DEUTSCHE BUNDESBANK
EUROSYSTEM

Technical Paper
A primer on optimal policy projections

01/2024

Thomas Dengler
Rafael Gerke
Sebastian Giesen
Daniel Kienzler
Joost Röttger
Alexander Scheer
Johannes Wacks

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet http://www.bundesbank.de

Non-technical summary

Bundesbank staff regularly calculate optimal policy projections (OPPs), allowing to derive recommendations about how to set the monetary policy instruments for a given macroeconomic outlook. As such, they can be a valuable input for the preparation of monetary policy meetings held by the Governing Council of the European Central Bank. An article in the Bundesbank's December 2023 Monthly Report (Deutsche Bundesbank, 2023) discusses the basic features of OPPs and how they can deliver valuable insights in practice. To remain accessible to a general audience, the article refrains from covering technical details about the computation of such projections. This note complements the article by providing a more technical treatment of the subject to interested readers.

The present note describes how to formally state and solve an optimal policy problem based on *(i)* a baseline projection, *(ii)* impulse response functions for monetary policy shocks at different time horizons and *(iii)* a loss function. The baseline projection provides the starting point for the OPP and describes how the economy is projected to evolve if monetary policy instruments follow a baseline, e.g. evolve as expected by financial markets. The impulse response functions capture how policymakers can affect the projected variables by selecting certain paths for their instruments, like the nominal interest rate or asset purchases. Together with a loss function, which allows to rank potential outcomes in terms of their desirability, one can then calculate how to best deviate from the baseline policy path, optimally trading off different goals over a specified policy horizon. The resulting optimal policy path and the associated paths for macroeconomic variables of interest constitute an optimal policy projection.

The note first shows in detail how to derive an OPP, using a textbook version of the New Keynesian model. It then illustrates various features of the OPP approach based on three examples. To highlight the versatility of the OPP approach, the note then shows results for two additional applications. The first application builds on a medium-scale DSGE model and a baseline projection taken from the Broad Macroeconomic Projection Exercise (BMPE) of the Eurosystem. The second application shows optimal policies for a heterogeneous-agent New Keynesian (HANK) model.

Nichttechnische Zusammenfassung

Fachleute der Bundesbank berechnen regelmäßig optimale Politikprojektionen (OPPs). Diese Analysen erlauben für einen vorgegebenen gesamtwirtschaftlichen Ausblick, Hinweise oder Empfehlungen hinsichtlich des Einsatzes der geldpolitischen Instrumente abzuleiten. Sie können daher einen wertvollen Beitrag für die Vorbereitung von geldpolitischen Sitzungen des EZB-Rats liefern. Ein Sonderaufsatz im Dezember 2023 des Monatsberichts der Bundesbank (Deutsche Bundesbank, 2023) beschreibt die wesentlichen Eigenschaften von OPPs und führt aus, wie sie wertvolle Einsichten für die geldpolitische Praxis liefern können. Um für eine breite Leserschaft zugänglich zu sein, sieht der Sonderaufsatz davon ab, technische Aspekte hinsichtlich der Berechnung von OPPs im Detail auszuführen. Die vorliegende Notiz ergänzt den Sonderaufsatz, indem sie interessierten Leserinnen und Lesern einen ersten Zugang zu technischen Details bietet.

Die vorliegende Notiz beschreibt wie man ein optimales Politikproblem auf Basis von drei Elementen formal aufschreibt und löst. Diese drei Elemente umfassen *(i)* eine Basislinie (Projektion), *(ii)* Impuls-Antwort-Funktionen für geldpolitische Schocks zu unterschiedlichen Zeithorizonten und *(iii)* eine Verlustfunktion. Die Basislinie projiziert, wie sich relevante makroökonomische Größen verhalten, wenn die geldpolitischen Instrumente bestimmten Annahmen folgen, z.B. sich entsprechend den Erwartungen an den Finanzmärkten entwickeln. Die Impuls-Antwort-Funktionen bilden ab, wie die Entscheidungsträger die projizierten Größen beeinflussen können, wenn sie bestimmte Pfade für die geldpolitischen Instrumente, wie den Nominalzins oder Anleihekäufe, wählen. Zusammen mit der Verlustfunktion, die es erlaubt, unterschiedliche Pfade für Zielgrößen miteinander zu vergleichen und zu bewerten, kann mit diesen Elementen berechnet werden, wie von der Basislinie in optimaler Weise über einen bestimmten Politikhorizont abgewichen werden sollte. Eine optimale Politikprojektion, die unterschiedliche Ziele gegeneinander abwägt, umfasst dann einen optimalen Pfad für die Politikinstrumente sowie einen sich ergebenen Pfad für die Zielgrößen.

Anhand einer Lehrbuchversion des neukeynesianischen Modells zeigt die vorliegende Notiz zunächst im Detail, wie eine OPP hergeleitet wird. Im Anschluss illustriert sie einige Eigenschaften des OPP-Ansatzes anhand von drei Beispielen. Um die Vielseitigkeit des OPP-Ansatzes zu verdeutlichen, werden ferner Ergebnisse von zwei weiteren Anwendungen gezeigt. Die erste Anwendung verwendet ein mittelgroßes DSGE-Modell und eine Basislinie aus der Broad Macroeconomic Projection Exercise (BMPE) des Eurosystems. Die zweite Anwendung zeigt optimale Politiken für ein neukeynesianisches Modell mit heterogenen Haushalten.

A primer on optimal policy projections*

THOMAS DENGLER† RAFAEL GERKE‡ SEBASTIAN GIESEN§
DANIEL KIENZLER¶ JOOST RÖTTGER‖ ALEXANDER SCHEER**
JOHANNES WACKS††

January 25, 2024

Abstract

Optimal policy projections (OPPs) offer a flexible way to derive scenario-based policy recommendations. This note describes how to calculate OPPs for a simple textbook New Keynesian model and provides illustrations for various examples. It also demonstrates the versatility of the approach by showing OPP results for simulations conducted using a medium-scale DSGE model and a New Keynesian model with heterogeneous households.

Keywords: Optimal monetary policy, macroeconomic projections, New Keynesian models, household heterogeneity

JEL Classification: C63, E20, E31, E47, E52, E58

---

*The views expressed in this paper are those of the authors and do not necessarily represent those of the Deutsche Bundesbank or the Eurosystem.

†Deutsche Bundesbank. Email: thomas.dengler@bundesbank.de.
‡Deutsche Bundesbank. Email: rafael.gerke@bundesbank.de.
§Deutsche Bundesbank. Email: sebastian.giesen@bundesbank.de.
¶Deutsche Bundesbank. Email: daniel.kienzler@bundesbank.de.
‖Deutsche Bundesbank. Email: joost.roettger@bundesbank.de.
**Deutsche Bundesbank. Email: alexander.scheer@bundesbank.de.
††Deutsche Bundesbank. Email: johannes.wacks@bundesbank.de.

1 Introduction

Optimal policy projections (OPPs) allow the derivation of policy recommendations for a given macroeconomic scenario. As such, they can be a valuable input for policymakers in central banks, who regularly face the difficult task of deciding how to best set their policy instruments given the current macroeconomic outlook.[1] Bundesbank staff regularly calculate OPPs in preparation for monetary policy meetings held by the Governing Council of the European Central Bank (ECB), as discussed in a recent article in the Bundesbank's December 2023 Monthly Report (Deutsche Bundesbank, 2023). By taking a high-level perspective, the Monthly Report article abstracts from technical details of OPP computation to focus on intuition and policy implications. This note provides a brief primer on OPPs that complements the article by offering a more technical treatment of the subject.

The note is organised as follows. Section 2 describes the main idea of the OPP approach, including its advantages and disadvantages. Using a simple textbook model, Section 3 presents the technical details needed to calculate an OPP. Section 4 provides some illustrative OPP simulations for that model. Section 5 shows OPP results for more complex macroeconomic models. Section 6 offers some final remarks.

2 Main idea

The computation of optimal policy projections requires three ingredients: *(i)* a baseline projection, *(ii)* impulse response functions for monetary policy shocks at different time horizons and *(iii)* a loss function.

