STATEC

Économie & Statistiques
Working papers

N°127 Update 08/2023

Inequality Constrained Output Gap*

Auteurs :
Vasja (STATEC)

Abstract

In this paper, we introduce an output gap model that uses inequality constraints to estimate it. Inequality constraints modify the output gap during periods of inconsistency with related macroeconomic variables but allow it to follow the output's dynamics freely during consistent periods. We evaluate our approach through simulation experiments and a large empirical analysis, comparing it to traditional methods like the simple Hodrick-Prescott filtered output gap and the European Commission's output gap estimates. Our findings show that inequality constraints enhance the correlation coefficient between related macroeconomic variables and the output gap, while incurring only a marginal increase in model complexity. Furthermore, our model can reduce revisions and improve the forecasting power of the output gap over more complex models, demonstrating that it is a promising alternative to existing models for estimating the output gap.

Keywords: potential output, output gap, fiscal policy, Kalman filter, production function, inequality constraints
JEL Codes : C32, E32, E37

*We are grateful to Alessandro Rossi, Peter Moffat, participants at the joint OGWG - ECFIN - JRC Conference: "Assessment of output gaps and potential output in the context of the COVID-19 pandemic and its aftermath" and seminar participants at STATEC Scientific Council and SEMECO series. Views and opinions expressed in this article are those of the author(s) and do not reflect those of STATEC and funding partners. The author(s) gratefully acknowledge the support of the Observatoire de la compétitivité, Ministre de l'Economie, DG Comptitivit, Luxembourg, and STATEC, the National Statistical Office of Luxembourg.
STATEC Research, 13 Rue Erasme, 1468 Luxembourg, Luxembourg, e-mail: vasja.sivec@statec.etat.lu

Inequality Constrained Output Gap*

Vasja Sivec†

This version: August 2023

Abstract

In this paper, we introduce an output gap model that uses inequality constraints to estimate it. Inequality constraints modify the output gap during periods of inconsistency with related macroeconomic variables but allow it to follow the output's dynamics freely during consistent periods. We evaluate our approach through simulation experiments and a large empirical analysis, comparing it to traditional methods like the simple Hodrick-Prescott filtered output gap and the European Commission's output gap estimates. Our findings show that inequality constraints enhance the correlation coefficient between related macroeconomic variables and the output gap, while incurring only a marginal increase in model complexity. Furthermore, our model can reduce revisions and improve the forecasting power of the output gap over more complex models, demonstrating that it is a promising alternative to existing models for estimating the output gap.

Keywords: potential output, output gap, fiscal policy, Kalman filter, production function, inequality constraints

JEL-Codes: C32, E32, E37

*We are grateful to Alessandro Rossi, Peter Moffat, participants at the joint OGWG - ECFIN - JRC Conference: "Assessment of output gaps and potential output in the context of the COVID-19 pandemic and its aftermath" and seminar participants at STATEC Scientific Council and SEMECO series. Views and opinions expressed in this article are those of the author(s) and do not reflect those of STATEC and funding partners. The author(s) gratefully acknowledge the support of the Observatoire de la Comptitivit, Ministre de lEconomie, DG Comptitivit, Luxembourg, and STATEC, the National Statistical Office of Luxembourg.

†STATEC Research, 13 Rue Erasme, 1468 Luxembourg, Luxembourg, e-mail: vasja.sivec@statec.etat.lu

1 Introduction

The output gap, which measures the difference between the actual output of an economy and its potential output, is an important concept in macroeconomic policy-making. It signals economic overheating or contractions and is used to guide economic policy, including government expenditures and monetary policy (Buti et al. (1998), Taylor (2000)). However, it is widely recognized that estimating the output gap has several limitations. The output gap is an unobservable variable, and estimates of it are subject to considerable uncertainty and error (Orphanides et al. (2000), Orphanides and Norden (2002), Cayen and Van Norden (2005), Orphanides and Van Norden (2005)). Another concern is that output gap estimates can be inconsistent with related macroeconomic variables, such as inflation, and lack forecasting power for such variables (Clark and McCracken (2006), Marcellino and Musso (2011), Sarwat and Ahmed (2013) and Kamber et al. (2017)). In this paper, we aim to address these shortcomings.

The Hodrick-Prescott filter (HP filter) is a popular and simple method for estimating the output gap (Hodrick and Prescott (1997)). It is often used as a benchmark for comparison with more complex output gap estimates. However, it suffers from a key limitation, it can produce estimates that are inconsistent with output-related macro variables, such as inflation or unemployment. To address this issue, we propose a modified version of the HP filter that introduces inequality constraints on the HP-filtered cycle. Our approach corrects the output gap estimate when it is inconsistent with related macro variables and leaves it unconstrained outside of such periods. Importantly, we do not impose strong structural assumptions on the exact relation between output gap and related macro variables. For example, we only require that the output gap be non-negative in the presence of inflationary pressures. This new approach improves upon existing methods. It provides an output gap better aligned with the output's dynamics that is nevertheless consistent with related macroeconomic variables, while still maintaining the simplicity of the HP filter.

Previous literature has attempted to correct the output gap by adding structural equations that link the output cycle to related macroeconomic variables, resulting in multivariate models with a high number of parameters and strong structural restrictions (Orphanides and Norden (2002), Marcellino and Musso (2011)). In contrast, we impose inequality constraints on the Hodrick-Prescott filtered cycle. This has several advantages that are explained below.

First, it reduces the number of parameters that need to be estimated, which is particularly useful for countries with short data and unstable relationships between the output gap and macroeconomic variables. Second, since inequality constraints are imposed the output gap estimate surely aligns with the chosen macroeconomic variable.[1] Thirdly, our approach does not necessitate introducing new structural equations to ensure consistency with the selected macroeconomic variables. We admit that there could be different structural forms that fit the data, but we avoid making strong structural or causality statements while remaining consistent with theoretical assumptions. This is advantageous when the true model is uncertain. Finally, our approach maintains the simplicity and speed of estimation of the HP filter.[2]

While our approach may produce inferior estimates when the multivariate model is the true data generating mechanism, this scenario is unlikely in practice due to the difficulty in identifying the true model. In fact, the variety of output gap models suggests that there is a lack of consensus on the true data generating mechanism. Therefore, we prefer the output gap to follow the output's dynamics and only impose restrictions when necessary. We do not burden the output gap with fitting the dynamics of related macroeconomic variables; instead, the inequality restrictions only attenuate output gap's magnitude. By avoiding structural restrictions, we allow the data to speak more freely. Furthermore, our Monte Carlo simulation demonstrates that our approach performs reasonably well even when the true model is the multivariate model. Finally, a simple model, such as ours, may be advantageous to work with under parameter and data uncertainty.

Additionally, in our empirical exercise, we compiled a comprehensive dataset from various archives, spanning from Autumn 2004 to Spring 2021, for each of the EU 28 countries. We use it to compare the performance of our approach with that of the HP filter model and the European Commission's (EC) model. We find that our approach increases the correlation

[1] In multivariate models, it is possible that the solution that maximizes the total model likelihood ignores the related macroeconomic variable or instead over-fits it and neglects the output's dynamics.

[2] Some multivariate models take several minutes to estimate. Our approach takes a second or less to

between the related variable and the output gap estimate, it can reduce output gap revisions, and improve its forecasting power relative to more complex models. These results suggest that it provides a valuable alternative to traditional output gap estimation methods, particularly in cases where the true data generating mechanism is unknown or difficult to identify.

The term "output gap" was first coined by Okun in 1962, where he defined it as the difference between actual output and potential output, with potential output being the maximum output an economy can produce under full employment conditions Okun (1963). Over time, various definitions and methods for estimating output gap have emerged. Gerlach (1999) categorizes these methods into three groups: (i) statistical, (ii) structural, and (iii) mixed approaches. Our proposed approach lies between the statistical and mixed approaches.

Statistical approaches, such as time-series filters, are commonly used to extract the trend and interpret it as potential output. The popular HP filter (Hodrick and Prescott (1997)) is widely used, although alternative filters exist (e.g., Baxter and King (1999), Butterworth (1930), or Christiano and Fitzgerald (2003)). Univariate filters, like the HP filter, are preferred due to their simplicity and low parameter uncertainty.[3][4] However, relying solely on univariate filters can produce counter-intuitive output gaps by ignoring information from other variables such as unemployment or inflation. In our *baseline model*, we enhance the HP filter with inequality constraints based on related variables to address this issue and retain the simplicity of the HP filter.

Structural approaches to estimating potential output include dynamic stochastic general equilibrium models (DSGE) and structural vector autoregressive models (SVAR). DSGEs are micro-founded models of the economy that define potential output as "output in the absence of rigidities." This concept is less intuitive to policymakers. Moreover, DSGEs impose strong assumptions on the structure of the economy, are sensitive to small changes in parameters, and do not consider trend estimation. Some examples of DSGE models include Juillard et al.

estimate.

[3]Another method is to regress output on a polynomial of time, as discussed in Gerlach and Smets (1999) and Canova (2020).

[4]Hamilton (2017) proposes a third approach based on forecasting. A recent paper by Canova (2020) offers an insightful critique of this and other approaches. According to Canova (2020), the polynomial filter performs best in recovering the output gap in many situations, although it should be noted that in his simulations, the HP filter also performs reasonably well.

(2006), Hirose et al. (2007), and Vetlov et al. (2011). In the SVAR approach, permanent supply and temporary demand shocks are identified, and potential output is assumed to be a sum of permanent shocks, while the output gap is assumed to be a sum of temporary shocks. However, the SVAR approach heavily relies on identifying assumptions to derive permanent and temporary shocks, which are not verifiable, making it less appealing for policymakers. Some notable examples of the SVAR approach include Blanchard and Quah (1988), Claus et al. (2000), Garratt et al. (2006), and Mitchell et al. (2008).

In the mixed approach, a structural equation is added to the univariate model. This added equation links the output cycle to a third variable, such as the unemployment rate (Okun's law (Okun (1963),Adams and Coe (1990) and Beneš et al. (2010)). Similar, the Phillips curve (Phillips (1958), Laxton and Tetlow (1992) and Kuttner (1994)) relates the output cycle to inflation. Some authors use alternative information, Laubach and Williams (2003) considers the interest rate, Borio et al. (2016) credit growth, and Dobrescu (2006) the external balance. The mixed approach is used by the EC and other international institutions in a so-called production function setting (see the Model section). In our *production function model*, we define inequality constraints consistent with Okun's law and the Phillips curve by considering variables such as the capacity utilization index and inflation.

The mixed approach is an improvement over univariate models as it incorporates additional variables to establish coherence between the output cycle and related variables. However, this approach has limitations and can result in over-fitting or ignoring important variables if not properly specified. It tends to feature complex models that are difficult to estimate on small samples available in practice. In contrast, our model avoids these issues as constraints are super-imposed to ensure consistency with economic theory. It is simple, parsimonious and does not require parameter estimation. It reduces the likelihood of bias due to parameter breaks or model misspecification. Overall, it provides a robust alternative for estimating potential output and the output gap.

The paper is organized as follows: In Section 2, we present our model and provide an illustrative example. Section 3 discusses model extensions. The estimation algorithm is described in Section 4, while Section 5 presents a Monte Carlo verification of the model's small sample properties. In Section 6, we introduce real data, which are then used in Section 7 to compare the performance of our model with that of the HP filter and the EC. Finally, Section 8 concludes the paper.

2 Model

Our baseline model is based on the HP-filter approach, but augmented with inequality constraints. Inequality constraints are used in engineering to correct implausible state estimates (see Simon (2010)). For example, due to noisy signals a gps positioning model might predict that a car is off-road, but inequality constraints ensure that projected car position is consistent with the road space (Simon and Chia (2002)). In our model, we constrain the output gap to be consistent with inflation, unemployment, or other desired variable. This does not require additional structural equations, making it parsimonious.

Before introducing the inequality constraints, we first present the HP-filter model (eq. (1)), as put forth in Harvey (1997) and Harvey and Trimbur (2008). The HP filter decomposes the output y_t into a trend component t_t and a cyclical component c_t:

$$y_t = t_t + c_t, \qquad c_t \sim NID(0, \sigma_c^2), \qquad t = 1, ..., T \tag{1}$$

where

$$t_{t+1} = t_t + \beta_t, \tag{2a}$$

$$\beta_{t+1} = \beta_t + \zeta_t, \qquad \zeta_t \sim NID(0, \sigma_\zeta^2) \tag{2b}$$

The trend t_t follows a double-integrated random walk (eq. (2a)-(2b)). The cycle c_t and slope disturbance ζ_t are assumed to be independent and normally distributed with zero mean and non-negative variances σ^2. t_t and c_t are unobservable, but can be estimated using the Kalman smoother. The solution for t_t is a weighted average of past and future observations of y_t, which is evident from re-expressing the model in the form of a minimization problem of a loss function:

$$\min_{t_1 \dots t_T} \sum_{t=1}^{T} \Big[\frac{1}{\sigma_c^2}\big(y_t - t_t\big)^2 + \frac{1}{\sigma_\zeta^2}\big(\Delta t_{t+1} - \Delta t_t\big)^2\Big] \tag{3}$$

The Kalman filter estimates a series (t_t) that simultaneously minimizes the deviations of the trend from the observable series $(y_t - t_t)$ and the volatility of the trend $(\Delta t_{t+1} - \Delta t_t)$. The model comprises of a single observable variable (y_t) and is characterized by two parameters (σ_c^2 and σ_ζ^2), thus a unique solution is only possible if we impose a restriction. To achieve this, we set a constraint on the ratio of the two variances, $\frac{1}{\lambda} = \frac{\sigma_\zeta}{\sigma_c}$. The value of λ determines the degree of smoothness of the trend. A higher λ places more weight on the volatility part, resulting in a smoother trend with greater weight on distant observations. Conversely, a lower value of λ places more emphasis on the deviations of the trend from y_t, producing a more volatile trend that closely tracks y_t.

