *Suppose there is no change in $mrs_t = 0$ and the path for mpl_t is negative for all $t \in [0, \infty)$. Then the impact responses at $t = 0$ are*

$$\pi_0 > \pi_0^w > 0.$$

[34] It is easy to prove that despite the presence of the adjustment component, wage inflation is always positive in this experiment. From (11) we get

$$\pi_t^w = \int_t^\infty e^{-\rho(\tau - t)} (mrs_\tau - \omega_\tau) \, d\tau,$$

and notice that $mrs_t = 0$ and $\omega_t < 0$ for all $t > 0$, from the phase diagram.

Suppose there is no change in $mpl_t = 0$ and the path for mrs_t is positive for all $t \in [0, \infty)$. Then the impact responses at $t = 0$ are

$$\pi_0^w > \pi_0 > 0.$$

Proof. Recall the expression for real wages from Proposition 4:

$$\omega_t = e^{r_1 t}\omega_0 + \int_0^\infty H_{s,t}\tilde{\omega}_s ds.$$

If $mpl_t < 0$ and $mrs_t = 0$ for all t, it follows that $\tilde{\omega}_s < 0$ for all s on the right-hand side, so $\omega_t < 0$ for all t. From equation (14), wage inflation at date 0 is then

$$\pi_t^w = -\Lambda_w \int_0^\infty e^{-\rho s}\omega_s ds > 0.$$

Moreover, from equation (25) at $t = 0$, we have

$$\dot{\omega}_0 = \Lambda \int_0^\infty e^{-r_2(\tau - t)}\tilde{\omega}_\tau d\tau < 0,$$

which then implies

$$\pi_0 = \pi_0^w - \dot{\omega}_0 > \pi_0^w > 0.$$

Symmetric derivations prove the other case. □

Proof of Proposition 1

The coefficient ψ in the statement of the proposition is defined as follows

$$\psi = \frac{r_2}{r_2 + \delta}\frac{-r_1}{-r_1 + \rho}.$$

We first derive the real wage path using (26) in the proof of Proposition 4. Solving the integrals gives

$$\begin{aligned}\omega_t =& \Lambda \int_0^t e^{r_1(t-s)} \int_s^\infty e^{-r_2(\tau - s)} e^{-\delta\tau} d\tau ds = \Lambda \frac{1}{r_2 + \delta}\int_0^t e^{r_1(t-s) - \delta s} ds = \\ =& \frac{e^{r_1 t} - e^{-\delta t}}{(r_2 + \delta)(r_1 + \delta)}\left(\Lambda_p mpl_0 + \Lambda_w mrs_0\right).\end{aligned}$$

Write price inflation as

$$\pi_t = \int_t^\infty e^{-\rho(\tau - t)}\left(\omega_\tau - mpl_\tau\right) d\tau,$$

substituting ω_t and integrating gives

$$\pi_t = \frac{1}{r_1 + \delta}\frac{1}{r_2 + \delta}\left(\frac{e^{r_1 t}}{\rho - r_1} - \frac{e^{-\delta t}}{\delta + \rho}\right)\left[\Lambda_p mpl_0 + \Lambda_w mrs_0\right] - \frac{e^{-\delta t}}{\rho + \delta} mpl_0.$$

We then get that $\pi_t > 0$ if and only if

$$\frac{1}{r_1+\delta}\frac{1}{r_2+\delta}\left(\frac{e^{r_1 t}}{\rho-r_1}-\frac{e^{-\delta t}}{\delta+\rho}\right)\left[\Lambda_p mpl_0+\Lambda_w mrs_0\right] > \frac{e^{-\delta t}}{\rho+\delta}mpl_0,$$

which can be rewritten using $-r_1 r_2 = \Lambda_p + \Lambda_w$ (from the proof of Proposition (4)), to get

$$\frac{r_2}{r_2+\delta}\frac{-r_1}{r_1+\delta}\left(\frac{e^{r_1 t}}{\rho-r_1}-\frac{e^{-\delta t}}{\delta+\rho}\right)\frac{\Lambda_p mpl_0+\Lambda_w mrs_0}{\Lambda_p+\Lambda_w} > \frac{e^{-\delta t}}{\rho+\delta}mpl_0.$$

$$\frac{1}{r_2+\delta}\frac{1}{\rho-r_1}\left(\Lambda_p mpl_0+\Lambda_w mrs_0\right) > mpl_0.$$