The baseline projection, which is the first key ingredient, is the starting point for the policy analysis. It provides the policymaker's outlook for the economy based on a given monetary policy path. If the available monetary policy instruments, like the short-term nominal interest rate and asset purchases, do not deviate from those paths assumed for the baseline, the economy evolves as projected by the baseline as well. The second key ingredient describes how monetary policy can affect the economy. This is captured through impulse response functions (IRFs) for macroeconomic variables of interest to contemporaneous and anticipated future (news) shocks to the monetary policy instruments. The IRFs are critical for the OPP approach, as they reflect how the economy responds to an announced change

[1] In principle, OPPs can also be used for fiscal or macro-prudential policies. The OPP concept has been around for quite some time (see e.g. Svensson and Tetlow, 2005) and is used by various central banks around the world (see Deutsche Bundesbank, 2023). Recently, the macroeconomic literature has shown a renewed interest in the subject, exploring new aspects, extensions and applications of the approach (see e.g. Bersson et al., 2019; de Groot et al., 2021; Harrison and Waldron, 2021; Hebden and Winkler, 2021; Barnichon and Mesters, 2023a,b; McKay and Wolf, 2023a,b).

in a policy instrument for today or future periods. Capturing the causal effect of policy changes at different horizons on outcomes is crucial for performing counterfactual policy simulations that are not subject to the Lucas critique (see e.g. McKay and Wolf, 2023b). The loss function is the third key ingredient. It reflects the objective of the policymaker, e.g. a preference for inflation and output gap stabilisation, and allows the ranking of different macroeconomic outcomes. Given such a loss function and policy IRFs, one can then calculate how to best deviate from the baseline policy path, optimally trading off different goals over a specified policy horizon. The resulting policy path and the associated paths for macroeconomic variables of interest constitute an optimal policy projection.

The OPP ingredients can come from the same source, but they do not have to. For instance, ingredients *(i)*-*(iii)* can all come from a single dynamic stochastic general equilibrium (DSGE) model with a Taylor rule that governs interest rate policy. Specifically, one can obtain ingredient *(i)* by simulating the evolution of the economy for a specific macroeconomic shock, ingredient *(ii)* based on shocks to the Taylor rule at different horizons, and ingredient *(iii)* as a quadratic approximation of social welfare for the model economy (see Galí, 2015).[2] For central bank policymakers, however, it is more natural for *(i)* to stem from a macroeconomic projection prepared by central bank staff. For the policy simulations conducted within the Bundesbank, the baseline projections come from the (Broad) Macroeconomic Projection Exercise, short (B)MPE (see Deutsche Bundesbank, 2023).[3] These projections are not the outcome of a single, particular model, but rather incorporate various sources of information. They assume a future path for monetary policy instruments, as given by market expectations at a specific cut-off date. Based on these projections, policymakers may ask what the interest rate path should look like to best fulfill the central bank mandate. For policy simulations performed by central bank staff, the loss function usually represents an operationalisation of the mandate rather than a measure of social welfare. For example, Bundesbank staff use a loss function that reflects the mandate of the Eurosystem for its policy simulations (see Deutsche Bundesbank, 2023). To represent policymakers' ability to affect the economy with their instruments in an empirically plausible way, the policy IRFs, i.e. OPP ingredient *(ii)*, typically come from estimated medium- to large-scale DSGE models.[4] This is also the case for the simulations performed

[2]As shown e.g. by de Groot et al. (2021), the optimal policy response calculated with the OPP approach coincides with the optimal Ramsey policy derived using a recursive Lagrangian approach (see Marcet and Marimon, 2019) and applied from a timeless perspective (see Woodford, 2003). Whereas the former solution is based on a sequence-space representation of the model, the latter is based on an augmented state-space representation. The latter representation includes natural state variables (such as capital) as well as additional "artificial" state variables, whose purpose is to enforce the policymaker's promises in future periods.

[3]Each quarter, macroeconomic projections are prepared as inputs for the decision making of the ECB Governing Council. The BMPE is conducted for June and December of each year by staff from the Eurosystem's national central banks (NCBs) and the ECB. The MPE is conducted by ECB staff for March and September of each year. For further details, see European Central Bank (2016).

[4]While the use of model-implied IRFs is the most common approach among central banks, conceptually, one could also use IRFs estimated from adequate data (see e.g. Barnichon and Mesters, 2023a,b; McKay and

by Bundesbank staff (see Deutsche Bundesbank, 2023).

For policymakers, a key advantage of the OPP approach is its flexibility. For example, given policy IRFs and a loss function, it is easy to compute OPPs for different macroeconomic scenarios, as captured by different baseline projections. Recent studies have also demonstrated the flexibility of the approach by showing that it does not rely on the assumption of a perfectly credible policymaker or completely rational agents, extending the standard OPP framework (see Svensson and Tetlow, 2005) along various dimensions (see e.g. de Groot et al., 2021; Harrison and Waldron, 2021; Hebden and Winkler, 2021). However, despite its flexibility, the OPP approach does rely on certain assumptions that limit the set of feasible applications in practice. For example, the validity of the OPP approach described in this note ultimately relies on linear (model) relationships. Although it is possible to incorporate some nonlinearities, such as occasionally binding constraints, asymmetric objectives or temporary regime changes (see de Groot et al., 2021; Harrison and Waldron, 2021; Hebden and Winkler, 2021), the OPP approach requires such nonlinearities to apply only temporarily, leaving the steady state of the model unaffected. Therefore, the approach is not well-suited for policy simulations with permanent effects, such as a transition to a permanently different central bank balance sheet. Likewise, the OPP approach cannot be applied if policy affects the information set of agents, for example when policy actions signal a policymaker's hidden type or private information in general (see McKay and Wolf, 2023b). As a result, the framework cannot capture, for example, a (de-)anchoring of inflation expectations that is driven by a feedback between agents' beliefs about the central bank's commitment to price stability and the policymaker's actions (see e.g. Lu et al., 2016). Discrete macroeconomic events, such as a sovereign default or a run on a large bank, and how they are affected by policy cannot be properly accounted for in the linear model structure either. Therefore, when applying OPPs in practice, one has to be aware of these limitations and use them only as one input for decision making, complemented by additional information and analyses (see Deutsche Bundesbank, 2023).

3 Computing an optimal policy projection

Using a simple textbook model, this section presents the technical details needed to compute an OPP.

Wolf, 2023b). However, doing so is a non-trivial task and only works as an approximation, as it requires estimates of anticipated future policy changes. McKay and Wolf (2023b) show that "empirical IRFs" do not provide a good approximation in case of an interest rate peg. Therefore, optimal policy projections that involve an occasionally binding effective lower bound on the nominal interest rate may likely be unreliable when based on such IRFs. More importantly, however, one can flexibly adapt a DSGE model to reflect a change in the economic environment, e.g. a steepening of the Phillips curve, to allow for additional (rather new) instruments, like asset purchases, or to incorporate expert judgment. Doing so for empirical IRFs is not as straightforward.

3.1 Model description

Consider the textbook linearised New Keynesian (NK) model (see e.g. Galí, 2015) summarised by the three-equation system

$$y_t = \mathbb{E}_t y_{t+1} - \sigma^{-1}\left(i_t - \mathbb{E}_t \pi_{t+1}\right) + u_t^d, \tag{1}$$

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t \pi_{t+1} + u_t^s, \tag{2}$$

$$i_t = \rho_i i_{t-1} + (1-\rho_i)\left(\phi_y y_t + \phi_\pi \pi_t\right) + \sum\nolimits_{k\geq 0} \varepsilon_{t|t-k}, \tag{3}$$

where π_t denotes inflation, y_t the output gap and i_t the nominal interest rate – all expressed in log-deviations from their steady-state values.[5]

The first equation is the dynamic IS curve, which governs the demand side of the economy and is subject to a (demand) shock u_t^d. The IS curve is derived from the Euler equation of a representative household with intertemporal elasticity of substitution σ^{-1}. It provides the connection between real economic activity and the expected real interest rate $i_t - \mathbb{E}_t\pi_{t+1}$. The second equation is the forward-looking NK Phillips curve (NKPC), linking current inflation π_t to current real economic activity y_t, expected future inflation $\mathbb{E}_t\pi_{t+1}$ and a cost-push shock u_t^s. The household discount factor is denoted as β, whereas κ denotes the slope of the NKPC. The third equation is a Taylor rule, which determines the current nominal interest rate i_t as a function of inflation and the output gap in the same period. The Taylor rule also depends on the lagged nominal rate i_{t-1} as well as contemporaneous and lagged exogenous shocks. Specifically, $\varepsilon_{t|t-k}$ denotes a shock to the policy rate announced in period $t-k$, with $k \geq 0$, but realised in period t. It is assumed that the model system has a unique solution.

