On the Welfare and Cyclical Implications of Moderate Trend Inflation☆

Guido Ascari[a], Louis Phaneuf[b], Eric R. Sims[c]
[a] University of Oxford and Pavia
[b] University of Quebec at Montreal [c] University of Notre Dame and NBER

Received March 24, 2016; Received in Revised Form Date May 19, 2018; Accepted Date

Abstract

The welfare and cyclical implications of moderate trend inflation are addressed in an augmented medium-scale DSGE model. Increasing trend inflation from 2 to 4 percent generates a consumption-equivalent welfare loss of about 4 percent. Welfare costs of this magnitude are driven by: staggered wage contracts, trend growth, extended borrowing, a roundabout production structure, and the interaction between trend inflation and shocks to the marginal efficiency of investment (MEI). A sticky-price model abstracting from these features generates much smaller losses. Moderate trend inflation also has important business-cycle implications, interacting much more strongly with MEI shocks than with productivity or monetary shocks.

JEL Codes: E31, E32.
Keywords: trend inflation; roundabout production; investment shocks; inflation costs; business cycles.

1. Introduction

In the aftermath of the Great Recession, a number of economists have argued that the Federal Reserve and other central banks should raise their inflation targets. At the time of the recession, the consensus was that the inflation target was about 2 percent annually. Economists like Blanchard et al. (2010), Ball (2013), and Krugman (2014) have since advocated for increases in the inflation target to 4 or even 5 percent. Implementing such proposals over a sufficiently long period of time would eventually lead to higher long-run or trend inflation. Proposals to raise the inflation target therefore naturally lead to the following question: how would the U.S. economy be affected by a moderate rise in trend inflation? Our paper offers a new perspective on this policy question. While doing so, it also provides new insights on the effects of moderate trend inflation on the welfare and business-cycle properties of medium-scale New Keynesian models more generally.

There exists a large literature that studies the macroeconomic consequences of non-zero trend inflation. The majority of models used for this purpose are small-scale sticky-price models (Ascari 2004; Hornstein and Wolman 2005; Kiley 2007; Levin and Yun 2007; Ascari and Ropele 2007; Coibion and Gorodnichenko 2011). By "small-scale" we mean that these models abstract from capital accumulation and most forms of real rigidity. By "sticky-price" we mean that these models typically assume flexible nominal wages. Two partial exceptions are Amano et al. (2007) and Amano et al. (2009). The former features a model with capital and convex capital adjustment costs, but abstracts from wage rigidity. The latter considers price and wage rigidity together, but omits capital and real rigidities.

☆We acknowledge Yuriy Gorodnichenko and Juan Rubio-Ramirez for helpful comments and suggestions at an early stage of this project, participants at the Royal Economic Society Conference 2016, at the "Rethinking Inflation Targeting" Conference at the Norges Bank 9-10 September 2016 and Johannes Wieland for the insightful discussion, Sylvain Leduc and an anonymous referee for useful comments on the current draft, and Jean-Gardy Victor for capable research assistance.

*Corresponding Author: Louis Phaneuf, P.O. Box 888, Downtown Station, Montreal, H3C 3P8, Canada, email: phaneuf.louis@uqam.ca, phone: 514-987-3000.

In this existing literature, a trend inflation rate of less than 4 percent generally has a modest impact on the properties of the standard New Keynesian model. For example, Ascari (2004) finds that an increase in trend inflation from 2 to 4 percent generates an additional steady-state output loss of about 0.5 percent. Amano et al. (2009) find that raising trend inflation from 2 to 4 percent generates an additional consumption-equivalent welfare loss of less than one percent. Using alternative versions of a sticky-price model, Amano et al. (2007), Ascari and Ropele (2007), and Ascari and Sbordone (2014) show that a trend inflation rate of less than 4 percent has a modest impact on business-cycle fluctuations.

We contribute to this literature by focusing on an augmented medium-scale model similar to Christiano et al. (2005) and Smets and Wouters (2007). Our paper is a first attempt to look at how moderate trend inflation affects the welfare and business cycle properties of these medium-scale models. The model we develop in this paper rests on several key features, some of which are relatively common ingredients and others which are less so. In addition to the usual nominal rigidities and real adjustment frictions, our model allows for non-zero trend inflation and trend growth in neutral and investment-specific productivity. An additional feature of our model is that it incorporates a roundabout production structure wherein firms use the outputs of other firms as a factor of production. We also consider an extended working capital channel wherein firms borrow funds from a financial intermediary to finance the costs of all of their variable inputs and not solely the wage bill.

The first part of the paper provides a quantitative evaluation of the welfare costs of positive trend inflation. When conducting our welfare analysis, we focus on a change in trend inflation from 2 to 4 percent, a scenario consistent with recent proposals to raise the inflation target. Our quantitative evaluation of the welfare costs of trend inflation rests on two measures: a consumption-equivalent welfare loss metric based on non-stochastic steady states and another metric calculated from stochastic means.

In our baseline model, the cost of raising trend inflation from 2 to 4 percent is 3.7 percent of each period's consumption based on non-stochastic steady states and 4.3 percent based on stochastic means. These losses are substantially higher than what much of the existing literature has found. It is therefore natural to ask what are the features of our model which are responsible for welfare costs of this magnitude. Our analysis points to five important factors: *i*) staggered wage contracts, *ii*) trend growth, *iii*) roundabout production, *iv*) extended working capital and *v*) the interaction between trend inflation and a persistent MEI shock. When our model is re-calibrated to abstract from these features, we find welfare costs of increasing the trend inflation rate from 2 to 4 percent of less than 0.25 percent of consumption, which is in line with the findings in much of the existing literature.

Staggered wage contracts are an important factor determining the welfare costs of trend inflation, significantly more than staggered price contracts. In particular, if we assume that wages are flexible the welfare cost of raising trend inflation from 2 to 4 percent would only be 1.0 percent of consumption based on stochastic means. Households with positive trend inflation would like to reset their wages each period, but only a fraction can. This leads to significant steady-state wage dispersion, which drives a wedge between aggregate labor supply and demand. It also results in higher wage markups on average, as updating households choose higher wages than they otherwise would to protect their future real wages from inflation. This higher average wage markup moves the economy further from the first best allocation, resulting in significant welfare losses. While qualitatively similar effects of trend inflation are also at work for price-setting, for plausible labor supply elasticities, the effect of trend inflation is substantially more important with wage-setting.

Trend growth in investment-specific and neutral technology, in conjunction with wage rigidity, also contributes significantly to the welfare costs of trend inflation. If there is no trend growth, then the welfare cost of moving from 2 to 4 percent inflation is significantly lower relative to our baseline. Positive trend growth means that households would like to adjust their wages each period even if trend inflation is zero. This results in steady state wage dispersion and higher than average wage markups than if trend growth were zero. Adding in positive trend inflation exacerbates these distortions, resulting in much larger welfare costs than if there were no trend growth.

Roundabout production plays a non negligible role in accounting for the welfare costs of positive trend inflation. The steady-state and mean consumption equivalent welfare costs of going from 2 to 4 percent trend inflation amount to 3.2 and 3.7 percent without roundabout production instead of 3.7 and 4.3 percent with it. Roundabout production has two effects in the model: it acts as an amplification source for real shocks and also is isomorphic to prices being stickier, because it introduces strategic complementarity into price-setting (Basu 1995 and Huang et al. 2004). Both of these features make trend inflation relatively more

costly.

The extended working capital channel also contributes significantly to our findings on welfare costs. Without extended borrowing the steady-state and mean welfare costs decline to 3.3 and 3.7 percent, respectively. Intuitively, working capital must raise the costs of trend inflation, because higher trend inflation raises the average nominal interest rate, which effectively represents a direct distortion on the first-order conditions for optimal inputs.

Finally, there are also potentially interesting interactions between the various shock sources and the consumption equivalent welfare losses based on stochastic means. When our model is re-calibrated so as to exclude MEI shocks, the mean welfare cost of going from 2 to 4 percent trend inflation falls from 4.3 percent of consumption to 3.9 percent.

Our second set of findings pertains to the cyclical implications of positive trend inflation. Whereas trend inflation has relatively minor effects on the dynamic responses of aggregate variables to productivity and monetary shocks, there are large interactions between trend inflation and MEI shocks. Contrary to other types of shocks, the interaction between trend inflation and the cyclical responses to MEI shocks depends heavily on the persistence of the shock. For moderate levels of shock persistence, output reacts more strongly to a MEI shock the larger is trend inflation. In particular, under our baseline parameterization the impulse response of output at a ten quarter horizon is about 15 percent larger with 4 percent trend inflation compared to 2 percent trend inflation. Interestingly, the interaction between trend inflation and the MEI shock flips signs at higher levels of persistence. When the shock is sufficiently persistent, higher trend inflation significantly dampens the response of output and other aggregate variables to MEI shocks.

The results of our paper have important implications for both policymakers and academics. On the policy front, the large welfare costs of trend inflation which we find represent a warning against policy proposals urging central banks to raise their inflation targets. In that respect, the message of our paper is complementary to Coibion et al. (2012), who weigh the benefits of a reduced incidence of zero lower bound episodes from higher trend inflation against the costs of higher trend inflation outside of periods where the zero lower bound binds. On the academic front, ours is the first paper to point out the large interaction between trend inflation and MEI shocks. An increasing body of research suggests that MEI shocks are a major driver of the business cycle – Justiniano et al. (2010, 2011) find that these shocks account for 50 percent or more of volatility in output. Conventional wisdom in the literature has been that trend inflation might matter in a normative sense, but that it is innocuous to ignore it for the purposes of understanding positive aspects of the business cycle. Our results suggest that this is not the case – trend inflation interacts strongly with MEI shocks.

