DNB Working Paper

No 818/November 2024

Scenario discovery to address deep uncertainty in monetary policy

Chamon Wieles, Jan Kwakkel, Willem L. Auping, Jan Willem van den End

DeNederlandscheBank
EUROSYSTEEM

Scenario discovery to address deep uncertainty in monetary policy

Chamon Wieles, Jan Kwakkel, Willem L. Auping and Jan Willem van den End*

* Views expressed are those of the authors and do not necessarily reflect official positions of De Nederlandsche Bank.

Working Paper No. 818

November 2024

De Nederlandsche Bank NV
P.O. Box 98
1000 AB AMSTERDAM
The Netherlands

Scenario discovery to address deep uncertainty in monetary policy

Chamon Wieles[a], Jan Kwakkel[a], Willem L. Auping[a], Jan Willem van den End[b,c*]

[a] Delft University of Technology, the Netherlands
[b] De Nederlandsche Bank, the Netherlands
[c] Vrije Universiteit Amsterdam, the Netherlands

November 2024

Abstract

We analyse shock and parameter uncertainty in a Dynamic Stochastic General Equilibrium (DSGE) model by exploratory modelling and analysis (EMA). This method evaluates in a novel way the performance of monetary policy under deep uncertainty about the shock and model parameters. Scenarios are designed based on the outcomes of interest for the policymaker. We assess the performance of different policies on their objectives in the scenarios. This maps out the policy trade-offs and supports the central bank in making robust policy decisions. We find that in response to a negative supply shock, policies with low interest rate smoothing and a strong response to inflation most obviously contribute to price stability under deep uncertainty.

JEL classification: E52, E58, G12
Keywords: Monetary policy, Scenarios, Exploratory modelling, Deep uncertainty.

---

* Corresponding author. Email address: w.a.van.den.end@dnb.nl
The authors wish to thank Maurice Bun, Servaas Storm and participants of the DNB research seminar (January 2024) for useful input. Views expressed are those of the authors and do not necessarily reflect official positions of De Nederlandsche Bank or the Eurosystem.

1. **Introduction**

We contribute to the growing literature on monetary policy under deep uncertainty by applying exploratory modelling and analysis (EMA) to a dynamic stochastic general equilibrium (DSGE) model. DSGE models are a common tool for central banks to explain economic dynamics and forecast them. EMA is commonly used for policy problems in environmental sciences when deep uncertainty plays a role. It differs from existing methods that deal with uncertainty in structural macroeconomic models, like DSGE models, in several ways. Exploratory modelling assumes that existing information is insufficient to specify a single model that accurately describes system behaviour. Instead, it generates a set or ensemble of models or model-based inferences that are plausible or interesting in the context of the analysis (Bankes, 1993; Bankes et al, 2013). Selecting a particular model out of an ensemble of plausible ones requires making suppositions about factors that are uncertain or unknown. EMA considers continuous ranges for the model parameters and shock values, identifies combinations of them in scenarios that are of interest to the policymaker, and applies a distinct way of evaluating the performance and robustness of different policy strategies. Multiple extensions of this EMA approach have been developed in recent years (see Bartholomew and Kwakkel, 2020). To the best of our knowledge, EMA has not yet applied to DSGE models before.

With regard to uncertainties in DSGE models, the literature usually considers uncertainty in only one or a few features of the model. Xiao et al. (2018), for example, incorporate predefined uncertainties, such as technology uncertainty and fiscal policy uncertainty, as a set of discrete values for each uncertainty. To deal with parameter uncertainty, a Bayesian econometric approach (which can be applied in a limited information setting) and Markov-Chain Monte Carlo analysis have been applied to DSGE models, for instance by Boivin and Giannoni (2006) and Christiano et al. (2010). Some studies apply a sensitivity analysis to DSGE models, such as Čapek et al. (2023) and Ratto (2008).

Next to parametric uncertainty, the central bank has to deal with shock uncertainty. Giannoni (2007) simulates the effects of a variety of exogenous shocks, instead of a single exogenous innovation, with a New Keynesian DSGE model. He considers the uncertainty of the relative importance of each shock, as well as uncertainty about the persistence in the shock processes. In a related study, Grassi et al. (2016) select shocks in a DSGE model to avoid the estimation bias arising from imposing non-observable shocks. These are generated by non-existing exogenous processes. Carceller and Van den End (2023) assess the robustness of monetary policy assuming deep uncertainty about the relative contribution of demand and supply shocks to inflation.

Also under deep uncertainty, policymakers have to design policies to achieve the desired policy goal. Kwakkel et al. (2016) state that deep uncertainty means that the various decision makers do not know or cannot agree on the system and its boundaries, the outcomes of interest and their relative importance. Moreover, the prior probability distribution for uncertain inputs to the system is unknown (Lempert et al., 2003; Walker et al., 2013), Moreover, under deep uncertainty, decisions are made over time in dynamic interaction with the system and cannot be assessed independently (Haasnoot et al., 2013; Hallegatte et al., 2012). This calls for robust policies that achieve desired outcomes also under deep uncertainty. Maier et al. (2016) states that robustness can be thought of as a measure of the insensitivity of the performance of a given policy strategy to future uncertain states of the world.

A well-known strategy for dealing with deep uncertainty in economic models is robust control. Dennis et al. (2019) and Olalla and Gómez (2011) apply robust control to a New Keynesian model to study the effect of model uncertainty in monetary policy. Another way to assess this uncertainty is info-gap theory, which aims at satisficing instead of optimising the outcome (Ben-Haim et al., 2018). Several studies investigate the design of robust policies based on DSGE models, such as Górajski et al. (2023). Robustness of policies is quantified on the basis of their average or maximum welfare loss for different assumed states of the world. Based on observed data they estimate distributions within which the parametric uncertainties can be located.

We apply EMA methods to address parametric and shock uncertainty in a DSGE model. The basic idea of exploratory modelling is to run a model that describes the policy problem a large number of times with different input values (i.e. values of model parameters and/or shocks that represent the uncertainties). The sampled input values span the range of the uncertainties. Thereby the model is used as a device to capture relevant uncertainties by enumerating a range of possible assumptions and systematically exploring the implications of them via large numbers of computational experiments (Moallemi et al., 2020). Subsequently, the set of model outcomes is analysed by various statistical techniques. This helps to understand the effects of specific uncertainties on the outcomes (following Kwakkel and Jaxa-Rozen, 2016).

As it is the first attempt to apply EMA to DSGE models, our study is exploratory by nature. We investigate what can be learned by applying EMA methods to DSGE models and how the central bank can implement EMA in its interest rate policy. We illustrate this by applying EMA to specific policy problems for which the central bank uses a DSGE model. In our framework the central bank steers the policy levers in order to robustly meet the policy objectives. This is done by a many-objective optimisation approach, which searches through the range of values of the policy levers (i.e., the parameters in the Taylor rule) and by examining potential trade-offs between different policy objectives.

An important feature of robust decision making is that the policy performance is evaluated in different states of the world (Homaei and Hamdy, 2020; Lempert et al., 2006). EMA provides for this by directed search. This method investigates the policies and sets of actions that the policymaker can take to achieve certain objectives under deep uncertainty (Kwakkel, 2017).This is useful because the policy that has the most desirable performance on a specific objective in one scenario is not necessarily the best policy. It may be possible that this policy underperforms in many of the other scenarios and hence is not robust optimal. We assess the performance of different policies by several metrics for robustness to deep uncertainty. The uncertain future states of the world are driven by scenarios that capture the uncertainty about the model parameters and shock values. In our case, we assume that the shock values relate to a negative supply shock. We investigate how the central bank should set the interest rate under those conditions to achieve its objectives. So, we find the robust optimal values for the policy levers: the smoothing parameter and inflation response parameter in the Taylor rule.

Based on EMA open exploration, we find that states with low price flexibility and high shock persistence determine the outcomes that are of interest to the central bank. Such states are associated with relatively large deviations of the output gap and inflation from steady state and thus with a failure to achieve the policy target. Based on EMA directed search, we find the values for the policy levers (i.e., the input parameters that the central bank can influence) that effectively influence the economic dynamics following a negative supply shock. Under deep uncertainty, a rather strong response to inflation contributes to the inflation objectives, while low interest rate smoothing contributes to both the inflation and output gap objectives. Both policies are also robust to uncertain states of the world.

The remainder of this paper is structured as follows. In section 2 we explain the DSGE models and the EMA framework, as well as the policy objectives of the central bank. In section 3 we apply the EMA open exploration methods and in section 4 the directed search methods to the monetary policy problem in the context of the DSGE model. In section 5 we discuss our approach in a broader context. Section 6 concludes.

