The negative mean output gap

Shekhar Aiyar and Simon Voigts [*]

February 14, 2022

Abstract

We demonstrate that in a New Keynesian model with labor search frictions, downward nominal wage rigidity causes the output gap to be negative on average. Since reducing wages in downturns is more difficult for employers than increasing wages in booms, employment adjusts more in downturns than booms; over the cycle this implies that negative deviations from potential output exceed positive deviations. By contrast, most methods for estimating the output gap incorporate a mean of zero by construction. To analyze the resulting bias, standard filtering techniques are applied to simulated data from the model, allowing a comparison of the model's true gap against estimates. While filters are fairly accurate during periods of moderate macroeconomic fluctuations, they underestimate potential output substantially during deep recessions. If standard filter estimates inform countercyclical policies, this can result in a weaker than optimal monetary or fiscal response to recessions, resulting in higher output losses than if the true output gap was known. For example, focusing on the 25% deepest recessions, premature policy tightening caused by the "symmetry bias" of filters amplifies output losses by roughly a third on average.

JEL codes: E32, E37, E61
Keywords: Output gap estimation, business cycles, countercyclical policy.

---

[*] Both authors are with the International Monetary Fund. Contact addresses: saiyar@imf.org and svoigts@imf.org. We thank Helge Berger, Larry Christiano, Alfredo Cuevas, Enrica Detragiache, Bernd Droge, Gita Gopinath, Bertrand Gruss, Li Lin, Julia Otten, Tobias Krahnke, Alexander Meyer-Gohde and Jeromin Zettelmeyer for valuable comments, and Xiaobo Shao for excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

1 Introduction

1.1 Motivation and Contributions

The output gap is a central feature of modern macroeconomics. Positive and negative deviations from potential output define the amount of slack in an economy, provide the environment within which wage bargaining and price-setting occur, and correspond to booms and downturns in the business cycle. The size and evolution of the output gap is at the heart of both monetary and fiscal policy. Rules-based methodologies for setting interest rates – such as the Taylor Rule – operate within an output gap framework, while the minutes of any advanced economy monetary policy committee bear ample testimony to its role in policy deliberations. Likewise, fiscal policy is usually premised on some notion of the structural budget balance, which requires an estimate of the output gap.

This paper's first contribution is to elaborate a basic claim about the output gap, namely that in the presence of downward nominal wage rigidity, it should have a negative mean over the long run. A negative mean output gap is featured by some existing models such as Benigno and Ricci (2011), Dupraz, Nakamura and Steinsson (2017) and Schmitt-Grohé and Uribe (2016). However, to the best of our knowledge and despite its obvious policy relevance, no paper has explicitly focused on the output gap and demonstrated that the gap in a New-Keynesian model – which is typically formulated in terms of distance from the fully flexible price allocation – is equivalent to the definition relevant to policymakers, which has to do with the buildup of inflationary or deflationary pressures and has an intellectual pedigree reaching back to Friedman (1968) and Phelps (1967). We make this link, arguing that a negative mean output gap follows in a very intuitive way from downward nominal wage rigidity, a feature of the world that has overwhelming empirical support.

Our second contribution pertains to the estimation of the output gap. We show that standard output gap estimation methods ('filters' from now on) embed a symmetric view of fluctuations and constrain the estimated series average to zero. This leads to systematic bias in an economy where the true output gap is non-zero. Exploiting the fact that the true output gap is observable in a model economy, we apply standard filters to artificial data obtained from a simulation of the Abbritti and Fahr (2013) model, demonstrating that potential output is underestimated relative to its true (model-implied) level. This "symmetry bias" is particularly pronounced during deep recessions. For example, during the 25% deepest recessions, the multivariate filter (which performs better than a simple Hodricks-Prescott filter) underestimates the negative gap by roughly a third of its true size. Embedding the biased estimates into simple fiscal and monetary policy rules, we show that the measurement error can be very costly. For example, during the 25% deepest contractions, a Taylor Rule embedding the biased estimates can double the cumulative output loss relative to a Taylor Rule using the true gap.

1.2 Output Gaps, Wage Rigidities and Filters

The output gap has a long intellectual history (see Congdon, 2008, for an excellent survey). The seminal work of Phillips (1958) documented an inverse relationship between the rate of change of wages and the unemployment rate, while Okun (1962) further defined deviations from potential output in terms of the gap between actual unemployment and 'full employment' (taken by Okun to be 4 percent). But the modern formulation of the output gap dates back to Friedman (1968) and Phelps (1967), who posited that the rate of wage increase was stable at only a single rate of unemployment, termed the 'natural rate'. This natural rate of unemployment is isomorphic to a zero output gap, where the rate of inflation is stable. If output (and employment) rises above this level, inflation accelerates; if it falls below this level, inflation declines.

In New Keynesian models (see, for example Clarida, Gali and Gertler, 1999), it is price rigidities that cause deviations of output from its natural level. The fully flexible price allocation corresponds to a closed output gap, at which point there is no tendency for inflation to either rise or fall. We show that when a New Keynesian model is extended by labor search and matching frictions (e.g. Pissarides, 2000), the output gap corresponds closely to the concept of Friedman and Phelps. A positive demand shock coupled with a failure of prices and wages to adjust instantaneously moves the economy above potential, wages increase over several periods, the unemployment rate drops below its structural level, and inflation accelerates. A negative demand shock, on the other hand, causes unemployment to rise above its structural level, and the rate of inflation declines.

Our central argument is that the asymmetry introduced by downward (but not upward) nominal wage rigidity shifts the mean of the output gap below zero. There is overwhelming empirical evidence for the proposition that nominal wage cuts are rare, making the distribution of nominal wage changes extremely non-normal. The International Wage Flexibility Project (IWFP)—a consortium of 40 researchers examining wage data in 16 countries—concludes that there is strong evidence of downward nominal wage rigidity, with some heterogeneity across countries (Dickens et al., 2007). Not only does a histogram of wage changes show very little weight below zero, but there is a sharp spike at zero, indicating a nominal bound beneath which wages cannot descend despite the 'demand' for lower real wages. Moreover, the extent to which nominal wage rigidity binds appears to become more acute in recessions. Daly et al. (2013) show that in the US, the share of workers whose wage changes were 'frozen' at zero jumped from 12% in 2006 to 16% in 2011, during the Great Recession. Notably, even during the very high unemployment environment of the Great Recession, very few workers experienced wage cuts, with the numbers edging up only slightly from 2006.

Figure 1: Evidence for downward nominal wage rigidity.

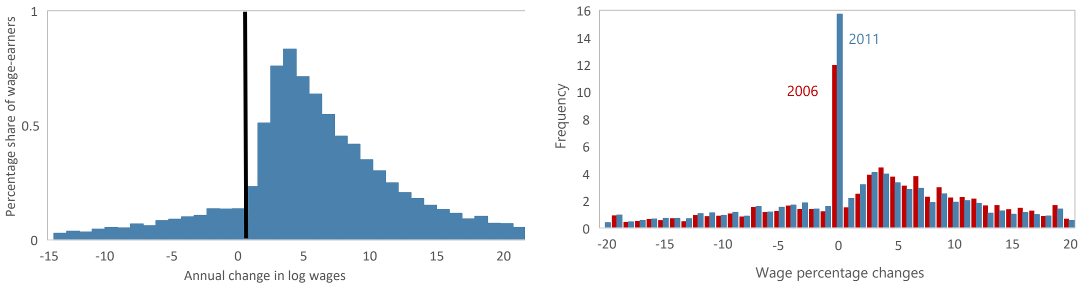

Sources: Dickens et al. (2007) and Daly et al. (2013) (left and right chart respectively).

Downward nominal wage rigidity (DNWR) results in different dynamics for negative deviations from potential output relative to positive deviations. The intuition is simple. Reducing the nominal wage is difficult or impossible for employers. So the response to a negative shock requires relatively greater adjustment of quantities (employment) to compensate for the relatively lesser adjustment of prices (wages). Assuming that positive and negative shocks to the economy are symmetrically distributed, it follows that cumulative negative deviations from potential output are greater than cumulative positive deviations. Symmetric shocks result in an asymmetric, negative-mean cycle.

DNWR is a central feature of the New-Keynesian (NK) quantitative business cycle model of Abbritti and Fahr (2013), which we borrow for our analysis and only extend by adding a simple fiscal sector as in Monacelli, Perotti and Trigari (2010). Our first result is that in this standard NK model with labor search frictions, DNWR induces asymmetry to the model's shock adjustment that causes the mean gap to be negative (about -0.5% under a calibration fitting a typical advanced economy). While the models in Dupraz, Nakamura and Steinsson (2019), Benigno and Ricci (2011), and Schmitt- Grohé and Uribe (2016) also feature a negative mean output gap, we corroborate this finding in a more standard framework.

Our second set of results pertains to the accuracy of output gap estimation methods, a classic subject in institutional and academic research. The literature can be divided into studies adopting a historical perspective based on real-world data and studies with an analytical, model-based perspective. The great majority[1] take a historical perspective and define a measurement error as the revision of the real-time estimate to the final estimate.[2] By this metric, small revisions are the mark of an estimation method's accuracy. While this approach is appealingly simple, its major drawback is that it does not speak to the plausibility of the final estimates themselves. As an illustrative thought experiment, consider an estimation method that estimates a constant, state-independent output gap. Despite its absurdity, this estimation method would be deemed perfect by the metric underlying the historical perspective: since real-time estimates always equal final estimates,the measurement error is always zero. By contrast, the analytical perspective for which we argue in this paper, derives the benchmark for true potential from a theoretical model. Estimation methods are applied to simulated data, allowing a comparison between the model's (observable) true potential and estimates. This approach is adopted by a much smaller set of studies (e.g. McCallum, 2001; Segal, 2017) and is only as meaningful as the underlying model is realistic. However, in contrast to the historical perspective, it allows an evaluation of the plausibility of final estimates, which is arguably more valuable for policymaking.

The vast majority of output gap estimation methods, including those widely used in policymaking institutions, generate estimates with a zero mean. We examine two illustrative cases, the simple but widely used HP filter and the IMF's more sophisticated multivariate filter (Blagrave et al., 2015). The restriction to zero-mean estimates makes them unable to accommodate the model's DNWR-induced negative mean gap of about -0.5%, so the gap is overestimated by 0.5% on average. Examining this bias in subsamples shows that during calm periods when no particularly large shocks materialize, the average true gap is close to zero and the zero-mean restriction does not cause a significant bias. During periods of economic distress, in contrast, the bias becomes highly significant. For example, during the 25% deepest recessions, the multivariate filter (which performs better than the simple HP) underestimates the negative gap by roughly a third of its true size. Furthermore, during the recovery periods, the estimated gap on average turns positive about 3 years ahead of the true gap.

Filters are thus highly inaccurate in precisely those periods when an accurate assessment of the gap is most crucial for policy. The explanation is straightforward. In the model, DNWR amplifies negative demand shocks, which causes the negative cumulative deviation of output from potential to be greater than the positive cumulative deviation (implying a negative mean gap). With estimated potential, in contrast, positive and negative cumulative deviations must cancel out (so that the mean estimate is zero), and the only way to achieve this is to 'drag down' estimated potential relative to true potential.

To the best of our knowledge, no other study establishes this "symmetry bias" in standard output gap filters for a standard NK model incorporating DNWR. This finding could help explain the widespread downward revisions to potential output in the wake of the Great Recession, a phenomenon noted by several authors (see e.g. Ball, 2014 or Fatás, 2018). It is also a possible

---

[1] See e.g. Orphanides and Norden (2002), Marcellino and Musso (2011), or Kempkes (2012), or Kangur et al. (2019), among many others.

[2] Real-time estimates are estimates for period T obtained in T, while final estimates are estimates for T obtained after T, i.e. from a historical perspective with a longer data sample available.

explanation for the observation in Coibion, Gorodnichenko and Ulate (2018) that estimates of potential output by policy-making institutions tend to move in line with cyclical shocks that have little long-run impact on the economy—in our simulation exercises, transitory demand shocks 'drag' estimated potential with them to maintain a zero-mean gap estimate.

Finally, we investigate how the symmetry bias can weaken the efficacy of countercyclical policy, to which output gap estimates are a key input. To this end, we incorporate the output gap in the model's monetary or fiscal policy rule, and compare simulations that use the true and the filter-estimated gap in the respective rule. Using a filter is inconsequential during periods of minor fluctuations, but leads policy to be less expansionary in deep recessions than under the true gap. For example, during the 25% deepest recessions, using a filter estimate instead of the true gap in the Taylor Rule comes close to doubling the cumulative output shortfall below potential. The results for a simple illustrative countercyclical fiscal rule are only mildly weaker. Output gap estimates appear to be a poor guide for policy during strong demand-induced contractions and their recovery phase, as potential output is strongly underestimated.