The remainder of the paper is organized as follows. Section 2 lays out our medium-scale DSGE model. Section 3 discusses some issues related to calibration. Section 4 examines the steady-state and mean welfare implications of moderate trend inflation. Section 5 studies the cyclical implications of trend inflation. Section 6 contains concluding remarks.

2. A Medium-Scale Macro Model with Trend Inflation

This section lays out our medium-scale DSGE model. As other recent New Keynesian models do, it embeds nominal rigidities in the form of Calvo (1983) wage and price contracts, habit formation in consumption, investment adjustment costs, variable capital utilization, and monetary policy governed by a Taylor rule.

Our model differs from most medium-scale models along the following dimensions. In addition to allowing for non-zero steady state inflation, we also incorporate trend output growth stemming from both neutral technology growth and growth in investment-specific technology (IST). The model also incorporates a roundabout production structure, a feature which Christiano (2016) refers to as "firms networking" after Acemoglu et al. (2015). Empirical evidence in support of roundabout production is discussed in Basu (1995), Huang et al. (2004), and Nakamura and Steinsson (2010). Another unique feature of the model is an extended working capital channel. Working capital has been a key feature of several macro models (Fuerst 1992; Christiano et al. 1997, 2005; Barth and Ramey 2002). We follow the approach in Phaneuf et al. (2018), who assume that firms need working capital in advance of production to cover the costs of all of their variable inputs and not only the wage bill.

Our baseline model abstracts from indexation of non-reoptimized prices and wages to lagged inflation. Backward-looking price or wage indexation is often included in models to help deliver hump-shaped, inertial impulse responses of inflation and marginal costs to monetary policy shocks. This assumption, however, has been often criticized in the literature as *ad hoc*. Moreover, it implies that every price and wage in the economy is changed in every period, which is not what we observe in micro data. The assumption of no price and, especially, no wage indexation are important for our results, so we will discuss it more below.

The subsections below lay out the key features of our model. Online Appendix A presents the full set of equilibrium conditions.

2.1. Good and Labor Composites

There are a continuum of firms, indexed by $j \in [0,1]$, who produce differentiated goods with the use of a composite labor input. The composite labor input is aggregated from differentiated labor supplied by a continuum of households, indexed by $i \in [0,1]$. The differentiated goods are bundled into a gross output good, X_t, by a competitive final good firm. Some of this gross output good is used as a factor of production by firms. Net output is gross output less intermediates. Households can either consume or invest the final net output good. Differentiated labor input is bundled into a final labor input by a competitive labor packer. The composite gross output and labor input respectively are:

$$X_t = \left(\int_0^1 X_t(j)^{\frac{\theta-1}{\theta}} dj \right)^{\frac{\theta}{\theta-1}}, \tag{1}$$

$$L_t = \left(\int_0^1 L_t(i)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}. \tag{2}$$

The parameters $\theta > 1$ and $\sigma > 1$ are the elasticities of substitution between goods and labor. Profit maximization yields conventional downward-sloping demand curves for varieties of intermediates and labor as well as aggregate price and wage indexes.

2.2. Households

Households supply differentiated labor input, facing a downward-sloping demand for their variety of labor. Each period there is a fixed probability, $(1-\xi_w)$, that households can adjust their nominal wage. As in Erceg et al. (2000), we assume that utility is separable in consumption and labor. State-contingent securities insure households against idiosyncratic wage risk arising from staggered wage-setting. With this setup, households will be identical along all dimensions other than labor supply and wages.

The problem of a typical household, omitting dependence on i except for these two dimensions, is:

$$\max_{C_t, L_t(i), K_{t+1}, B_{t+1}, I_t, Z_t} \quad E_0 \sum_{t=0}^{\infty} \beta^t \left(\ln\left(C_t - bC_{t-1}\right) - \eta \frac{L_t(i)^{1+\chi}}{1+\chi} \right), \tag{3}$$

subject to the following budget constraint,

$$P_t \left[C_t + I_t + \left(\gamma_1 (Z_t - 1) + \frac{\gamma_2}{2} (Z_t - 1)^2 \right) \frac{K_t}{\varepsilon_t^{I,\tau}} \right] + \frac{B_{t+1}}{1+i_t} \leq W_t(i) L_t(i) + R_t^k Z_t K_t + \Pi_t + B_t, \tag{4}$$

the physical capital accumulation process,

$$K_{t+1} = \vartheta_t \varepsilon_t^I \left(1 - \frac{\kappa}{2} \left(\frac{I_t}{I_{t-1}} - g_I \right)^2 \right) I_t + (1-\delta) K_t, \tag{5}$$

and a downward-sloping demand for labor:

$$L_t(i) = \left(\frac{W_t(i)}{W_t} \right)^{-\sigma} L_t. \tag{6}$$

P_t is the nominal price of goods, C_t is consumption, I_t is investment measured in units of consumption, K_t is the physical capital stock, and Z_t is the level of capital utilization.[1] $W_t(i)$ is the nominal wage paid to labor of type i, and R_t^k is the common rental price on capital services (the product of utilization and physical capital). Π_t denotes distributed dividends from firms. B_t is a stock of nominal bonds with which a household enters a period. γ_1 and γ_2 are parameters relating to a resource cost of capital utilization. This cost is measured in units of physical capital. κ is a parameter governing a cost to adjusting investment growth relative to trend growth, $g_I \geq 1$. i_t is the nominal interest rate. $0 < \beta < 1$ is a discount factor, $0 < \delta < 1$ is a depreciation rate, and $0 \leq b < 1$ is a parameter for internal habit formation. χ is the inverse Frisch labor supply elasticity.

$\varepsilon_t^{I,\tau}$ measures the level of investment specific technology (IST). We assume that it follows a deterministic trend with no stochastic component. The deterministic trend is necessary to match the observed downward trend in the relative price of investment goods in the data. The exogenous variable ϑ_t is a stochastic marginal efficiency of investment (MEI) shock. Justiniano et al. (2011) distinguish between these two types of investment shocks, showing that IST shocks map one-to-one into the relative price of investment goods, while MEI shocks do not impact the relative price of investment. They find that MEI shocks are critical for business cycles, while stochastic shocks to IST are not. These findings form the basis for our modeling the MEI component as stochastic while the IST term as deterministic.

A household given the opportunity to adjust its wage in period t will choose a "reset wage," w_t^*. It is straightforward to show that under our assumptions all updating households adjust to the same reset wage.

2.3. Firms

The production function for a typical producer j is:

$$X_t(j) = A_t \Gamma_t(j)^{\phi} \left(\widehat{K}_t(j)^{\alpha} L_t(j)^{1-\alpha} \right)^{1-\phi} - \Upsilon_t F, \tag{7}$$

where F is a fixed cost, and production is required to be non-negative. Υ_t is a growth factor, to be discussed later. Given Υ_t, F is chosen to keep profits zero along a balanced growth path, so the entry and exit of firms can be ignored. $\Gamma_t(j)$ is the amount of intermediate input, and $\phi \in (0,1)$ is the intermediate input share. Intermediate inputs come from aggregate gross output, X_t. $\widehat{K}_t(j)$ is capital services (the product of utilization and physical capital), while $L_t(j)$ is labor input.

The firm gets to choose its price, $P_t(j)$, as well as quantities of intermediates, capital services, and labor input. Each period there is a $(1-\xi_p)$ probability that a firm can re-optimize its price. Regardless of whether a firm is given the opportunity to adjust its price, it will choose inputs to minimize total cost, subject to the constraint of producing enough to meet demand. The cost minimization problem of a typical firm is:

$$\min_{\Gamma_t, \widehat{K}_t, L_t} \quad (1 - \psi_\Gamma + \psi_\Gamma(1+i_t)) P_t \Gamma_t + (1 - \psi_K + \psi_K(1+i_t)) R_t^k \widehat{K}_t + (1 - \psi_L + \psi_L(1+i_t)) W_t L_t \tag{8}$$

s.t.

$$A_t \Gamma_t^{\phi} \left(\widehat{K}_t^{\alpha} L_t^{1-\alpha} \right)^{1-\phi} - \Upsilon_t F \geq \left(\frac{P_t(j)}{P_t} \right)^{-\theta} X_t.$$

Here ψ_l, $l = \Gamma, K, L$, is the fraction of payments to a factor that must be financed at the gross nominal interest rate, $1 + i_t$. Assuming $\psi_l = 1$ for all l means that all factor payments are financed through working capital, so that the factor prices relevant for firms are the product of the gross nominal interest rate and the factor price. We refer to this case as extended borrowing (EB). With $\psi_l = 0$ for all l, firms do not have to borrow to pay any of their factors. To economize on notation, we define $\Psi_l = (1 - \psi_l + \psi_l(1+i_t))$ for $l = \Gamma, K, L$.

Assume that $\psi_l = 1$ for $l = \Gamma, K, L$. Applying some algebraic manipulations to the first order conditions for cost-minimization yields the following expression for real marginal cost, v_t, which is common across firms:

[1] The relative price of investment goods to consumption goods is $1/\varepsilon_t^I$. Hence, if $\widehat{I}_t$ is physical units of investment, then $\varepsilon_t^I I_t = \widehat{I}_t$. Writing the accumulation equation in terms of investment measured in consumption units yields (5).

$$v_t = (1+i_t)\left(\frac{1}{1-\phi}\right)^{1-\phi}\left(\frac{1}{\phi}\right)^{\phi}\tilde{v}_t^{1-\phi}, \tag{9}$$

where $\tilde{v}_t$ is the standard real marginal cost given a Cobb-Douglas production function without roundabout production ($\phi = 0$) and extended borrowing:

$$\tilde{v}_t = \left(\frac{1}{1-\alpha}\right)^{1-\alpha}\left(\frac{1}{\alpha}\right)^{\alpha}(r_t^k)^{\alpha}(w_t)^{1-\alpha}A_t^{-1}. \tag{10}$$

Relative to the basic case in the literature, roundabout production reduces the sensitivity of real marginal cost to factor prices by a factor of $1-\phi$. Second, the nominal interest rate is a direct component of real marginal cost.