2. Methodology

2.1 DSGE model

DSGE models are a common tool for central banks to explain economic dynamics and forecast them. The models are also used for storytelling and policy simulations (Del Negro and Schorfheide, 2013). The basic New Keynesian DSGE model consists of three equations: one for demand, supply, and the policy rule (Galí, 2015). Many extensions of DSGE models have been developed, featuring different economic structures and additional variables such as investment, employment, and real wages (Smets and Wouters, 2003). For specifying a DSGE model, the policymaker has to make assumptions about the economic state and the behaviour of economic agents and translate those into model parameters. It is common for policymakers to take the economic model with certainty, considering fixed values of the parameters and the median or mean of a posterior distribution (Górajski et al., 2023; Taylor and Williams, 2010).

The basic three equations DSGE model describes the behaviour of households, firms, and the central bank in setting the interest rate. The model is calibrated around a steady state with 2% inflation, and, after deviating from this steady state following a shock, it will return to this equilibrium. We assume that the main goals of the central bank are price stability and balanced economic conditions. These are operationalized by policy objectives, like the objective to minimise the deviation of inflation from the steady-state level. The central bank aims at these objectives by setting the interest rate, which follows a conventional Taylor rule. The parameters of the Taylor rule are the policy levers in the model.

As benchmark model we take the simple New Keynesian DSGE model described by Galí (2015). It comprises three equations that model the output gap, inflation, and interest rate in a dynamic fashion.

$$\tilde{y}_t = E(\tilde{y}_{t+1}) - \frac{1}{\sigma}(i_t - E(\pi_{t+1}) - r_t^n) - (\rho_d - 1)d \tag{1}$$

$$\pi_t = \beta\, E\,(\pi_{t+1}) + \kappa\,\gamma\,\tilde{y}_t + \mu \tag{2}$$

$$i_t = \rho_r i_{t-1} + (1 - \rho_r)(\varphi_\pi \pi_t + \varphi_y\,\tilde{y}_t) + v_t \tag{3}$$

Where $\tilde{y}_t$ and π_t represent output and inflation as deviations from their steady state values and i_t the nominal policy interest rate. Furthermore, d represents a preference (demand) shock and μ a mark-up (supply) shock, which both follow an AR(1) process with persistence parameters ρ_u and ρ_d. *E(·)* is the expectations operator, assuming agents have rational expectations. Parameter σ is the inverse intertemporal elasticity of substitution, γ a time preference parameter, and r_t^n the equilibrium interest rate. *K* is the degree of price stickiness, with $\kappa = (1 - \omega)(1 - \omega)\beta/\omega$ and ω the Calvo parameter. The latter reflects the share of firms that does not adjust their price (prices are sticky), implying that a low value of ω is associated with more flexible prices.

The central bank reacts with full and credible commitment according to a standard Taylor rule, with smoothing parameter ρ_r and with φ_π and φ_y being the policy response parameters for inflation and the output gap (i.e., the weight the central bank puts on inflation and the output gap respectively in the reaction function). These parameters are set by the central bank and represent the policy levers in the model.

We assume that the economy is hit by a negative supply shock μ (e.g. a cost-push shock), which leads to higher costs of production. This decreases the aggregate supply of goods. If aggregate demand remains constant while supply falls, the price level rises, and the economy experiences stagflation (falling output and rising prices). This is a policy problem, since the combination of falling output and rising prices confronts the central bank with a potential trade-off between the objectives of price stability and balanced economic conditions. Applying EMA to this problem reveals which uncertainties are associated with this trade-off and which stance of the policy levers is most effective to address it.

2.2. EMA framework

We explore the use of EMA for DSGE models and identify steps that a central bank has to take to incorporate EMA into the policy practice. Two methods can be distinguished in EMA: open exploration and directed search. In open exploration, the model is run for a large number of different combinations of uncertainties (the scenarios) and, optionally, policies. Directed search investigates the policies and sets of actions that the policymaker can take to achieve certain objectives (Kwakkel, 2017).

2.2.1. Open exploration

EMA dissects the model and the policy problem into different components: uncertainties, scenarios, levers, policies, outcomes, objectives, and robustness metrics. The uncertainties are the uncertain model parameters and shock values (which we both assign ranges of values to) and their distribution. We take the parametric uncertainties and the shock uncertainties of the DSGE model into account. The different uncertainties together form the uncertainty space, as shown in Figure 1. A scenario is a point in this uncertainty space: it comprises a combination of values for the different uncertainties.

The (policy) levers are parameters or components of the model that the policymaker can influence. We also assign a range and a distribution to these parameters. The levers form a decision space (comparable to the uncertainty space), and a policy is represented by a point in this decision space. In our case, the central bank controls the interest rate by the policy levers in the Taylor rule and the different policies thus refer to the interest rate strategies of the central bank.

Running the model for one scenario and one policy is referred to as an experiment, and for each experiment, a set of model outcomes is obtained. When designing a policy, it is important to identify the interests of the central bank. These are translated into objectives, i.e., outcomes of the model the policymaker wants to minimise or maximise, and robustness metrics; the criteria on which we assess the robustness. Like the policy levers and uncertainties, the outcomes and objectives form their own spaces: the outcome space and the objective space.

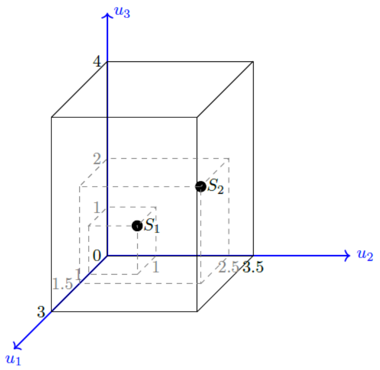

Figure 1. Uncertainty space formed by three theoretical continuous uncertainties u_1 (range (0,3)), u_2 (range (0,3.5)), and u_3 (range (0,4)). Two possible scenarios are shown: S_1(u_1= u_2= u_3= 1), and S_2 (u_1=1.5; u_2=2.5; u_3=2).

In theory, the whole uncertainty and/or decision space can be explored. This can be conducted by picking scenarios and policies at random, for instance, by Latin Hypercube sampling (McKay et al., 1979).[1] This method first divides the ranges of each parameter into separate segments, thus dividing the uncertainty and decision space into separate spaces, or cells, all having the same probability of being chosen (McKay et al., 2000). The latter implies that a uniform distribution is imposed on the parameter and shock values, which fits with deep uncertainty in the sense that simulated values are equally likely in a uniform distribution. The sampler then picks a random value in a randomly chosen cell.[2]

[1] The number of experiments (or samples) in Latin Hypercube sampling varies based on the complexity of the problem and the dimensionality of the parameter space. The number should balance computational efficiency with the need for accurate statistical estimation.
[2] Latin Hypercube sampling can also be used with non-uniform distributions. Each grid cell will then keep the same probability, but the grid cell size/volume will differ. We assume a uniform distribution

By open exploration, we conduct experiments that span the range of the uncertainties in the uncertainty space (Kwakkel and Pruyt, 2015). For this, we assign a range and a uniform distribution to the uncertain parameters, based on their expected true value. Then, we analyse the outcome space to see how the outcomes behave for different combinations of parameter and policy lever values. To understand the effect of uncertainties on the outcome space, scenario discovery is applied (Bryant and Lempert, 2010; Kwakkel and Jaxa-Rozen, 2016). Scenario discovery defines a subspace of the uncertainty space that illuminates the vulnerabilities of a proposed policy (Dalal et al., 2013). In other words, with scenario discovery, common input space properties are identified across ensembles of exploratory model runs (Steinmann et al., 2020).

From the outcome space, a specific feature is selected that the policymaker wants to explore because this feature is of interest to him. Such a feature can be an outcome that exceeds a certain critical value and thereby conflicts with a policy objective. Experiments that share this feature can be compared. To illustrate this, in Figure 2, the dots in the uncertainty space (right) represent scenarios (combinations of uncertain parameters in the absence of policies), and the dots in the outcome space have their corresponding outcomes. The area of the outcome space of interest for the policymaker is marked by the red dashed rectangle. In this area, including the larger yellow dots, there are 11 experiments (the larger blue dots in the uncertainty space.) with o1 being greater than 3. So the policymaker can identify regions in the uncertainty space that are of interest in terms of the outcome (Bryant and Lempert, 2010; Kwakkel et al., 2013). This means that a certain feature of the outcomes forms the basis of scenario discovery.

since we first want to know what could happen and what the impact of is before we deal with the likelihood of this.

