Working Papers

FEDERAL RESERVE BANK OF BOSTON™

No. 19–19

Output Hysteresis and Optimal Monetary Policy

Vaishali Garga and Sanjay R. Singh

Abstract:

We analyze the implications for monetary policy when deficient aggregate demand can cause a permanent loss in potential output, a phenomenon we term *output hysteresis*. In the model, the incomplete stabilization of a temporary shortfall in demand reduces the return to innovation, thus reducing total factor productivity growth and generating a permanent loss in output. Using a purely quadratic approximation to welfare under endogenous growth, we derive normative implications for monetary policy. Away from the zero lower bound (ZLB), optimal commitment policy sets interest rates to eliminate output hysteresis. A strict inflation targeting rule implements the optimal policy. However, when the nominal interest rate is constrained at the ZLB, strict inflation targeting is suboptimal and admits output hysteresis. A new policy rule that targets output hysteresis returns output to its pre-shock trend and approximates the welfare gains under optimal commitment policy. A central bank that is unable to commit to future policy actions suffers from hysteresis bias, as the bank's inconsistent policy does not offset past losses in potential output.

Keywords: endogenous growth, zero lower bound, output hysteresis, optimal monetary policy

JEL Classifications: E52, E61, O41

Vaishali Garga is an economist in the research department at the Federal Reserve Bank of Boston. Her email address is vaishali.garga@bos.frb.org. Sanjay R. Singh is an assistant professor of economics at the University of California, Davis. His email address is sjrsingh@ucdavis.edu.

We owe an invaluable debt to Gauti Eggertsson for continuous advice and extensive comments on various versions of this paper. We thank Alexei Abrahams, Francesco Bianchi, James Cloyne, Mark Gertler, David Glancy, Violeta Gutkowski, David Hémous, Peter Howitt, Òscar Jordà, Peter Karadi, Todd Keister, Mrityunjay Kothari, Neil Mehrota, Pascal Michaillat, Matthias Paustian, Carl Walsh, and David Weil. We also thank seminar participants at the Boston Fed, Brown University, the Cleveland Fed, the Federal Reserve Board, Rutgers University, the Reserve Bank of India, the University of California, Davis, and the University of California, Santa Cruz and conference participants at CEF Fordham, the ECB 2017 workshop on monetary policy in nonstandard times, the fifth ECB/CBRT conference, Mid-West Macro 2017, the 2017 North American Summer Meetings of the Econometric Society, and the 2018 Society for Economic Dynamics for helpful comments and suggestions. All errors are our own.

This version: December, 2019 **https://doi.org/10.29412/res.wp.2019.19**

1 Introduction

The Federal Reserve's objectives of maximum employment and price stability do not, by themselves, ensure a strong pace of economic growth or an improvement in living standards. The most important factor determining living standards is productivity growth ... a portion of the relatively weak productivity growth we have recently seen may be the result of the recession itself. ... In particular, investment in research and development has been relatively weak... Federal Reserve actions to strengthen the recovery may not only help bring our economy back to its productive potential, but it may also support the growth of productivity and living standards over the longer run.
Janet L. Yellen, Former Chair of the Board of Governors of the Federal Reserve System (2015)

In the aftermath of the Great Recession, the US economy experienced its slowest post-recession recovery since World War II. Ten years in, real GDP was still approximately 15 percent below the pre-recession trend (Figure 1). Similar trajectories have been observed in other OECD countries as well (see Martin, Munyan, and Wilson 2015). One of the primary drivers of this shortfall has been slow productivity growth (Hall 2016; Stock and Watson 2016), the source of which has been a subject of extensive debate. Fernald (2014a) and Cette, Fernald, and Mojon (2016) show that total factor productivity (TFP) growth started slowing in 2004, three years before the US recession started. Thus, they say that the slower growth which followed the recession may not have been due to the recession itself. On the other hand, Decker, Haltiwanger, Jarmin, and Miranda (2014) show that the Great Recession accelerated the slowdown in startup entry, which is a significant channel for productivity growth. Similarly, investment in research and development (R&D), considered to be another important contributor to TFP growth, fell considerably during the recent recession. These facts underscore Chair Yellen's concerns as cited above.[1]

The standard theoretical treatment of monetary policy is largely silent on the interaction of monetary policy with the economy's productive potential.[2] In this paper, we construct a model in which there is such an interaction. We embed a model of Schumpeterian growth, along the lines of Aghion and Howitt (1992) and Grossman and Helpman (1991), in a New Keynesian (NK) setting. A contraction in aggregate demand reduces the incentives for firms to invest in R&D, which leads to lower innovation. This results in an endogenous slowdown in TFP growth, which accumulates into a persistent output gap. Thus, following

[1]More recently, Yellen (2016) remarked: "Are there circumstances in which changes in aggregate demand can have an appreciable, persistent effect on aggregate supply? Prior to the Great Recession, most economists would probably have answered this question with a qualified "no." ... This conclusion deserves to be reconsidered in light of the failure of the level of economic activity to return to its pre-recession trend in most advanced economies. This post-crisis experience suggests that changes in aggregate demand may have an appreciable, persistent effect on aggregate supply–that is, on potential output." (October 14, 2016)

[2]There is a recent synchronous literature that explores these interactions, including Anzoategui et al. (2019), Bianchi, Kung, and Morales (2019) and Benigno and Fornaro (2018). Ours is the first paper to analyze the interaction of optimal monetary policy at the ZLB, aggregate demand,, and endogenous growth. We discuss this at length in the related literature section.

Figure 1: Real Gross Domestic Product (GDP)

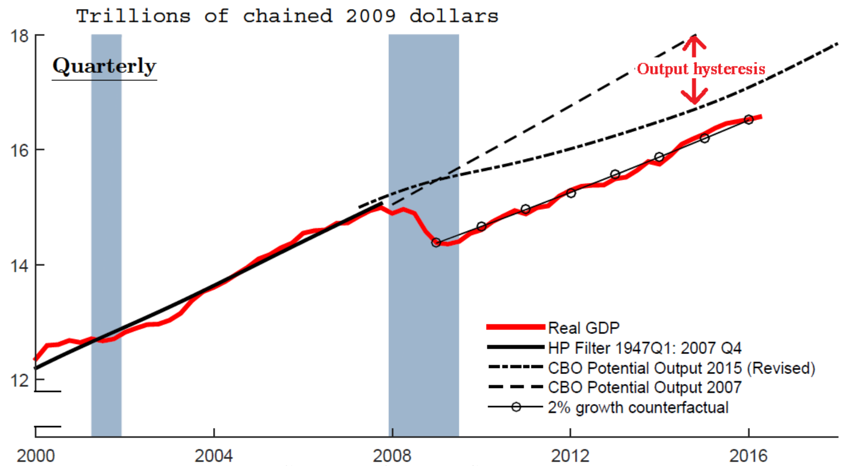

Source: Authors' calculations using quarterly real GDP data from the St. Louis FRED database and 2007 and 2015 potential output taken from the Congressional Budget Office's February 2016 releases.
Note: The trend line up to 2007:Q4 is estimated on quarterly data from 1947:Q1 to 2007:Q4 using the Hodrick-Prescott filter with a smoothing parameter of 1600. The solid black line with circles is constructed using a 2 percent annual growth rate starting in 2009. The shaded areas represent recessions dated by the National Bureau of Economic Research.

a recession, unemployment returns to its natural rate while output remains below its pre-recession trend. In this framework, monetary policy can affect the economy's long-run potential output. This is in contrast to the traditional NK models which do not incorporate endogenous productivity, and thus incorrectly predict a recovery of output to its pre-recession trend.

Using this framework, we ask whether it is optimal for monetary policy to engineer a recovery back to the pre-recession trend. Optimal policy analysis is the focus and main contribution of this paper. In order to analyze the normative implications for the conduct of monetary policy, we derive a closed-form expression for the linear-quadratic approximation of the representative agent's lifetime utility function. This expression generalizes the approximation derived by Benigno and Woodford (2004) to the endogenous growth environment and nests the exogenous growth as a special case. In particular, we decompose the stabilization objectives of the social planner into three key market distortions: a wage inflation gap, a labor efficiency gap, and a productivity growth rate gap. Of these, the productivity growth rate gap is novel to the endogenous growth framework and provides an additional rationale for the stabilization of short-run fluctuations.

We use this framework to study an economy hit with a temporary shortfall in demand. While our quadratic approximation is general, we focus the discussion on liquidity demand and monetary policy shocks because the model exhibits *divine coincidence* under these shocks. This coincidence implies that monetary policy can completely negate these shocks and maintain the economy at the *first-best level.* One implication of this property is that while the natural rate of interest, *r-star*, is exogenous, the level of potential output becomes an endogenous object. Hence, these two shocks allow us to tractably study monetary policy with endogenous growth. In this environment, we define *output hysteresis* as the gap between the actual output and its initial deterministic trend. We obtain the following three sets of results.

First, away from the zero lower bound (ZLB), an optimizing policymaker with the ability to commit to future policy actions (optimal commitment policy) sets interest rates to offset the permanent output gap. A standard textbook prescription of a strict inflation targeting rule implements the optimal policy. If the central bank strictly targets inflation and the nominal interest rate is away from the ZLB, there is no output hysteresis. Although a strict inflation targeting rule implements optimal policy away from the ZLB, it is unable to stabilize aggregate demand when the ZLB becomes a binding constraint. As a result, a strict inflation targeting rule admits output hysteresis after a ZLB episode. On the other hand, there exist policy rules which, if credibly communicated to the public, could prevent output hysteresis following recessions induced by a shortage of aggregate demand, whether or not the ZLB is binding. A new rule whereby the central bank targets zero output hysteresis emerges in the endogenous growth framework. The central bank commits to keeping interest rates lower until output returns to the initial trend. This *hysteresis targeting rule* rule signals the central bank's ex-ante commitment to running a *high-pressure economy* in the future when there is no slack in employment. Thus, we find that *output hysteresis is contingent on the monetary policy* specification of the central bank.

While the hysteresis targeting rule can eliminate output hysteresis, it raises the question of whether it is desirable to run a *high-pressure economy* using this rule. Our second set of results speak to this concern. At the ZLB, the optimal policy response is to credibly commit to keeping future interest rates low in order to incentivize recovery *close* to the pre-recession trend level. A zero output-hysteresis targeting policy rule eliminates all the persistent effects resulting from constrained monetary policy, thereby *closely* replicating the welfare gains achieved under optimal policy for a feasible range of parameters. This particular rule has the relative advantage of easily communicating the central bank's policy stance to the public, unlike most other optimal policy rules studied in the literature.

Third, and importantly, we uncover a new dynamic inconsistency problem. A policymaker unable to commit to future policy actions (discretionary policy) does not find in her interest to undo permanent output gaps following a ZLB period. This means that it is suboptimal ex-post for policy to be redesigned in order to offset the existing output hysteresis. We label this dynamic inconsistency as the *hysteresis bias* of a discretionary policymaker. This dynamic inconsistency problem complements our first finding that hysteresis is a consequence of a central bank's policy constraints, in particular, its inability to credibly commit to future policy actions, and not due to inept or irrational behavior on part of the central bank.

On the technical front, the hysteresis bias result may be surprising to business cycle scholars, given that we operate in an environment with an endogenous state variable (the level of productivity) influenced by policy levers. Because of the linearity assumption in production functions in a broad class of endogenous growth models (see Jones 2005 for details), past productivity losses do not affect the current allocation of resources between investment and consumption, and are *bygones* from the policymaker's perspective. The lack of credible monetary policy tools results in permanent output shortfalls. This long-run consequence of policy constraints provides a reason for policymakers to pursue aggressive stabilization policy through implementable rules during times of severe demand shortfalls.

The rest of the paper paper is organized as follows. We conclude this section with a brief review of the related literature. Section 2 proposes a production economy with nominal wage rigidities augmented with endogenous growth. Section 3 discusses the main theoretical results under liquidity demand shocks. A purely quadratic approximation of the household's utility function allows us to decompose the policymaker's objectives into key market distortions/wedges. In Section 4, we briefly discuss optimal monetary policy under discount rate shocks and supply shocks. We also summarize the findings from a quantitative medium-scale model that illustrates the potency of monetary policy in offsetting output hysteresis. An extended discussion of the quantitative model, including estimated structural impulse responses, is relegated to the online appendix to keep the main paper focused on the optimal monetary policy analysis. Section 5 concludes.

Related Literature

Our paper is closely related to the recent work of Anzoategui, Comin, Gertler, and Martinez (2019), Benigno and Fornaro (2018), Bianchi, Kung, and Morales (2019), Garcia-Macia (2015), Guerron-Quintana and Jinnai

(2019), Moran and Queraltó (2018) and Queralto (2019), all of whom integrate endogenous growth into a standard business cycle framework. Among these papers, our framework is most similar to that of Benigno and Fornaro (2018), who identify the possibility of an economy entering a phase characterized by a persistent liquidity trap and low TFP growth due to pessimistic expectations. While our model is closest to the one by Benigno and Fornaro (2018), our paper should be seen as complementary to it. We complement their elegant analysis by studying optimal monetary policy in response to shocks to economic fundamentals, while Benigno and Fornaro (2018) study the possibility that the economy is trapped in the ZLB equilibrium.[3] To our best knowledge, ours is the first paper to analyze the desirability of admitting permanent output gaps in the presence of severe demand shocks, particularly relevant once the ZLB is binding. The analytical result on *hysteresis bias* is new to the literature and has important implications for central bank policy.[4]

Moran and Queraltó (2018) provide empirical evidence in support of the mechanism by which monetary policy shocks affect R&D investment, which in turn affects TFP growth. In related work, Jordà, Singh, and Taylor (2019), using panel data for seventeen advanced countries over 1890–2015, provide causal evidence of the persistent effects of monetary policy.[5] Moran and Queraltó (2018) also build a quantitative endogenous growth model with nominal rigidities and emphasize the importance of the ZLB constraint on TFP. Our work complements their analysis by characterizing optimal policy at and away from the ZLB under endogenous growth. Our finding regarding the hysteresis bias emphasizes that the central bank's lack of commitment tools can lead to a persistent drop in output.

We contribute to the optimal monetary policy literature by providing an analytically tractable generalization of the textbook optimal policy problem with nominal rigidities (Woodford 2003; Benigno and Woodford 2004). Recently, a number of papers have explored the implications for optimal monetary policy in a hysteresis-prone environment. Blanchard (2018) provides a detailed survey of the empirical and the-

[3] Our framework cannot rule out the stagnation traps identified by Benigno and Fornaro (2018) without making further assumptions. For example, one can construct fiscal strategies following Benhabib, Schmitt-Grohé, and Uribe (2002) that rule out such traps for a given monetary policy rule. Our analysis implicitly assumes that there exists an equilibrium selection device that allows us to focus on the fundamentals-driven equilibrium.

[4] Stadler (1990) and Fatás (2000) are also important precursors to our paper and this recent literature.