To simplify the introduction of inequality constraints, we express the model in a state-space form, following the notation in Durbin and Koopman (2012):

$$y_t = Z_t \times \alpha_t + \epsilon_t, \tag{4a}$$

$$\alpha_{t+1} = T \times \alpha_t + R \times \eta_t, \qquad \eta_t \sim N(0, Q) \tag{4b}$$

where y_t is a $n \times 1$ vector of dependent variables, where in our case, $n = 1$. Other vectors and matrices are defined as:

$$\alpha_t = \begin{bmatrix} t_t \\ t_{t-1} \\ c_t \end{bmatrix}, \quad Z_t = \begin{bmatrix} 1 & 0 & 1 \end{bmatrix}, \quad T = \begin{bmatrix} 2 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$R = \begin{bmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix}, \quad \eta_t = \begin{bmatrix} \zeta_t \\ c_t \end{bmatrix}, \quad Q = \begin{bmatrix} \sigma_\zeta & 0 \\ 0 & \sigma_c \end{bmatrix}$$

We next modify the model presented in (4a)-(4b) by adding inequality constraints. The augmented system will have to meet the following inequality constraint(s):

$$\begin{aligned} &D_t \times \alpha_t + \nu_t \geqslant d_t \qquad \text{if} \qquad (x_t > \bar{x}_t) \ \wedge \ (c_t < 0), \quad where \quad \nu_t \sim NID(0, \sigma_\nu) \\ &D_t \times \alpha_t + \nu_t \leqslant d_t \qquad \text{if} \qquad (x_t < \bar{x}_t) \ \wedge \ (c_t > 0) \end{aligned} \tag{5}$$

The matrix D_t is a known $s \times n$ matrix, where s represents the number of restrictions to be imposed. The vector d_t is a pre-set $s \times 1$ vector of values, and we will explain the purpose of ν_t shortly.

Each row of the matrix D_t serves to select a particular state from the vector α_t, or a combination of states, that will be subject to the imposed constraint. For instance (see the first part of constraint in (5)), to enforce non-negativity on the cycle ($c_t \geq 0$) when a related variable x_t exceeds a certain threshold $\bar{x}_t$, we can set D_t equal to the row vector $[0, 0, 1]$. This constraint will remain inactive as long as the value of the cycle is non-negative. However, if during a given period, the cycle evaluates to a negative value ($c_t < 0$), then the constraint will be violated and will become active. During such periods, if $\nu_t = \sigma_\nu = 0$, D_t selects the cycle ($D_t \times \alpha_t = c_t$) and evaluates it to d_t. If d_t is set to zero, then the cycle will be forced to close ($c_t = d_t = 0$). This inequality constraint ensures that the cycle remains non-negative whenever x_t exceeds the average threshold ($\bar{x}_t$).

The inequality constraint described thus far is referred to as a "hard" constraint because it forces a counter-intuitive cycle to a value of zero, effectively closing the output gap. However, instead of using hard constraints, we can use "soft" constraints that do not completely close the gap, but instead reduce its size. Soft constraints are implemented using a variable ν_t, which is defined as a zero-mean random noise. This variable controls the tightness of the constraints. When σ_ν is set to a positive value, the constraint becomes "soft", allowing the cycle to be evaluated outside the valid space, but still close to it. In this way, we can introduce either hard or soft constraints depending on the desired level of stringency.

Inequality-constrained models can be estimated with two methods: i) the active set method (see Simon (2010), Alouani and Blair (1993), Sircoulomb et al. (2008)) and 2) the interior point approach (see Bell et al. (2009)). We use the active set method because it is more computationally efficient. According to this method, only constraints active at the solution affect optimality conditions and, when active, they can be expressed as equality constraints. At each iteration, the algorithm solves an unconstrained optimization problem. If a previously inactive constraint becomes active, we replace the binding inequality constraint with an equality constraint:

$$D_t \times \alpha_t + v_t = d_t \tag{6}$$

Optimality properties for equality constrained estimator are discussed in Simon and Simon (2003) and Gupta and Hauser (2007).

Simon (2010) presents four methods for estimating state-space models with equality constraints. The perfect measurements approach entails augmenting the model presented in (4a)-(4b) with the inclusion of (6), during periods characterized by binding constraints. The estimate projection approach involves estimating the unconstrained state variables using a standard Kalman filter, and then projecting them onto the constraint surface. The gain projection approach involves modifying the Kalman gain to take into account the constraints. Finally, the probability density truncation approach involves truncating the probability density function of the state variables to satisfy the constraints. We employ the perfect measurements approach due to the mathematical equivalence of state estimates among the four approaches, as well as its computational efficiency and ease of implementation. Specifically, we expand the model in the following manner:

$$\begin{bmatrix} y_t \\ d_t \end{bmatrix} = \begin{bmatrix} Z \\ D_t \end{bmatrix} \times \alpha_t + \begin{bmatrix} 0 \\ \nu_t \end{bmatrix}, \qquad \text{where} \qquad \nu_t \sim NID(0, \sigma_\nu^2) \tag{7}$$

When a constraint is binding, the cycle c_t is selected and assigned the value d_t, using $D \times \alpha_T$. In our example, we have set d_t to zero. If the variance σ_ν^2 is positive, the estimated value of the cycle decreases, leading to a reduction in the output gap. When the variance is zero, the cycle evaluates to zero, indicating that the output gap closes. The equation (7) represents our baseline model.

The constraint imposed in equation (7) is redundant when the inequality constraints are not binding. During such periods, we can use either the model in (4a-4b), or we can treat d_t as a missing value to prevent it from affecting the state vector α_t. This is achieved by either: i) setting $D_t = 0$ (Mariano and Murasawa (2010)), ii) setting $\sigma_\nu^2 = \infty$ (Giannone et al. (2008)), or iii) reducing the model dimensions as explained in Durbin and Koopman (2012).

To contrast our proposed model with the HP filter, we can once again represent the model as a problem of minimizing a loss function:

$$\sum_{t=1}^{T} \Big[\frac{1}{\sigma_c^2}\big(y_t - t_t\big)^2 + \frac{1}{\sigma_\zeta^2}\big(\Delta t_{t+1} - \Delta_t\big)^2 + \frac{1}{\sigma_\nu^2}(\underbrace{D_t\alpha_t - d_t}_{=c_t - d_t})^2\Big] \tag{8}$$

When constraints are binding the term $(D_t\alpha_t - d_t)$ enters the loss function. When d_t is set to zero, the term $(D_t\alpha_t - d_t)$ is minimized when c_t is set to zero as well and the gap is reduced. The extent of this correction is determined by the value of σ_ν^2, which acts as an importance weight for the term $(D_t\alpha_t - d_t)$. If σ_ν^2 is set to zero, the weight on $(D_t\alpha_t - d_t)$ becomes infinite, resulting in a gap evaluation of zero. On the other hand, when σ_ν^2 is positive, the gap is reduced to some extent.

When estimating σ_ν with maximum likelihood, it can be challenging to obtain plausible estimate due to its tendency to be either very close to zero or unreasonably large for certain countries.A similar issue arises when estimating the smoothing constant $\lambda = \frac{\sigma_\zeta}{\sigma_c}$ of the HP filter using maximum likelihood, as demonstrated in studies such as Hamilton (2017) and Harvey and Jaeger (1993). Harvey and Jaeger (1993) notes that the smoothing constant is to be fixed in order to obtain sensible results. Therefore, we prefer to calibrate this parameter to a reasonable value based on economic intuition. Calibration offers the advantage of simplifying interpretation and reducing the number of estimable parameters. To achieve this, we introduce a parameter κ that links σ_ν to the cycle's variance:

$$\kappa = \frac{\sigma_\nu}{\sigma_c} \tag{9}$$

The parameter κ is defined as the ratio between the variance of the constraint and the variance of the cycle, and it determines the level of stringency of the constraints. Higher values of κ indicate weaker confidence in the restrictions imposed. When κ is set to 1, the model produces a solution that is close to the unconstrained problem, the HP filter solution. On the other hand, when κ is set to 0, the model produces the solution with hard constraints.

2.1 Estimation

The following steps describe the estimation algorithm:

- Set the values for the parameters $\{\lambda, \kappa\}$ and initialize vector d_t to missing values for all time points t: $d_t = ., \forall t$.

- Estimate the state vector α_t, which represents the trend and cycle component of the time series, using a Kalman smoothing algorithm. The estimation is based on the observed data y_t and the current state vector d_t. Since d_t is initialized as missing values on the first iteration, the smoothed state will be equivalent to the trend component estimated by the Hodrick-Prescott (HP) filter.

- Verify whether the estimated state vector α_t violates any constraints imposed on the model, which are defined by Equation (5). The algorithm progresses forward in time from $t = 1$ to T. When the first violation is detected, the corresponding d_t value is set to zero and the algorithm immediately returns to Step 2 to re-estimate the state vector[5]. The algorithm stops if no new violations are detected.

So far, we introduced the baseline model and discussed its estimation. The next section will present an illustrative example and two popular model extensions that are often used in practice.

2.2 Illustrative example

Figure (3) shows the estimated output gap for EU27 (1995-2021)[6] using the Autumn 2019 vintage, with a signal-to-noise ratio set to $\lambda = 100$. To save space, we will only present figures for the EU27 area. Inequality constraints were derived from the cycle of the GDP deflator (HP filtered and shown in Figure 4). The black dotted line in figure (3) is the EC's output gap estimate, while the black solid line with square markers represents the output gap estimated using our model with hard constraints ($\kappa = 0$). The periods in which the inequality constraints were binding are marked by a blue background. To avoid instabilities due to rounding errors, we set κ to $1e - 8$.

To illustrate the effect of relaxing the tightness of constraints, we can observe the estimated output gaps for different values of κ. As κ increases, the constraints become less binding and the estimates move closer to the unconstrained output gap (depicted with light-grey line with triangle markers). In the figure, the output gaps with soft inequality constraints are shown in grey and tend to evaluate between the output gaps with hard constraints and the unconstrained output gap. However, it is important to note that hard constraints can affect the trajectory of the output gap in periods adjacent to the period in which the constraint binds, which may lead to deviations from this pattern.

Figure (3) shows that inequality restrictions attenuate the output gap during periods

[5]This is done because imposing a constraint on the state vector at time t may change the trajectory of the state vector in neighboring time points, which often makes constraints further in time redundant.

[6]EU27 includes: Belgium, Bulgaria, Czech Republic, Denmark, Germany, Estonia, Ireland, Greece, Spain, France, Croatia, Italy, Cyprus, Latvia, Lithuania, Luxembourg, Hungary, Malta, Netherlands, Austria, Poland, Portugal, Romania, Slovenia, Slovakia, Finland and Sweden.

when it is inconsistent with the signaling variable. For example, the unrestricted output gap for the year 2013 is negative (indicated by the grey line with the triangle mark in Figure (3)). However, in that year, the cycle of the GDP deflator was still positive, which is inconsistent with a negative output gap. The imposition of inequality constraints reduces the size of the output gap (soft constraints) or eliminates it altogether (hard constraints) in that year.

It is worth noting that hard constraints can also impact the trajectory of the output gap in periods adjacent to when the constraint first becomes binding. As observed, the output gap with hard constraints remains closed not only in the years 2007 and 2013 but also in the subsequent years 2014 and 2015. Figure (5) demonstrates how these restrictions determine potential output. Specifically, they imply lower potential immediately before and during the crisis, consistent with overheating before the crissis and stronger scarring effects during the crisis.

However, we stress that this is an illustrative example and is not intended to establish a general truth. Its purpose is to illustrate the mechanics of our approach. A better-informed estimator would dedicate significantly more attention to signaling variable selection and the selection of related parameters.

2.3 Model extensions

2.3.1 Model with autoregressive cycle

The model presented can be extended with an autoregressive cycle (see for example Watson (1986), Clark (1987), or Morley et al. (2003)). Additional details can be found in section Model extension (1) in the appendix, where we also provide an example estimate for the EU27 (see Figure (6)).

Upon analyzing the impact of incorporating an autoregressive cycle (limited to selected countries), we observed that it did not significantly alter the results presented in the empirical section of this study. This finding aligns with previous research, such as Harvey and Jaeger (1993), which highlights the difficulty in distinguishing the HP-filter cycle from the cycle estimated in a model with an autoregressive component. As a result, we focus our further analysis on the models without the autoregressive cycle.

2.3.2 Model consistent with the production function approach to output gap estimation

The production function approach to output gap estimation involves estimating potential output from the inputs to production rather than modeling the output series. Specifically, it includes applying various models to the series that enter the Cobb-Douglas production function, such as labor, capital and total factor productivity (see Cobb and Douglas (1928)). The IMF, OECD, and the EC are among the organizations that utilize this approach for output gap estimation.

We will compare our *production function model* with the EC's. This will allow us to assess the effectiveness of our approach within a production function setting. The production function model is further described in section Model extension (2) in the appendix. Here, we mention that we apply the *baseline model* to estimate the trend of total factor productivity (TFP) and the trend of the unemployment rate. The trend of TFP is estimated using inequality restrictions derived from the capacity utilization index (CUBS), while for the unemployment rate, we use either the real GDP deflator or wages to derive the inequality restrictions. The remaining components of our model adhere to the EC's specification (see Havik et al. (2014)), which is detailed in the appendix as well.

2.4 Monte Carlo verification of small sample properties

In order to evaluate the performance of our method, we conduct a Monte Carlo study. First, we estimate a multivariate trend-cycle model on the EU15 data and use it as the true model to simulate new data.[7] [8] Next, we estimate three different models using simulated data. The first model is a multivariate model with the same structure as the true model. This model is expected to perform the best, as it matches the model that generated the data.

[7]EU15 countries are: Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, Netherlands, Portugal, Spain, Sweden and United Kingdom.

[8]European Commission publishes country-specific CUBS indicators. We constructed EU15 CUBS indicator by weighting country-specific CUBS indicators by their rGDP shares. For many EU27 countries CUBS is not available from 1996 onward.

The second model is the univariate HP filter model, which neglects the multivariate nature of the true model. The last model is our baseline model, which uses inequality constraints to incorporate information from multiple variables, but does not have the exact same structure as the true model. We first demonstrate that while the inequality constrained model may not correspond exactly to the true data generating process, it still performs well and improves over the HP filter model.