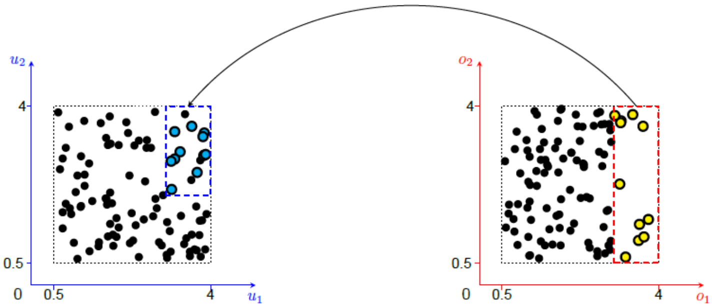

Figure 2. Uncertainty space and outcome space of a model with 2 uncertainties, 2 outcomes and no policy levers. For illustration, 100 scenarios (experiments) (dots in the uncertainty space) and their corresponding results (dots) in the outcome space are visualized. Our region of interest in the outcome space (red, dashed rectangle) consists of the results (yellow-filled larger dots) of 11 experiments (blue-filled dots in uncertainty space). These experiments lie in the blue, dashed area, but not all experiments in this area result in $o_1 > 3$.

A method to find scenarios (i.e., combinations of the uncertainties) with common characteristics of the outcome variables is behaviour-based scenario discovery. This method identifies clusters of the inflation and output gap that have common characteristics. The clusters are formed by time series clustering methods that use similarity metrics, as described by Steinmann et al. (2020), which clusters time series based on their behaviour across time.[3] The patient rule induction method (PRIM) then identifies the ranges of the parameter and shock values that distinguish the outcome dynamics in each cluster (see Friedman and Fisher, 1999 and Kwakkel et al., 2016). PRIM identifies the range of parameter values of the uncertainties that determine the outcome clusters. It does so by searching for a set of subspaces (boxes) within the uncertainty space with an outcome value that significantly differs from its average outcome value in the domain (Kwakkel and Jaxa-Rozen, 2016). The subspaces of the uncertainty space provide a range of the parameter values that form a box of an outcome cluster that is significantly different from parameter ranges of other outcome clusters.

So, scenario discovery by time series clustering and PRIM identify the combination of parameters and shock values (i.e., the scenarios) that are of interest to the policymaker. It reveals clusters of the outcome variables that may conflict with the

[3] Clustering on the full time series rather than individual time points ensures that each cluster represents coherent temporal patterns, reflecting the time-dependent nature of the data and retaining the meaningful structure of how the series changes over time.

policy objectives. Thereby, these methods provide useful information for designing policies to prevent such outcomes.

2.2.2 Directed search

Directed search investigates the uncertainty space to find policy strategies that satisfy the objectives (Kwakkel and Pruyt, 2015). The same scenario can, under different policies, result in different outcomes. In Figure 3, the uncertainty space of two uncertainties (u1 and u2) results in three different outcome spaces (o1 and o2) due to three different policies (p1, p2, and p3). For instance, policy 1 creates an outcome space that spans over a relatively large range of outcome 1 and has relatively high values for outcome 2. Note that, whereas we know the distributions of the uncertainties (we assume a uniform distribution), we do not know the distributions of the outcomes in the different outcome spaces. This distribution of outcomes can differ per policy.

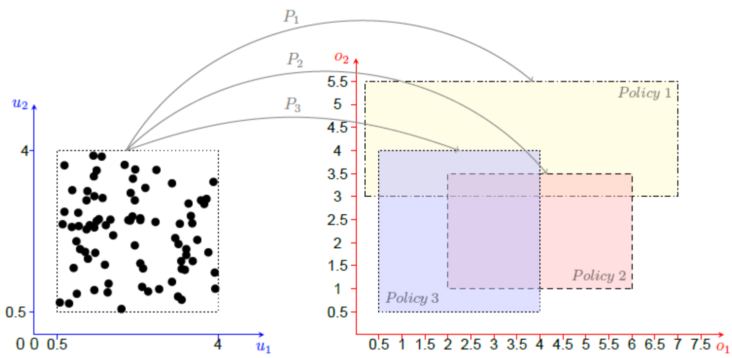

Figure 3. Three outcome spaces associated with the same uncertainty space under different policies (P_1, P_2, and P_3).

An optimal policy can be designed by searching for combinations of policy levers that optimise the objectives. The outcomes of these Pareto optimal policies in the objective space is the Pareto front. Each point on this front is an outcome from which the policymaker cannot deviate to improve the performance of one outcome without reducing the performance of another outcome. The exact Pareto front is unknown but can be approximated by optimisation methods. We refer to these approximations as Pareto optimal solutions. Finding these optimal policies is a multi-objective optimisation problem (Statnikov and Matusov, 2012).

It is important to assess policies on their robustness in different states of the world. To evaluate the robustness, we use the taxonomy of robustness frameworks of Herman et al. (2015). It first generates the optimal combinations of the policy levers by many-objective search. Then, it generates states of the world and conducts a vulnerability analysis. We base the comparison of the different policies in terms of their robustness on three robustness metrics: regret, satisficing and descriptive metrics (Kwakkel et al., 2016).

2.3 Application to DSGE model

We apply both open exploration and directed search to a monetary policy problem in the context of the DSGE model specified in section 2.1. EMA allows us to assess the DSGE model under deep uncertainty in different ways compared to existing approaches. For instance, EMA does not use a welfare-loss function, which is a common approach in DSGE models to quantify losses in terms of deviations of inflation and the output gap from target (as the canonical loss function in Galí, 2015). Instead, EMA considers specific policy objectives that are of interest to the policymaker and compares the policy performance and robustness for each objective. Robustness metrics reflect the sensitivity of the policy performance to future uncertain states of the world (McPhail et al., 2018). In addition, exploratory modelling analyses the (combined) input parameters, i.e., the scenarios, based on their outcome. This allows for selecting scenarios based on outcomes that are policy relevant. This differs from common ways to select scenarios for simulating policy problems in macroeconomic models, which usually start from choosing input parameters and then analyse the outcomes.

We assume that policy objectives of interest to the central bank are price stability and balanced economic conditions. A central bank with a price stability objective will aim at inflation reverting back to steady state following the shock, in the medium-term (assumed to be the policy horizon of three years). In addition, we assume that the central bank wants to avoid deflation (Krugman and Wells, 2021). Moreover, in our approach, the central bank tries to keep inflation and output close to their steady state levels over the entire policy horizon and to limit deviations across time. This should contribute to keeping inflation expectations anchored. Lastly, for the benefit of trust in the central bank, it aims at a monetary policy response that achieves the policy target in different states of the world. The latter is typically assessed by robustness metrics (to be discussed in section 4.2).

The policy objectives are summarized as follows:

1. Annual inflation rate around 2% at the end of the policy horizon (12 quarters);
2. Limit deviations of inflation from steady state level;
3. Limit period-to-period changes in inflation;
4. Limit output gap deviations;
5. Limit period-to-period changes in the output gap;
6. Avoid deflation (minimise periods with deflation);
7. Similar performance of the policy response, irrespective of state of the world (robustness).

The policy objectives are analysed by directed search in section 4.1, based on the model outcomes the central bank wants to minimise or maximise. Finding optimal policies is pursued by optimising the objectives through a many-objective optimisation approach. The robustness of these optimal policies is assessed by robustness metrics in section 4.2.

The stylized framework in Figure 4 connects the DSGE model to EMA. It distinguishes the exogenous uncertainties (X), policy levers (L), the relations I, and the measures (M) in an XLRM framework.

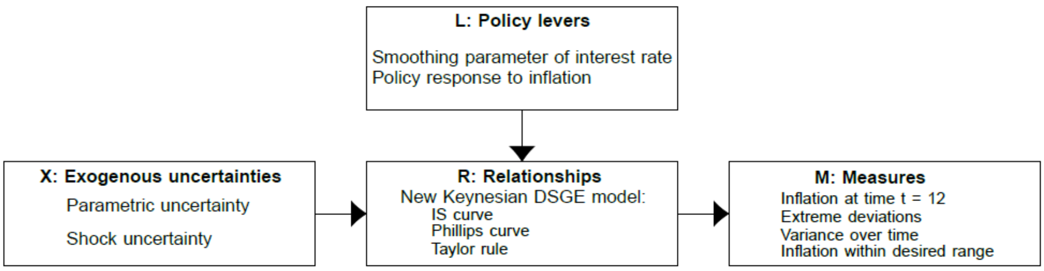

Figure 4. XLRM framework of the problem. X represents the exogenous uncertainties, L the policy levers, R the internal relations, and M the measures (or objectives) of the model.