3.2 How to use impulse response functions for policy (news) shocks

IRFs for model variables of interest to policy shocks at different time horizons capture how policy changes affect the economy's behaviour (see Figure 1).[6] Let T denote the policy horizon over which the optimal policy will be calculated. Define the policy shock vector $\varepsilon_t \equiv \left(\varepsilon_{t|t}, \varepsilon_{t+1|t}, \varepsilon_{t+2|t}, \ldots, \varepsilon_{t+T|t}\right)'$, which collects policy shocks announced in period t for periods $t+k$, $0 \leq k \leq T$. In addition, define the vector $d^{x,k} \equiv \left(d_0^{x,k}, d_1^{x,k}, d_2^{x,k}, \ldots, d_T^{x,k}\right)'$, which contains the impulse response coefficients for variable $x \in \{\pi, i, y\}$ to a policy shock $\varepsilon_{t+k|t}$ that is announced today to take place in $0 \leq k \leq T$ periods from now.[7] The length of

[5] While the simple textbook version of the NK model is used to describe the details for the OPP computation, it is straightforward to generalise it and apply it to any linearised DSGE model (see e.g. de Groot et al., 2021; Hebden and Winkler, 2021) – as highlighted by the applications shown in Section 5.

[6] The depicted impulse responses use the model parametrisation from Section 4.

[7] Note that $d_s^{x,k}$ is not necessarily equal to zero, even if $k > s$, i.e. even if the considered horizon ends before the period in which the shock actually materialises. This is because of anticipation effects: A news

Figure 1: Impulse responses to anticipated policy shocks at different time horizons

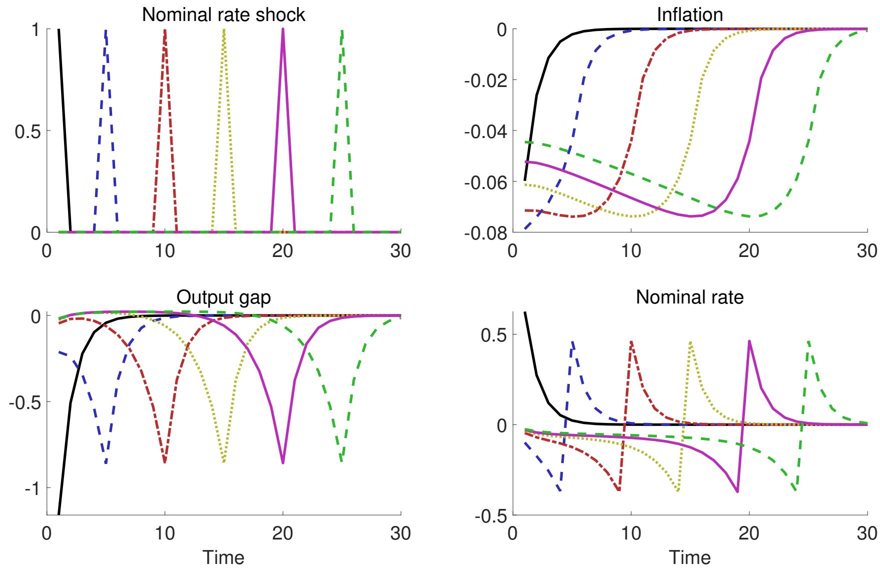

Notes: All variables are expressed in percentage deviations from steady state.

this IRF vector is thus $T+1$. To illustrate the relationship between the IRF coefficients and the model variables for a given period $t+s$, suppose the economy starts in period t in the model's steady state. Ceteris paribus, a shock $\varepsilon_{t+k|t}$ then implies that variable x takes on the value $x_{t+s} = d_s^{x,k} \times \varepsilon_{t+k|t}$ in period $t+s$. More generally, by exploiting the linear model structure, one can calculate x_{t+s} upon announcement of policy shocks ε_t in period t as

$$x_{t+s} = \sum_{0 \leq k \leq T} d_s^{x,k} \varepsilon_{t+k|t}. \tag{4}$$

Alternatively, one can express it in matrix notation as

$$x_{t+s} = \left(d_s^{x,0}, d_s^{x,1}, d_s^{x,2}, \ldots, d_s^{x,T}\right) \begin{pmatrix} \varepsilon_{t|t} \\ \varepsilon_{t+1|t} \\ \varepsilon_{t+2|t} \\ \vdots \\ \varepsilon_{t+T|t} \end{pmatrix}. \tag{5}$$

Let the vector $X_t \equiv (x_t, x_{t+1}, x_{t+2}, \ldots, x_{t+T})'$ denote the entire time path for variable x from period t to $t+T$, implied by the policy shocks ε_t. Each element of this vector, x_{t+s}, captures the response of variable x in period $t+s$ to contemporaneous and anticipated future shocks

shock about the future can have an impact on variable x already today. Interest rate forward guidance is one manifestation of this phenomenon. This can be seen in Figure 1, which shows IRFs for shocks to the interest rate rule that are all announced in the first period but realise at different dates. Only the policy shock that is realised in the first period (solid black line) constitutes a contemporaneous shock without anticipation effects.

$\varepsilon_t = \left(\varepsilon_{t|t}, \varepsilon_{t+1|t}, \varepsilon_{t+2|t}, \dots, \varepsilon_{t+T|t}\right)'$. In order to calculate X_t, construct the coefficient matrix D^x based on the individual vectors $d^{x,k}$,

$$D^x \equiv \left[d^{x,0}, d^{x,1}, d^{x,2}, \dots, d^{x,T}\right] = \begin{bmatrix} d_0^{x,0} & d_0^{x,1} & d_0^{x,2} & \cdots & d_0^{x,T} \\ d_1^{x,0} & d_1^{x,1} & d_1^{x,2} & \cdots & d_1^{x,T} \\ d_2^{x,0} & d_2^{x,1} & d_2^{x,2} & \cdots & d_2^{x,T} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ d_T^{x,0} & d_T^{x,1} & d_T^{x,2} & \cdots & d_T^{x,T} \end{bmatrix}. \tag{6}$$

Given the shock vector $\varepsilon_t = \left(\varepsilon_{t|t}, \varepsilon_{t+1|t}, \varepsilon_{t+2|t}, \dots, \varepsilon_{t+T|t}\right)'$, which collects policy (news) shocks known in period t to take place in periods $t, t+1, t+2, \dots, t+T$, one can now conveniently write

$$X_t = D^x \varepsilon_t, \tag{7}$$

to compute the entire path $X_t = (x_t, x_{t+1}, x_{t+2}, \dots, x_{t+T})'$ for variable x, or

$$X_t = \begin{bmatrix} d_0^{x,0} & d_0^{x,1} & d_0^{x,2} & \cdots & d_0^{x,T} \\ d_1^{x,0} & d_1^{x,1} & d_1^{x,2} & \cdots & d_1^{x,T} \\ d_2^{x,0} & d_2^{x,1} & d_2^{x,2} & \cdots & d_2^{x,T} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ d_T^{x,0} & d_T^{x,1} & d_T^{x,2} & \cdots & d_T^{x,T} \end{bmatrix} \begin{pmatrix} \varepsilon_{t|t} \\ \varepsilon_{t+1|t} \\ \varepsilon_{t+2|t} \\ \vdots \\ \varepsilon_{t+T|t} \end{pmatrix} = \begin{pmatrix} \sum_{0\le k\le T} d_0^{x,k} \varepsilon_{t+k|t} \\ \sum_{0\le k\le T} d_1^{x,k} \varepsilon_{t+k|t} \\ \sum_{0\le k\le T} d_2^{x,k} \varepsilon_{t+k|t} \\ \vdots \\ \sum_{0\le k\le T} d_T^{x,k} \varepsilon_{t+k|t} \end{pmatrix}. \tag{8}$$

3.3 How to incorporate a baseline scenario

So far, the model's steady state has served as the implicit baseline scenario. That is, if all policy shocks are equal to zero, the economy would simply stay in the steady state forever. However, due to the linear model structure, one can express the time path of model variable $x \in \{\pi, i, y\}$ for an arbitrary baseline,

$$B_t^x \equiv \left(x_t^b, x_{t+1}^b, x_{t+2}^b, \dots, x_{t+T}^b\right)', \tag{9}$$

and policy shocks ε_t simply as

$$X_t = B_t^x + D^x \varepsilon_t. \tag{10}$$

It is easy to see that without policy shocks, i.e. $\varepsilon_t = (0, 0, 0, \dots, 0)'$, the dynamics of variable x are given by the baseline B_t^x.