We estimate the following multivariate trend-cycle model for the EU15 data of the 2021 Spring vintage:

$$
\begin{aligned}
y_t &= t_t + c_t && (10)\\
t_t &= 2 \times t_{t-1} + t_{t-2} + \epsilon_t && (11)\\
CUBS_t &= \beta \times c_t + u_t && (12)
\end{aligned}
$$

The model decomposes output y_t into trend t_t and cycle c_t. The trend component is modeled as a second-order random walk process in (11). The cycle component is constructed to capture the business cycle by loading it on the capacity utilization rate ($CUBS_t$) in (12), which is a widely used indicator of the state of the economy. The root mean squared error (RMSE) is used to compare the performance of the three models, which is calculated as the difference between the simulated cycle and the re-estimated cycle. We conduct 1000 Monte Carlo replications and fix the ratio between the two variances (σ_ϵ, σ_c) to 10 or 100. The sample size is varied between $T = [10, 25, 50]$. We also experiment with different values of tightness of restrictions for the inequality constrained model.

Figure (1) displays the results of this experiment, which includes twelve figures. The four columns represent varying levels of restriction tightness ($\kappa = 1e - 8, 0.05, 0.5, 1$), and the three rows correspond to different sample sizes ($T = 10, 25, 50$). The vertical axis of each figure represents the root mean square errors (RMSEs). The horizontal axis represents the value of the smoothness parameter ($\lambda = 10, 100$). The RMSE for the (true) multivariate model is represented by *red square*, while the HP filter and inequality constrained filters are represented by *blue circle* and *yellow triangle*, respectively.

In this experiment, the multivariate model (red square) achieved the lowest RMSE, as it accurately represents the true data generating process. The HP filter (blue circle), which disregards the information contained in the CUBS variable, achieved the highest RMSE in most cases. The RMSE of the inequality constrained filter (yellow triangle) tends to lie between these two estimates, indicating that it incorporates some information from the CUBS variable, though not as much as the multivariate model.

We first comment on the model with soft constraints ($\kappa = 0.05, 0.5, 1$), reported in columns (2)-(4) and visually distinguished from the model with hard constraints by a vertical dashed line. It is worth noting that as the restrictions are relaxed (indicated by an increase in κ from 0.05 to 1), the RMSE of the inequality constrained filter approaches that of the HP filter. Visually, this can be observed as the yellow triangle moves closer to the blue dot. When the tightness of restrictions is set to 1, the solution of the inequality constrained filter resembles that of the HP filter. Alternatively, when the tightness of restrictions is set to 0.05, the inequality constrained filter closes the RMSE-gap between the true and the HP filter solution by about half. This demonstrates how tightness of constraints (κ) governs the extent to which the associated variable influences the solution, leading to improved accuracy in estimation.

Furthermore, the sample size (T) has a limited impact on the accuracy of our model and other models. While the root mean square error (RMSE) decreases slightly with an increase in sample size, the improvement is not substantial. However, all of the investigated samples are small by conventional standards, as are the yearly data used in practice to estimate the output gap. The smoothness of the trend (λ) seems to have a more pronounced impact, affecting all models to a similar degree. Increasing the smoothness parameter leads to higher RMSE values across all models. This relationship arises because a smoother trend results in a more volatile cycle and a larger discrepancy between the true and estimated cycles.

We now refer to the model with hard constraints, which is reported in column (1). It is evident that the application of hard constraints ($\kappa = 1^{-8}$) can impede the performance of the inequality constrained filter. This effect is particularly noticeable in the cases with $T = 50$

and $\lambda = 100$, as well as $T = 25$ and $\lambda = 100$, where the inequality constrained filter exhibits the poorest performance. The cause is evident from Equation (12) which contains a noise term (u_t), indicating that the signaling variable ($CUBS_t$) may not always deliver the correct signal for the output gap. Based on this finding, we discourage the use of hard constraints in practical applications.

In the next experiment we demonstrate that the inequality constrained filter performs better than the multivariate filter in a situation where an irrelevant variable is mistakenly included in the model. To illustrate this, we begin by simulating output based on the model presented in equations (11)-(12). Subsequently, we substitute the simulated draws of $CUBS_t$ with simulated draws of the real effective exchange rate, which is not part of the actual data generating process. In particular, we employ an autoregressive (AR) process to generate draws for the real effective exchange rate of the EU15. This scenario serves as an analogy for a practitioner mistakenly employing a loosely related or a highly noisy variable to estimate the output gap.

Figure (2) illustrates the results of this experiment. When constraints are soft, the multivariate model exhibits the poorest performance, while the HP filter demonstrates the best performance, aligning with expectations. The RMSE of the inequality constrained model typically falls between these two extremes. Notably, as we progressively relax the restrictions (κ increases), the RMSE of the multivariate model gradually converges towards that of the HP filter model which performs best. Additionally, we observe that in this experiment, the sample size plays a more significant role in determining the root mean square error (RMSE). This holds especially for the multivariate model. With an adequate sample size, the multivariate model effectively learns the true value of β, which is near zero in this scenario, resulting in improved performance. This is evident in the case of $T = 100$ and $\kappa = 0.25$. In cases, where κ is set to 0.5 or higher the inequality constrained filter retains its superior performance. Therefore, we recommend setting κ to values higher or equal to 0.25 when practitioners are uncertain about whether a particular variable drives the output gap.

These two experiments demonstrated that the inequality constrained filter performs relatively well in small samples, regardless of whether the signaling variable accurately captures the underlying cycle or when the assumed model is flawed. These findings highlight the robustness of the inequality constrained filter in practical scenarios.

3 Empirical application

We apply our estimation procedure to real-world data by collecting two vintage datasets for EU countries. We use them to estimate two models, namely the *baseline model* and the *production function model*. We then compare these models to alternative estimation procedures, such as the HP filter and the EC's model. Additionally, we evaluate these models across various dimensions that are essential for policymakers. Finally, we conclude this chapter with a section on the best-performing models.

3.1 Data

We estimated our models using two real-time datasets collected from the archives of the EC's website for the Output Gap Working Group.[9]

The first dataset contains the series entering our *baseline model*. It starts with the vintage Autumn 2004 and concludes with the vintage Spring 2021. As the European Commission (EC) generates biannual forecasts, in Autumn and in Spring, this dataset encompasses a total of 34 vintages for 28 European countries.[10][11] These vintages include forecasts up to T+2,[12] which makes our estimates subject to endpoint bias.[13] Based on these data, the inequality restrictions were derived from the consumer price index, wage inflation, unemployment rate, or the capacity utilization index.

In the early vintages, not all the series necessary for the *production function model* are available and the forecasts are limited to T+2.[14] The first vintage that includes all the series

[9]The data is available on: "circabc.europe.eu" →Browse categories →Economic and Financial Affairs →Output Gaps →Library. Folder "99. Archives".

[10]The data includes the UK.

[11]Occasionally the EC also produces an interim Winter or Summer forecast.

[12]We follow the EC's convention of defining T. In Autumn, T is the current calendar year and in Spring T is the past calendar year.

[13]Endpoint bias arises from the tendency of the HP filter (and other smoothing filters) to close the cycle at the end of the sample.

[14]With the working age population (15-75 years old) most often being the missing series.

entering our production function model is Spring 2014. Therefore, we estimate the production function model for 15 vintages from Spring 2014 until Spring 2021. These vintages include forecasts for up to T+5, reducing the impact of endpoint bias on our estimates.[15] For this model and data, the inequality restrictions for the total factor productivity were derived using the capacity utilization index, while those for the unemployment rate cycle were derived from the real GDP deflator or wage inflation (see section Model extension (2) in the appendix).

The data for the old member states is primarily accessible from the 1960s onward, while for the new member states, it generally begins in the 90s. Due to the relatively limited duration of the data series for the new member states, the comparison statistics (which will be introduced later) will be computed from 2000 onwards to ensure a consistent analysis. A more comprehensive overview of the datasets can be found in the online appendix (see section Data).

3.2 Results

In subsequent sections, we will compare the *baseline model* and the *production function model* with the alternative approaches. We conduct a comparison for both models as they each come with their advantages and disadvantages. A more precise comparison can be made for the baseline model as it features a larger number of vintages (34). However, the baseline model is susceptible to endpoint bias because it was estimated using data that include forecasts up to the time period T+2. On the other hand, the production function model is estimated using data that include forecasts up to T+5, and it is expected to exhibit a smaller endpoint bias. Unfortunately, its estimation is based on a small number of vintages (15), which introduces uncertainty in its comparison.

We demonstrate the effectiveness of our approach by performing three types of comparisons[16]. First, we use a correlation analysis to show that the constraints incorporate information from a signaling variable into the output gap. Second, we compare different

[15] It should be noted that endpoint bias can be reduced only if forecasts are accurate.

[16] Early version of the paper also included an analysis of output gap sign reversals. We omit it for space considerations and because most sign reversals were of little practical importance (most occur due to a narrow crossing of the zero-gap value (e.g.) from -0.05% to +0.05%.

output gap estimates based on the size of their revisions over time. Finally, we evaluate the forecasting power of different gap models. Our results show that the inequality constraints can produce smaller revisions compared to the EC's or HP filter estimates, depending on the restricting variable and model settings. Additionally, our method delivers output gaps with good forecasting performance.

We will consider the model estimated on the latest vintage (Spring 2021) as a proxy for the "true" output gap due to its utilization of a more extensive data set. We will then compare the estimates from earlier vintages to the Spring 2021 estimate.

In accordance with economic theory and from a policy perspective, several variables could indicate counter-intuitive gaps. We derive inequality constraints from four variables: *inflation*[17], *wage inflation*, *unemployment rate*, or *capacity utilization index*. Inequality restrictions can only be derived from stationary variables. To make them stationary, we consider three different approaches. Firstly, we detrend them with an *(a) HP filter*. However, the HP filter is sensitive to the arrival of new data and this could lead to instabilities. As such, we also consider demeaning variables with their own *(b) 5-year moving average (5y MA)* and by *(c) differencing them.*

The output gap of country i for year t and data vintage v is denoted by $OG_{i,t}^{v}$. This output gap is estimated using a signaling variable x, with inequality constraints of tightness κ and a signal-to-noise ratio of λ. Thus, we can represent the output gap as:

$$OG_{i,t,v}^{x,\kappa,\lambda} = \frac{Y_{i,t,v} - \overline{Y}_{i,t,v}^{x,\kappa,\lambda}}{\overline{Y}_{i,t,v}^{x,\kappa,\lambda}} \tag{13}$$

The notation introduced defines the output gap of country i in year t, data vintage v, using the signaling variable x, constraint tightness κ, and signal-to-noise ratio λ. Here, Y denotes the actual output and $\bar{Y}$ denotes the potential output. The vintage v is represented by the year and season, where season can be either Spring (1) or Autumn (2). For instance, the Spring 2014 vintage is denoted by $v =$ '$2014, 1'$. For clarity purposes we suppress superscripts

[17]Expressed from real GDP deflator.

'x, κ, λ' in the remained of the text.

The estimates are analyzed in two ways: in-sample and real-time. In the *in-sample analysis*, we compare estimates based on observable data from each vintage with the estimates obtained using the final vintage data-set (Spring 2021) for the same period. For example, if we use the Spring 2014 vintage ($v =$ '$2014, 1'$), data is considered in-sample when $t \leq y = 2014$ because only data up to the year $t = 2013$ is observable.

In the *real-time analysis*, the output gap is calculated for the fiscal year $T+1$ and according to the EC's terminology. For instance, for the vintage $v =$ '$2014, 1'$, data is considered $T + 1$ when $y = 2014$. The output gap estimate is then based on the 1-year ahead forecast. However, for the vintage $v = 2013, 2'$, data is also considered $T + 1$ in $\tilde{y} = 2014$ because governments prepare budget plans for the fiscal year 2014 in Autumn 2013. In Autumn, the output gap estimate is based on the 2-year ahead forecasts. In this way, the real-time comparison reflects the use of the output gap in practice[18].

We used the EC's (European Commission) forecasts throughout the analysis. This choice was made to prevent differences in forecasts from affecting the comparison, and to make the analysis more realistic by reducing endpoint bias.

In the subsequent sections of this paper, we will compare three output gap estimates: our Inequality Constrained Output Gap (IC GAP), the HP Filtered Output Gap (HP GAP), and the Output Gap estimated by the Economic Commission (EC GAP). We will assess their performance based on five criteria: the strength of the in-sample correlation with the signaling variable used for estimating the IC GAP (CORR), their correlation in real-time (CORR T+1), magnitude of in-sample revisions (REV), magnitude of real-time revisions (REV T+1), and the effectiveness of the output gap in forecasting the signaling variable (FCSE). Similar comparisons have been conducted in previous studies, such as Orphanides and Norden (2002), Bundesbank (2014), and Mc Morrow et al. (2015).

We will examine these criteria across various signaling variables (e.g., real GDP deflator, wage inflation, unemployment rate, and capacity utilization rate), employing different techniques to induce stationarity (e.g., HP filter cycle, demeaning with a 5-year moving average,

[18]The results do not change if we instead define $T + 1$ as the first forecasted data point.

and differencing), employing distinct smoothing constants ($\lambda = 10, 100$), and varying the level of constraint stringency ($\kappa = 0.05, 0.1, 0.25, 0.5, 1$). For an overview of the estimated models, please refer to section Models in the appendix.

Given the substantial volume of results, including them in the paper would require a significant amount of space. Therefore, we have included them in the online appendix ([LINK]). However, Table (1) provides a representative excerpt from these results.

The left panel of Table (1) presents the outcomes of the baseline model, which utilized the unemployment rate, made stationary by differencing, to derive the constraints. The smoothing constant was set to 10. In the right panel, the results of the production model are displayed. This model employed the GDP deflator, demeaned using a 5-year moving average, to derive the constraints. The smoothing constant was also set to 10.

These two models were chosen due to their good performance across all five comparison criteria, demonstrating the advantages of our proposed approach (refer to the "Best Performing Models" section). Throughout the rest of this paper, we will frequently refer to Table (1).

3.2.1 Correlation analysis

We use inequality constraints to restrict the output gap when it contradicts the signaling variable, with the aim of increasing the correlation between the estimated output gap and the signaling variable. To assess the effectiveness of these constraints, we compare the correlation coefficients between the output gaps and the signaling variables.