[5] In Appendix G.1, we identify quantitatively similar effects of monetary policy shocks on R&D at the firm-level in Compustat, monthly number of business incorporations from the Survey of Current Business, new quarterly establishments births from National Private Sector Business Employment Dynamics data, utilization-adjusted TFP series from (Fernald, 2014b) using various external instruments from Jordà (2005) local projections techniques. To keep our analysis focused on optimal monetary policy, we relegate the discussion of those results and their robustness to the appendix. Meier and Reinelt (2019) also find that monetary policy shocks affect aggregate TFP. They provide a quantitative assessment in support of a misallocation channel through markup dispersion.

oretical advances in the hysteresis literature. Here we only highlight the papers that are the most closely related to ours from the optimal monetary policy perspective. Acharya, Bengui, Dogra, and Wee (2018) analyze optimal monetary policy in an environment with permanent skill-loss resulting from temporary unemployment at the ZLB. Galí (2016) solves for optimal policy in an insider-outsider model of labor markets (Blanchard and Summers, 1986). Erceg and Levin (2014) evaluate monetary policy rules in an environment where workers may exit the labor force, in order to reconcile the lower labor force participation rates in the economy. We complement these various analyses by allowing contractions in demand to negatively affect long-run supply via endogenous productivity growth. Because of the linearity assumption in the production function of the endogenous growth model, our setup is analytically tractable. Annicchiarico and Pelloni (2016) study Ramsey policy and Ikeda and Kurozumi (2014) study the use of simple operational rules in an endogenous TFP growth setting away from the ZLB.

We also contribute to the literature on stabilization policy, where DeLong and Summers (2012) and Fatás and Summers (2015) argue that these permanent deviations in output can be avoided using appropriate policy tools. These two papers focus on fiscal policy as the appropriate mechanism to counteract the permanent negative effects, while our analysis carves out a role for monetary policy as suggested by Yellen (2016). Our theoretical analysis offers an implementable policy rule for the central bank that approximates the welfare gains achieved under optimal policy. In regard to fiscal policy, we show in Appendix F that investment tax credits are expansionary; in related work, it has been shown that debt-financed fiscal policy can be self-financing in hysteresis-prone environments (see Eggertsson, Mehrotra, Singh, and Summers (2016)). However, our focus is on monetary policy.

Finally, our paper adds to the Hansen/Summers/Gordon secular stagnation literature. While our model cannot generate permanent recessions (as in Eggertsson and Mehrotra 2015; Guerrieri and Lorenzoni 2017) due to the representative agent setup, it formalizes how demand-side and supply-side secular stagnation are related. These other papers employ a permanent shock to the borrowing limit, and hence to the natural interest rate, *r-star*. As a result, output is permanently depressed. In our setting, a temporary shock to *r-star* propagates through a slowdown in TFP growth to generate a permanent effect on the level of output. Our paper formally demonstrates that secular stagnation may be a consequence of policy constraints.

2 A New Keynesian Model with Endogenous Growth

We integrate a textbook model of endogenous growth into a New Keynesian (NK) environment. Households set nominal wages in staggered contracts following Calvo (1983).[6] On the production side, we use a discrete time version of the Schumpeterian growth model of Aghion and Howitt (1992), following Aghion and Howitt (2008, Ch. 4). There is a continuum of intermediate goods, each of which is produced by a sector-specific monopolist. Economics growth results from innovations that raise productivity by improving the quality of products. These innovations are undertaken by profit-maximizing entrepreneurs in every sector, who spend final output in research. We assign the task of mitigating the effects of nominal rigidities to the monetary authority, while fiscal policy is responsible for offsetting the distortions associated with imperfect competition.

There are six main actors in our model —households, wage unions, firms, entrepreneurs, the fiscal authority, and the central bank, described below.

2.1 Households and Wage Setting

2.1.1 Households

There is a continuum of monopolistically competitive households (indexed on the unit interval), each of which supplies a differentiated labor service to the production sector. As is standard, we assume perfect risk sharing within the household. A household derives utility from consuming a final consumption good, disutility from supplying labor, and utility from holding a risk-free bond. The household's utility function is given by:

$$\mathbb{E}_t \Sigma_{s=0}^{\infty} \beta^j \left[\log(C_{t+s}) - \frac{\omega}{1+\nu} \int_0^1 L_{t+s}(j)^{1+\nu} dj + \xi_t \frac{B_{t+1}}{P_t} \right]$$

where $\nu > 0$ is the inverse Frisch elasticity of labor supply, $\omega > 0$ is a parameter that pins down the steady-state level of hours, and the discount factor β satisfies $0 < \beta < 1$.

We use this particular specification of the household utility function augmented with a preference for holding risk-free bonds in order to introduce the liquidity demand shock, ξ_t. Fisher (2015) models this shock

[6] Our results do not depend on assuming this specific form of nominal rigidity. Results apply with a quadratic adjustment costs or menu cost frictions that generate money non-neutrality.

as a micro-foundation for the risk-premia shock considered by Smets and Wouters (2007). The primary reason we prefer to model this shock, as we prove shortly in Proposition 2, is that ξ_t allows us to maintain *divine coincidence* (Blanchard and Galí, 2007).[7] That is, a central bank following an optimal policy rule does not face a tradeoff in stabilizing fluctuations in output and inflation arising from such shocks. This shock is an example of a *purely intertemporal* shock considered by Eggertsson (2008), and we think using it is desirable in order to generalize the study of optimal monetary policy in an environment with endogenous TFP growth.

Labor income, $W_t L_t$, is subsidized at a fixed rate, τ^w. Households own an equal share of all firms, and thus receive Γ_t dividends from profits, and pay taxes τ^b on their income from risk-free bonds. Finally, each household receives a lump-sum government transfer, T_t. A household's budget constraint in period t states that consumption expenditure plus asset accumulation must equal the household's disposable income:

$$P_t C_t + B_{t+1} = (1-\tau^b)B_t(1+i_t) + (1+\tau^w)W_t L_t + \Gamma_t + T_t. \tag{1}$$

Utility maximization delivers the first-order condition linking the household's inter-temporal consumption smoothing to the marginal utility of holding the risk-free bond:

$$1 = \beta \mathbb{E}_t \left[\frac{C_{t+1}^{-1}}{C_t^{-1}} (1+i_t) \frac{P_t}{P_{t+1}} (1-\tau^b) \right] + \xi_t C_t. \tag{2}$$

The stochastic discount factor by which financial markets discount nominal income in period $t+1$ is given by:

$$Q_{t,t+1} = \beta \frac{C_{t+1}^{-1}}{C_t^{-1}} \frac{P_t}{P_{t+1}}.$$

The household does not directly choose its hours. Rather, each type of worker is represented by a wage union that sets wages on a staggered basis. Consequently, the household supplies labor at the posted wages

[7] We assume that the household cannot issue any risk-free debt, B_{t+1}. Anzoategui et al. (2019) also use the same specification for the demand shock because this shock induces a co-movement in investment and consumption. This is a relevant feature for our setting. This shock also has a standard interpretation of a shock to money in the household's utility function if the central bank paid interest on reserves. In Section 4, we show the results for standard preference shocks to the household's utility as employed in Eggertsson (2008). Alternately, we could have introduced these shocks through the budget constraint of the household. Amano and Shukayev (2012) show that such shocks are important ingredients for building models with binding ZLB. We prefer introducing them as shocks to the "wealth in the utility" function. Intrinsic desirability for wealth is a conventional assumption. See, for instance, Michaillat and Saez (2014) for more references.

as demanded by firms.

2.1.2 Wage Setting

Wage setting follows the modeling of Erceg, Henderson, and Levin (2000). Perfectly competitive labor agencies combine j-type labor services into a homogeneous labor composite, L_t, according to a Dixit-Stiglitz aggregator:

$$L_t = \left[\int_0^1 L_t(j)^{\frac{1}{1+\lambda_{w,t}}} dj\right]^{1+\lambda_{w,t}},$$

where $\lambda_{w,t} > 0$ is the (time-varying) nominal wage markup. Labor unions representing workers of type-j set wages (with indexation) on a staggered basis following Calvo (1983), taking as given the demand for their specific labor input:

$$L_t(j) = \left(\frac{W_t(j)}{W_t}\right)^{-\frac{1+\lambda_{w,t}}{\lambda_{w,t}}} L_t, \qquad \text{where } W_t = \left[\int_0^1 W_t(j)^{\frac{-1}{\lambda_{w,t}}} dj\right]^{-\lambda_{w,t}}.$$

In particular, with probability $1-\theta_w$, the type-j union is allowed to re-optimize its wage contract and it chooses W_t^* to minimize the disutility of working for a type-j laborer, taking into account the probability that the union will not get to reset wage in the future. If a union is not allowed to optimize its wage rate, it indexes wages to the steady state wage inflation, $\bar{\Pi}^w$. Workers supply whatever amount of labor is demanded at the posted wage. The first-order condition for this problem is given by:

$$\mathbb{E}_t \sum_{s=t}^{\infty} (\beta\theta_w)^{s-t} C_s^{-1} \left[\frac{W_t^*(j)(\bar{\Pi}^w)^{s-t}}{P_s} - (1+\lambda_{w,t})\omega L_s^{\nu}(j) C_s\right] L_s(j) = 0. \tag{3}$$

By the law of large numbers, the probability of resetting the nominal wage corresponds to the fraction of types that actually change their wage. Consequently, the nominal wage evolves according to:

$$W_t^{\frac{-1}{\lambda_{w,t}}} = (1-\theta_w)\, W_t^{*\,\frac{-1}{\lambda_{w,t}}} + \theta_w (W_{t-1}\bar{\Pi}^w)^{\frac{-1}{\lambda_{w,t}}}. \tag{4}$$

2.2 Firms

2.2.1 Final Good Producer

Households consume the *final* good, which is produced by perfectly competitive firms. These firms use an identical production technology employing a homogeneous labor composite supplied by the wage union and a CES composite of intermediate goods weighted by their productivity:[8]

$$Y_t^G = M_t^{1-\alpha} L_t^{1-\alpha} \int_0^1 A_{it}^{1-\alpha} x_{it}^{\alpha} di, \tag{5}$$

where each x_{it} is the flow of intermediate product i used at time t, the productivity parameter; A_{it} reflects the quality of that product; and M_t is the stationary (aggregate) productivity shock.

The firms choose L_t and $\{x_{it}\}_{i\in[0,1]}$ to maximize profits, taking as given both the wage index W_t and the prices of the intermediate goods, $\{p_{it}\}_{i\in[0,1]}$. The inverse demand for the labor composite and the intermediate good i are given by the following first-order conditions:

$$\begin{aligned} \frac{W_t}{P_t} &= (1-\alpha) M_t^{1-\alpha} L_t^{-\alpha} \int_0^1 A_{it}^{1-\alpha} x_{it}^{\alpha} di \\ \frac{p_{it}}{P_t} &= \alpha M_t^{1-\alpha} L_t^{1-\alpha} A_{it}^{1-\alpha} x_{it}^{\alpha-1}. \end{aligned} \tag{6}$$

2.2.2 Intermediate Goods Producer

There is a continuum of intermediate products indexed by $i \in [0,1]$, each of which is produced by an index-specific monopolist. The monopolist uses one unit of the final good to produce one unit of her own good. As a result, every monopolist faces a marginal cost of P_t. Each intermediate monopolist sets prices flexibly to maximize her firm's profits, taking as given the final sector's demand for its product. In particular, she solves for

$$\max_{p_{it}} (1-\tau^p) p_{it} x_{it} - P_t x_{it} \quad s.t. \quad \text{inverse demand in equation 6}, \tag{7}$$

[8] We denote gross output by Y_t^G, to keep it distinct from Y_t (defined shortly after), which we refer to as the GDP analog in our model.

where τ^p is a sales tax/subsidy imposed on the monopoly price. Further, we assume that there is a competitive fringe in every sector who can produce the intermediate good with quality $\frac{A_{it}}{\gamma}$, where $\gamma > 1$ is the step-size of innovation and captures the quality distance between the frontier and laggard firms within a sector. As a result, the intermediate monopolist cannot charge a price higher than $p_{it} = \gamma^{1-\alpha} P_t$. In equilibrium, the monopolist charges a price given by:

$$p_{it} = \zeta P_t \equiv \min \left(\gamma^{1-\alpha}, \frac{1}{(1-\tau^p)\alpha} \right) P_t.$$

The linearity in the use of rival goods in the final good's production function is an important ingredient here. It makes an intermediate firm's profits linear in the labor demanded by the firm producing the final good and the productivity of the firm producing the intermediate good.[9] Increasing its productivity enables the intermediate firm to capture a larger share of the demand for the final good. Profits are given by:

$$\begin{aligned} \Gamma_t(A_{it}) &= \chi^m P_t M_t L_t A_{it} \\ \text{where } \chi^m &= (\zeta - 1) \left(\frac{\alpha}{\zeta} \right)^{\frac{1}{1-\alpha}} . \end{aligned} \tag{8}$$

2.2.3 Entrepreneurs

There is a single entrepreneur in each sector who invests $RD(z_{it})A_{it}$ of the final good in research and development in period t, where $RD' > 0$, $RD'' > 0$.[10] The dependence on productivity A_{it} is assumed for stationarity. With probability z_{it}, the entrepreneur is successful in making a process improvement. The productivity in sector i goes up by a factor of $\gamma > 1$ (step size of innovation) and she gets the monopoly rights (*patent*) over production of the intermediate good in the *following* period. If she fails to innovate, the incumbent monopolist continues to produce with productivity A_{it} until replaced by a successful entrant.

[9] This linear dependence on productivity is central to endogenous growth models. Jones (2005, Sec 6.2) formalizes this argument as "*any* model of sustained exponential growth requires that a particular differential equation is linear in some sense".

[10] We follow Aghion, Akcigit, and Howitt (2014) in this discrete time analog of their classic Schumpeterian model, but extend it to allow for a more general innovation production function that allows decreasing returns to R&D. Benigno and Fornaro (2018) use a similar model but with $RD'' = 0$. Assuming $RD'' > 0$ introduces decreasing returns to innovation, which is a feature stressed regularly in the innovation literature. As we will argue in Section 6, this curvature is a crucial parameter that regulates the quantitative implications of several endogenous growth models.

Hence,

$$A_{it+1} = \begin{cases} \gamma A_{it} \text{ with probability } z_{it} \\ \\ A_{it} \text{ otherwise.} \end{cases} \tag{9}$$

The cost of research is increasing according to the innovation intensity chosen by the entrepreneur and the existing level of technology in the intermediate good's sector in which the entrepreneur operates. Specifically, we assume that $RD(z_{it}) = \delta z_{it}^{\varrho}$, where $\delta > 0$ and and $\varrho > 1$ is the inverse elasticity of innovation intensity to R&D expenses. τ^r is a research subsidy provided by the government to the entrepreneur. The entrepreneur in every sector chooses z_{it} to maximize her expected discounted profits (from the patent):

$$\max_{z_{it}\in[0,1]} \{z_{it}\mathbb{E}_t Q_{t,t+1} V_{t+1}(\gamma A_{it}) - (1-\tau^r) P_t RD(z_{it}) A_{it}\}, \tag{10}$$

where value of the patent is given by:

$$V_t = \Gamma_t + (1 - z_{it})\mathbb{E}_t Q_{t,t+1} V_{t+1}. \tag{11}$$

The value function is linear in productivity due to the linearity in the production function (see Appendix A). Writing the normalized value function as $\tilde{V}_{it} \equiv \frac{V_{it}}{P_t A_{it}}$ and focusing on the symmetric equilibrium, we solve for interior solution (where $z_t > 0$):

$$\varrho z_t^{\varrho-1} = \beta \mathbb{E}_t \left(\frac{C_{t+1}}{C_t}\right)^{-1} \frac{\gamma \tilde{V}_{t+1}}{(1-\tau_t^r)\delta}. \tag{12}$$

According to equation (12), the entrepreneur chooses the innovation intensity so that the discounted marginal revenue of an additional unit of innovation intensity is equal to the marginal cost of this unit. An increase in demand for the final good increases the value of obtaining the patent. This is because of the *market size effect*—for a given cross-sectional distribution of productivities, an increase in the demand for the final good requires higher quantities of the intermediate goods to fulfill that demand. Since a monopolist's profits are increasing in the quality of its product, she can capture a higher share of the increased market with a successful innovation.