All parameters, excluding policy levers, are considered uncertainties. We assume that according to the central bank's preferences, only the interest rate smoothing parameter and the policy response to inflation are used as policy levers.

3. Results open exploration

This section presents the results of the open exploration of the DSGE model. We analyse the effects of a negative supply shock, as an example of how EMA can be applied in monetary policy analysis. A negative supply shock is an interesting case to analyse, as this shock implies a trade-off for monetary policy between the inflation and output objectives. We take the uncertainties into account by simulating 15,000 experiments of the parameter and shock values. Latin Hypercube is used for the randomization of the values. This method requires a large number of experiments to cover the whole uncertainty space. We balance covering the whole uncertainty space with the computational time by conducting 15,000 experiments by Latin Hypercube sampling.

The uncertainties and their ranges included in the open exploration are shown in Table 1. The simulation of the values is based on uniform distributions, with mean values based on values commonly used in (calibrated) DSGE models.

Table 1. Model uncertainties for open exploration.

Parameter	Distribution	Lower limit	Upper limit
Discount factor	Uniform	0.95	1
Inverse of marginal elasticity of substitution	Uniform	1	2.5
Calvo parameter	Uniform	0.6	0.95
Autoregressive term of demand shock	Uniform	0.5	0.95
Autoregressive term of supply shock	Uniform	0.5	0.95
Inverse Frisch elasticity of labor supply	Uniform	0.6	1.5
Targeted share of impatient agents	Uniform	0.2	0.5
Policy response to inflation	Uniform	1	2.5
Smoothing parameter policy interest rate	Uniform	0.5	1

Taking into account the uncertainties, the paths of the output gap $\tilde{y}_t$ and inflation π_t shows a large dispersion following a negative supply shock (Figure 1). Whereas most experiments result in paths of the output gap that approach the steady state from below (associated with a negative deviation in line with economic intuition), the initial responses of the output gap to the shock differ a lot. There are also many scenarios in which the output gap shows a positive response initially. This is in line with the common finding in the literature that output gap estimates are subject to a wide range of uncertainty owing to data revisions and the difficulty of distinguishing between cycle and trend in real-time (Grigoli et al., 2015). In contrast to the results for the output gap, the responses of inflation to the negative supply shock show almost no negative deviations, while the positive deviations have a much wider range. The positive response of inflation to the negative supply shock is in line with economic intuition.

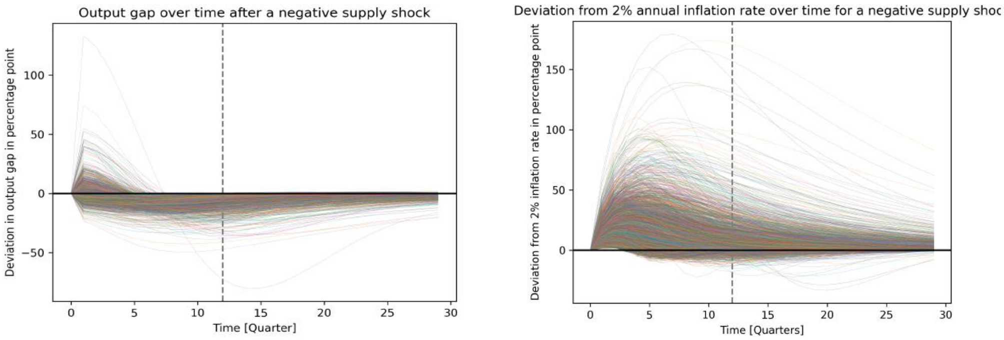

Figure 1. Output gap (left panel) and inflation (right panel) after negative supply shock, including uncertainty.

Behaviour based scenario discovery identifies clusters of inflation and the output gap that have common characteristics. The clusters are formed by time series clustering methods, as explained in the previous section. A certain minimal set of scenarios is needed in each cluster to provide reliable information, and hence, we only take into account clusters that contain at least 500 experiments to ensure that the results are significant. This results in two clusters for both the output gap and inflation. The clustering of the output gap reveals that the outcomes in cluster 1 are characterised by relatively small deviations from steady state, while the outcomes in cluster 2 display somewhat larger deviations (Figure 2). Although there is not much difference between the behaviour of the output gap in both clusters, the central bank would be most interested in cluster 2, taking into account the objective of limiting output gap deviations.

In the case of inflation, the behaviour in the clusters is more distinct (Figure 2). Inflation in cluster 2 displays more positive deviations from steady state than in cluster 1. Inflation in cluster 2 is also more persistent, as it takes more time for inflation to revert to the steady state. Hence, the central bank would be most interested in cluster 2, taking into account the objectives of limiting deviations of inflation from steady state and limiting period-to-period changes in inflation.

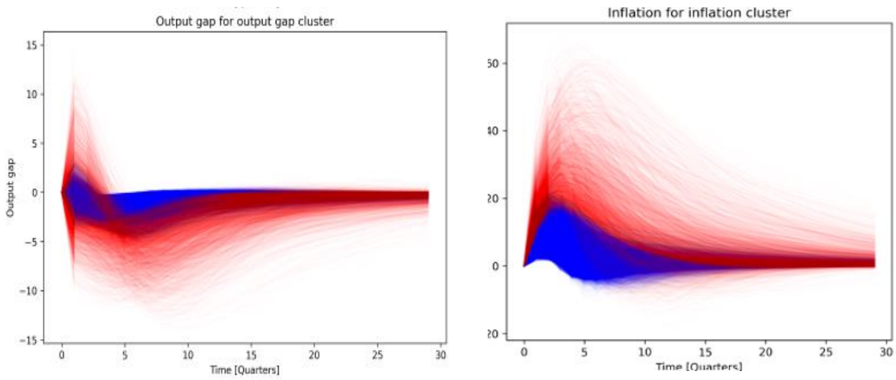

Figure 2. Two clusters for the output gap (left panel) and inflation (right panel) following a negative supply shock. Cluster 1 is coloured blue and cluster 2 red.

We apply PRIM to identify the parameter ranges that distinguish the clusters of the output gap and inflation dynamics. Each subspace, or box, in the uncertainty space, is uniquely formed by the parameters that best describe this subspace. The subspaces are bounded by a parameter range, of which the quasi-p-value indicates whether the restriction to describe the box (i.e., the number of uncertain parameters imposed) is statistically significant.

Applying PRIM to the two clusters of the output gap results in the parameter ranges shown in Table 2. With the exception of the Calvo parameter, the parameter ranges in clusters 1 and 2 are almost similar to the whole range of uncertainties. It implies that the parameter ranges in the different clusters largely overlap, which limits the use of this exercise for the output gap. From the Calvo parameter in cluster 1 (which is characterised by small output gap deviations), we infer that output gap deviations are likely small if prices are more flexible (which is the case when the Calvo parameter is low).

Table 2. Uncertainty ranges for the output gap clusters (the quasi-p-value indicates whether the restriction to describe the box is statistically significant).

Parameter	Range	Cluster 1	Cluster 2
Calvo parameter	[0.6 - 0.95]	[0.60 - 0.77***]	[0.84*** - 0.95]
Smoothing parameter	[0.5 - 1]	[0.50 - 0.95***]	[0.68*** - 0.93]
Autoregressive factor of supply shock	[0.5 - 0.95]	[0.5 - 0.89***]	
inverse Frisch elasticity of labor supply	[0.6 - 1.5]	[0.65 - 1.5*]	
Coverage		0.454	0.627
Density		1.00	0.655

* indicates a p-value smaller than 0.05
** indicates a p-value smaller than 0.025
*** indicates a p-value smaller than 0.01

Applying PRIM to the inflation clusters results in the parameter ranges shown in Table 3. Like for the output gap, several of the parameter ranges are largely similar to the full range of uncertainties. The most obvious distinctions between the two clusters show up in the Calvo parameter and the persistence of the supply shock. In cluster 1, prices are more flexible (lower Calvo parameter) than in cluster 2. Figure 2 shows that in cluster 1, inflation more likely has smaller deviations from steady state than in cluster 2. Combining this with the PRIM results implies that larger price flexibility limits upward risks to inflation in case of a negative supply shock. Inflation risks are also contained if the shock persistence is lower. This conclusion is drawn from the inflation outcomes in cluster 1 (in which the deviations from steady state are contained) and the PRIM outcome in Table 3, which shows that the value of the autoregressive term in the supply shock equation is lower in cluster 1.