In-sample results: To assess the general in-sample performance of our approach, we first calculate the average correlation coefficient across all country-vintage pairs.

$$\overline{r} = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{V}\sum_{v=1}^{V} r_{i,v} = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{V}\sum_{v=1}^{V}\frac{\sum_{t=t_1}^{T}(OG_{i,t,v} - \overline{OG}_{i,v})(x_{i,t,v} - \overline{x}_{i,v})}{\sqrt{\sum_{t=t_1}^{T}(OG_{i,t,v} - \overline{OG}_{i,v})^2(x_{i,t,v} - \overline{x}_{i,v})^2}} \tag{14}$$

Here, $r_{i,v}$ represents the correlation coefficient of country i and vintage v estimated on the interval $[t_1, T]$. $r_{i,v}$ measures the extent of correlation between the country-specific output gap

($OG_{i,t,v}$) and the corresponding signaling variable ($x_{i,t,v}$). Individual correlation coefficients were calculated for the period from $t_1 = 2000$ to $T = y - 1$ (or $t = 2000, ..., y - 1$). The value of $\bar{r}$ is the average correlation coefficient across all country-vintage pairs. N and V denote the total number of countries and vintages, respectively.

Rows marked as *CORR* in Table (1) depict the average correlation coefficients between the output gap ($OG_{i,t,v}$) and the selected signaling variable ($x_{i,t,v}$). The left panel presents the correlation coefficients for the baseline models, while the right panel displays those for the production function models. Within each panel, we present the average correlation coefficient for the inequality constrained output gap (IC Gap), followed by the HP gap and the EC's gap. The analysis also explores various levels of restriction tightness, ranging from strong restrictions ($\kappa = 0.05$) to increasingly looser restrictions ($\kappa = 0.10, 0.25, 0.50, 1.00$)[19]. These are listed in the first column labeled κ. It is important to note that constraint stringency only affects the inequality constrained output gap.

We first examine the correlation coefficients for the baseline model. When restrictions are strong ($\kappa = 0.05$), the correlation coefficient between the growth of the unemployment rate and the inequality constrained gap is 0.67.[20] However, when the tightness of constraints decreases, the correlation coefficient decreases to 0.53 when $\kappa = 1$, which is close to the unrestricted model's correlation coefficient of 0.48 (see the *HP gap* column). This pattern is consistent for other variables and other signal-to-noise ratios. In essence, the inequality constraints work to reduce the size of the output gap $OG_{i,t,v}$ when it does not align with the signaling variable $x_{i,t,v}$, thereby increasing the correlation between the two (as evidenced by equation (14)). This also demonstrates how constraint stringency κ can increase or decrease the correlation.

The right panel displays the results for the production function model. Here as well, the inequality constraints increase the correlation between the signal and the output gap, although the correlation is lower in magnitude. This was expected since this model also includes variables whose trend does not depend on inequality constraints.

[19] As mentioned in the Monte Carlo section of the paper, hard restrictions (with $\kappa = 0$) are often overly restrictive for practical purposes.

[20] Unemployment rate is negatively correlated with the output gap. For ease of comparison, we have

Table (1) demonstrates the results for two selected variables. However, to asses general properties of the method the experiment was repeated over a range of variables, stationarity inducing transformations and smoothing constants. Table (2) reports summary results. It displays the share of instances where the inequality constrained output gap outperformed the HP filter or the EC's output gap for each of the comparison criteria and demeaning procedure. Note that these shares are calculated across all countries, vintages, smoothing constants λ and constraint stringency settings κ.

The row labeled CORR in Table (1) indicates the percentage of instances where the inequality constrained output gap displayed a stronger correlation with the signaling variable compared to the HP filter or the EC's output gap. We observe that the inequality constrained filter consistently generates output gaps that exhibit higher correlation than the HP filter, regardless of the specific variable considered (e.g., rGDP deflator, wage inflation, unemployment rate, capacity utilization), the demeaning procedure employed (cycle, demeaned, differenced), or the model type (baseline, production function), even if constraints are weak, such as when $\kappa = 1$. The reported percentage of output gaps better correlated with the signaling variable than the HP filter gap is always 100

We next compare the inequality constrained output gap to the EC's output gap. The reported percentages are now lower, indicating that the EC's output gap can be better correlated with the signaling variable than the inequality constrained output gap. For example, when data is differenced the inequality constrained output gap is better correlated with the signaling variable only in 33% of cases, for the baseline model. However, note that the reported 33% includes models with low stringency of restrictions (i.e. $\kappa = 0.5, 1$). Typically, the inequality constrained output gap becomes better correlated with related variables when we impose stronger restrictions ($\kappa = 0.05, 0.1$). In fact, for our most restricted setting, when $\kappa = 0.05$, the inequality restricted gap is better correlated with signaling variables than the EC's output gap for all variables, except when capacity utilization is used to derive the restrictions. This holds despite the simplicity of our model, which does not require the estimation of additional equations.

reversed its sign in the table.

Based on the results presented, the inequality constrained output gap is always better correlated with the signaling variable when compared to the HP gap. It tends to be better correlated compared to the EC's output gap when restrictions are sufficiently tight ($\kappa = 0.05, 0.1, 0.25$).

Finally, we have not been selective in presenting these averages with respect to signaling variables or parameter settings. We have simply reported averages across all combinations. However, being selective could further improve the performance of the approach.

Real Time results: To assess the performance of our approach as it would be applied in practice, or in real-time, we conducted a similar analyses as before but we calculated the correlation coefficient between the inequality-constrained output gap and the restricting variable for the upcoming fiscal year:

$$\overline{CORR} = \frac{1}{N}\sum_{i=1}^{N} r_i = \frac{1}{N}\sum_{i=1}^{N} \frac{\sum_{v=1}^{V}(OG_{i,t=y,v} - \overline{OG}_i^v)(x_{i,t=y(v),V} - \overline{x}_i^V)}{\sqrt{\sum_{v=1}^{V}(OG_{i,t=y,v} - \overline{OG}_i^v)^2 (x_{i,t=y(v),V} - \overline{x}_i^V)^2}} \tag{15}$$

In this equation, unlike eq. (14), the correlation is calculated for periods with $t = y$ (before it was $t = 1...y-1$), where y derives from $v = y - s$. For example, for vintage $v = 2014 - 1$', it refers to year $y = 2014$ and for vintage $v = 2014 - 2$' it refers to year $y = 2015$. A significant drawback of the real-time comparison is the limited number of available real-time observations. We have access to only 34 real-time observations for the baseline model and 15 for the production function model, which corresponds to the number of available vintages. This introduces uncertainty into these estimates. In addition, both the output and signaling variables are forecasted, which can introduce a bias. Finally, the average gap ($\overline{OG}_i^v$) is calculated using the real-time gaps at $t = y$, while the average value of the signaling variable $\overline{x}_i^V$ is based on the final vintage of data.

Table (1) present real-time correlation coefficients for our example models (CORR T+1). We first note that, similar as for the in-sample correlation, the correlation between the signaling variable and the output gap is higher when we impose more stringent restrictions.

This holds for the baseline model as well as for the production function model. We also mention that in rare cases, a less stringent restrictions can lead to a higher correlation coefficient in $T+1$,[21] due to a small number of observations or due to a biased forecast.

We next note that, in our example data in Table (1), the baseline model output gap tends to be better correlated with the signaling variable than the HP filter or the EC's model. This tends to hold for most variables and especially when we use more stringent constraints. Table (2) shows that, for the baseline model, the inequality constrained gap is better correlated than the HP filter gap in T+1 in 37%/100%/75%, and in 50%/83%/85% compared to the EC's gap, when the signal is rendered stationary by using the HP filtering/demeaning with five year moving average/differencing.

These shares are lower for the production function model. Exception being the case where the signal is demeaned by differencing, where it is better correlated than the HP or the EC output gap, regardless of the strength of restrictions (100%).

This section showed that the imposition of inequality constraints increased the correlation between the output gap and desired variable in-sample, and tended to increase it in real-time, especially in the case of the baseline model.

3.2.2 Revisions

The value of output gap is subject to uncertainty due to factors such as arrival of new data, revisions of old data, forecast uncertainty, and model uncertainty.[22] Since the final estimate of the output gap can differ from the real-time estimate used for policymaking, this section compares the magnitude of such revisions.

In-sample analysis: We use the normalized root mean squared error (nRMSE) to express the average size of revisions:

$$\overline{nRMSE} = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{V}\sum_{v=1}^{V}\frac{RMSE_{i,v}}{\sigma_{OG_{i,t_1\ldots T,V}}} = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{V}\sum_{v=1}^{V}\frac{\sqrt{\frac{1}{T-t_1}\sum_{t=t_1}^{T}\left(OG_{i,t,v}-OG_{i,t,V}\right)^2}}{\sigma_{OG_{i,t_1\ldots T,V}}} \quad (16)$$

[21]See, for example, online appendix, Table 6, variable LUR with smoothness constant set to 100.

[22]See for example Orphanides and Norden (2002), Orphanides and Van Norden (2005), Marcellino and

Note that we normalize the revisions by the standard deviation of the final vintage output gap ($\sigma_{OG_{i,t_1,\ldots,T,V}}$). Various models have substantial variation in the size of the cyclical component they produce due to differences in imposed parameters.[23] By scaling the revisions, we account for these differences. For example, for the inequality constrained output gap, binding constraints can reduce the output gap's amplitude. Without this normalization, the output gap with hard constraints would tend to perform better in terms of RMSE solely due to its smaller amplitude. The drawback is that the standard deviation of the final vintage output gap is itself uncertain and a small standard deviation tends to inflate the normalized RMSE.

In row REV of Table (1), we present the demonstrative nRMSE coefficients. Among the considered models, the HP filtered output gap exhibits the smallest revisions. This observation holds not only in the provided example but also in general. The HP filter model is characterized by its parsimonious nature, as it does not require parameter estimation. However, this simplicity occasionally results in inconsistencies with related macroeconomic variables. Our proposed approach addresses this limitation by introducing marginal complexity. Consequently, the revisions it produces are typically only slightly higher than those of the HP filter.

Note also that an increase in the strength of the restrictions tends to correspond to larger revisions. This phenomenon can be attributed to two factors. Firstly, as the sample changes, certain restrictions may become either non-binding or binding over time, and stronger binding restrictions result in more substantial corrections to the output gap. Secondly, we normalize the revisions by dividing them by the standard deviation of the final vintage output gap. Consequently, stronger restrictions lead to a reduction in the standard deviation that is used as a denominator in the calculation of nRMSE.

Row REV in Table (2) reveals that the inequality-constrained output gap consistently exhibited lower stability compared to the HP filter output gap, irrespective of the chosen demeaning procedure, smoothing constant, or strength of restrictions (indicated by the values

Musso (2011), Bundesbank (2014) and Mc Morrow et al. (2015).

[23]Such as the smoothing constant λ or calibrated shock variances in the case of the EC's model.

with 0% in row REV). Conversely, in approximately 70-75% of cases for the baseline model and 40-50% of cases for the production function model, our model demonstrated greater stability than the EC's output gap. Our model combines the stability inherent in a simple model like the HP filter with the incorporation of information from additional variables, addressing a limitation of the HP filter model. Furthermore, our model yields smaller revisions compared to the EC's model, as it remains relatively uncomplicated while maintaining consistency with these additional variables.

Real-time analysis: To evaluate the performance of our model in real-time, we estimate the nRMSE as follows:

$$\overline{nRMSE} = \frac{1}{N}\sum_{i=1}^{N}\frac{RMSE_i}{\sigma_{OG_{i,V}}} = \frac{1}{N}\sum_{i=1}^{N}\frac{\sqrt{\frac{1}{V}\sum_{v=1}^{V}\left(OG_{i,t=y,v} - OG_{i,t=y(v),V}\right)^2}}{\sigma_{OG_{i,V}}} \tag{17}$$

we express the average difference between the real-time estimate of the output gap ($OG_{i,t=y(v),v}$) and the estimate for the same year based on the final vintage ($OG_{i,t=y(v),V}$) with the root mean squared error ($RMSE_i$). We normalize this difference by dividing it by the standard deviation of the output gap based on the final vintage ($\sigma_{OG_{i,V}}$). However, an important limitation of this normalization process is the high level of uncertainty associated with $\sigma_{OG_{i,V}}$, as it is estimated on a sample of 17 observations for the baseline model and only 7 observations for the production function model.[24] Furthermore, due to its smaller value, $\sigma_{OG_{i,V}}$ has a tendency to inflate the normalized root mean squared error (nRMSE) of output gap estimates that incorporate strong inequality restrictions. Therefore, we advise exercising caution when interpreting these results.

In row REV T+1 of Table (1), we display demonstrative real-time nRMSE coefficients. We note that, for the baseline model with the unemployment rate as a signal, all three models perform similar, with coefficient values about 0.8. Inequality constrained gap achieves a

[24] The number of observations corresponds approximately to the number of available vintages divided by two (since there are two vintages per year).

slightly lower nRMSE of 0.79 for the models with loose restrictions ($\kappa = 0.25, 0.5, 1$). For the production function model, the coefficients are more diverse. The HP filter model achieves the lowest nRMSE of 0.75. For loose restrictions, the inequality constrained output gap, achieves a nRMSE which lies between the lowest nRMSE of the HP filter gap and the EC's output gap (min. 0.76). For strong restrictions, the size revisions, measured with nRMSE, becomes higher (max. of 0.83) than that of the EC (0.75). This is because more stringent restrictions reduce the denominator of nRMSE ($\sigma_{OG_{i,V}}$), which tends to inflate the normalized revisions.

Row REV T+1 in Table (2) displays the share of inequality constrained models, that have smaller revisions than the HP filter output gap or the EC's output gap, respectively. The inequality constrained gap never outperforms it if HP filter cycle of the signaling variable is used to derive the restrictions (0%). It outperforms it in 20% to 30% of cases when the signal is demeaned or differenced. Similar percentages hold when we compare it to the EC's output gap. When it comes to the production function model, it almost never outperforms the HP filter model (0%-5%) but it often outperforms the EC's model (70%). In sum, the HP filter performs best, and the inequality constrained filter performs similar to the EC's output gap. Again, note that in this comparison we have not selected the best performing settings.

