

Böhl, Gregor

Working Paper

Ensemble MCMC sampling for robust Bayesian inference

IMFS Working Paper Series, No. 177

Provided in Cooperation with:

Institute for Monetary and Financial Stability (IMFS), Goethe University Frankfurt am Main

Suggested Citation: Böhl, Gregor (2022) : Ensemble MCMC sampling for robust Bayesian inference, IMFS Working Paper Series, No. 177, Goethe University Frankfurt, Institute for Monetary and Financial Stability (IMFS), Frankfurt a. M.

This Version is available at:

<https://hdl.handle.net/10419/268226>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

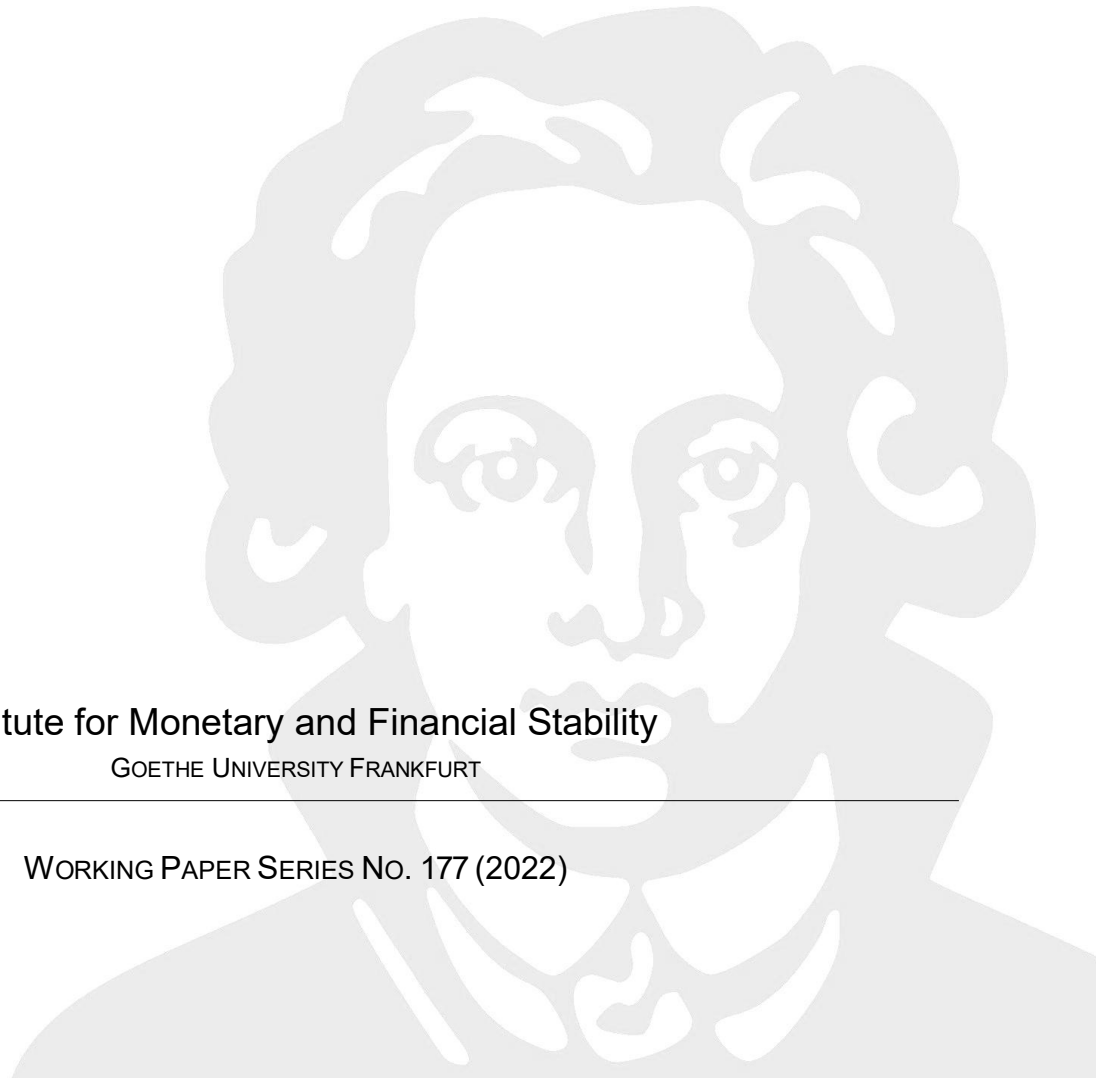
If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

GREGOR BOEHL

Ensemble MCMC Sampling for Robust Bayesian Inference

Institute for Monetary and Financial Stability
GOETHE UNIVERSITY FRANKFURT

WORKING PAPER SERIES NO. 177 (2022)



This Working Paper is issued under the auspices of the Institute for Monetary and Financial Stability (IMFS). Any opinions expressed here are those of the author(s) and not those of the IMFS. Research disseminated by the IMFS may include views on policy, but the IMFS itself takes no institutional policy positions.

The IMFS aims at raising public awareness of the importance of monetary and financial stability. Its main objective is the implementation of the “Project Monetary and Financial Stability” that is supported by the Foundation of Monetary and Financial Stability. The foundation was established on January 1, 2002 by federal law. Its endowment funds come from the sale of 1 DM gold coins in 2001 that were issued at the occasion of the euro cash introduction in memory of the D-Mark.

The IMFS Working Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Institute for Monetary and Financial Stability

Goethe University Frankfurt

House of Finance

Theodor-W.-Adorno-Platz 3

D-60629 Frankfurt am Main

www.imfs-frankfurt.de | info@imfs-frankfurt.de

Ensemble MCMC Sampling for Robust Bayesian Inference

Gregor Boehl

University of Bonn

November 23, 2022

Abstract

This paper proposes a Differential-Independence Mixture Ensemble (DIME) sampler for the Bayesian estimation of macroeconomic models. It allows sampling from particularly challenging, high-dimensional black-box posterior distributions which may also be computationally expensive to evaluate. DIME is a “Swiss Army knife”, combining the advantages of a broad class of gradient-free global multi-start optimizers with the properties of a Monte Carlo Markov chain. This includes (i) fast burn-in and convergence absent any prior numerical optimization or initial guesses, (ii) good performance for multimodal distributions, (iii) a large number of chains (the “ensemble”) running in parallel, (iv) an endogenous proposal density generated from the state of the full ensemble, which (v) respects the bounds of the prior distribution. I show that the number of parallel chains scales well with the number of necessary ensemble iterations. DIME is used to estimate the medium-scale heterogeneous agent New Keynesian (“HANK”) model with liquid and illiquid assets, thereby for the first time allowing to also include the households’ preference parameters. The results mildly point towards a less accentuated role of household heterogeneity for the empirical macroeconomic dynamics.

Keywords: Bayesian Estimation, Monte Carlo Methods, Heterogeneous Agents, Global Optimization, Swiss Army Knife

JEL: C11, C13, C15, E10

1 Introduction

Since the pioneering work of Geweke (1999) and Schorfheide (2000), Bayesian methods have found their way into the toolboxes of researchers at universities and central banks.

*Address: Institute for Macroeconomics and Econometrics, University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany. I am grateful to Christian Bayer, Flora Budianto, Cees Diks, Dominik Hecker, Alexander Meyer-Gohde, Felix Strobel and participants of several conferences and seminars for discussions and helpful comments on the contents of this paper. Part of the research leading to the results in this paper has received financial support from the Alfred P. Sloan Foundation under the grant agreement G-2016-7176 for the Macroeconomic Model Comparison Initiative (MMCI) at the Institute for Monetary and Financial Stability. I also gratefully acknowledge financial support by the Deutsche Forschungsgemeinschaft (DFG) under CRC-TR 224 (projects C01 and C05) and under project number 441540692.

Email address: gboehl@uni-bonn.de

URL: <https://gregorboehl.com>

They are used extensively to bring (semi-)structural models to the data, such as modern New Keynesian-type DSGE models or various forms of Bayesian vector autoregression models. Thereby, they allow to evaluate the empirical performance of these models, to quantify the effects of macroeconomic policy, and to sensibly assess the uncertainty surrounding the economic analysis. Yet, the application of Bayesian methods can be very challenging in practice as they require the *identification* and *sampling* from the high-probability density region of the associated posterior distribution of parameters. This posterior distribution is a potentially highly complex topology with large dimensionality and ex-ante unknown properties. Standard optimization and sampling tools often do not perform well on such distributions. However, the quality of the estimation – and thereby their usefulness for economic analysis – crucially depends on our ability to precisely pin-down this posterior distribution.

This paper substantiates this ability by introducing a novel sampling algorithm: the *differential-independence mixture ensemble* (DIME) Monte Carlo Markov chain method. Motivated by recent advances in the field of computational astrophysics, the DIME sampler seeks to satisfy five practical requirements:

- i) Fast burn-in to the high probability density region of the posterior, absent any prior posterior mode density optimization or initial guesses.
- ii) Good performance for high-dimensional, multimodel and complex distributions.
- iii) Convergence speed scales well with the number of chains, allowing for efficient parallelization.
- iv) An endogenous proposal distribution generated from the current state of all chains.
- v) The proposal distribution that respect the bounds of the prior distribution.

Point i) is desirable because a common practice in economic applications is to treat *finding* (or identifying) the high probability density region separately from the problem of actual *sampling* from it. A significant amount of processing time is hence, instead of posterior sampling, spend on numerical optimization routines for mode finding. However, the posterior of DSGE models is often not only high-dimensional, but also non-monotonic and discontinuous due to e.g. issues with indeterminacy, the various cross-equation restrictions, or a misalignment of prior and likelihood distributions.¹ Numerical optimizers tend to behave unstable for such functions, can show strong dependence on initial guesses, and may, if at all, converge to local maxima only. Likewise, conventional samplers often are not well suited to deal with such functions reliably. As it is usually very difficult to ex-post determine whether the result is a reliable sample from the posterior or not, this motivates requirement ii).

Point iii) is worthwhile because current DSGE models are becoming more and more expensive to evaluate, either because of severe nonlinearities or because agents are het-

¹A forward looking model, such as a DSGE model, is called *indetermined* if there exists no unique rational expectations solution. For linearized models this is linked to the eigenvalues of the associated system of difference equations, which are a function of the model parameters. An additional problem with nonlinear models constitutes through the process of nonlinear filtering, which is often also based on sampling. This can lead to noisy likelihood estimates. Due to the particular complications that arise with these models, I here focus on DSGE models rather than Bayesian vector autoregression models, for which DIME MCMC is applicable nonetheless.

erogeneous across multiple dimensions.² As multi-core architectures have successively become more affordable – contemporary laptop computers often already come with 8 processors, and build-in graphics processing units (GPUs) may feature many more – researchers want to take advantage of this development by being able to run estimations in parallel, thereby significantly reducing total runtimes. Requirement iv) is important because ex-ante, the posterior distribution is a black box to the researcher: the specific features of the distribution are unknown, leaving researchers unable to make good choices on the setup of the sampling algorithm on their own. Point v) helps to avoid large rejection rates due to draws falling outside the support of the prior distribution.

To meet these requirements, instead of using a single or small number of recursive chains (such as e.g. the random-walk Metropolis algorithm), DIME relies on an *ensemble* of a large number of chains which jointly evolve over time. For each iteration, proposals are generated based on the current state of the full ensemble and, as the ensemble evolves over time, proposal steps naturally adapt direction and scale of the estimated posterior distribution. After convergence, the invariant distribution of *all* chains corresponds to target distribution. The sampler is mixing between a local and a global transition kernel: the local kernel explores the direct proximity of a particular chain. The global kernel, in contrast, reshuffles chains over the complete domain of the current approximation of the posterior distribution. This means that DIME MCMC is equally efficient for posterior sampling and for quick convergence to the high-density region of the posterior (called *burn-in*), and makes no difference between the two. The sampler can hence be seen as a “Swiss Army knife” for structural econometric analysis.

The local transition kernel builds on the *differential evolution* (DE) concept developed in the literature on global optimization. DE optimizes a function by maintaining a population of candidate solutions and creating new candidate solutions by combining existing ones, and then keeping whichever candidate solution has the best fitness on the optimization problem at hand. This can be turned into an MCMC method by exchanging the last step by the Metropolis-Hastings algorithm.³ A major problem with this MCMC version of DE is that, although proposals are state-dependent and adaptive, chains evolve ex-ante independently. This frequently causes the dispersion among chains to increase over time, which deteriorates the quality of proposals, thereby leading to slow overall convergence of the full ensemble. The algorithm also does not perform well with multimodal distributions because chains are unlikely to switch modes.

In contrast, at the core of the *global transition kernel* lies an ensemble version of a modified adaptive Independence Metropolis-Hastings method, where candidates are created based on a proposal distribution that is independent of the state of a single chain.⁴ This attribute makes the global transition kernel fully robust against odd-shaped and multimodal distributions. However, independence Metropolis-Hastings performs only

²See, e.g. Boehl and Strobel (2020), for the estimation of medium scale DSGE models with the zero-lower bound on nominal interest rates as an example for nonlinear estimation, or Bayer et al. (2020) for the estimation of heterogeneous agent models.

³See Storn and Price (1997) for the global optimizer and Ter Braak (2006); Nelson et al. (2013) for Differential Evolution MCMC.

⁴For independence Metropolis-Hastings see, e.g., Tierney (1994). For specifications of adaptive independence Metropolis-Hastings see for example Haario et al. (2001) and Roberts and Rosenthal (2007). A similar algorithm from the global optimization literature is the covariance matrix adaptation evolution strategy (CMA-ES, Igel et al., 2007).

well if the proposal distribution is stationary and close to the target distribution. Since this requirement is very difficult to be met ex-ante, the algorithm had limited appeal for many target distributions relevant in practice. To circumvent this problem, I develop a time-varying proposal distribution that adjusts to new ensembles based on their average posterior density. This puts decaying weights on early samples but guarantees convergence to a stationary proposal distribution once the average density of candidates is stationary. While this improves performance considerably, the global transition kernel alone still converges slowly, and is not robust to multimodality.

DIME MCMC exploits the complementarity of the local and global transition kernel: the combination of the two kernel dispels the individual weaknesses. In a mixture, the global kernel occasionally reshuffles some of the chains, which counteracts dispersion of the ensemble and ensures that individual chains do not “get stuck” in local maxima. This, in turn, also increases the quality of DE proposals from the local transition kernel. The independent proposals from the global transition kernel also make sure that chains switch between modes for multimodal distributions. On the other side, the differential evolution heritage of the local transition kernel enables to use the sampler on target distributions (or objective functions) that are discontinuous or noisy. This local search generates good proposal candidates during burn in, and provides updates for the proposal distribution of the global transition kernel. The algorithm is also similar to multi-start optimizers as it searches the complete relevant function domain. DIME hence combines the advantages of a broad class of global optimizers with the properties of a Monte Carlo Markov chain (MCMC) sampler, thereby satisfying requirements i) and ii).