The implications of output gap uncertainty for policymaking have been studied extensively. However, in contrast to our study, the existing literature typically understands output gap uncertainty as the presence of a zero-mean measurement error, introduced as a stochastic variable or process (e.g. Orphanides et al., 2000; Boehm and House, 2019; Smets, 2002; McCallum, 2001). Segal (2017) models uncertainty by embedding an HP estimate in the Taylor Rule as we do, but the resulting measurement error is zero-mean as in the other studies since his model is symmetric. With a zero-mean measurement error, output gap uncertainty does not affect the first moments of a model, but can only lead to an increase in the second moments (causing a negligible welfare loss). In our model, non-linearity gives rise to a systematic estimation bias that can significantly impact the model's first moments, allowing us to derive a broader set of policy-relevant results.

The remainder of the paper is structured as follows. Section 2 introduces the model, its calibration, and provides intuition for its adjustment to shocks. After a technical digression on the risk channel of uncertainty, Section 3 presents stochastic model simulations and measures the bias of the two standard output gap filters. Section 4 turns to the policy simulations and illustrates how the symmetry bias constrains the effectiveness of countercyclical policies. Section 5 concludes.

2 Model

Abbritti and Fahr (2013) develop a New-Keynesian DSGE model with search frictions in the labor market and DNWR in the form of asymmetric wage adjustment costs. It is designed to capture business cycle asymmetries and largely replicates empirically documented skewness in key macroeconomic variables. Following the example of Monacelli, Perotti and Trigari (2010), we add a simple fiscal sector to allow an assessment of the implications of fiscal policy under different estimates of the output gap (given in Section 4). Since the Abbritti and Fahr (2013) model is only extended slightly, we present here a short summary highlighting non-standard elements and the fiscal sector, while referring the reader to the original paper for a full exposition.

2.1 Summary of the model

Potential output Y_t^{flex} in NK models is output under price and wage flexibility (see for example Clarida, Gali and Gertler, 1999). In the simulation exercises we will however use a modified measure of potential, $\tilde{Y}_t^{flex} = Y_t^{flex} - \Lambda$, where Λ is a constant 'correction term'. As detailed in Section 3.1, uncertainty effects originating from the presence of nominal rigidity reduce actual output $Y_t^{rigidity}$ by a constant amount (nominal rigidity amplifies macroeconomic volatility and with it uncertainty, leading agents to alter their behavior in a way that lowers output). This constant drag on actual output shows up in the difference $Y_t^{rigidity} - Y_t^{flex}$ since uncertainty effects play no role for potential output (which is obtained under nominal flexibility). This would confound our results, since we are narrowly interested in how DNWR causes a negative mean gap, rather than in the broader implications of uncertainty for the output gap. To abstract from the latter, Λ is set to the size of the constant drag on actual output. This way, uncertainty effects also reduce our measure of potential, and thus drop out in the output gap given by

$$gap_t = \frac{Y_t^{rigidity} - \tilde{Y}_t^{flex}}{\tilde{Y}_t^{flex}}. \tag{1}$$

The level of unemployment under nominal flexibility (associated with $\tilde{Y}_t^{flex}$) is interpreted as the NAIRU. It is positive as a result of labor search frictions combined with a positive probability for matches to be destroyed in each period. For $gap_t = 0$, inflation is stable and the unemployment rate equals the NAIRU, while a positive (negative) gap accelerates (slows down) inflation and is associated with an unemployment below (above) the NAIRU.

The labor market is characterized by search and matching frictions. The matching function is standard and has constant returns to scale. The model assumes contemporaneous hiring, so matches between workers and firms create employment in the same period. Match destruction occurs exogenously with a positive probability.

Households are large families with members distributed along the unit interval, who pool consumption and insure each other against employment risk. $\mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \log C_t$ is a representative household's discounted present value of utility from consumption, where $log()$ is the utility function, C_t a Dixit–Stiglitz consumption bundle and β the discount factor. The period budget constraint is:

$$C_t + \frac{B_t}{P_t R_t \exp(\varepsilon_t^b)} = n_t w_t + (1 - n_t) b + \frac{B_{t-1}}{P_t} + D_t - T_t$$

where w_t is the real wage earned by employed family members (mass n), and b_t the income of unemployed members. P_t denotes the aggregate price level, D_t are profits from the ownership of firms and T_t are lump sum taxes. Holding risk-free nominal bonds B_t yields the gross nominal interest rate $R_t \exp\left(\varepsilon_t^b\right)$, where $\varepsilon_t^b = \rho^b \varepsilon_{t-1}^b + \eta_t^b$ (with $\eta_t^b \sim N\left(0, \sigma^b\right)$) is a zero-mean serially correlated risk premium shock. Its purpose is to introduce demand fluctuations (following Smets and Wouters, 2007) in a way that makes potential output invariant to them.[3] Optimization yields a standard consumption Euler equation.

There are two types of firms. Wholesale firms use capital and labor to produce homogeneous intermediate goods with a standard production function, in which total factor productivity $\exp(Z_t)$ with $Z_t = \rho^z Z_{t-1} + \eta_t^z, \quad \eta_t^z \sim N(0, \sigma^z)$ is subject to mean-zero productivity shocks. Because of labor search frictions, wholesale firms have to post vacancies to attract workers. Investment adjustment costs are symmetric, while wage adjustments give rise to asymmetric costs c_t^w described in greater detail below. The first order condition for vacancy posting is given by

$$\frac{\kappa}{\lambda_t q_t} = J_t = \alpha \varphi_t \frac{Y_t}{n_t} - w_t \left(1 + c_t^w\right) + (1 - s) \mathbb{E}_t \left[\beta_{t,t+1} \frac{\kappa}{\lambda_{t+1} q_{t+1}}\right],$$

---

[3] If demand disturbances were alternatively introduced by shocks to the time discount factor, they would also change potential. With risk premium shocks, in contrast, potential is exclusively determined by the history of productivity shocks, in line with interpreting potential as summary of the economy's productive capacity.

where the LHS represents the expected costs of filling a vacancy (κ, λ_t, q_t are posting costs, the Langrange mutliplier on the household budget constraint, and labor market tightness, respectively). The RHS describes the expected value of a new worker, consisting of additional output $\alpha Y_t/n_t$ (expressed in units of the final good by multiplying the relative price of wholesale goods φ_t), minus wages including adjustment costs c_t^w, plus the continuation value. Optimal investment implies that the product of capital plus the continuation value of an additional unit equal expected costs including capital adjustment costs. Wholesale firms sell their goods to retail goods firms on a competitive market.

There is a continuum of retail firms producing differentiated goods that are imperfect substitutes in the consumption bundle. Each firm buys wholesale goods at price $P_t\varphi_t$ and transforms them one-to-one and at no costs into a differentiated good variety. Monopolistic competition allows firms to charge a markup. Optimal price setting subject to standard convex adjustment costs yields the following Philipps Curve

$$\Gamma_t'\Pi_t = \epsilon\left(\varphi_t + \Gamma_t\right) - (\epsilon - 1) + \mathbb{E}_t\left[\beta_{t,t+1}\frac{Y_{t+1}}{Y_t}\Gamma_{t+1}'\Pi_{t+1}\right],$$

where Π_t is inflation, ϵ is the elasticity of substitution between retail good varieties, Γ_t adjustment costs and $\beta_{t,t+1}$ the stochastic discount factor.

Asymmetric wage adjustment costs c_t^w give rise to downward nominal wage rigidity, the central non-standard feature of the model. It is assumed that adjustment costs c_t^w are proportional to the overall wage bill and that they are identical for new hires and existing workers (which can be interpreted as the extension of wage agreements). Wage adjustment costs have the same functional form as in Fahr and Smets (2010):

$$c_t^w\left(\pi_t^w\right) = \frac{\phi^w - 1}{2}\left(\pi_t^w - \pi_{ss}^w\right)^2 + \frac{1}{\psi^2}\left\{\exp\left[-\psi\left(\pi_t^w - \pi_{ss}^w\right)\right] + \psi\left(\pi_t^w - \pi_{ss}^w\right) - 1\right\} \tag{2}$$

where π_t^w is wage inflation, π_{ss}^w is steady state wage inflation[4], ϕ^w governs the degree of convexity, and ψ the degree of asymmetry around π_{ss}^w. For positive values of ψ, reducing wage inflation below steady state is more costly than raising it above. This asymmetry disappears as $\psi \to 0$, since the asymmetric last term then goes to zero. Figure 2 shows adjustment costs for three illustrative cases of no asymmetry, weak asymmetry, and strong asymmetry.

---

[4] π_{ss}^w equals the long-run deterministic output growth rate γ, but long-run growth is neglected in the simulation exercises presented later.

Figure 2: Illustrative wage adjustment costs.

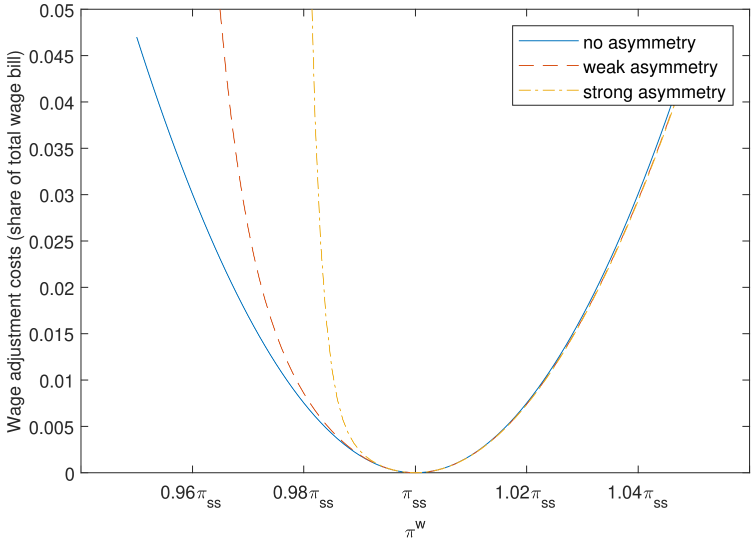

Wage determination is governed by Nash bargaining. The only non-standard element is that wage adjustment costs affect the value of a job for the firm. As shown in detail in Arseneau and Chugh (2008), the effective bargaining power of workers is given by $\omega_t = \frac{\eta}{\eta+(1-\eta)\tau_{t,t+1}}$, with η denoting the exogenous part and

$$\tau_{t,t+1} = 1 + c_t^w + \frac{\delta c_t^w}{\delta W_t} W_t + (1-s)\,\mathbb{E}_t \beta_{t,t+1} \left(\frac{\delta c_t^w}{\delta W_t} \right) \frac{W_{t+1}}{\Pi_{t+1}}$$

the state-dependent part.[5] The effective bargaining power of workers deteriorates when wages rise ($\delta c_t^w / \delta W_t > 0$), while it increases when wages decline. Due to the asymmetry of wage adjustment costs the bargaining power rises by more in recessions (when downward wage adjustment leads to high costs) than it drops in expansions (when wages adjust upwards at comparably little costs). As stressed in Abbritti and Fahr (2013), this mechanism is an important magnifier for the impact of downward wage rigidity on the adjustment of the economy.

The resulting wage is given by

$$w_t = \omega_t \left(\alpha \varphi_t \frac{Y_t}{n_t} - c_t^w w_t + (1-s)\,\mathbb{E}_t \left[\beta_{t,t+1} J_{t+1} \right] \right) + (1-\omega_t) \left(b_t - (1-s)\,\mathbb{E}_t \left[\beta_{t,t+t} \left(1 - f_{t+1}\right) \tilde{N}_{t+1} \right] \right),$$

where $\tilde{N}_t$ is the value of employment for a family, J_t the value of a filled vacancy for a firm, and s the job destruction rate. $\omega_t c_t^w w_t$ is the deadweight loss of wage adjustment costs, which lowers the surplus of a match and thereby the wage. In this setup, the state-dependency of workers' effective bargaining power ω_t mitigates fluctuations in the wage bill. When there is upward pressures on wages, the resulting decline in bargaining power mitigates the wage adjustment. By the same token, downward pressure goes along with an increase in bargaining power, dampening the magnitude of the wage reduction.

---

[5] The optimal sharing rule is given by $\omega_t J_t = (1-\omega_t)\,\tilde{N}_t$, with $\tilde{N}_t$ and J_t denoting the value of employment for worker and the firm respectively. It nests the standard rule when there are no wage adjustment costs, i.e. for $c_t^w = 0$ and therefore $\tau_{t,t+1} = 1$.

Monetary policy is described by a standard Taylor rule with interest rate smoothing. The output gap matters for monetary policy only if $\omega_{\text{gap}} > 0$, which is the case in some policy scenarios discussed in Section 4.

$$i_t = i_{t-1}^{\omega_r} \left(\frac{1}{\beta} \pi_t^{\omega_\pi} \text{gap}_t^{\omega_{\text{gap}}} \right)^{\omega_r - 1} \tag{3}$$

The fiscal sector is modelled in a simple manner following the example of Monacelli, Perotti and Trigari (2010), with government spending G_t assumed to be waste. G_t is fully financed by lump sum taxes each period, $G_t = T_t$, and is governed by the rule

$$G_t = \phi \; gap_t \; . \tag{4}$$

For $\phi < 0$, government spending has the opposite sign than the output gap, implying that it cushions deviations of actual output from potential. This is the case in Section 4, whereas $\phi = 0$ in all other exercises.