3.2.3 Forecasting

We will now compare the models based on their forecasting ability. To do so, we estimate an autoregressive model with the output gap as an exogenous regressor:

$$x_{i,v,t} = \alpha + \sum_{p=1}^{p_{max}} \beta_p x_{i,v,t-p} + \gamma y_{i,v,t-1} + \epsilon_t \qquad t = t_1, ..., y-1 \tag{18}$$

$$\hat{x}_{i,v,y} = \hat{\alpha} + \sum_{p=1}^{p_{max}} \hat{\beta}_p x_{i,v,y-1} + \hat{\gamma} y_{i,v,y-1} \qquad t = y \tag{19}$$

$$\overline{RMSE} = \frac{1}{N}\sum_{i=1}^{N} RMSE_i = \frac{1}{N}\sum_{i=1}^{N}\sqrt{\frac{1}{V}\sum_{v=1}^{V}\left(\hat{x}_{i,v,y} - x_{i,V,y}\right)} \tag{20}$$

where $x_{i,v,t}$ represent the signaling variable for country i in vintage v at period t. To ensure stationarity, we apply differencing to the signaling variable.[25] The model incorporates a constant α, a lag polynomial of order p for the dependent variable $x_{i,v,t-p}$, and the lagged value of the output gap $y_{i,v,t-1}$ as an exogenous predictor. Due to limited data availability, only one lag of the output gap is included. The appropriate length of the autoregressive polynomial is determined using the Bayesian Information Criteria (BIC), which favors parsimonious models and performs well in small sample sizes. The maximum lag length is set to 4, and the model is estimated for each country-vintage pair using in-sample data. Subsequently, the estimated model is used to predict the real-time value of $\hat{x}_{i,v,y}$ for vintage $v = y, s$ using equation (19).

The imposition of inequality constraints has been found to increase the contemporaneous correlation between the output gap and the signaling variable. Moreover, this increase in correlation tends to be higher for models with tighter restrictions. This higher correlation could improve the forecasting power of the output gap. On the other hand, since imposition of inequality constraints increases correlation between $x_{i,v,t-p}$ and $y_{i,v,t-1}$, which both enter eq. 18 as regressors, this could decrease the forecasting performance. It could increase multicolinearity and impose duplicate information into the forecasting model.[26]

In row FCSE of 1 we display RMSE for our example models. We first focus on the baseline model. The HP filter model performs the best, with the lowest RMSE of 12.03. The EC's model exhibits the highest revisions, resulting in an RMSE of 14.4. The RMSE of the inequality-constrained models falls between the HP filter and EC's output gap. As the strength of restrictions increases, the RMSE tends to approach the EC's RMSE (13.43), while relaxing the restrictions brings it closer to the HP filter's RMSE (12.28). For the Production function model, the lowest RMSE is achieved with tight restrictions ($\kappa = 0.05$), and the highest with loose restrictions ($\kappa = 1$). However, exceptions to this pattern are discussed in the subsequent section on the Best Kappa.

On average, the HP filter model exhibited favorable forecasting performance, and the

[25] We forecast the first difference of the inflation rate, the growth rate of unemployment, and the level of the capacity utilization rate.

[26] To elaborate, the imposition of constraints enhances the correlation between the output gap and the selected variable. However, when conditioning on the lag of the variable, the partial correlation between constrained output gap and the variable may decrease.

inequality-constrained filter tended to inherit these desirable properties. Specifically, when the signal was differenced using the HP filter, its own 5-year moving average, or by differencing, the inequality-constrained gap yielded smaller revisions compared to the HP filter gap in 27%, 20%, and 15% of cases, respectively (refer to Table (2)). In stark contrast, when compared to the EC's model, these percentages were 57%, 93%, and 95%.

3.2.4 Signal-to-noise ratio - λ

Table (3) presents a comparison of the performance, with respect to two signal-to-noise ratios commonly used in practice ($\lambda = 10$ and $\lambda = 100$). The table includes results for the inequality constrained gap, with rows corresponding to different combinations of the signaling variables tightness of restrictions κ. The columns show the performance criteria (correlations, revisions, forecast performance), the type of analysis (in-sample, real-time), and the model type (baseline model - BM, production function model - PF). A value of 1 is assigned to a specific combination if the model with $\lambda = 10$ outperforms the same model with $\lambda = 100$, and zero otherwise. The last row reports the average share of instances where the model with $\lambda = 10$ outperformed the model with $\lambda = 100$.

We focus on in-sample statistics, since they are more reliable and are less likely to suffer from end point bias.

Correlations coefficients are on average higher when the smoothing constant is equal to 10 (74% for the baseline model and 93% for the production function model). Setting λ to a low value results in a more volatile trend that captures most of the long-term and medium-term frequency dynamics of the output. The cycle then consists of high-frequency part of the output dynamics. It could be that the association between the signal is higher in the high frequency spectrum, which is why an output gap with a low signal-to-noise ratio produces higher correlation coefficients.

Models with a high signal-to-noise ratio (100) generally exhibit superior performance in terms of revisions compared to models with a low signal-to-noise ratio (10). For instance, in the case of baseline models, those with a high signal-to-noise ratio outperform models with a low signal-to-noise ratio in 82

The results regarding forecasts are mixed. In the production function model, a low signal-to-noise ratio proves to be more effective, surpassing the high signal-to-noise ratio in 97

When the cycle is used to demean the signal, a low signal-to-noise ratio demonstrates superior performance. On the other hand, when employing the own 5-year moving average to detrend the signal, a high signal-to-noise ratio works best. However, the results become mixed when we utilize differencing as a means to detrend the signal.

Overall, a low signal-to-noise ratio tends to lead to stronger correlation between the signal and the output gap, while a high signal-to-noise ratio aids in minimizing revisions. The relationship between the signal-to-noise ratio and forecast accuracy is unclear.

3.2.5 Tightness of restrictions - κ

Table (4) displays best performing settings for the tightness of restrictions parameter κ. In rows, we vary the signaling variable and signal to noise ratio λ. Columns denote comparison criteria (correlations, revisions and forecasts), type of model (baseline and production function model) and type of comparison (in sample or real time).

Strong restrictions with $\kappa = 0.05$ generally yield output gaps with the highest correlation to the signaling variable, particularly in-sample for both the baseline and production function models. However, soft restrictions with $\kappa = 1$ occasionally result in better correlations in real-time, albeit with limited sample certainty. It could also be the case that the utilization of forecasted signaling variables to impose restrictions on the output gap, along with the utilization of their true (final vintage) values to compute real-time correlations, favors models with soft restrictions in instances where the forecasts prove to be incorrect.

When considering revisions, utilizing soft constraints with $\kappa = 1$ is preferable. However, this preference is likely influenced by our choice to employ the normalized RMSE instead of the conventional RMSE for measuring the magnitude of revisions. In the Revisions section, it is clarified that the normalized RMSE was selected to avoid favoring models that simply reduce the amplitude of the output gap.

Finally, the findings regarding the forecasting properties are inconclusive. Approximately two-thirds of the variables demonstrate a preference for soft constraints with $\kappa = 1$ in the baseline model, while one-third of the variables benefit from harder constraints ($\kappa = 0.05, 0.1, 0.05$). In contrast, the production function model exhibits the opposite pattern, with hard constraints leading to more accurate forecasts for three-quarters of the variables.

3.3 Best performing models

Based on the findings above, we conclude that softer restrictions ($\kappa = 1$) tend to result in smaller revisions and better forecasting accuracy, which occurs at the expense of a decreased correlation between the output gap and related variables. However, for specific model combinations, stronger restrictions promote stability and forecast accuracy.

In order to identify the models that perform well across all of the criteria considered, we compiled a table (Table (5)). It contains the models that most frequently outperformed either the EC's or the HP filter model. The first column of the table lists the different comparison criteria, the second column states the object of comparison (EC's output gap or HP filter gap). The first three rows define the restricting variable (x), smoothness constant (λ) and the strength of restricitons (κ).

The cells within the table contain either a zero or a one, with ones indicating that a given model outperformed either the EC's model or the HP filter. The left panel presents baseline models and the right panel the production function models. The bottom row of the table provides the total share of cases in which a out model outperformed the two alternatives. The table includes models which performed best according to this criteria.

The best performing baseline models are those which utilize the unemployment rate to derive the signal for the output gap. The unemployment rate was rendered stationary by either differencing (LUR(1)) or by demeaning with own five year moving average (lur(-) 5yMA). Smoothness constant does not appear to be important since half of the models feature low level (10) and half of the models a high level of smoothness (100). Top performing models feature intermediate strength of restrictions (from 0.1 to 0.5), likely because the stronger restrictions improve to correlation properties but harm the revision and forecast properties. They outperform competing models in 80% of cases. We also note that these models outperformed the EC's model according to all the criteria. They could not outperform the HP filter model according to in-sample revisions and out-of-sample forecasts.

In contrast, the right panel of the figure displays the results for the Production function model. In the production function models the nominal variables affect the unemployment rate's cycle. We did not find a clear advantage w.r.t. the signaling variable as the best performing models include GDP deflator as well as wage inflation. Best performing models featured a low tightness of restrictions ($\kappa = 0.5, 1$), likely reflecting good performance of the HP filter output gap when it comes to revisions and forecasting properties.

The overall findings suggest that for the baseline model, and according to the investigated criteria, using unemployment-based signaling variables, made stationary by differencing or by demeaning with own five year moving average, with medium tightness of restrictions tends to produce best results. For the production function model, which tends to perform worse than the baseline model, the results are less clear.

4 Conclusion

This paper proposed a novel method for estimating the output gap that addresses some of the drawbacks associated with existing methods. Specifically, our approach involves imposing inequality restrictions on the cycle of output, which allows for the output gap estimate to be corrected only when it is inconsistent with related macroeconomic variables, such as inflation or unemployment. For example, if there are inflationary pressures in the economy, the output gap estimate will be constrained to be non-negative. In other periods, the output gap estimate is left unrestricted.

This method offers several advantages over traditional output gap estimation methods. Firstly, it corrects for inconsistencies in the output gap estimate only when necessary, reducing the potential for overfitting and imposition of uncertain model assumptions. Additionally, the method is easy to implement and interpret. Furthermore, the inequality restrictions provide a more flexible and intuitive framework for modeling the output gap, as they allow for different constraints to be imposed depending on the economic context.

Previous literature often addresses the inconsistency of output gap estimates by adding structural equations to the output gap model, which can increase complexity and require additional parameters. This can lead to the need for arbitrary restrictions to obtain reasonable estimates. Our approach, on the other hand, imposes inequality constraints on the output gap, reducing the number of estimated parameters and ensuring that the resulting estimate is consistent with the chosen macroeconomic variable. This allows for a simpler and more stable estimation of the output gap.

The main limitation of our approach is that it may lead to inferior estimates if the true data generating mechanism is similar to the more complex models that have been proposed in the literature. However, in practice, it is often challenging to identify the true data generating mechanism, as evidenced by the variety of output gap models that are available. Furthermore, we demonstrate in a Monte Carlo simulation that even when a multivariate model is the true data generating mechanism, our approach performs reasonably well.

Based on a large set of vintage data sets for the 28 EU countries, we conducted an empirical investigation of the properties of our new approach. Our results indicate that the inequality restrictions significantly increase the correlation coefficient between the variable used for the restriction and the resulting output gap estimate. This suggests that the inequality constraints successfully incorporate information from related macroeconomic variables into the output gap estimate. Additionally, we find that it can effectively reduce the revisions in the output gap estimates, compared to more complex models. It can also improve its forecasting power. Overall, empirical findings provide support for the usefulness of the proposed approach in practical macroeconomic analysis and policy-making.

References

Charles Adams and David T Coe. A systems approach to estimating the natural rate of unemployment and potential output for the united states. *Staff Papers*, 37(2):232–293, 1990.

Ali T Alouani and William D Blair. Use of a kinematic constraint in tracking constant speed, maneuvering targets. *IEEE Transactions on Automatic Control*, 38(7):1107–1111, 1993.

Marianne Baxter and Robert G King. Measuring business cycles: approximate band-pass filters for economic time series. *Review of economics and statistics*, 81(4):575–593, 1999.

Bradley M Bell, James V Burke, and Gianluigi Pillonetto. An inequality constrained nonlinear kalman–bucy smoother by interior point likelihood maximization. *Automatica*, 45(1): 25–33, 2009.

Jaromír Beneš, Kevin Clinton, Roberto García-Saltos, Marianne Johnson, Douglas Laxton, Peter Manchev, and Troy Matheson. Estimating potential output with a multivariate filter. 2010.

Olivier J Blanchard and Danny Quah. The dynamic effects of aggregate demand and supply disturbances, 1988.

Claudio Borio, Piti Disyatat, and Mikael Juselius. Rethinking potential output: Embedding information about the financial cycle. *Oxford Economic Papers*, 69(3):655–677, 2016.

Deutsche Bundesbank. On the reliability of international organisations estimates of the output gap. *Monthly Report April*, 2014.

Marco Buti, Daniele Franco, and Hedwig Ongena. Fiscal discipline and flexibility in emu: the implementation of the stability and growth pact. *Oxford Review of Economic Policy*, 14(3):81–97, 1998.

Stephen Butterworth. On the theory of filter amplifiers. *Wireless Engineer*, 7(6):536–541, 1930.

Fabio Canova. Faq: How do i measure the output gap? 2020.

Jean-Philippe Cayen and Simon Van Norden. The reliability of canadian output-gap estimates. *The North American Journal of Economics and Finance*, 16(3):373–393, 2005.

Lawrence J Christiano and Terry J Fitzgerald. The band pass filter. *international economic review*, 44(2):435–465, 2003.

Peter K Clark. The cyclical component of us economic activity. *The Quarterly Journal of Economics*, 102(4):797–814, 1987.

Todd E Clark and Michael W McCracken. The predictive content of the output gap for inflation: resolving in-sample and out-of-sample evidence. *Journal of Money, Credit and Banking*, pages 1127–1148, 2006.

Iris Claus, Paul Conway, and Alasdair Scott. *The output gap: measurement, comparisons and assessment*. Number 44. Reserve Bank of New Zealand, 2000.