To address point iii), the ensemble structure makes DIME MCMC “embarrassingly parallelizable”. While there is some trade-off between the quality of proposals (which increases with the number of chains) and the number of iterations (lesser chains increase convergence rate per function evaluation), I show in Section 4 that the sampler scales quite well. The method is essentially self-tuning and, solely requires setting the number of chains as the only metaparameter, thereby satisfying requirement iv).⁵ Finally, to address point v) I distinct between parameter and proposal space by introducing a bijective mapping between these two. This allows to let the sampling algorithm run in a unbounded space, which boosts acceptance ratios, thereby again increasing robustness and sampling efficiency.⁶

In this paper I asses the performance of DIME MCMC on three very distinct use cases. I first evaluate the algorithm’s capability to deal with high-dimensional and bimodal distributions with ex-ante known properties. I document that the sampler performs well on such distributions, even when the two modes are fully disconnected. I then test the performance of the sampler on the estimation exercise from Smets and Wouters (2007). DIME MCMC returns the original parameter estimates independently of the number of chains used. For the given example, I find that convergence times roughly scale well with the number of chains, which suggests that the losses through parallelization mainly amount to the computational overhead of serialization.

Finally, I estimate a heterogeneous agents New Keynesian model, including the house-

⁵Sections 3 and 4 discuss the role of the number of chains and of the kernel mixing probability, and provide sane defaults for these parameters.

⁶A similar concept, so-called *bijectors*, is applied in the literature on neural networks where they are used to create proxy-posteriors which feature a more favorable geometry. See, e.g., Dillon et al. (2017).

holds preference parameters. These parameters may be of particular relevance on their own as they govern the steady state distribution of assets. This exercise was so far deemed impossible due to the large computational costs associated with solving for the steady state distribution for each single likelihood evaluation of the model, and is only enabled by the fact that DIME MCMC is trivial to parallelize. The estimation results point towards a rather attenuated role of portfolio choice for macroeconomic dynamics, with the parameter that determines the magnitude of the liquidity friction being identified significantly below its prior mean. The degree of idiosyncratic income risk is also estimated to be below its prior mean, but still in the range of values used in the literature.

This paper comes with reference implementations of DIME MCMC in Python and Julia programming languages, and for matlab. The implementations for Python and Julia can directly be installed through the official software repositories, and are actively developed at Github. The Python package integrates into the well-established emcee-package, which is a collection of (ensemble) MCMC samplers (Foreman-Mackey et al., 2013).⁷

Literature

The workhorse of Bayesian estimations in many economic applications is the random walk Metropolis Hastings (RWMH) algorithm, which dates back to the seminal work of Metropolis et al. (1953) and Hastings (1970). The shortcomings of RWMH are well documented (e.g. Chib and Ramamurthy, 2010; Herbst and Schorfheide, 2015). The main issue is that convergence of RWMH to the posterior distribution can be extremely slow, and sampling from ill-shaped or multimodal distributions is hardly possible in practice. To circumvent the first problem, numerical optimization routines are frequently used to find a good initialization values for RWMH. These routines are often slow as well, and not very robust when applied to more complicated posterior distributions. In particular, they tend to “getting stuck” at local maxima. Another problem with RWMH as well as with most numerical optimizers is that they are not parallelizable due to their recursive nature. They therefore can not benefit from multi-core architectures, which is important if the posterior density is computationally expensive to evaluate.

Well-known alternatives to RWMH include Gibbs and slice sampling (Geman and Geman, 1984; Damlén et al., 1999), which perform better on high-dimensional distributions. They are, however, not robust to multimodal distributions and do not perform well during burn-in and convergence to the high probability density region. Also, these methods can not trivially be parallelized. A recent innovation from the econometrics community is sequential Monte Carlo (SMC) method introduced in Herbst and Schorfheide (2014). The core idea is to run many RWMH chains in parallel interrupted by occasional resampling stages to ensure that all chains converge to the high probability density region. In order to prevent convergence to local optima, the authors develop a tempering scheme for SMC. By construction, this circumvents many of the shortcomings of standard RWMH and, given the right choice of a tempering scheme, can also performs well on multimodal

⁷Documentation and downloads for the Python package can be found at <https://github.com/gboehl/emcwrap>. The standalone version in Julia programming language is located at <https://github.com/gboehl/DIMESampler.jl>. The matlab implementation can be found at <https://github.com/gboehl/dime-mcmc-matlab>.

distributions. SMC is also reported to work well on vector autoregression (VAR) models (Bognanni and Herbst, 2018). The combination of tempering with RWMH chains may have the disadvantage of relatively slow convergence. Additionally, the method has relatively many degrees of freedom in the choice of metaparameters, which may determine overall performance. In contrast, the proposal density of DIME is endogenous and, through the adaptation extensions, chains converge more quickly. Like SMC, DIME can straightforwardly be applied to VAR models.

Research in the field of astrophysics has recently made considerable progress on the frontier of Monte Carlo sampling. Ensemble MCMC is conceptionally introduced by Goodman and Weare (2010). The authors develop the idea of an ensemble of Markov chains which, based on the current state of all chains, generates proposals inspired by the numerical optimization method of Nelder and Mead (1965). They show that such sampler is affine invariant and “uniformly effective over all the convex bodies of a given dimension regardless of their shape”, thereby significantly outperforming RWMH. The success of Ensemble MCMC methods is accelerated by its excellent implementation in the open source packet *emcee* (Foreman-Mackey et al., 2013).⁸ As shown in Section 4, Goodman and Weare (2010) indeed performs well in terms of sampling efficiency but, at least for the models considered here, is rather slow to converge to the posterior distribution. As acknowledged by the authors, the method by construction does not perform well for multimodal distributions.

Building on the differential evolution MCMC sampler of Ter Braak (2006), Vrugt et al. (2009) follow some similar ideas to the ones developed here. A common problem with multi-chain methods such as DE-MCMC is that when considering interacting vectors, the entire ensemble has to be considered as a whole, which increases n -fold the dimension of the target and may thus significantly impact convergence. To address this problem, the authors add a scheme to resample DE-MCMC chains that are stuck, which disburdens the problem of overly dispersed chains. While the replacement of malperforming chains is an important issue that is also pointed out in ter Braak and Vrugt (2008), the proposed heuristic for outlier detection may not work well for all distributions in practice. To allow better support for multimodal distributions the authors also add random proposals for which the jump distance is unity. While this is a practical workaround, it is likely to also slow down convergence, in particular for more challenging distributions. A nice addition is a crossover step to decrease autocorrelation similar to the *snooker move* introduced in ter Braak and Vrugt (2008). An extension of DIME along these lines indeed increases convergence speed and decreases autocorrelation times, but comes at the expense of being less robust to multimodal distributions.

The recent rise of frameworks allowing for *automatic differentiation*⁹ (AD) has renewed interest in the Hamiltonian Monte Carlo (HMC) method (Duane et al., 1987; Childers et al., 2022). HMC proposals are based on the Jacobian of the posterior distribution, which normally is expensive to evaluate (e.g. via finite difference methods). However, AD provides computationally efficient means to calculate Jacobians. HMC clearly outperforms RWMH in terms of sampling efficiency and in its capability to sam-

⁸Emcee is implemented in the Python language and can be found at <https://github.com/dfm/emcee>. The package also provides routines for efficient parallelization.

⁹AD is e.g. available through the Python packages *JAX* or *TensorFlow*, or in the new *Julia* programming language.

ple from more complex distributions. A descendant of HMC is e.g. implemented in the well-known *STAN* framework. Drawbacks of HMC are that it does not necessary provide fast burn-in, and does not perform well for multimodal distributions if the modes are sufficiently disconnected. Further, HMC requires that the parameter space is continuous (c.f. Neal et al., 2011), which is generally not the case for DSGE models due to parameter combinations for which the model is indetermined or explosive. Lastly, HMC requires the implementation of the likelihood function – and hence the complete model and filtering routines – in a framework that allows for AD, which may require a major programming effort. Section 6 briefly touches upon a mixture sampler of DIME with the HMC method.

The rest of the paper is structured as follows. Section 2 explains the basic DIME algorithm. Section 3 studies the performance of the algorithm on a high dimensional bimodal distribution. In section 4 the sampler is used on the Smets-Wouters model and in section 5 it is applied to the estimation of a large-scale HANK model. Section 6 concludes.

2 Mixture Ensemble MCMC Sampling

Let $\pi(x)$ be the probability density of a target distribution with $x \in \mathbb{R}^n$. In practice, $\pi(x)$ is often the posterior density $\pi(x) = p(x|Y)$, which for given data Y and model x equals

$$p(x|Y) = \frac{p(Y|x)p(x)}{p(Y)}. \quad (1)$$

$p(Y|x)$ is the likelihood which, provided (x, Y) , can be calculated using various Bayesian filtering techniques. I especially consider cases where the evaluation of $p(Y|x)$ may be very computationally expensive. Let me assume that the prior $p(x)$ is specified such that it is straightforward to evaluate and to sample from, and

$$p(Y) = \int p(Y|x)p(x)dx \quad (2)$$

is an unknown constant for given data Y . We then wish to draw a sufficiently large number of samples from π in order to approximate some quantity

$$E_\pi [h(x)] = \int h(x)\pi(x)dx \approx \frac{1}{N} \sum_i h(x_i) \quad (3)$$

while minimizing the number of necessary likelihood evaluations.

2.1 Random Walk Metropolis-Hastings

As a reference point, let me briefly sketch the classic random walk Metropolis-Hastings algorithm (Hastings, 1970, RWMH). Start with a *single* parameter vector X_t at iteration t . A new replacement candidate is generated by $\hat{X}_t = X_t + \varepsilon_t$ where $\varepsilon_t \sim \mathcal{N}(0, \Sigma)$ is called the RWMH *proposal distribution*, which is often assumed to follow a multivariate normal distribution. The replacement candidate \hat{X}_t is accepted with the Metropolis acceptance probability

$$P(X_{t+1} = \hat{X}_t) = \min \left\{ 1, \frac{\pi(\hat{X}_t)}{\pi(X_t)} \right\}. \quad (4)$$

If it is accepted, set $X_{t+1} = \hat{X}_t$. Otherwise, set $X_{t+1} = X_t$. A large literature discusses the properties of RWMH, see e.g. Sokal (1997) or Roberts and Rosenthal (2001).

The practical performance of the algorithm crucially depends on the choice of the proposal distribution, i.e. here on the covariance matrix Σ . This may be problematic since Σ has $\frac{d(d+1)}{2}$ degrees of freedom and it is challenging to determine ex-ante which choice of Σ will maintain a high acceptance ratio while still exploring the target distribution to a satisfactory degree. To maintain a sufficiently large acceptance ratio, Σ is often scaled down to relatively small values. Consequently, RWMH is very slow to converge to the high probability density region of the posterior (so-called *burn-in* or *thermalization*). To speed up computation, RWMH is thus often used subsequent to a numerical optimization routine, which is supposed to provide better starting values. As discussed above, such numerical optimization routines may also have severe limitations.

DIME MCMC uses a different approach. It combines the characteristics of a broad class of global optimizers with the properties of a MCMC sampler. The first feature is that the sampler suggests local proposals – a replacement candidate that for each chain is relative to its previous state – as well as global proposals that are independent of the state of a single chain. Both proposal densities adapt to the state of the complete ensemble, explicitly for the global transition kernel and implicitly for the local transition kernel. The coexistence of local and global proposals prevents single chains from “getting stuck” at local maxima, speeds up convergence, and eases sampling from distributions with two or more modes, even if these are fully separated. The second feature is the separation of proposal space from parameter space, which ensures that any proposed replacement candidate has a positive prior probability. This increases acceptance rates notably.

2.2 Proposal space vs. parameter space

The prior distribution often has bounded support. Naturally, replacement candidates beyond these bounds are always rejected. This is problematic in particular for medium- and large-scale DSGE models as these frequently feature exogenous AR(1) processes with roots close to a unity. Since the prior of these roots is bounded by $(0, 1)$, estimates close to unit roots will often cause poor sampling performance because any proposal with values of the AR-coefficient larger one will be rejected. A model with several AR(1) processes close to unit roots will hence feature a rather low acceptance fraction during MCMC sampling.

To circumvent this problem define the *parameter space* $\mathbb{X} : x \in \mathbb{X} \Leftrightarrow p(x) > 0$ to be the space of all parameter combinations for which the prior density is positive. Let the *proposal space* $\mathbb{Z} = \mathbb{R}^n$ be unbounded and f_b be a *bijective map*

$$f_b : \mathbb{R}^n \rightarrow \mathbb{X} \tag{5}$$

such that for any $x \in \mathbb{X}$ there exists a unique $z \in \mathbb{R}^n$ for which $f_b(z) = x$. f_b then is always uniquely invertible, and by definition, f_b maps within the bounds of the prior distribution whereas its domain is unbounded. Below, f_b will be used to ensure that every sample has a positive prior density.

In the spirit of Goodman and Weare (2010) consider an *ensemble*

$$\mathbf{X}_t = (X_{t,1}, \dots, X_{t,n_c}), \tag{6}$$

of n_c individual *chains* $X_{t,i}$ (or *particles*, in SMC terminology) indexed by $i = 1, 2, \dots, n_c$ running in parallel at each iteration t . While \mathbf{X}_t holds the ensemble in parameter space, let

$$\mathbf{Z}_t = (Z_{t,1}, \dots, Z_{t,n_c}) = (f_b^{-1}(X_{t,1}), \dots, f_b^{-1}(X_{t,n_c})) \quad (7)$$

be its complementary representation in proposal space. As in Herbst and Schorfheide (2014), initialize the ensemble with n_c draws from the prior distribution

$$\mathbf{X}_0 \stackrel{n_c}{\sim} p(x). \quad (8)$$

Initializing the ensemble with the prior distribution ensures that the full set of prior information on the relevant parameter space is considered, independently of potential multimodality or possible discontinuities.¹⁰

A straightforward choice for the functional form of the bijective transform f_b is to chose $x_q = \exp(z_q) + \underline{b}$ for priors that are bounded below by \underline{b} (e.g. following a gamma and inverse gamma distribution), and the standard logistic function $x_q = \frac{\bar{b}-b}{1+\exp(-z_q)} + \underline{b}$ for priors with two-sided bounds (\underline{b}, \bar{b}) (e.g. the beta distribution).¹¹

2.3 Strategy mixture

In each iteration and for each chain the transition kernel is a mixtures of a local and a global proposals.¹² Let the global proposal kernel be selected with probability χ and, respectively, the local proposal kernel be chosen with probability $1 - \chi$.

Each iteration t of a DIME MCMC run then comprises:

1. Update the proposal distribution for the global transition kernel based on \mathbf{Z}_t .
2. To each chain i randomly assign a transition kernel $K_{t,i} \in \{G, L\}$ with probabilities $(\chi, 1 - \chi)$.
3. For each chain i , propose a replacement candidate vector $\hat{Z}_{t,i}$ based on the assigned transition kernel.
4. For each chain i calculate the factor weight $w_{t,i}$.
5. For each chain i , apply the bijective transform to obtain $\hat{X}_{t,i} = f_b^{-1}(\hat{Z}_{t,i})$.
6. For each chain i , evaluate the posterior density $\pi(\hat{X}_{t,i})$ of the candidate.
7. For each chain i , generate $Z_{t+1,i}$ by accepting $\hat{Z}_{t,i}$ with a Metropolis acceptance probability of

$$P(Z_{t+1,i} = \hat{Z}_{t,i}) = \min \left\{ 1, \frac{\pi(\hat{X}_{t,i})}{\pi(X_{t,i})} w_{t,i} \right\}, \quad i = 1, 2, \dots, n_c, \quad (9)$$

¹⁰Goodman and Weare (2010) suggest to initialize the ensemble as a small ball around some initial value. However, if the posterior is oddly shaped – e.g., if it is bimodal –, this bears the risk that the ensemble can not fully unfold. It may also be unclear which initial value to choose, in particular if we seek to avoid nonlinear optimization routines.