2.3 Aggregation and Market Clearing

The aggregate behavior of the economy depends on the aggregate (which also corresponds to the average in this case) productivity index, defined as:

$$A_t = \int_0^1 A_{it} di. \tag{13}$$

Because of the linear production function, we can aggregate the firm-level variables to form aggregate composites. Specifically, $RD_t = \int RD_{it} di$ is the total R&D expenditure, and $X_t = \int X_{it} di$ is the aggregate intermediate good produced in the economy. We can rewrite the aggregate output and nominal wage purely in the form of aggregates:

$$Y_t^G = \left(\frac{\alpha}{\zeta}\right)^{\frac{\alpha}{1-\alpha}} M_t L_t A_t \tag{14}$$

$$W_t = (1-\alpha)\left(\frac{\alpha}{\zeta}\right)^{\frac{\alpha}{1-\alpha}} M_t A_t P_t. \tag{15}$$

The growth rate of output in the economy is equal to the growth rate of aggregate productivity:

$$g_{t+1} = \frac{A_{t+1} - A_t}{A_t}. \tag{16}$$

In any period, innovations occur in z_t sectors, and $1 - z_t$ sectors use the previous period's production technology. Aggregating across all the sectors, we get the following equation governing the dynamics of aggregate productivity:

$$A_{t+1} = \int_0^1 [z_t \gamma A_{it} + (1 - z_t) A_{it}] di = A_t + z_t(\gamma - 1)A_t. \tag{17}$$

This means that the growth rate of the economy in period $t + 1$ is determined in period t and equals the number of innovating sectors times the step-size of innovation:

$$g_{t+1} = z_t (\gamma - 1). \tag{18}$$

The number of innovating sectors z_t may be interpreted as new entrants, since the incumbent firms do not undertake R&D investment in our model. The final output produced in the economy is used for consumption, research, and the production of intermediate goods:

$$Y_t^G = C_t + RD_t + X_t. \tag{19}$$

Henceforth, we define $Y_t^G - X_t = (1 - \frac{\alpha}{\zeta})Y_t^G \equiv Y_t$ as GDP.

2.4 Fiscal and Monetary Policy

To close the model, we assume a net zero supply of risk-free bonds:[11]

$$B_t = 0.$$

The government's budget is balanced every period, so total lump-sum transfer payments are equal to taxes on intermediat goods, labor, and R&D:

$$P_t T_t = \tau^p \int_0^1 p_{it} x_{it} di + \tau^r P_t RD_t + \tau^w \int_0^1 W_t(h) L_t(h) dh. \tag{20}$$

An independent central bank follows a Taylor rule in setting the nominal interest rate in the economy:

$$1 + i_t = \max\left(1, (1 + i_{ss}) \left(\frac{\Pi_{W,t}}{\bar{\Pi}_W}\right)^{\phi_\pi} \left(\frac{L_t}{\bar{L}}\right)^{\phi_y} \varepsilon_t^i\right); \quad \phi_\pi > 1, \phi_y \geq 0. \tag{21}$$

The nominal interest rate is set in order to target deviations of wage inflation and employment at respective steady state targets, as long as the implied nominal interest rate is non-negative. ε_t^i is defined as a monetary policy shock.[12]

[11] To keep the analysis focused on optimal monetary policy, we do not allow government debt to offset the increased demand for liquidity. In a recent work, Del Negro, Eggertsson, Ferrero, and Kiyotaki (2017) assess the importance of government liquidity facilities during the Great Recession in stabilizing the economy. Similar analysis under endogenous growth and sticky prices, following the quantitative model of Guerron-Quintana and Jinnai (2019), is left for future work.

[12] From equation 14, we wish to emphasize that this rule is analogous to the Taylor rule used to represent the monetary policy response in an exogenous growth model. Once we normalize Y_t by the level of productivity A_t, there is a one-to-one mapping between employment and normalized output. In the presence of liquidity demand and monetary policy shocks, rules targeting employment or normalized output are identical.

2.5 Equilibrium

We formally define the economy's competitive equilibrium in Appendix A. In order to arrive at a stationary system of equations, we normalize the equilibrium equations by dividing the non-stationary variables such as consumption, output, and the real wage by the level of productivity. We define $c_t = \frac{C_t}{A_t}$ as the normalized (productivity-adjusted) consumption and so forth. This allows us to solve for the steady state.

We find the steady state by imposing restrictions on the parameters such that the steady state satisfies a) $z \in (0,1)$, b) consumption is non-negative and c) nominal interest rates are non-negative.[13] In Appendix B, we analytically characterize the existence and uniqueness of the steady state by imposing additional assumptions on the nature of wage rigidity and the duration of patents granted to the monopolist.

Steady State Efficiency

Because of the presence of research externalities and monopoly distortions, the private sector equilibrium is inefficient. We define the *efficient* steady state as the one in which the welfare of the representative household is maximized subject to the production technology of the final consumption good (equation 14), the law of motion of aggregate productivity (equation 17), and the economy's resource constraint (equation 19) for a given initial productivity level. The complete system of equations is provided in Appendix D.

Proposition 1 states that the steady state of the competitive equilibrium allocation is inefficient. This is due to the presence of three *static* distortions in our setup: (i) monopoly power in each intermediate goods sector, (ii) monopolistic competition in the labor market, and (iii) inter-temporal research externalities. Whereas the first two distortions are common in the business cycle literature, the third distortion is specific to the endogenous growth literature. The entrepreneur is unable to reap all the benefits of her technology advance because she gets replaced with positive probability by a new entrant (*surplus appropriability effect*). This makes her under-invest in R&D. On the other hand, an entrant replaces the incumbent to profit from the full step size of innovation, γ, rather than the incremental gain in knowledge, $\gamma - 1$. This *business stealing effect* (Aghion and Howitt, 1992) incentivizes the entrepreneur to over-invest in R&D. As a result of these two opposing forces, private investment in research can be higher or lower than the first-best allocation.

[13] In our numerical simulations, we verify that innovation probability is bounded, $z_t \in (0,1)$.

We assume that the fiscal authority has access to lump-sum taxes, and so the first-best allocation in the steady state can be implemented by a set of constant taxes elaborated in the following proposition:

Proposition 1 (Steady State Efficiency)**.** *Assuming the policymaker has access to non-distortionary lump-sum taxes, the steady state of the competitive equilibrium can be made efficient using the following three taxes:*

a) sales subsidy, $\tau^p = 1 - \frac{1}{\alpha}$

b) labor tax, $\tau^w = \frac{1-\lambda_w}{\lambda_w}$ *, and*

c) research tax, $\tau^r = 1 - \left[\left(\frac{\gamma l^*(1-\alpha)\alpha^{\frac{\alpha}{1-\alpha}}}{1-\beta(1-z^*)}\right)\left(\frac{1-\beta}{(\gamma-1)c^*}\right)\right]$*, where terms with* $*$ *denote the efficient steady state values.*

Proof. See Appendix E. □

It is commonly argued in the endogenous growth literature that the private sector under-invests in R&D (Jones and Williams 1998), and therefore the productivity growth rate is higher in the efficient steady state. These distortions would imply that in the absence of relevant fiscal instruments, monetary policy could affect the growth rate of output in the long run. We follow the monetary economics literature and suppose that the average productivity growth rate is optimal and independent of monetary policy. As shown by Woodford (2003) and Benigno and Woodford (2004), the linear-quadratic approximation to the social welfare function around the non-stochastic efficient variables is justified if there are no distortions under price stability. In the parlance of the literature, there are no permanent differences between the efficient and the natural rate of interest. Here the idea is to disassociate the welfare losses from fluctuations in the growth rate from those losses arising from suboptimal growth solely due to monopoly distortions and research externalities. We make the following assumption:

Assumption 1. *The fiscal authority provides the set of constant subsidies described in Proposition 1, such that the competitive equilibrium is efficient in the steady state.*

The crucial difference to note from the earlier monetary economics literature is that *monetary policy in our setting has a bearing on the long-run level of output* even though we do not allow monetary policy to influence the steady state distortions. We log-linearize the competitive equilibrium around the efficient-steady state and define the following approximate equilibrium:[14]

Definition 2.1 (Approximate Equilibrium)**.** The approximate competitive equilibrium in this economy with endogenous growth is defined as a sequence of variables $\{\hat{\pi}_t^w,\, \hat{c}_t,\, \hat{y}_t,\, \hat{g}_{t+1},\, \hat{i}_t,\, \hat{L}_t,\, \hat{w}_t,\, \hat{\pi}_t,\, \hat{V}_t\}$ which satisfy the following equations, for a given sequence of exogenous shocks $\{\hat{\xi}_t,\, \hat{M}_t,\, \hat{\epsilon}_t^i,\, \hat{\lambda}_{wt}\}$.[15]

Aggregate Demand Equation:

$$-\left(\mathbb{E}_t\hat{c}_{t+1} - \hat{c}_t + \hat{g}_{t+1}\right) + \hat{i}_t - \mathbb{E}_t\hat{\pi}_{t+1} + \hat{\xi}_t = 0. \tag{22}$$

Endogenous Growth Equations:

$$(\varrho - 1)\eta_g\hat{g}_{t+1} = -(\mathbb{E}_t\hat{c}_{t+1} - \hat{c}_t + \hat{g}_{t+1}) + \mathbb{E}_t\hat{V}_{t+1} \tag{23}$$

$$\hat{V}_t = \eta_y\hat{y}_t - \eta_z\hat{g}_{t+1} - \eta_q(\mathbb{E}_t\hat{c}_{t+1} - \hat{c}_t + \hat{g}_{t+1}) + \eta_q\mathbb{E}_t\hat{V}_{t+1}, \tag{24}$$

where $\eta_g = \frac{1+g}{g} > 1$, $\eta_y = 1 - \frac{(1-z)\beta}{1+g} > 0$, $\eta_z = \frac{\beta}{\gamma-1} > 0$, $\eta_q = \frac{(1-z)\beta}{1+g} > 0$.

Market Clearing Equations:

$$\frac{c}{y}\hat{c}_t + \frac{\mathbb{R}}{Y}\varrho\eta_g\hat{g}_{t+1} = \hat{y}_t \tag{25}$$

$$\hat{y}_t = \hat{M}_t + \hat{L}_t. \tag{26}$$

Wage Setting Equations:

$$\hat{\pi}_t^w = \tilde{\beta}\mathbb{E}_t\hat{\pi}_{t+1}^w + \kappa_w[\hat{c}_t + \nu\hat{L}_t - \hat{w}_t] + \kappa_w\hat{\lambda}_{wt} \tag{27}$$

$$\hat{w}_t = \hat{M}_t \tag{28}$$

$$\hat{\pi}_t^w = \hat{w}_t - \hat{w}_{t-1} + \hat{\pi}_t + \hat{g}_t, \tag{29}$$

where $\kappa_w \equiv \frac{(1-\theta_w)(1-\beta\theta_w)}{\theta_w(1+\nu\left(1+\frac{1}{\lambda_w}\right))} > 0$.

Monetary Policy Rule:

$$\hat{i}_t = \max\left(-\frac{\bar{i}}{1+\bar{i}}, \phi_\pi\hat{\pi}_t^w + \phi_y\hat{L}_t + \hat{\varepsilon}_t^i\right). \tag{30}$$

The equations for aggregate demand, market clearing, and the wage Phillips curve are familiar to scholars

[14] In the appendix, we show the necessary condition for local determinacy in the model with perfect wage rigidity and a single period patent granted to the monopolist. This derivation generalizes the local determinacy condition derived by Benigno and Fornaro (2018) to include diminishing returns to R&D investment. However, a closed form representation of determinacy is not possible in our benchmark setup with Calvo wage rigidities and stochastic patent duration. With this caveat in mind, we proceed by assuming local determinacy to characterize theoretical implications for policy. In all numerical simulations, we verify the local determinacy of the efficient steady state.

[15] For any variable x, $\hat{x}_t = log\left(\frac{x_t}{\bar{x}}\right)$, where $\bar{x}$ is the efficient/non-distortionary steady state. With few exceptions: $\hat{g}_{t+1}$ is the deviation of gross growth rate from the steady state value that is $\hat{g}_{t+1} \equiv \log\left(\frac{1+g_{t+1}}{1+\bar{g}}\right)$. We verify that the growth rate is always non-negative in numerical simulations. The log-linearized representation of liquidity demand shock $\hat{\xi}_t \equiv cA_t\xi_t$ since the steady state value of the shock ξ is assumed to be 0. $\hat{\xi}_t$ is scaled-version of the non-linear liquidity demand shock, where the scaling variable is the balanced growth level of consumption under flexible wages. AR(1) shock to ξ_t maps one-to-one to a shock to $\hat{\xi}_t$.

of the NK business cycle literature. The new ingredient is the endogenous growth block (equations 23–24). It is a log-linear transformation of the entrepreneur's profit-maximization condition.

The endogenous growth condition (equation 23) states that the entrepreneur makes her R&D investment decision based on the expected present discounted value of the future profits. Thus there are two forces governing her decision: the rate at which she discounts the future, and the expected value of future profits. In our model, the household sector retains ownership over all firms. Therefore, the rate at which a firm discounts the future is given by the household's stochastic discount factor. A higher stochastic discount factor increases the entrepreneur's incentive to innovate (*discounting effect*) because of the lower discounting of future profits. Second, higher expected future output increases her incentive to invest in innovation because of the *market size effect*, discussed above. Furthermore, a 1 percentage point change in innovation investment translates into $\frac{1}{\varrho \eta_g}$ percentage point change in the gross productivity growth rate, where $\frac{1}{\varrho}$ is the elasticity of innovation intensity, and ϱ is assumed to be greater than 1 following the innovation literature (see Acemoglu and Akcigit 2012). This implies decreasing returns to investment in innovation—a higher value of ϱ signifies the lower responsiveness of innovation success (and productivity growth rate as a result) to innovation investment.

Equilibrium Concepts and Policy Instruments

We now provide a brief discussion of the natural rate allocation, the first-best allocation, and the pre-recession trend allocation under the endogenous growth setting. Importantly, we underscore the subtleties introduced in the endogenous growth setup relative to the exogenous growth setting. This allows us to formally define potential output, which becomes an endogenous object under endogenous growth.

We assume that the (normalized) economy is in the efficient steady state at beginning of time $t = 0$. The *first-best allocation* is the competitive equilibrium allocation under flexible wages such that the fiscal authority utilizes (non-distortionary) time-varying taxes in order to maximize the representative agent's welfare. The *natural-rate* allocation (or interchangeably, the flexible-wage allocation) is the competitive equilibrium allocation under flexible wages such that the fiscal authority provides (non-distortionary) constant tax instruments, as outlined in Proposition 1.

Under liquidity demand and monetary policy shocks, the natural-rate and the first-best allocations coincide. This coincidence implies that the natural rate of interest, *r-star*, is exogenous even in the presence of endogenous growth. Moreover, these shocks do not affect the flexible-wage equilibrium output, consumption, and investment in R&D. The economy stays on the initial balanced growth path (BGP), also referred to as the *pre-recession trend*. Thus, any change in output in the sticky-wage economy emerges solely because of nominal (wage) rigidities. This property helps to isolate the role of monetary policy. For the sake of transparency, our paper's main analysis is conducted with these two shocks. We provide a brief discussion of discount rate shocks, cost push shocks, and stationary TFP shocks in Section 4.[16]

Whether potential output is endogenous or not depends on the precise definition, of price/wage *flexibility* in the presence of a pre-determined state variable. One is the Neiss and Nelson (2003) definition of flexible wages, under which wages have been set flexibly since time zero and remain flexible indefinitely. Wages set under this concept are called time-zero flexible wages. The second concept of flexibility is the definition in Woodford (2003), where wages are set flexibly in the current and future periods, taking as given the evolution of the state variable. Wages set under this concept are called time-t flexible wages. Based on these two concepts of flexible wages, there are *time-zero first-best*, *time-zero natural rate*, *time-t first-best* and *time-t natural rate* allocations. We provide formal definitions in Appendix D.9. To avoid notational clutter, henceforth we will use the *first-best* allocation for the time-zero first-best allocation and *natural rate* for the time-zero natural rate allocation whenever possible without introducing ambiguity. For the ease of exposition, we refer to time-zero flexible wages as *flexible wages*. We define *potential output* as the level of output that coincides with the time-t first-best allocation. We believe this is more appealing than the definition based on the time-zero concept because the time-t concept coincides with the maximum non-inflationary output an economy can produce at a given time.