3.4 How to implement an arbitrary interest rate path

In general, one can use the policy IRFs to assess the consequences of arbitrary interest rate counterfactuals for the economy (see Laseen and Svensson, 2011).[8] Specifically, one can evaluate the consequences of an arbitrary interest path $\bar{I}_t$ by following a two-step procedure. First, one has to find policy shocks ε_t^*, such that $\bar{I}_t = B_t^i + D^i \varepsilon_t^*$ holds, where B_t^i is the baseline path for the interest rate. Fortunately, one can accomplish this analytically due to the linearity of the model:

$$\varepsilon_t^* = \left(D^i\right)^{-1} \left(\bar{I}_t - B_t^i\right). \tag{11}$$

Given ε_t^*, one can then calculate the counterfactual path for variable x conditional on the interest path $\bar{I}_t$ as

$$\overline{X}_t = B_t^x + D^x \varepsilon_t^*. \tag{12}$$

3.5 How to compute the optimal interest rate path

An optimal interest rate path for the economy selects policy shocks such that the implied paths for the economy minimise a given loss function. While counterfactual policy projections based on specific interest rate paths can provide valuable insights, policymakers are usually interested in the optimal interest rate path. To determine what is optimal, a loss function is required to represent the policymaker's preferences, allowing the ranking of different paths for the economy. Specifically, the intertemporal loss function $\sum_{0\leq s\leq T} \beta^s L_{t+s}$ is considered with quadratic period loss function

$$L_t = \frac{(\pi_t)^2 + \lambda(y_t)^2 + w_i\left(\Delta_t^i\right)^2}{2}, \tag{13}$$

where $\Delta_t^i \equiv i_t - i_{t-1}$ denotes the difference between the interest rate in period t and the previous period $t-1$.

As will be shown in detail below, when combined with a linear system of equations, which represent the (equality) constraints faced by the policymaker, a quadratic loss function permits a formulation of the policy problem as a simple linear-quadratic programming problem. The constraints reflect the relationship between policy shocks (choices) and outcomes, generating potential trade-offs for the policymaker. How such trade-offs are resolved then critically depends on the assumed specification of the loss function. A key

[8] Analogously to shocks to the short-term policy rate, as presented here, one can proceed with shocks to central bank asset holdings in models that explicitly model such asset holdings by the central bank (see Section 4.3 and Appendix A).

advantage of a linear-quadratic programming problem is that it admits an analytical solution if no further types of constraints, such as inequality constraints, are imposed. And even if such constraints are imposed, the problem can typically still be solved efficiently by numerical means.

For simple DSGE models, it is often possible to obtain microfounded relative weights for the loss function based on a second-order welfare approximation. For instance, the simple textbook model implies the weights $\lambda = \kappa/\theta$ and $w_i = 0$, with θ denoting the demand elasticity for intermediate goods (see e.g. Galí, 2015). In this case, fluctuations of inflation and the output gap around their long-run values are costly for households in the economy, with λ capturing the relative importance of the output gap. A policymaker would then aim at stabilising these two variables, which may involve a trade-off between stabilising inflation and the output gap, within and across periods (see Section 4). According to this microfoundation, the third term in the period loss function does not matter. Indeed, a positive weight w_i is usually assumed by central bank practitioners to avoid large swings in interest rates (see Svensson and Tetlow, 2005), e.g. to capture an aversion to financial market volatility in a reduced form. Since the mandate of a central bank usually explicitly states price stability as a key objective, inflation stabilisation naturally enters the loss function. A positive relative output gap weight λ may then reflect a dual mandate (Federal Reserve System) or a medium-term focus (Eurosystem).[9]

As mentioned above, without further restrictions on the policy instrument, such as an effective lower bound (ELB) on the policy rate, one can compute the optimal policy path analytically. To do so, first define the vector

$$Z_t \equiv \begin{bmatrix} \Pi_t \\ Y_t \\ I_t - I_{t-1} \end{bmatrix}, \tag{14}$$

which stacks the time paths for the three variables relevant to the loss function and is therefore of dimension $3(T+1) \times 1$. Next, define the vector

$$B_t \equiv \begin{bmatrix} B_t^{\pi} \\ B_t^{y} \\ B_t^{\Delta_i} \end{bmatrix}, \tag{15}$$

which is of the same dimension as Z_t and contains the stacked baseline paths for the vari-

[9]See Deutsche Bundesbank (2023) for details on the interpretation of the loss function from the perspective of the Eurosystem.

ables, and the matrix

$$D \equiv \begin{bmatrix} D^{\pi} \\ D^{y} \\ D^{\Delta_i} \end{bmatrix}, \tag{16}$$

which is of size $3(T+1) \times (T+1)$ and contains the impulse response coefficient vectors for different shock horizons. One can now conveniently express Z_t as

$$Z_t = B_t + D\varepsilon_t. \tag{17}$$

The policy problem can then be written as

$$\min_{\varepsilon_t} \left\{ \frac{1}{2} Z_t{}' W Z_t \right\}, \tag{18}$$

subject to (17). The matrix W is of size $3(T+1) \times 3(T+1)$ and reflects discounting and the relative weights in the loss function L_t. Specifically, it is given by

$$W = \begin{bmatrix} W_\beta & 0 & 0 \\ 0 & \lambda W_\beta & 0 \\ 0 & 0 & w_i W_\beta \end{bmatrix}, \tag{19}$$

for the example in this section, with diagonal matrix $W_\beta \equiv diag\left(1, \beta, \beta^2, \ldots, \beta^T\right)$.

By substituting out Z_t in the minimisation problem above and ignoring constant terms irrelevant to the policy choice, one can simplify the problem to

$$\min_{\varepsilon_t} \left\{ \frac{1}{2} \varepsilon_t' D' W D \varepsilon_t + B_t' W D \varepsilon_t \right\}. \tag{20}$$

The analytical solution to this problem is given by the first-order condition

$$\varepsilon_t^* = -\left(D' W D\right)^{-1} (B_t' W D)', \tag{21}$$

and the optimal policy projection for variable x can then be calculated as

$$X_t^* = B_t^x + D^x \varepsilon_t^*. \tag{22}$$

Although it can no longer be achieved analytically, it is feasible to solve the minimisation problem also subject to an ELB constraint, $i_t = B_t^i + D^i \varepsilon_t \geq \log\left(i_{ELB}/i_{SS}\right)$, where i_{ELB} denotes the ELB on the gross nominal rate and i_{SS} the steady-state value of the gross nominal rate.[10]

[10]In MATLAB, one could, for instance solve this constrained linear-quadratic programming problem by using the built-in function `quadprog.m`. To use the function, rewrite the constraint as $-D^i \varepsilon_t \leq B_t^i - \log\left(i_{ELB}/i_{SS}\right)$.

4 Illustrative examples

This section illustrates the OPP approach for the texbook model based on three examples.

4.1 Optimal policy response to a cost-push shock

To illustrate the OPP approach, this section starts by simulating the optimal policy responses to a transitory and a persistent cost-push shock (see e.g. Galí, 2015). To compute the respective responses, a baseline projection is required. For this application, the baseline, shown by the dashed red lines in Figures 2 and 3, is given by the response of the textbook NK model to a cost-push shock u_t^s, which follows an AR(1) process with persistence parameter $0 \leq \rho_s < 1$ and is shown in the upper left panels. The interest rate response is therefore based on the Taylor rule in this case, with coefficients $\rho_i = 0.7$, $\phi_\pi = 1.5$ and $\phi_y = 1$. The remaining parameter values for the model are $\beta = 0.98$, $\kappa = 0.03$, $\sigma = 1$ and $\theta = 6$.[11] The loss function weights are $\lambda = \kappa/\theta$ and $w_i = 0$. The solid blue lines in Figures 2 and 3 display the optimal policy paths. Whereas Figure 2 shows the optimal response to a transitory cost-push shock ($\rho_s = 0$), Figure 3 does the same for a persistent one ($\rho_s = 0.5$).