¹¹Another natural choice would be $x_q = \Phi^{-1}(F^q(z_q))$, where Φ^{-1} is the quantile function of the standard normal distribution and F^q is the CDF of prior q . This effectively transforms the prior to a multivariate Gaussian, which may be beneficial for the approximation of the proposal distribution of the global transition kernel.

¹²See Tierney (1994) for a theoretical discussion of mixture kernels. An alternative approach would be to draw the transition kernel not *per chain and iteration* but *for all chains per iteration*.

or reject $\hat{Z}_{t,i}$ and set $Z_{t+1,i} = Z_{t,i}$ with probability

$$P(Z_{t+1,i} = Z_{t,i}) = 1 - P(Z_{t+1,i} = \hat{Z}_{t,i}). \quad (10)$$

Section 4 investigates the question of the optimal number of chains n_c contra the number of iterations T in detail. As documented there, a value of $n_c \in (4n, 6n)$ is often a good choice, where larger ensembles help to tackle more irregular posterior distributions, as e.g. bimodal distributions, but fewer chains may speed up burn-in.

2.4 Local proposal kernel

The local kernel is *local* in the sense that for each chain j to which the local transition kernel is assigned, the candidate is proposed relative to the current state $X_{t,j}$ of j . At its core, the random-walk proposal distribution of RWMH is replaced with a proposal that follows the *differential evolution* concept of Ter Braak (2006).

More formally, for each iteration t and each chain $j : K_{t,j} = L$ draw two chains $\{k, l\} \in \{1, 2, \dots, n_c\}$ with $k \neq j$ and $l \neq k \wedge l \neq j$. Take the difference of the state of these two chains as a displacement vector which is added to the state $Z_{t,j}$ of chain j . The replacement candidate for chain $X_{t,j}$ is then $\hat{X}_{t,j} = f_b(\hat{Z}_{t,j})$ with

$$\hat{Z}_{t,j} = Z_{t,j} + \gamma(Z_{t,k} - Z_{t,l}) + \epsilon_{t,j}, \quad \forall j \in \{j : K_{t,j} = L\} \quad (11)$$

where γ is a scaling factor and $\epsilon_{t,i}$ is some (very) small noise.

As the ensemble evolves over time, proposal steps naturally adapt direction and scale of the current estimate of the posterior distribution. When the ensemble converges to the posterior distribution, so does the proposal distribution. Note that probability to draw the displacement vector $Z_{t,k} - Z_{t,l}$ is exactly equal to drawing the displacement vector $Z_{t,l} - Z_{t,k}$. Denoting the respective Metropolis-Hastings proposal distribution by g it thus holds that

$$g(\hat{Z}_{t,j} | Z_{t,j}) = g(Z_{t,j} | \hat{Z}_{t,j}) \quad (12)$$

and Equation (10) implies *detailed balance* for $w_{t,j} = 1$.

Theorem 1 in Ter Braak (2006) shows that the unique stationary distribution of DE-MCMC has the PDF $\pi(x)$. Intuitively, in the stationary distribution, the proposal distribution is the γ -scaled difference of two draws from the posterior distribution, which by itself is a stationary and symmetric distribution. This result on \mathbb{Z} applies one-to-one to the stationary distribution of \mathbb{X} . The interested reader is redirected to ter Braak and Vrugt (2008) for a complete proof that the stationary distribution of differential evolution MCMC equals the target distribution.¹³

From this intuition it also follows that, if $\pi(x)$ follows a Gaussian distribution, after convergence each individual proposal $\hat{X}_{t,j}$ is of the same form as an RWMH proposal.¹⁴ Under the assumption that the target distribution is near-Gaussian, the optimal choice

¹³As pointed out in ter Braak and Vrugt (2008), the original proof in Ter Braak (2006) contained an error.

¹⁴This can be seen by acknowledging that, if $\pi(x)$ is Gaussian, each draw $X_{t,j}$ is also Gaussian, and the difference between two chains hence also follows a Gaussian distribution.

for the scale γ is $\gamma = \frac{2.38}{\sqrt{2n}}$ from the RWMH literature (e.g. Roberts and Rosenthal, 2001), which is expected to give an acceptance probability of 23% for high-dimensional posteriors, i.e. for large n . Throughout this paper I set γ to this default value.

2.5 Global proposal kernel

In contrast to the local proposal kernel, the global proposal is *global* in so far as candidate proposal only depends on the global state of the ensemble (and its history), but not directly on the current state of a single chain.

For each chain j in iteration t with $K_{t,j} = G$ the displacement vector is drawn from an independent but adaptive proposal distribution. The distribution adapts such that it roughly corresponds to the current estimate of the posterior. A natural choice for such proposal distribution is the multivariate t -distribution with fixed degrees of freedom ν . This distribution is especially useful because for $\nu > 2$ it can be parameterized over its mean and covariance (μ_t, Σ_t) , and exhibits fat tails.¹⁵ In each iteration t the ensemble \mathbf{Z}_t is used to update the parameters (μ_t, Σ_t) with weights proportional to the average posterior density of \mathbf{Z}_t . This choice of weights ensures flexibility of the proposal distribution during burn-in but also stationarity after convergence.

More formally, define the *absolute weight* of the ensemble \mathbf{X}_t in iteration t on the proposal distribution as

$$w_t = a_t \sum_i^{n_c} \pi(X_{t,i}), \quad (13)$$

where $a_t = \frac{1}{n_c} \sum_i^{n_c} \mathbb{1}_{\{X_{t,i} \neq \hat{X}_{t-1,i}\}}(X_{t,i})$ is the mean acceptance ratio in t . Let W_t be the *cumulative weights* in t initialized with $W_0 = 0$. Denote by $(\mu_t^{\mathbf{Z}}, \Sigma_t^{\mathbf{Z}})$ the sample mean and sample covariance matrix of the current ensemble \mathbf{Z}_t . Then for each chain $j : K_{t,j} = G$ the proposal is given by

$$\hat{Z}_{t,j} \sim t_\nu \left(\mu_t, \frac{\nu - 2}{\nu} \Sigma_t \right), \quad (14)$$

with

$$\mu_t = \left(\frac{W_{t-1}}{W_t} \right) \mu_{t-1} + \left(\frac{w_t}{W_t} \right) \mu_t^{\mathbf{Z}}, \quad (15)$$

$$\Sigma_t = \left(\frac{W_{t-1}}{W_t} \right) \Sigma_{t-1} + \left(\frac{w_t}{W_t} \right) \Sigma_t^{\mathbf{Z}}, \quad (16)$$

$$W_t = W_{t-1} + w_t. \quad (17)$$

Note that the proposal density is independent of chain j . To again satisfy *detailed balance* in (10) set $w_{t,j} = \frac{f^t(Z_{t,j})}{f^t(\hat{Z}_{t,j})}$, where f^t is the density function of the multivariate t -distribution as defined in (14).

The idea behind this proposal is similar to Haario et al. (2001) but with the weighted updating of new ensembles. This has the strong advantage that during burn-in, newer

¹⁵The benefits of a fat tailed proposal distribution for adaptive independence MCMC is also pointed out by Holden et al. (2009).

updates have more weight than old ones and the proposal distribution adapts quickly to the current shape of the estimated target distribution. However, once the chains converge we have, for sufficiently large n_c , that $\sum_i^{n_c} \pi(X_{t,i}) \approx \sum_i^{n_c} \pi(X_{t+s,i})$ for $s = 1, 2, \dots$ and new ensembles have decaying weights.¹⁶ This implies that the procedure converges to the specification in Haario et al. (2001) and enjoys the same convergence properties therein. See Roberts and Rosenthal (2007) for a detailed proof of adaptive independence Metropolis Hastings.

A natural choice for the degrees of freedom ν of the multivariate t distribution is to pick rather low values, which imply fatter tails of the proposal distribution. All results of this paper are rather insensitive to the choice of ν , and throughout the following sections I use $\nu = 10$. While the above specification of the mean/covariance updating could be tweaked with a number of additional parameters (e.g. a tempering scheme for the probability weights), this is unnecessary in practice. If a researcher wishes to replace *less* chains it is sufficient to decrease the probability χ for the global transition kernel and vice versa. This leaves the specification of the global transition kernel essentially parameter free, and in total requires the user only to specify n_c and χ as the necessary parameters for DIME MCMC. Throughout this paper I set $\chi = 0.1$, which provides a good compromise between fast burn-in and robustness for multimodal distributions.

3 A high-dimensional bimodal toy distribution

This section studies the performance of DIME MCMC on a distribution with known properties. I focus on a class of high dimensional bimodal distributions where the two modes may be disconnected and can have different density masses. Such distributions are known to be challenging for MCMC samplers. The probability density of the random variable M is given by the multivariate Gaussian mixture

$$\pi(x) = \lambda P(X = x) + (1 - \lambda)P(Y = x) \quad (18)$$

where $X \sim \mathcal{N}_n(\mu_X, \sigma I_n)$ and $Y \sim \mathcal{N}_n(\mu_Y, \sigma I_n)$ are both n -dimensional Gaussian distributions with the same covariance, which is the identity matrix scaled by the scalar $\sigma > 0$. $\lambda \in (0, 1)$ is a weighting parameter and $\mu_X = (m/2, 0, \dots, 0)'$ and $\mu_Y = (-m/2, 0, \dots, 0)'$ are both vectors of zeroes apart from the first entries, which are $m/2$ and $-m/2$ respectively. The distribution of M is then bimodal whenever $m \neq 0$ and the distance between the two modes is given by $|m|$. When keeping σ fix, increasing m complicates Monte Carlo sampling because the modes are less connected. Corresponding with the typical size of a target distribution when estimating medium-scale DSGE models, let M be in $n = 35$ dimensions.¹⁷

Figure 1 illustrates this exercise graphically by marginalizing over the first dimension. The shaded areas mark the 2.5%-percentile and the median of the first dimension. Each ensemble is initialized with a sample from $\mathcal{N}_n(\mathbf{0}_n, \sqrt{2}I_n)$. The initial ensemble is hence distributed across the domain of M , with relatively more chains closer to the origin

¹⁶Note that it is not an actual requirement that the cumulative density is approximately equal across ensembles. Σ_t converges even if $\sum_i^{n_c} \pi(X_{t,i})$ varies a lot as long as it is stationary.

¹⁷The posterior of the model of Smets and Wouters (2007) has 36 dimensions while the posterior of the HANK model estimated in Section 5 has 31 dimensions.

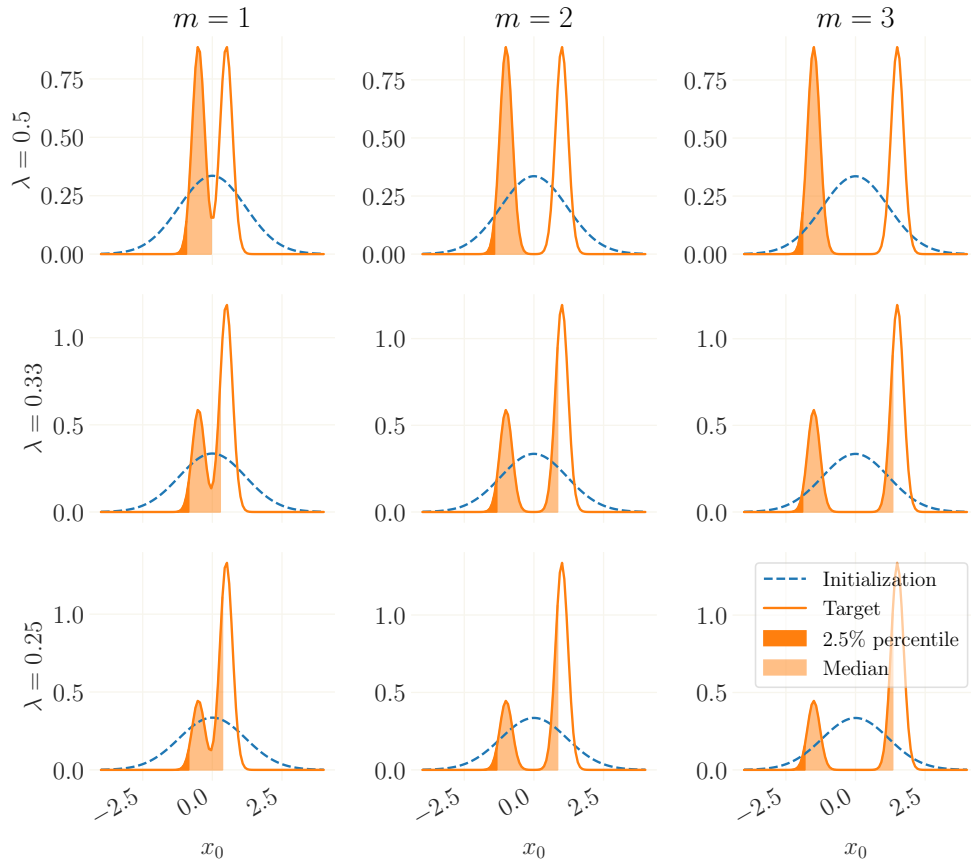


Figure 1: A 35 dimensional multivariate Gaussian mixture, marginalized over the first dimension (orange line). The dashed line depicts the initialization distribution of the ensemble. The frontier between the dark orange and light orange shaded area is the 2.5%-percentile and the frontier between the light shaded area and no shade is the median of the distribution.

(dashed blue line in Figure 1). Calculations are done for $\sigma = 0.05$ and distances of $m \in \{1, 2, 3\}$ (the columns of figure 1).¹⁸ The first row shows the target distribution for $\lambda = 0.5$ where both modes peak at the same maximum density. For $m = 1$ both modes are connected, meaning that for any point between the modes the density is still reasonably large (that is, larger than 0.1 for the cases considered here). For $m = 2$ the trough between the modes is relatively short in distance, but the minimum density is already close to zero. The gap for which the density is zero again increases considerably when setting $m = 3$, for which the modes are fully disconnected. The challenge for MCMC sampling lies in the fact that the chains must be able to bridge this gap, which for conventional samplers is unlikely once the density is close to zero.

¹⁸These values are chosen to demonstrate the frontier of what is possible with DIME, without additional adjustments of the algorithm.

	$m = 1$		$m = 2$		$m = 3$	
	2.5% HDI	median	2.5% HDI	median	2.5% HDI	median
$\lambda = 0.5$	0.00827	0.08253	0.00960	0.58256	0.01239	1.08592
$\lambda = 0.33$	0.00946	0.01337	0.01004	0.01709	0.01453	0.02222
$\lambda = 0.25$	0.01253	0.00944	0.01308	0.01148	0.01897	0.01592

Table 1: RMSEs of the estimated 2.5%-percentile and the median of the first dimension of the target distributions. Results obtained from 100 batches.

For each of the nine exercises, I conduct 100 batches of 210 chains each (correspondingly, $n_c = 6n$), let each batch run for 2000 iterations with $\chi = 0.1$, and then calculate the 2.5%-percentile and the median over the first dimension. Table 1 presents the root mean squared errors (RMSE) of these target measures over all batches. As the table suggests, DIME MCMC performs very well over all nine exercises. Even when M is fully disconnected ($m = 3$), the sampling error only increases marginally. The only exception is the estimate of the median for the first row where $\lambda = 0.5$. This finding can, however, safely be ignored: for $m = 2$ and $m = 3$ the posterior density in the gap between the two modes is almost zero because both modes have the exact same density masses. Correspondingly, this region contains no chains and a precise quantification of the median is impossible unless we let the number of iterations go to infinity.