We emphasize that the distinction between the two natural rate concepts defined here is different from the distinction imposed in exogenous growth environments with capital investment (Edge 2003). In our benchmark endogenous growth model, the natural rate of interest is always same under the two concepts of

[16]The distinction between the natural-rate, the first-best, and the pre-recession trend allocations will become crucial in Section 4. This is because under discount rate and supply shocks, there can be a divergence among these concepts. Thereafter, in order to provide a data counterpart for output hysteresis, we define it as the deviation of output under a competitive equilibrium allocation from the pre-recession trend. Under liquidity demand and monetary policy shocks, it does not matter whether the output hysteresis is defined as deviation from the (time-zero) first-best, the (time-zero) natural rate. or the pre-recession trend output.

flexibility. Only the levels of productivity and output differ. Importantly, this difference in levels may be permanent, depending on the central bank's policy rule. In contrast, the introduction of capital investment introduces a temporary difference in the levels of capital and output, as well as in the interest rate, depending on the nature of flexibility assumed. In those models, there is no medium- or long-run difference between the various concepts, as capital always returns to its initial steady state value. Hence, in contrast to the setups in Neiss and Nelson (2003) and Woodford (2003), the economy's potential output is an endogenous object even in the long run, as will become clear in the next section.

Sticky wage allocation is the equilibrium allocation under staggered (nominal) wages such that the fiscal authority provides (non-distortionary) constant tax instruments, as outlined in Proposition 1. We refer the reader to Appendices D.9.1, D.9.2, and D.9.3 for a formal definition of these equilibria concepts.

Proposition 2. *The (time-zero) natural rate allocation coincides with the (time-zero) first-best allocation under liquidity demand and monetary policy shocks.*

Proof. See Appendix E. □

Proposition 2 implies that the representative agent's welfare is maximized if the policymaker could replicate the natural rate allocation. This outcome is always possible if the policymaker has access to time-varying tax instruments (see for example Correia, Nicolini, and Teles 2008, and Correia, Farhi, Nicolini, and Teles 2013). In Appendix D.6, we illustrate how the first-best allocation can be implemented by using appropriate state-contingent fiscal instruments, even at the ZLB. Henceforth, we assume that the policymaker does not have access to these time-varying fiscal instruments: the fiscal authority satisfies Assumption 1, and adjusts lump-sum taxes every period to balance the budget. The central bank sets the nominal interest rate i_t on the risk-free (nominal) bond B_t subject to the ZLB constraint:

$$i_t \geq 0 \quad \forall t. \tag{31}$$

The nominal interest rate is the central bank's only policy instrument.

Calibration and Impulse Responses

Our approximate equilibrium is linearized around a locally determinate steady state. We can analytically solve for the impulse responses under the assumption of an AR(1) process for shocks. The exact solution is provided in Appendix C for the case of liquidity demand shocks. However, in order to illustrate the dynamics for the benchmark model, we calibrate the model with parameters reported in Table 1, with quarterly time periods. There are eight parameters–we calibrate five of these using values standard in the NK literature. The discount factor β equals 0.99. The labor share $1-\alpha$ is set to 0.67. Preferences are logarithmic in consumption and the inverse Frisch elasticity ν is set at 2. The wage adjustment probability is set such that wages are reset once every four quarters and the steady state wage markup is 10%. Monetary policy is assumed to follow a standard Taylor rule (equation 30) with $\phi_\pi = 1.5$ and $\phi_y = 0.5$.

Table 1: Parameters for Welfare Analysis in the Simple Benchmark Model

Standard Parameters	*Formula*	*Value*
Labor Share	$1-\alpha$	0.67
Discount Rate	β	0.99
Steady State Wage Markup	λ_w	0.10
Calvo Probability of Wage Adjustment	$(1-\theta_w)$	$1-0.75$
Inverse Frisch Elasticity	ν	2
Innovation Step Size	γ	1.20
Inverse Innovation Elasticity	ϱ	1.47
Innovation Cost Parameter	δ	22.6

We choose the remaining three innovation parameters: step size of innovation γ, the (inverse of the) innovation elasticity ϱ, and cost parameter in R&D investment δ such that the model replicates an annual steady state growth rate of 2%, annual firm entry rate of 10%, and an R&D-to-GDP ratio of 25%.[17] In the data, the private R&D-to-GDP ratio is 2% (NIPA 1953–2007). We do not have a data counterpart to this ratio under an efficient steady state. As shown in Proposition 1, the steady state in the private sector equilibrium may feature over-/under-investment in R&D because of the research externalities. Jones and Williams (1998) estimated that the social return on R&D investment is at least four times more than the private return on R&D. In a more general setup than ours, they find that forces, such as the business-stealing effect, are quantitatively dominated by a loss in profits to innovating firms due to technological spillovers to other firms. Because agents cannot appropriate the full extent of profits from innovation, they under-invest

[17]The probability of innovation success is interpreted as the firm entry rate, consistent with the "creative destruction" literature.

in R&D. We choose a higher value for the R&D-to-GDP ratio for the ease of illustrating our results. We show that our results are robust to choosing a grid of R&D-to-GDP ratio in the range of 10% –30% later in Section 3.3.[18] In Section 4.5, we show a quantitatively realistic calibration of model, away from an efficient steady state, that can replicate key variable moments in the data.

Figure 2: Model-Based Impulse Response Functions

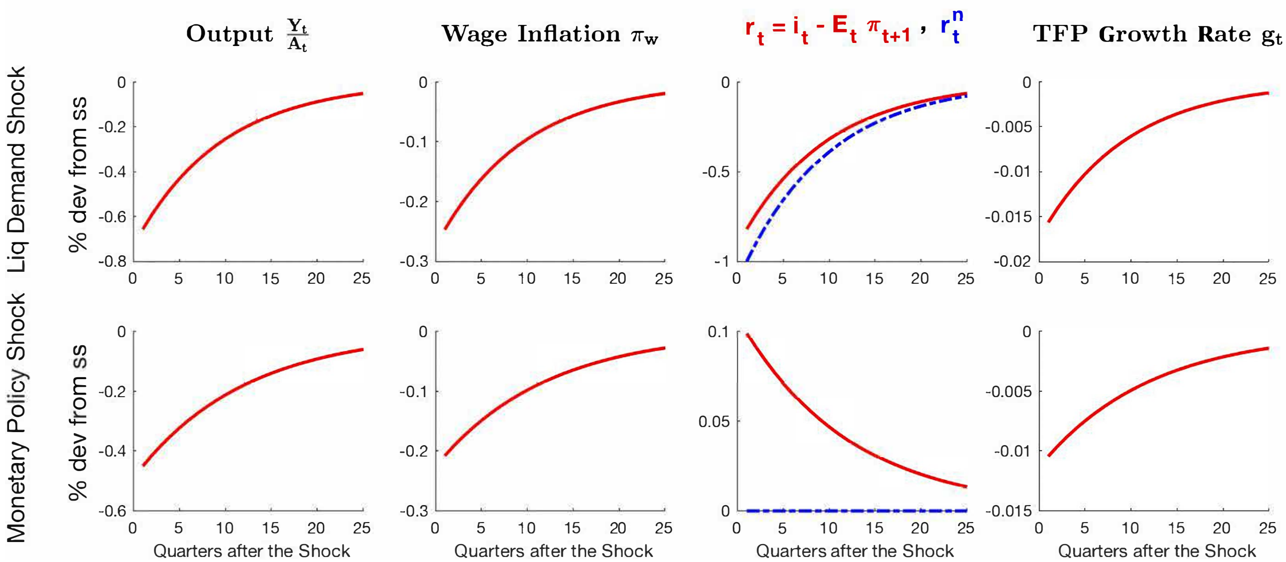

Source: Authors' calculations.
Note: The figures illustrate the impulse response functions (IRFs) from the benchmark model presented in Section 2. The IRFs are plotted in response to liquidity demand shock, and monetary policy shock, with persistence 0.9 and 0.92, respectively. Also, ss = steady state.

Rows 1 and 2 in Figure 2 plot the impulse responses for normalized output, wage inflation, the real interest rate, and the productivity growth rate for a positive shock to liquidity demand, ξ_t, and a contractionary monetary policy shock, ε_t, each following an AR(1) process with persistence of 0.90 and 0.92, respectively. A positive liquidity demand shock corresponds to a fall in the annualized *natural interest rate* of 1 percentage point. It increases the household's desire for saving in the risk-free bond and thus diverts the resources away from consumption. The lower anticipated aggregate demand reduces investment in R&D by entrepreneurs, exerting a drag on productivity growth. Furthermore, a positive liquidity demand shock reduces the household's stochastic discount factor, for a given nominal interest rate. This is equivalent to an increase in the entrepreneur's "borrowing cost" for investment in innovation. These two forces act in the same direction to reduce investment in innovation. Hence, the productivity growth rate is lower following a contraction in demand induced by the liquidity demand shock.

[18]See Table 3.

Similarly, a surprise contractionary monetary policy shock (annualized 68 basis points) implies an increase of 10 basis points in the real interest rate and tends to lower the nominal wage. Due to the stickiness of nominal wages, aggregate demand adjusts downwards. The equilibrium increase in the real interest rate, combined with expectations of lower future aggregate demand, leads to a reduction in R&D investment and, therefore, in TFP growth.

3 Normative Implications for the Conduct of Monetary Policy

Next, we analyze the *normative* implications for the conduct of monetary policy. In order to do so, we derive a quadratic approximation of the household's welfare function. We use this result to analyze optimal monetary policy in response to aggregate demand shocks. We highlight three results. One, away from the ZLB, optimal policy is equivalent to a strict inflation targeting rule, and does not involve permanent shifts in output. Two, at the ZLB, optimal policy commits to keeping interest rates lower in the future. Such a policy returns the economy *close* to the pre-shock trend. Three, a discretionary policy (time-consistent policy) at the ZLB involves excessive output hysteresis relative to the commitment policy. We label this as the *hysteresis bias* of discretionary policy. This result implies that output hysteresis is an artifact of policy constraints faced by the central bank rather than an outcome of irrational or inept behavior on part of the central bank. The central bank's lack of credibility tools is sufficient to generate hysteresis. Numerically, we show that a new strict hysteresis targeting policy closely replicates optimal monetary policy, thereby implying significant welfare gains over discretionary policy. This is true for a range of values for the key parameter. ϱ, which determines the innovation sensitivity to investment fluctuations.

3.1 Quadratic Approximation of Welfare

One primary contribution of our paper is that we derive a purely quadratic approximation of welfare for the representative household under endogenous growth. This expression can enable researchers to solve for optimal policy in a tractable manner. It generalizes the quadratic objective derived by Benigno and Woodford (2004) to an endogenous growth setting.

Proposition 3. *Assume that the economy is at the efficient steady state at time $t = 0$, with given productivity level A_0. Under the sticky wage allocation, the quadratic approximation of representative agent's lifetime utility function $\mathbb{W}_0$ around the non-stochastic efficient steady state is given by*

$$\frac{\mathbb{W}_0 - \mathbb{W}_0^*}{U_{c_{ss}} y_{ss}}$$
$$= -\frac{1}{2} \sum_{t=0}^{\infty} \beta^t \left[\lambda_y \underbrace{\left(\hat{y}_t - \frac{\beta}{1-\beta} \frac{1}{\nu + \frac{y}{c}} \hat{g}_{t+1} \right)^2}_{(i)} + \lambda_g \underbrace{\hat{g}_{t+1}^2}_{(ii)} + \lambda_\pi \underbrace{(\hat{\pi}_t^w)^2}_{(iii)} \right] + \mathcal{O}(||\hat{\xi}_t, \hat{\varepsilon}_t^i||^3) + t.i.p. \qquad (32)$$

(i) : labor efficiency gap, (ii): productivity growth rate gap, and (iii): wage inflation gap,

where $\lambda_y = \left(\nu + \frac{y}{c}\right) > 0$, $\lambda_g = \frac{c}{y} \frac{\beta}{1-\beta} \left[\frac{\nu}{\nu + \frac{y}{c}} \frac{\beta}{1-\beta} + [(\varrho - 1)\eta_g + 1] \right] > 0$, $\lambda_\pi = \frac{1+\lambda_w}{\lambda_w} \frac{1}{\kappa_w} > 0$, $\kappa_w \equiv \frac{(1-\theta_w)(1-\beta\theta_w)}{\theta_w(1+\nu(1+\frac{1}{\lambda_w}))} > 0$, $\eta_g = \frac{1+g}{g} > 1$ *and t.i.p. stands for "terms independent of policy".* $\mathbb{W}^*$ *denotes welfare under the (time-0) first-best allocation. The approximation is scaled by the constant* $U_{c_{ss}} y_{ss} = \frac{y_{ss}}{c_{ss}}$ *(evaluated at the efficient steady state).*

Proof. See Appendix E. □

This approximation is composed of three gaps/wedges–(i) the labor efficiency gap, (ii) the productivity growth rate gap, and (iii) the wage inflation gap. These are the three stabilization goals for a planner who wants to maximize social welfare. The first and the third terms are standard in a textbook NK model.

The first term, the labor efficiency gap, is the difference between the marginal product of labor and the marginal rate of substitution between consumption and leisure for the representative household,

$$(i) = mrs_t - mpn_t,$$

where these terms denote deviations from the respective steady state values. Since we do not model price setting frictions in this simple benchmark model, and do not consider price-markup shocks, mpn_t corresponds to the (productivity-adjusted) real wage. Thus the labor efficiency gap captures the time-varying wedge in the household's disutility from supplying labor at a pre-set nominal wage.

The third term, the wage inflation gap, describes the loss in efficiency resulting from the dispersion in wages across the members of the household. Wage dispersion, similar to price dispersion in standard NK models, is costly because firms hire a different amount of hours from various members of the household, causing the marginal disutility of labor to vary within the household. Under flexible wages, both the labor inefficiency gap and the wage inflation gap go to zero.

The second term, the *productivity growth rate gap*, is a key new ingredient of the endogenous growth model. Investment in R&D in a given period contributes to an increase in productivity, which persists into the indefinite future. These inter-temporal spillovers of R&D investment may not be internalized by the private agents and may result in the investment response being too high or too low relative to the first-best allocation. Starting from a productivity level A_0, the growth rate gap in equation (32) captures the suboptimality of deviations from the first-best level of productivity given by $A_t^* = A_0(1+g_{ss})^t$ at all times when $t > 0$. Under nominal (wage) rigidities, as discussed in previous section, demand shocks may induce this permanent gap, thus leaving the agent permanently worse off. This gap disappears under the exogenous growth assumption, and the quadratic approximation simplifies to the setting discussed in the graduate textbook treatment of Galí (2015, Ch. 4).