Table 3. Uncertainty ranges for the inflation clusters (the quasi-p-value indicates whether the restriction to describe the box is statistically significant).

Parameter	Range	Cluster 1	Cluster 2
Calvo parameter	[0.6 - 0.95]	[0.6 - 0.81***]	[0.88*** - 0.95]
Autoregressive factor of supply shock	[0.5 - 0.95]	[0.5 - 0.83***]	[0.78*** - 0.93*]
Policy reaction to interest	[1 - 2.5]	[1.2*** - 2.5]	
Smoothing parameter	[0.5 - 1]		[0.5 - 0.98]
Discount factor	[0.95 - 1]		[0.95 - 0.99***]
Coverage		0.454	0.299
Density		1	0.953

* indicates a p-value smaller than 0.05
** indicates a p-value smaller than 0.025
*** indicates a p-value smaller than 0.01

To conclude, the scenario discovery indicates that the outcomes of interest for the central bank are located in inflation and output gap clusters 2. In those clusters, the output gap and inflation deviate relatively more from steady state, which may conflict with the policy objectives. The PRIM exercise shows that cluster 2 outcomes are associated with relatively low price flexibility and (with regard to inflation) high shock persistence. This suggests that under deep uncertainty, the central bank should focus its interest rate policy on states of the world that comprise low price flexibility and high shock persistence to prevent outcomes that conflict with the policy objectives.

4. Results directed search

4.1 Optimal policies

This section discusses the results of the directed search, which is the search for Pareto optimal policies after a supply shock. The search is first conducted in the baseline scenario (assuming no uncertainty), for the policy objectives specified in section 2.3. We search over the range of the policy levers embedded in the Taylor rule: i.c. the response to inflation φ_π and the interest rate smoothing parameter ρ_r (these levers are no uncertainties but values chosen by the central bank). Then, we introduce deep uncertainty by examining how the policies perform in the 15,000 scenarios for the uncertainties in the parameter and shock values. Finally, we assess the robustness of chosen policy strategies.

The performance of the different policies under the baseline scenario assuming no uncertainty is visualised in Figure 3 (the coloured curves connecting the vertical lines represent different policies). The Figure shows the trade-off between optimising the inflation-related objectives and optimising the output gap-related objectives. This comes to the fore in the crossing of the curves between ’largest change inflation’ and ’y max’. It shows that the objective of minimising the period-to-period change in inflation conflicts with aiming at limiting output gap deviations. This reflects the typical trade-off for monetary policy when it faces a negative supply shock.

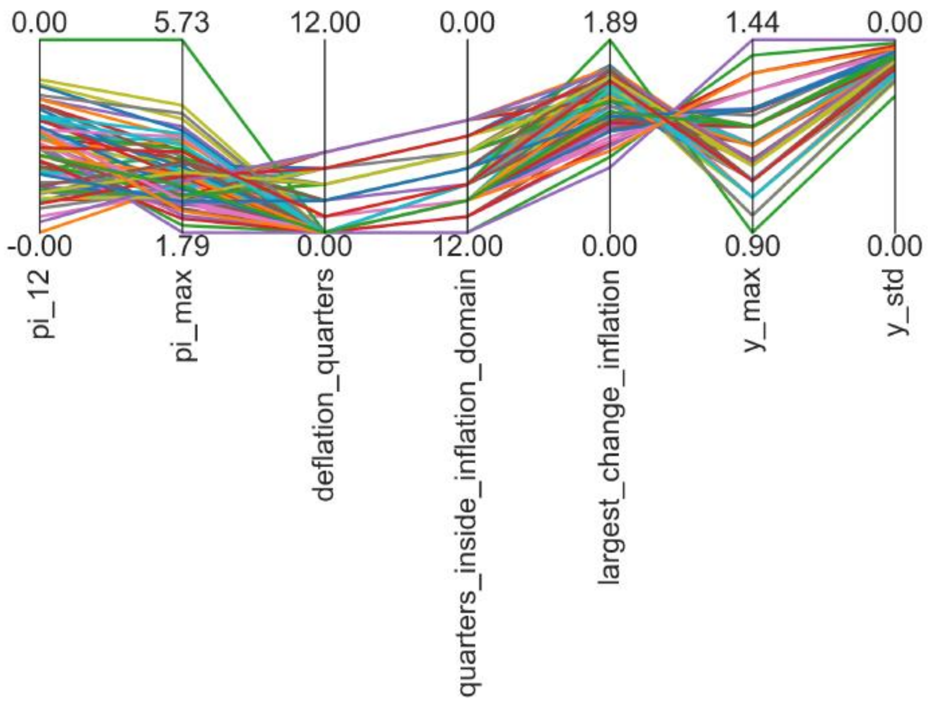

Figure 3. The vertical line of each policy objective represents the range of all outcomes for that objective, for all policies. The objecties are: inflation rate around 2% at end of policy horizon (pi_12), limit deviations of inflation from steady state level (pi_max), minimise periods with deflation (deflation_quarters), maximise quarters with inflation being between 2% and 4% (quarters_inside_inflation_domain), limit changes in inflation (largest_change_inflation), limit output gap deviations (y_max), limit changes in the output gap (y_std). The coloured curves connecting the vertical lines represent different policies, which cross the vertical lines at the value of the policy outcome for that objective. The crossing of the curves at the vertical lines indicate the performance of the optimal policies. For an objective that is minimised (maximised), the lower (upper) boundary of the outcome range is at the bottom (top) of the vertical line. The range of the vertical line is adjusted if there is a certain minimum or maximum value, such as the number of quarters in which deflation occurs and quarters in which inflation lies within a certain domain (both have a range of [0,12]).

For different combinations of policy levers (φ_π and ρ_r) the performance on an objective can be compared.

Figure 4 plots the performance on an inflation objective (i.e. limiting deviations of inflation from steady state, in the left panel) and an output gap objective (i.e. limiting deviations of the output gap, in the right panel) for different optimal policies. The larger and darker the dots, the better the policy performance. The trade-off clearly shows up, as the preferred policies for one objective are the least preferred policies for the other objective. In addition, the optimal policies form a triangle, implying that there are no Pareto optimal policies that combine a relatively high φ_π (ie a strong policy response to inflation which is desirable for the inflation objective) and low ρ_r (ie low interest rate smoothing which is desirable for the output gap objective).

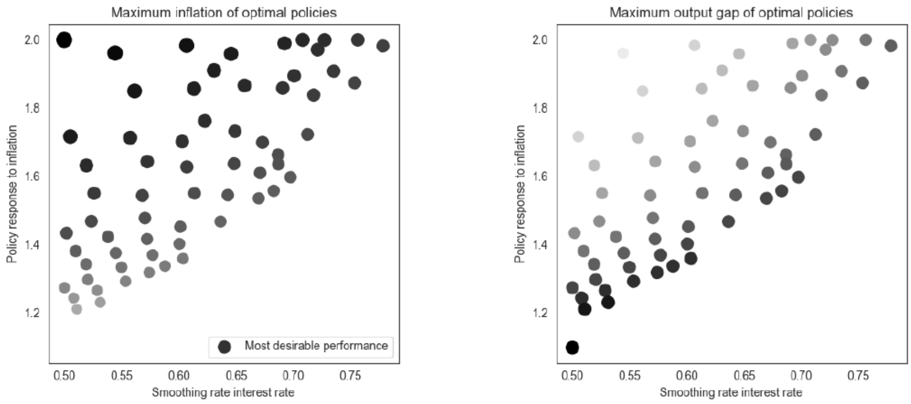

Figure 4. Policy levers (on the x and y axes) and policy performance (on the z-axis). Inflation objective (limiting deviations of inflation from steady state) in the left panel. Output gap objective (limiting output gap deviations) in the right panel. The policymaker prefers a small value for both objectives. The policy outcome for the objective is reflected in the size of the dot (bigger dot is preferred option) and colour (darker colour is preferred). The trade-off emerges as the colours and sizes in both panels are almost inverted.

Based on their performance on the objectives, we end up with four policies that are optimal for a specific objective. Table 4 summarizes the values of the policy levers φ_π and ρ_r that distinguish the four policies, next to the baseline (policy 0). As shown in

Figure 4, the policies that optimise inflation objectives have a higher value of φ_π, while the objective of limiting output gap deviations requires a lower value of ρ_r.