For the simulations in the second and third row of Figure 1 and Table 1, λ is set to 0.33 and 0.25, respectively. This example is more challenging because some chains must “jump” between the modes in order to correctly reflect the different density masses of the modes. In practice any single-particle sampler is likely to “get stuck” in either of the mode, thereby ultimately misrepresents the posterior distribution. Yet, also for this example RMSEs are very small and acceptance ratios are in the desired range between 20-25%. This success crucially depends on the global transition kernel, which allows to reshuffle chains between the two modes. Consequently, DIME performs less good if the probability of drawing the global transition kernel χ is set larger than 20%. In that case, too many chains are reshuffled too early, thereby causing estimates of the proposal distribution to ignore the second mode. This corroborates the previous recommendation of setting $\chi = 0.1$ for black-box distributions. For this setup, DIME MCMC seems to be able to reliably sample from high-dimensional and bimodal distributions, even if the modes are fully disconnected.

4 The Smets-Wouters model

A common benchmark case for the Bayesian estimation of DSGE models is the work of Smets and Wouters (2007, henceforth SW), who pioneered the use of Bayesian methods for bringing medium-scale DSGE models to the data. I use this prominent reference in three exercises. First, I assess whether DIME MCMC is able to exactly recover the posterior distribution from the original paper. Secondly, I use the model of SW to benchmark DIME MCMC against two other ensemble MCMC samplers as well as against the modified adaptive independence Metropolis-Hastings proposal from the global transition kernel. Lastly, I use their model to evaluate the trade-off of more chains versus longer chains.

For each of the exercises exactly the same model specification, priors, data and data treatment as in the original paper are used. All estimations are done on a workstation with 40 Intel Xeon CPUs with 3.1GHz each and a total of 32GB RAM. I use the package *pydsge* for model parsing and solving, and to calculate the likelihood using the standard Kalman filter.¹⁹

4.1 Comparison with the original estimates

To reproduce the estimation from SW I let an ensemble of 200 chains run for 3000 iterations, of which 500 are kept as the posterior. The original estimation relies on 250.000 samples (of which 50.000 are discarded) obtained using RWMH after running an optimization procedure from pre-optimized starting values. Table A.4 in Appendix A shows summary statistics over the posterior distribution of the estimation together with posterior statistics from SW. Overall, the DIME MCMC estimates and the posterior values from the original estimation of SW are very closely aligned. Notable differences are the estimate of the standard deviation of the risk premium shock, σ_u , which is substantially larger than the SW estimate, as well as in the estimate of the steady state labor supply \bar{l} . Judging from the standard deviation of the latter estimate, the parameter seems not well identified. In summary, the estimates indicates that the DIME MCMC can fully recover the original values of SW. The table also shows the marginals over the proposal distribution of the global transition kernel which, in terms of mean and standard deviation and as expected, is very closely aligned to the posterior distribution.

4.2 Comparison of different Ensemble MCMC samplers

Turn now to the comparison of the performance of DIME MCMC with a selection of alternative (ensemble) MCMC samplers. I let each sampler run ten times over different random seeds. For each sampler and seed I chose the same initialization and, again, let 200 chains run for 3000 periods. To allow for fair comparison, the bijective mapping between proposal and parameter space is applied for *all* samplers. Figure 2 plots the log-density of each *single* chain over time. The dashed lines mark the mode i.e., the maximum posterior density value. The different colors in each panel correspond to different ensemble batch runs. The top-left panel plots the batches over the DIME sampler. Chains converge quickly towards the high density region of the posterior, reaching the 68% set roughly in period 300 and the 97.5% set in about iteration 500. Although convergence is difficult to asses, it seems as if all DIME chains across all ensembles have converged to the posterior roughly at about iteration 700. Throughout the convergence period the single chains remain relatively close to each other, both within and across ensembles.

The panel at the top-right plots the performance of the “Stretch” move of Goodman and Weare (2010). This proposal kernel is quite popular in the field of astrophysics. Across batches, initial convergence of the ensembles is relatively rapid, but convergence then slows down. The 68% and 97.5% sets are reached between iterations 1500 and 2000 and after 2250, respectively. Correspondingly, the chains do not converge to the

¹⁹Pydsge is a toolbox to solve, filter, and estimate DSGE models in Python language, which is presented in Boehl and Strobel (2022a). The package is available in the official Python repositories and developed and maintained at GitHub: <https://github.com/gboehl/pydsge>.




Figure 2: Using different ensemble MCMC methods to estimate the model of Smets and Wouters (2007). Each panel shows the course of the log-likelihood of several ensembles over time, using different Ensemble MCMC methods. Each colors represent a different ensemble with different random seed. For each ensemble all individual chains are plotted. The scaling of all panels is the same.

posterior before iteration 2500. Convergence behavior differs slightly across batches. The bottom-left panel in Figure 2 shows ensembles following the differential evolution MCMC (DE-MCMC) method of Ter Braak (2006) in the implementation of Foreman-Mackey et al. (2013). The graphic suggests that burn-in for DE-MCMC is slow and the ensemble does not converge to the prior distribution within the given 3000 iterations. An apparent problem seems to be that dispersion in log-density across chains in each ensemble is very large. The likely reason is that each chain moves ex-ante independently, i.e. state of the complete ensemble is only used for relative repositioning of each chain. When ensemble dispersion is high, the quality of replacement proposals deteriorates and convergence slows down even further, thereby causing single chains to converge extremely slowly. The bottom-right panel plots the dynamics of an ensemble following the global transition kernel taken alone, i.e. without mixing with the local transition kernel. As apparent from the figure, the adaptive independence Metropolis-Hastings approach *alone* performs quite badly and convergence is very slow. Performance during burn-in varies much over random seeds.

DIME MCMC is a mixture kernel of DE-MCMC and the global transition kernel. Figure 2 clearly illustrates their individual weaknesses. The DE-MCMC ensemble is

overdispersed, which causes unfavorable individual proposals and, in turn slow converge. In contrast, the ensemble of the global transition kernel is very narrow because chains with a higher probability density have a larger weight in the proposal distribution. Thus, candidate proposals will lie in the immediate proximity of the current ensemble, again causing slow convergence. However, since these individual weaknesses cancel out both kernels are strongly complementary: when mixing the two kernels, the “fairly good” proposals from the global transition kernel are sufficient to reshuffle chains that would otherwise (i.e. with DE-MCMC alone) be stuck in regions with lower probability density.²⁰ This reshuffling is efficient to decrease the dispersion of the ensemble and, consequently, the proposals of the local transition kernel improve, which helps to explore the neighborhood of the current state of the ensemble well.

4.3 The number of chains n_c

Next, let me benchmark the sensitivity of the estimation results with respect to the number of chains n_c . Figure 3 illustrates burn-in speed and convergence dynamics in terms of the number of total function evaluations. For the chosen range of $n_c \in (4n, 6n)$ it seems that no setup emerges which is to be strongly preferred. I start with $n_c = 2n$, which is the minimum number of chains suggested by Foreman-Mackey et al. (2013). As depicted in top-left panel, convergence is slower than for a larger number of chains and the course of the different ensembles shows larger variation. For more chains ($n_c = 4n$, top-right panel) convergence is faster, with no significant difference to $n_c = 6n$ in the bottom-left panel. For larger ensembles ($n_c = 8n$, bottom-right) convergence per function iteration is again marginally slower whereas individual ensembles are almost indistinguishable.

This exercise reveals a mild trade-off between the number of iterations and the quality of the proposal candidates. For only few chains per ensemble, each iteration requires only few function evaluations. However, the relatively small number of chains produces less favorable replacement proposals. When, in contrast, ensembles are large, each iteration is relatively costly and fewer iterations can be done for a given number of function evaluations. However, for a large range of n_c between $4n$ and $6n$ a larger number of chains approximately compensates one-to-one for fewer iterations. For this range, not the number of iterations is central, but the total number of function evaluations across chains. Importantly, this suggests that estimations can be scaled very well when parallelizing chains using computers with a larger number of processing units. An increase in the number of chains does always require less chain iterations. Even if this relationship would not be one-to-one, this implies that it is advisable to use at least as many chains as numbers of processing units.

Appendix B provides Gelman and Rubin (1992) statistics over the number of chains. The Gelman-Rubin statistic is a measure of convergence. The results substantiate the finding that the optimal number of chains lie between $6n$ and $8n$, with a lower number of iterations (i.e., a larger number of chains) causing higher Gelman-Rubin coefficients. This indicates that it is more important to run many iterations than to run a large number of chains. However, average integrated autocorrelation times (e.g. Sokal, 1997)

²⁰Following a similar intuition, Vrugt et al. (2009) use the inter-quartile range to discover potential outlier chains, which are then replaced with the current best chain.




Figure 3: Using DIME MCMC to estimate the model of Smets and Wouters (2007). Each panel shows the course of the log-likelihood of several ensembles over time, using different numbers n_c of chains. Each colors represent a different ensemble with different random seed. For each ensemble all individual chains are plotted. The scaling of all panels is the same.

across chains, parameters and different ensemble sizes are relatively constant around 40. This in turn suggests again that the ensemble size does not have a major influence on sampling quality.

Overall I suggest to chose n_c to be a multiple of the number of available processors which lies in the range $5n$ and $6n$, and to set $\chi = 10\%$. It is advisable to monitor the trace plot of the likelihood function (e.g. as in Figure 3) and the histogram of the posterior. For very rugged or multimodal distributions the number of chains should be increased. In such cases it is additionally expedient to decrease χ to prevent chains from getting reshuffled too early. In contrast, if for some reason the posterior is expected to be rather near-Gaussian, a larger value of χ can be chosen, which will decrease autocorrelation times and hence requires fewer ensemble iterations during the sampling stage, i.e. after burn-in.

5 Full estimation of HANK

To explore the full potential of DIME MCMC I use the sampler to estimate a Heterogeneous-Agent New Keynesian (HANK) model with portfolio choice and all the

features of a standard medium-scale DSGE model. The central novelty relative to the literature is that I include the households' preference parameters in the set of estimated parameters, which increases the complexity of the calculations significantly.

HANK models are a relatively new class of models (see, e.g., Gornemann et al. (2012) and Kaplan et al. (2018a)) that combine the New Keynesian paradigm with household heterogeneity and incomplete financial markets. This allows, for example, to study the impact of economic inequality on macroeconomic aggregates and vice versa. While the estimation of HANK models is pioneered by Winberry (2018), Bayer et al. (2020, henceforth BBL) and Auclert et al. (2021), these papers do not estimate the parameters of households' preferences that govern the households' optimization problem. The reason for excluding these parameters is that they alter the model's steady state, which would then have to be re-evaluated for every posterior draw. As the re-evaluation of the steady state involves finding a stationary distribution such that all equilibrium conditions are satisfied, this comes at large computational costs. Consequently, BBL and Auclert et al. (2021) both opt to calibrate all parameters which affect the model's steady state and focus on estimating the remaining parameters.

However, the households' preference parameters have the potential to form central attributes of the ergodic distribution of assets and income, and may thus qualitatively and quantitatively determine the magnitude of the novel channels exposed by this class of models. Hence, these parameters could potentially affect the macroeconomic dynamics of this class of models fundamentally. Since at the same time, their inclusion in the estimation is computationally expensive – finding the steady state and the stationary distribution takes about 10 seconds for the implementation considered here – it is a perfect use case for the DIME sampler.

5.1 Model and Data

The model is the fusion of a two-asset HANK model with a medium-scale DSGE model. The HANK core shares many features with the model of Auclert et al. (2021) and Kaplan et al. (2018b). This core is extended by several frictions in the spirit of Smets and Wouters (2007), which, among other features, allow for additional endogenous persistence in response to aggregate shocks.²¹ To ease comparison with the DSGE literature I use the priors of Smets and Wouters (2007). Accordingly, some of the functional forms (e.g. capital adjustment costs and Calvo pricing) are adapted from there. In the following I discuss only those equations that deviate from Auclert et al. (2021) and refer the interested reader to Appendix C for further details on the model.

Households supply labor and have access to a liquid and an illiquid asset. Importantly, they face borrowing constraints on both assets, and adjustment costs on the illiquid asset. Firms accumulate capital, and staggered price setting results in a conventional Phillips curve. Adding ad-hoc price indexation with parameter ι_p , inflation π_t is determined by

$$\pi_t - \bar{\pi} = \frac{\beta}{1 + \beta\iota_p} (E_t\pi_{t+1} - \bar{\pi}) + \frac{\iota_p}{1 + \beta\iota_p} (\pi_{t-1} - \bar{\pi}) + \kappa_p \left(\widehat{MC}_t - \frac{1}{\mu} \right) + \epsilon_{p,t}, \quad (19)$$

²¹It is well known that endogenous persistence is a crucial feature to replicate the hump-shaped empirical responses that are reported in the VAR literature.

where $\bar{\pi}$ is the steady state inflation. $\epsilon_{p,t}$ is assumed to follow an AR(1) process around its zero mean and the slope of the Phillips curve is given by $\kappa_p = \frac{1-\zeta_p\beta}{1+\iota_p\beta} \frac{1-\zeta_p}{\zeta_p}$. Labor unions set nominal wages which are also subject to staggered pricing, which gives rise to a Phillips curve for wages. Adding wage indexation with parameter ι_w , this yields

$$\begin{aligned} \pi_t^w - \bar{\pi} &= \frac{\beta}{1 + \beta\iota_w} (E_t \pi_{t+1}^w - \bar{\pi}) + \frac{\iota_w}{1 + \beta\iota_w} (\pi_{t-1}^w - \bar{\pi}) \\ &\quad + \kappa_w \left(\varphi N_t^{1+\nu} - \frac{(1 - \tau_t)w_t N_t}{\mu_t^w} \int e_{it} c_i t^{-\sigma} di \right) + \epsilon_{w,t}, \end{aligned} \quad (20)$$

where $\epsilon_{w,t}$ as well follows an AR(1) process and $\kappa_w = \frac{1-\zeta_w\beta}{1+\iota_w\beta} \frac{1-\zeta_w}{\zeta_w}$. Monetary policy sets the nominal interest rate r_t following a conventional monetary policy rule,

$$r_t^n - r^n = \rho (r_{t-1}^n - r^n) + (1 - \rho) [\phi_\pi (\pi_t - \bar{\pi}) + \phi_y \Delta \ln Y_t] + \epsilon_{r,t}, \quad (21)$$

with $\epsilon_{r,t}$ as an exogenous AR(1) process representing monetary policy surprises. Note that in order to remain agnostic about the central bank's welfare objective, a traditional measure of output gap is absent in this equation. The setup of capital adjustment costs is as in Smets and Wouters (2007) and yields the following expressions for Tobin's Q and the firm's investment decisions

$$R_{t+1}q_t = (1 - \delta)E_t q_{t+1} + \alpha E_t \left\{ Z_{t+1} \frac{N_{t+1}}{K_t} {}^{1-\alpha} \widehat{MC}_{t+1} \right\} \quad (22)$$

$$1 = \exp(\epsilon_{i,t})q_t \left[1 - S \left(\frac{I_t}{I_{t-1}} \right) - S' \left(\frac{I_t}{I_{t-1}} \right) \frac{I_t}{I_{t-1}} \right] + E_t \left\{ \exp(\epsilon_{i,t+1}) \frac{q_{t+1}}{R_{t+1}} S' \left(\frac{I_{t+1}}{I_t} \right) \left(\frac{I_{t+1}}{I_t} \right)^2 \right\}, \quad (23)$$

where R_t is the gross real interest rate on liquid assets, $S(x) = \frac{1}{2S''}(x-1)^2$ is a quadratic adjustment cost function, and $\epsilon_{i,t}$ is an exogenous AR(1) process on the marginal productivity of investment. Finally, labor income taxation is progressive with parameter Ξ such that after-tax labor income y_{jt} is given by

$$y_{jt} = y_{jt}^p {}^{1-\Xi} + \int p(e_{jt}) \left(y_{jt}^p - y_{jt}^p {}^{1-\Xi} \right), \quad (24)$$

with pretax income $y_{jt}^p = (1 - \tau_t)w_t N_t e_{it}$.