In Corollary 1, we show the conditions under which the welfare loss resulting from these productivity growth rate deviations is larger than the loss that arises due to changes in the labor efficiency gap. We provide a sufficient condition for the growth rate gap to be of higher importance for stabilization than the labor efficiency wedge. We argue that this condition is likely to be satisfied even for the extreme values of parameters considered in the literature. This condition highlights the importance of stabilizing the productivity growth rate around the first-best allocation.

Corollary 1 (Importance of Growth Stabilization)**.** *The relative weight on growth rate gap is higher than the relative weight on labor efficiency wedge if*

$$\frac{\beta}{1-\beta} > \frac{y}{c}\left(\nu + \frac{y}{c}\right). \tag{33}$$

Proof. See Appendix E. □

The common calibration values of the discount rate, β, at a quarterly frequency lie in the range of

$[0.98, 1)$. This implies a lower bound on the left-side of equation (33) at 49. We bound the right-side as follows: the consumption-to-output ratio in the US has fluctuated between 0.54 and 0.66 from 1960 to 2014 (BEA). Estimates of Frisch elasticity of labor, $1/\eta$, in the micro literature lie between 0.1 and 0.5 (Chetty et al. 2016) while the macro literature uses estimates in the range of (2,4). Using a value of 0.1 for η^{-1} and 0.54 for c/y ratio, this implies an upper bound on the right-side at 22. Hence, for a wide range of the parameter estimates used in the macroeconomics literature, the welfare loss from a given growth rate deviation is higher than the welfare loss from a similar change in the labor efficiency gap. Intuitively, a given deviation in the productivity growth rate from steady state has long run, potentially permanent, effects. On the other hand, fluctuations in labor efficiency pertain to welfare losses only in the period that these losses are encountered.

3.2 Away from the Zero Lower Bound

Optimal Policy Away from ZLB

We now turn to investigating the implications for the conduct of monetary policy in our model and show the main results outlined at the beginning of this section. First, we show that optimal policy involves setting the nominal interest rate in order to perfectly stabilize output and productivity along the first-best allocation.

Proposition 4 (Optimal Pilicy Away from the ZLB)**.** *Given a process for liquidity demand and monetary policy shocks, optimal monetary policy under a sticky wage allocation tracks the natural rate of interest when the zero lower bound constraint is slack.*

Proof. See Appendix E. □

From Proposition 2, we know that the flexible wage allocation coincides with the first-best allocation. Under a sticky wage allocation, setting the nominal interest rate to track the natural interest rate implements the flexible wage allocation, thereby replicating the first-best allocation. This result implies that the output follows a trend-stationary process since the normalized output and productivity growth rate are always at the steady state. Hence, the following corollary follows:

Corollary 2. *When the ZLB is slack, the time series of output under optimal policy is a trend stationary process (integrated of order zero); that is,*

$$\log Y_t = a + b * t,$$

where $a = \log Y_0$ is the initial level of output, and $b = \log(1 + g_{ss})$ is the steady state productivity growth rate.

Proof. See Appendix E. □

We established that permanent output gaps are an undesirable feature of the endogenous growth economy in response to temporary demand shocks. The optimal monetary policy does not allow for these hysteresis effects. Next, we characterize how policy rules used in the standard NK model fare in this environment.

Policy Rules Away from the ZLB

Assume that the central bank follows the Taylor rule shown in equation 30. Given local determinacy, we can derive the deviations in the level of productivity and output from the respective levels under flexible wages as:

$$\log A_t - \log A_t^e = \sum_{s=0}^{t-1} \psi_g^i \epsilon_s^i; \quad \log Y_t - \log Y_t^e = \hat{y}_t + \sum_{s=0}^{t-1} \psi_g^i \epsilon_s^i,$$

where $\psi_g^i > 0$ (the detailed expression is derived in Appendix C) and ϵ_t^i is the liquidity demand shock or the monetary policy shock at time t. We refer to the permanent deviation in output from the flexible wage benchmark as the *output hysteresis* (or alternately as the *permanent gap*). Then we can show the following proposition, generalizing the standard NK model results to an endogenous growth environment:

Proposition 5 (Output Hysteresis). *Given the monetary policy rule (equation 30) and in the absence of a ZLB constraint on the nominal interest rate, transitory (modeled as an AR(1) process) liquidity demand shocks or monetary policy shocks induce a permanent gap in the time series of output from the counterfactual (flexible wage) level of output if and only if monetary policy does not employ a strict targeting rule, i.e.*

$$Y_T \neq Y_T^e \iff \{\phi_\pi, \phi_y > 0 : \phi_\pi \not\to \infty \cup \phi_y \not\to \infty\},$$

where $1 < T < \infty$ *such that* $y_T = y$ *(the steady state value) and* $y_T \equiv \frac{Y_T}{A_T}$ *is the normalized (or stochastically detrended) output.*

Proof. See Appendix E. □

Intuitively, as long as there is an incomplete stabilization of normalized output, i.e. $\hat{y}_t \neq 0 \; \forall t$, permanent gaps emerge in this economy. This is a consequence of a standard monetary policy specification assumed in equation 30. Normalized output (and the growth rate of productivity) exhibits a monotonic response to the shocks, and the response approaches zero as the shocks die out. Thus, the sum of the productivity growth rate deviations from the steady state accumulate to the *output hysteresis*, denoted henceforth by $h_t \equiv \sum_{s=1}^{t} \hat{g}_s = \hat{g}_t + h_{t-1}$. This result generalizes the textbook results (Galí, 2015, Ch. 3) to an endogenous productivity environment.

Since entrepreneurs are forward-looking, the expectation of low future demand depresses investment in innovation. This then causes a slowdown in productivity growth, which is not offset by the monetary policy rule. Hence, the potential output is permanently lower relative to the flexible wage economy. As inflation and employment approach the economy's steady state, output tends to revert to this permanently lower level of potential output. Had the monetary policy followed a strict inflation targeting rule, these permanent effects would not have emerged. Note that under the demand shocks considered, the property of *divine coincidence* (Blanchard and Galí, 2007) holds. This result implies that the central bank faces no tradeoff in stabilizing output and inflation. Setting the nominal interest rate so as to track the natural interest rate leads to the perfect stabilization of the economy, and therefore there are no long-lasting supply effects from the demand shocks. The central bank's inability to perfectly track the natural interest rate is what gives rise to permanent supply side deviations following demand shocks. This is the second key implication of our framework and formalizes the concept of an *Inverse Say's law*, recently put forward in Summers (2015).

However, it may not be possible for the central bank to implement the optimal policy due to a binding ZLB constraint. As a result, under the standard monetary policy rule, temporary contractions in aggregate demand may result in permanent downward shifts in output.

3.3 At the Zero Lower Bound

We first show that a policy rule which perfectly stabilizes the economy when the nominal interest rate is away from the ZLB may fail to do so when monetary policy is constrained by a lower bound on the nominal interest rate. Thus, output hysteresis arises with policies that are optimal away from the ZLB in the endogenous growth environment.

To illustrate this, we follow Eggertsson and Woodford (2003) in setting up a two-state Markov chain for the natural interest rate, $\hat{r}_t^n$, in the endogenous growth economy.[19] Structurally, a negative shock to the natural interest rate is an increase in the demand for risk-free bonds that represents the *flight to safety* aspects of the financial crisis of 2007–2009 (Krishnamurthy and Vissing-Jorgensen, 2012). We assume that the economy hits the ZLB unexpectedly in period 1; that is, the nominal interest rate that is consistent with the stable inflation target breaches a policy lower bound constraint, $r_t^n < i^{LB}$ (assume $i^{LB} = 0$):

$$\text{A1a} \quad \hat{r}_t^n = \hat{r}_S < 0 \quad \forall\, 1 \le t < T^e. \tag{34}$$

With probability μ, the economy continues to stay in the low state and with complementary probability, the shock returns to the steady state. We assume that the economy is back at the no-deflation steady state after a stochastic but finite time $T^e < \infty$:

$$\text{A1b} \quad \hat{r}_t^n = (1 - \beta) > 0 \quad \forall\, t \ge T^e. \tag{35}$$

Further, we assume restrictions on parameters such that the equilibrium is locally determinate around the no-deflation steady state (Assumption A2). We calibrate the expected duration of the ZLB at 4.6 quarters (about 14 months) and the natural interest rate at –3% (annual). This calibration implies a 5% drop in (normalized) output and 1% drop in nominal wage inflation relative to the target. The central bank is assumed to follow a strict inflation targeting rule.

Proposition 6 (Output Hysteresis at the ZLB)**.** *Under a strict inflation targeting rule ($\phi_\pi \to \infty$ in equation 30), a positive shock to liquidity demand, such that the ZLB is binding for finite time T^e, results in a*

[19] In the notation of our framework, $\hat{r}_t^n = -\xi_t + (1 - \beta)$. $\xi > 1 - \beta$ makes the ZLB binding.

permanent gap in output from the flexible wage counterfactual.

Proof. See Appendix E. □

This result follows from the fact that a) when the ZLB ($t < T^e$) is binding, there is wage deflation and low output along the equilibrium path, and b) after time $t \geq T^e$ (when the ZLB is non-binding), the monetary authority raises the nominal interest rate to the level consistent with the wage inflation target and full employment. While the economic indicators of employment and wage inflation return to full capacity levels, the economy's productive potential is permanently lower relative to the counterfactual path, in which the ZLB is not binding. Such losses in potential output can be sizable for reasonable durations of a binding ZLB constraint.

Figure 3: Strict Targeting Policy at Zero Lower Bound

Source: Authors' calculations.
Note: The figures reports one realization of output, inflation, and the nominal interest rate from a two-state Markov chain for the natural interest rate under alternate policy equilibria. In period 1, the natural interest rate becomes negative, stays there for 28 quarters, and returns back to the full employment steady state. The realizations under a strict inflation targeting rule and under hysteresis targeting rule are shown. Wage inflation is plotted in deviation from steady state. Output in period –1 is normalized at 1. Black line in the output graph plots evolution of deterministic trend at an annual 2% steady state growth rate.

While we leave a thorough quantitative analysis to Section 4.5, here we illustrate the extent of output hysteresis in our model at the efficient steady state. In Figure 3, we plot output (solid line) when the ZLB is binding for 28 quarters. Output falls on impact by 5%, and in the subsequent periods productivity continues to grow at a rate slower than its annual steady state growth rate of 2% because investment in R&D is reduced during the recessionary period. This underinvestment results in a persistent output shortfall.

In Section 3.2, we proved that the strict inflation targeting rule implements the optimal monetary policy under endogenous growth when the nominal interest rate is away from the ZLB. After a ZLB episode, such a rule prescribes raising interest rates as soon as deflationary pressures subside and employment is back to full capacity, leading to a persistent output shortfall. In our framework, the ensuing long-lasting supply effects of demand shocks suggest a role for monetary policy based on an inertial rule. Reifschneider and Williams (2000), Eggertsson and Woodford (2003), and others have shown that optimal policy at the ZLB involves some form of history dependence. The key new result in our setting is that an inertial rule is needed in order to offset negative supply side effects at the ZLB.

Instead of strict inflation (or output) targeting, the central bank can target the history of productivity growth rate deviations due to current and past shocks, which we refer to as the *output hysteresis targeting* rule. Specifically, if the central bank follows a hysteresis-augmented Taylor rule of the form:

$$\hat{i}_t = \max\left(-\frac{\bar{i}}{1+\bar{i}}, \phi_\pi \hat{\pi}^w_t + \phi_y \hat{L}_t + \phi_h h_{t+1} + \hat{\varepsilon}^i_t\right), \tag{36}$$

which incorporates an additional target of the cumulative sum of all deviations in productivity growth rate h_{t+1} resulting from the history of shocks until time t, the central bank could avoid the permanent gaps by committing to maintaining a path of interest rates until output is restored to the counterfactual path of output. When $\phi_h \to \infty$, we label the rule as the *strict output hysteresis targeting rule.*

The dashed line in Figure 3 tracks the level of output under the hysteresis-augmented Taylor rule. This is an inertial policy which signals the central bank's commitment to maintaining a path of nominal interest rates consistent with reversing past policy constraints/mistakes. A positive liquidity demand shock results in a drop in normalized output and wage inflation. However, since the central bank is committed to undoing any permanent gaps in output, it is willing to tolerate excess wage inflation (Figure 3, panel B). This reduces the real interest rate gap, which results in lower growth rate deviations on impact, and allows subsequent growth rate overshooting to undo past constraints on policy. Thus, the hysteresis targeting policy embeds a *forward guidance* mechanism, credibly signaling the central bank's intention to tolerate excess inflation.[20]

Should monetary policy offset these hysteresis effects at the ZLB? We take up this question next. Our normative analysis at the ZLB retains the assumptions (A1 and A2) regarding the exogenous dynamics of

[20] Note that our use of "targeting" is distinct from that in the delegation literature. See Vestin (2006) for examples and references. That literature specifies a target for the monetary authority in that the monetary authority chooses an instrument in order to maximize a welfare-objective with a quadratic term for the target. In most cases, this welfare objective is different from the societal welfare objective function. Instead, we simply augment the central bank's policy rule with an additional objective following Chung, Herbst, and Kiley (2015). Strict targets may be implemented without an explicit instrument rule, as in Chung, Herbst, and Kiley (2015). For example, they implement a nominal GDP target with an equation that sums the price level and output gaps (from flexible level) to zero. We leave the extension of our framework to delegation problems, as in Vestin (2006), to future work.

natural interest rate and local determinacy.

Optimal Policy at the ZLB

We first solve the optimal commitment policy, when the central bank can credibly commit to future state-contingent policy actions. At the ZLB, the economy is characterized by deflation and a drop in output. By committing to pursuing an accommodative policy in the future, the central bank manages the expectations of private agents regarding the future path of inflation. This commitment policy achieves two objectives: (i) it reduces the severity of the economic contraction during the ZLB, and (ii) it allows aggregate demand to overshoot the steady state level after the the ZLB is no longer binding. While the first (*forward guidance*) channel reduces the drop in output from the trend through the reduced contraction in demand, the second channel tends to reverse past drops in output that occurred during the ZLB. The key takeaway from this analysis is that the optimal monetary policy returns the economy *close* to the pre-recession trend. In the baseline calibration, the strict inflation targeting rule rule admits a permanent output gap of 0.88 percent. On the other hand, the optimal policy involves a permanent gap of only 0.085 percent.

The policymaker maximizes the lifetime utility of the household subject to assumption 1 and the competitive equilibrium conditions: (i) Euler Equation (equation 22), (ii) Wage Setting Block (equations 27–29), (iii) Endogenous growth block (equations 23–24), (iv) resource constraints and market clearing conditions (equations 25–26), and (v) the lower bound on the nominal interest rate (equation 31).

Since the first-order conditions involve a complementary slackness condition, the solution to the optimal policy problem does not have a closed form. We solve it numerically for each state-contingent realization of the shock. We provide the first-order conditions in the appendix. The solution method is a version of shooting algorithm outlined in Eggertsson and Woodford (2003).