When given the opportunity to adjust its price, a firm will maximize the expected discounted value of profits discounted by the stochastic discount factor of households. It is straightforward to show that all updating firms will update to a common reset price.

Given marginal cost, v_t, the optimality conditions related to price-setting are standard.

2.4. Monetary Policy

Monetary policy follows a Taylor rule:

$$\frac{1+i_t}{1+i} = \left(\frac{1+i_{t-1}}{1+i}\right)^{\rho_i}\left[\left(\frac{\pi_t}{\pi}\right)^{\alpha_\pi}\left(\frac{Y_t}{Y_{t-1}}g_Y^{-1}\right)^{\alpha_y}\right]^{1-\rho_i}\varepsilon_t^r. \tag{11}$$

The interest rate responds to deviations of inflation from trend, π, and to deviations of output growth from its trend, g_Y. ε_t^r is an exogenous policy shock. The parameter ρ_i governs the smoothing effect on nominal interest rates while α_π and α_y are control parameters.

2.5. Shock Processes

Neutral productivity obeys a process with both a trending and stationary component. A_t^τ is the deterministic trend component, where g_A is the gross growth rate:

$$A_t = A_t^\tau \widetilde{A}_t, \qquad A_t^\tau = g_A A_{t-1}^\tau. \tag{12}$$

The initial level in period 0 is normalized to unity. The stationary component of neutral productivity follows an AR(1) process, with innovation, u_t^A, drawn from a mean zero normal distribution with known standard deviation equal to s_A:

$$\widetilde{A}_t = \left(\widetilde{A}_{t-1}\right)^{\rho_A}\exp\left(s_A u_t^A\right), \quad 0 \le \rho_A < 1, \tag{13}$$

The IST term follows a deterministic trend, where g_{ε^I} is the gross growth rate and the initial level in period 0 is normalized to unity:

$$\varepsilon_t^{I,\tau} = g_{\varepsilon^I}\varepsilon_{t-1}^{I,\tau} \tag{14}$$

The MEI shock follows a stationary AR(1) process, with innovation drawn from a mean zero normal distribution with standard deviation s_I:

$$\vartheta_t = (\vartheta_{t-1})^{\rho_I}\exp(s_I u_t^I), \quad 0 \le \rho_I < 1 \tag{15}$$

The only remaining shock in the model is the monetary policy shock, ε_t^r. We assume that it is drawn from a mean zero normal distribution with known standard deviation s_r.

2.6. Growth

Most variables in the model inherit trend growth from the deterministic trends in neutral and investment-specific productivity. Let this trend factor be Υ_t. Output, consumption, investment, intermediate inputs, and the real wage will all grow at the rate of this trend factor on a balanced growth path. The capital stock will grow faster due to growth in investment-specific productivity, with $\widetilde{K}_t \equiv \frac{K_t}{\Upsilon_t \varepsilon_t^{I,\tau}}$ being stationary.

Given our specification of preferences, labor hours will be stationary. The full set of equilibrium conditions re-written in stationary terms can be found in Online Appendix A.

The trend factor that induces stationarity among transformed variables is:

$$\Upsilon_t = (A_t^{\tau})^{\frac{1}{(1-\phi)(1-\alpha)}} \left(\varepsilon_t^{I,\tau}\right)^{\frac{\alpha}{1-\alpha}}. \tag{16}$$

This reverts to the conventional trend growth factor when $\phi = 0$. A higher value of ϕ amplifies the effects of trend growth in neutral productivity on output and its components. For a given level of trend growth in neutral productivity, the economy will grow faster the larger is the share of intermediates in production.

3. Parameterization and Model Fit

In this section we describe the parameterization of the model and speak to the model's empirical fit by focusing on some empirical moments. Parameter values are summarized in Table 1. Many of the model's parameters are set to conventional values within the literature. $\beta = 0.99$ is the discount factor, $b = 0.8$ is the habit formation parameter, $\chi = 1$ is the inverse Frisch elasticity, and $\eta = 6$ is the weight on disutility of labor set so that steady-state labor hours are around 1/3. θ and σ are the elasticities for goods and labor which are both set at 6 following Rotemberg and Woodford (1997) and Liu and Phaneuf (2007). The Calvo price and wage probabilities, ξ_p and ξ_w, are set at 2/3. This implies a median duration of prices of about 5 months, which is consistent with the evidence in Bils and Klenow (2004). Setting $\xi_w = 2/3$ is rather conservative, broadly consistent with the estimate in Christiano et al. (2005), but somewhat lower than the micro evidence offered by Barattieri et al. (2014) and also lower than what is estimated in Justiniano et al. (2010, 2011).

We set the share of capital services parameter to $\alpha = 1/3$. In Nakamura and Steinsson (2010), the weighted average revenue share of intermediate inputs in the U.S. private sector using Consumer Price Index (CPI) expenditure weights is roughly 51% in 2002. This revenue share times the markup yields the cost share of intermediate inputs, i.e., ϕ. Our calibration of θ implies a markup of 1.2, thus, our estimate of the weighted average cost share of intermediate inputs, ϕ, is roughly 61%.[2]

The parameter $\delta = 0.025$ is the depreciation rate on physical capital and $\kappa = 3$ is the investment adjustment cost parameter, consistent with the estimates in Christiano et al. (2005) and Justiniano et al. (2010). γ_1 is set so that steady state utilization is 1, and γ_2 is set to five times γ_1, consistent with the estimates provided in Justiniano et al. (2010, 2011).

We assume that firms have to fully finance the costs of all variable inputs of production, so that $\psi_L = \psi_K = \psi_\Gamma = 1$. This is based on the analysis in Phaneuf et al. (2018), who show that this form of extended borrowing can help models generate hump-shaped inflation dynamics conditional on a monetary policy shock without relying on backward price and wage indexation. The parameters of the Taylor rule are $\rho_i = 0.8$, $\alpha_\pi = 1.5$, and $\alpha_y = 0.2$.

We calibrate parameters relating to trend growth and inflation to fit observable features of the data. Online Appendix B describes how we construct time series of output, consumption, and investment corresponding to the concepts in our model. Over the period 1960:I-2007:III, the average growth rate of the relative price of investment is -0.00472. This implies a value of $g_{\varepsilon^I} = 1.00472$. The average growth rate of the implied price index for output over the same period is 0.008675, implying a value $\pi^* = 1.008$, or 3.52 percent at an annual frequency. The average growth rate of output per capita over this period is 0.005712, or 2.28 percent annualized. Given the values of g_{ε^I} and ϕ, g_A is chosen so that $g_A^{1-\phi} = 1.0022$, which generates steady state growth in output in the model equal to the observed counterpart in the data. This calibration implies an annualized growth rate of TPF of about 1 percent per year.[3]

Our model includes three stochastic disturbances – neutral productivity, marginal efficiency of investment, and monetary policy. It is common to specify medium-scale models with a number of disturbances (e.g. Smets

[2]The steady-state price markup is for a trend inflation of zero. We find that this markup is almost insensitive to trend inflation between 0 and 4 percent leaving ϕ unaffected as trend inflation rises.

[3]This is a lower average growth rate of TFP than would obtain in a conventional growth accounting exercise. This is due to the fact that our model includes roundabout production, which means that a traditional growth accounting exercise ought to overstate the growth rate of TFP relative to true productivity.

and Wouters 2007). The proliferation of shocks in these models has been criticized by some researchers, notably Chari et al. (2009) who argue that only three shocks – the productivity shock, the investment shock, and the monetary policy shock – can be considered structural. They argue that other shocks are "dubiously structural" and lack a clear economic interpretation.

Neutral productivity shocks are typically estimated to be quite persistent. This finding emerges both in structural Bayesian estimations of fully-specified DSGE models as well as in univariate growth accounting exercises. Accordingly, we set the autoregressive parameter of the neutral productivity shock at 0.95. There is less compelling evidence on the persistence of the marginal efficiency of investment shock. We follow Justiniano et al. (2011) and set the autoregressive parameter of the MEI process at 0.81. We later asses sensitivity of our results for higher or lower values of this parameter, which ends up being crucial for the cyclical implications of trend inflation.

To pin down the standard deviations of the three shocks in our model, we proceed as follows. We target a size of shocks s_I, s_A, and s_r, for which our baseline model exactly matches the actual volatility of output growth observed in our data (0.0078) for a quarterly average trend inflation equal to its observed value during the postwar period ($\pi^* = 1.0088$). To determine the exact numbers for s_I, s_A, and s_r, we assign to each type of shock a target percentage contribution to the unconditional variance decomposition of output growth. In particular, we target a 50 percent share of the variance of output growth due to the MEI shock, 35 percent to the productivity shock, and 15 percent to the monetary shock. This implies values of $s_I = 0.0276$, $s_A = 0.0030$, and $s_r = 0.0020$. Our targets for the contribution of the three shocks to the variance of output growth are based on empirical consensus from the recent literature. In this literature, investment shocks are the main driver behind business-cycle fluctuations, followed by neutral technology shocks. Recent papers finding a large role for investment shocks include Fisher (2006), Justiniano and Primiceri (2008), Justiniano et al. (2010, 2011), and Altig et al. (2011).

To assess the empirical relevance of our baseline model, we analyze some basic business cycle moments and compare them to moments from the data. The model is solved via second order perturbation about the non-stochastic steady state. The moments are summarized in Table 2. The reported volatility and correlation statistics are for variables measured in growth rates or as deviations from stochastic trends obtained using the HP filter.