Table 4. Selected policies and the values of their levers. Policy 3(4) aims at limiting the deviations of inflation (output gap) from steady state.

Policy	Description	Policy response to inflation	Smoothing rate interest rate
0	Baseline	1.5	0.8
1	All quarters inside inflation domain	1.961096	0.544366
2	Lowest inflation at time t = 12	1.981912	0.778116
3	Best inflation	2.000000	0.500000
4	Best output gap	1.098116	0.500231

Table 4 shows that policies 1 and 3 have comparable values for φ_π and ρ_r. Policies 1, 2, and 3, which are optimised for inflation objectives, are most responsive to inflation (high φ_π). Policy 4, which is optimised for the output gap objective, has a substantially lower inflation response parameter. Both policies 3 and 4 have a relatively low level of interest rate smoothing (ρ_r), which – according to the finding in

Figure 4 – serves the output gap objective.

Figure 5 shows what the different values of the policy levers imply for the interest rate paths. The rate paths of policies 1 and 3 largely overlap. The interest rate path under policy 4 – the output gap optimising policy – is markedly higher than under the other policies. This is explained by the effect of φ_π and ρ_r on inflation in the DSGE model. A low value of these parameters, as in policy 4, causes large deviations of inflation from the steady state, because low values of both policy levers imply that the interest rate initially is set too low to stabilise inflation after a negative supply shock. As a consequence, the central bank has to raise the interest rate quite aggressively to stabilise inflation over the policy horizon.

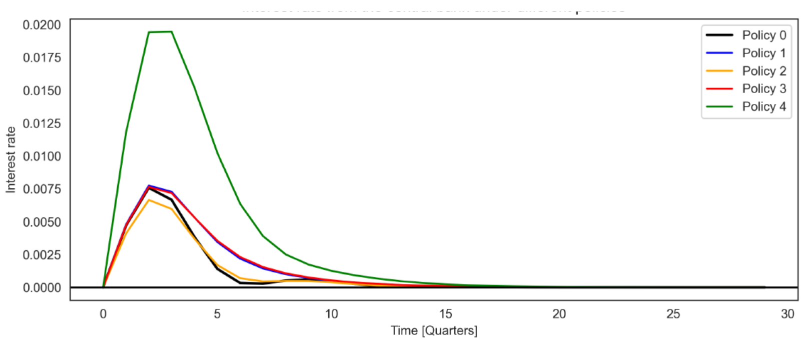

Figure 5. Interest rate paths under different policies that are optimised for a specific objective.

Next, we assess how the four policies perform under deep uncertainty, i.e., in the 15,000 scenarios for the parameter and shock values. Figure 6 shows the areas of the outcomes for inflation (right-hand panels) and the output gap (left-hand panels). The blue areas reflect the outcomes of the baseline policy, and the red areas reflect the outcomes of the other four policies. The dashed lines capture the outer boundaries of 90% of the results.[4]

The results show that policies 1, 2 and 3 (that are optimal for the inflation objectives) generally yield a lower inflation level than the baseline policy. Compared to the baseline, inflation is clearly higher with policy 4. This indicates that a low response to inflation (φ_π), as in policy 4, is not desirable for a central bank focusing on price stability under deep uncertainty. Important to note is that inflation at the end of the policy horizon (t=12) is within an acceptable range for all policies.

Policy lever ρ_r (interest rate smoothing) is low in case of policies 1, 3 and 4. Policy 2 has a higher ρ_r than policies 1 and 3, but a similar value of φ_π. Since the inflation deviations under policies 1 and 3 are somewhat smaller than under policy 2, the outcomes suggest that low interest rate smoothing contributes to the price stability objectives. This result differs from the finding of Orphanides and Williams (2007), who conclude that a more aggressive response to inflation and higher degree of interest rate smoothing would be optimal for price stability in a situation with imperfect knowledge about certain model parameters.

With regard to the output gap, policy 4 results in lower deviations from steady state compared to the baseline policy. Also, compared to policies 1, 2 and 3, policy 4 is more effective in limiting output gap deviations. This indicates that monetary policy characterised by a low response to inflation and low interest rate smoothing contributes to the output gap objectives.

[4] Because the distributions of the uncertainties are not validated, the this 90% statistic is less meaningful than in case of a known underlying distribution.

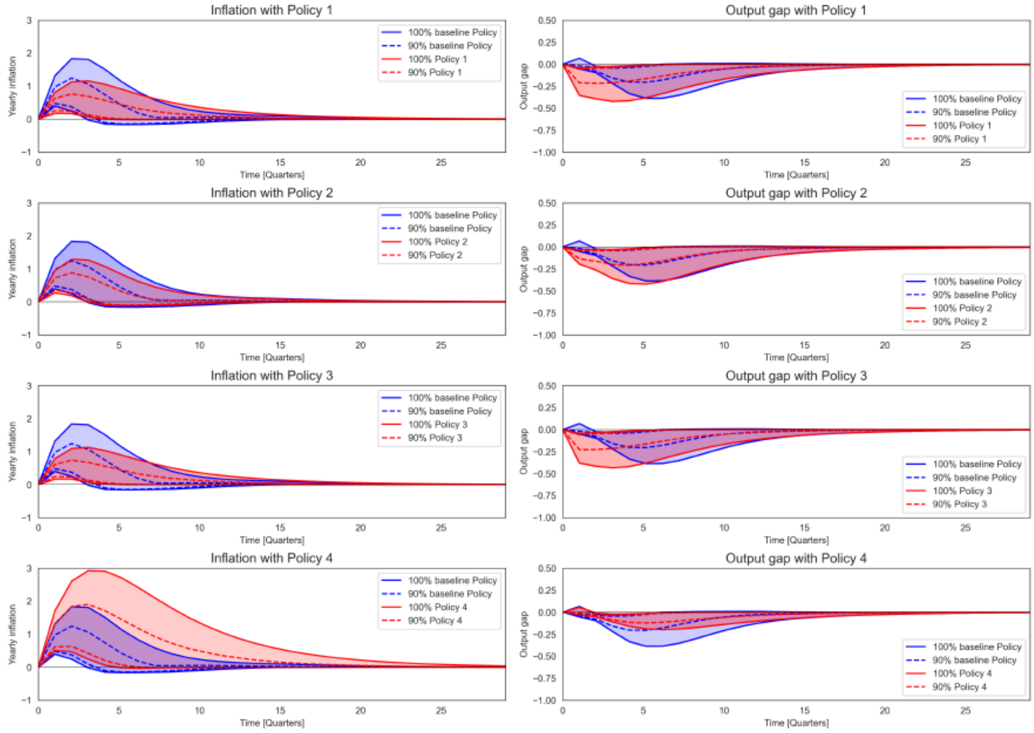

Figure 6. Ranges for inflation (left panels) and the output gap (right panels), for different policies under 15,000 scenarios.

4.2 Robustness

In this section, we assess the robustness of the four policies introduced in the previous section. The policies were optimised conditional on the absence of uncertainty. However, assuming deep uncertainty makes it imperative to evaluate the robustness of the policies in different states of the world. This assessment is based on various robustness metrics. For objectives with a continuous range, a regret metric is used. This metric measures how well the policy performs compared to best performing policy for each scenario. For objectives with a discrete range or that include a threshold, we use a satisficing metric. This metric reflects how many experiments have a desired performance. Lastly, we use some descriptive measures to assess the robustness of the policies.

The overall finding based on the robustness metrics is that no policy clearly outperforms the others on the different objectives. Even though policy 3 (high response to inflation, low interest rate smoothing) has the least regret for the inflation objectives, it has the largest regret on output gap objectives (Table 5). If the central bank focuses mainly on the inflation rate at the end of the policy horizon (t=12), then policy 2 (relative high interest rate smoothing) has the lowest regret for this objective. However, policy 2 has a higher regret than other policies for the objectives limiting deviations of inflation (i.e., maximum inflation) and period-to-period change in inflation.

Table 5. Regret metrics for different policies. Least regret per metric in italics. The objective maximum inflation (output) refers to limiting the deviation of inflation (output gap) from the steady state level.