Figure 4 shows the equilibrium output, inflation, and the nominal interest rate under a realization of the shock binding for 28 quarters. A central bank with the ability to credibly commit offsets the permanent output gap by promising to keep interest rates lower after the ZLB is no longer binding. Under optimal policy, the central bank minimizes total losses in welfare by trading welfare losses during the ZLB against the welfare losses from the policy that arise after the ZLB is no longer binding. By committing to keeping interest

Figure 4: Optimal Policy at the Zero Lower Bound

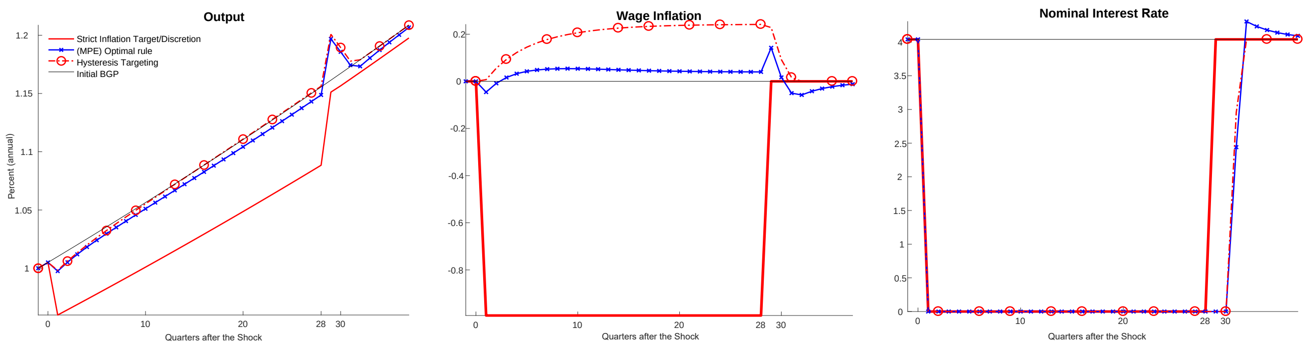

Source: Authors' calculations.
Note: The figure reports one realization of output, inflation, and the nominal interest rate from a two-state Markov chain for the natural interest rate under alternate policy equilibria. In period 1, the natural interest rate becomes negative and stays there for 28 quarters, and returns back to the full employment steady state. The realizations under a Taylor rule, Markov-Perfect Equilibrium (or discretionary) optimal policy, optimal commitment policy, and hysteresis targeting are shown. Wage inflation is plotted as a deviation from the steady state. Output in period –1 is normalized at 1. Black line in the output graph plots evolution of deterministic trend at an annual 2% steady state growth rate.

rates lower upon exiting the ZLB, the central bank creates anticipation of a boom, which lowers the real interest rate during the ZLB. This has the effect of reducing the impact of the shock relative to a discretionary policy. On impact, the drop in wage inflation and output are only 0.04% and 1.23%, respectively.

Upon exit from the ZLB, the central bank keeps the nominal interest rate lower for two additional quarters to follow through with its promise and thus creates a boom in output and inflation. Because of the procyclicality of investment in innovation, the boom in output allows for the growth rate to overshoot its target. Hence, the permanent output gap is reduced substantially for two reasons: a) *the forward guidance* channel of optimal policy, and b) the accommodation of excess wage inflation upon exit from the ZLB. In the steady state, output is only 0.085 percent below the (time-zero) efficient path of output (crossed blue line). In our numerical example, we have a two-state Markov chain for the shock process with an expected duration for the ZLB of 4.6 quarters. On average, agents expect the central bank to keep interest rates lower for two quarters beyond the 4.6 quarters implied by the shock. While we illustrate one realization of the ZLB binding for 28 quarters, we emphasize that the expansionary effects of commitment do not arise because agents at time zero expect the central bank to keep interest rates lower after 28 quarters.

Note that this is the optimal policy subject to the binding ZLB constraint.[21] It is possible to avoid the permanent output gap altogether by a commitment to accommodating even higher inflation after the ZLB

[21] If the policymaker had access to time-varying proportional tax instruments such that she could replicate the flexible wage allocation, then the first-best allocation can be implemented (as shown in Appendix D.6) However, the optimal monetary policy, in the absence of these time-varying taxes, trades off welfare losses during the ZLB episode against welfare losses in the future in the absence of appropriate time-varying tax instruments.

period. Such a policy would be optimal had the social planner put a higher weight on growth rate stabilization relative to the "true" welfare weight in equation (32) (as shown in row 3 of Table 2). However, under the "true" welfare weights, the policymaker allows some permanent output gap because perfectly neutralizing the permanent output gap comes at the expense of higher wage dispersion inefficiency upon exiting from the ZLB. Thus, *the ZLB introduces a short-run versus a long-run tradeoff for the central bank, even when we have assumed away the initial steady state distortions* (assumption 1).

Comparison with the Exogenous Growth Benchmark at the Zero Lower Bound

How does this optimal policy compare to the policy when the central bank does not internalize that it can influence the productivity growth rate? That is, a policymaker solves the optimal policy problem as before except she does not choose the productivity growth rate. The optimal policy under this non-internalizing scenario does not allow the central bank to accommodate as high an inflation rate after a ZLB episode as the optimal policy considered above would allow. Consequently, the permanent output gap is somewhat larger. Figure 5 shows the optimal policy under this "misspecified" setting and compares it to the optimal policy when the central bank internalizes the consequence of its actions on TFP growth rate.

Figure 5: Exogenous Productivity Comparison

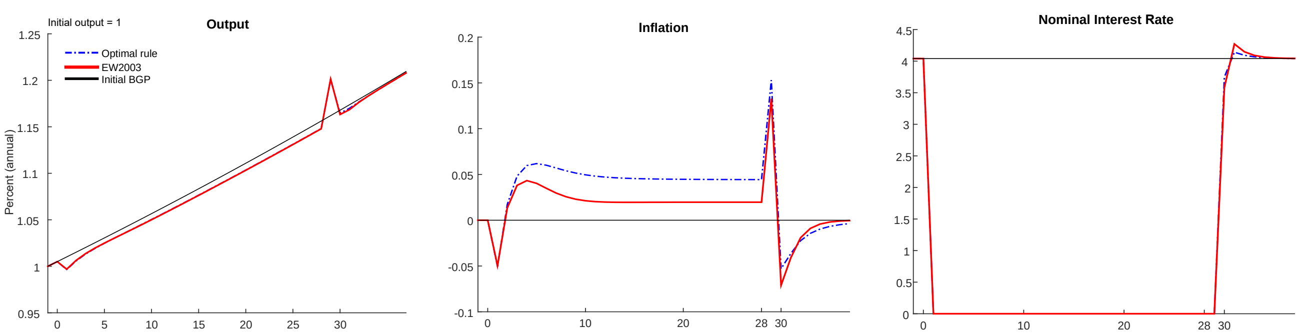

Source: Authors' calculations.
Note: The figure reports one realization of output, inflation, and the nominal interest rate from a two-state Markov chain for the natural interest rate under alternate policy equilibria. In period 1, the natural interest rate becomes negative, stays there for 28 quarters, and returns back to the full employment steady state. EW2003 denotes optimal "misspecified" policy when the central bank does not choose the productivity growth rate. The optimal rule (dashed) denotes the optimal commitment equilibrium allocation. Wage inflation is plotted in deviation from the steady state. Output in period −1 is normalized at 1. The black line in the output graph plots the evolution of the deterministic trend at an annual 2% steady state growth rate.

Quantitatively, this difference in the optimal policies is negligible (the permanent output loss is 0.09% under the misspecified problem, compared to 0.085% under the fully optimal rule). This is because the key problem in this economy is deficient aggregate demand. Since the R&D investment is pro-cyclical under liquidity demand shocks, stabilizing inflation also stabilizes aggregate output and hence R&D investment. The main implication of this analysis is that while optimal commitment policy prescriptions are not quantitatively different under the two policy environments, the cost of not adhering to optimal commitment rules is elevated because of permanent output gaps. The key insight that we illustrate next is that we do not need a vastly suboptimal rule to generate output gaps. *A minimum departure from a fully optimal policy by introducing a lack of credibility is sufficient to generate permanent output gaps.*

Markov Perfect Policy at the ZLB and the Hysteresis Bias

Next, we analyze the optimal policy when the policy maker is *unable* to commit to future policy actions. Such a policy is referred to as the *discretionary* policy and the resulting equilibrium as the *Markov Perfect Equilibrium* (MPE, formally defined in Maskin and Tirole 2001). The key result here is that the discretionary policy is characterized by a new dynamic inconsistency (Kydland and Prescott 1977) problem that we label as the *hysteresis bias*: once the ZLB is no longer binding, the nominal interest rate is set without any intention to offset the long-run effects of past contractions in aggregate demand. Hence, a policy of committing to lower future interest rates is not time-consistent because the central bank would increase the interest rate as soon as employment recovers back to full employment. The discretionary policymaker treats past productivity losses as *bygones*.

The policymaker sets the current short-term nominal interest rate in order to maximize the quadratic approximation of the welfare function (equation 32) subject to assumption 1 and the constraints: (i) Euler equation (equation 22), (ii) the wage setting block (equations 27-29), (iii) the endogenous growth block (equations 23-24), (iv) resource constraints and market clearing conditions (equations 25-26), and (v) the lower bound on the nominal interest rate (equation 31). The problem is similar to the optimal commitment problem, except the policymaker cannot commit to future policy actions.

Proposition 7 (Optimal Discretionary Policy at the ZLB)**.** *If Assumptions A1 and A2 hold, then for a given level of productivity at time zero, A_0, the Markov Perfect Equilibrium is characterized by:*

$$\log A_1 = \log A_0 + \log(1 + g_{ss});$$

for $0 < t < T^e$

$$\hat{y}_t = \psi_y r_S^n < 0; \hat{\pi}_t^w = \psi_p r_S^n < 0; \hat{g}_t = \psi_g r_S^n < 0$$

$$\log A_{t+1} = \log A_t + \psi_g r_S^n,$$

and when $t \geq T^e$

$$\hat{y}_t = \hat{\pi}_t^w = \hat{g}_t = 0$$

$$\log A_{t+1} = \log A_{t+1}^* + (T^e - 1)\psi_g r_S^n < \log A_{t+1}^*,$$

where $\psi_y = \frac{(1-\beta\mu)\eta_C^{-1}}{(1-\beta\mu)(1-\mu)-\kappa_w(\nu+\eta_C)\mu\eta_C^{-1}} > 0$, $\psi_p = \frac{\kappa_w(\nu+\eta_C)}{1-\mu\beta}\psi_y > 0$, *and* $\psi_g = \frac{1-\frac{c}{y}\eta_C}{\frac{\mathbb{R}}{y}\varrho\eta_g}\psi_y > 0$. A_{t+1}^* *is the (time-0) first-best output at time* $t+1$.

Proof. See Appendix E.1.2 □

Since the policymaker is unable to commit to future actions, optimal policy involves setting interest rates such that the economy returns to the normalized steady state as soon as the shock abates. This leads to excessive deflation during the ZLB relative to the commitment policy that involves $\hat{\pi}_{T^e}^w > 0$. This dynamic inconsistency problem, identified as the *deflation bias* by Eggertsson (2006), is also present in our setup. The new feature is that when the ZLB stops binding at stochastic time T^e, the discretionary policymaker does not offset the difference in the level of productivity from the first-best allocation. MPE thus admits a unit root in the time-series of productivity and hence output. This is the *hysteresis bias* we identify. An absence of credibility is sufficient to generate a permanent output shortfall.

Under discretionary policy, the policymaker re-optimizes every period, hence past deviations of the growth rate from the steady state are no longer under the influence of a policymaker at time T^e onwards. In order to bring the output back to the first-best allocation output, the policymaker needs to incentivize excess investment in R&D after the economy has returned to full employment. Such an allocation is not desirable from the policymaker's perspective from time T^e onwards. This can be seen by directly looking at the first-order conditions of discretionary equilibrium. Once the shock to the natural interest rate is over, the policymaker sets the real interest rate equal to the natural real interest rate, which implies that there is zero slack in the economy. Intuitively, this happens because even though the level of productivity is an endogenous state variable, it only affects the absolute level of the stochastically-trending variables. The efficiency of the resource allocation in the normalized economy is independent of the level of productivity. As soon as the central bank is able to set the normalized variables to their steady state values, it does so. Past deviations of productivity growth rate enter the welfare-loss function as additive inefficiencies that cannot be influenced by policymaker optimizing at time $t \geq T^e$. In other words, what is relevant for stabilization at time t is the gap from the time-t first-best allocation. Once the ZLB is no longer binding, setting interest rates such that employment goes back to the efficient steady state implements the time-t first-best allocation.

Figure 4 plots the path of output under the MPE. There is an unanticipated shock at time $t = 1$. Output falls by 5% and continues to grow at a slower pace. When the shock stops binding in period $T^e = 28$, the economy is permanently at a lower output trajectory. This also corresponds to the policy under a strict inflation targeting rule specification discussed above. The output in the new steady state is permanently lower by 0.88 percent. Compare the equilibrium evolution of variables under discretionary policy to that under optimal commitment policy. The discretionary policy leads to excessive deflation and slack in the economy during the ZLB. Since the discretionary policy does not offset output hysteresis, it also leads to a larger permanent output gap.

This hysteresis bias of discretionary policy thus strengthens the result from Proposition 6 that *output hysteresis is an artifact of policy constraints faced by the central bank* and does not arise because of irrational or inept behavior on part of the central bank. An implication of the hysteresis bias, we emphasize, is that it is suboptimal for the central bank to redesign policy ex-post in order to offset past output hystereses. Hence, if the central bank could *credibly commit to being irresponsible*, as suggested by Krugman (1998), it could not only reduce the deflation experienced during the ZLB periods, but also minimize the permanent output gap. This raises the stakes for optimal commitment policy: the central bank must credibly communicate this policy to the public ex-ante.

Alternative Policy Rules at the ZLB

Eggertsson and Woodford (2003) have underscored the complex nature of the optimal commitment policy in that it may not be feasible to properly communicate the policy stance to the public even if full credibility can be achieved. On the other hand, we showed that the discretionary policymaker suffers from *hysteresis bias* and does not offset past inefficiencies. In this regard, alternate simple policy rules that have a built-in commitment to reverse past policy mistakes assume importance. Such policy rules are presumably easier to communicate to the public; for example, a commitment to keep interest rates low until the permanent output gap is filled may be more readily understood. Earlier, we illustrated the potency of this strict output hysteresis targeting rule, given by:

$$h_{t+1} \equiv \sum_{s=1}^{t+1} \hat{g}_s = 0,$$

where the growth rate g_{t+1} is determined at time t.

Table 2: Policy Rules at the ZLB: Welfare Comparison

Policy Rule	*Welfare Loss*	*Permanent Output Gap*
Optimal Rules		
Discretion (MPE)	100	-0.88
Commitment	0.043	-0.085
Commitment with higher wt on $\hat{g}_t$	0.11	0
Simple Rules		
Strict Inflation Target	100	-0.88
Hysteresis Targeting	0.049	0
Wage Level Targeting	0.053	-0.30
W × Y targeting	0.311	-0.37

Notes: These values report the conditional welfare loss (in percent) starting from an efficient steady state. Losses are expressed in consumption equivalent units relative to discretionary rule. Computation details are given in Appendix E.2. The true relative weight on growth rate gap is 3.94. Under a weight of 165, the permanent output gap is zero.

The central bank ex-ante announces that it will set interest rates in order to completely eliminate permanent losses in output. Such a rule is fully optimal in the absence of the ZLB. At the ZLB, though not fully optimal, this rule may have a relative advantage in being easily communicated to the public. Figure 6 plots the nominal interest rate, output, and wage inflation under such a rule, contrasting it with the realized paths of these variables under the optimal commitment rule. The central bank keeps the interest rate low for an additional two quarters as in the optimal policy. The forward guidance element, through anticipation of higher inflation, leads to a reduction in the real interest rate, which implies a lower drop in inflation and normalized output on impact. In the calibrated experiment, output drops by 1.17% on impact.