The volatility of output growth equals the actual volatility in the data by construction. The model slightly over-predicts the volatility of HP-filtered log output relative to the data. The model does a very good job matching the volatility of consumption in the data, whether measured in growth rates or HP filtered log-levels. The volatility of investment is somewhat overestimated by our model, but remains plausible. The volatility of first-differenced hours is somewhat higher in the model than in the data, while the model slightly underestimates the volatility of HP-filtered log hours. The baseline model somewhat underestimates the variability of inflation in the data.

The correlation between the growth rates of consumption and output predicted by our baseline model fits the data quite well (0.63 vs 0.75), while our model slightly underpredicts the cyclicality of consumption with output when measured in HP filtered log levels (0.59 versus 0.91). Our model accurately delivers a strong correlation between investment and output, whether measured in growth rates or filtered log-levels. Labor hours are procyclical in our model, though slightly less so when compared to the data.

The first-order autocorrelation of inflation predicted by the model is high at 0.82, and remains high up to lags of several quarters. Our model predicts that inflation is highly persistent in spite of the fact that it abstracts from wage and price indexation. The model also generates a positive first-order autocorrelation of output growth of 0.65 compared to 0.36 in the data, which according to Cogley and Nason (1995) is a useful test of the strength of the endogenous business-cycle propagation mechanisms embedded in a particular model.

4. The Welfare Costs of Trend Inflation

This section examines the normative implications of moderate trend inflation. We begin by discussing issues related to the measurement of welfare costs. Then we present and discuss our main results. We conclude with a discussion related to the role of indexation.

4.1. Measuring Welfare Costs

Welfare is defined as the sum of the present discounted value of flow utilities across households. It is implicitly a function of the realization of the vector of state variables, S_t:

$$V(S_t) = \ln(C_t - bC_{t-1}) - \eta v_t^w \frac{L_t^{1+\chi}}{1+\chi} + \beta E_t V(S_{t+1}). \tag{17}$$

The variable v_t^w is a measure of wage dispersion and arises due to aggregating differentiated labor inputs across households. As in Schmitt-Grohé and Uribe (2004), this recursive measure of welfare can be included as an equilibrium condition in a second order solution.[4] The model is solved for different levels of trend inflation. Let a subscript B denote a "base" scenario (e.g. two percent trend inflation) and A be an alternative scenario (e.g. four percent trend inflation).

The level of trend inflation affects both the policy functions, $V(\cdot)$, as well as realizations of the state vector, S_t. There are several ways in which one can compare welfare across different levels of trend inflation. One is to compute consumption-equivalent differences based on steady state welfare. Let λ denote the fraction of consumption that would need to be sacrificed in each period in the base case to yield the same welfare as in the alternative case. Based on non-stochastic steady states, it is given by:

$$\lambda_{ss} = 1 - \exp\left[(1-\beta)\left(V_A(S_A^*) - V_B(S_B^*)\right)\right] \tag{18}$$

An alternative welfare comparison is based on stochastic means rather than non-stochastic steady states. An advantage of this measure is that it incorporates information on how the level of trend inflation interacts with stochastic disturbances. It is given by:

$$\lambda_m = 1 - \exp\left[(1-\beta)\left(E\left[V_A(S_{A,t})\right] - E\left[V_B(S_{B,t})\right]\right)\right] \tag{19}$$

The consumption equivalent metrics defined in (18)-(19) do not account for transition dynamics. That is, if an economy were to transition from a world with two percent trend inflation (B) to a world with four percent inflation (A), either of the above metrics might overstate the welfare cost because it takes time for the state variables to transition to the higher inflation regime. An alternative metric would be to condition on exactly the same realization of the state as in the base model. For example, starting from the steady state of the base model, this metric would be:

$$\lambda_{tr} = 1 - \exp\left[(1-\beta)\left(V_A(S_B^*) - V_B(S_B^*)\right)\right] \tag{20}$$

For our analysis in the text, we focus on the welfare metrics based on the steady state and stochastic means. The welfare losses computed using (20) are slightly smaller, but overall very similar, to the metrics based on the steady state.[5]

4.2. Main Results

Table 3 reports the welfare costs implied by the benchmark model, with panel (a) showing welfare losses computed from non-stochastic steady states while panel (b) presents welfare losses computed from stochastic means. This table focuses on various different changes in the inflation target, though in the text we focus mostly on the costs of going from a two to four percent (annualized) target, a scenario consistent with many recent policy proposals.

According to our benchmark model, increasing trend inflation from 2 to 4 percent would generate a consumption-equivalent welfare loss of 3.7 percent conditioned on non-stochastic steady states and 4.3 percent conditioned on stochastic means.[6] The gap between the welfare loss based on the stochastic mean and the

[4] Note that we take a second order approximation to all equilibrium conditions. This means that we need not make dubious assumptions such as the existence of a subsidy to undo steady state distortions associated with monopoly power to derive a closed form welfare cost measure.

[5] For example, the welfare loss of moving from 2 to 4 percent trend inflation based on the non-stochastic steady is 3.73 percent. Based on (20), the consumption equivalent welfare loss is 3.69 percent.

[6] See Amano et al. (2007) for an analysis of how trend inflation affects the steady-state and mean values of key macroeconomic variables. That the welfare loss based on stochastic means is higher than the one based on the deterministic steady state is consistent with their analysis.

one based on the deterministic steady state depends on the properties of the stochastic processes as well as other features of the model, points to which we return in more depth below. Based on stochastic means, the cost of going from 0 to 2 percent trend inflation is 2.2 percent, whereas that of going from 2 to 4 percent is nearly twice as large. This non-linearity is important when thinking about policies to raise the inflation target in light of the zero lower bound, as in Coibion et al. (2012). While very small amounts of trend inflation might be desirable to reduce the frequency of ZLB episodes, going from 2 to 4 percent trend inflation would result in substantially larger welfare costs.

Our metrics for the welfare loss of trend inflation are larger than most reported values in the existing literature. While extremely high levels of trend inflation can imply large costs, it is generally found that modest amounts of trend inflation (say, between 0 and 4 percent annualized) have small welfare costs. An exception is Amano et al. (2009), who find that increasing trend inflation from its optimal level (slightly negative in their model) to 4 percent results in a mean welfare cost of about two percent of consumption. In comparison, the welfare costs we find are almost three times larger than that – in our model, going from 0 percent trend inflation to 4 percent implies a welfare cost of about 6.4 percent based on stochastic means. Our model shares two important features with theirs – the coexistence of both price and wage rigidity as well as trend output growth – but includes several features from which their model abstracts. In addition to capital accumulation, our model also features a number of real rigidities and frictions. Some of these are relatively standard in the literature while some others are not as common. These features include extended borrowing, roundabout production, and important stochastic shocks to the marginal efficiency of investment. We discuss in turn the roles played by all of these features in driving our results.

Table 4 shows both the steady state and mean welfare losses of going from 2 to 4 percent trend inflation in a variety of alternative specifications of our model. Deviations from our benchmark specification are described in the left column. When changing any feature of the model, the standard deviations of the shocks are re-calibrated to match the observed volatility of output growth as well as our specified variance decomposition. This is done to facilitate comparisons with our benchmark case.

The first row considers the case where wages are flexible, $\xi_w = 0$. Here the welfare costs of trend inflation are substantially smaller than in our benchmark analysis – the mean cost of going from 2 to 4 percent trend inflation is only about 1 percent of consumption. The next row considers the case of no trend growth. For this exercise we set the trend growth rates of both the IST and neutral productivity terms to zero. Here the mean cost of going from two to four percent trend inflation is about 2.5 percent of consumption, also substantially smaller than our benchmark results. It does not make much difference whether trend growth comes from neutral productivity or investment specific technical change. In row (iii) we consider the case in which trend growth in output comes exclusively from neutral productivity, and in row (iv) the one in which trend growth comes exclusively from IST. The steady state and mean welfare costs of higher trend inflation are about the same in both cases. Row (v) considers the case where there is no trend growth in output and wages are flexible. The welfare costs in this case are similar to the case when there is trend growth but wages are flexible. Row (vi) presents welfare costs of higher trend inflation when prices are flexible. Here the welfare costs of higher trend inflation based on the deterministic steady state are slightly smaller than in our baseline analysis, but the costs based on stochastic means are actually somewhat higher when prices are flexible. This arises due to the fact that we re-calibrate the shock sizes to match the observed volatility of output growth for each iteration of the model. Taken together, rows (i) and (vi) indicate that wage rigidity plays a far stronger role in driving the welfare costs of trend inflation compared to price rigidity.

We next turn to an analysis of the role of real features in the model in accounting for the welfare costs of trend inflation. We first consider the case in which there is no roundabout production, i.e. $\phi = 0$. This is shown in row (vii) of Table 4. This results in a mean welfare cost of going from 2 to 4 percent trend inflation of 3.7 percent, or about 0.6 percentage points lower than our baseline welfare cost. In row (viii) we consider the role of our assumption of extended borrowing, wherein firms must borrow to finance all variable inputs. The absence of this feature in the model leads to a similar reduction in the welfare cost of trend inflation. In row (ix) we consider the case in which only the wage bill must be borrowed in advance, which is a fairly common assumption in the literature (e.g., Christiano et al., 1997, 2005; Ravenna and Walsh, 2006). Here the mean cost of increasing the trend inflation rate from 2 to 4 percent is 3.8 percent, also lower than our benchmark case.