Since the cost-push shock is inflationary and contractionary at the same time (see e.g. the dashed red lines), it introduces a trade-off for a central bank that cares about inflation and output gap stabilisation. To counter the inflationary pressure, the central bank can only hike the nominal rate and depress real economic activity even further. Compared to a simple Taylor rule, the optimal policy under commitment can, however, optimally balance this trade-off not only within a given period but also intertemporally. In the case of a purely transitory shock (see Figure 2), the optimal policy smooths the interest rate increases and the resulting output gap losses over time, which necessitates a less strong output gap reaction on impact.[12] For the persistent shock (see Figure 3), whose size was calibrated to yield a similar inflation response in the first period as in the case of the transitory shock, the optimal policy smooths the response even further. Indeed, the hiking response is delayed and counterbalanced by an expansionary response in earlier periods, which reduces the contractionary impact of the shock on the output gap in earlier periods.

4.2 Optimal policy response in a liquidity trap

To demonstrate that the OPP approach can also accommodate an occasionally binding ELB constraint, the optimal policy response to a persistent negative demand shock is sim-

[11] All of these parameter values are kept throughout the remainder of this section.

[12] The response is also stronger compared to the baseline case, which is based on a Taylor rule with a rather high output gap reaction coefficient. Note that the assumed Taylor rule coefficients (or even its functional form) have no bearing on the calculated OPP.

Figure 2: OPP for a transitory cost-push shock

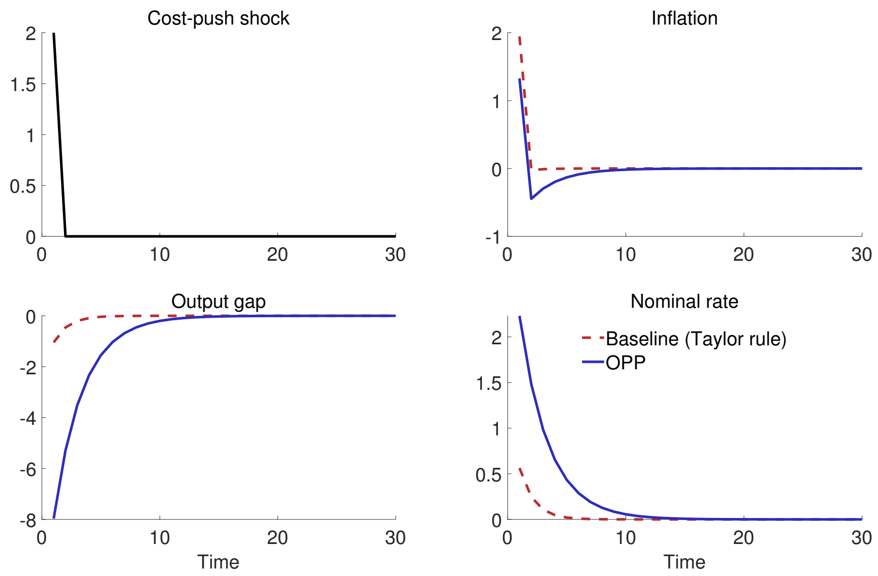

Notes: All variables are expressed in percentage deviations from steady state.

Figure 3: OPP for a persistent cost-push shock

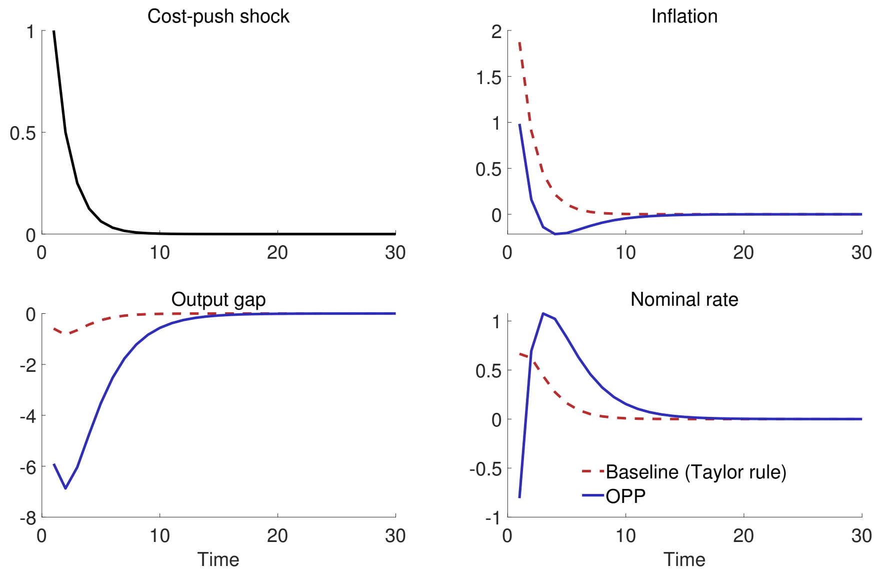

Notes: All variables are expressed in percentage deviations from steady state.

ulated.[13] For such a demand shock u_t^d, Figure 4 shows time paths for the nominal rate, the inflation rate and the output gap for three cases. The dashed red lines display the baseline scenario (a version of the Taylor rule (3) that is subject to an ELB), the solid blue lines the unconstrained OPP and the dashed-dotted yellow lines the ELB-constrained OPP.[14] As shown by the solid blue lines, an unconstrained policymaker can stabilise inflation and the output gap entirely by appropriately setting the nominal rate, reflecting a "divine coincidence" (see Blanchard and Galí, 2007). By contrast, the constrained OPP cannot do so, displaying a lower-for-longer element to reduce the impact of the shock on the economy in earlier periods (see e.g. Eggertsson and Woodford, 2003).

4.3 Optimal policy response with two instruments

At the ELB, asset purchases have become an accepted policy tool that can – to some extent – substitute for additional interest rate cuts (see e.g. Gerke et al., 2022). For the same demand shock depicted in Figure 4, Figure 5 shows time paths for the nominal rate, the inflation rate and the output gap if the central bank can also choose to buy long-term government bonds.[15] By using asset purchases, the central bank can provide additional stimulus at the ELB, which raises output and inflation compared to the interest-rate-only scenario in Figure 4. Moreover, monetary policy does not rely as much on a lower-for-longer interest rate policy in this case. As a result, inflation and output only slightly overshoot their target values in later periods to provide stimulus during earlier ELB periods. Notice that without an ELB, the optimal policy does not make use of the asset purchases, reflecting the assumption that $w_i = 0$ holds in this section, whereas the use of asset purchases is costly ($w_q > 0$).[16] This simulation illustrates that it is straightforward to extend the OPP analysis to applications with multiple instruments.[17]

[13] An AR(1) process is assumed for u_t^d, with persistence parameter $\rho_d = 0.85$.

[14] The baseline path for the interest rate implied by the Taylor rule with an occasionally binding ELB constraint, $i_t \geq \log(i_{ELB}/i_{SS})$, is computed based on the policy IRFs and the method proposed by Holden (2016). Specifically, the built-in MATLAB function `intlinprog.m` is used in that case, solving a mixed-integer linear programming problem. This baseline is only shown as a reference for the optimal policy, which does not depend on whether the ELB is imposed for the baseline or not. It only matters whether the policy problem faces the constraint.

[15] This analysis is conducted in Hills et al. (2021) based on an augmented state-space model representation. Harrison (2017) provides the microfoundation for this model version. Compared to the textbook model from before, the IS curve is now slightly different and given by $y_t = \mathbb{E}_t y_{t+1} - \sigma^{-1}(i_t - \mathbb{E}_t \pi_{t+1} - \gamma q_t) + u_t^d$, where $q_t \geq 0$ denotes the fraction of government bonds held by the central bank. Central bank asset holdings give rise to two additional terms in the loss function, $L_t = 0.5\left[(\pi_t)^2 + \lambda(y_t)^2 + w_i\left(\Delta_t^i\right)^2 + w_q(q_t)^2 + w_{\Delta q}\left(\Delta_t^q\right)^2\right]$, where $\Delta_t^q \equiv q_t - q_{t-1}$. The law of motion for asset holdings is given by the exogenous process $q_t = \sum_{k\geq 0} \varepsilon_{t|t-k}^q$, where $\varepsilon_{t|t-k}^q$ denotes a shock to asset holdings announced in period $t-k$, with $k \geq 0$, and realised in period t.