The commitment to this simple rule implies that the central bank accommodates excess wage inflation up to 0.25% before it starts to raise interest rates gradually. Such a policy is relatively more accommodative than the optimal policy. Rows 2 and 5 in Table 2 show that the hysteresis targeting policy achieves most of the welfare gains under optimal policy relative to a strict inflation targeting (or a discretionary) policy, conditional on the ZLB being binding in period 1. An optimal commitment policy with a higher weight placed on the output gap can also close the output gap (as shown in row 3 of Table 2) but it results in somewhat higher welfare losses compared to the strict hysteresis targeting rule.

Figure 6: Alternate Rules at the Zero Lower Bound

Source: Authors' calculations.
Note: The figure reports one realization of output, inflation, and the nominal interest rate from a two-state Markov chain for the natural interest rate under alternate policy equilibria. In period 1, the natural interest rate becomes negative and stays there for 28 quarters, and returns to the full employment steady state. The realizations under a Taylor rule, Markov Perfect Equilibrium (or discretionary) optimal policy, optimal commitment policy, hysteresis targeting, and nominal wage level targeting rule are shown. Wage inflation is plotted in deviation from steady state. Output in period −1 is normalized at 1. The black line in the output graph plots evolution of the deterministic trend at an annual 2% steady state growth rate.

Contrast this policy with the policy of nominal wage level targeting (the analog of a simple price level targeting rule), where the central bank ex-ante announces its intention to set interest rates in order to attain a particular level w^* for the normalized output (y_t) adjusted nominal wages w_t^n:

$$w_t^n + \lambda y_t = w^*; \quad \text{where } \lambda \equiv \frac{1+\lambda_w}{\lambda_w}.$$

Figure 6 shows the realized paths of output, inflation, and the nominal interest rate under wage level targeting against those obtained under the optimal commitment policy. This simple policy also approximates the welfare gains achieved under the optimal commitment policy (as seen in row 6 of Table 2) relative to the discretionary policy, but results in a permanent output gap of 0.3 percent, given that it is not as accommodative as the optimal policy.

Compared with wage level targeting, the hysteresis targeting rule requires the central bank to be more tolerant of higher inflation upon exiting from the ZLB. But it may have an advantage in communication and operationalization over a policy of wage-level targeting. A central bank's commitment to keeping the interest rate lower until output has been restored to the pre-shock trend is perhaps more readily observable and implementable, assuming that achieving credibility is not a constraint for the central bank. Such a policy of hysteresis targeting is equivalent to a real GDP targeting rule because:

$$\log Y_t - \log Y_t^e = h_t,$$

where Y_t^e denotes the counterfactual path of output under time-zero flexible-wage allocation.

A third simple targeting rule is the nominal GDP (NGDP) targeting rule (see Woodford 2012 and references therein). Since our benchmark model features only nominal wage frictions, a comparison with a conventional NGDP targeting rule may not be useful. The analog of NGDP targeting in this simple framework is the $W \times Y$ rule:

$$W_t \times Y_t = W_t^e \times Y_t^e,$$

where W_t^e is the counterfactual path of nominal wages under a time-zero flexible-wage allocation. The central bank commits to adjusting interest rates in order to achieve this target relationship whenever possible. As shown in row 6 of Table 2, this $W \times Y$ rule also implies significant welfare gains and smaller persistent output shortfalls than are achieved under discretionary policy.

Table 3 compares the permanent output gaps and welfare losses achieved under these three operational rules against the optimal commitment policy for a range of innovation elasticity parameters. *Ceteris paribus*, we vary ϱ (the inverse of innovation intensity elasticity) and δ (the R&D cost parameter) in order to hit a 2% annual growth, a 10% firm entry rate and an annual R&D-to-GDP ratio in the range of 10% to 30%, while keeping all other parameters fixed at the values described in Table 1. Hysteresis targeting policy approximates the welfare gains achieved under optimal monetary policy for this range of parameters. This analysis highlights that a new operational rule that approximates welfare gains achieved under optimal policy is available for implementation in our framework. Since the standard NK models feature exogenous

Table 3: Policy Rules at the ZLB: Welfare Comparison for Range of ϱ

Innovation Intensity ϱ	1.02	1.09	1.20	Benchmark 1.47	1.50	1.71	2.78
Permanent Output Gap (Percent)							
Discretion (MPE)	–1.74	–1.58	–1.25	–0.88	–0.867	–0.695	–0.346
Commitment	0.0149	–0.073	–0.085	–0.085	–0.084	–0.076	–0.044
Hysteresis Targeting	0	0	0	0	0	0	0
Wage Level Targeting	0.048	–0.291	–0.354	–0.297	–0.29	–0.246	–0.129
W × Y targeting	0.026	–0.11	–0.136	–0.145	–0.143	–0.132	–0.08
Welfare Loss (Percent)							
Commitment	0.021	1.17	2.26	4.33	4.48	5.48	7.62
Hysteresis Targeting	0.031	1.27	2.53	4.87	5.02	6.09	8.13
Wage Level Targeting	0.073	5.90	13.83	5.25	23.41	26.67	32.63
W × Y targeting	0.04	1.36	2.88	6.43	6.66	8.56	13.13

Notes: These values report the conditional welfare loss starting from an efficient steady state. Losses are expressed in consumption equivalent units relative to the discretionary rule. Only two parameters are adjusted: Innovation intensity elasticity, $(1/\varrho)$, and research cost, δ, to target a 2% annual growth rate.

productivity, this rule is not available to the policymaker in those environments.

4 Discussion: Optimal Monetary Policy Under Alternate Shocks

In the analysis so far, we have focused on shocks such that the economy exhibits divine coincidence. The virtue of this exercise was that it did not matter whether output hysteresis was defined as a deviation from the (time-zero) first-best allocationn, (time-zero) natural rate allocation, or the pre-recession trend output. In this section, we consider alternate demand and supply shocks, and analyze the optimal monetary policy response in each case. Now the distinction between the three equilibrium concepts will become crucial.

4.1 Discount Rate Shocks

Discount rate shocks are modeled as shocks to the household's discount rate. A positive shock to the discount rate temporarily makes the household more patient. This transmits to innovation through two opposing channels: One, the lower discounting of future profits increases the present discounted value of innovation, thereby increasing investment in R&D. Two, in the presence of nominal wage rigidities, increased patience lowers aggregate consumption demand. If the aggregate demand channel is strong enough, output falls,

Figure 7: Path of GDP under TFP and Wage Markup Shocks

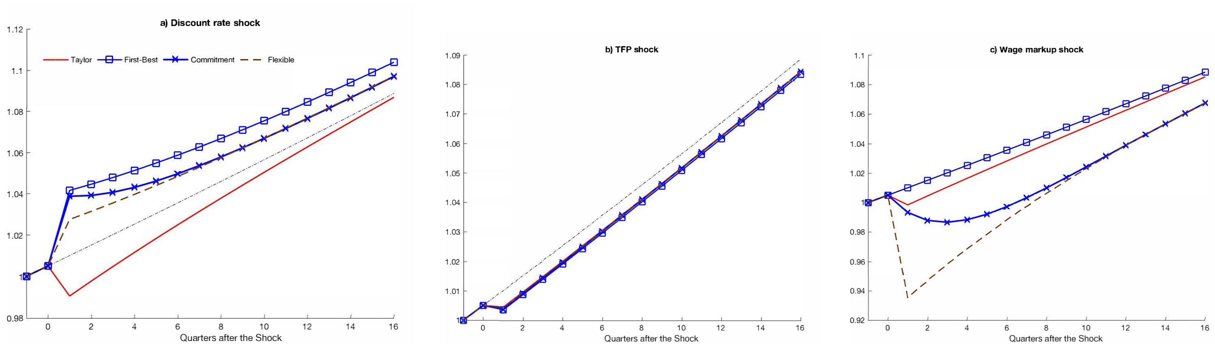

Source: Authors' calculations.
Note: The figure reports model-based evolution of GDP under discount rate (panel a), TFP (panel b), and wage markup shocks (panel c). Shocks are parametrized such that output falls by 1 percent on impact. For illustration, the shock persistence is chosen to equal 0.9. Output in period –1 is normalized at 1. Black line plots evolution of deterministic trend at an annual 2% steady state growth rate.

thereby reducing the investment in R&D due to a shrunken market (*aggregate demand effect*). Under the first-best allocation, however, prices are flexible, so there is no negative aggregate demand channel effect. This leads to an increase in R&D investment relative to the pre-recession trend (Figure 7a, the squared-blue line). In the presence of nominal rigidities, however, the overall effect on R&D investment is determined by two opposing forces, as described above. In our calibration, the aggregate demand channel dominates and investment in R&D and hence, TFP growth rate and output fall under a standard Taylor rule (Figure 7a, the red line).

The response of the first-best allocation and the flexible-wage allocation (Figure 7a, the red line) differ because of the breakdown in divine coincidence under discount rate shocks. Compared to the social planner, the entrepreneurs do not internalize the long-run benefits of innovation, despite the presence of an efficient steady state (Nuño, 2011). Replicating the flexible wage allocation is no longer an optimal policy. In fact, the natural rate of interest, *r-star*, is an endogenous object in this environment. Under the optimal commitment equilibrium, the policymaker lowers the real interest rate in order to closely replicate the welfare gains under the first-best allocation. This results in an overshooting of output relative to both the flexible-price GDP and Taylor rule GDP (Figure 7a, the crossed-blue line).

4.2 Stationary TFP shocks

A negative productivity shock shrinks the resources available for consumption and R&D investment. It is optimal to reduce R&D investment in response to a temporary reduction in the level of TFP. Temporary lower productivity growth, as a result of low investment, accumulates to generate a permanent output gap relative to the pre-shock trend. Hence, the time-zero first-best allocation features a unit-root process for output (Figure 7b, the squared-blue line). Since optimal monetary policy approximates the first-best allocation, the optimal commitment solution also admits output hysteresis (Figure 7b, the crossed-blue line).

Table 4: Welfare Comparison of Alterantive Policy Rules: Permanent Output Gap (Percent)

Policy Rule	*Discount Rate Shock*	*Markup Shock*	*Productivity Shock*	*Liq Demand Shock*	*MP Shock*
Optimal Rules					
Commitment	0.0022	0.18	0.00009	0	0
Discretion	0.0032	0.753	0.0001	0	0
Simple Rules					
Taylor Rule (eq 27)	0.0253	2.04	0.0003	0.022	0.019
Hysteresis Targeting	0.0024	4.61	0.0015	0	0
Wage Level Targeting	0.0024	0.2881	0.00009	0	0
W $\times$ Y targeting	0.0024	4.6	0.0015	0	0

4.3 Wage Markup Shocks

In the presence of cost-push shocks, the central bank faces a *tradeoff in stabilizing short-term inflation and long-run output*. The optimal commitment allocation (Figure 7c, the crossed-blue line) admits a permanent output gap. This result is a generalization of the short-run tradeoff in the exogenous growth NK model. With exogenous growth, the central bank counters a positive wage markup shock by committing to generating a negative output gap in the future. The same commitment under endogenous growth implies a reduction in market size for entrepreneurs and hence a reduced incentive to undertake R&D investment. Thus, in a bid to reduce current wage inflation, the central bank keeps output permanently below the time-zero first-best allocation. The inflation stabilization objective generates a long-run tradeoff for the central bank.[22]

4.4 Welfare Analysis

In Table 4 we report the consumption equivalent welfare losses, conditional on starting from an efficient steady state. These losses are computed as an average over 10,000 simulations with each starting at the same efficient steady state. The hysteresis targeting rule is of the form $h_{t+1} + y_t - y_t^f = 0$ rule, where the

[22] Note that the time-zero first-best allocation is a trend stationary process (Figure 7c, the squared-blue line). This is because we assume that the social planner has access to time-varying taxes to counter these shocks (Correia et al., 2013).

superscript f denotes the flexible wage allocation, h_t is the (log) hysteresis determined at time $t-1$, and y_t is the (log) stationarized output. The wage level targeting rule is implemented as $W_t + y_t - y_t^f = 0$, where W_t is the (log) nominal wage. W $\times$ Y targeting takes the form: $W_t + h_{t+1} + y_t - y_t^f = 0$. In response to demand shocks, hysteresis targeting closely replicates the welfare gains achieved under optimal commitment. In response to supply shocks, however, it is an order of magnitude more costly (in terms of welfare) to implement hysteresis targeting relative to the optimal policy. This highlights the importance of correctly identifying the source of business cycle fluctuationswhen designing optimal monetary policy.

4.5 Quantitative Assessment

In Appendix G, we build and calibrate a quantitative model, extending the Smets and Wouters (2007) model with an endogenous (Schumpeterian) growth mechanism, to evaluate the quantitative import of output hysteresis. Since the model is fairly standard, we relegate its full presentation to this appendix (see also Guerron-Quintana and Jinnai (2019) for an estimated endogenous growth business cycle model).[23] The key finding here is that the quantitative magnitude of the output hysteresis in our model depends on the elasticity of the innovation intensity, measured by inverse of parameter ρ. A lower value for ρ implies higher sensitivity of innovation success to a given change in R&D investment, which in turn allows the model to generate large changes in the productivity growth rate and hence the level of GDP. We illustrate this result with a simulation of the model under a liquidity demand shock calibrated to replicate the US Great Recession. Under a standard Taylor rule, the two calibrations of ρ of 1.07 and 3.08 generate a permanent output gap relative to the pre-recession trend of 1.25% and 0.08%, respectively. In the innovation literature, these two values of ρ lie in the range of the empirically estimated values. Hence, our exercise suggests the plausibility of quantitatively significant hysteresis under a commonly assumed monetary policy rule. Furthermore, we show that a hysteresis targeting rule also buffers the negative impact on output from the liquidity demand shock (an immediate output drop of 0.3% compared to 2.6% under a Taylor rule for $\rho = 3.08$, and a similar difference for low ρ) via its built-in commitment mechanism to keeping interest rates lower for longer. More details are provided in the Appendix G.

[23] Guerron-Quintana and Jinnai (2019) also show that endogenous growth can generate positive co-movement of equity prices and investment in recent models of financial frictions with adverse asset liquidity shocks.

5 Conclusion

This paper undertakes an analysis of optimal monetary policy in an environment where the economy's long-run potential output is endogenous to short-run fluctuations in demand. At the ZLB, an optimizing policymaker commits to keeping interest rates lower in order to offset the long-run effects of a contraction in aggregate demand. However, a policymaker who is unable to commit to future interest rates does not offset the permanent output gaps following a ZLB episode. This is the hysteresis bias of discretionary policy that we formalize in our NK framework with endogenous growth.

There are, however, certain shortcomings in our analysis that we now highlight. Our modeling assumption is that a new innovation gets adopted with certainty in the following period. This is clearly unrealistic. Comin and Hobijn (2010) and others have found that firms adopt new technology with a lag of up to seven years on average. As long as a contraction in demand results in lower investment in *knowledge creation*, the model of output hysteresis presented in this paper has insights for the conduct of monetary policy. The key elasticity determining the long-run effect of suboptimal monetary policy is the elasticity of innovation to R&D expenditure. We have discussed robustness to calibrating various parameterizations of this elasticity. However we leave the investigation of optimal monetary policy in a richer model with implementation lags and technology diffusion (see, for example, Anzoategui et al. 2019) for future work.

While the empirical evidence on the interaction between monetary policy and long-term investment in research is still scant, there is a large literature emphasizing the potency of tax credits for spurring R&D growth. Time-varying fiscal instruments in the presence of non-distortionary lump-sum taxation can replicate the first-best outcome in our framework (see Appendix D.6). However, in this paper we limit our focus to time-varying use of monetary policy instrument. We leave the analysis of optimal fiscal policy for future research.

References

Acemoglu, Daron, and Ufuk Akcigit. 2012. "Intellectual Property Rights Policy, Competition and Innovation." *Journal of the European Economic Association* 10(1): 1–42.