Next, we consider the role of marginal efficiency of investment shocks, which in our calibration account for half of the variance of output growth. To our knowledge, no previous study has examined the interaction

between trend inflation and MEI shocks. For this particular exercise, we set the variance of the MEI shock to zero and re-calibrate the magnitudes of the productivity and monetary policy shocks so as to explain 75 and 25 percent of the variance of output growth, respectively. Results are shown in row (x) of the table. Although the absence of MEI shocks does not have any impact on the steady state costs of trend inflation, the consumption equivalent welfare loss based on stochastic means is about 0.4 percentage points lower than in our baseline analysis at 3.9 percent. When all three of the aforementioned features are excluded from the model – roundabout production, extended borrowing, and MEI shocks (row (xi))– the mean welfare cost of going from 2 to 4 percent trend inflation is a full percentage point lower than in our baseline, at 3.3 percent of consumption. The important conclusion of these exercises is that real rigidities and frictions, which are often included to give a model more realistic amplification and propagation, tend to magnify the welfare costs of trend inflation.

Row (xii) considers the case where there is no roundabout production, no extended borrowing, and no MEI shocks, and in addition turns off wage rigidity and trend output growth. Here the welfare costs are very small, amounting to about 0.2 percent of consumption. Comparing these results to row (v), roundabout production, extended borrowing, and MEI shocks substantially amplify the welfare costs of trend inflation even without wage stickiness and trend growth – the welfare cost is about 1 percent with these features, or five times larger than when these features are absent and wages are flexible and there is no trend growth. Row (xiii) turns off the productivity shock but otherwise assumes the same parameterization and structure of our baseline model. For this exercise we parameterize the standard deviations of the MEI and monetary shocks to account for 75 and 25 percent of the variance of output growth, respectively. This alteration of our parameterization increases the mean welfare cost to 4.5 percent of consumption, or about 0.2 percent more than in our baseline. This exercise further confirms that the interaction between trend inflation and MEI shocks is relatively costly when compared with the two other shocks in our model.

An important parameter governing how wage rigidity interacts with trend inflation is the Frisch labor supply elasticity. In our model the inverse Frisch elasticity is given by the parameter χ. Row (xiv) considers the case where this parameter is set to 0, implying an infinite Frisch elasticity, as in the indivisible labor models of Rogerson (1988) and Hansen (1985). Here we see that the welfare costs of trend inflation are substantially smaller than in our baseline analysis, with a cost of going from 2 to 4 percent trend inflation based on stochastic means of 1.5 percent. The intuition for the effect of χ on the welfare costs of trend inflation is straightforward. Trend inflation distorts the relative allocation of labor across households through an effect on wage dispersion. With curvature in preferences over labor, this misallocation can be quite costly. But if this curvature is absent, as is the case with $\chi = 0$ (or more generally relatively unimportant for low values of χ), misallocated labor arising from wage dispersion has much smaller effects on welfare.

4.3. Indexation

Our baseline model abstracts from indexation of non-reoptimized prices and wages to the previous period's inflation rate or to trend inflation, as sometimes assumed in the literature. Nevertheless, the use of indexation has been criticized by a number of researchers. Regarding price indexation, Woodford (2007, pg. 204), for instance, argues that "the model's implication that prices should continuously adjust to changes in prices elsewhere in the economy flies in the face of the survey evidence." Cogley and Sbordone (2008, pg. 2101) mention that backward-looking price indexation lacks "a convincing microeconomic foundation." Chari et al. (2009, pg. 261) state that "this feature is inconsistent with microeconomic evidence on price setting." Finally, Christiano (2016, pg. 354) argues that the "no-indexation assumption is suggested by the same microeconomic observations that motivate price setting frictions in the first place." Price indexation also has the undesirable property that it delivers inertial inflation responses to non-monetary shocks, a finding which is generally at odds with available VAR evidence. Phaneuf et al. (2018) show that a model similar to the one considered here can account for inertial inflation responses to monetary shocks and non-inertial responses to non-monetary shocks without relying on indexation. Row (xv) of Table 4 augments our model to allow for non-updated prices to be fully indexed to the previous period's inflation rate. This has only a modest effect on the welfare cost of trend inflation, with the consumption equivalent change in welfare based on stochastic means of going from 2 to 4 percent inflation about 3.8 percent. This is consistent with our earlier findings that price rigidity is relatively unimportant in driving the large welfare costs of trend inflation.

Regarding "the troubling wage indexation assumption," Christiano et al. (2016, pg. 1524) argue that there "is little evidence that this type of indexation is widely used" (see also footnote 4 therein). hof (2012)

show that the degree of wage and price indexation has varied over time in the US, being substantial during the Great Inflation period. The evidence is, however, mixed. Rabanal and Rubio-Ramírez (2005) suggest that indexation has been relatively stable throughout the postwar period.[7] Using Bayesian methods they estimate a model that features backward-looking indexation for both prices and wages. They report a coefficient of price indexation of 0.76 and one for wages of 0.25 for the sample 1960:I-2001:IV. Their estimates are very similar for the subsample 1982:IV-2001:IV. In any case, there is no evidence that moderate changes in the level of trend inflation (up to 4 percent), as considered in this paper, bring about a substantial change in the level of wage indexation.

The previous results about the importance of wage rigidity suggest that the assumption of no wage indexation is important for our results. In row (xvi) of Table 4 we augment the model to allow for complete indexation to last period's inflation rate. Indeed, this results in a substantially lower welfare cost of trend inflation, of a similar magnitude to the case when wages are perfectly flexible. To explore the role of wage indexation further, Figure 1 plots the consumption equivalent welfare cost (based on stochastic means) of going from 2 to 4 percent inflation as a function of a wage indexation parameter, γ_w, ranging from zero to one. As one might naturally expect, the welfare costs are decreasing in the degree of wage indexation. Nevertheless, the welfare costs remain high (compared to most of the existing literature) for a range of plausible values of this parameter. For example, at $\gamma_w = 0.5$, the welfare cost of moving from 2 to 4 percent trend inflation remains about 2 percent of consumption.

Given the importance of this parameter for our results, it useful to look at the the empirical evidence from a micro level. Barattieri et al. (2014) provide an exhaustive and up-to-date study of micro data from the US relating to wage-setting. They report that the probability of a quarterly wage change for any reason is between 20 and 25 percent. Although they are unable to distinguish between re-optimized wages and wages mechanically adjusted due to indexation, their estimated hazard rate is not consistent with indexation being important. In Europe, the Eurosystem Wage Dynamic Network research project has collected data about the micro behavior of wages for twenty-five European countries. Babeck et al. (2010) show that some degree of wage indexation is present (17 percent of firms are affected) in continental European countries. There are, however, substantial difference across countries. Indexation mechanisms are especially prevalent in Belgium and Spain, whereas fewer than 5 percent of firms use indexation in Estonia and Italy. Le Bihan et al. (2012) provide a detailed study of these micro data for French firms. While their estimated probability of a wage change is somewhat higher than estimates by Barattieri et al. (2014), they conclude that "results from all specifications indicate the overall degree of indexation is weak"(Le Bihan et al., 2012, pg. 28). Sigurdsson and Sigurdarottir (2011) study wage-setting using micro data from Iceland. Their results concerning the probability of wage adjustment are broadly in line with the estimates in Barattieri et al. (2014) for the US.

5. The Cyclical Effects of Trend Inflation

This section analyzes the positive implications of moderate trend inflation. We focus on impulse responses of key macroeconomic variables to the three shocks in the model for different levels of trend inflation. An important objective of this section is to highlight the strong but heretofore unknown (to our knowledge) interaction between MEI shocks and the level of trend inflation.

Figure 2 plots impulse responses of output and inflation to the three shocks in the model for three different levels of trend inflation. Solid lines are for the case of zero trend inflation, dashed lines for two percent (annualized) trend inflation, and dotted lines are for four percent trend inflation. There is not much of an effect of trend inflation on the responses of output and inflation to a productivity shock. For higher levels of trend inflation, the output response is slightly smaller at short forecast horizons and somewhat higher at longer horizons, but the effect is not very strong. Inflation falls by a little less for higher levels of trend inflation, but again the effect of trend inflation on the response is rather weak.

Output responds more to an expansionary monetary policy shock for higher levels of trend inflation. These effects are more noticeable in the graphs than for the productivity shock. Interestingly, inflation jumps up by more on impact but is less persistent for higher levels of trend inflation. Note also that the response of inflation to the monetary shock is hump-shaped in spite of the fact that our baseline model

[7] Rabanal and Rubio-Ramírez (2008) obtain essentially the same results for a model of the Euro area.

features no price or wage indexation. While the interaction between trend inflation and the output response to a policy shock is qualitatively the same as in Ascari and Sbordone (2014), the effect on inflation is different. In their baseline New Keynesian model, inflation reacts less on impact but is more persistent for higher levels of trend inflation after a monetary policy shock. In our model the reverse is true.

We next turn attention to the interaction between trend inflation and the impulse responses to a MEI shock. Visually it is clear that the interaction between trend inflation and the MEI shock is stronger than for either the monetary or productivity shocks. Higher levels of trend inflation are associated with a larger output response to the MEI shock at all forecast horizons. The effect of trend inflation is largest at long forecast horizons. At a twenty quarter forecast horizon, for example, the output response is more than twice as large with four percent trend inflation than it is with zero trend inflation. Inflation reacts more strongly to the MEI shock for higher levels of trend inflation, but this is not nearly as noticeable as with the output response.

There is a strong and interesting relationship between the persistence of the shock and the effect of trend inflation. Figure 3 plots the impulse response of output to a MEI shock for three different levels of trend inflation. In the different panels we consider different values of ρ_I. For relatively small values of ρ_I, positive trend inflation results in a stronger output response to a MEI shock at all forecast horizons. But for higher values of ρ_I, the effect of trend inflation on the response to the MEI shock is the opposite. For $\rho_I = 0.95$, for example, the response of output to the MEI shock at four percent inflation is roughly one-third its value with zero trend inflation at medium forecast horizons. The shape of the response is also slightly different. For $\rho_I = 0.99$, these effects are exacerbated – the response of output to the MEI shock when trend inflation is four percent is negative for several years. In other words, it seems that at low levels of persistence higher rates of trend inflation make the MEI shock more expansionary, while the reverse is true at high levels of persistence in the shock.