[16] The parameter values $w_q = 1/16000$ and $w_{\Delta q} = 0$ are assumed for the simulation in this section (see Hills et al., 2021).

[17] Details about the policy problem with two instruments can be found in Appendix A.

Figure 4: OPPs for a persistent negative demand shock

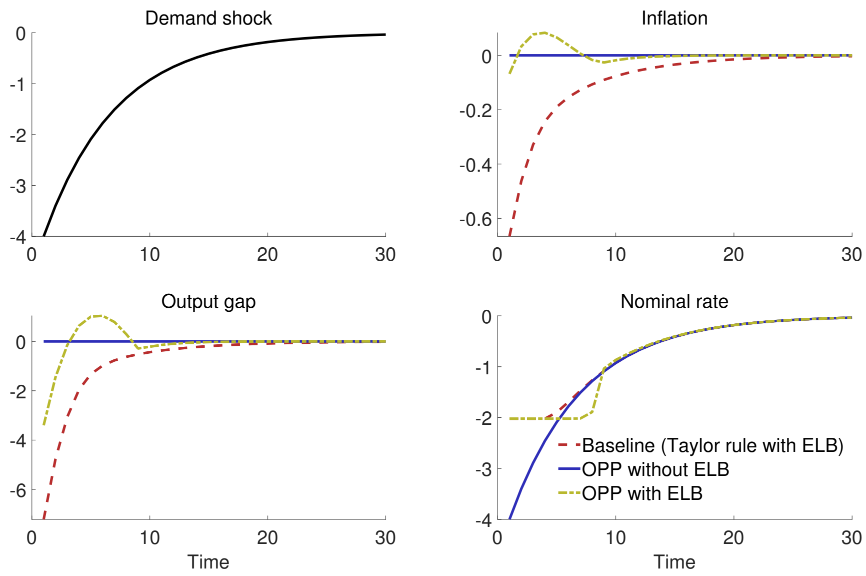

Notes: All variables are expressed in percentage deviations from steady state.

Figure 5: OPPs for a persistent negative demand shock (with asset purchases)

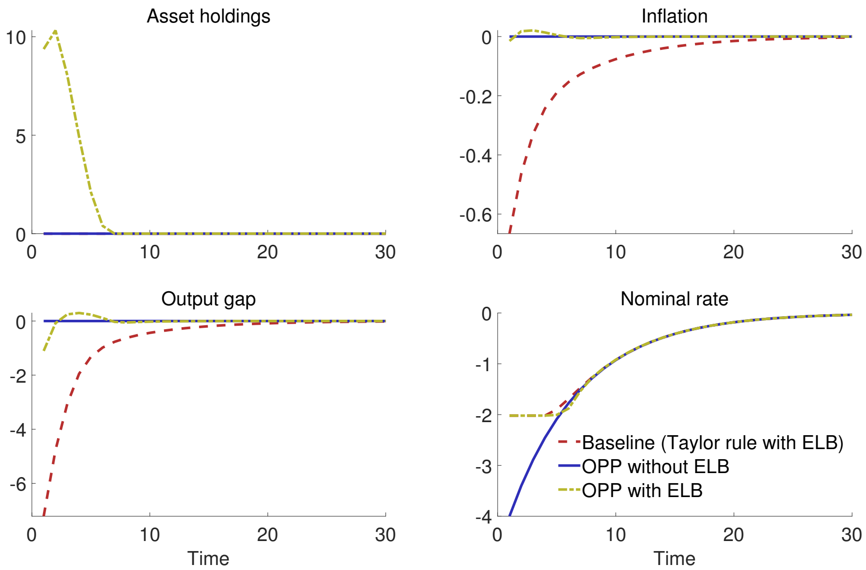

Notes: All variables are expressed in percentage deviations from steady state.

5 Applications

This section highlights how the OPP approach can be used in practice for applications that go beyond the simple textbook model considered so far. To this end, Section 5.1 considers a medium-scale DSGE model and an externally provided baseline, whereas Section 5.2 shows results for an application with a heterogeneous-agent NK model.

5.1 Medium-scale DSGE model and external baseline

Conceptually, it is straightforward to use the OPP approach also for more complex models or to use external sources for the baseline projections. For instance, the policy simulations performed by Bundesbank staff use a state-of-the-art medium-scale New Keynesian DSGE model that can capture the transmission of conventional and unconventional monetary policy measures (see Gerke et al., 2022). As in the previous section, the simulations allow policymakers to jointly optimise over the paths for interest rates and central bank asset purchases. Moreover, the simulations rely on the (B)MPE for the baseline projection (see Deutsche Bundesbank, 2023). To be suitable for policy analysis, a macroeconomic model must certainly be more complex than the simple textbook model considered so far. In the case of the Bundesbank, the policy model features – amongst other things – two household types, physical capital, consumption habits, nominal price and wage rigidities, financial frictions and bounded rationality. This added complexity is necessary to obtain a good fit for the model estimation, which in turn is important to capture the transmission of monetary policy in an empirically plausible way.

From a computational perspective, what matters for the OPPs is the availability of IRFs for all available policy instruments. Furthermore, the model needs to include those variables that enter the loss function. While the computation of the IRFs is more involved for more complex models, the model complexity does not affect the computation of the OPPs once the IRFs are available.[18] Although the calculation of OPPs based on an externally provided baseline projection, such as the (B)MPE, does not come with additional technical difficulties per se, certain adjustments might be necessary in practice.[19] Figure 6 shows OPP results based on a version of the policy model with a detailed fiscal policy sector.[20] This

[18]The results shown in this section were computed using the toolkit provided by de Groot et al. (2021). The toolkit allows to compute OPPs for DSGE models that can be formulated in Dynare (see Adjemian et al., 2011). As a result, it cannot be used for the application shown in the next section.

[19]For instance, it might be necessary to transform the time series from the baseline projection in a way that makes them consistent with the model-implied variables. This may involve transforming annual into quarterly data or de-trending certain empirical time series.

[20]For these simulations, only optimal interest rate policy is considered, i.e. the policymaker does not optimise over two instruments. See Deutsche Bundesbank (2023) for a discussion of how different evolutions of central bank asset holdings might affect an optimal interest rate path. Since there is no choice about asset

Figure 6: OPP for a medium-scale DSGE model based on the June 2023 BMPE

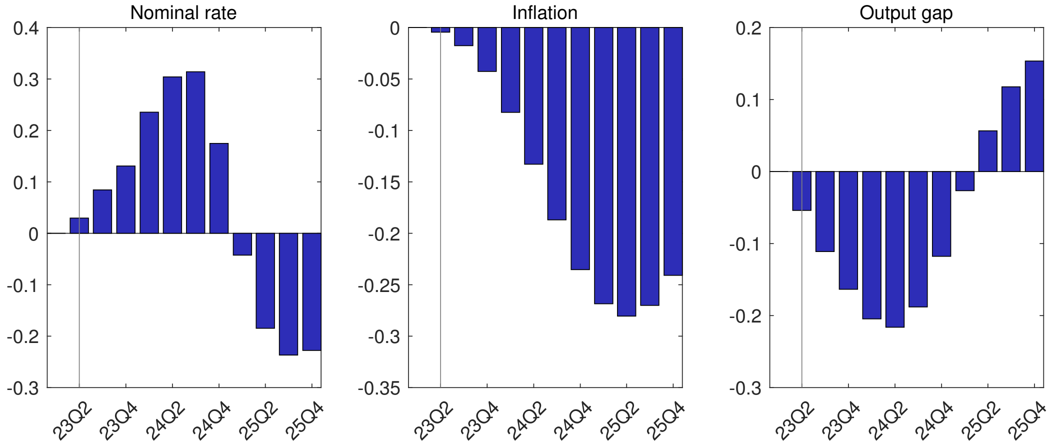

Notes: All variables are shown for the OPP and expressed relative to the June 2023 BMPE baseline. The nominal rate is expressed in absolute deviations of the OPP from the baseline (both in annual percentage points). Inflation is expressed in cumulative absolute deviations of the OPP from the baseline (both in annual percentage points). The output gap is expressed in absolute deviations of the OPP from the baseline (both in percent).

model version is estimated on recent euro area data and able to capture how monetary-fiscal interactions affect the optimal monetary policy. The OPP is calculated based on the BMPE baseline scenario prepared by Eurosystem staff for June 2023 (see European Central Bank, 2023). Compared to the interest rate path envisaged by financial market participants at that time, the OPP-implied interest rate is higher during the first quarters of the projection horizon. This policy reflects that, from the perspective of June 2023 (see vertical dark grey lines), the baseline projection foresaw inflation remaining too high for too long, suggesting that a tighter than expected monetary policy was justified for the next quarters.