Setting $t=0$ and rearranging gives the condition for $\pi_0 > 0$ in the statement of the proposition, with

Write wage inflation as

$$\pi_t^w = \int_t^\infty e^{-\rho(\tau-t)}\left(mrs_\tau-\omega_\tau\right)d\tau.$$

Similar steps as those above yield the following condition for $\pi_t^w > 0$

$$\frac{r_2}{r_2+\delta}\frac{-r_1}{r_1+\delta}\left(\frac{e^{r_1 t}}{\rho-r_1}-\frac{e^{-\delta t}}{\delta+\rho}\right)\frac{\Lambda_p mpl_0+\Lambda_w mrs_0}{\Lambda_p+\Lambda_w} < \frac{e^{-\delta t}}{\rho+\delta}mrs_0.$$

The last statement follows because

$$\frac{1-\alpha\psi}{(1-\alpha)\psi} > \frac{\alpha\psi}{1-(1-\alpha)\psi}.$$

Deriving Π^{Spiral}

Recall the definition

$$\Pi^{Spiral} = \int_0^\infty \pi_t dt.$$

Using the decomposition of π_t in the text

$$\int_0^\infty \pi_t = \int_0^\infty \Pi_t^{Conflict}dt + \int_0^\infty \dot{\omega}_t dt$$

and since $\omega_t \to 0$ the second term on the right-hand side is zero. Moreover, since

$$\Pi_t^{Conflict} = \frac{\Lambda_p\Lambda_w}{\Lambda_p+\Lambda_w}\frac{1}{\rho+\delta}\left(mrs_0-mpl_0\right)e^{-\delta t}$$

the expression in the text follows.

A.3 Optimal Policy

Quadratic Approximation of the Welfare Function

The welfare of the representative consumer is

$$\int_0^\infty e^{-\rho t}\left(\frac{1}{1-\sigma}Y_t^{1-\sigma}-\frac{1}{1+\eta}N_t^{1+\eta}\right)dt. \tag{27}$$

We will first derive the expression in parenthesis in terms of relative price distortions in prices and wages.

Labor demand for variety j is

$$L_{jt}=\left(\frac{W_{jt}}{W_t}\right)^{-\varepsilon_L}L_t$$

and imposing market clearing in the labor market we obtain

$$N_t=\int_0^1 L_{jt}dj=\Delta_t^w L_t, \tag{28}$$

where

$$\Delta_t^w\equiv\int_0^1\left(\frac{W_{jt}}{W_t}\right)^{-\varepsilon_L}dj, \tag{29}$$

which is a measure of allocative distortions due to wage dispersion.

Demand for variety i is

$$\frac{Y_{it}}{Y_t}=\left(\frac{P_{it}}{P_t}\right)^{-\varepsilon_C}.$$

Since all goods producers use the same inputs ratio X_{it}/L_{it}, we also have

$$Y_{it}=\left[a_L+a_X\left(\frac{X_t}{L_t}\right)^{1-\frac{1}{\varepsilon}}\right]^{\frac{1}{1-\frac{1}{\varepsilon}}}\frac{L_{it}}{L_t}L_t.$$

Combining these conditions we obtain

$$\left(\frac{P_{it}}{P_t}\right)^{-\varepsilon_C}=\frac{Y_{it}}{Y_t}=\left[a_L+a_X\left(\frac{X_t}{L_t}\right)^{1-\frac{1}{\varepsilon}}\right]^{\frac{1}{1-\frac{1}{\varepsilon}}}\frac{L_{it}}{L_t}\frac{L_t}{Y_t}.$$

Integrating both sides, using $\int L_{it}=L_t$ and rearranging, yields

$$Y_t\Delta_t=\left[a_L L_t^{1-\frac{1}{\varepsilon}}+a_X X_t^{1-\frac{1}{\varepsilon}}\right]^{\frac{1}{1-\frac{1}{\varepsilon}}}, \tag{30}$$

where

$$\Delta_t=\int_0^1\left(\frac{P_{it}}{P_t}\right)^{-\varepsilon_C}di \tag{31}$$

is a measure of allocative distortions due to price dispersion.