For the estimation I use a subset of the data used in Boehl et al. (forthcoming) which amounts to a relatively conventional setup for medium scale models: growth rates of consumption, investment, output and wages, together with inflation, labor hours and the federal funds rate. The data is at quarterly frequency and ranges from 1983:I to 2008:IV. As in Justiniano et al. (2010), investment and consumption time series are adjusted such that investment also includes durables consumption. In the model, those seven observables are matched by seven economic shocks, which are all defined in percentage deviations from the steady state: the two markup shocks, the monetary policy shock, a government spending shock on G_t , a discount factor shock on β_t and the shock on the marginal efficiency of investment, $\epsilon_{i,t}$. Further details can be found in Appendix D.

5.2 Estimation methodology

Model solution and likelihood inference is done following the methodology introduced in Auclert et al. (2021).²² In brief, let y_t be the time t vector of model variables (including disaggregated variables) and let the sequence of first-order conditions and market clearing conditions, up to some distant point T periods in the future, be

$$F = \{f(y_{t-1}, y_t, E_t y_{t+1}; x)\}_{t=0}^T = 0, \quad (25)$$

which depends on the parameter vector x . Denote by $Y_t \subset y_t$ only the aggregated variables and by $Z_t \subset y_t$ those variables that are purely exogenous. The authors propose a novel and computationally efficient procedure of finding the steady state Jacobian matrix of F with respect to $\{Y_t\}_{t=0}^T$ and $\{Z_t\}_{t=0}^T$. These sequence-space Jacobians (SSJ) can then be used to calculate impulse responses to aggregate shocks up to a first order approximation. Notably, this works for the broad class of models for which it is not required to explicitly keep track of any of the disaggregated distribution variables on a global domain. Simulations are based on the sequence space rather than, as in BBL, the *state space*. The authors show that the first-order sequence space representation can be used directly for likelihood inference, without the need for using the Kalman filter (which would require a state space representation). In their application, the authors are able to re-use (parts of) the Jacobians depending on the types of parameters to be estimated. In contrast, in my application each Jacobian has to be calculated from scratch due to the re-calculation of the steady state.

In a deterministic setup, the steady state \bar{y} must satisfy

$$f(\bar{y}, \bar{y}, \bar{y}; x) = 0. \quad (26)$$

Given a guess for the steady state values of aggregated variables \bar{Y} , the stationary distribution of idiosyncratic variables can be found by solving for the stationary decision rules via backward iteration, and solving for the stationary distribution via forward iteration. Hence, there exists a known mapping $\bar{Y} \rightarrow \bar{y}$, and finding \bar{Y} can be done using conventional root finding methods. Often, the size of this root finding problem can further be reduced to only searching a subset $\bar{K} \subset \bar{Y}$ since \bar{Y} can be expressed in terms of this subset. Still, finding \bar{y} is relatively time consuming and must be repeated for any parameter draw x if the households' micro parameters change.²³

5.3 Estimation results

As usual, some parameters are fixed prior to the estimation. Most of these parameters configure the technical setup of the estimation (e.g. the number of grid points), and can be found in Table C.6 in Appendix C. All other parameters are estimated using the priors

²²The authors provide their set of methods as a Python toolbox maintained at GitHub: <https://github.com/shade-econ/sequence-jacobian>.

²³For any numerical root finding method a good initial guess is crucial. This also holds for finding the steady state. For bad initial guesses, the root search may either diverge, crash due to numerical errors when solving for the stationary distribution, or simply take up a very long time. This is problematic because it also prohibits the calculation of the likelihood for cases in which a likelihood actually exists. In practice, for every draw I use the steady state values for the prior mean as the initial guess, which causes about 2/3 of all parameter vectors sampled from the prior distribution to be accepted.

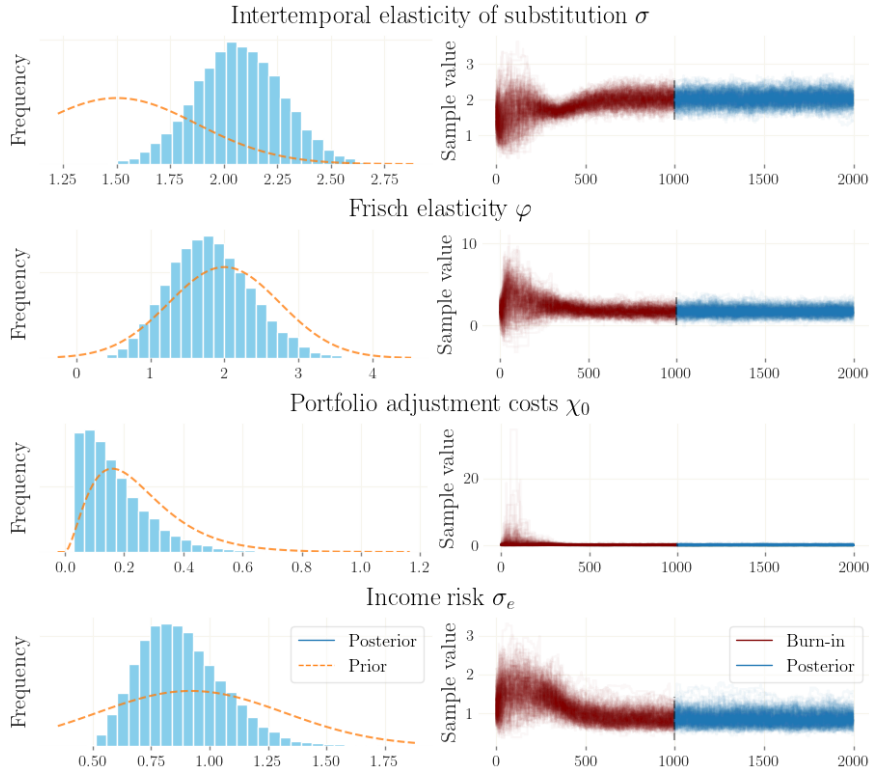


Figure 4: Posterior histogram and prior distribution (left side, blue and orange) of selected households' parameters with trace plots of the ensemble over iterations (right side). The part colored in red is discarded as burn-in.

presented in the first three columns in Tables 2 and 3, which follow the specification of Smets and Wouters (2007). Exceptions are the portfolio adjustment cost parameter χ_0 , tax progressively parameter Ξ , and the standard deviation of the AR(1) process for idiosyncratic labor productivity σ^e , which are specific to the HANK model. For these parameters I opt for generally flat priors. I let the prior mean of $\sigma^e = 0.92$, which is the values used in Auclert et al. (2021). For the same reason, the prior mean of χ_0 is set to 0.25.

For the estimation I run a DIME MCMC ensemble with a total of $n_c = 192$ chains for 2000 iterations. The last 1000 iterations are kept as a sample from the posterior. The number of chains is the number of available CPUs (48) times 4 and corresponds with $n_c/n \approx 5.33$, which lies in the range recommended in the previous section. The ensemble converges to the high-density region of the posterior after about 800 iterations and the full estimation takes 84 hours on the machine with 48 cores. Note that doubling the number of cores would result in half the estimation time. Respectively, on a machine with 192 cores each chain would have a dedicated processor and the estimation would take a bit less than a day. Using a machine with more cores requires a larger number of chains, which would still scale close-to one-to-one (see Section 4). The figures F.6 and

		Prior			Posterior		
		distribution	mean	std.	mean	std.	mode
σ	intertemporal elasticity of substitution	normal	1.500	0.375	2.043	0.202	1.850
φ	Frisch elasticity	normal	2.000	0.750	1.738	0.562	1.805
ζ_p	Calvo parameter for price setting	beta	0.500	0.100	0.590	0.050	0.592
ζ_w	Calvo parameter of wage setting	beta	0.500	0.100	0.416	0.069	0.413
ι_p	price inertia	beta	0.500	0.150	0.331	0.129	0.335
ι_w	wage inertia	beta	0.500	0.150	0.322	0.147	0.303
S''	derivative capital adjustment costs	gamma	4.000	2.000	2.279	0.702	1.725
ϕ_π	monetary policy coefficient inflation	gamma	1.500	0.250	2.322	0.219	2.198
ϕ_y	monetary policy coefficient output	gamma	0.125	0.050	0.222	0.063	0.205
ρ	monetary policy persistence	beta	0.750	0.100	0.652	0.052	0.680
\bar{y}	trend output	normal	0.400	0.100	0.438	0.026	0.434
\bar{n}	steady state labor hours	normal	0.000	2.000	-0.047	1.961	1.469
π^*	inflation target	gamma	0.625	0.100	0.596	0.051	0.624
i^*	steady state nominal interest rate	gamma	1.250	0.100	1.239	0.089	1.259
χ_0	portfolio adjustment costs (scale)	gamma	0.250	0.150	0.153	0.118	0.094
Ξ	tax progressivity	beta	0.200	0.100	0.089	0.059	0.071
σ^e	standard deviation of labor productivity	normal	0.920	0.400	0.860	0.185	1.064

Table 2: Estimation results for HANK: model parameters

F.7 to F.11 in Appendix F graphically illustrate convergence of the ensemble. Tables 2 and 3 show summary statistics of the posterior distribution. Figure 5 shows impulse response functions to a monetary policy and a TFP shock of the estimated model.

This paper focusses on the performance of the DIME sampler instead of the economic dynamics of the estimated HANK model. For this reason I deem an in-debt analysis of the economic implications of the estimated model out of the scope of this paper and leave it as a promising endeavour for future research. Nevertheless, a cursory comparison of the parameter estimates from the HANK model with those of Smets and Wouters (2007) – for a somewhat smaller sample – reveals some surprising differences.²⁴ In HANK, the inverse elasticity of substitution, σ , is relatively large. This differs to the estimate of SW and the findings documented in Boehl and Strobel (2022b,a) for US data until 2019, who report values close-to unity. This estimate is likely to be related to the fact that HANK model features the additional the precautionary savings channel due to the assumption of incomplete financial markets.

An interesting finding is that in the HANK model both the price and the wage Phillips curve are identified to be relatively steep, which is reflected by Calvo adjustment probabilities ζ_p and ζ_w to be estimated relatively low. This stands in contrast to many more recent estimates which find rather large values for these parameters, which suggests a very flat Phillips curve. While this effect may come from different data samples and slightly different specifications of the Phillips curves, it calls for further investigation. The relatively lower estimate of S'' is also documented in BBL and may indicate that capital adjustment costs play a smaller role in the HANK model, which may be due to the fact that in the HANK model, portfolio adjustment represent a additional friction that actively influences the capital investment decision. The other parameters in Table 2,

²⁴The estimation of Smets and Wouters (2007) is replicated in Appendix A.

		Prior			Posterior		
		distribution	mean	std.	mean	std.	mode
ρ_z	AR coefficient technology shock	beta	0.500	0.200	0.957	0.018	0.960
ρ_r	AR coefficient MP shock	beta	0.500	0.200	0.619	0.070	0.640
ρ_g	AR coefficient gov. spending shock	beta	0.500	0.200	0.993	0.006	0.993
ρ_w	AR coefficient wage MU shock	beta	0.500	0.200	0.985	0.006	0.980
ρ_p	AR coefficient price MU shock	beta	0.500	0.200	0.911	0.028	0.917
ρ_i	AR coefficient investment shock	beta	0.500	0.200	0.837	0.042	0.836
ρ_β	AR coefficient interest wedge shock	beta	0.500	0.200	0.962	0.030	0.991
σ_z	standard dev. technology shock	inv.gamma	0.100	0.250	0.415	0.036	0.430
σ_r	standard dev. MP shock	inv.gamma	0.100	0.250	0.130	0.020	0.113
σ_g	standard dev. gov. spending shock	inv.gamma	0.100	0.250	1.148	0.088	1.105
σ_w	standard dev. wage MU shock	inv.gamma	0.100	0.250	2.662	0.732	2.448
σ_p	standard dev. price MU shock	inv.gamma	0.100	0.250	0.201	0.047	0.188
σ_i	standard dev. investment shock	inv.gamma	0.100	0.250	1.289	0.215	1.350
σ_β	standard dev. interest wedge shock	inv.gamma	0.100	0.250	0.047	0.017	0.029

Table 3: Estimation results for HANK: parameters of exogenous processes

which govern the monetary policy rule and the steady state values of the observables, are well-aligned with the original estimates in SW. These parameters are likely identified independently of the model’s setup of the household sector.

The estimate of the portfolio adjustment cost parameter χ_0 is well below its respective prior mean, pointing towards a less accentuated role of the households’ portfolio choice problem. Complementary, the standard deviation of the idiosyncratic labor productivity, σ_z , is also slightly below to its prior value. Both of these values correspond to the parameters chosen by Kaplan et al. (2018a) rather than those of Auclert et al. (2021). While by no means this evidence can be used to evaluate the role of idiosyncratic income risk or portfolio choice, it also calls for further investigation. Lastly, the estimates of the parameters that govern the exogenous autoregressive processes are much in line with conventional estimates, where technology, government spending and investment specific shocks are usually highly autocorrelated.

In Appendix E I repeat the estimation but letting the households’ state space being represented on a smaller grid (480 nodes instead to 2625 nodes). This reduction cuts the estimation time about one-third to 50 hours in total. As the reported estimates suggest, the reduction in the number of approximation nodes does not have a significant impact on the estimation results. Consequently, it may be possible to obtain reliable results from using a smaller representation of the households’ state space.

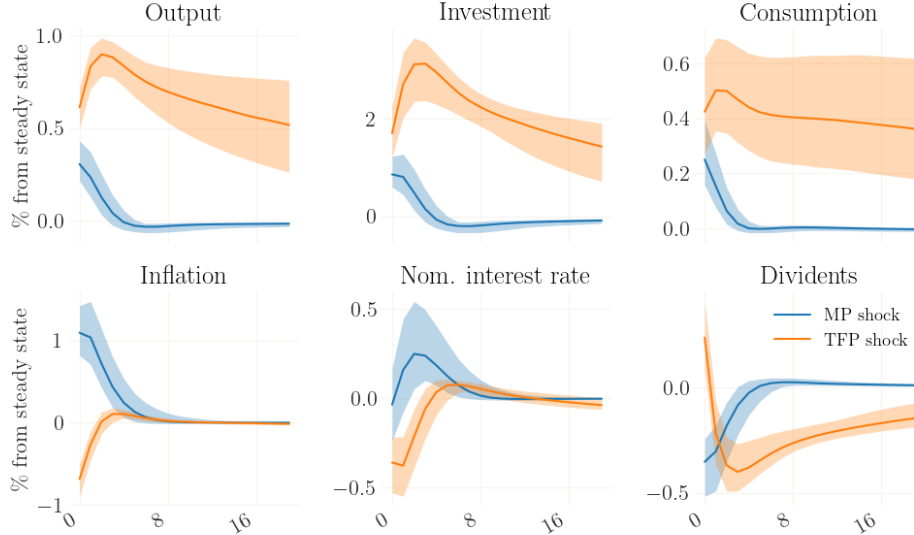


Figure 5: Impulse response functions to a monetary policy shock (blue) and a shock to TFP (orange). Responses and credible sets correspond to 1000 simulations drawn from the posterior distribution. The measures are annualized where applicable.