Charles W Cobb and Paul H Douglas. A theory of production. *The American Economic Review*, 18(1):139–165, 1928.

Emilian Dobrescu. Double-conditioned potential output. *Journal for Economic Forecasting*, 3(1):32–50, 2006.

James Durbin and Siem Jan Koopman. *Time series analysis by state space methods*, volume 38. OUP Oxford, 2012.

Anthony Garratt, Donald Robertson, and Stephen Wright. Permanent vs transitory components and economic fundamentals. *Journal of Applied Econometrics*, 21(4):521–542, 2006.

Stefan Gerlach and Frank Smets. Output gaps and monetary policy in the emu area1. *European Economic Review*, 43(4-6):801–812, 1999.

Domenico Giannone, Lucrezia Reichlin, and David Small. Nowcasting: The real-time informational content of macroeconomic data. *Journal of monetary economics*, 55(4):665–676, 2008.

Nachi Gupta and Raphael Hauser. Kalman filtering with equality and inequality state constraints. *arXiv preprint arXiv:0709.2791*, 2007.

James D Hamilton. Why you should never use the hodrick-prescott filter. *Review of Economics and Statistics*, (0), 2017.

Andrew Harvey. Trends, cycles and autoregressions. *The Economic Journal*, 107(440):192–201, 1997.

Andrew Harvey and Thomas Trimbur. Trend estimation and the hodrick-prescott filter. *Journal of the Japan Statistical Society*, 38(1):41–49, 2008.

Andrew C Harvey and Albert Jaeger. Detrending, stylized facts and the business cycle. *Journal of applied econometrics*, 8(3):231–247, 1993.

Karel Havik, Kieran Mc Morrow, Fabrice Orlandi, Christophe Planas, Rafal Raciborski, Werner Röger, Alessandro Rossi, Anna Thum-Thysen, Valerie Vandermeulen, et al. The production function methodology for calculating potential growth rates and output gaps. Technical report, Directorate General Economic and Financial Affairs (DG ECFIN), European Commission, 2014.

Yasuo Hirose, Saori Naganuma, et al. Structural estimation of the output gap: A bayesian dsge approach for the us economy. *Bank of Japan, working paper*, 2007.

Robert J Hodrick and Edward C Prescott. Postwar us business cycles: an empirical investigation. *Journal of Money, credit, and Banking*, pages 1–16, 1997.

Michel Juillard, Ondřej Kameník, Michael Kumhof, Douglas Laxton, et al. *Measures of potential output from an estimated dsge model of the united states*. CNB, 2006.

Güneş Kamber, James Morley, and Benjamin Wong. Intuitive and reliable estimates of the output gap from a beveridge-nelson filter. *Review of Economics and Statistics*, (00), 2017.

Kenneth N Kuttner. Estimating potential output as a latent variable. *Journal of business & economic statistics*, 12(3):361–368, 1994.

Thomas Laubach and John C Williams. Measuring the natural rate of interest. *Review of Economics and Statistics*, 85(4):1063–1070, 2003.

Douglas Laxton and Robert Tetlow. *A simple multivariate filter for the measurement of potential output*. Number 59. Bank of Canada Ottawa, 1992.

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

Roberto S Mariano and Yasutomo Murasawa. A coincident index, common factors, and monthly real gdp. *Oxford Bulletin of economics and statistics*, 72(1):27–46, 2010.

Kieran Mc Morrow, Werner Röger, and Karl Havik. An assessment of the real time reliability of the output gap estimates produced by the eu's production function methodology. Technical report, European Commission, Output GAP Working Group Workshop, September 2015.

James Mitchell, Gian Luigi Mazzi, and Filippo Moauro. Structural var based estimates of the euro area output gap: theoretical considerations and empirical evidences. Technical report, Working paper, 2008.

James C Morley, Charles R Nelson, and Eric Zivot. Why are the beveridge-nelson and unobserved-components decompositions of gdp so different? *Review of Economics and Statistics*, 85(2):235–243, 2003.

Arthur M Okun. *Potential GNP: its measurement and significance*. Yale University, Cowles Foundation for Research in Economics New Haven, 1963.

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

Athanasios Orphanides and Simon Van Norden. The reliability of inflation forecasts based on output gap estimates in real time. *Journal of Money, Credit and Banking*, pages 583–601, 2005.

Athanasios Orphanides, Richard D. Porter, David Reifschneider, Robert Tetlow, and Frederico Finan. Errors in the measurement of the output gap and the design of monetary policy. *Journal of Economics and Business*, 52(1-2):117–141, 2000. URL https://ideas.repec.org/a/eee/jebusi/v52y2000i1-2p117-141.html.

Alban W Phillips. The relation between unemployment and the rate of change of money wage rates in the united kingdom, 1861–1957 1. *economica*, 25(100):283–299, 1958.

Christophe Planas, Werner Roeger, Alessandro Rossi, et al. Does capacity utilisation help estimating the tfp cycle? Technical report, Directorate General Economic and Financial Affairs (DG ECFIN), European Commission, 2010.

Jahan Sarwat and Saber Mahmud Ahmed. What Is the Output Gap?, 2013. URL https://www.imf.org/external/pubs/ft/fandd/2013/09/basics.htm.

Dan Simon. Kalman filtering with state constraints: a survey of linear and nonlinear algorithms. *IET Control Theory & Applications*, 4(8):1303–1318, 2010.

Dan Simon and Tien Li Chia. Kalman filtering with state equality constraints. *IEEE transactions on Aerospace and Electronic Systems*, 38(1):128–136, 2002.

Dan Simon and Donald L Simon. Aircraft turbofan engine health estimation using constrained kalman filtering. In *ASME Turbo Expo 2003, collocated with the 2003 International Joint Power Generation Conference*, pages 485–492. American Society of Mechanical Engineers, 2003.

Vincent Sircoulomb, Ghaleb Hoblos, Houcine Chafouk, and José Ragot. State estimation under nonlinear state inequality constraints. a tracking application. In *Control and Automation, 2008 16th Mediterranean Conference on*, pages 1669–1674. IEEE, 2008.

John B Taylor. Reassessing discretionary fiscal policy. *Journal of economic Perspectives*, 14 (3):21–36, 2000.

Igor Vetlov, Tibor Hlédik, Magnus Jonsson, Kucsera Henrik, and Massimiliano Pisani. Potential output in dsge models. 2011.

Mark W Watson. Univariate detrending methods with stochastic trends. *Journal of monetary economics*, 18(1):49–75, 1986.

Appendix

Model extension (1) - autoregressive cycle

The model presented in eq. (1)-(2b) can be extended with an autoregressive cycle. Following Harvey and Jaeger (1993), we write the the model for the cycle c_t as:

$$\begin{bmatrix} c_t \\ c_t^* \end{bmatrix} = \rho \begin{bmatrix} cos\lambda_c & sin\lambda_c \\ -sin\lambda_c & cos\lambda_c \end{bmatrix} \begin{bmatrix} c_{t-1} \\ c_{t-1}^* \end{bmatrix} + \begin{bmatrix} \kappa_t \\ \kappa_t^* \end{bmatrix}, \qquad t = 1, ..., T \tag{1}$$

where κ_t and κ_t^* are two normally independent white noise disturbances with zero mean and common variance σ_κ^2. λ_c represents the frequency of the cycle in radians.[27] We restricted the cycle periodicity to a range of $(2, 16)$ years.[28][29]

The autoregressive coefficient ρ is also restricted to $0 \leq \rho < 1$.[30] This restriction ensures the model's stationarity. The cycle has a zero mean with a variance of $\sigma_c^2 = \frac{\sigma_\kappa^2}{1-\rho^2}$.

The model in Eq. (1)-(2b), with the cycle as described above (Eq. (1)), is expressed in a state-space format as follows:

$$\alpha_t = \begin{bmatrix} t_t \\ t_{t-1} \\ c_t \\ c_t^* \end{bmatrix}, \quad Z_t = \begin{bmatrix} 1 & 0 & 1 & 0 \end{bmatrix}, \quad T = \begin{bmatrix} 2 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & \rho cos\lambda_c & \rho sin\lambda_c \\ 0 & 0 & -\rho sin\lambda_c & \rho cos\lambda_c \end{bmatrix}$$

[27] The following identity relates radians (λ_c) to periodicity measured in years: $\lambda_c = \frac{2\pi}{y}$, where y is the period in years.

[28] This restrictions is based on the assumption that the average business cycle lasts approx. 8 years.

[29] To implement this restriction, we express λ_c as: $\lambda_c = U_c + \frac{(L_c - U_c) \times \tilde{\lambda_c}}{1+\tilde{\lambda_c}}$, where $L_c = \frac{2\pi}{2}$ represents the lower bound and $U_c = \frac{2\pi}{16}$ represents the upper bound. The model is maximized over $\tilde{\lambda_c}$.

[30] We implement this restriction by expressing ρ as: $\rho = U_\rho + \frac{(L_\rho - U_\rho) \times \tilde{\rho}}{1+\tilde{\rho}}$, where $L_c = 0$ represents the lower bound and $U_c = 0.9999$ represents the upper bound. We maximize over $\tilde{\rho}$.

$$R = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \eta_t = \begin{bmatrix} \zeta_t \\ \kappa_t \\ \kappa_t^* \end{bmatrix}, \quad Q = \begin{bmatrix} \sigma_\zeta^2 & 0 & 0 \\ 0 & \sigma_\kappa^2 & 0 \\ 0 & 0 & \sigma_\kappa^2 \end{bmatrix}$$

where κ_t and κ_t^* are two mutually independent white noise disturbances with a common variance σ_κ. This model can be unstable in empirical applications. Restrictions on σ_κ are (often) needed when series are short (less than 20 observations).

As with the HP filter model, to obtain a meaningful solution, we need to restrict the signal-to-noise ratio. Harvey and Trimbur (2008) showed that augmenting the HP filter model with a cycle cuts out more high-frequency movements compared to when the cycle is assumed to be an irregular disturbance, like in the HP filter model. Therefore, the original signal-to-noise ratio λ needs to be adjusted. They propose to set the noise-to-signal ratio in the model with the cycle, at a value that causes its gain at 0.5 to match the gain of the HP filter at 0.5. Gain is the factor by which the filter amplifies frequencies. Frequencies with gain values after 0.5 tend to be attenuated, while frequencies before it tend to be amplified. An HP filter trend with $\lambda = 10/\lambda = 100$ tends to amplify frequencies at cycle lengths longer than 11/19.8 years and tends to attenuate higher frequencies. They show that the corresponding signal-to-noise ratio for the model with the cycle is:

$$\frac{1}{\lambda_c} = q_c = \left[2sin(\frac{\lambda_{0.5}}{2})\right]^2 \frac{\left(1+\rho^2-2\rho cos(\lambda_c)cos(\lambda_{0.5})\right)\left(1-\rho^2\right)}{1+\rho^4+4\rho^2 cos^2(\lambda_c)-4\left(\rho-\rho^3\right)cos(\lambda_c)cos(\lambda_{0.5})^{+}2\rho^2 cos(2\lambda)} \tag{2}$$

where λ_c, the frequency at which HP filter gain is 0.5, is:[31]

$$\lambda_{0.5} = 2sin^{-1}\frac{\sqrt[4]{\frac{1}{\lambda}}}{2}. \tag{3}$$

Eq. (2) approximates the HP filter's signal-to-noise ratio λ with the signal-to-noise ratio for a model with an autoregressive cycle. It holds in the absence of inequality constraints.

[31] Period, or cycle length, is then $\frac{2\pi}{\lambda_{0.5}}$

Therefore, this approximation is biased for the model with inequality constraints; nevertheless, we found that it creates cycles similar to the HP filter solution. The results for the output gap obtained with the above-defined signal-to-noise ratio are shown in Figure (6). As observed from the figure, they closely resemble the model where the cycle is irregular (see Figure (3)).

An alternative option to calibrate λ is to follow Borio et al. (2016). They calibrate the signal-to-noise ratio by estimating the implied λ for the resulting HP filter estimates:

$$\lambda_{emp} = var(y_t - t_t^{HP})/var(\delta t_t^{HP} - \delta t_{t-1}^{HP}) \tag{4}$$

They note that empirical λ_{emp} tends to be higher because the cycle component ϵ_t is autocorrelated, which in small samples creates a bias. This bias is less present in a model that includes a more complex process for the cycle. The reduced form for the model for the cycle above is ARMA(2,1) model and therefore likely does not include a bias on yearly data.

Model extension (2): The production function model

Our specification of the production function model follows the European Commission's approach:

$$
\begin{aligned}
Y_t &= TFP_t \times L_t^{\alpha} \times K_t^{1-\alpha} && (5)\\
TFP_t &= \frac{Y_t}{L_t^{\alpha} \times K_t^{1-\alpha}} && (6)\\
L_t &= POPWA_t \times PART_t \times (1 - UR_t) \times HOURS_t && (7)\\
PART_t &= 100 \times \frac{\frac{EMP_t}{1-UR_t}}{POPWA_t} && (8)
\end{aligned}
$$

The production function model specifies that the output (Y_t) is a function of labor (L_t), capital (K_t), and total factor productivity (TFP_t). Labor and capital are observable and total factor productivity is calculated as a residual term: $TFP = Y_t/(L^{\alpha} \times K^{1-\alpha})$. In a perfectly competitive market with constant returns to scale, the elasticity of output to labor (α) is equal to the share of labor income in total output. This simple and transparent function captures the fundamental concepts of production theory and is commonly used by international institutions for comparative analysis.

Labor input is decomposed into several components, including the population of working age ($POPWA_t$), the labor market participation rate ($PART_t$), the unemployment rate (UR_t), and the average annual hours worked per employee ($HOURS_t$).

The output gap (GAP_t) at time t is defined as the relative difference between the actual output (Y_t) and the potential output ($YPOT_t$) at time t, expressed as a percentage of potential output:

$$GAP_t = \frac{Y_t - YPOT_t}{YPOT_t} \tag{9}$$

Potential output is obtained by replacing inputs with their estimated trend values (with (6) and (7) inserted in (5)):

$$YPOT_t = \overline{TFP}_t \times (POPWA_t \times PART_t^{HP} \times (1 - \overline{UR}_t) \times HOURS_t^{HP})^{\alpha} \times K_t^{1-\alpha} \quad (10)$$

The EC estimates the trend for total factor productivity ($\overline{TFP}_t$) using a dampened trend model with an AR(2) model for the cycle. However, in some cases, this model fails to provide sensible estimates and alternative models are used. Additionally, the cycle of total factor productivity is used to explain the capacity utilization index (CUBS) in a country-specific manner (see Planas et al. (2010) for more details). In our analysis, we will use a *baseline model* for the TFP trend, which includes inequality constraints derived from the CUBS.