Substituting (28) and (30) in the welfare function (27), we can then write it as

$$\int_0^\infty e^{-\rho t}\mathcal{U}\left(l_t, x_t, \delta_t, \delta_t^w\right) dt \tag{32}$$

where

$$\mathcal{U}\left(l, x, \delta, \delta^w\right) \equiv U\left(\frac{F\left(L^{ss}e^l, X^{ss}e^x\right)}{e^\delta}\right) - V\left(e^{\delta^w}L^{ss}e^l\right)$$

and

$$U\left(C\right) \equiv \frac{1}{1-\sigma}C^{1-\sigma},$$
$$F\left(L, X\right) \equiv \left(a_L L^{1-\frac{1}{\varepsilon}} + a_X X^{1-\frac{1}{\varepsilon}}\right)^{\frac{1}{1-\frac{1}{\varepsilon}}},$$
$$V\left(L\right) \equiv \frac{1}{1+\eta}L^{1+\eta}.$$

Consider the second-order approximation (first-order in δ and δ^w as they are already second-order variables), where arguments are omitted for brevity on the right-hand side.

$$\mathcal{U}\left(l, x, \delta, \delta^w\right) \approx \frac{1}{2}\left[\mathcal{U}_{ll}l^2 + 2U_{lx}xl + \frac{1}{2}U_{xx}x^2\right] + \mathcal{U}_\delta\delta + \mathcal{U}_{\delta^w}\delta^w.$$

By definition, the natural level of employment satisfies the first-order condition $\mathcal{U}_l = 0$, and, taking a first order approximation to this first-order condition near the steady state, the linearized natural level of employment $l^*\left(x\right)$ must satisfy, by the implicit function theorem,

$$\mathcal{U}_{ll}l^*\left(x\right) + \mathcal{U}_{lx}x = 0.$$

Therefore, the approximation for $\mathcal{U}$ above can be rewritten as

$$\frac{1}{2}\left[\mathcal{U}_{ll}l^2 + 2U_{ll}l^*l + \frac{1}{2}U_{xx}x^2\right] + \mathcal{U}_\delta\delta + \mathcal{U}_{\delta^w}\delta^w,$$

and ignoring constant terms not controlled by the planner, as

$$\frac{1}{2}\mathcal{U}_{ll}\left(l - l^*\right) + \mathcal{U}_\delta\delta + \mathcal{U}_{\delta^w}\delta^w.$$

To derive the values of the terms $\mathcal{U}_{ll}, \mathcal{U}_\delta, \mathcal{U}_{\delta^w}$, proceed as follows. Differentiating $\mathcal{U}$ gives

$$\mathcal{U}_l = U'\left(\frac{F\left(L, X\right)}{\Delta}\right)\frac{F_L\left(L, X\right)}{\Delta}L - V'\left(L\right)\Delta^w L,$$

$$\mathcal{U}_{ll} = U''\left(\frac{F(L,X)}{\Delta}\right)\left(\frac{F_L(L,X)}{\Delta}L\right)^2 + U'\left(\frac{F(L,X)}{\Delta}\right)\frac{F_{LL}(L,X)}{\Delta}L^2 + $$
$$+ U'\left(\frac{F(L,X)}{\Delta}\right)\frac{F_L(L,X)}{\Delta}L - \Delta^w V''(L)L^2 - \Delta^w V'(L)L$$

$$\mathcal{U}_\delta = -U'\left(\frac{F(L,X)}{\Delta}\right)\frac{F(L,X)}{\Delta^2}\Delta,$$
$$\mathcal{U}_{\delta^w} = -V'(L)\Delta^w L.$$

Computing these expressions at frictionless steady state gives $\mathcal{U}_l = 0$ and

$$\mathcal{U}_{ll} = U''(Y)(F_L L)^2 + U' F_{LL} L^2 - V'' L^2,$$
$$\mathcal{U}_\delta = -U'(Y)Y,$$
$$\mathcal{U}_{\delta^w} = -V'(L)L.$$

Rearranging the expression for $\mathcal{U}_{ll}$

$$\mathcal{U}_{ll} = U' F_L L\left[\frac{U''}{U'}Y\frac{F_L L}{F} + \frac{F_{LL}}{F_L}L - \frac{V''}{V'}L\right] = U' F_L L\left[-\sigma s_L + \frac{F_{LL}}{F_L}L - \eta\right] =$$
$$= -U' F_L L\left[\sigma s_L + \frac{s_X}{\epsilon} + \eta\right].$$