To get a better sense of the relationship between the persistence of the MEI shock and the impact of trend inflation, Figure 4 plots the (absolute) difference of the output response to a MEI shock at a ten quarter horizon for inflation rates of four and two percent. This difference is increasing for values of ρ_I up until about 0.75 and is positive for values of ρ_I less than 0.9. For values of ρ_I greater than 0.9, trend inflation exerts a negative effect on the response of output to a MEI shock. This effect is particularly large at values of ρ_I near one.

The interaction between the persistence of the MEI shock, trend inflation, and the output response to the MEI shock suggests that the persistence of the shock might have important effects on the welfare costs of trend inflation. To examine this, we return to the exercises of the previous section. We $\rho_I = 0.95$ and re-solve the model. Panel (xvii) of Table 4 shows the mean welfare costs of increasing trend inflation from two to four percent with this higher persistence in the shock. The mean welfare cost is more than twice as large as under our baseline, with the cost of trend inflation amounting to more than 11 percent of consumption. In panel (xviii), we consider the case in which the MEI shock is less persistent, with $\rho_I = 0.66$. Here the mean welfare cost of increasing trend inflation from two to four percent is 0.0396, or about 0.3 percentage points less than in our baseline.

What is the intuition for the "sign flip" in the effect of trend inflation on the responses to MEI shocks for different levels of trend inflation? In the context of a textbook New Keynesian model with sticky prices only, Ascari (2004) shows that trend inflation both "flattens" the Phillips Curve and makes current inflation more sensitive to expected inflation. The MEI shock is an aggregate demand shock, raising current demand for goods relative to supply and pushing output and inflation in the same direction. Holding expected inflation fixed (i.e. considering a sufficiently transitory shock so as to ignore any effect on expectations of the future), a flatter Phillips Curve due to higher trend inflation should imply that a positive demand shock ought to have a bigger effect on output. This is exactly what we observe for moderate levels of persistence in the MEI shock – output rises by more the higher is trend inflation. Holding the slope of the Phillips Curve fixed, in contrast, a heightened sensitivity to the future due to higher trend inflation ought to result in a shock which raises expected inflation having smaller effects on output. When the MEI shock is sufficiently persistent, this expectations channel is quantitatively more important than the flattening of the Phillips Curve, and trend inflation exerts a dampening effect on the output response to a MEI shock.

To our knowledge, ours is the first paper to point out both the large interaction between trend inflation and the MEI shock and the dependence of that interaction on the persistence of the shock. There is a growing body of work suggesting that "investment shocks," broadly defined, may play a key role in short run

business cycle fluctuations. Fisher (2006) argues that IST shocks are an important source of fluctuations. Justiniano et al. (2010, 2011) argue that MEI shocks are perhaps the primary driver of business cycle fluctuations. Justiniano et al. (2011) argue that the MEI shock might proxy for a more fundamental shock to the functioning of the financial sector, which fits well with conventional wisdom concerning the cause and propagation of the recent Great Recession. Combining the empirical findings of the importance of investment shocks with our results about the important interaction of MEI shocks with trend inflation, we argue that it is important from a positive perspective to take into account the effects of trend inflation in DSGE business cycle models.

6. Conclusion

Economists have recently debated the desirability for central banks of the world to raise their inflation targets. This debate is the result of economic pain experienced during and after the Great Recession. Ireland (2011), for instance, argues that because of the zero lower bound on nominal interest rates, the Federal Reserve was prevented from stabilizing the U.S. economy as it previously did. With the same flexibility the Federal Reserve had in the two previous recessions, the last one might have been shorter and less severe.

Proposals to raise the inflation target are built on the premise that it would not be costly to permanently increase trend inflation by a moderate amount, say from 2 to 4 percent at annualized rate. Despite the practical importance of this question, few efforts have been devoted to address this question in the context of the empirically realistic medium-scale DSGE models that central bankers and academics use to study the macroeconomy. The evidence we have provided here offers a comprehensive benchmark against which these costs can be gauged in future research, and serves as a cautionary warning that the welfare costs of increasing an inflation target may be substantially higher than many think.

Another important side to our findings is their implications for the business cycle. Moderate trend inflation can have strong distorting effects conditioned on shocks to the marginal efficiency of investment. The strong interaction between trend inflation and MEI shocks has been heretofore overlooked in the literature. Given that much recent research ascribes an important role to MEI shocks, our analysis suggests that it is not innocuous to ignore trend inflation and that trend inflation may have larger effects on business cycle dynamics than previously thought.

References

, 2012. Time variation in u.s. wage dynamics. Journal of Monetary Economics 59 (8), 769 – 783.

Acemoglu, D., Akcigit, U., Kerr, W., 2015. Networks and the macroeconomy: An empirical exploration. NBER Macroeconomics Annual.

Altig, D., Christiano, L., Eichenbaum, M., Lindé, J., 2011. Firm-specific capital, nominal rigidities, and the business cycle. Review of Economic Dynamics 14(2), 225–247.

Amano, R., Ambler, S., Rebei, N., 2007. The macroeconomic effects of non-zero trend inflation. Journal of Money, Credit and Banking 39 (7), 1821–1838.

Amano, R., Moran, K., Murchison, S., Rennison, A., 2009. Trend inflation, wage and price rigidities, and productivity growth. Journal of Monetary Economics 56, 353–364.

Ascari, G., 2004. Staggered prices and trend inflation: Some nuisances. Review of Economic Dynamics 7, 642–667.

Ascari, G., Ropele, T., 2007. Optimal monetary policy under low trend inflation. Journal of Monetary Economics 54, 2568–2583.

Ascari, G., Sbordone, A. M., 2014. The macroeconomics of trend inflation. Journal of Economic Literature 52(3), 679–739.

Babeck, J., Caju, P. D., Kosma, T., Lawless, M., Messina, J., Rm, T., 2010. Downward nominal and real wage rigidity: Survey evidence from european firms. The Scandinavian Journal of Economics 112 (4), 884–910.

Ball, L., 2013. The case for four percent inflation. Central Bank Review 13 (2), 17–31, Research and Monetary Policy Department, Central Bank of the Republic of Turkey.

Barattieri, A., Basu, S., Gottschalk, P., 2014. Some evidence on the importance of sticky wages. American Economic Journal: Macroeconomics 6(1), 70–101.

Barth, M., Ramey, V., 2002. The cost channel of monetary transmission. NBER Macroeconomics Annual.

Basu, S., 1995. Intermediate goods and business cycles: Implications for productivity and welfare. American Economic Review 85(3), 512–531.

Bils, M., Klenow, P. J., 2004. Some evidence on the importance of sticky prices. Journal of Political Economy 112(5), 947–985.

Blanchard, O., Dell'Ariccia, G., Mauro, P., 2010. Rethinking macroeconomic policy. Journal of Money, Credit, and Banking 42(s1), 199–215.

Calvo, G. A., 1983. Staggered prices in a utility-maximising framework. Journal of Monetary Economics 12, 383–398.

Chari, V. V., Kehoe, P. J., McGrattan, E. R., 2009. New keynesian models: Not yet useful for policy analysis. American Economic Journal: Macroeconomics 1(1), 242–266.

Christiano, L. J., 2016. Comment on "Networks and the macroeconomy: An empirical exploration" in: Martin Eichenbaum and Jonathan Parker (eds.), NBER Macroeconomics Annual 2015, Volume 30, Chicago: University of Chicago Press., 346–373.

Christiano, L. J., Eichenbaum, M., Evans, C. L., 1997. Sticky price and limited partecipation models of money: A comparison. European Economic Review 41, 1201–1249.

Christiano, L. J., Eichenbaum, M., Evans, C. L., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. Journal of Political Economy 113 (1), 1–45.

Christiano, L. J., Eichenbaum, M., Trabandt, M., 2016. Unemployment and Business Cycles. Econometrica 84, 15231569.

Cogley, T., Nason, J. N., 1995. Output dynamics in real-business-cycle models. American Economic Review 85, 492–511.

Cogley, T., Sbordone, A. M., 2008. Trend inflation, indexation and inflation persistence in the New Keynesian Phillips Curve. American Economic Review 98 (5), 2101–2126.

Coibion, O., Gorodnichenko, Y., 2011. Monetary policy, trend inflation and the great moderation: An alternative interpretation. American Economic Review 101 (1), 341–370.

Coibion, O., Gorodnichenko, Y., Wieland, J. F., 2012. The optimal inflation rate in New Keynesian models: Should central banks raise their inflation targets in light of the ZLB? The Review of Economic Studies 79, 1371–1406.

Erceg, C. J., Henderson, D. W., Levin, A. T., 2000. Optimal monetary policy with staggered wage and price contracts. Journal of Monetary Economics 46, 281–313.

Fisher, J., 2006. The dynamic effects of neutral and investment-specific technology shocks. Journal of Political Economy 114(3), 413–451.

Fuerst, T., 1992. Liquidity, loanable funds, and real activity. Journal of Monetary Economics 29(1), 3–24.

Hansen, G. D., 1985. Indivisible labor and the business cycle. Journal of Monetary Economics 16, 309–327.

Hornstein, A., Wolman, A. L., 2005. Trend inflation, firm-specific capital and sticky prices. Federal Reserve Bank of Richmond Economic Quarterly Review 91, 57–83.