Policy	Regret maximum inflation	Regret inflation at time t = 12	Adjusted regret inflation at time t = 12
Policy 0	0.468	0.012	0.012
Policy 1	0.029	0.022	0.022
Policy 2	0.152	*0.004*	*0.004*
Policy 3	*0.000*	0.024	0.024
Policy 4	1.057	0.157	0.157

Policy	Regret change inflation	Regret maximum output	Regret standard deviation output gap
Policy 0	0.226	0.075	0.000
Policy 1	0.012	0.095	0.000
Policy 2	0.113	0.079	0.000
Policy 3	*0.000*	0.105	0.000
Policy 4	0.173	*0.000*	0.000

Table 6 shows the outcomes of the satisficing metrics, which are the fractions of cases out of 15,000 scenarios that satisfy certain thresholds. The metrics are all inflation-oriented, since the output gap objectives were not formulated in terms of a threshold. Policy 3, which responds strongly to inflation and has low interest rate smoothing, scores best on the satisficing metrics. Policy 1 scores second best. Interestingly, policy 4, which optimises the output gap objectives, has a high satisficing score for the objective of price stability at the end of the horizon (inflation at t=12). This suggests that, although policy 4 does not perform well on this objective (according to Figure 6 in the previous section), it is quite robust for achieving price stability at t=12 in different states of the world.

Table 6. Satisficing metrics (fraction of cases out of 15,000 scenarios that satisfy the threshold) for different policies. Highest outcomes in italics. For quarters without deflation, the threshold is 0 quarters; for 12 quarters in the inflation domain, the threshold is 12 quarters; and for inflation at time t=12, the threshold is whether inflation is between 2% and 4% annually.

Policy	Quarters without deflation	12 quarters in inflation domain	Inflation at time t = 12 in domain
Policy 0	0.062	0.062	0.726
Policy 1	0.815	0.815	0.964
Policy 2	0.168	0.168	0.785
Policy 3	*0.930*	*0.930*	*0.990*
Policy 4	0.798	0.760	0.908

The descriptive statistics display some more variation among the policies.[5] Robust policies are associated with a low standard deviation. According to this criterion, policy 2 (relative high interest rate smoothing) is most robust with regard to the objective to limit the deviation of inflation at the end of the horizon (t=12). Policy 3 (high inflation response, low smoothing) is most robust concerning the objective to limit large inflation deviations over time. Policy 4 is most robust with regard to the output gap objectives (cf., Table 7 and

Table 8).

Table 7. Descriptive statistics for two inflation objectives. Preferred values in italics. The objective Largest change inflation aims at limiting changes in inflation, as measured by the mean change (pi_change_means) and the standard deviation of the change in inflation (pi_change_std). The statistics are based on the 15,000 scenarios.

Policy	Inflation at time t = 12		Maximum inflation		Largest change inflation	
	Mean	Standard deviation	Mean	Standard deviation	pi_change_means	pi_change_std
Policy 0	0.010	0.017	0.759	0.240	0.302	0.025
Policy 1	0.016	0.026	0.471	0.162	0.126	0.028
Policy 2	*0.008*	*0.013*	0.562	0.171	0.216	*0.024*
Policy 3	0.017	0.027	*0.450*	*0.158*	*0.114*	0.028
Policy 4	0.065	0.108	1.143	0.397	0.249	0.022

Table 8. Descriptive statistics for output gap objectives, based on the 15,000 scenarios. Preferred values in italics.

Policy	Maximum output gap		Deviation in output gap	
	Mean	Standard deviation	Mean [$\times 10^{-3}$]	Standard deviation [$\times 10^{-3}$]
Policy 0	0.105	0.052	0.373	0.176
Policy 1	0.120	0.053	0.424	0.190
Policy 2	0.106	0.054	0.389	0.174
Policy 3	0.127	0.055	0.436	0.197
Policy 4	*0.065*	*0.028*	*0.220*	*0.900*

[5] We note that the distributions of the uncertainties are not validated. Hence, the descriptive metrics are only valid if the uncertainties follow the uniform distribution that we have assigned to them.

5. Discussion

DSGE models differ from models to which EMA methods are usually applied. Such models mostly describe environmental problems, which may include dis-equilibrium dynamics and do not necessarily revert to a steady state following a shock. Also in economics, large and unprecedented shocks, such as the Global Financial Crisis of 2008, may cause output and inflation to deviate from the steady state for a longer period of time than the horizon of monetary policy (Cerra et al., 2021). Moreover, DSGE models tend to converge to the steady state that is known ex-ante. Maier et al. (2016) show that a large part of the uncertainty in the models they discuss relates to the existence of uncertain futures. Also for these reasons EMA is a useful tool to analyse economic policy problems.

Moreover, while policy trade-offs can show up in simulations with a DSGE model, such as the trade-off between inflation and output objectives, these do not present a wicked problem, as is sometimes the case in environmental sciences (Rittel and Webber, 1973). A wicked problem is characterised by complex interactions between interacting agents, bounded rationality and network effects linked to complex interconnectedness and systemic risks. Wicked problems can be addressed by complex adaptive systems approaches. DSGE models, like most structural macro models, take stylised assumptions that improve tractability, but do so at the expense of sacrificing some of the richness of the behaviour of the modelled systems (Kirman et al., 2020).

Our application of EMA can also contribute to the expanding research on enhancing DSGE models with complexities which capture uncertainties. For instance, recent research on HANK type DSGE models combine heterogeneous agent models (macroeconomists' workhorse framework for studying income and wealth distributions) with New Keynesian models (the basic framework for studying monetary policy and movements in aggregate demand, as we have used in our approach), see Kaplan et al. (2023).

Furthermore, economic agents will adjust their behaviour over time, also in response to policy measures (Lucas critique). For instance, in a situation where the central bank is very responsive to inflation with interest rate measures, agents may behave differently than in a situation where the central bank is more reluctant in changing the interest rate. This could imply that for each policy strategy a different distribution or range of uncertain parameters might have to be considered. This could be considered a priori, by changing the distribution for each policy when evaluating its performance and robustness, or a posteriori by including this in the interpretation of the results.

Another criticism on structural macro models like DSGE models, is the assumption of time invariant parameters, also in the Taylor rule. Hurtado (2014) shows that the parameters in the DSGE models, including those assumed to be structural, change over time. From this, Storm (2021) concludes that users of DSGE models should be aware of the dynamics and be cautious about the effect of them on the outcomes. Our approach partly fills this awareness gap by assuming that - although the model parameters are time-invariant - they differ in each experiment.

6. Conclusion

We introduce EMA into economics and provide a novel way to address deep uncertainty in structural economic models, like DSGE models used by central banks. They usually make assumptions about the values of the parameters and shocks when applying those models. This introduces parametric and shock uncertainty. We show how EMA is useful to analyse how the model outcomes are affected by these uncertainties and we show the use of EMA to design policies that are robust to future unknown states of the world. Under those conditions, it is important to know the scenarios that are policy relevant; in other words, which combinations of parameters drive these outcomes of interest. Open exploration by EMA enables such a scenario discovery, by identifying which parameters and shock values are associated with the model outcome of interest.

Applied to a negative supply shock in a basic DSGE model, the EMA open exploration shows that states of the world with low price flexibility and high shock persistence drive model outcomes that are of interest to the central bank. Such states are associated with relatively large deviations of the output gap and inflation from steady state, by which the scenarios may lead to outcomes that conflict with the policy objectives. Moreover, a directed search finds that the policy levers, the parameters in the Taylor rule that the central bank can influence, also have a large influence on the economic dynamics following a negative supply shock. We find that, under deep uncertainty, a rather strong response to inflation contributes to the price stability objectives, while low interest rate smoothing contributes to both the inflation and output gap objectives. Both policies are also robust to uncertain states of the world.

The results show that EMA can provide useful insights into DSGE models and support the central bank in designing robust policies. However, more steps are needed before the central bank can incorporate EMA in its toolkit and policy process. Our approach is based on a theoretical case and uses the most basic and stylized DSGE model. In practice, central banks use more complex DSGE models and/or have other policy interests than we have explored. That may require other objectives and robustness metrics.

To apply EMA in the policy practice of central banks, collaborative research is needed to combine economic insights about the (practical) use of DSGE models with the EMA modeler's experience to find the best suitable algorithms and methods to solve actual monetary policy problems. Future research is recommended to optimise the EMA methods and algorithms, such as adjusting the clustering algorithm for the time series after a supply shock, that best serve DSGE models and their functionalities in a situation of deep uncertainty.

Literature

Bankes, S. C. (1993), Exploratory Modeling for Policy Analysis, Operations Research, 41(3), 435-449.

Bankes, S. C., Walker, W. E. and Kwakkel, J. H. (2013). Exploratory Modeling and Analysis. In G. S. and M. C. Fu (Eds.), Encyclopedia of Operations Research and Management Science (3rd ed.). Berlin, Germany: Springer.