The second order approximation of $\mathcal{U}$ is then

$$-\frac{1}{2}U' Y s_L\left(\sigma s_L + \frac{s_X}{\epsilon} + \eta\right)(l - l^*)^2 - U' Y \delta - V'(L)L\delta^w.$$

In steady state, we have $V'(L) = F_L U'(Y)$, or, multiplying both sides by L and using the definition of s_L,

$$V'(L)L = \frac{F_L L}{Y}U'(Y)Y = s_L U'(Y)Y.$$

Finally, using

$$y = s_L l + s_X x, \quad y^* = s_L l^* + s_X x,$$

we have

$$y - y^* = s_L(l - l^*).$$

Therefore, removing the multiplicative constant $U'(Y)Y$, the second order approximation of the social welfare function takes the form

$$\int_0^\infty e^{-\rho t}\left[-\frac{1}{2}\left(\sigma + \frac{1}{\epsilon}\frac{s_X}{s_L} + \frac{\eta}{s_L}\right)(y_t - y_t^*)^2 - \delta_t - s_L\delta_t^w.\right]$$

Our last step is to express the last two terms, in δ and δ^w, in terms of inflation rates.

Differentiating (31) with respect to time gives

$$\frac{\dot{\Delta}_t^p}{\Delta_t^p} = \varepsilon_C \pi_t + \lambda_p \left[\frac{\left(P_t^* / P_t\right)^{-\varepsilon_C}}{\Delta_t^p} - 1 \right]. \tag{33}$$

The exact relation between P_t^* / P_t and price inflation is

$$\pi_t = \frac{\dot{P}_t}{P_t} = \frac{\lambda_p}{1 - \varepsilon_C} \left(\left(\frac{P_t^*}{P_t} \right)^{1-\varepsilon_C} - 1 \right),$$

which can be rewritten as

$$\frac{P_t^*}{P_t} = \left(1 + \frac{1 - \varepsilon_C}{\lambda_p} \pi_t \right)^{\frac{1}{1-\varepsilon_C}}.$$

Substituting in (33) and using the notation $\delta_t^p = \log \Delta_t^p$ gives

$$\dot{\delta}_t^p = \varepsilon_C \pi_t + \lambda_p \left[e^{-\delta_t^p} \left(1 + \frac{1 - \varepsilon_C}{\lambda_p} \pi_t \right)^{-\frac{\varepsilon_C}{1-\varepsilon_C}} - 1 \right].$$

A second order Taylor approximation of the right-hand side at $\pi_t = 0$ and $\Delta^p = 1$ yields

$$\dot{\delta}_t^p = -\lambda_p \delta_t^p + \frac{1}{2} \frac{\varepsilon_C}{\lambda_p} \pi_t^2,$$

where we approximate to the first order in δ_t and to the second order in π_t. Solving this equation backward in time starting at $\delta_0 = 0$, we obtain

$$\delta_t^p = \frac{1}{2} \int_0^t \frac{\varepsilon_C}{\lambda_p} e^{-\lambda_p (t-s)} \pi_s^2 ds.$$

Computing the present value of distortions gives

$$\int_0^\infty e^{-\rho t} \delta_t^p dt = \frac{1}{2} \frac{\varepsilon_C}{\lambda_p \left(\lambda_p + \rho \right)} \int_0^\infty e^{-\rho t} \pi_t^2 dt.$$

Analogous steps for wage distortions gives

$$\int_0^\infty e^{-\rho t} \delta_t^w dt = \frac{1}{2} \frac{\varepsilon_L}{\lambda_w \left(\lambda_w + \rho \right)} \int_0^\infty e^{-\rho t} \left(\pi_t^w \right)^2 dt.$$

In conclusion, the social welfare function is approximated to the second order by the expression

$$-\frac{1}{2} \int_0^\infty e^{-\rho t} \left[\left(\sigma + \frac{1}{\epsilon} \frac{s_X}{s_L} + \frac{\eta}{s_L} \right) \left(y_t - y_t^* \right)^2 + \frac{\varepsilon_C}{\lambda_p \left(\lambda_p + \rho \right)} \pi_t^2 - \frac{s_L \varepsilon_L}{\lambda_w \left(\lambda_w + \rho \right)} \pi_t^w \right] dt,$$

which corresponds to the quadratic objective in the text with coefficients

$$\begin{aligned}\Phi_y &= \sigma + \frac{1}{\epsilon}\frac{s_X}{s_L} + \frac{\eta}{s_L}, \\ \Phi_p &= \frac{\varepsilon_C}{\lambda_p\left(\lambda_p+\rho\right)}, \\ \Phi_w &= \frac{s_L\varepsilon_L}{\lambda_w\left(\lambda_w+\rho\right)}.\end{aligned}$$