The resource constraint states that total output must cover the sum of private consumption, government spending, investment, and adjustment costs.

2.2 Calibration

All parameter values other than those of exogenous shocks processes and the monetary and fiscal policy rules are taken from Abbritti and Fahr (2013). Selected parameter values related to nominal rigidity and the labor market are shown in Table 1. In a nutshell, the labor market calibration implies a steady state unemployment of 10% and a quarterly job finding rate of 0.35.

We use a different calibration for exogenous shocks because the standard deviation (SD) of gap_t in the original calibration is only 0.55 percent of potential, which is unreasonably small compared to the amplitudes of estimates (typically 3 to 4 times larger in advanced economies). To improve realism, parameters governing the stochastic profile of shocks processes are calibrated such that the SD and first-order autocorrelation (AC) of potential output and of the output gap are in line with IMF estimates. This strategy does not impose a restriction on the mean output gap—which is determined by the model—but relies on institutional estimates to calibrate the magnitude and persistence of its fluctuations. IMF estimates embed the judgment of practitioners with a detailed understanding of the respective country, and are generally well in line with those of other institutions.[6]

In particular, we examine potential output and GDP data from 1981 to 2017, obtained from the 2017 World Economic Outlook. To focus on advanced economies we use a non-weighted average of G7 countries, which yields a SD of the output gap of 1.77% and an AC of 0.61. To reproduce these moments, σ^b and ρ^b (governing the dynamics of demand shocks) are set to 0.33 and 0.85 respectively. There is a clear mapping from estimated output gap dynamics to the calibration of demand shocks because only demand shocks drive output gaps in the model (productivity shocks are neutral to the gap, mainly owing to the search-and-matching frictions on the labor market, as explained below). For productivity shocks we use the standard persistence parameter $\rho^z = 0.95$, but chose σ^z so as to replicate deviations of estimated potential around its long-run trend. The latter must be constructed because IMF estimates do not distinguish between trend productivity growth and transitory fluctuations. To do so, we first run the multiple breakpoint test embedded in eViews (using the global information criteria) to detect up to 5 breaks in a regression of estimated potential on a time trend and a constant. Then, we

---

[6] For example, the average standard deviation of output gap estimates for G7 countries (over a comparable time span starting at the earliest available date) is 1.77% for the IMF and 1.98% for the OECD.

allow for changes in trend growth at those break points. The standard deviation of the resulting residuals is 0.3% percent of potential (averaged over the G7 countries), which is replicated in the model for $\sigma^z = 0.05$. Estimated potential fluctuations map into productivity shocks because only this type of shock moves potential output. Overall, in our calibration, productivity shocks induce comparably small, low-frequency movements of output driven by potential. Demand shocks cause stronger fluctuations of output around potential at a higher frequency, with overall output fluctuations well in line with the data.[7]

In the monetary and fiscal policy rules, $\omega_{gap} = 0$ and $\phi = 0$ mute any response to the output gap in all exercises other than those presented in Section 4. In this section, we use the standard value $\phi = 0.125$ in the Taylor Rule (see e.g. Gopinath et al., 2020), and $\phi = 0.5$ as an illustrative parameter in the fiscal rule.

Table 1: Key Aspects of Model Calibration.

Parameter	Value	Motivation
Nominal rigidity		
ϕ^w: price rigidity	60.5	empirical price duration (about 3 quarters)
ϕ^w and ψ: symmetric and asymmetric wage rigidity	37.6 and 24100	empirical volatility and skewness of nominal wage inflation
Labor market		
$\bar{m}$: matching efficiency	0.56	st.st. job filling rate of 0.9 (Ravenna and Walsh, 2011)
κ: vacancy posting costs	0.214	st.st. aggregate hiring costs 1% GDP (Blanchard and Galí, 2010)
b: non-work income	1.25	determined by st.st. relationships
ζ: elasticity of matches	0.5	matching function estimates (Petrongolo and Pissarides, 2001)
η: worker bargaining power	0.5	as in Blanchard and Galí (2010)
Monetary policy		
ω_w: persistence parameter	0.85	standard
ω_π: weight inflation	1.5	standard
ω_{gap}: weight output gap	0 or 0.125	deactivated or standard value
Fiscal policy		
ϕ: output gap coefficient	0, 0.25 or 0.5	deactivated or illustrative low and high responsiveness
Shock processes		
σ^b: SD risk premium	0.33	empirical SD of output gap estimates
ρ^b: persist. risk premium	0.85	empirical AC of output gap estimates
σ^z: SD productivity	0.05	empirical SD of potential fluctuations around trend
ρ^z: persist. prod.	0.95	standard

2.3 Shock adjustment

Demand shocks. The dashed lines in Figure 3 depict the impact of a negative (i.e. expansionary) risk premium shock of one standard deviation (33 basis points in the first period). The shock reduces the return on assets held by households below the interest rate set by the central bank and thereby increases consumption. At the same time it raises investment by lowering the cost of capital. The surge in demand is met by an increase in output caused by a rise in

---

[7] Averaging over the G7 countries, the standard deviation of annual real GDP around its HP-filtered trend (using $\lambda = 6.25$ and a sample from 1980-2019) is 1.23 percent of the trend. Applying the same calculus to annualized GDP in the model yields a standard deviation of 1.02 percent.

the number of vacancies and therefore employment. The upward adjustment of wages is not constrained by the asymmetric component of adjustment costs and exceeds inflation, implying a mild increase in real wages. The central bank reacts by raising the nominal interest rate to partly offset the shock's impact on households' return on assets and firms' cost of capital.

Solid lines depict the adjustment to a positive (i.e. contractionary) shock of the same size. The induced decline in output is larger and more persistent than its increase under the expansionary shock. Accordingly, the scaling-back of vacancy postings is stronger and more persistent than the hike in vacancy postings following the expansionary shock. This asymmetry is accounted for by the impact of downward nominal wage rigidity on real wage adjustment. When the economy expands, nominal wages can rise by enough to elevate the real wage, which dampens firms' incentives to post vacancies and therefore *mitigates* the upward adjustment of output. When the economy contracts there is a crucial difference. In this case also the real wage increases, which now reinforces firms' willingness to reduce vacancy postings, and therefore *amplifies* the reduction in output. The real wage increases because downward nominal wage rigidity prevents nominal wages from falling by more than prices.[8]

Turning to the output gap, note that since the real allocation under nominal flexibility is invariant to demand shocks (when induced via changes in the risk premium), potential output and the NAIRU remain unchanged.[9] With constant potential, deviations of output shown in the first panel correspond to changes in the output gap, and positive (negative) deviations in the employment panel indicate that unemployment drops below (raises above) the NAIRU.[10] Owing to the asymmetry of nominal wage rigidity, a contractionary demand shock causes a greater and more persistent output gap (and deviation of unemployment from the NAIRU) than an expansionary shock of the same size.

---

[8] Abbritti and Fahr (2013) show that rising real wages are not unusual in historical recessionary periods.

[9] Intuitively, any pressure to move demand away from the constant level of potential would lead to infinitely strong price adjustments when there is no rigidity, leading the central bank to adjust rates until the shock is perfectly offset.

[10] Since the labor force is constant, a given increase in employment corresponds to a decline in unemployment by the same amount.

Figure 3: Adjustment to positive (dashed lines) and negative
(solid lines) risk premium shock of one standard deviation.

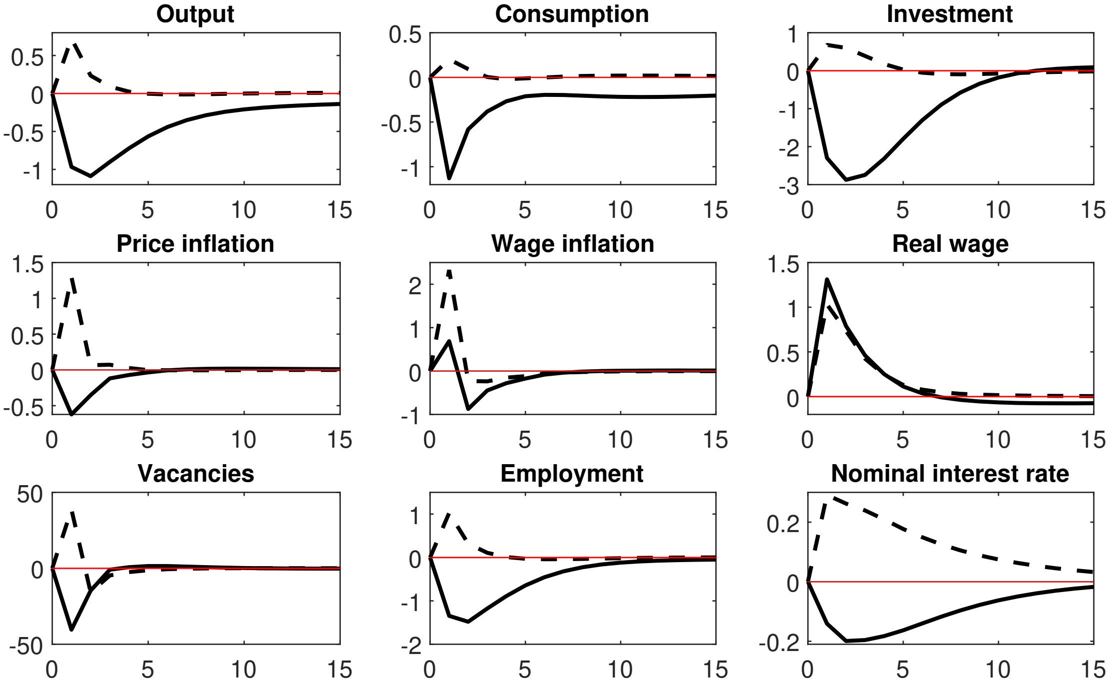

Supply shocks. The solid (dashed) black lines in Figure 4 depict the adjustment to a positive (negative) productivity shock of one standard deviation. Compared to demand shocks, the impact on output is small but highly persistent. A positive shock raises productivity and increases the value of a worker for the firm, which leads firms to post more vacancies and increase employment. Investment surges with higher productivity, and consumption rises due to a positive wealth effect. The increase in the real wage—resulting from wage inflation combined with a mild decline in prices—offsets some of the uptick in vacancy postings caused by the productivity gain. The adjustment of the nominal interest rate is negligible due to the small magnitude of price adjustments. We skip the discussion of the negative productivity shock because its impact is almost perfectly symmetric, for the reason that nominal rigidity—and thus also DNWR—only plays a negligible role, as explained in the following.

Grey lines marked with crosses depict the shock's impact when prices and wages are flexible. The output panel then shows potential, which rises (declines) persistently when the economy's productive capacity is enhanced (reduced) by a higher (lower) productivity. The output gap, which corresponds to the distance between the grey and black lines in the output panel, has a negligible response to both positive and negative productivity shocks. Absent nominal rigidity, employment rises (declines) in response to a positive (negative) shock, mirroring a decrease (increase) of the NAIRU. The changes in employment with and without nominal rigidity are very much in line, consistent with the virtually closed output gap and the negligibly weak wage and price adjustment pressures.

The fact that the economy's adjustment is not significantly constrained by nominal rigidity, and thus that productivity shocks do not cause significant output gaps, can be explained by the search-and-matching specification of the labor market. Wage changes are irrelevant for the *stock* of workers already in employment, but matter only for the *inflow* of new workers by affecting vacancy posting (the value of a worker to the firm depends on the wage). This makes delays in wage adjustments resulting from nominal rigidity less consequential for labor supply than in a simple neoclassical labor market, where labor supply is directly linked to the real wage. Furthermore, when a productivity shock is realized, the dominant driver in the adjustment of the value of a worker to the firm is the change in their productivity itself; wage adjustments, including potential delays caused by nominal rigidity, play a negligible role.