6 Conclusion

This paper develops a differential-independence mixture ensemble (DIME) MCMC sampler. The sampler is a Swiss Army knife that can be used for posterior sampling and global optimization problems alike. I show that the method perform well for high-dimensional and multimodal distributions. The proposal density of DIME is generated endogenously from the state of an ensemble of many chains, thereby automatically adapting to the shape of the current estimate of the posterior distribution. A separation of parameter space and proposal space guarantees that proposals respect the bounds of prior distribution, which results in significantly higher acceptance rates and, consequently, in higher sampling efficiency.

Mixing between local and global proposal leads to very fast burn-in and convergence to the high density region of the posterior. I show that DIME MCMC is easy to parallelize, where the number of iterations required until convergence decreases almost one-to-one with the number of chains. This makes the method feasible for large-scale problems with models that are computationally expensive to simulate.

The DIME sampler allows, for the first time, to include the households' micro parameters when estimating a HANK model with portfolio choice. These parameters strongly affect the households' decision problem and, thereby, determine the endogenous distribution of assets. The detailed analysis of the estimated model, e.g. by putting the resulting parameter estimates in relation to estimates from micro data, is a promising endeavour for future research.

A natural extension to DIME, also for future research, is to replace the differential-evolution proposal in the local transition kernel by a HMC proposal for applications in

which automatic differentiation is feasible. In such setting, HMC supersedes differential evolution MCMC: if the Jacobian can be evaluated at low computational costs, proposals can readily be well-adopted to the actual shape of the posterior. Yet, the mixture with the global transition kernel would remain powerful as it can speed up burn-in and enables sampling from very challenging multimodal distributions.

References

- Auclert, A., Bardóczy, B., Rognlie, M., Straub, L., 2021. Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica* 89, 2375–2408.
- Bayer, C., Born, B., Luettticke, R., 2020. Shocks, Frictions, and Inequality in US Business Cycles. CEPR Discussion Papers 14364.
- Boehl, G., Goy, G., Strobel, F., forthcoming. A structural investigation of quantitative easing. *Review of Economics and Statistics* .
- Boehl, G., Strobel, F., 2020. US business cycle dynamics at the zero lower bound. Bundesbank Discussion Papers 65/2020. Deutsche Bundesbank.
- Boehl, G., Strobel, F., 2022a. Estimation of DSGE Models with the Effective Lower Bound. CRC 224 Discussion Papers. University of Bonn and University of Mannheim, Germany.
- Boehl, G., Strobel, F., 2022b. The Empirical Performance of Financial Frictions Since 2008. CRC 224 Discussion Papers. University of Bonn and University of Mannheim, Germany.
- Bognanni, M., Herbst, E., 2018. A sequential monte carlo approach to inference in multiple-equation markov-switching models. *Journal of Applied Econometrics* 33, 126–140.
- ter Braak, C.J., Vrugt, J.A., 2008. Differential evolution markov chain with snooker updater and fewer chains. *Statistics and Computing* 18, 435–446.
- Chib, S., Ramamurthy, S., 2010. Tailored randomized block mcmc methods with application to dsge models. *Journal of Econometrics* 155, 19–38.
- Childers, D., Fernández-Villaverde, J., Perla, J., Rackauckas, C., Wu, P., 2022. Differentiable State-Space Models and Hamiltonian Monte Carlo Estimation. Technical Report. National Bureau of Economic Research.
- Damlen, P., Wakefield, J., Walker, S., 1999. Gibbs sampling for bayesian non-conjugate and hierarchical models by using auxiliary variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 61, 331–344.
- Dillon, J.V., Langmore, I., Tran, D., Brevdo, E., Vasudevan, S., Moore, D., Patton, B., Alemi, A., Hoffman, M., Saurous, R.A., 2017. Tensorflow distributions. arXiv preprint arXiv:1711.10604 .
- Duane, S., Kennedy, A.D., Pendleton, B.J., Roweth, D., 1987. Hybrid monte carlo. *Physics letters B* 195, 216–222.
- Edge, R.M., Gürkaynak, R.S., Kisacikoglu, B., 2013. Judging the DSGE model by its forecast. Technical Report. mimeo.
- Flegal, J.M., Haran, M., Jones, G.L., 2008. Markov chain monte carlo: Can we trust the third significant figure? *Statistical Science* , 250–260.
- Foreman-Mackey, D., Hogg, D.W., Lang, D., Goodman, J., 2013. emcee: the mcmc hammer. *Publications of the Astronomical Society of the Pacific* 125, 306.
- Gelman, A., Rubin, D.B., 1992. Inference from iterative simulation using multiple sequences. *Statistical science* , 457–472.
- Geman, S., Geman, D., 1984. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence* , 721–741.
- Geweke, J., 1999. Using simulation methods for bayesian econometric models: inference, development, and communication. *Econometric reviews* 18, 1–73.
- Goodman, J., Weare, J., 2010. Ensemble samplers with affine invariance. *Communications in applied mathematics and computational science* 5, 65–80.
- Gornemann, N., Kuester, K., Nakajima, M., 2012. Monetary policy with heterogeneous agents. Working Papers 12-21. Federal Reserve Bank of Philadelphia. URL: <http://ideas.repec.org/p/fip/fedpwp/12-21.html>.
- Haario, H., Saksman, E., Tamminen, J., 2001. An adaptive metropolis algorithm. *Bernoulli* , 223–242.
- Hastings, W.K., 1970. Monte carlo sampling methods using markov chains and their applications .
- Herbst, E., Schorfheide, F., 2014. Sequential monte carlo sampling for dsge models. *Journal of Applied Econometrics* 29, 1073–1098.

- Herbst, E.P., Schorfheide, F., 2015. Bayesian estimation of dsge models, in: Bayesian Estimation of DSGE Models. Princeton University Press.
- Holden, L., Hauge, R., Holden, M., 2009. Adaptive independent metropolis–hastings. *The Annals of Applied Probability* 19, 395–413.
- Igel, C., Hansen, N., Roth, S., 2007. Covariance matrix adaptation for multi-objective optimization. *Evolutionary computation* 15, 1–28.
- Justiniano, A., Primiceri, G.E., Tambalotti, A., 2010. Investment shocks and business cycles. *Journal of Monetary Economics* 57, 132–145. URL: <https://ideas.repec.org/a/eee/moneco/v57y2010i2p132-145.html>.
- Kaplan, G., Moll, B., Violante, G.L., 2018a. Monetary policy according to HANK. NBER Working Papers 3. National Bureau of Economic Research, Inc. URL: <https://ideas.repec.org/p/nbr/nberwo/21897.html>.
- Kaplan, G., Moll, B., Violante, G.L., 2018b. Monetary policy according to hank. *American Economic Review* 108, 697–743.
- Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E., 1953. Equation of state calculations by fast computing machines. *The journal of chemical physics* 21, 1087–1092.
- Neal, R.M., et al., 2011. Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo* 2, 2.
- Nelder, J.A., Mead, R., 1965. A simplex method for function minimization. *The computer journal* 7, 308–313.
- Nelson, B., Ford, E.B., Payne, M.J., 2013. Run dmc: an efficient, parallel code for analyzing radial velocity observations using n-body integrations and differential evolution markov chain monte carlo. *The Astrophysical Journal Supplement Series* 210, 11.
- Roberts, G.O., Rosenthal, J.S., 2001. Optimal scaling for various metropolis-hastings algorithms. *Statistical science* 16, 351–367.
- Roberts, G.O., Rosenthal, J.S., 2007. Coupling and ergodicity of adaptive markov chain monte carlo algorithms. *Journal of applied probability* 44, 458–475.
- Schorfheide, F., 2000. Loss function-based evaluation of dsge models. *Journal of Applied Econometrics* 15, 645–670.
- Smets, F., Wouters, R., 2007. Shocks and frictions in us business cycles: A bayesian dsge approach. *American Economic Review* 97, 586–606.
- Sokal, A., 1997. Monte carlo methods in statistical mechanics: foundations and new algorithms, in: *Functional integration*. Springer, pp. 131–192.
- Storn, R., Price, K., 1997. Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *Journal of global optimization* 11, 341–359.
- Ter Braak, C.J., 2006. A markov chain monte carlo version of the genetic algorithm differential evolution: easy bayesian computing for real parameter spaces. *Statistics and Computing* 16, 239–249.
- Tierney, L., 1994. Markov chains for exploring posterior distributions. *the Annals of Statistics* , 1701–1728.
- Vrugt, J.A., Ter Braak, C., Diks, C., Robinson, B.A., Hyman, J.M., Higdón, D., 2009. Accelerating markov chain monte carlo simulation by differential evolution with self-adaptive randomized subspace sampling. *International journal of nonlinear sciences and numerical simulation* 10, 273–290.
- Winberry, T., 2018. A method for solving and estimating heterogeneous agent macro models. *Quantitative Economics* 9, 1123–1151.

Appendix A Posterior distribution of the estimation of the Smets-Wouters model

	Prior			Proposal		SW mean	Posterior DIME MCMC		
	distribution	mean	std./df	mean	sd.		mean	sd.	mode
σ_c	normal	1.500	0.375	1.357	0.132	1.38	1.354	0.134	1.434
σ_l	normal	2.000	0.750	1.965	0.579	1.81	1.947	0.570	1.871
β_{tpr}	gamma	0.250	0.100	0.140	0.063	0.16	0.140	0.054	0.132
h	beta	0.700	0.100	0.699	0.051	0.71	0.700	0.053	0.702
S''	normal	4.000	1.500	5.443	1.098	5.74	5.462	1.112	6.835
ι_p	beta	0.500	0.150	0.252	0.103	0.25	0.246	0.103	0.187
ι_w	beta	0.500	0.150	0.571	0.135	0.58	0.574	0.136	0.528
α	normal	0.300	0.050	0.183	0.018	0.19	0.182	0.018	0.195
ζ_p	beta	0.500	0.100	0.664	0.063	0.66	0.658	0.066	0.609
ζ_w	beta	0.500	0.100	0.728	0.071	0.70	0.726	0.068	0.705
Φ_p	normal	1.250	0.125	1.579	0.078	1.60	1.578	0.079	1.544
ψ	beta	0.500	0.150	0.542	0.121	0.54	0.547	0.122	0.487
ϕ_π	normal	1.500	0.250	2.053	0.178	2.04	2.052	0.175	2.068
ϕ_y	normal	0.125	0.050	0.095	0.023	0.08	0.094	0.023	0.106
ϕ_{dy}	normal	0.125	0.050	0.231	0.028	0.22	0.230	0.028	0.201
ρ	beta	0.750	0.100	0.817	0.026	0.81	0.817	0.026	0.810
ρ_r	beta	0.500	0.200	0.113	0.079	0.15	0.112	0.061	0.104
ρ_g	beta	0.500	0.200	0.982	0.011	0.97	0.983	0.008	0.980
ρ_z	beta	0.500	0.200	0.963	0.013	0.95	0.964	0.011	0.968
ρ_u	beta	0.500	0.200	0.264	0.140	0.95	0.259	0.146	0.231
ρ_p	beta	0.500	0.200	0.900	0.070	0.89	0.903	0.072	0.946
ρ_w	beta	0.500	0.200	0.976	0.017	0.96	0.975	0.033	0.989
ρ_i	beta	0.500	0.200	0.728	0.063	0.71	0.727	0.059	0.670
μ_p	beta	0.500	0.200	0.767	0.171	0.69	0.742	0.134	0.664
μ_w	beta	0.500	0.200	0.881	0.061	0.84	0.880	0.066	0.923
ρ_{gz}	normal	0.500	0.250	0.503	0.092	0.52	0.502	0.090	0.515
σ_g	inv.gamma	0.100	2.000	0.532	0.030	0.53	0.532	0.030	0.531
σ_u	inv.gamma	0.100	2.000	1.833	0.615	0.23	1.828	0.486	1.871
σ_z	inv.gamma	0.100	2.000	0.460	0.028	0.45	0.460	0.029	0.459
σ_r	inv.gamma	0.100	2.000	0.243	0.015	0.24	0.243	0.015	0.233
σ_p	inv.gamma	0.100	2.000	0.151	0.027	0.14	0.149	0.032	0.120
σ_w	inv.gamma	0.100	2.000	0.249	0.023	0.24	0.249	0.023	0.276
σ_i	inv.gamma	0.100	2.000	0.448	0.048	0.45	0.448	0.048	0.493
$\bar{\gamma}$	normal	0.400	0.100	0.419	0.020	0.43	0.419	0.020	0.428
\bar{l}	normal	0.000	2.000	0.938	1.168	0.53	0.971	1.196	0.906
$\bar{\pi}$	gamma	0.625	0.100	0.673	0.104	0.78	0.670	0.102	0.730

Table A.4: Replication and comparison of the estimation of (Smets and Wouters, 2007, SW) using DIME MCMC. The inverse gamma distribution is parameterized in terms of degrees of freedom as in dynare. The marginals from the proposal distribution is obtained by sampling from the respective multivariate t -distribution in proposal space and then applying the bijective transformation. The mean values of the original estimation (column SW) are obtained from the original paper.

Appendix B Benchmarking against the number of chains

Table B.5 shows Gelman-Rubin coefficients for different ensemble sizes and numbers of cumulative function evaluations. Each measure is the average over the mean across parameters and over ten batches. For each given number of function evaluations n_f (the columns), the sample length is split in half and only the second half is used to calculate the coefficient. E.g., for a given number of function evaluations n_f the sample from iteration $\frac{n_f/n_c}{2}$ to iteration n_f/n_c is used for calculation. The reason is that the Gelman-Rubin coefficient is sensitive to sample length, i.e. the calculation of the Gelman-Rubin coefficient requires much longer chains than a typical sample from the posterior. Note that some more recent work has cast doubt on the reliability of the coefficient to study convergence of MC Markov chains (Flegal et al., 2008).

	1e+05	2e+05	3e+05	4e+05	5e+05	6e+05	7e+05	8e+05	9e+05	1e+06
$n_c = 2n$	1.111 (0.045)	1.094 (0.016)	1.069 (0.006)	1.052 (0.011)	1.044 (0.010)	1.041 (0.015)	1.041 (0.020)	1.036 (0.016)	1.030 (0.012)	1.026 (0.008)
$n_c = 4n$	1.283 (0.105)	1.247 (0.083)	1.223 (0.185)	1.170 (0.155)	1.146 (0.168)	1.127 (0.156)	1.098 (0.103)	1.088 (0.098)	1.084 (0.103)	1.078 (0.104)
$n_c = 6n$	1.515 (0.191)	1.488 (0.146)	1.355 (0.141)	1.373 (0.289)	1.319 (0.251)	1.288 (0.231)	1.302 (0.291)	1.265 (0.242)	1.227 (0.216)	1.206 (0.216)
$n_c = 8n$	1.676 (0.144)	1.633 (0.214)	1.570 (0.307)	1.486 (0.322)	1.436 (0.360)	1.361 (0.299)	1.294 (0.225)	1.264 (0.202)	1.244 (0.194)	1.227 (0.189)

Table B.5: Gelman-Rubin coefficients over different numbers of function evaluations (per column) and numbers of chains n_c per ensemble. Values are means over the means across parameters over 10 batches. Standard deviations across batches are given in brackets.

Appendix C Details on the HANK model

This part of the model is by large adopted from Auclert et al. (2021).