The trend unemployment rate is modeled as a double random walk with an AR(2) cycle, which is equivalent to the HP filter, according to the EC's specification. In their model, the cycle component of unemployment rate is used to explain either the acceleration of inflation in the NAIRU model or the acceleration of unit labor costs in the NAWRU model. However, we will use our *baseline model* for the trend unemployment rate, which incorporates inequality constraints based on various variables. Trends for variables superscripted with HP are estimated with an HP filter. Population of working age is not filtered.

Models

(1) Model Name	(2) Description	(3) Potential Signalling Variable	(4) Possible stationarity inducing transformation	(5) Signal-to-noise ratio (λ) (6) Strength of restrictions (κ)
Baseline	Filter is applied to the real GDP series.	GDP deflator (DEF)	HP filtering/demeaning with 5y MA/I(2) differencing	10/100 0.05/0.1/0.25/0.5/1
		Index of wages (WAGE)	HP filtering/demeaning with 5y MA/I(2) differencing	10/100 0.05/0.1/0.25/0.5/1
		Unemployment rate (LUR)	HP filtering/demeaning with 5y MA/I(1) differencing	10/100 0.05/0.1/0.25/0.5/1
		Capacity Utilization Index (CUBS)	/	10/100 0.05/0.1/0.25/0.5/1
Production function	Filter is applied to TFP (CUBS as signaling variable) and unemployment rate (signaling variable is in the djacent column).	GDP deflator (DEF)	HP filtering/demeaning with 5y MA/I(2) differencing	10/100 0.05/0.1/0.25/0.5/1
		Index of wages (WAGE)	HP filtering/demeaning with 5y MA/I(2) differencing	10/100 0.05/0.1/0.25/0.5/1

This table summarizes the baseline and production function models that were estimated. Each model is briefly described, along with their related settings and parameters, such as the signaling variable, stationarity-inducing transformations, signal-to-noise ratios, and strength of inequality constraints. Additional details on these models can be found in sections Model, Model extensions, and Empirical application.

Table 1: Demonstrative results

Criteria	κ	Baseline Model x = LUR, differenced I(1), $\lambda = 10$ Shadow Gap	HP gap	EC gap	Production Function Model x = DEF, 5y MA, $\lambda = 10$ Shadow Gap	HP gap	EC gap
CORR	0.05	0.67	0.48	0.55	0.42	0.33	0.36
	0.1	0.64	0.48	0.55	0.41	0.33	0.36
	0.25	0.59	0.48	0.55	0.39	0.33	0.36
	0.5	0.56	0.48	0.55	0.37	0.33	0.36
	1	0.53	0.48	0.55	0.36	0.33	0.36
CORR T+1	0.05	0.56	0.50	0.41	0.21	0.19	0.29
	0.10	0.56	0.50	0.41	0.21	0.19	0.29
	0.25	0.54	0.50	0.41	0.21	0.19	0.29
	0.50	0.53	0.50	0.41	0.20	0.19	0.29
	1.00	0.52	0.50	0.41	0.20	0.19	0.29
REV	0.05	0.18	0.15	0.21	0.28	0.20	0.22
	0.10	0.18	0.15	0.21	0.26	0.20	0.22
	0.25	0.17	0.15	0.21	0.23	0.20	0.22
	0.50	0.16	0.15	0.21	0.22	0.20	0.22
	1.00	0.16	0.15	0.21	0.21	0.20	0.22
REV T+1	0.05	0.80	0.80	0.80	0.83	0.75	0.80
	0.1	0.80	0.80	0.80	0.81	0.75	0.80
	0.25	0.79	0.80	0.80	0.78	0.75	0.80
	0.5	0.79	0.80	0.80	0.77	0.75	0.80
	1	0.79	0.80	0.80	0.76	0.75	0.80
FCSE	0.05	13.43	12.03	14.04	1.116	1.133	1.219
	0.1	13.10	12.03	14.04	1.119	1.133	1.219
	0.25	12.67	12.03	14.04	1.124	1.133	1.219
	0.5	12.46	12.03	14.04	1.127	1.133	1.219
	1	12.28	12.03	14.04	1.130	1.133	1.219

This table presents demonstrative results for two models chosen for their strong performance across all five comparison criteria (see section Best performing models). The left panel shows outcomes of the baseline model, using the differenced unemployment rate for constraints. The right panel displays results of the production model, using a 5-year moving average demeaned GDP deflator for constraints (with λ set to 10). The first column lists the comparison criteria: in-sample correlation (**CORR**, see eq. (14)), real-time correlation (**CORR T+1**, see eq. (15)), in-sample output gap revisions (**REV**, see eq. (16)), real-time output gap revisions (**REV T+1**, see eq. (17)), and forecast errors (**FCSE**, see eq. (20)). κ represents the tightness of restrictions for the **IC gap**. Results are also shown for the European Commission's (**EC gap**) and HP filter gap (**HP gap**). We report average expressed from all the EU28 countries and vintages (*baseline model*: from Spring 2004 to Spring 2021, *production function model*: from Autumn 2014 to Spring 2021). The models used data from 1980 onwards or later, depending on availability.

Table 2: Summary of results

Criteria	Type	Baseline Model IC gap >HP gap	IC gap > EC gap	Production Function Model IC gap > HP gap	IC gap > EC gap
CORR	Cycle	100%	100%	100%	100%
	Demeaned	100%	43%	100%	35%
	Differenced	100%	33%	100%	65%
CORR T+1	Cycle	37%	50%	25%	5%
	Demeaned	100%	83%	40%	0%
	Differenced	75%	85%	100%	100%
REV	Cycle	0%	73%	0%	40%
	Demeaned	0%	73%	0%	50%
	Differenced	0%	75%	0%	50%
REV T+1	Cycle	0%	3%	0%	70%
	Demeaned	30%	30%	5%	70%
	Differenced	20%	31%	5%	70%
FCSE	Cycle	27%	57%	75%	100%
	Demean	20%	93%	35%	95%
	Differentiated	15%	95%	15%	15%

The table presents percentages of cases where the inequality constrained output gap exhibited better performance compared to the HP filter (**IC gap>HP gap**) or the European Commission's output gap (**IC gap>EC gap**), based on the following criteria: the in-sample correlation coefficient between the signaling variable and the output gap (**CORR**, see eq. (14)), the real-time correlation coefficient (**CORR T+1**, see eq. (15)), the size of in-sample output gap revisions (**REV**, see eq. (16)), the size of real-time output gap revisions (**REV T+1**, see eq. (17)), and the size of forecast errors (**FCSE**, see eq. (20)). These results represent the average across all considered smoothness parameters ($\lambda = 10, 100$), restricting variables (CPI inflation, wage inflation, unemployment rate, or capacity utilization index), and the tightness of restrictions ($\kappa = 0.05, 0.10, 0.25, 0.50, 1.00$). These averages were expressed from the data for all the EU28 countries and vintages (*baseline model*: from Spring 2004 to Spring 2021, *production function model*: from Autumn 2014 to Spring 2021). The analysis separately presents results when the restricting variable was detrended by HP filtering (**Cycle**), demeaned with its own 5-year moving average (**Demeaned**), or differenced (**Differenced**).

Table 3: Performance of IC Gap estimates - by signal-to-noise ratio (λ) $\lambda = 10 > \lambda = 100$

CRITERIA		Correlations				Revisions				Forecasts	
MODEL		BM		PF		BM		PF		BM	PF
TYPE		In sample	Real-time	In sample	Real-time	In sample	Real-time	In sample	Real-time	Real-time	Real-time
x	κ										
CPI cycle	0.05	1	1	1	1	0	0	0	0	1	1
CPI cycle	0.1	1	1	1	1	0	0	0	0	1	1
CPI cycle	0.25	1	0	1	0	0	0	0	0	1	1
CPI cycle	0.5	1	0	1	0	0	0	0	0	1	1
CPI cycle	1	1	0	1	0	0	0	0	0	1	1
WAGE cycle	0.05	1	0	1	1	0	0	0	0	1	1
WAGE cycle	0.1	1	0	1	0	0	0	0	0	1	1
WAGE cycle	0.25	1	0	1	0	0	0	0	0	1	1
WAGE cycle	0.5	1	0	1	0	0	0	0	0	1	1
WAGE cycle	1	1	0	1	0	0	0	0	0	1	1
LUR cycle	0.05	0	0			1	0			1	
LUR cycle	0.1	0	0			1	0			1	
LUR cycle	0.25	0	0			1	0			1	
LUR cycle	0.5	0	0			1	0			1	
LUR cycle	1	0	0			1	0			1	
CPI (-) 5yMA	2	1	0	1	0	0	0	0	0	0	1
CPI (-) 5yMA	3	1	0	1	0	0	0	0	0	0	1
CPI (-) 5yMA	4	1	0	1	0	0	0	0	0	0	1
CPI (-) 5yMA	5	0	0	1	0	0	0	0	0	0	1
CPI (-) 5yMA	6	0	0	1	0	0	0	0	0	0	1
WAGE (-) 5yMA	2	1	0	1	0	1	0	0	0	0	1
WAGE (-) 5yMA	3	1	0	1	0	0	0	0	0	0	1
WAGE (-) 5yMA	4	1	0	1	0	0	0	0	0	0	1
WAGE (-) 5yMA	5	1	0	0	0	0	0	0	0	0	1
WAGE (-) 5yMA	6	0	0	0	0	0	0	0	0	0	1
LUR (-) 5yMA	2	0	0			0	0			0	
LUR (-) 5yMA	3	0	0			0	0			0	
LUR (-) 5yMA	4	0	0			0	0			0	
LUR (-) 5yMA	5	0	0			0	0			0	
LUR (-) 5yMA	6	0	0			0	0			0	
CPI I(2)	0.05	1	0	1	1	0	0	0	0	1	0
CPI I(2)	0.1	1	0	1	1	0	0	0	0	1	1
CPI I(2)	0.25	1	1	1	1	0	0	0	0	0	1
CPI I(2)	0.5	1	1	1	1	0	0	0	0	0	1
CPI I(2)	1	1	1	1	1	0	0	0	0	0	1
WAGE I(2)	0.05	1	1	1	1	0	0	0	0	0	1
WAGE I(2)	0.1	1	1	1	1	0	0	0	0	0	1
WAGE I(2)	0.25	1	1	1	1	0	0	0	0	0	1
WAGE I(2)	0.5	1	1	1	1	0	0	0	0	0	1
WAGE I(2)	1	1	1	1	1	0	0	0	0	0	1
LUR I(1)	0.05	1	1			1	0			0	
LUR I(1)	0.1	1	1			0	0			1	
LUR I(1)	0.25	1	1			0	0			1	
LUR I(1)	0.5	1	1			0	0			1	
LUR I(1)	1	1	1			0	0			1	
CUBS	0.05	1	1			0	0			0	
CUBS	0.1	1	1			0	0			1	
CUBS	0.25	1	1			0	0			1	
CUBS	0.5	1	1			1	0			1	
CUBS	1	1	1			1	0			1	
AVERAGE		74%	40%	93%	43%	18%	0%	0%	0%	50%	97%

In this table we compare the performance of the inequality constrained output gap model with different signal-to-noise ratios (λ). The table includes comparison criteria, type of analysis (in-sample or real-time), and the model used (*baseline model* or *production function model*). The first two columns specify the restricting variable and its transformation (x) and the tightness of restrictions (κ). Entries with 1 indicate that the model with $\lambda = 10$ performed better than the model with $\lambda = 100$, while 0s indicate the opposite. The last row reports the average share when the model with $\lambda = 10$ outperformed the model with $\lambda = 100$. The average is calculated over EU28 countries and vintages (*baseline model*: from Spring 2004 to Spring 2021, *production function model*: from Autumn 2014 to Spring 2021).

Table 4: Performance of IC Gap estimates - by tightness of restrictions (κ) best performing κ

CRITERIA		Correlations				Revisions				Forecasts	
MODEL		BM		PF		BM		PF		BM	PF
TYPE		In sample	Real-time	In sample	Real-time	In sample	Real-time	In sample	Real-time	Real-time	Real-time
x	λ										
DEF cycle	10	0.05	0.05	0.05	0.05	1.00	1.00	1.00	1.00	1.00	0.05
WAGE cycle	10	0.05	0.05	0.05	1.00	1.00	1.00	1.00	1.00	0.50	0.05
LUR cycle	10	0.05	1.00			1.00	1.00			1.00	
DEF cycle	100	0.05	0.10	0.05	1.00	1.00	1.00	1.00	1.00	0.10	0.05
WAGE cycle	100	0.05	0.50	0.05	1.00	1.00	1.00	1.00	1.00	0.05	0.05
LUR cycle	100	0.05	1.00			1.00	1.00			1.00	
DEF (-) 5yMA	10	0.05	0.05	0.05	0.10	1.00	1.00	1.00	1.00	0.10	0.05
WAGE (-) 5yMA	10	0.05	0.05	0.05	1.00	1.00	1.00	1.00	1.00	1.00	1.00
LUR (-) 5yMA	10	0.05	0.05			1.00	1.00			1.00	
DEF (-) 5yMA	100	0.05	0.05	0.05	1.00	1.00	1.00	1.00	1.00	0.05	0.05
WAGE (-) 5yMA	100	0.05	0.05	0.05	1.00	1.00	1.00	1.00	1.00	1.00	1.00
LUR (-) 5yMA	100	0.05	0.05			1.00	0.50			1.00	
DEF I(2)	10	0.05	0.05	0.05	0.05	1.00	1.00	1.00	1.00	0.10	0.50
WAGE I(2)	10	0.05	0.05	0.05	0.10	1.00	1.00	1.00	1.00	0.50	0.05
LUR I(1)	10	0.05	0.05			1.00	0.50			1.00	
CUBS	10	0.05	1.00			1.00	1.00			1.00	
DEF I(2)	100	0.05	0.05	0.05	0.05	1.00	1.00	1.00	1.00	1.00	1.00
WAGE I(2)	100	0.05	0.05	0.05	0.05	1.00	1.00	1.00	1.00	1.00	0.05
LUR I(1)	100	0.05	0.05			1.00	0.50			1.00	
CUBS	100	0.05	1.00			1.00	1.00			1.00	

In this table we compare the performance of the inequality constrained output gap using different values of the tightness of restrictions parameter (κ). The first two columns specify the restricting variable and its transformation (x), along with the signal-to-noise ratio (λ). The first three rows indicate the comparison criteria, the type of analysis (in-sample or real-time), and the model (*baseline model* or the *production function model*). Numeric entries in the table indicate the optimal value of κ for each variable, signal-to-noise ratio, criteria, and model.