Acharya, Sushant, Julien Bengui, Keshav Dogra, and Shu Wee. 2018. "Slow Recoveries and Unemployment Traps: Monetary Policy in a Time of Hysteresis." Federal Reserve Bank of New York Staff Reports, 831. Available at https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr831.pdf.

Aghion, Philippe, Ufuk Akcigit, and Peter Howitt. 2014. "The Schumpeterian Model." Lecture Notes, Chapter 5. Available at http://www.brown.edu/Departments/Economics/Faculty/Peter_Howitt/2070-2015/Ch5-Schumpeter_140521.pdf.

Aghion, Philippe, George-Marios Angeletos, Abhijit Banerjee, and Kalina Manova. 2010. "Volatility and growth: Credit constraints and the composition of investment." *Journal of Monetary Economics* 57(3): 246–265.

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth through Creative Destruction." *Econometrica* 60(2).

Aghion, Philippe, and Peter W. Howitt. 2008. *The Economics of Growth*. MIT Press.

Amano, Robert, and Malik Shukayev. 2012. "Risk Premium Shocks and the Zero Lower Bound on Nominal Interest Rates." *Journal of Money, Credit and Banking* 44(8): 1475–1505.

Annicchiarico, Barbara, and Alessandra Pelloni. 2016. "Innovation, Growth and Optimal Monetary Policy." *CEIS Working Paper*.

Anzoategui, Diego, Diego Comin, Mark Gertler, and Joseba Martinez. 2019. "Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence." *American Economic Journal: Macroeconomics* 11(3): 67–110. Available at http://www.aeaweb.org/articles?id=10.1257/mac.20170269.

Barnichon, Régis, and Christian Brownlees. 2019. "Impulse Response Estimation by Smooth Local Projections." *Review of Economics and Statistics* 101(3): 522–530.

Benhabib, Jess, Stephanie Schmitt-Grohé, and Martin Uribe. 2002. "Avoiding liquidity traps." *Journal of Political Economy* 110(3): 535–563.

Benigno, Gianluca, and Luca Fornaro. 2018. "Stagnation Traps." *Review of Economic Studies*.

Benigno, Pierpaolo, and Michael Woodford. 2004. "Optimal Monetary and Fiscal Policy: A Linear-Quadratic Approach." In *NBER Macroeconomics Annual 2003, Volume 18*, 271–364. The MIT Press.

Bianchi, Francesco, Howard Kung, and Gonzalo Morales. 2019. "Growth, Slowdowns, and Recoveries." *Journal of Monetary Economics* 101: 47–63.

Blanchard, Olivier. 2018. "Should We Reject the Natural Rate Hypothesis?" *Journal of Economic Perspectives* 32(1): 97–120.

Blanchard, Olivier, and Jordi Galí. 2007. "Real Wage Rigidities and the New Keynesian Model." *Journal of Money, Credit and Banking* 39(s1): 35–65.

Blanchard, Olivier J., and Lawrence H. Summers. 1986. "Hysteresis and the European Unemployment Problem." In *NBER Macroeconomics Annual 1986, Volume 1*, 15–90. MIT Press.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb. 2017. "Are Ideas Getting Harder to Find?" Technical Report. National Bureau of Economic Research.

Brown, James R., Steven M. Fazzari, and Bruce C. Petersen. 2009. "Financing Innovation and Growth: Cash Flow, External Equity, and the 1990s R&D Boom." *The Journal of Finance* 64(1): 151–185.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12(3): 383–398.

Cette, Gilbert, John Fernald, and Benoît Mojon. 2016. "The Pre-Great Recession Slowdown in Productivity." *European Economic Review*.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber. 2016. "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities." University of Chicago Press.

Christiano, L. J., and Jonas D. M. Fisher. 1998. "Stock Market and Investment Good Prices: Implications for Macroeconomics." Technical Report.

Christiano, L.J., M. Eichenbaum, and C.L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113(1).

Chung, Hess, Edward Herbst, and Michael T. Kiley. 2015. "Effective Monetary Policy Strategies in New Keynesian Models: A Reexamination." *NBER Macroeconomics Annual* 29(1): 289–344.

Comin, Diego, and Bart Hobijn. 2010. "An Exploration of Technology Diffusion." *American Economic Review* 100(5): 2031–2059.

Correia, Isabel, Emmanuel Farhi, Juan Pablo Nicolini, and Pedro Teles. 2013. "Unconventional Fiscal Policy at the Zero Bound." *American Economic Review* 103(4): 1172–1211.

Correia, Isabel, Juan Pablo Nicolini, and Pedro Teles. 2008. "Optimal Fiscal and Monetary Policy: Equivalence Results." *Journal of Political Economy* 116(1): 141–170.

Decker, Ryan, John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2014. "The Secular Decline in Business Dynamism in the US."

Del Negro, Marco, Gauti Eggertsson, Andrea Ferrero, and Nobuhiro Kiyotaki. 2017. "The Great Escape? A Quantitative Evaluation of the Feds Non-Standard Policies." *The American Economic Review.*

DeLong, Bradford J., and Lawrence H. Summers. 2012. "Fiscal Policy in a Depressed Economy." *Brookings Papers on Economic Activity* 1: 233–297.

Eberly, Janice C. 1997. "International Evidence on Investment and Fundamentals." *European Economic Review* 41(6): 1055–1078.

Edge, Rochelle. 2003. "A Utility-Based Welfare Criterion in a Model with Endogenous Capital Accumulation." Technical Report. Board of Governors of the Federal Reserve System (US).

Eggertsson, Gauti B. 2006. "The Deflation Bias and Committing to Being Irresponsible." *Journal of Money, Credit and Banking* 38(2): 283–321.

Eggertsson, Gauti B. 2008. "Great Expectations and the End of the Depression." *American Economic Review* 98(4): 1476–1516.

Eggertsson, Gauti B. 2011. "What Fiscal Policy is Effective at Zero Interest Rates?" In *NBER Macroeconomics Annual 2010, Volume 25*, 59–112. University of Chicago Press.

Eggertsson, Gauti B., and Neil R. Mehrotra. 2015. "A Model of Secular Stagnation." Technical Report. Mimeo, Brown University.

Eggertsson, Gauti B., Neil R. Mehrotra, Sanjay R. Singh, and Lawrence H. Summers. 2016. "A Contagious Malady? Open Economy Dimensions of Secular Stagnation." Technical Report. National Bureau of Economic Research.

Eggertsson, Gauti B., and Michael Woodford. 2003. "Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* 2003(1): 139–233.

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin. 2000. "Optimal Monetary Policy with Staggered Wage and Price Contracts." *Journal of Monetary Economics* 46(2): 281–313.

Erceg, Christopher J., and Andrew T. Levin. 2014. "Labor Force Participation and Monetary Policy in the Wake of the Great Recession." *Journal of Money, Credit and Banking* 46(S2): 3–49.

Fatás, Antonio. 2000. "Do Business Cycles Cast Long Shadows? Short-Run Persistence and Economic Growth." *Journal of Economic Growth* 5(2): 147–162.

Fatás, Antonio, and Lawrence H. Summers. 2015. "The permanent effects of fiscal consolidations."

Fernald, John G. 2014a. "Productivity and Potential Output before, during, and after the Great Recession." *NBER Macroeconomics Annual* 29(1): 1–51.

Fernald, John G. 2014b. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." *Federal Reserve Bank of San Francisco Working paper* 19: 2014.

Fisher, Jonas D. M. 2015. "On the Structural Interpretation of the Smets–Wouters Risk Premium Shock." *Journal of Money, Credit and Banking* 47(2-3): 511–516.

Galí, Jordi. 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and its Applications.* Princeton, NJ: Princeton University Press.

Galí, Jordi. 2016. "Hysteresis and the European Unemployment Problem Revisited." Technical Report. National Bureau of Economic Research.

Garcia-Macia, Daniel. 2015. "The Financing of Ideas and the Great Deviation."

Gertler, Mark, and Peter Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7(1): 44–76.

Gordon, Robert J. 2015. "Secular stagnation: A Supply-Side View." *American Economic Review* 105(5): 54–59.

Griliches, Zvi. 1990. "Patent Statistics as Economic Indicators: A Survey." *Journal of Economic Literature* 28(4): 1661–1707.

Grossman, Gene M., and Elhanan Helpman. 1991. "Quality Ladders in the Theory of Growth." *Review of Economic Studies* 58(1): 43–61.

Guerrieri, Veronica, and Guido Lorenzoni. 2017. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *The Quarterly Journal of Economics* 132(3): 1427–1467.

Guerron-Quintana, Pablo A., and Ryo Jinnai. 2019. "Financial Frictions, Trends, and the Great Recession." *Quantitative Economics* 10(2): 735–773.

Gürkaynak, Refet S., Brian Sack, and Eric Swanson. 2005. "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models." *The American Economic Review* 95(1): 425–436.

Hall, Bronwyn, Jacques Mairesse, and Pierre Mohnen. 2010. "Measuring the Returns to R&D." In *Handbook of Economics of Innovation, Volume 2*, chap. Chapter 24, 1033–1082. Elsevier.

Hall, Robert E. 2016. "Why Has the Unemployment Rate Fared Better than GDP Growth?" *September 27, 2016 Draft for the Boston Federal Reserve's conference, " The Elusive 'Great Recovery': Causes and Implications for Future Business Cycle Dynamics".*

Hansen, Alvin H. 1939. "Economic Progress and Declining Population Growth." *American Economic Review* 29(1): 1–15.

Howitt, Peter. 2000. "Endogenous Growth and Cross-Country Income Differences." *American Economic Review* 829–846.

Howitt, Peter, and Philippe Aghion. 1998. "Capital Accumulation and Innovation as Complementary Factors in Long-Run Growth." *Journal of Economic Growth* 3(2): 111–130.

Ikeda, Daisuke, and Takushi Kurozumi. 2014. "Post-Crisis Slow Recovery and Monetary Policy." Technical Report.

Jones, Charles I. 2005. "Growth and Ideas." *Handbook of Economic Growth* 1: 1063–1111.

Jones, Charles I., and John C. Williams. 1998. "Measuring the Social Return to R&D." *Quarterly Journal of Economics* 1119–1135.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses Local Projections." *The American Economic Review* 95(1): 161–182.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2019. "The Effects of Quasi-Random Monetary Experiments." *Journal of Monetary Economics.*

Jordà, Òscar, Sanjay R. Singh, and Alan M. Taylor. 2019. "The Long-Run Effects of Monetary Policy." *Mimeo, UC Davis.*

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2013. "Is there a Tradeoff Between Inflation and Output Stabilization?" *American Economic Journal: Macroeconomics* 5(2): 1–31.

Kortum, Samuel. 1993. "Equilibrium R&D and the Patent–R&D ratio: US evidence." *American Economic Review* 83(2): 450–457.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen. 2012. “The Aggregate Demand for Treasury Debt.” *Journal of Political Economy* 120(2): 233–267.

Krugman, Paul R. 1998. “It’s Baaack: Japan’s Slump and the Return of the Liquidity Trap.” *Brookings Papers on Economic Activity* 2: 137–205.

Kydland, Finn E., and Edward C. Prescott. 1977. “Rules Rather than Discretion: The Inconsistency of Optimal Plans.” *Journal of Political Economy* 85(3): 473–492.

Martin, Robert, Teyanna Munyan, and Beth Anne Wilson. 2015. “Potential Output and Recessions: Are We Fooling Ourselves?” *FRB International Finance Discussion Paper #1145*.

Maskin, Eric, and Jean Tirole. 2001. “Markov Perfect Equilibrium: I. Observable Actions.” *Journal of Economic Theory* 100(2): 191–219.

Meier, Matthias, and Timo Reinelt. 2019. “Monetary Policy, Markup Dispersion and Aggregate TFP.” *Mimeo, Universitat Mannheim.*

Mertens, Karel, and Morten O. Ravn. 2013. “The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States.” *American Economic Review* 103(4): 1212–1247.

Michaillat, Pascal, and Emmanuel Saez. 2014. “An Economical Business-Cycle Model.” Technical Report. National Bureau of Economic Research.

Moran, Patrick, and Albert Queraltó. 2018. “Innovation, Productivity, and Monetary Policy.” *Journal of Monetary Economics* 93(C): 24–41.

Nakamura, Emi, and Jón Steinsson. 2008. “Five Facts about Prices: A Reevaluation of Menu Cost Models.” *The Quarterly Journal of Economics* 123(4): 1415–1464.

Neiss, Katharine S., and Edward Nelson. 2003. “The Real-Interest-Rate Gap as an Inflation Indicator.” *Macroeconomic Dynamics* 7(02): 239–262.

Nuño, Galo. 2011. “Optimal Research and Development and the Cost of Business Cycles.” *Journal of Economic Growth* 16(3): 257–283.

Ottonello, Pablo, and Thomas Winberry. 2017. “Financial Heterogeneity and the Investment Channel of Monetary Policy.” Technical Report. Working paper, July.

Queralto, Albert. 2019. “A Model of Slow Recoveries from Financial Crises.” *Journal of Monetary Economics.*

Reifschneider, David, and John C. Williams. 2000. “Three Lessons for Monetary Policy in a Low-Inflation Era.” *Journal of Money, Credit and Banking* 936–966.

Romer, Christina D., and David H. Romer. 2004. “A New Measure of Monetary Shocks: Derivation and Implications.” *The American Economic Review.*

Smets, Frank, and Rafael Wouters. 2007. “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach.” *American Economic Review* 97(3): 586–606.

Stadler, George W. 1990. “Business Cycle Models with Endogenous Technology.” *American Economic Review* 763–778.

Stock, James H. 2008. “What’s New in Econometrics: Time Series, Lecture 7.” *Short course lectures, NBER Summer Institute.*

Stock, James H., and Mark W. Watson. 2012. “Disentangling the Channels of the 2007–09 Recession.” *Brookings Papers on Economic Activity.*

Stock, James H., and Mark W. Watson. 2016. “Why Has GDP Growth Been So Slow to Recover?” *Draft for the Boston Federal Reserve’s conference, “ The Elusive ‘Great Recovery’: Causes and Implications for Future Business Cycle Dynamics”.*

Stock, James H., and Mark W. Watson. 2017. “Identification and Estimation of Dynamic Causal Effects in Macroeconomics.” *Working Paper.*

Summers, Lawrence H. 2015. "Demand Side Secular Stagnation." *American Economic Review* 105(5): 60–65.

Taylor, John B. 1993. "Discretion versus Policy Rules in Practice." In *Carnegie-Rochester conference series on public policy*, vol. 39, 195–214.

Vestin, David. 2006. "Price-Level versus Inflation Targeting." *Journal of Monetary Economics* 53(7): 1361–1376.

Wieland, Johannes F., and Mu-Jeung Yang. 2016. "Financial Dampening." Technical Report. National Bureau of Economic Research.

Wong, Arlene. 2019. "Refinancing and The Transmission of Monetary Policy to Consumption." Technical Report. Mimeo.

Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton, NJ: Princeton University Press.

Woodford, Michael. 2012. "Methods of Policy Accommodation at the Interest-Rate Lower Bound." In *Proceedings-Economic Policy Symposium-Jackson Hole.*

Yellen, Janet L. 2015. "The Outlook for the Economy." *Speech at the Providence Chamber of Commerce, Providence, Rhode Island, May.*

Yellen, Janet L. 2016. "Macroeconomic Research After the Crisis." *Speech at the "The Elusive 'Great' Recovery: Causes and Implications for Future Business Cycle Dynamics" conference, Federal Reserve Bank of Boston, October.*