Figure 4: Adjustment to positive and negative productivity shock of one standard deviation. *Solid (dashed) black lines depict the positive (negative) shock. Grey lines with crosses show the adjustment under nominal flexibility.*

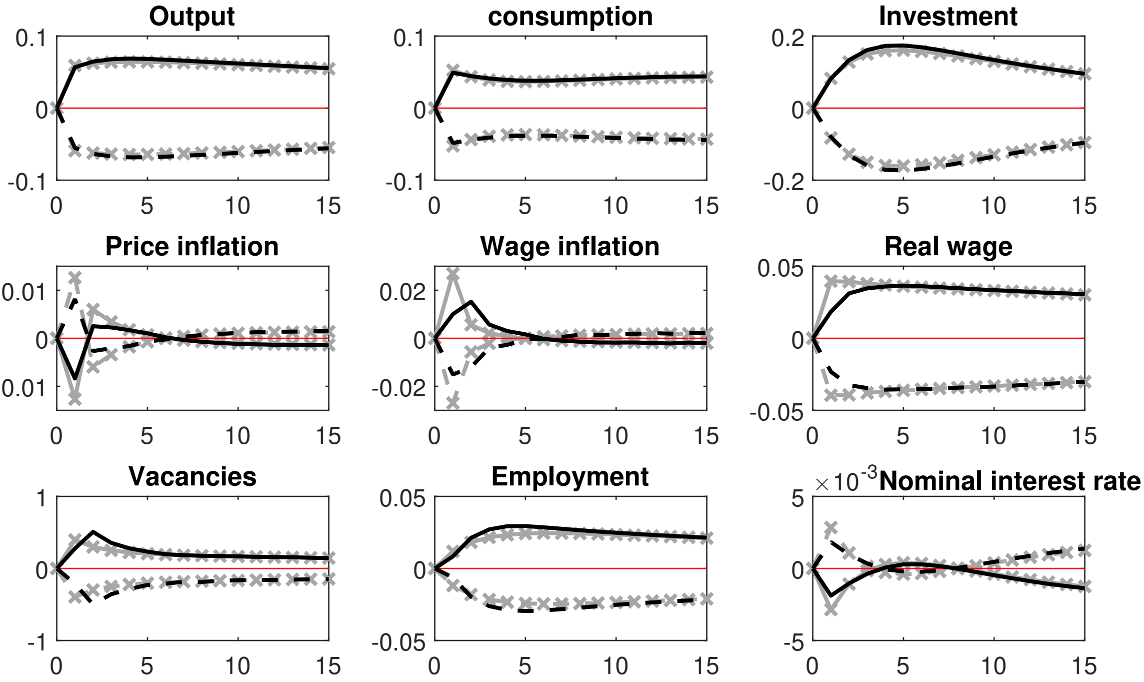

3 Simulations on the mean output gap and filter accuracy

Before discussing simulation results, we digress to the use of a "corrected" measure of potential.

3.1 Controlling for the uncertainty-channel of nominal rigidity

This study examines how DNWR amplifies the impact of adverse demand shocks and thereby causes a negative mean output gap. However, asymmetry in wage adjustment is not the only driver of the gap's negative mean, since as a significant share results from uncertainty effects. Uncertainty effects describe how uncertainty perceived by agents changes their decision-making in a way that typically lowers mean output (see e.g. Coeurdacier, Rey and Winant, 2011). These effects weigh more on actual output than on potential output, because nominal rigidity (which is only relevant for actual output), increases uncertainty by amplifying macroeconomic volatility.[11] This impact of uncertainty on the New-Keynesian output gap would confound our narrow analysis of the implications of DNWR on the gap. As outlined in the following, we therefore use the correction term Λ in equation (1) to artificially reduce potential output by the same amount as rigidity-driven uncertainty effects reduce actual output—with the effect that they drop out in the difference between actual and potential.

We run an auxiliary model simulation with a calibration where the asymmetric component of wage adjustment costs is deactivated, while the symmetric component is adjusted such that overall wage rigidity (measured by the volatility of wages) is kept constant.[12] Since there is no DNWR, the resulting negative mean output gap is fully attributed to uncertainty effects, which, since overall nominal rigidity is kept constant, operate in the same strength as in the baseline calibration with DNWR. We then set Λ equal to the difference between mean output in this simulation and the mean of Y_t^{flex}, i.e. output when there is no nominal rigidity at all (and thus no resulting amplification of volatility and uncertainty). The value of Λ thus equals the time-invariant[13] drag on output resulting from nominal rigidity-driven uncertainty. By deducting Λ from potential in the definition of the output gap (1), we introduce the adverse implications of uncertainty also to potential, so that it drops out in the output gap. gap_t can thus be fully attributed to DNWR-induced asymmetry in the model's shock adjustment.

This correction is quantitatively significant: relative to non-corrected potential (for $\Lambda = 0$), the mean output gap is -1.08% in the baseline calibration with DNWR and -0.62% when wage adjustment is symmetric. The difference of 0.46% is the isolated contribution of DNWR. Our choice of Λ reduces potential by 0.62%, so that the gap is zero for symmetric wage adjustment, and exclusively captures the -0.46% resulting from DNWR under the baseline calibration.

Although we abstract from uncertainty effects in the rest of this paper, we briefly discuss the economics behind them. In our model with labor search frictions, the quantitatively most relevant implication of uncertainty is that it gives rise to an option value of not hiring workers (see Leduc and Liu, 2016). As firms are unable to lay off workers until a match is exogenously destroyed by a random event, hiring workers entails the risk of not being able to cut employment when it is desirable. The associated option-value of not hiring disincentives employment and thereby lowers output. An additional simulation exercise shows that this channel is behind roughly 3/4 of the -0.62% gap resulting from uncertainty effects.[14]

3.2 Mean output gap in the model

This section presents simulations of the output gap. We run two simulations, both based on the same random draw of time series of supply and demand shock. First, actual output is obtained under the model's baseline calibration. Second, for the potential series, we run the model without nominal rigidities, i.e. with price and wage adjustment costs of zero. 30,000

---

[11] Nominal rigidity increases the standard deviation of output by a factor of about six. The main reason is that it is only under rigidity that demand shocks affect the real allocation (see Section 2.3).

[12] Asymmetry is deactivated by setting $\psi \to 0$ in equation (2). To keep overall wage rigidity constant, we follow Abbritti and Fahr (2013) in re-adjusting the convexity parameter of adjustment costs such that nominal wage volatility remains unchanged. In our calibration this requires adjusting ϕ^w from 37.6 to 5.

[13] Uncertainty effects materialize as a constant reduction of output, as their strength is determined by the model's unconditional second moments that are fixed in the model calibration.

[14] We compare the auxiliary calibration with symmetric wage adjustment costs with an alternative calibration that approximates a spot labor market and thereby suppresses the option-value of not hiring. Following Leduc and Liu (2016), a spot labor market is approximated by using negligibly small vacancy posting costs combined with a 90% breakup probability for matches every period. Since this modification amplifies the volatility of output, we simultaneously reduce the volatility of shocks by about 2/3. This keeps the standard deviation of output divided by mean output constant across the two calibrations, allowing us to control for the degree of uncertainty. When the option-value is suppressed, the mean gap is -0.15%, suggesting that option-value drives about 3/4 of the headline figure of -0.63%. We do not investigate the nature of the remaining -0.15%.

quarters of data are generated using a second order accurate perturbation method to retain asymmetries (following Abbritti and Fahr, 2013).

Figure 5 shows an illustrative subsample of 100 quarters. As we abstract from a deterministic long-run productivity trend, actual and potential output (shown in the upper panel) fluctuate around the steady state output level of around 2. Productivity shocks induce mild variations of potential at a low frequency, while demand shocks cause larger deviation of output around potential at a higher frequency. Deviations of actual output above potential tend to be smaller than those below potential, as DNWR only amplifies the impact of contractionary demand shocks. Accordingly, the output gap (shown in the lower panel, expressed as percentage of potential) appears to be negative more often than positive, and the negative deviations tend to be larger in magnitude. Over the full sample, the gap relative to corrected potential (abstracting from uncertainty effects), averages to -0.46%.

Figure 5: Simulated actual and potential.

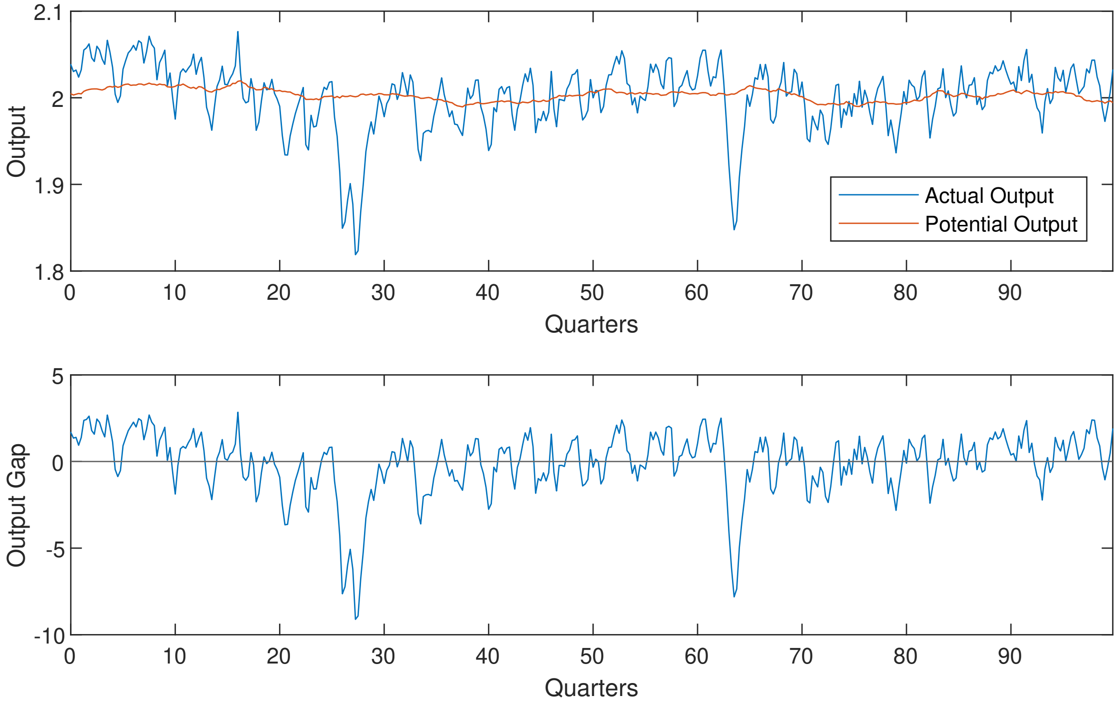

The full sample-mean of -0.46% is predominantly driven by rare and especially deep contractions. To shed light on the periods driving the negative mean gap, we apply a simple procedure to partition the sample into expansionary and contractionary subsamples. It is based on the Harding-Pagan business cycle dating algorithm and illustrated in Figure 6. First, the algorithm is applied on the simulated data to identify peaks and troughs (depicted by P1 to P4 and T1 to T4 respectively).[15] Contractionary (expansionary) subsamples are then defined as periods

---

[15] The algorithm is calibrated such that the average duration of a full cycle is about 8 years. When annualized data is used in the context of the MVF, the parameters are adjusted accordingly.

in proximity to a trough (peak). That is, expansionary subsamples (illustrated by green bars) consist of periods whose closest extrema is a peak, while contractionary subsamples (red bars) consist of periods closer to a trough than to a peak. With this partitioning procedure, expansionary and contractionary subsamples jointly cover the full sample. However, in order to focus only on the most or least extreme contractions or expansion, we also consider alternative partitionings where parts of the sample are discarded. For example, when we focus on the 25% deepest contractions and highest expansion, the expansionary (contractionary) subsample would only consist of P3 (T2). If we considered the 50% most extreme periods, the expansionary (contractionary) subsample is made up of P2 and P3 (T1 and T2).

Figure 6: Partitioning into expansionary and contractionary subsamples.

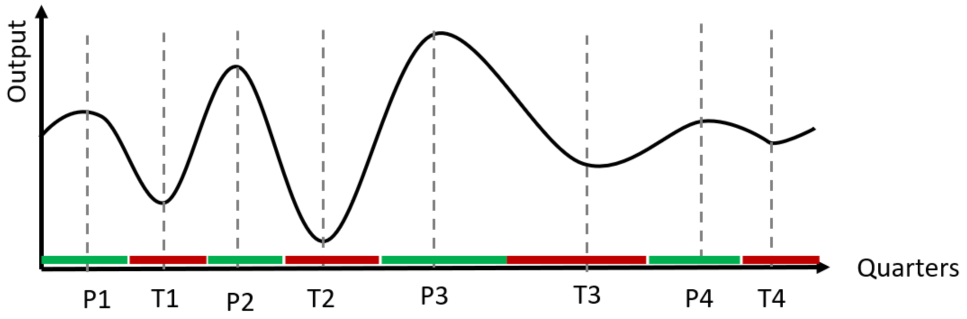

Table 2 presents the mean output gaps for different partitionings. The full sample-mean of -0.46% is shown in the first row, where all contractions and expansions are included. Averaging over subsamples that include only a share of the most extreme contractions and expansions leads to larger negative mean gaps, up to a mean of -1.12% for the most extreme 10%. Thus the asymmetry of the model is most pronounced for large shocks, which follows from the functional form of wage adjustment costs shown in Figure 2. The asymmetry between the costs of positive and negative wage adjustments increases in the size of the adjustment. As a result, large negative demand shocks are amplified disproportionately, resulting in larger downward adjustments to employment and output. This is in line with the long-standing notion of inflation 'greasing the wheels of the labor market', e.g. as formulated in Tobin (1972). In a shallow downturn, a mild decline in the real wage can be achieved by nominal wage inflation falling behind price inflation, in which case DNWR is not binding. In contrast, in a deep recession, the scope for real wage adjustment is more likely to be insufficient, especially when it is narrowed by price inflation falling steeply or even turning negative. As a result, DNWR becomes more binding and forces a greater adjustment in quantities.

Table 2: Average output gap in subsamples.