Optimal Policy Problem

Define the coefficients

$$\begin{aligned}\xi_p &= \frac{1}{\epsilon}\frac{s_X}{s_L}, \\ \xi_w &= \frac{\sigma s_L + \eta}{s_L}.\end{aligned}$$

and let hats denote deviations from first best allocations. The optimal policy problem can then be written compactly as follows

$$\min \int_0^\infty e^{-\rho t}\frac{1}{2}\left[\Phi_y \hat{y}_t^2 + \Phi_p \pi_t^2 + \Phi_w\left(\pi_t^w\right)^2\right] dt$$

subject to

$$\rho\pi_t = \Lambda_p\left(\hat{\omega}_t + \xi_p\hat{y}_t\right) + \dot{\pi}_t, \tag{34}$$
$$\rho\pi_t^w = \Lambda_w\left(\xi_w\hat{y}_t - \hat{\omega}_t\right) + \dot{\pi}_t^w, \tag{35}$$
$$\dot{\hat{\omega}}_t = \pi_t^w - \pi_t - \dot{\omega}_t^*, \tag{36}$$

taking ω_0 as given. Form the Hamiltonian

$$\begin{aligned}-\frac{1}{2}\left[\Phi_y \hat{y}_t^2 + \Phi_p \pi_t^2 + \Phi_w\left(\pi_t^w\right)^2\right] + \\ -\Lambda_p\left(\hat{\omega}_t + \xi_p\hat{y}_t\right)\nu_t + \pi_t\dot{\nu}_t + \\ -\Lambda_w\left(\xi_w\hat{y}_t - \hat{\omega}_t\right)\nu_t^w + \pi_t^w\dot{\nu}_t^w + \\ \left(\pi_t^w - \pi_t - \dot{\omega}_t^*\right)\mu_t - \rho\omega_t\mu_t + \omega_t\dot{\mu}_t,\end{aligned}$$

where ν_t, ν_t^w, μ_t denote the Lagrange multipliers on the three constraints, and derive the following first-order conditions for $\hat{y}_t, \hat{\omega}_t, \pi_t, \pi_t^w$:

$$-\Phi_y\hat{y}_t - \Lambda_p\xi_p\nu_t - \Lambda_w\xi_w\nu_t^w = 0, \tag{37}$$
$$-\Lambda_p\nu_t + \Lambda_w\nu_t^w - \rho\mu_t + \dot{\mu}_t = 0, \tag{38}$$
$$-\Phi_p\pi_t + \dot{\nu}_t - \mu_t = 0, \tag{39}$$
$$-\Phi_w\pi_t^w + \dot{\nu}_t^w + \mu_t = 0. \tag{40}$$

Given that the initial inflation rates π_0 and π_0^w are free variables we have

$$\nu_0 = \nu_0^w = 0,$$

which, using (37) and (38), implies

$$\hat{y}_0 = 0,$$

and

$$\dot{\mu}_0 = \rho\mu_0. \tag{41}$$

We will use these two as initial conditions for $\hat{y}_0$ and μ_0 and derive the paths of these two variables from the following two ODEs, that come from differentiating (37) and (38) with respect to time:

$$\begin{aligned} \Phi_y \dot{\hat{y}}_t + \Lambda_p \xi_p \dot{\nu}_t + \Lambda_w \xi_w \dot{\nu}_t^w &= 0, \\ -\Lambda_p \dot{\nu}_t + \Lambda_w \dot{\nu}_t^w - \rho\dot{\mu}_t + \ddot{\mu}_t &= 0. \end{aligned}$$

Using (39) and (40) to substitute for $\dot{\nu}_t$ and $\dot{\nu}_t^w$, the ODEs above become:

$$\begin{aligned} \Phi_y \dot{\hat{y}}_t + \Lambda_p \xi_p \left(\Phi_p \pi_t + \mu_t\right) + \Lambda_w \xi_w \left(\Phi_w \pi_t^w - \mu_t\right) &= 0, \\ \ddot{\mu}_t - \rho\dot{\mu}_t - \left(\Lambda_p + \Lambda_w\right)\mu_t - \Lambda_p \Phi_p \pi_t + \Lambda_w \Phi_w \pi_t^w &= 0. \end{aligned} \tag{42}$$