Appendix C.1 Households

The Bellman equation of households is given by

$$V_t(e_{it}, l_{it-1}, a_{it-1}) = \max_{c_{it}, b_{it}, a_{it}} \left\{ \frac{c_{it}^{1-\sigma}}{1-\sigma} - \varphi \frac{N_t^{1+\nu}}{1+\nu} + \beta E_t V_{t+1}(e_{it+1}, b_{it+1}, a_{it}) \right\} \quad (\text{C.1})$$

such that

$$c_{it} + a_{it} + b_{it} = \frac{(1 - \tau_t) w_t N_t}{\int P(e_{jt}) e_{jt}^{1-\Xi} dj} e_{it}^{1-\Xi} + (1 + r_t^a) a_{it-1} + (1 + r_t^b) b_{it-1} - \Phi_t(a_{it}, a_{it-1}), \quad (\text{C.2})$$

$$a_{it} \geq 0, \quad (\text{C.3})$$

$$b_{it} \geq \bar{b}, \quad (\text{C.4})$$

where $\Phi_t(\cdot)$ is the portfolio adjustment cost function

$$\Phi_t(a_{it}, a_{it-1}) = \frac{\chi_1}{\chi_2} \left| \frac{a_{it} - (1 + r_t^a)a_{it-1}}{(1 + r_t^a)a_{it-1} + \chi_0} \right|^{\chi_2} [(1 + r_t^a)a_{it-1} + \chi_0], \quad (\text{C.5})$$

with $\chi_0, \chi_1 > 0$ and $\chi_2 > 1$. Individual labor productivity e_{it} is assumed to follow a random walk process with coefficient ρ_e and a standard deviation of the innovations of σ_t^e , which is by itself assumed to follow an exogenous AR(1) process on an aggregate level.

Appendix C.2 Financial market

No arbitrage at the financial market requires that

$$1 + E_t r_{t+1} = \frac{1 + i_t}{1 + E_t \pi_{t+1}} = \frac{E_t[d_{t+1} + p_{t+1}]}{p_t} = 1 + E_t R_{t+1}^a = 1 + E_t r_{t+1}^b + \omega, \quad (\text{C.6})$$

with ω the parameter governing the cost for liquidity transformation charged by the financial intermediary. Ex-post returns are subject to surprise inflation and capital gains

$$1 + r_t = \frac{1 + i_{t-1}}{1 + \pi_t} = 1 + r_t^b + \omega \quad (\text{C.7})$$

and

$$1 + r_t^a = \Theta_p \left(\frac{d_t + p_t}{p_{t-1}} \right) + (1 - \Theta_p)(1 + r_t), \quad (\text{C.8})$$

where Θ_p denotes the share of equity in the illiquid portfolio.

Appendix C.3 Firms

Firms have a production function

$$y_{jt} = F(k_{jt-1}, n_{jt}) = Z_t k_{jt-1}^\alpha n_{jt}^{1-\alpha}, \quad (\text{C.9})$$

and aggregate marginal costs are given by

$$\widehat{MC}_t = w_t / F_N(\cdot), \quad (\text{C.10})$$

which enter the Phillips curve (19). Z_t is the aggregate level of TFP which follows an AR(1) process around its steady state value. Aggregate investment is given by

$$I_t = K_t - (1 - \delta)K_{t-1} + S \left(\frac{I_t}{I_{t-1}} \right), \quad (\text{C.11})$$

with the quadratic capital adjustment cost function $S(x) = \frac{1}{2S''}(x - 1)^2$ as in the main body, and $\delta > 0$ the parameter for capital depreciation. Dividends are defined as

$$d_t = Y_t - w_t - I_t - \psi_t. \quad (\text{C.12})$$

Tobin's Q and the capital investment decisions follow equations (22) and (23) from the main body.

Appendix C.4 Market clearing

The optimality condition for labor unions is (20) and the monetary policy rule is given by (21). Balanced budget requires

$$\tau_t w_t N_t = r_t B^g + G_t, \quad (\text{C.13})$$

and market clearing requires

$$Y_t = \int c_{it} di + G_t + I_t + \psi_t + \omega b_{it} di, \quad (\text{C.14})$$

$$p_t + B^g = \int a_{it} + b_{it} di. \quad (\text{C.15})$$

Appendix C.5 Fixed parameters

The parameters that are not estimated are set as in table C.6.

Parameter		Value	Target
β	time preference parameter	–	r^*
χ_1	portfolio adj. cost scale	–	$B = 1.04Y$
\bar{b}	borrowing constraint	0	
ρ_e	autocorrelation of earnings	0.966	
ν	disutility of labor	–	$N = 1$
μ_p	steady state markup	–	$p + B^g = 14Y$
μ_w	steady state wage markup	1.1	
Z	TFP	0.468	$Y = 1$
α	capital share	0.33	$K = 10Y$
ω	steady state liquidity premium	0.1	
G	steady state government spending	0.2	
B^g	bond supply	2.8	
n_e	points for Markov chain of e	3	
n_b	points for liquid asset grid	25	
n_a	points for illiquid asset grid	35	

Table C.6: Parameters fixed for the estimation of HANK.

Appendix D Data

The following measurement equations are used for the HANK estimation:

$$\begin{aligned}\text{Real GDP growth} &= \bar{\gamma} + (y_t - y_{t-1}), \\ \text{Real consumption growth} &= \bar{\gamma} + (c_t - c_{t-1}), \\ \text{Real investment growth} &= \bar{\gamma} + (i_t - i_{t-1}), \\ \text{Real wage growth} &= \bar{\gamma} + (w_t - w_{t-1}), \\ \text{Labor hours} &= \bar{n} + n_t, \\ \text{Inflation} &= \bar{\pi} + \pi_t, \\ \text{Federal funds rate} &= 100 \left(\frac{\bar{\pi}}{\beta\gamma^{-\sigma_c}} - 1 \right) + r_t,\end{aligned}$$

The observables are constructed as follows:

- GDP: $\ln(\text{GDP}/\text{GDPDEF}/\text{CNP16OV_ma}) * 100$
- CONS: $\ln((\text{PCEC}-\text{PCEDG}) / \text{GDPDEF} / \text{CNP16OV_ma}) * 100$
- INV: $\ln((\text{GPDI}+\text{PCEDG}) / \text{GDPDEF} / \text{CNP16OV_ma}) * 100$
- LAB: $\ln(13 * \text{AWHNONAG} * \text{CE16OV} / \text{CNP16OV_ma}) * 100$
- INFL: $\ln(\text{GDPDEF}) * 100$
- WAGE: $\ln(\text{COMPENFB} / \text{GDPDEF}) * 100$
- FFR: $\text{FEDFUNDS}/4$

Due to artificial dynamics in the civilian noninstitutional population series that arise from irregular updating (Edge et al., 2013), I use a 4-quarter trailing moving average from Boehl et al. (forthcoming), denoted CNP16OV_ma, to calculate per capita variables.

- GDP: GDP - Gross Domestic Product, Billions of Dollars, Quarterly, Seasonally Adjusted Annual Rate, FRED
- GDPDEF: Gross Domestic Product: Implicit Price Deflator , Index 2012=100, Quarterly, Seasonally Adjusted , FRED
- CNP16OV: Civilian noninstitutional population, Thousands of Persons, Quarterly, Seasonally Adjusted, FRED
- CNP16OV_ma: a four-quarter trailing average of CNP16OV
- PCEC: Personal Consumption Expenditures, Billions of Dollars, Quarterly, Seasonally Adjusted Annual Rate, FRED
- PCEDG: Personal Consumption Expenditures: Durable Goods, Billions of Dollars, Quarterly, Seasonally Adjusted Annual Rate, FRED

- GPDI: Gross Private Domestic Investment, Billions of Dollars, Quarterly, Seasonally Adjusted Annual Rate, FRED
- AWHNONAG: Average Weekly Hours of Production and Nonsupervisory Employees: Total private, Hours, Quarterly, Seasonally Adjusted, FRED
- CE16OV: Employment Level, Thousands of Persons, Quarterly, Seasonally Adjusted, FRED
- COMPNFB, Nonfarm Business Sector: Compensation Per Hour, Index 2012=100, Quarterly, Seasonally Adjusted, FRED
- FEDFUNDS: Effective Federal Funds Rate, Percent, FRED

Appendix E Estimation of HANK on a smaller grid

Tables E.7 and E.8 present the estimation results of HANK using a smaller grid than in Section 5. In particular, the number of grid points for the liquid asset is set to $n_b = 10$ (relative to $n_b = 25$ before) and the number of grid points for the illiquid asset is $n_a = 16$ (compared to 35 before). Finally, the state space of capital is represented by 4 nodes instead of 25 nodes. This implies a smaller grid of 480 nodes instead of the 2625 nodes before, which reduces the estimation time about 34 hours to 70 hours.

	Prior			Large grid			Small grid		
	distribution	mean	std.	mean	std.	mode	mean	std.	mode
σ	normal	1.500	0.375	2.043	0.202	1.850	2.153	0.201	2.035
φ	normal	2.000	0.750	1.738	0.562	1.805	1.712	0.518	1.790
ζ_p	beta	0.500	0.100	0.590	0.050	0.592	0.584	0.052	0.608
ζ_w	beta	0.500	0.100	0.416	0.069	0.413	0.422	0.065	0.371
ι_p	beta	0.500	0.150	0.331	0.129	0.335	0.307	0.127	0.254
ι_w	beta	0.500	0.150	0.322	0.147	0.303	0.330	0.145	0.252
S''	gamma	4.000	2.000	2.279	0.702	1.725	2.079	0.632	2.050
ϕ_π	gamma	1.500	0.250	2.322	0.219	2.198	2.204	0.213	2.253
ϕ_y	gamma	0.125	0.050	0.222	0.063	0.205	0.228	0.064	0.272
ρ	beta	0.750	0.100	0.652	0.052	0.680	0.627	0.055	0.613
\bar{y}	normal	0.400	0.100	0.438	0.026	0.434	0.434	0.025	0.435
\bar{n}	normal	0.000	2.000	-0.047	1.961	1.469	-0.005	1.947	-1.629
π^*	gamma	0.625	0.100	0.596	0.051	0.624	0.594	0.051	0.610
i^*	gamma	1.250	0.100	1.239	0.089	1.259	1.243	0.086	1.218
χ_0	gamma	0.250	0.150	0.153	0.118	0.094	0.118	0.121	0.032
Ξ	beta	0.200	0.100	0.089	0.059	0.071	0.107	0.069	0.126
σ^e	normal	0.920	0.400	0.860	0.185	1.064	0.664	0.107	0.651

Table E.7: Estimation results for HANK with small grid: model parameters

	Prior			Large grid			Small grid		
	distribution	mean	std.	mean	std.	mode	mean	std.	mode
ρ_z	beta	0.500	0.200	0.957	0.018	0.960	0.957	0.016	0.958
ρ_r	beta	0.500	0.200	0.619	0.070	0.640	0.603	0.065	0.629
ρ_g	beta	0.500	0.200	0.993	0.006	0.993	0.991	0.006	0.983
ρ_w	beta	0.500	0.200	0.985	0.006	0.980	0.989	0.004	0.992
ρ_p	beta	0.500	0.200	0.911	0.028	0.917	0.947	0.026	0.959
ρ_i	beta	0.500	0.200	0.837	0.042	0.836	0.788	0.047	0.832
ρ_β	beta	0.500	0.200	0.962	0.030	0.991	0.941	0.044	0.951
σ_z	inv.gamma	0.100	0.250	0.415	0.036	0.430	0.412	0.036	0.396
σ_r	inv.gamma	0.100	0.250	0.130	0.020	0.113	0.138	0.020	0.141
σ_g	inv.gamma	0.100	0.250	1.148	0.088	1.105	1.139	0.089	1.195
σ_w	inv.gamma	0.100	0.250	2.662	0.732	2.448	2.740	0.751	3.270
σ_p	inv.gamma	0.100	0.250	0.201	0.047	0.188	0.205	0.052	0.173
σ_i	inv.gamma	0.100	0.250	1.289	0.215	1.350	1.183	0.204	1.209
σ_β	inv.gamma	0.100	0.250	0.047	0.017	0.029	0.060	0.026	0.051

Table E.8: Estimation results for HANK with small grid: parameters of exogenous processes

Appendix F Details on the estimation of HANK




Figure F.6: Traceplot of the log-likelihood of all chains

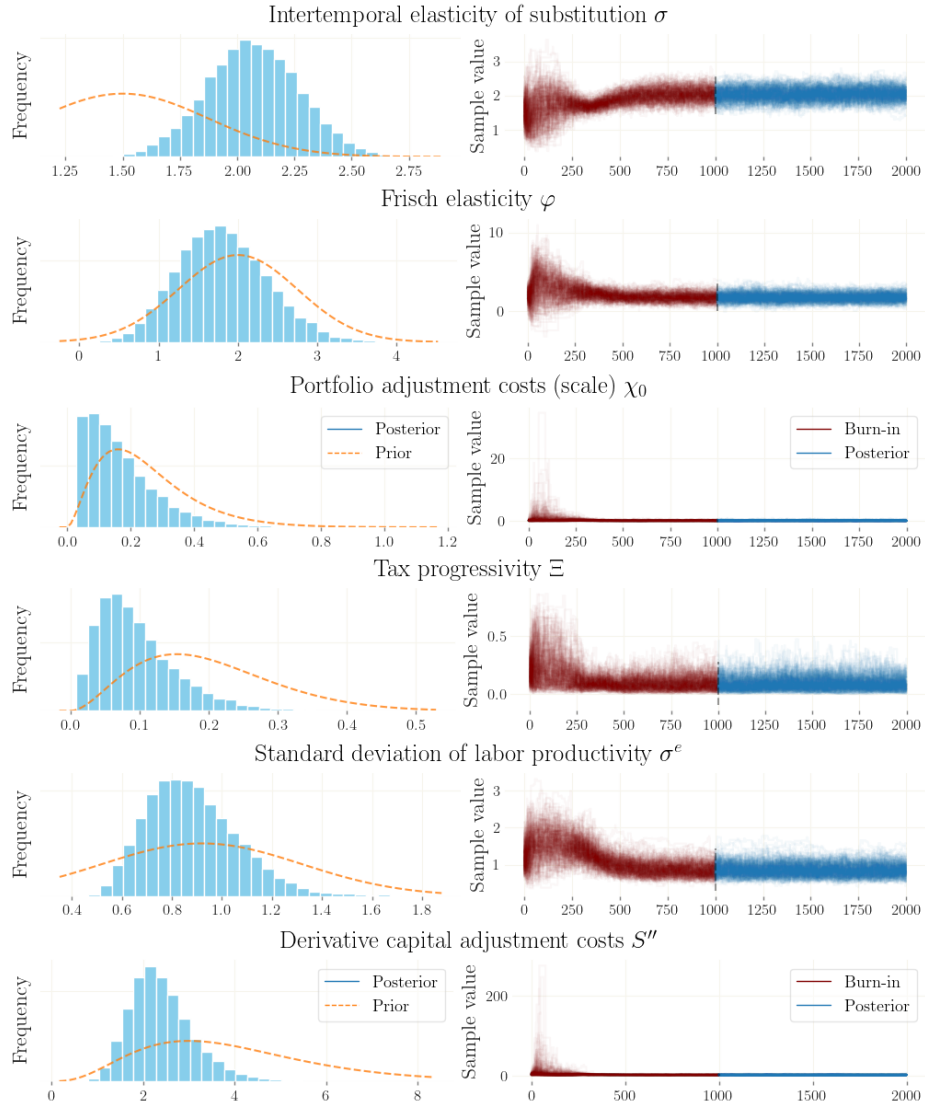


Figure F.7: Traceplots of the 192 DIME chains for the HANK estimation from Section 5. The left panels shows histograms of the marginal distribution over single parameter values. The dashed line plots the respective prior density. The right panels displays the trace of all chains over time, as corresponding to the parameters.