Included share of deepest contractions and highest expansions	Mean output gap in % of potential
Full sample	-0.46%
50%	-0.65%
25%	-0.90%
10%	-1.12%

3.3 Accuracy of output gap filters

This section demonstrates that for two illustrative cases, output gap estimates from standard estimation methods (‘filters’) have a mean of zero. This inevitably introduces an underestimation of potential output on average, as the true data generating process is characterized by a negative mean. The intuition is straightforward: DNWR only amplifies contractionary demand shocks, so deviations of output below true potential sum to a greater integral than deviations of output above true potential. Filter estimates do not accommodate this asymmetry, since they estimate potential such that output fluctuation around it sum up to zero. One can think of estimated potential being ’shifted down’ relative to true potential—which shrinks the integral of negative output deviations while increasing the integral of positive deviations—to an extent that makes negative and positive deviations cancel out. The asymmetry in the model’s shock adjustment from DNWR is thus picked up as a spurious reduction in estimated potential.

This "symmetry bias" translates into a bias in the economic interpretation of output swings, in that demand-driven slumps tend to be misinterpreted as resulting from a slowdown in potential.[16] As shown below, the extent of this misinterpretation increases with the severity and duration of an adverse demand shock. The intuition follows directly from the zero-mean property of filters: the larger and the more sustained a shortfall of output below true potential, the greater is the error when estimated potential is ‘dragged down’ such that the estimated gap averages to zero.[17]

The downward bias in potential estimates can help to explain the finding of Coibion, Gorodnichenko and Ulate (2018). The authors examine potential output estimates by various institutions and report that shocks that have only transitory effects on output are incorporated in potential estimates, to an extent that increases over time. For example, monetary shocks and shocks to government spending are reported to have transitory effects on GDP, but nevertheless lead to a gradual adjustment of estimated potential in the same direction. In our model analysis, the zero-mean property forces potential estimates to gradually track demand shocks that do not affect the economy’s productive potential.

We now turn to two illustrative filters for potential output, the Hodrick-Prescott (HP) filter and the IMF’s multivariate filter (MVF henceforth). The first is the most widely used among the category of purely statistical filters, while the latter will serve as an example of filters that incorporate structural economic relationships into the estimation. However, our findings apply in principle to all zero-mean filters, including e.g. the band-pass filter or production function based methods that use zero-mean filters to estimate trend components for input factors and TFP.[18] The bias is assessed by exploiting the observability of potential output in the model. That is, we apply output gap filters to the simulated data from the previous section and compare the resulting estimates with the true output gap series. As outlined in the introduction, this approach sets our study apart from conventional studies in that it scrutinizes the plausibility of final output gap estimates, instead of the convergence of real-time estimates.

---

[16] The downward bias in potential estimates also skews the interpretation of upswings towards being driven by transitory demand fluctuations. However, as discussed later on, this is quantitatively less relevant than the misinterpretation of slowdowns.

[17] As an illustration consider two periods with a negative demand shock in *t=1* and a positive demand shock of the same size in *t=2* . Due to asymmetric shock adjustment, the resulting output gaps are -2% in *t=1* and +1% in *t=2* (potential remains constant), so the integral between output and true potential is -1%. For the integral between output and estimated potential to be zero, filters attribute some of the output shortfall in *t=1* to an adverse supply shock lowering potential. For example, when half of it is attributed to a supply shock, the gap in *t=1* is estimated at -1%, and the integral sums to zero.

[18] Band-pass filtering of historical real GDP data for G7 countries (with samples from the earliest available data on FRED to 2018Q2, and using 6 to 32 quarters as business cycle frequency) yields an average output gap of -0.05%. The European Commission’s production function based estimates for all countries in the AMECO database average to -0.16% (from the earliest available data to 2019).

3.4 The Hodrick-Prescott Filter

When output is decomposed using the HP filter, the trend component is interpreted as potential output, and its distance to actual output—the cyclical component—as the output gap. We prove in the appendix that the average of the cyclical component is zero for any sample size. The upper panel in Figure 7 again shows an illustrative subsample of simulated actual and potential output, as well as the trend component (i.e. estimated potential) of an HP-filtering using $\lambda = 1600$ for quarterly frequency (as suggested by Ravn and Uhlig, 2002). The figure illustrates that during periods of significant and sustained shortfalls of output below true potential, filter-estimated potential tends to undershoot true potential. That is, potential is underestimated during demand-driven slumps so that it cuts through actual output such that deviations average to zero. The reverse holds in periods of sustained positive gaps when the estimate overshoots true potential, as e.g. between periods 20 and 60. But positive gaps tend to be smaller and, owing to DNWR, less sustained than negative gaps, so that potential output is underestimated on average.

Figure 7: Accuracy of HP filter.

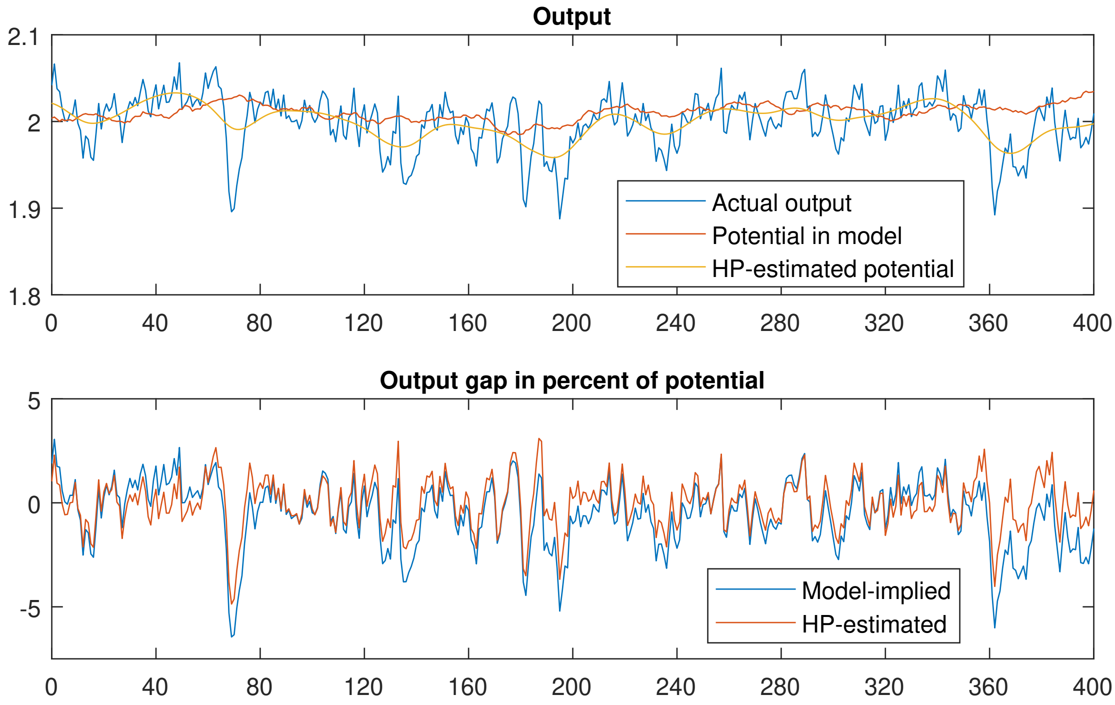

The first three columns of Table 3 show true and estimated mean gaps, as well as the resulting mean bias, in expansionary and contractionary subsamples. The latter are selected with the algorithm introduced in section 3.2, with the difference that averages are reported for expansions and contractions separately. For example, the true mean gap is 0.19% for the set of all expansions and -1.13% for the set of all contractions (averaging over both yields a full sample-mean of -0.46% for the true gap, 0% for the estimated gap, and accordingly -0.46% for the bias). During contractions, the filter underestimates potential output and with it the size of the negative gap. The bias is larger for deeper contractions, averaging a very substantial

2.69% during the 10% of deepest recessions. That is, during those periods when an accurate assessment of the cyclical position is most relevant for policymakers, the filter performs poorly. For expansions the bias is smaller and has no clear sign across subsamples, which is explained by DNWR only amplifying negative deviations of output from potential. The estimation bias also surfaces in a difference between the points in time when the true and the estimated gaps turn positive in the recovery from a contraction, or negative in the slowdown following an expansion. The last column reports the number of quarters by which the filter mis-estimates the closing of the gap. The bias is again mild for expansions (staying below one year even during the most extreme 10% of expansions), but much more substantial for contractions. On average, the filter indicates that the economy has reached potential about one year too soon. During the most severe (and long-lived) 10% of contractions, the filter signals that potential has been reached about 3.5 years before the true gap closes.

Table 3: HP estimation bias.

	True gap	Estimated gap	Bias	Estimated gap closing
Expansionary Subsamples				
All expansions	0.19%	0.34%	0.14%	0.4 quarters too soon
50% highest	0.67%	0.25%	-0.41%	2.0 quarters too soon
25% highest	0.88%	0.22%	-0.66%	3.1 quarters too soon
10% highest	1.04%	0.19%	-0.85%	3.5 quarters too soon
Contractionary Subsamples				
All recessions	-1.13%	-0.34%	0.78%	4.1 quarters too soon
50% deepest	-1.98%	-0.43%	1.54%	8.2 quarters too soon
25% deepest	-2.68%	-0.53%	2.14%	11.4 quarters too soon
10% deepest	-3.28%	-0.58%	2.69%	14.0 quarters too soon

A positive bias value indicates underestimation of potential.

3.5 The IMF's multivariate filter

The IMF's multivariate filter (MVF), described in Blagrave et al. (2015), is an example of a more elaborate filtering methodology that is not purely statistical but features economic structure. The MVF embeds an Okun's law relationship as well as a Phillips curve to inform the estimation of the output gap. In contrast to the HP filter, estimates obtained from the MVF do not mechanically force a mean of zero for all possible data samples. However, MVF estimates have a mean of virtually zero (0.006%) for our simulated data, and also for historical data for the G7 countries.[19] The output gap in the filter is modelled as as $gap_t = \phi \, gap_{t-1} + \epsilon_t$ ($0 < \phi < 1$) with a prior distribution of ϕ that is typically centred around a mode of 0.6. As a result, gap_t tends to gravitate strongly towards its unconditional mean of zero.

Since the filter is designed for an annual frequency, we annualized the simulated quarterly time series for the three variables required to run the filter: output, the unemployment rate, and inflation.[20] Figure 8 again shows an illustrative subsample, this time including MVF-estimated

---

[19] We applied the MVF to historical data for the G7 countries, with samples starting at the earliest available observation and going to 2018. The average output gap estimate is -0.08%.

[20] Every four consecutive quarters of simulated data are averaged into one year of annual data.

potential and the output gap. As in the case of the HP filter, MVF-estimated potential tends to be below true potential, especially when there are large and persistent negative output gaps. Table 4 shows the mean estimation bias in the same way as the previous table.[21] While the average full-sample bias is the same as for the HP filter—both estimate an average gap of close to zero—the MVF generally performs significantly better, suggesting that incorporating economic structure improves the estimate's accuracy. However, the bias is still substantial during severe contractions: during the deepest 10% of contractions, roughly a third of the true output gap of -3.10% is not picked up by the filter. The same holds for the mis-estimation of the point in time when the economy has fully recovered from a contraction. For example, for the deepest 25% of contractions, filter estimates show the output gap closing about 3 years too soon.

Figure 8: Accuracy of IMF's multivariate filer.

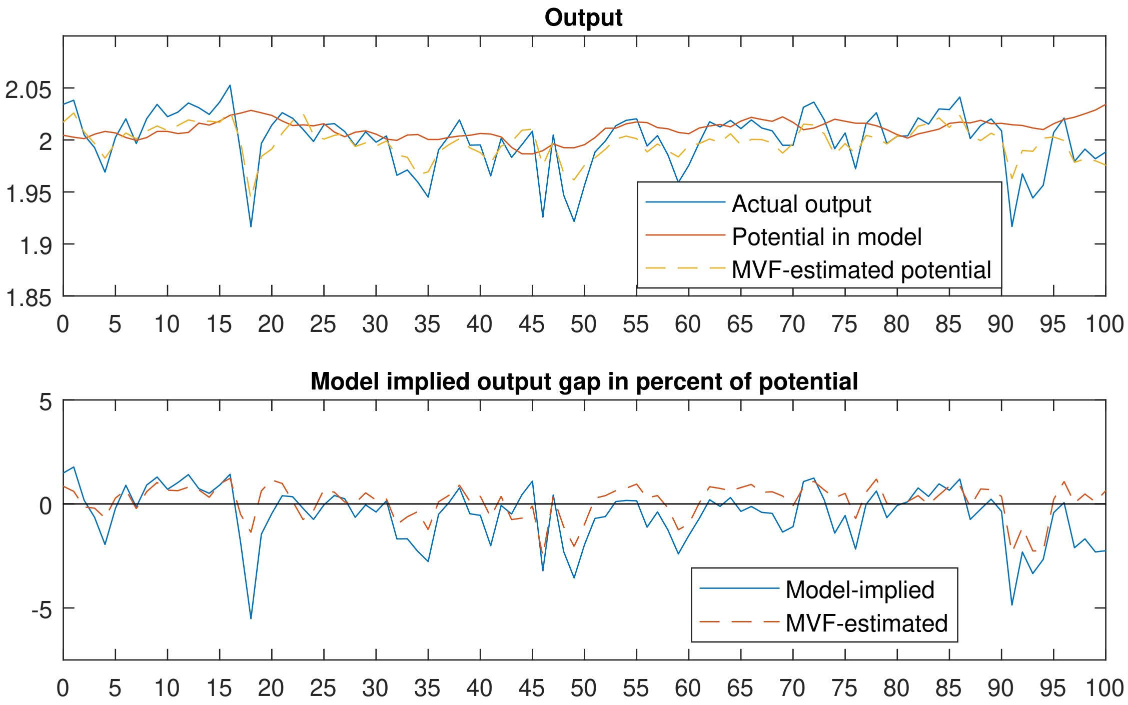

---

[21] Small differences in the true gap averages between Table 3 and 4 result from the annualization of the data.