A useful observation here is that the ODE for the Lagrange multiplier μ_t is a second order ODE with exactly the same structure as the ODE for the real wage, analyzed in Proposition 4. Therefore, by analogy with (26), the solution can be written as follows

$$\mu_t = e^{r_1 t}\mu_0 + \int_0^\infty H_{s,t}\left(\Lambda_w \Phi_w \pi_t^w - \Lambda_p \Phi_p \pi_t\right) ds. \tag{43}$$

To derive μ_0, we use the analog of equation 25 in Proposition 4, evaluated at time $t = 0$

$$\dot{\mu}_0 = r_1 \mu_0 + \int_0^\infty e^{-r_2 t}\left(\Lambda_w \Phi_w \pi_t^w - \Lambda_p \Phi_p \pi_t\right) dt$$

and the initial condition (41), to obtain

$$\mu_0 = \frac{1}{\rho - r_1}\int_0^\infty e^{-r_2 t}\left(\Lambda_w \Phi_w \pi_t^w - \Lambda_p \Phi_p \pi_t\right) dt. \tag{44}$$

In summary, an optimal policy is found finding a pair of paths $\{\mu_t, \hat{y}_t, \pi_t, \pi_t^w, \omega_t\}_{t=0}^\infty$ that satisfy $\hat{y}_0 = 0$, the optimality conditions (42), (43), and (44) and the equilibrium conditions (34)-(35).

Algorithm

This algorithm solves for optimal policy exploiting a simple finite difference method to express all differential and integral equations as linear equations.

Choose a vector $\mathbf{t} = (t_1, t_2, ..., t_K)$ of K equi-spaced dates in the interval $[0, T]$ for some horizon T. Notice that

$$\Lambda_p mpl_t + \Lambda_w mrs_t = \left(\Lambda_p + \Lambda_w\right) \omega_t^* + \left(\Lambda_w \xi_w - \Lambda_p \xi_p\right) \hat{y}_t.$$

Real wages are then obtained from the following matrix version of (26):

$$\boldsymbol{\omega} = \omega_0 e^{r_1 \mathbf{t}} + \mathbf{H}\left(\left(\Lambda_p + \Lambda_w\right) \boldsymbol{\omega}^* + \left(\Lambda_w \xi_w - \Lambda_p \xi_p\right) \hat{\mathbf{y}}\right),$$

where boldface variables represent vectors of the corresponding variable at the times $\mathbf{t}$ and $\mathbf{H}$ is a matrix with elements $H_{t_i, t_j} \Delta t$ (using the expression for $H_{s,t}$ in Proposition 4). The inflation equations can also be written in matrix form

$$\boldsymbol{\pi} = \Lambda_p \mathbf{A}\left(\boldsymbol{\omega} - \boldsymbol{\omega}^* + \xi_p \hat{\mathbf{y}}\right),$$
$$\boldsymbol{\pi}^w = \Lambda_w \mathbf{A}\left[\xi_w \hat{\mathbf{y}} - \left(\boldsymbol{\omega} - \boldsymbol{\omega}^*\right)\right],$$

where the matrix $\mathbf{A}$ has elements A_{t_i, t_j} with

$$A_{t,s} = e^{-\rho(s-t)} \Delta t \text{ if } s \geq t,$$

and

$$A_{t,s} = 0 \text{ if } s < t.$$

Substituting the solution for $\boldsymbol{\omega}$ in the inflation equations above gives, after some rearranging

$$\boldsymbol{\pi} = \omega_0 \Lambda_p \mathbf{A} e^{r_1 \mathbf{t}} + \Lambda_p \mathbf{A}\left[\xi_p \mathbf{I} - \left(\Lambda_p \xi_p - \Lambda_w \xi_w\right) \mathbf{H}\right] \hat{\mathbf{y}} + \Lambda_p \mathbf{A}\left[\left(\Lambda_p + \Lambda_w\right) \mathbf{H} - \mathbf{I}\right] \boldsymbol{\omega}^* \quad (45)$$

$$\boldsymbol{\pi}^w = -\omega_0 \Lambda_w \mathbf{A} e^{r_1 \mathbf{t}} + \Lambda_w \mathbf{A}\left[\xi_w \mathbf{I} + \left(\Lambda_p \xi_p - \Lambda_w \xi_w\right) \mathbf{H}\right] \hat{\mathbf{y}} - \Lambda_w \mathbf{A}\left[\left(\Lambda_p + \Lambda_w\right) \mathbf{H} - \mathbf{I}\right] \boldsymbol{\omega}^* \quad (46)$$