Huang, K. X., Liu, Z., Phaneuf, L., 2004. Why does the cyclical behavior of real wages change over time? American Economic Review 94(4), 836–856.

Ireland, P. N., 2011. A new keynesian perspective on the great recession. Journal of Money, Credit, and Banking 43(1), 31–54.

Justiniano, A., Primiceri, G., 2008. The time varying volatility of macroeconomic fluctuations. American Economic Review 98(3), 604–641.

Justiniano, A., Primiceri, G., Tambalotti, A., 2010. Investment shocks and business cycles. Journal of Monetary Economics 57(2), 132–145.

Justiniano, A., Primiceri, G., Tambalotti, A., 2011. Investment shocks and the relative price of investment. Review of Economic Dynamics 14(1), 101–121.

Kiley, M. T., 2007. Is moderate-to-high inflation inherently unstable? International Journal of Central Banking 3 (2), 173–201.

Krugman, P., 2014. Inflation targets reconsidered, Draft paper for ECB Sintra Conference.

Le Bihan, H., Montorns, J., Heckel, T., July 2012. Sticky wages: Evidence from quarterly microeconomic data. American Economic Journal: Macroeconomics 4 (3), 1–32.

Levin, A., Yun, T., 2007. Reconsidering the natural rate hypothesis in a New Keynesian framework. Journal of Monetary Economics 54 (5), 1344–1365.

Liu, Z., Phaneuf, L., 2007. Technology shocks and labor market dynamics: Some evidence and theory. Journal of Monetary Economics 54(8), 2534–2553.

Nakamura, E., Steinsson, J., 2010. Monetary non-neutrality in a multi-sector menu cost model. Quarterly Journal of Economics 125 (3), 961–1013.

Phaneuf, L., Sims, E., Victor, J. G., 2018. Inflation, output and markup dynamics with purely forward-looking wage and price setters. European Economic Review 105, 115–134.

Rabanal, P., Rubio-Ramírez, J. F., 2005. Comparing New Keynesian models of the business cycle: A Bayesian approach. Journal of Monetary Economics 52 (6), 1151–1166.

Rabanal, P., Rubio-Ramírez, J. F., Mar 2008. Comparing New Keynesian models in the Euro area: A Bayesian approach. Spanish Economic Review 10 (1), 23–40.

Ravenna, F., Walsh, C. E., 2006. Optimal monetary policy with the cost channel. Journal of Monetary Economics 53(2), 199–216.

Rogerson, R., 1988. Indivisible labor, lotteries, and equilibrium. Journal of Monetary Economics 21(1), 3–16.

Rotemberg, J. J., Woodford, M., 1997. An optimization-based econometric framework for the evaluation of monetary policy. In: Rotemberg, J. J., Bernanke, B. S. (Eds.), NBER Macroeconomics Annual 1997. Cambridge, MA, The MIT Press, pp. 297–346.

Schmitt-Grohé, S., Uribe, M., 2004. Solving dynamic general equilibrium models using a second order approximation to the policy function. Journal of Economic Dynamics and Control 28, 755–775.

Sigurdsson, J., Sigurdarottir, R., 2011. Evidence from nominal wage rigidity and wage setting from icelandic microdata, Central Bank of Iceland Working Paper 55.

Smets, F., Wouters, R., 2007. Shocks and frictions in US business cycles: A bayesian DSGE approach. American Economic Review 97 (3), 586–606.

Woodford, M., 2007. Interpreting inflation persistence: Comments on the conference 'quantitative evidence on price determination'. Journal of Money, Credit, and Banking 39(2), 203–210.

Table 1: Parameter Values

Parameter	Value or Target	Description
β	0.99	Discount factor
α	1/3	Capital production function
ϕ	0.61	Intermediate production function
δ	0.025	Depreciation rate
η	6	Labor disutility
χ	1	Inverse Frisch elasticity
b	0.8	Habit formation
κ	3	Inv. adjustment cost
γ_1	$u^* = 1$	Linear term utilization cost
γ_2	$5\gamma_1$	Quadratic term utilization cost
θ	6	Elasticity of sub. goods
σ	6	Elasticity of sub. labor
ψ_L	1	Borrowing labor
ψ_K	1	Borrowing capital
ψ_Γ	1	Borrowing intermediates
F	$\Pi^* = 0$	Fixed cost
ξ_p	0.66	Price stickiness
ξ_w	0.66	Wage stickiness
ρ_i	0.8	Taylor rule AR(1)
α_π	1.5	Taylor rule inflation
α_y	0.2	Taylor rule output growth
π^*	1.0088	Trend inflation (gross)
g_A	$1.0022^{1-\phi}$	Trend productivity growth (gross)
g_{ε^I}	1.0047	Trend IST growth (gross)
ρ_A	0.95	AR(1) productivity
ρ_I	0.81	AR(1) MEI
s_r	0.0020	S.D. monetary shock
s_A	0.0030	S.D. productivity shock
s_I	0.0276	S.D. MEI shock

Note: this table lists the values of calibrated parameters or the target used in the calibration.

Table 2: Moments

	$E(\Delta Y)$	$\sigma(\Delta Y)$	$\sigma(\Delta I)$	$\sigma(\Delta C)$	$\sigma(\Delta L)$	$\sigma(Y^{hp})$	$\sigma(I^{hp})$
Model	0.0057	0.0078	0.0251	0.0040	0.0095	0.0166	0.0544
Data	(0.0057)	(0.0078)	(0.0202)	(0.0047)	(0.0079)	(0.0156)	(0.0302)
	$\sigma(C^{hp})$	$\sigma(L^{hp})$	$\sigma(\pi)$	$\rho(\Delta Y, \Delta I)$	$\rho(\Delta Y, \Delta C)$	$\rho(\Delta Y, \Delta L)$	$\rho(Y^{hp}, I^{hp})$
Model	0.0082	0.0143	0.0044	0.9361	0.6314	0.3803	0.9317
Data	(0.0083)	(0.0171)	(0.0065)	(0.9192)	(0.7542)	(0.6313)	(0.9701)
	$\rho(Y^{hp}, C^{hp})$	$\rho(Y^{hp}, L^{hp})$	$\rho_1(\Delta Y)$	$\rho_1(\pi)$			
Model	0.5918	0.5754	0.6557	0.8161			
Data	(0.9053)	(0.8750)	(0.3634)	(0.9071)			

Note: this table shows selected moments generated from the baseline model. These moments are generated using the parameter values shown in the tables above with annualized trend inflation of 3.52 percent. "σ" denotes standard deviation, "Δ" refers to the first difference operator, ρ_1 is a first order autocorrelation coefficient, and a superscript "hp" stands for the HP detrended component of a series using smoothing parameter of 1600. The variables Y, I, C, and L are the natural logs of these series; π is quarter-over-quarter inflation. Moments in the data are computed on the sample 1960:I-2007:III and are shown in parentheses.

Table 3: Welfare Costs of Trend Inflation

	π^*	1.0000→	1.0200→	1.0352→
(a) Steady State				
	1.0000	0	n/a	n/a
	1.0200	0.0191	0	n/a
	1.0352	0.0449	0.0263	0
	1.0400	0.0557	0.0373	0.0113
(b) Means				
	1.0000	0	n/a	n/a
	1.0200	0.0222	0	n/a
	1.0352	0.0519	0.0303	0
	1.0400	0.0643	0.0430	0.0131

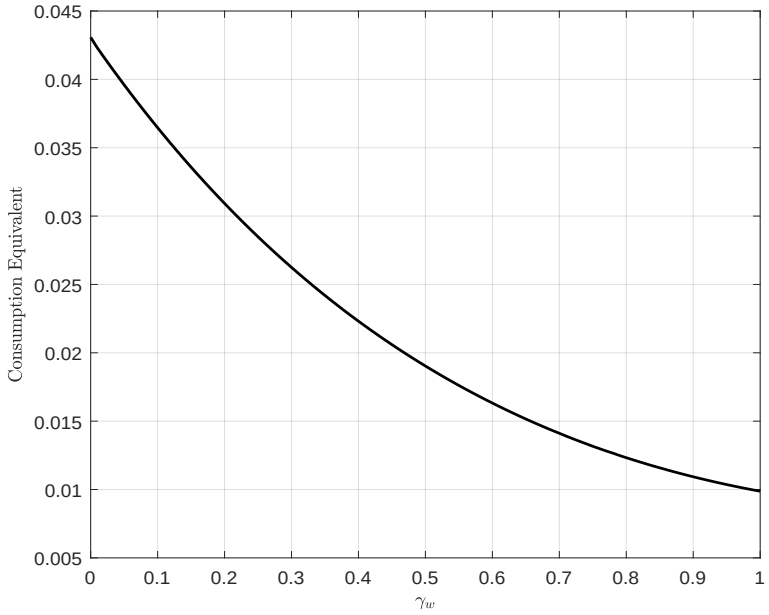
Note: this table shows consumption equivalent welfare losses from increasing the trend inflation rate using the benchmark parameterization of our model. Panels marked (a) show losses based on the non-stochastic steady state, while panels marked (b) present losses based on stochastic means. A number in the table has the interpretation as the fraction of consumption the representative household would be willing to give up to avoid changing the trend inflation rate from the level in the columns to the level shown in the rows.