Table 4: MVF estimation bias.

	True gap	Estimated gap	Bias	Estimated gap closing
Expansionary Subsamples				
All expansions	0.22%	0.69%	0.48%	0.62 years too late
50% highest	0.70%	1.06%	0.36%	0.53 years too late
25% highest	0.89%	1.21%	0.31%	0.40 years too late
10% highest	1.12%	1.38%	0.26%	0.15 years too late
Contractionary Subsamples				
All recessions	-1.16%	-0.68%	0.48%	1.54 years too soon
50% deepest	-1.95%	-1.33%	0.62%	2.59 years too soon
25% deepest	-2.54%	-1.81%	0.73%	3.02 years too soon
10% deepest	-3.10%	-2.24%	0.86%	3.48 years too soon

A positive bias value indicates underestimation of potential.

4 Policy implications of the bias

Closing the output gap—often formulated equivalently as closing the unemployment gap with respect to the NAIRU—is in many countries part of the central bank's mandate. The gap can also matter for fiscal policy, either when it is directed at taming the cycle, or in the context of fiscal rules such as the European Union's Stability and Growth Pact, which is formulated in terms of cyclically adjusted fiscal balances. When the output gap informs countercyclical policy, the symmetry bias of filtering techniques can potentially impede its calibration and hence its effectiveness. To explore to what extent this is the case, we compare simulations with different monetary and fiscal policies rules that take into account either the true gap or a real-time estimate.

Since almost all output gap filters use some type of forecast as input, taking a real-time perspective when using a gap estimate in a policy rule requires making an assumption on how this forecast is formed. Instead of relying on one particular assumption, we report our results for two illustrative polar cases. In the first, to obtain a real-time estimate for a period T, a 2-sided HP filter is applied on the full simulated output series. As the filter's input encompasses data from after T, this corresponds to a perfectly accurate output forecast. This perfect-foresight case is contrasted with the use of a 1-sided HP filter which takes as input only data up to period T. The results obtained for both cases are broadly in line with each other.

A technical difficulty arising when real-time estimates are used in a policy rule is that these estimates are endogenous: they affect the policy stance and thereby output, i.e. the estimation sample. Estimates must therefore be solved for jointly with the rest of the model. Instead of augmenting our model by the respective filter, we employ a simple iterative procedure that achieves the same outcome.[22] Simulating policy implications for the MVF is beyond the scope

---

[22] The procedure involves multiple model simulations based on the same randomly generated shock series. Step 1 simulates the model with flexible prices and wages to obtain an initial series of true potential (since there is no gap, there is also no countercyclical fiscal policy and the Taylor Rule is based on inflation only). In the next simulation in step 2, this series of true potential is used in the output gap in the respective policy rule (either a fiscal rule or an extended Taylor Rule). At the beginning of step 3, the simulated output series from the last step is filtered to obtain the next series of potential output. It is then used in the output gap estimate in the

of this paper, as this filter is constructed for annual data while the model is quarterly.

4.1 Monetary policy

Parameter ω_{gap} in Taylor Rule (3) is set to 0.125, so that the monetary policy stance is jointly determined by inflation and the output gap. To shed light on the implications of using an estimate instead of the true gap, we first examine an illustrative subsample of the generated data, then consider peak and trough events, and finally present descriptive statistics. Figure 9 shows the illustrative subsample. Green lines depict the case when the gap in the Taylor Rule uses true potential, while dashed red (dashed orange) lines indicate the use of potential estimates from a 1-sided (2-sided) HP filter. The different panels share the horizontal time axis, in which each fourth quarter is marked by a dotted vertical line. The bottom panel shows the policy-relevant output gap used in the respective Taylor Rule, and the middle panel depicts the resulting nominal interest rates. The top panel shows actual output for the three cases, as well as potential output (identical across the three scenarios). The demand-induced recession beginning around quarter 1815 and lasting roughly 5 years illustrates the bias's impact on policy calibration. In the bottom panel, the bias surfaces in the policy-relevant output gaps in the five years between 1820 and 1840. While the true output gap remains negative throughout this period, the filters' zero-mean property forces estimated potential to shift downwards (and estimated gaps upwards) to the point of bringing about periods of positive gap estimates. This bias weakens the expansionary monetary policy stance (middle panel) and thereby the cushioning of the recession, relative to the policy stance under an unbiased estimate (top panel).

---

respective policy rule in this step's simulation.

Step 3 is then repeated, each time taking a filtering of last simulation's output as the potential output estimate on which the output gap estimate in the policy rule is based. In each iteration, the distance between the newly filtered output series and last iteration's filtered output becomes smaller, and convergence is reached after roughly 10 iterations. When convergence is achieved the potential estimate underlying the policy rule is consistent with a filtering of equilibrium output–that is, the model allocation validates a real-time estimate of the output gap.

Figure 9: Impact of bias on monetary policy.

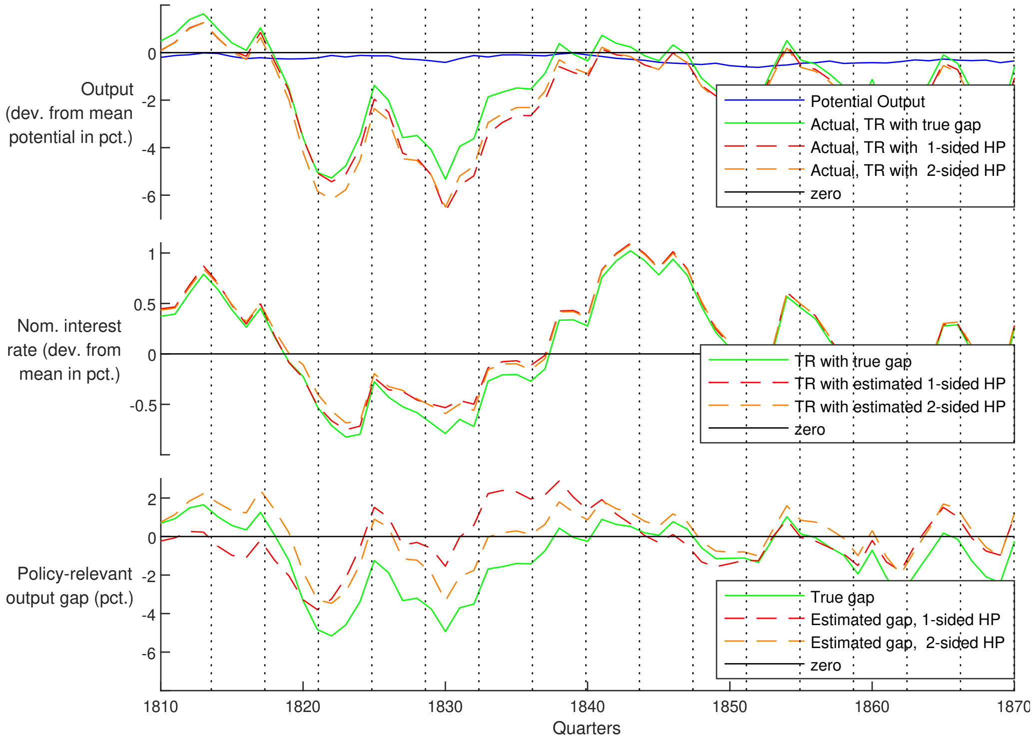

To quantify policy implications more systematically , we conduct an event-style analysis of average business cycle peaks and troughs. In the first step, we again apply the Harding-Pagan algorithm to identify peaks and troughs in the simulated data. To compute the peak event, the sample is partitioned into subsamples, each ranging from one trough to the next (so that they all contain one peak). All subsamples are then centered around their peaks (so that peaks occur at t=0 in all them). The final step averages across all subsamples. By the same token, for the trough event, the sample is partitioned into subsamples ranging from peak to peak, which are then centered around their troughs before we average across them. As in Section 3.3, we analyze not only the average peak and trough in the full sample, but also analyze the average peak (trough) corresponding to a certain share of the largest expansions (deepest contractions) in the sample.

Figure 10 depicts the same variables (in the same colors) as the previous figure. The left (right) column shows the average peak (trough) for the subset of the 25% largest expansions (deepest contractions). The lower right panel shows that these deep troughs are associated with a true output gap of about -5%, which gradually closes over the course of about 20 quarters. The filter bias again causes an erroneous downward-shift of estimated potential, thereby pushing up gap estimates. Here, the 1-sided HP filter is highly inaccurate in that the gap is estimated to turn positive shortly after the trough, while the 2-sided filter is slightly more accurate, but also estimates the gap switching to positive after roughly a year. The nominal interest rate panel shows that these biases dampen the policy's expansionary stance, and the top panel depicts the resulting output costs—monetary policy forgoes its chance to significantly accelerate the average recovery from rare and deep recessions. The left panel shows that the bias is quantitatively less significant around peaks. With DNWR, output gaps caused by expansionary shocks are comparably weaker and short-lived, so the misestimation of potential is smaller and less consequential for policy.

Figure 10: Monetary policy, 25% most extreme peaks (left) and troughs (right).

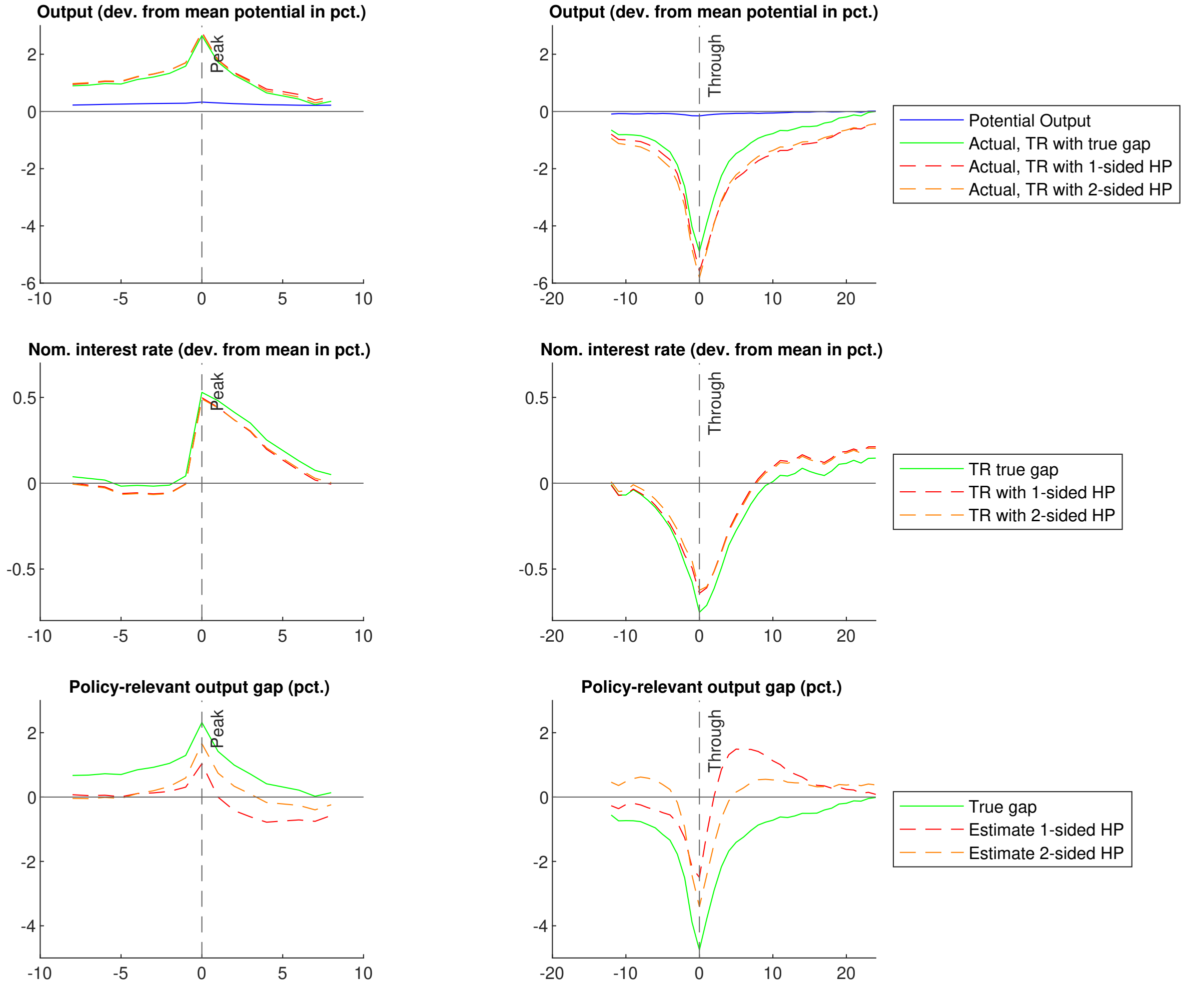

Next we compute a statistic to quantify how strongly recoveries are impeded by the bias. From the perspective of the trough period $t = 0$, we compute the discounted PV of the output gap going forward, i.e. the average deviation of actual output from potential during the time between t=0 and the closing of the output gap. This is a measure of "lost output", appropriately discounted. These PVs are shown in Table 5 and are reported as a share of mean potential output. We present values for the full sample as well as for the deepest 10%, 25% and 50% of contractions. When all troughs are considered, the bias only mildly impedes recoveries, raising the average PV of the gap from -7% if the true gap was known to -10.9% (-10.3%) if it is estimated by a 1-sided (2-sided) HP filter. However, the consequences of relying on biased estimates become substantially more grave when we consider the deepest half of troughs. Here the average PV is -17.4% under the true gap, and increases to roughly -30% for the two filters.