5.2 Household heterogeneity

The OPP approach also permits the computation of optimal policies for models with heterogeneous households (see McKay and Wolf, 2023a,b), which capture how monetary policy transmission is affected by movements in wealth and income inequality (see e.g. Kaplan and Violante, 2018).[21] As argued in the previous section, if IRFs for the policy instruments are available, obtaining OPPs for such models is no more difficult than for the textbook NK

purchases in this simulation, the loss function (13) is used with weights $\lambda = 0.25$ and $w_i = 3$. However, instead of quarterly inflation, annual (year-on-year) inflation enters the loss function as an argument in this case. Similarly, we use annualised instead of quarterly interest rates.

[21] Note that, as with prior OPP applications, optimal policy does not affect the steady state, i.e. only optimal stabilisation policies around a given steady state can be studied this way. Under certain assumptions, it is possible to also derive a microfounded loss function for a model with household heterogeneity, which includes terms related to distributional concerns (see McKay and Wolf, 2023a). By contrast, this section again uses the ad-hoc loss function (13), with weights $\lambda = 0.25$ and $w_i = 0.1$.

Figure 7: OPPs for a cost-push shock in model versions with heterogeneous households

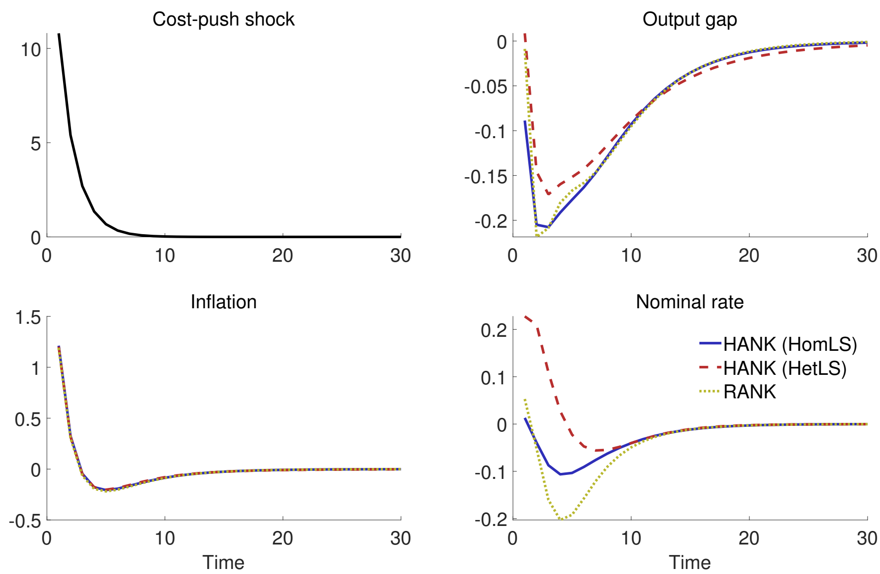

Notes: All variables are expressed in percentage deviations from steady state. HANK (HomLS): model version with heterogeneous households and homogeneous labour supply. HANK (HetLS): model version with heterogeneous households and heterogeneous labour supply. RANK: model version with representative household.

model. However, getting there, i.e. computing these policy IRFs, is usually substantially more difficult (see e.g. Dobrew et al., 2023; Gerke et al., 2024).[22] Figure 7 shows OPPs for a cost-push shock in a sticky-wage heterogeneous-agent New Keynesian (HANK) model and its representative-agent New Keynesian (RANK) counterpart.[23] For the HANK model, two cases are distinguished (see Gerke et al., 2024, for details). In the first case, the now common assumption is made that all households work the same amount of hours (HomLS case). In the second case, individual hours can vary across households, taking the individual income and wealth situation into account (HetLS case). The OPPs predict inflation to behave very similarly across all three model versions, reflecting a very flat New Keynesian wage Phillips curve. The output gap behaves similarly in all three cases as well, but differences are more noticeable. For the nominal rate, the OPP-implied path is similar only in the RANK model and the HANK model version with homogeneous labour supply.

[22]It is more difficult due to the need to set up and use more advanced numerical solution techniques and the additional costs in terms of computational resources. Simply writing down a model in Dynare, as is possible for the application in the previous section, is usually not feasible anymore. For details on how to compute IRFs for such models, see e.g. Reiter (2009), Bayer and Luetticke (2020) and Auclert et al. (2021).

[23]To make the OPPs easier to compare across model versions, the impulse responses of the model variables to the cost-push shock in the RANK model are also used as the baseline projection for the HANK case. The loss function and its parametrisation are the same across model versions. While the calibrations of the different model versions feature different values for the household discount factor, we assume the same discount factor for the policymaker in all cases ($\beta = 0.9951$). Since the baseline and the loss function are the same, differences between the OPPs reflect differences in the monetary transmission mechanism, which in turn shapes the trade-offs faced by the policymaker.

As shown in Gerke et al. (2024), the transmission of monetary policy shocks is dampened when labour supply is allowed to vary across the income distribution. As a result, to stabilise inflation and real economic activity, the policymaker has to be more restrictive to achieve a similar outcome as under homogeneous labour supply.

6 Final remarks

Optimal policy projections offer a flexible way to derive monetary policy recommendations for a given scenario and are a useful tool to stimulate discussions. The OPP results shown in this note for a medium-scale DSGE model and a HANK model have furthermore highlighted the versatility of this approach. At the same time, it is important to keep in mind the assumptions that underlie the OPP approach. Not all of these can be (completely) overcome by extensions or workarounds, in particular when it comes to certain forms of nonlinearities, such as a (de-)anchoring of inflation expectations or financial (in)stability. This underlines that the OPP recommendations must always be taken with a grain of salt and complemented by other sources of information and additional analyses.

References

Adjemian, S., H. Bastani, M. Juillard, F. Karamé, J. Maih, F. Mihoubi, W. Mutschler, G. Perendia, J. Pfeifer, M. Ratto, and S. Villemot (2011): "Dynare: Reference Manual, Version 4," Dynare Working Papers 1, CEPREM.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021): "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 89, 2375–2408.

Barnichon, R. and G. Mesters (2023a): "Evaluating Policy Institutions —150 Years of US Monetary Policy—," Mimeo.

——— (2023b): "A Sufficient Statistics Approach for Macro Policy," *American Economic Review*, 113, 2809–2845.

Bayer, C. and R. Luetticke (2020): "Solving Discrete Time Heterogeneous Agent Models with Aggregate Risk and Many Idiosyncratic States by Perturbation," *Quantitative Economics*, 11, 1253–1288.

Bersson, B., P. Hürtgen, and M. Paustian (2019): "Expectations Formation, Sticky Prices, and the ZLB," Deutsche Bundesbank Discussion Paper No. 34/2019.

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

de Groot, O., F. Mazelis, R. Motto, and A. Ristiniemi (2021): "A Toolkit for Computing Constrained Optimal Policy Projections (COPPs)," ECB Working Paper Series 2555.

Deutsche Bundesbank (2023): "Model-Based Recommendations for Monetary Policy Decision-Making," Monthly Report, December 2023, 37-54.

Dobrew, M., R. Gerke, S. Giesen, and J. Röttger (2023): "Make-Up Strategies with Incomplete Markets and Bounded Rationality," Deutsche Bundesbank Discussion Paper No 01/2023.

Eggertsson, G. and M. Woodford (2003): "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 34, 139–211.