Bartholomew, E. and J.H. Kwakkel (2020), On considering robustness in the search phase of Robust Decision Making: A comparison of Many-Objective Robust

Decision Making, multi-scenario Many-Objective Robust Decision Making, and Many Objective Robust Optimisation, Environmental Modelling & Software, 127.
Ben-Haim, Y., M. Demertzis and J.W. van den End (2018), Evaluating monetary policy rules under fundamental uncertainty: an info-gap approach, Economic Modelling 73, 55-70.
Boivin, J. and M. P. Giannoni (2006), Has Monetary Policy Become More Effective?, The Review of Economics and Statistics, 88(3), p. 445-462.
Bryant, B. P. and R. J. Lempert (2010), Thinking Inside the Box: a participatory computer-assisted approach to scenario discovery, Technological Forecasting and Social Change, 77(1), 34-49.
Čapek, J., J. Crespo Cuaresma, N. Hauzenberger and V. Reichel (2023), Macroeconomic forecasting in the euro area using predictive combinations of DSGE models, International Journal of Forecasting, 39(4), p. 1820-1838.
Carceller del Arco, M. and J. W. van den End (2023), Robust monetary policy under shock uncertainty, DNB working paper, 793.
Cerra, V., M. Hakamada and R. Lama (2021), Financial Crises, Investment Slumps, and Slow Recoveries, IMF working papers, 2021/170.
Christiano, L. J., M. Trabandt and K. Walentin (2010), DSGE Models for Monetary Policy Analysis, NBER working paper, 16074.
Dalal, S., B. Han, R. Lempert, A. Jaycocks and A. Hackbarth (2013), Improving scenario discovery using orthogonal rotations, Environmental Modelling & Software, 48, 49–64.
Del Negro, M. and F. Schorfheide (2013), DSGE Model-Based Forecasting, Chapter 2 in Handbook of Economic Forecasting, 2, p. 57-140, Elsevier.
Dennis, R., K. Leitemo and U. Söderström (2009), Methods for robust control, Journal of Economic Dynamics and Control, 33(8), p. 1604-1616.
Friedman, J. H. and N.I. Fisher (1999), Bump hunting in high-dimensional data, Statistics and computing, 9(2), 123–143.
Galí, J. (2015), Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications, Economics Books, Princeton University Press, edition 2.
Giannoni, M.P. (2007), Robust Optimal Monetary Policy in a Forward-Looking Model with Parameter and Shock Uncertainty, Journal of Applied Econometrics , 22(1), 179-213.
Górajski, M, Z. Kuchta and A. Leszczyńska-Paczesna (2023), Price-setting heterogeneity and robust monetary policy in a two-sector DSGE model of a small open economy, Economic Modelling, 122(C).
Grassi, S., M. Leon-Ledesma and F. Ferroni (2016), Fundamental shock selection in DSGE models, Meeting Papers 47, Society for Economic Dynamics.
Grigoli, F. (2015), A Hybrid Approach to Estimating the Efficiency of Public Spending on Education in Emerging and Developing Economies, Applied Economics and Finance, 2(1), 19-32.
Haasnoot, M., J. H. Kwakkel, W. E. Walker and J. Ter Maat (2013), Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world. Global environmental change, 23(2), 485–498.
Hallegatte, S., A, Shah, C. Brown, R. Lempert and S. Gill (2012). Investment decision making under deep uncertainty–application to climate change. World Bank, Working paper, 6193.

Herman, J. D., P. M. Reed, H. B. Zeff and G. W. Characklis (2015), How should robustness be defined for water systems planning under change? Journal of Water Resources Planning and Management, 141(10), 04015012.
Homaei, S. and M. Hamdy (2020). A robustness-based decision making approach for multi-target high performance buildings under uncertain scenarios. Applied Energy, 267, 114868.
Hurtado, S. (2014), DSGE models and the Lucas critique, Economic Modelling, 44, S12–S19.
Kaplan, G., B. Moll and G.L. Violante (2023), The very model of modern monetary policy, IMF Finance and developments, March 2023.
Kirman, A., Dellink, R., Chateau, J., Barnes, S., Rovenskaya, E. , Dieckmann, U., van Ruijven, B. and K. Riahi (2020), Methodologies and Tools for Integrated Systems Modelling. In: Systemic Thinking for Policy Making: The Potential of Systems Analysis for Addressing Global Policy Challenges in the 21st Century, p. 29-38, Paris, France: New Approaches to Economic Challenges, OECD.
Krugman, P. and R. Wells (2021), Macroeconomics. Worth Publishers, Macmillan Learning.
Kwakkel, J. H., W.L. Auping and E. Pruyt (2013), Dynamic scenario discovery under deep uncertainty: The future of copper, Technological Forecasting and Social Change, 80(4), 789–800.
Kwakkel, J. H. and E. Pruyt (2015), Using system dynamics for grand challenges: The esdma approach, Systems Research and Behavioural Science, 32(3), 358–375.
Kwakkel, J. H., M. Haasnoot and W. E. Walker (2016), Comparing robust decision-making and dynamic adaptive policy pathways for model-based decision support under deep uncertainty. Environmental Modelling & Software, 86, p. 168-183.
Kwakkel, J. H. and M. Jaxa-Rozen (2016), Improving scenario discovery for handling heterogeneous uncertainties and multinomial classified outcomes, Environmental Modelling & Software, 79, 311–321.
Kwakkel, J. H. (2017), The exploratory modeling workbench: An open source toolkit for exploratory modeling, scenario discovery, and (multi-objective) robust decision making, Environmental Modelling & Software, 96, 239–250.
Lempert, R. J., S.W. Popper and S.C. Bankes (2003), Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis, Technological Forecasting and Social Change, RAND.
Lempert, R. J., D. G. Groves, S. W. Popper and S. C. Bankes (2006), A General, Analytic Method for Generating Robust Strategies and Narrative Scenarios. Management Science, 52(4), 514-528.
Maier, H. R., J. H. Guillaume, H. van Delden, G.A. Riddell, M. Haasnoot and J. H. Kwakkel (2016), An uncertain future, deep uncertainty, scenarios, robustness and adaptation: How do they fit together? Environmental modelling & software, 81, p. 154 -164.
McKay, M. D., R.J. Beckman and W.J. Conover (1979), A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code, Technometrics, 21(2), 239-245.
McPhail, C., H. R. Maier, J. H. Kwakkel, E. Giuliani, A. Castelletti and S. Westra (2018), Robustness metrics: How are they calculated, when should they be used and why do they give different results? Earth's Future.
Moallemi, E.A., J. Kwakkel, J. Fjalar, J. de Haan and B. A. Bryan (2020), Exploratory modeling for analysing coupled human-natural systems under uncertainty, Global Environmental Change, 65.

Olalla, M. G. and A. R. Gómez (2011), Robust control and central banking behaviour, Economic Modelling, 28(3), 1265-1278.
Orphanides, A. and J. C. Williams (2007), Robust monetary policy with imperfect knowledge, Journal of Monetary Economics, 54(5), 1406-1435.
Ratto, M. (2008), Analysing DSGE Models with Global Sensitivity Analysis, Computational Economics, 31(2), p. 115-139.
Rittel, H. W. J. and M. W. Webber (1973), Dilemmas in a General Theory of Planning, Policy Sciences, 4(2), 155-169.
Smets, F. and R. Wouters (2003), An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area, Journal of the European Economic Association, MIT Press, 1(5), p. 1123-1175.
Statnikov, R. B. and J. B. Matusov (2012). Multicriteria optimisation and engineering. Springer Science & Business Media.
Steinmann, P., W.L. Auping, and J.H. Kwakkel (2020). Behaviour-based scenario discovery using time series clustering. Technological Forecasting and Social Change, 156, 120052.
Storm, S. (2021), Cordon of conformity: Why DSGE models are not the future of macroeconomics, International Journal of Political Economy, 50(2), 77–98.
Taylor, J.B. and J. C. Williams (2010), Simple and robust rules for monetary policy, Working Paper 2010-10, Federal Reserve Bank of San Francisco.
Xiao B., Y. Fan and X. Guo (2018), Exploring the macroeconomic fluctuations under different environmental policies in China: a DSGE approach, Energy Economics, 76, p. 439–456.
Walker, W., V. Marchau and J. Kwakkel (2013), Uncertainty in the framework of policy analysis, in W. Walker and W. Thissen (Eds.), Public policy analysis : New developments, p. 215–261, Springer.

DeNederlandscheBank
EUROSYSTEEM

De Nederlandsche Bank N.V.
Postbus 98, 1000 AB Amsterdam
020 524 91 11
dnb.nl