The implications of the bias are even more extreme when only the steepest quarter of recessions are considered—raising the PV from about -24% to roughly -40%. For rare deep recessions defined as the worst 10%, the bias comes close to doubling the PV of the gap.

Table 5: Output gap PVs for true and estimated gaps in the Taylor Rule.

Share deepest troughs	True gap	1-sided HP	2-sided HP
Full sample	-7.0	-10.9	-10.3
50%	-17.4	-30.0	-29.1
25%	-24.1	-40.7	-39.5
10%	-32.0	-55.3	-53.7

The present values of output gaps are computed from the perspective of the trough period and are expressed in % of mean potential.

4.2 Fiscal policy

The previous analysis is now repeated for fiscal policy. Under the illustrative countercyclial fiscal rule in equation (4), government spending is zero when the output gap is closed, and moves in the opposite direction from the gap when it opens. Although spending is financed by lump-sum taxes in every period, the fiscal multiplier is sufficiently large to give fiscal policy a role in stabilizing the economy. As discussed in Monacelli, Perotti and Trigari (2010), the fiscal multiplier in models with labor search and matching frictions is boosted by the introduction of imperfect competition and nominal rigidity. The main channel is that rigidity introduces counter-cyclicality in price markups that affects the surplus of new worker-firm matches via the marginal product of labor. Figure 12 in the appendix shows that in our model, the peak fiscal multiplier for short-lived changes in government spending is about 0.5. We consider an illustrative calibration of the fiscal rule's responsiveness of $\phi = 0.5$, and assume that the Taylor Rule is based on inflation only ($\omega_{\text{gap}} = 0$). Figure 13 in the appendix repeats the analysis in Figure 9 and reaches a similar conclusion: the symmetry bias of filters dilutes the expansionary stance in the illustrative sample period, and prematurely induces a tightening when the true output gap is still negative. This policy mistake exaggerates the depth of the recession.

Turing to the event-style analysis in Figure 11 (again only considering the 25% of most extreme contractions / expansions), the picture is similar to monetary policy. The bias has only minor implications for peaks, but it amplifies the depth of troughs and impedes the recoveries. The reason is that the gap is estimated to turn positive prematurely, triggering a fiscal tightening when output is still below potential. The present values shown in Table 6 are again broadly similar to those for monetary policy, although the general implications of the bias are mildly weaker. Once again, the bias matters most for the deepest recessions.

Figure 11: Fiscal policy, 25% most extreme peaks (left) and troughs (right).

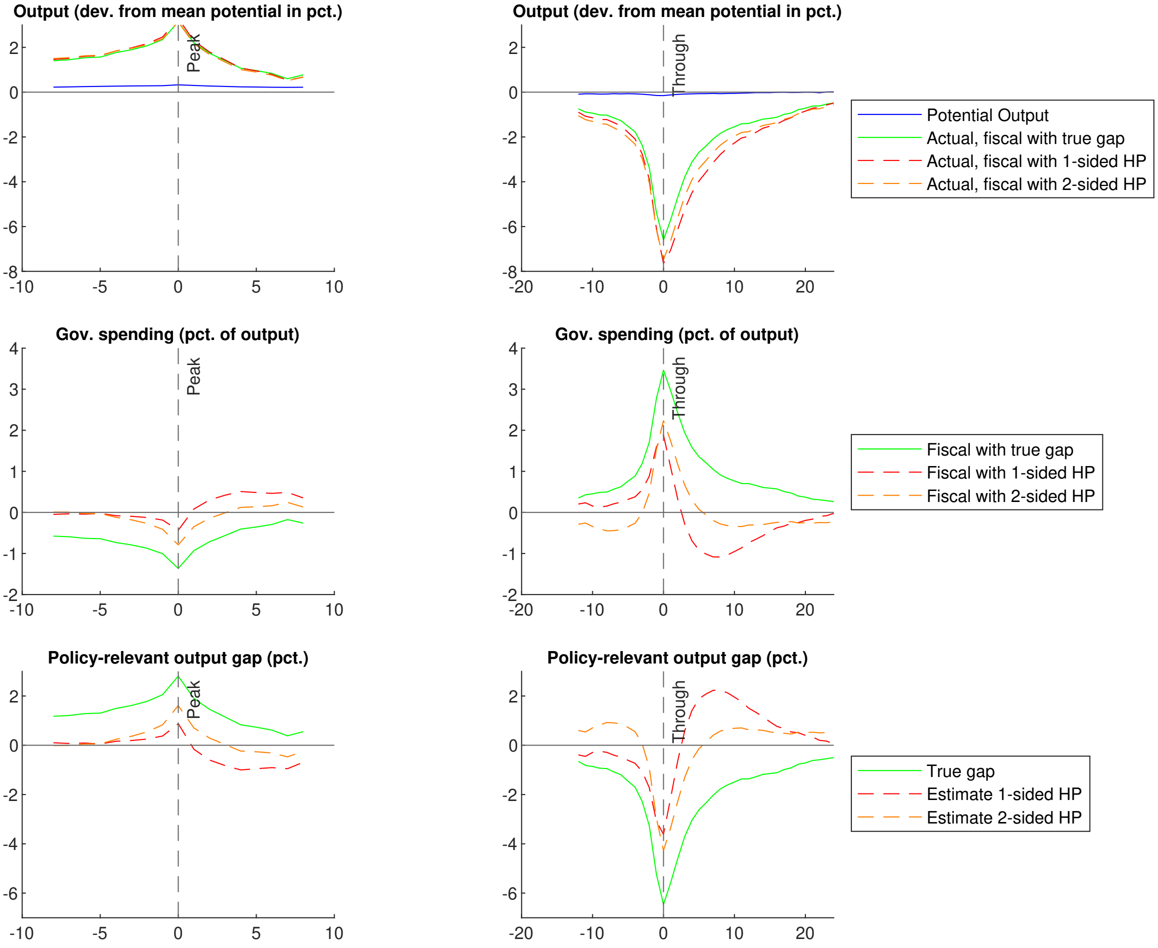

Table 6: Output gap PVs for true and estimated gaps in the fiscal rule.

Share deepest troughs	True gap	1-sided HP	2-sided HP
Full sample	-10.5	-15.5	-13.0
50%	-28.8	-38.3	-34.4
25%	-44.6	-58.5	-53.7
10%	-63.2	-83.6	-77.8

The present values of output gaps are computed from the perspective of the trough period and are expressed in % of mean potential.

5 Conclusion

There is no theoretical reason to insist that positive deviations of output from potential must on average equal negative deviations. Indeed, Okun's 1962 paper conceived of the output gap (the 'GNP gap' as he called it) as intrinsically negative; with employment sometimes falling below full employment but not above it. The same holds in Milton Friedman's 'plucking model' that was recently formalized by Dupraz, Nakamura and Steinsson (2019). This simply enshrined the intuition that while unemployment is a commonplace, observable phenomenon, 'overemployment' is a much more theory-laden concept. We would argue that this remains the case. Presumably there are sharp physical limits to the number of extra hands that can usefully be employed with a given stock of capital, capping the magnitude of positive deviations from potential output relative to negative deviations.

In this paper, however, we make no appeal to such arguments. Instead we demonstrate that DNWR in a standard NK model with labor search frictions inevitably leads to greater adjustment of output and employment in downturns than in booms. If positive and negative demand shocks are distributed symmetrically, this implies a negative mean output gap.

We further show that common filtering techniques are intrinsically biased in a world characterized by asymmetric shock adjustment and a negative mean output gap. Both purely statistical filters such as the HP filter, and other techniques that incorporate greater economic structure, such as the IMF's MVF, in practise impose a zero mean on the estimated gap. This causes the filters to erroneously underestimate potential output, especially during severe downturns, an error that we call "symmetry bias".

Our work suggests that the best way forward for estimating cyclical deviations from potential output lies in methods that do not impose a zero mean. For example, Coibion, Gorodnichenko and Ulate (2018) suggest stacking up identified supply shocks (using the identification strategy of Blanchard and Quah, 1989) to construct a series for potential that is not stationary around actual output. This approach is not subject to our critique as it can accommodate a negative mean gap, but requires choosing a starting point in time when the output gap is closed.

To shed light on the implications of the symmetry bias for countercyclical monetary and fiscal policy, we compare simulations where policy calibration is based on the true gap with simulations where it is based on a filter estimate. The bias results in an inappropriately weak policy response to contractions—especially severe contractions—and premature tightening, both substantially aggravating the contraction's output costs. Existing studies on the policy implications of output gap uncertainty report much weaker costs, because they implicitly rule out any impact on the model's first moments through their reliance on symmetric models.

In general, our results indicate an urgent need for research on output gap estimation and on optimal policy-setting in a negative mean output gap world, as opposed to the zero mean output gap world implicitly or explicitly assumed by policymaking institutions. This raises a host of questions: for example, if economies spend more time below full employment than above it, a fiscal policy aiming to be symmetrical around full employment will be expansionary more often than contractionary, and may therefore imply explosive debt dynamics. As another example, Taylor Rules for monetary policy are likely to differ significantly if time series for the output gap are allowed to have negative means rather than having a zero mean imposed through filtering.

More fundamentally, our work here illustrates the perils for macroeconomic policy of relying too heavily on output gap estimates. If the economy is characterized by a negative mean output gap, as we argue, then standard filtering technique will be intrinsically biased, and the bias will be greater during strong contractions when stabilization policies are most required. This at least suggests that policy should place greater reliance on observables, such as inflation, unemployment and the level of public debt.

One caveat of our analysis is that it abstracts from hysteresis effects, which could cause potential output to decline as a result of long-lasting negative demand shocks, and thereby validate the bias. But it would require an extremely unlikely coincidence for hysteresis effects to exactly cancel out the asymmetry induced by downward nominal wage rigidity, leaving the output gap centered at precisely zero. Indeed, our work shows that the repeated downward revisions of estimated potential since 2008 are plausibly a consequence of the zero-mean property of filters; this should be taken into account when the revisions are instead interpreted as evidence of hysteresis.

Appendix: Proof for a zero-mean HP cyclical component

We first show that the HP filter can be written in a linear form. The filter computes the trend component $\{\hat{m}_t\}$ of a time series $\{y_t\}$ as solution to the following minimization problem:

$$\arg\min_{\{m_t\}} \sum_{t=1}^{T} (y_t - m_t)^2 + \lambda \sum_{t=3}^{T} (m_t - 2m_{t-1} + m_{t-2})^2$$

which, in vector notation, can be written as

$$\arg\min_{\mathbf{m}} (\boldsymbol{y} - \boldsymbol{m})^{'} (\boldsymbol{y} - \boldsymbol{m}) + \lambda (\boldsymbol{B}\boldsymbol{m})^{'} (\boldsymbol{B}\boldsymbol{m})$$

with $\underset{(T\times 1)}{\boldsymbol{y}} = \begin{pmatrix} y_1 \\ \vdots \\ y_T \end{pmatrix}$, $\underset{(T\times 1)}{\boldsymbol{m}} = \begin{pmatrix} m_1 \\ \vdots \\ y_T \end{pmatrix}$ and $\underset{(T-2\times T)}{\boldsymbol{B}} = \begin{pmatrix} 1 & -2 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & -2 & 1 & 0 & \cdots & 0 \\ \vdots & & & & \ddots & & \vdots \\ 0 & & \cdots & 0 & 1 & -2 & 1 \end{pmatrix}$.

Setting the first derivative to zero yields the following solution for the vector of the estimated trend, $\hat{\boldsymbol{m}}$:

$$\hat{\boldsymbol{m}} = \left(\boldsymbol{I} + \lambda \boldsymbol{B}^{'} \boldsymbol{B}\right)^{-1} \boldsymbol{y} \tag{5}$$

Note that $\left(\boldsymbol{I} + \lambda \boldsymbol{B}^{'} \boldsymbol{B}\right)^{-1}$ is symmetric because $\boldsymbol{B}$ is.