European Central Bank (2016): "A Guide to the Eurosystem/ECB Staff Macroeconomic Projection Exercises," July 2016.

——— (2023): "Eurosystem Staff Macroeconomic Projections for the Euro Area," June 2023.

Galí, J. (2015): *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton, NJ: Princeton University Press.

Gerke, R., S. Giesen, M. Lozej, and J. Röttger (2024): "On Household Labour Supply in Sticky-Wage HANK Models," Deutsche Bundesbank Discussion Paper No 01/2024.

Gerke, R., D. Kienzler, and A. Scheer (2022): "On the Macroeconomic Effects of Reinvestments in Asset Purchase Programmes," Bundesbank Discussion Paper No 47/2022.

Harrison, R. (2017): "Optimal Quantitative Easing," Bank of England Staff Working Paper 678.

Harrison, R. and M. Waldron (2021): "Optimal Policy with Occasionally Binding Constraints: Piecewise Linear Solution Methods," Bank of England Staff Working Paper 911.

Hebden, J. and F. Winkler (2021): "Impulse-Based Computation of Policy Counterfactuals," Finance and Economics Discussion Series 2021-042.

Hills, T., T. Nakata, and T. Sunakawa (2021): "A Promised Value Approach to Optimal Monetary Policy," *Oxford Bulletin of Economics and Statistics*, 83, 176–198.

Holden, T. (2016): "Computation of Solutions to Dynamic Models with Occasionally Binding Constraints," Mimeo.

Kaplan, G. and G. Violante (2018): "Microeconomic Heterogeneity and Macroeconomic Shocks," *Journal of Economic Perspectives*, 32, 167–94.

Laseen, S. and L. Svensson (2011): "Anticipated Alternative Policy Rate Paths in Policy Simulations," *International Journal of Central Banking*, 7, 1–35.

Lu, Y., R. King, and E. Pasten (2016): "Optimal Reputation Building in the New Keynesian Model," *Journal of Monetary Economics*, 84, 233–249.

Marcet, A. and R. Marimon (2019): "Recursive Contracts," *Econometrica*, 87, 1589–1631.

McKay, A. and C. Wolf (2023a): "Optimal Policy Rules in HANK," Mimeo.

——— (2023b): "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?" *Econometrica*, 91, 1695–1725.

Reiter, M. (2009): "Solving Heterogeneous-Agent Models by Projection and Perturbation," *Journal of Economic Dynamics and Control*, 33, 649–665.

Svensson, L. and R. Tetlow (2005): "Optimal Policy Projections," *International Journal of Central Banking*, 1, 177–207.

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

A Appendix: Optimal policy projection with two instruments

This section presents details about the optimal policy problem for the modified textbook model with two policy instruments from Section 4.3. In this case, policy shocks involve shocks to the interest rate, now denoted as $\varepsilon^i_{t+k|t}$, and shocks to asset holdings, $\varepsilon^q_{t+k|t}$, for $0 \leq k \leq T$. To reflect this, define the overall policy shock vector as

$$\varepsilon_t \equiv \begin{pmatrix} \varepsilon^i_t \\ \varepsilon^q_t \end{pmatrix}, \tag{23}$$

with $\varepsilon^j_t \equiv \left(\varepsilon^j_{t|t}, \varepsilon^j_{t+1|t}, \dots, \varepsilon^j_{t+T|t}\right)'$, $j \in \{i, q\}$. The policy shock vector ε_t thus now stacks two instrument-specific shock vectors below each other, doubling the size of ε_t to $2(T+1)$ elements compared to Section 3.

The use of two instruments also requires two instrument-specific IRF vectors. Let the vector $d^{x,k}_j \equiv \left(d^{x,k}_{0,j}, d^{x,k}_{1,j}, \dots, d^{x,k}_{T,j}\right)'$ contain the impulse response coefficients for variable $x \in \{\pi, i, q, y\}$ and the instrument-specific policy shock $\varepsilon^j_{t+k|t}$, $j \in \{i, q\}$. For the shock vector ε_t, the value of variable x in period $t+s$ can then be written as

$$x_{t+s} = \sum\nolimits_{j \in \{i,q\}} \sum\nolimits_{0 \leq k \leq T} d^{x,k}_{s,j} \varepsilon^j_{t+k|t}, \tag{24}$$

or as

$$x_{t+s} = \left(d^{x,0}_{s,i}, d^{x,1}_{s,i}, \dots, d^{x,T}_{s,i}, d^{x,0}_{s,q}, d^{x,1}_{s,q}, \dots, d^{x,T}_{s,q}\right) \begin{pmatrix} \varepsilon^i_{t|t} \\ \varepsilon^i_{t+1|t} \\ \vdots \\ \varepsilon^i_{t+T|t} \\ \varepsilon^q_{t|t} \\ \varepsilon^q_{t+1|t} \\ \vdots \\ \varepsilon^q_{t+T|t} \end{pmatrix}, \tag{25}$$

in matrix notation.

To write the entire path for variable x, given by $X_t = (x_t, x_{t+1}, x_{t+2}, ..., x_{t+T})'$, in a convenient manner, first define the instrument-specific coefficient matrix

$$D^x_j \equiv \left[d^{x,0}_j, d^{x,1}_j, ..., d^{x,T}_j\right] = \begin{bmatrix} d^{x,0}_{0,j} & d^{x,1}_{0,j} & \cdots & d^{x,T}_{0,j} \\ d^{x,0}_{1,j} & d^{x,1}_{1,j} & \cdots & d^{x,T}_{1,j} \\ \vdots & \vdots & \ddots & \vdots \\ d^{x,0}_{T,j} & d^{x,1}_{T,j} & \cdots & d^{x,T}_{T,j} \end{bmatrix}. \tag{26}$$

Now, define the $(T+1) \times 2(T+1)$-dimensional overall coefficient matrix D^x by stacking the instrument-specific coefficient matrices next to each other, i.e. $D^x \equiv \left[D_i^x, D_q^x\right]$. As in the one-instrument case, the time path X_t can then be computed as $X_t = D^x \varepsilon_t$. However, the individual elements of X_t are now given by equation (24).

The next step is to define the outcome vector

$$Z_t \equiv \begin{bmatrix} \Pi_t \\ Y_t \\ I_t - I_{t-1} \\ Q_t \\ Q_t - Q_{t-1} \end{bmatrix}, \tag{27}$$

the baseline vector

$$B_t \equiv \begin{bmatrix} B_t^{\pi} \\ B_t^{y} \\ B_t^{\Delta_i} \\ B_t^{q} \\ B_t^{\Delta_q} \end{bmatrix}, \tag{28}$$

and the coefficient matrix

$$D \equiv \begin{bmatrix} D^{\pi} \\ D^{y} \\ D^{\Delta_i} \\ D^{q} \\ D^{\Delta_q} \end{bmatrix}. \tag{29}$$

Since the relationship $Z_t = B_t + D\varepsilon_t$ from Section 3 still applies, so does the formulation of the policy problem, $\min_{\varepsilon_t} \left\{ \frac{1}{2} Z_t{}' W Z_t \right\}$ s.t. $Z_t = B_t + D\varepsilon_t$.

However, the matrix W is now given as

$$W = \begin{bmatrix} W_\beta & 0 & 0 & 0 & 0 \\ 0 & \lambda W_\beta & 0 & 0 & 0 \\ 0 & 0 & w_i W_\beta & 0 & 0 \\ 0 & 0 & 0 & w_q W_\beta & 0 \\ 0 & 0 & 0 & 0 & w_{\Delta_q} W_\beta \end{bmatrix}. \tag{30}$$

Performing the same steps as in Section 3, the solution to the policy problem is again given by

$$\varepsilon_t^* = -\left(D^{'}WD\right)^{-1}(B_t^{'}WD)^{'},$$

for the unconstrained case, with ε_t^* now including shocks for two policy instruments.

For the constrained case, the solution can again be computed numerically by solving a linear-quadratic programming problem with inequality constraints. This case is relevant for the illustrative example from Section 4.3 because asset purchases only matter when there is a binding ELB constraint (see Figure 5). Furthermore, a non-negativity constraint for asset holdings, $q_t \geq 0$, is imposed in the example, which can be written as $-D^q\varepsilon_t \leq B_t^q$ in matrix notation.