Translating in matrix forms condition (43) we obtain

$$\boldsymbol{\mu} = \mu_0 e^{r_1 \mathbf{t}} + \mathbf{H}\left(\Lambda_w \Phi_w \boldsymbol{\pi}^w - \Lambda_p \Phi_p \boldsymbol{\pi}\right)$$

and substituting for μ_0 using the matrix version of (44), we can write

$$\boldsymbol{\mu} = \mathbf{M}_w \boldsymbol{\pi}^w - \mathbf{M}_p \boldsymbol{\pi} \quad (47)$$

where

$$\mathbf{M}_p = \Lambda_p \Phi_p \mathbf{M},$$
$$\mathbf{M}_w = \Lambda_w \Phi_w \mathbf{M},$$

and

$$\mathbf{M} = \frac{1}{\rho - r_1} \begin{bmatrix} e^{r_1 t_1} \\ e^{r_1 t_2} \\ e^{r_1 t_3} \\ \dots \\ e^{r_1 t_K} \end{bmatrix} \begin{bmatrix} e^{-r_2 t_1} & e^{-r_2 t_2} & e^{-r_2 t_3} & \dots & e^{-r_2 t_K} \end{bmatrix} \Delta t + \mathbf{H}.$$

Finally the ODE for $\hat{y}_t$ in equation (42) can be written in matrix form as

$$\Phi_y \mathbf{D} \cdot \hat{\mathbf{y}} + \Lambda_p \xi_p \left(\Phi_p \boldsymbol{\pi} + \boldsymbol{\mu} \right) + \Lambda_w \xi_w \left(\Phi_w \boldsymbol{\pi}^w - \boldsymbol{\mu} \right) = 0$$

where

$$\mathbf{D} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ & \dots & & \dots \\ & & & \end{bmatrix} \frac{1}{\Delta t},$$

(the form of the first row imposes the condition $y_0 = 0$). Substituting for $\boldsymbol{\mu}$, from (47), we can rewrite the ODE for $\hat{y}_t$ as:

$$\Phi_y \mathbf{D} \cdot \hat{\mathbf{y}} + \mathbf{G}_p \boldsymbol{\pi} + \mathbf{G}_w \boldsymbol{\pi}^w = 0,$$

where

$$\begin{aligned} \mathbf{G}_p &= \Lambda_p \xi_p \Phi_p \mathbf{I} - \left(\Lambda_p \xi_p - \Lambda_w \xi_w \right) \mathbf{M}_p, \\ \mathbf{G}_w &= \Lambda_w \xi_w \Phi_w \mathbf{I} + \left(\Lambda_p \xi_p - \Lambda_w \xi_w \right) \mathbf{M}_w. \end{aligned}$$

Substituting for $\boldsymbol{\pi}$ and $\boldsymbol{\pi}^w$, from (45)-(46), we obtain a linear equation in $\hat{\mathbf{y}}$ and exogenous variables, which can be written as

$$\mathbf{C} \hat{\mathbf{y}} = \mathbf{b},$$

where

$$\mathbf{C} = \Phi_y \mathbf{D} + \mathbf{G}_p \Lambda_p \mathbf{A} \left[\xi_p \mathbf{I} - \left(\Lambda_p \xi_p - \Lambda_w \xi_w \right) \mathbf{H} \right] + \mathbf{G}_w \Lambda_w \mathbf{A} \left[\xi_w \mathbf{I} + \left(\Lambda_p \xi_p - \Lambda_w \xi_w \right) \mathbf{H} \right],$$

and

$$\mathbf{b} = \left(\Lambda_w \mathbf{G}_w - \Lambda_p \mathbf{G}_p \right) \left\{ \omega_0 \mathbf{A} e^{r_1 \mathbf{t}} + \mathbf{A} \left[\left(\Lambda_p + \Lambda_w \right) \mathbf{H} - \mathbf{I} \right] \boldsymbol{\omega}^* \right\}.$$