Table 4: Welfare Costs of Trend Inflation: Alternative Specifications and Parameter Robustness

	π^*: 1.02 → 1.04	
Alternative Specification	Steady State	Mean
(i) Flexible wages	0.0093	0.0102
(ii) No trend growth	0.0210	0.0247
(iii) All growth from neutral productivity	0.0375	0.0436
(iv) All growth from IST	0.0371	0.0423
(v) Flexible wages, no growth	0.0095	0.0107
(vi) Flexible prices	0.0332	0.0506
(vii) No RP	0.0316	0.0367
(viii) No EB	0.0327	0.0369
(ix) Only wage subject to borrowing	0.0332	0.0377
(x) No MEI shocks	0.0373	0.0392
(xi) No RP, EB, or MEI shocks	0.0302	0.0325
(xii) No RP, EB, or MEI Shocks, flexible wages, no growth	0.0017	0.0022
(xiii) No productivity shocks	0.0373	0.0445
(xiv) Infinite Frisch elasticity ($\chi = 0$)	0.0145	0.0149
(xv) Full price indexation	0.0332	0.0377
(xvi) Full wage indexation	0.0093	0.0099
(xvii) More persistent MEI shock ($\rho_I = 0.95$)	0.0373	0.1162
(xviii) Less persistent MEI shock ($\rho_I = 0.66$)	0.0373	0.0396

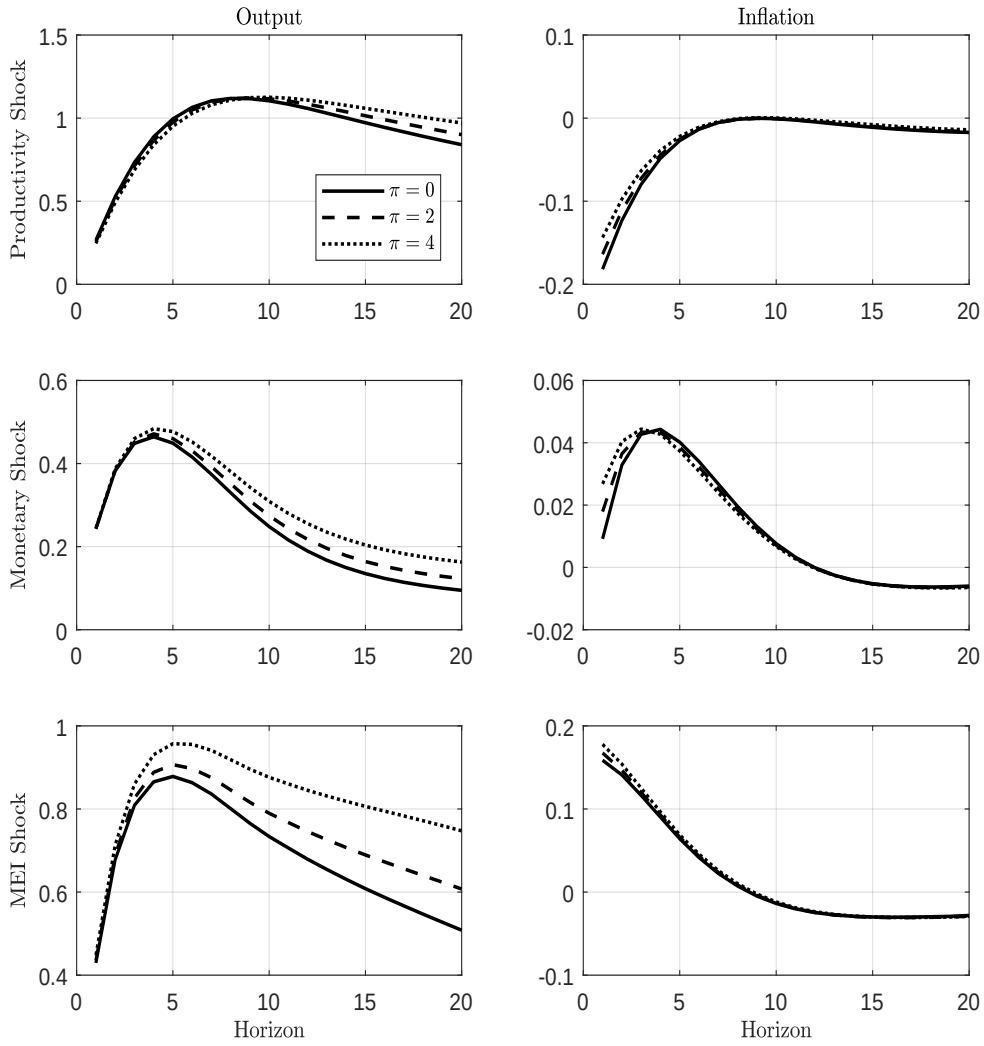
Note: this table presents consumption equivalent welfare losses from higher trend inflation under different specifications of our model. These specifications are as indicated in the rows.

Figure 1: Wage Indexation and the Welfare Costs of Trend Inflation



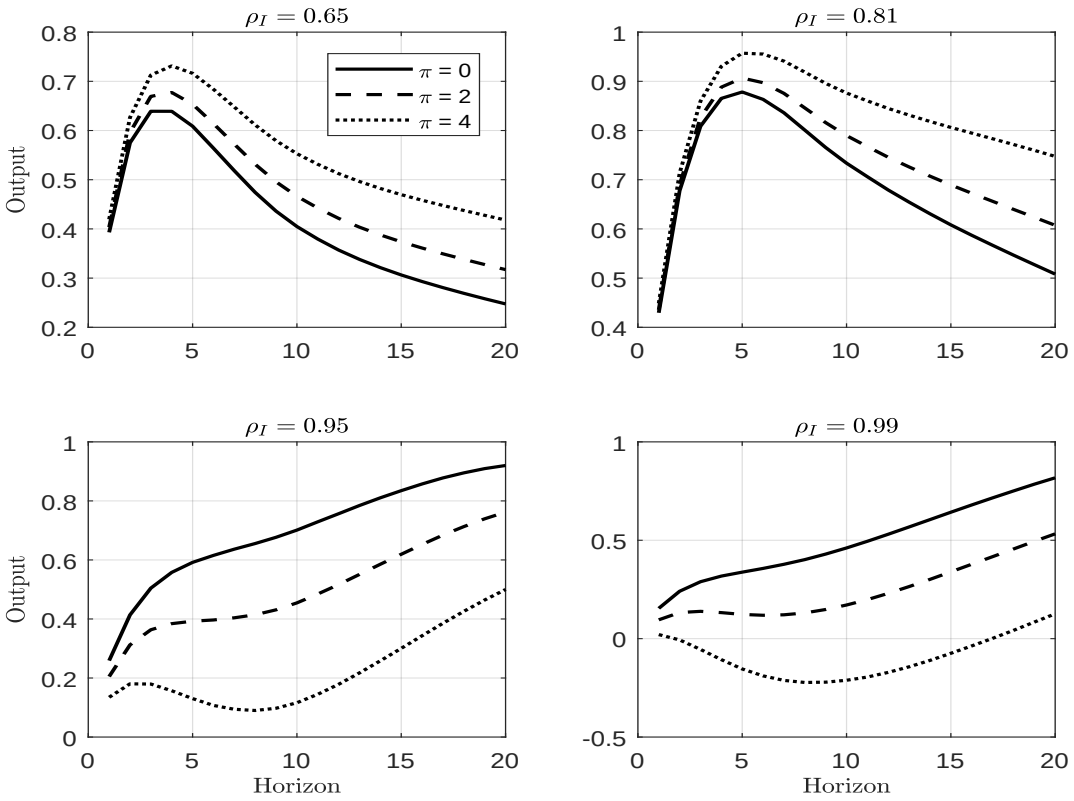
Note: this figure plots the consumption equivalent welfare cost of moving from 2 to 4 percent trend inflation (based on stochastic means) for values of the backward indexation parameter for non-optimized wages, γ_w, ranging from 0 to 1.

Figure 2: Trend Inflation and Impulse Responses



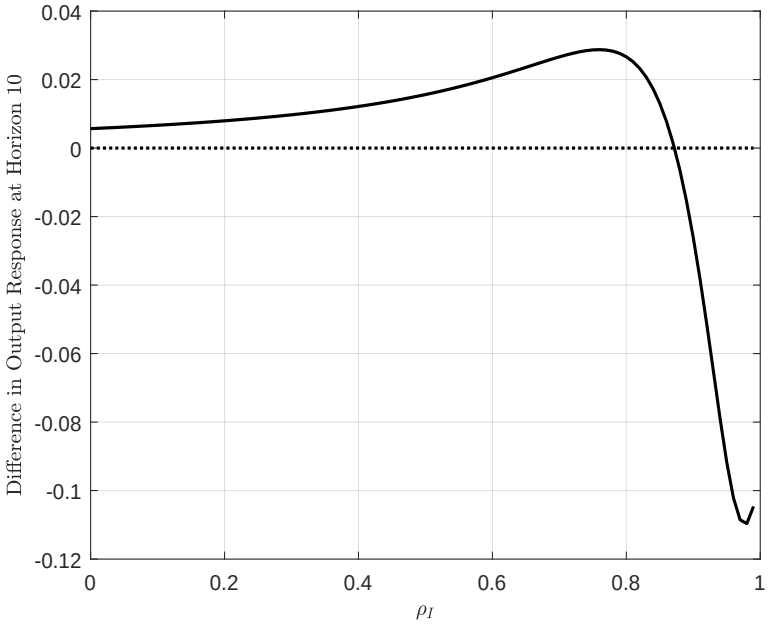
Note: this figure plots impulse responses of output and inflation to the three shocks in our model – productivity, marginal efficiency of investment (MEI), and the monetary policy shock – for three different annualized rates of trend inflation.

Figure 3: Trend Inflation and the Persistence of the MEI Shock



Note: this figure plots impulse responses of output to the MEI shock for three different levels of trend inflation. The different panels consider different values of the persistence parameter for the MEI shock, ρ_I.

Figure 4: Interaction Between Trend Inflation and the Persistence of the MEI Shock



Note: this figure plots the absolute difference between the impulse response of output at a ten quarter horizon with four percent trend inflation and zero percent trend inflation against different values of ρ_I.