Table 5: Best performing models

		Baseline Model				Production Function Model		
	x	lur (-) 5yMA	lur (-) 5yMA	LUR I(1)	LUR I(1)	DEF cycle	WAGE cycle	DEF (-) 5yMA
	λ	100	100	10	10	10	100	10
	κ	0.1	0.25	0.25	0.5	1	1	0.5
CORR	$IC > HP$	1	1	1	1	1	1	1
	$IC > EC$	1	1	1	1	1	1	1
CORR T+1	$IC > HP$	1	1	1	1	1	0	1
	$IC > EC$	1	1	1	1	0	1	0
REV	$IC < HP$	0	0	0	0	0	0	0
	$IC < EC$	1	1	1	1	1	1	1
REV T+1	$IC < HP$	1	1	1	1	0	0	0
	$IC < EC$	1	1	1	1	1	1	1
FCSE	$IC < HP$	0	0	0	0	1	1	1
	$IC < EC$	1	1	1	1	1	1	1
AVERAGE		80%	80%	80%	80%	70%	70%	70%

In this table, we present the models that performed the best on average, according to the considered performance criteria (*CORR*, *CORR T+1*, *REV*, *REV T+1*, *FCS*). The variable x indicates the signaling variable, λ represents the smoothness constant, and κ denotes the tightness of restrictions. The left-most column indicates the comparison criteria and models being compared (IC - inequality constrained output gap, HP - HP filter output gap, EC - European Commission output gap). The last row labeled $AVERAGE$ shows the share of instances in which the selected model outperformed the EC's and the HP filter model, denoted by entries with a value of 1.

Figures

Figure 1: Monte Carlo experiment, correct signal

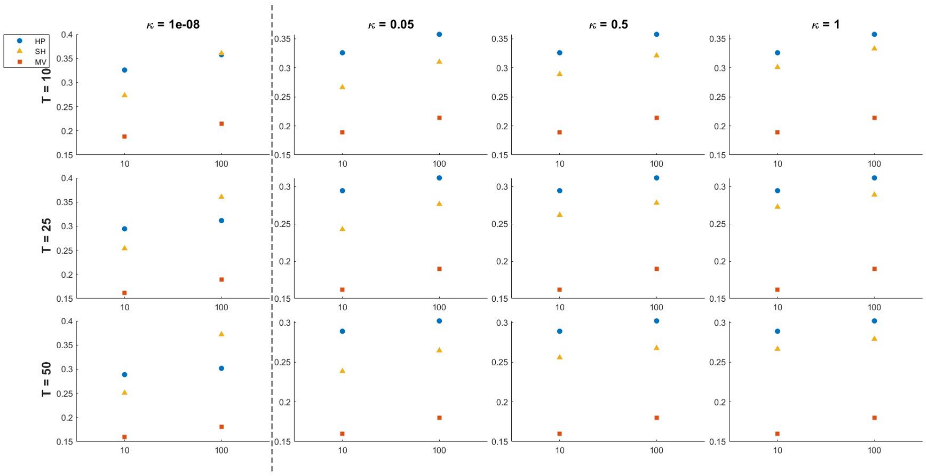

Figure displays RMSEs (root mean square errors) for three models: the multivariate model (**red** circles, see eq. (10-12)), HP filter (**blue** circles, see eq. (1-2b)) and the inequality constrained model (**orange** circles, see eq. (7). Multivariate model is the true model which generated the data. Data generating mechanism is described in section Small sample properties - Monte Carlo verification. RMSE is a measure of the differences between the true (simulated) cycle and the cycle re-estimated from simulated data.

Figure 2: Monte Carlo experiment, false signal

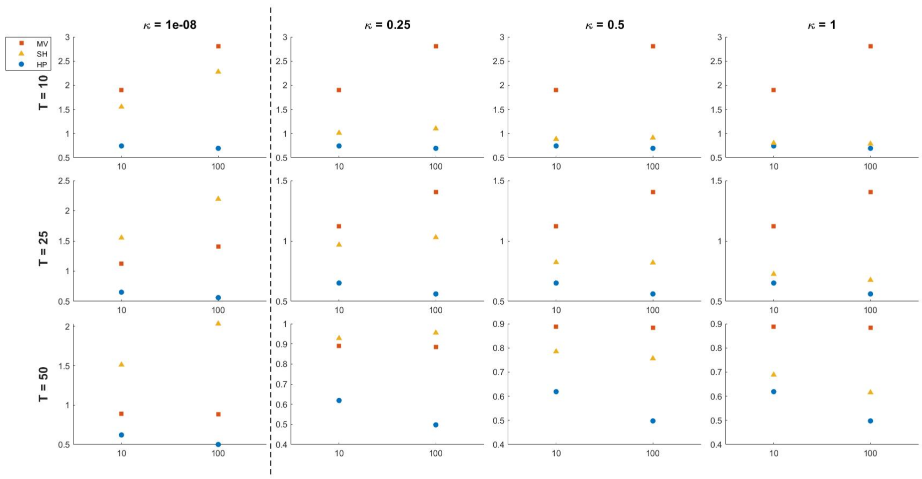

Figure displays RMSEs (root mean square errors) for three models: the multivariate model (**red** circles, see eq. (10-12)), HP filter (**blue** circles, see eq. (1-2b)) and the inequality constrained model (**orange** circles, see eq. (7). For all models we replaced the variable which generated the data with an irrelevant variable (the false signal case). Details on the data generating mechanism is described in section Small sample properties - Monte Carlo verification. RMSE is a measure of the differences between the true (simulated) cycle and the cycle re-estimated from simulated data.

Figure 3: Output Gap, EU 27, Autumn 2017

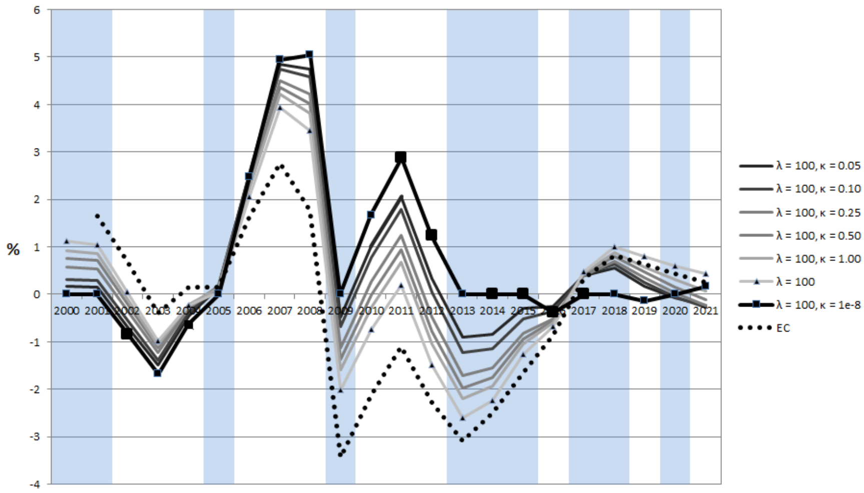

Figure displays output gaps with varying degrees of tightness of constraints. The restricting variable is the cycle of GDP deflator. λ stands for signal-to-noise ratio, κ for the strength of constraint. Lower κ implies a tighter constraint (see eq. (9) for more details). We plot the following lines: **dotted black line** is the EC's output gap estimate (vintage Autumn 2019), **black line with square marker** is the output gap with (hard) inequality constraint, **gray line** is the output gap with (soft) inequality constraint (κ defines the degree of tightness), **gray line with triangle marker** is the unconstrained output gap (it is equivalent to an HP filter gap with signal-to-noise ratio $\lambda = 100$). **Shaded areas** are periods in which inequality constraint was binding the value of the output gap with (hard) inequality constraint. Data vintage is Autumn 2019, yearly data from 1995-2021 for EU27.

Figure 4: GDP deflator and cycle, EU 27, Autumn 2017

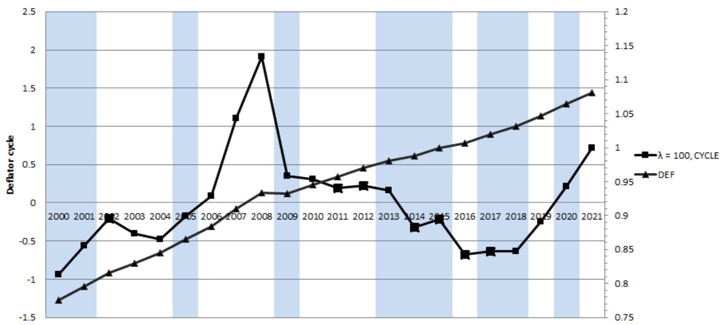

The Figure displays the GDP deflator and its cycle estimated using the HP filter with a signal-to-noise ratio of $\lambda = 100$. The shaded areas represent periods where the inequality constraint was binding in the case of the output gap with a hard inequality constraint. The data vintage is Autumn 2019 and the yearly data covers the period from 1995-2021 for the EU27.

Figure 5: Potential output, EU 27, Autumn 2017

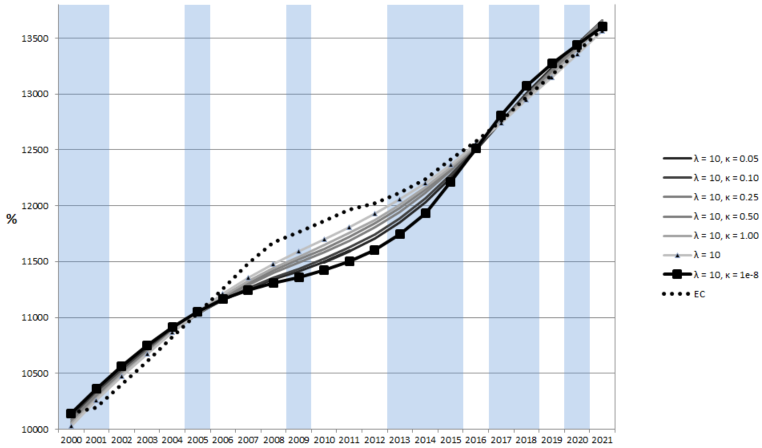

The figure illustrates the impact of varying the degree of tightness of the inequality constraint on potential output, with the cycle of GDP deflator as the restricting variable. The signal-to-noise ratio is denoted by λ, the strength of the constraint κ. The following lines are shown in the plot: the **dotted black line** represents the European Commission's estimate of potential output (vintage Autumn 2019), the **black line with square marker** represents potential output with a hard inequality constraint, the **gray line** represents potential output with a soft inequality constraint (where the degree of tightness is defined by κ), and the **gray line with triangle marker** represents the unconstrained potential output (equivalent to an HP filter trend with signal-to-noise ratio $\lambda = 100$). **Shaded areas** indicate periods where the inequality constraint was binding for the output gap with a hard inequality constraint. The data used in the plot is based on the Autumn 2019 vintage, with yearly data from 1995-2021 for the EU27.

Figure 6: Output Gap model with an AR cycle, EU 27, Autumn 2019

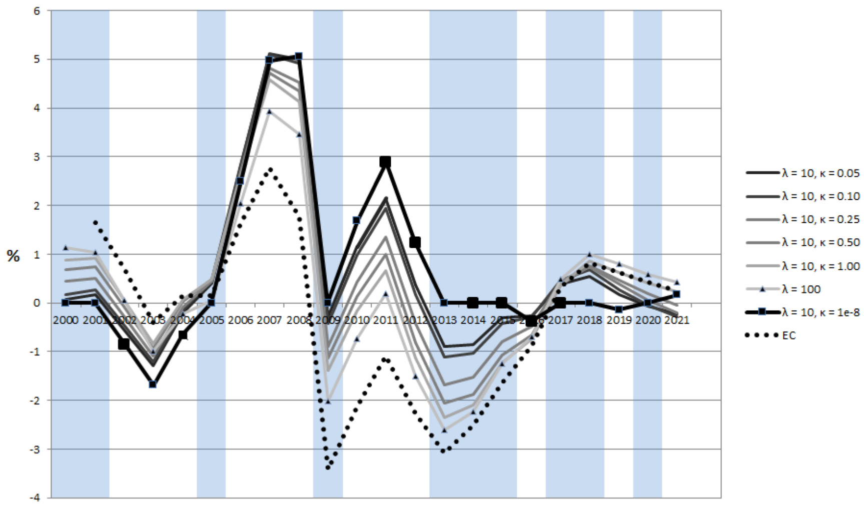

This figure shows the output gap under varying degrees of inequality constraint. The cycle of GDP deflator is the variable that is constrained, with λ representing the signal-to-noise ratio and κ the strength of constraint. Lower values of κ indicate a tighter constraint (see eq. (9) for more details). The lines are as follows: the **dotted black line** shows the EC's output gap estimate (vintage Autumn 2019), the **black line with square marker** shows the output gap with (hard) inequality constraint, the **gray line** shows the output gap with (soft) inequality constraint (with κ defining the degree of tightness), and the **gray line with triangle marker** shows the unconstrained output gap (equivalent to an HP filter gap with signal-to-noise ratio $\lambda = 100$). The **shaded areas** indicate periods in which the inequality constraint was binding for the case of the output gap with (hard) inequality constraint. The data is based on the vintage Autumn 2019 and covers yearly data from 1995-2021 for EU27.