In the following we prove that the estimated cyclical components $\hat{u}_t = y_t - \hat{m}_t$ have a mean of zero. Let $\mathbf{1} = (1 \cdots 1)'$ denote the unit vector and note that $\boldsymbol{B}\mathbf{1} = \mathbf{0}$ because all rows in $\boldsymbol{B}$ sum up to zero. It also holds that $\left(\boldsymbol{I} + \lambda \boldsymbol{B}'\boldsymbol{B}\right)\mathbf{1} = \mathbf{1}$, which can be verified by rewriting the LHS as $\boldsymbol{I}\mathbf{1} + \lambda \boldsymbol{B}'\boldsymbol{B}\mathbf{1} = \mathbf{1} + \lambda \boldsymbol{B}'\mathbf{0} = \mathbf{1}$. Multiplying from the left by $\left(\boldsymbol{I} + \lambda \boldsymbol{B}'\boldsymbol{B}\right)^{-1}$ and transposing yields

$$\mathbf{1}' = \mathbf{1}' \left(\boldsymbol{I} + \lambda \boldsymbol{B}'\boldsymbol{B}\right)^{-1}. \tag{6}$$

The mean of estimated trend components is given by $\bar{\hat{m}}_t = \frac{1}{T}\sum_{t=1}^{T} \hat{m}_t = \frac{1}{T}\mathbf{1}'\hat{\boldsymbol{m}}$. Using (5) it can be written as $\bar{\hat{m}}_t = \frac{1}{T}\mathbf{1}' \left(\boldsymbol{I} + \lambda \boldsymbol{B}'\boldsymbol{B}\right)^{-1} \boldsymbol{y}$, which simplifies to $\bar{\hat{m}}_t = \frac{1}{T}\mathbf{1}'\boldsymbol{y} = \bar{\hat{y}}_t$ using (6). It follows that the mean of the estimated cyclical component is zero: $\bar{\hat{u}} = \frac{1}{T}\sum_{t=1}^{T} \hat{u}_t = \frac{1}{T}\sum_{t=1}^{T} (y_t - \hat{m}_t) = \frac{1}{T}\sum_{t=1}^{T} y_t - \frac{1}{T}\sum_{t=1}^{T} \hat{m}_t = \bar{\hat{y}}_t - \bar{\hat{y}}_t = 0$.

Figure 12: Model Adjustment to a transitory 1% GDP government spending shock.

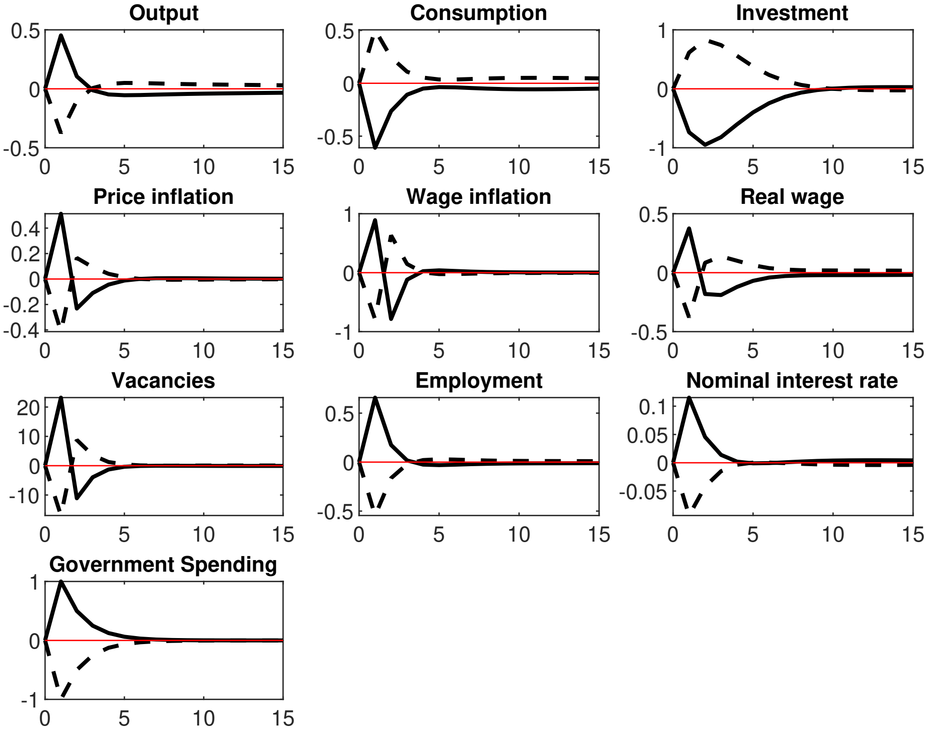

Figure 13: Impact of bias on fiscal policy.

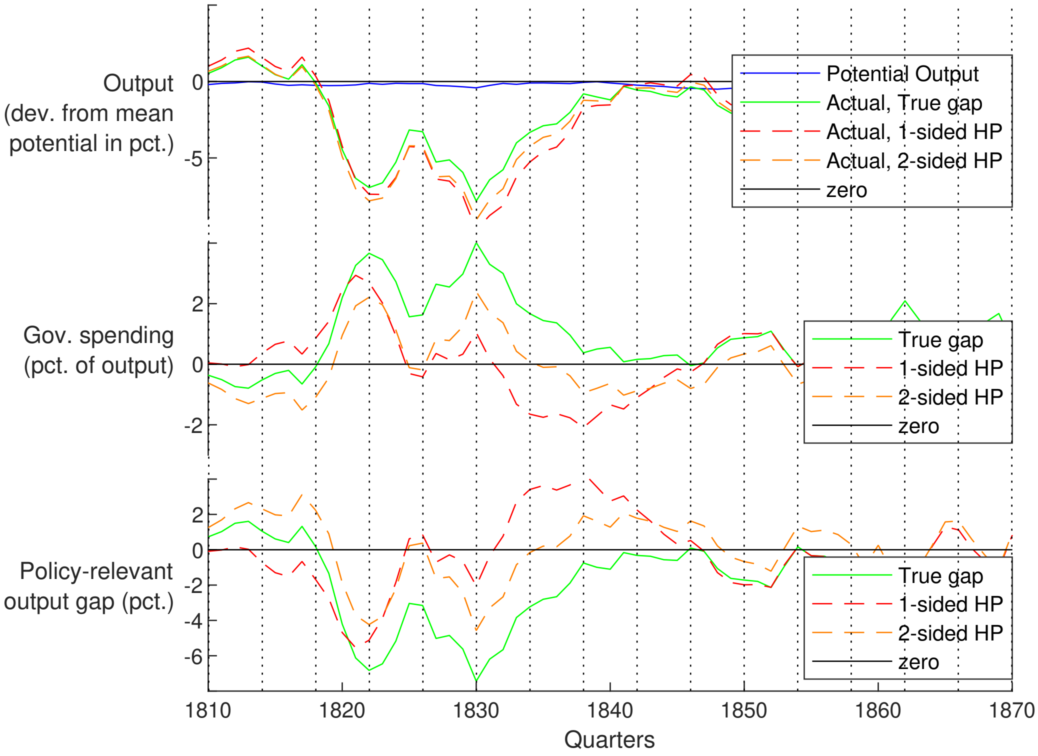

In the first half of the illustrative recession, the filters underestimate the extent of the negative gap, while in the second half of the recession, the zero-mean property forces estimated potential to shift downward to an extent that implies positive output gap estimates. As a result, fiscal policy in the first half is less expansionary than it would be if the true gap was known, while in the second half it even turns contractionary. Both exaggerates the depth of the recession.

References

Abbritti, Mirko, and Stephan Fahr. 2013. "Downward wage rigidity and business cycle asymmetries." *Journal of Monetary Economics*, 60(7): 871–886.

Arseneau, David M, and Sanjay K Chugh. 2008. "Optimal fiscal and monetary policy with costly wage bargaining." *Journal of Monetary Economics*, 55(8): 1401–1414.

Ball, Laurence. 2014. "Long-term damage from the Great Recession in OECD countries." *European Journal of Economics and Economic Policies: Intervention*, 11(2): 149–160.

Blagrave, Patrick, Mr Roberto Garcia-Saltos, Mr Douglas Laxton, and Fan Zhang. 2015. "A simple multivariate filter for estimating potential output." , (WP/15/79).

Blanchard, Olivier, and Danny Quah. 1989. "The Dynamic Effects of Aggregate Demand and Supply Disturbances." *The American Economic Review*, 79(4): 655–673.

Blanchard, Olivier, and Jordi Galí. 2010. "Labor markets and monetary policy: A New Keynesian model with unemployment." *American economic journal: macroeconomics*, 2(2): 1–30.

Boehm, Christoph E, and Christopher L House. 2019. "Optimal Taylor rules when targets are uncertain." *European Economic Review*, 119: 274–286.

Clarida, Richard, Jordi Gali, and Mark Gertler. 1999. "The science of monetary policy: a new Keynesian perspective." *Journal of economic literature*, 37(4): 1661–1707.

Coeurdacier, Nicolas, Helene Rey, and Pablo Winant. 2011. "The risky steady state." *American Economic Review*, 101(3): 398–401.

Coibion, Olivier, Yuriy Gorodnichenko, and Mauricio Ulate. 2018. "The cyclical sensitivity in estimates of potential output." *Brookings Papers on Economic Activity.*

Daly, Mary C, Bart Hobijn, Timothy Ni, et al. 2013. "The path of wage growth and unemployment." *FRBSF Economic Letter*, 20: 1–5.

Dickens, William T, Lorenz Goette, Erica L Groshen, Steinar Holden, Julian Messina, Mark E Schweitzer, Jarkko Turunen, and Melanie E Ward. 2007. "How wages change: micro evidence from the International Wage Flexibility Project." *Journal of Economic Perspectives*, 21(2): 195–214.

Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson. 2019. "A plucking model of business cycles." *National Bureau of Economic Research Working Paper.*

Fahr, Stephan, and F Smets. 2010. "Downward wage rigidities in a monetary union with optimal monetary policy." *The Scandinavian Journal of Economics*, 112(4): 812–840.

Fatás, Antonio. 2018. "Fiscal Policy, Potential Output, and the Shifting Goalposts." *IMF Economic Review*, 1–19.

Gopinath, Gita, Emine Boz, Camila Casas, Federico J Díez, Pierre-Olivier Gourinchas, and Mikkel Plagborg-Møller. 2020. "Dominant currency paradigm." *American Economic Review*, 110(3): 677–719.

Kangur, Alvar, Koralai Kirabaeva, Jean-Marc Natal, and Simon Voigts. 2019. "How Informative Are Real Time Output Gap Estimates in Europe?" *IMF Working Paper, forthcoming.*

Kempkes, Gerhard. 2012. "Cyclical adjustment in fiscal rules: Some evidence on real-time bias for EU-15 countries."

Leduc, Sylvain, and Zheng Liu. 2016. "Uncertainty shocks are aggregate demand shocks." *Journal of Monetary Economics*, 82: 20–35.

Marcellino, Massimiliano, and Alberto Musso. 2011. "The reliability of real-time estimates of the euro area output gap." *Economic Modelling*, 28(4): 1842–1856.

McCallum, Bennett T. 2001. "Should monetary policy respond strongly to output gaps?" *American Economic Review*, 91(2): 258–262.

Monacelli, Tommaso, Roberto Perotti, and Antonella Trigari. 2010. "Unemployment fiscal multipliers." *Journal of Monetary Economics*, 57(5): 531–553.

Orphanides, Athanasios, and Simon van Norden. 2002. "The unreliability of output-gap estimates in real time." *Review of economics and statistics*, 84(4): 569–583.

Orphanides, Athanasios, Richard D Porter, David Reifschneider, Robert Tetlow, and Frederico Finan. 2000. "Errors in the Measurement of the Output Gap and the Design of Monetary Policy." *Journal of Economics and Business*, 52(1-2): 117–141.

Petrongolo, Barbara, and Christopher A Pissarides. 2001. "Looking into the black box: A survey of the matching function." *Journal of Economic literature*, 39(2): 390–431.

Ravenna, Federico, and Carl E Walsh. 2011. "Welfare-based optimal monetary policy with unemployment and sticky prices: A linear-quadratic framework." *American Economic Journal: Macroeconomics*, 3(2): 130–62.

Ravn, Morten O, and Harald Uhlig. 2002. "On adjusting the Hodrick-Prescott filter for the frequency of observations." *Review of economics and statistics*, 84(2): 371–376.

Segal, Guy. 2017. "To Respond or Not to Respond: Measures of the Output Gap in Theory and in Practice." *49th issue (June 2017) of the International Journal of Central Banking.*

Smets, Frank. 2002. "Output gap uncertainty: does it matter for the Taylor rule?" *Empirical Economics*, 27(1): 113–129.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and frictions in US business cycles: A Bayesian DSGE approach." *American economic review*, 97(3): 586–606.

Tobin, James. 1972. "Inflation and Unemployment." *American Economic Review*, 62.