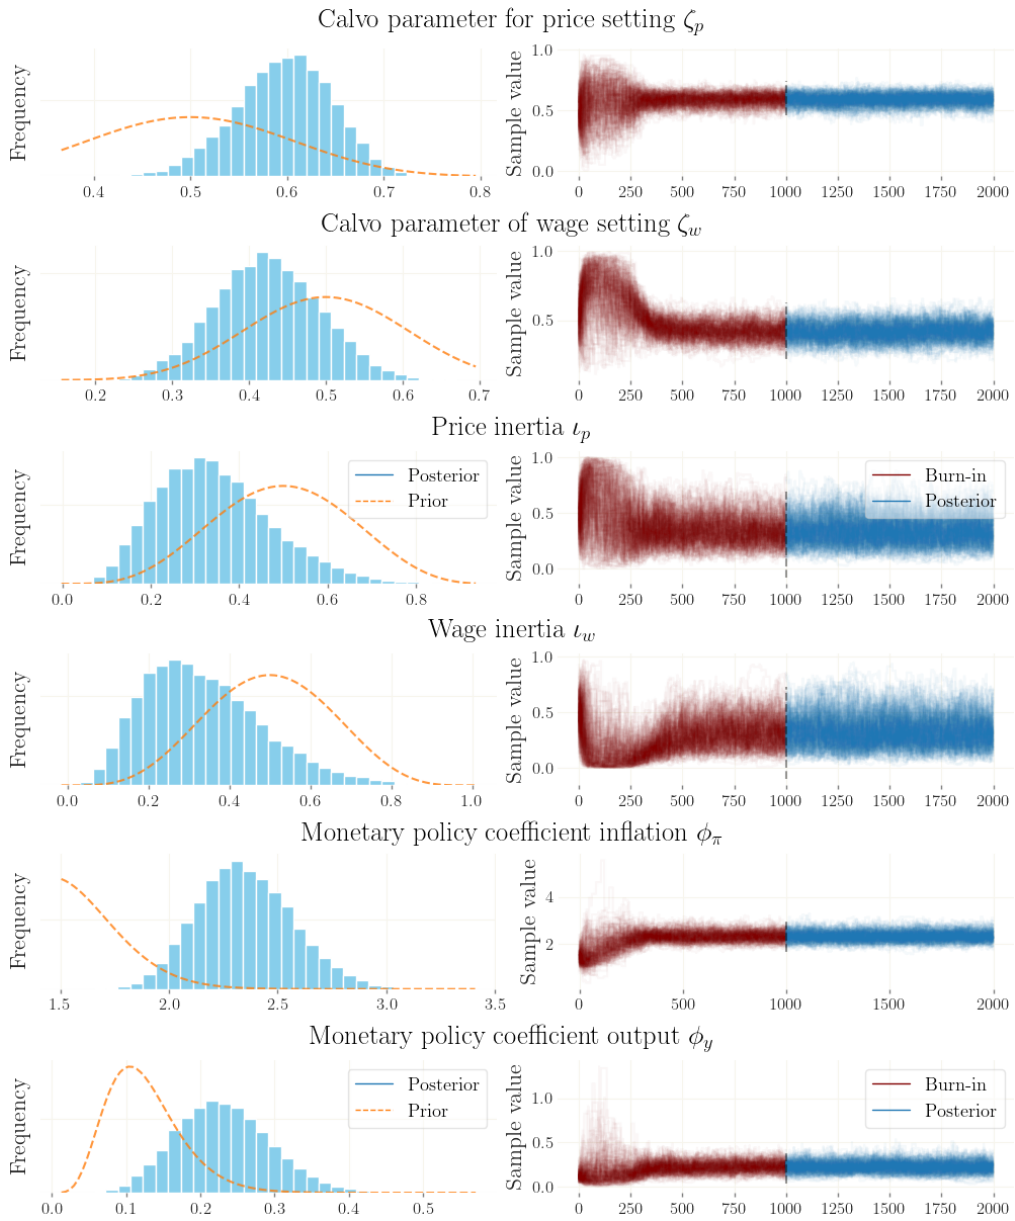


Figure F.8: Traceplots of the 192 DIME chains for the HANK estimation from Section 5. The left panels shows histograms of the marginal distribution over single parameter values. The dashed line plots the respective prior density. The right panels displays the trace of all chains over time, as corresponding to the parameters.

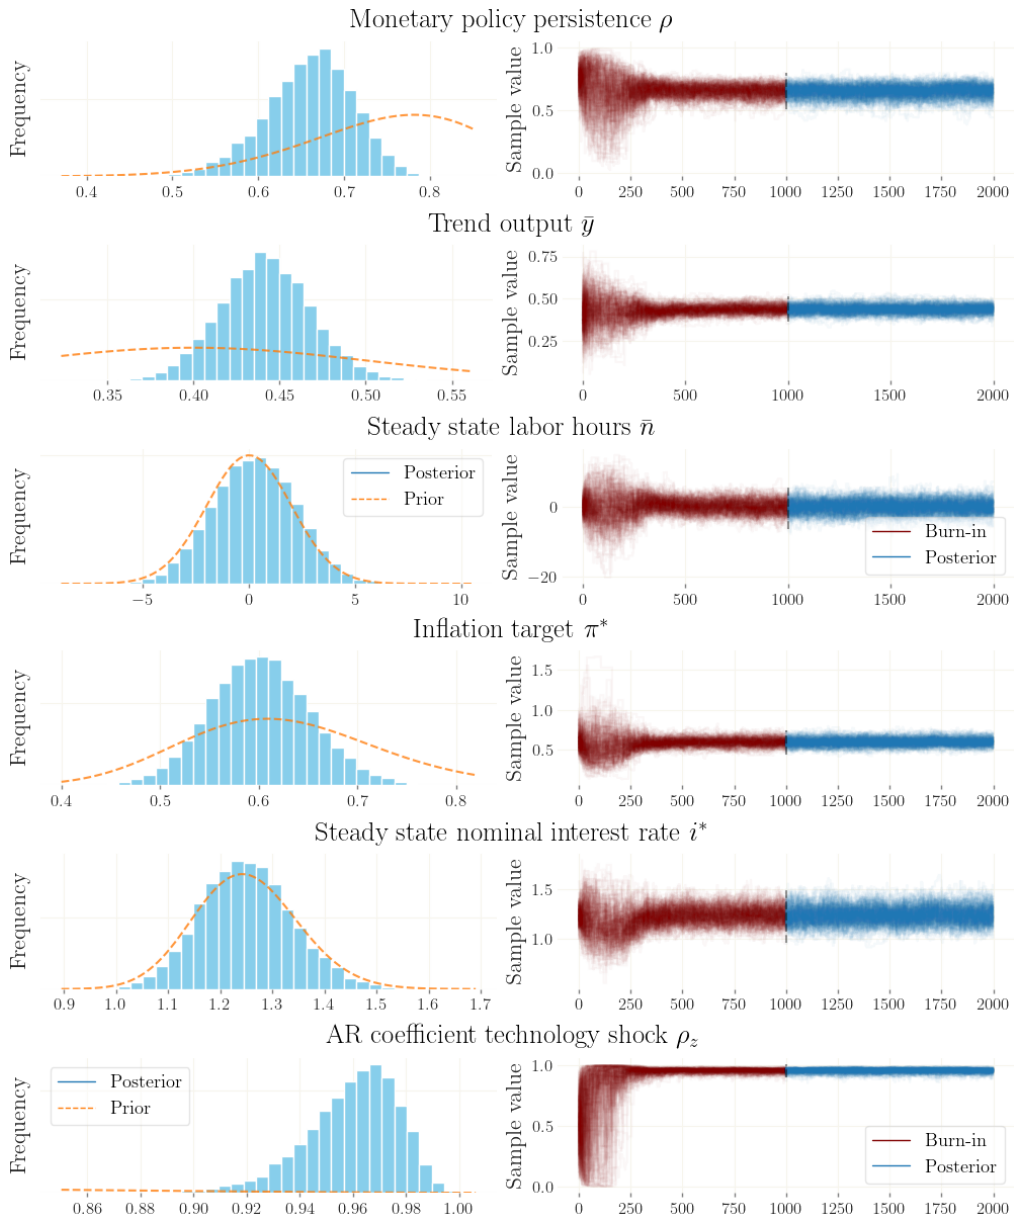


Figure F.9: Traceplots of the 192 DIME chains for the HANK estimation from Section 5. The left panels shows histograms of the marginal distribution over single parameter values. The dashed line plots the respective prior density. The right panels displays the trace of all chains over time, as corresponding to the parameters.

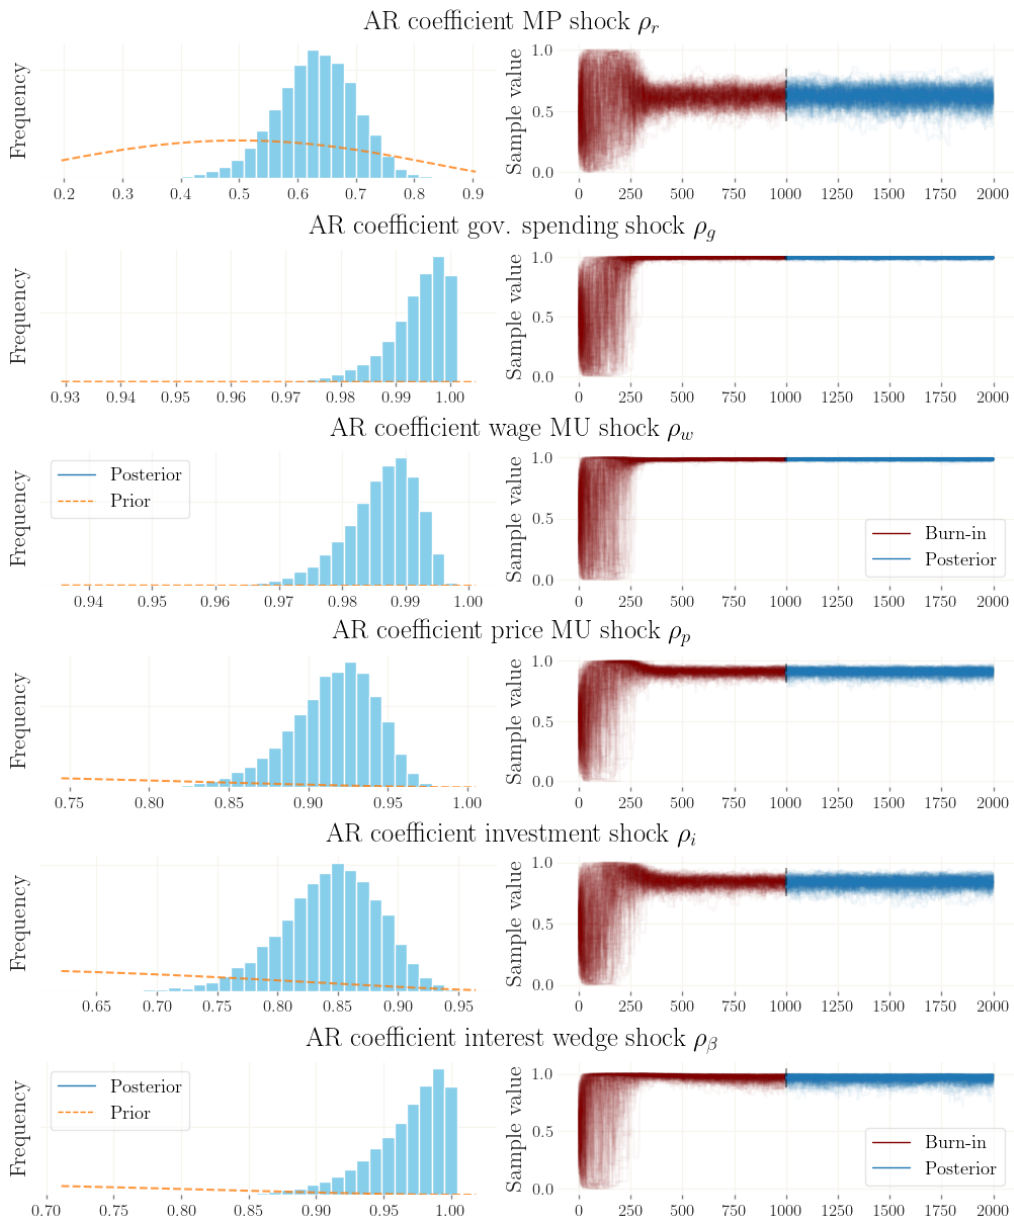


Figure F.10: Traceplots of the 192 DIME chains for the HANK estimation from Section 5. The left panels shows histograms of the marginal distribution over single parameter values. The dashed line plots the respective prior density. The right panels displays the trace of all chains over time, as corresponding to the parameters.




Figure F.11: Traceplots of the 192 DIME chains for the HANK estimation from Section 5. The left panels shows histograms of the marginal distribution over single parameter values. The dashed line plots the respective prior density. The right panels displays the trace of all chains over time, as corresponding to the parameters.

IMFS WORKING PAPER SERIES

Recent Issues

176 / 2022	Michael D. Bauer Carolyn Pflueger Adi Sunderam	Perceptions about Monetary Policy
175 / 2022	Alexander Meyer-Gohde Ekaterina Shabalina	Estimation and Forecasting Using Mixed-Frequency DSGE Models
174 / 2022	Alexander Meyer-Gohde Johanna Saecker	Solving linear DSGE models with Newton methods
173 / 2022	Helmut Siekmann	Zur Verfassungsmäßigkeit der Veranschlagung Globaler Minderausgaben
172 / 2022	Helmut Siekmann	Inflation, price stability, and monetary policy – on the legality of inflation targeting by the Eurosystem
171 / 2022	Veronika Grimm Lukas Nöh Volker Wieland	Government bond rates and interest expenditures of large euro area member states: A scenario analysis
170 / 2022	Jens Weidmann	A new age of uncertainty? Implications for monetary policy
169 / 2022	Moritz Grebe Peter Tillmann	Household Expectations and Dissent Among Policymakers
168 / 2022	Lena Dräger Michael J. Lamla Damjan Pfajfar	How to Limit the Spillover from an Inflation Surge to Inflation Expectations?
167 / 2022	Gerhard Rösl Franz Seitz	On the Stabilizing Role of Cash for Societies
166 / 2022	Eva Berger Sylwia Bialek Niklas Garnadt Veronika Grimm Lars Othér Leonard Salzmänn Monika Schnitzer Achim Truger Volker Wieland	A potential sudden stop of energy imports from Russia: Effects on energy security and economic output in Germany and the EU
165 / 2022	Michael D. Bauer Eric T. Swansson	A Reassessment of Monetary Policy Surprises and High-Frequency Identification

164 / 2021	Thomas Jost Karl-Heinz Tödter	Reducing sovereign debt levels in the post-Covid Eurozone with a simple deficit rule
163 / 2021	Michael D. Bauer Mikhail Chernov	Interest Rate Skewness and Biased Beliefs
162 / 2021	Magnus Reif Mewael F. Tesfaselassie Maik Wolters	Technological Growth and Hours in the Long Run: Theory and Evidence
161 / 2021	Michael Haliassos Thomas Jansson Yigitcan Karabulut	Wealth Inequality: Opportunity or Unfairness?
160 / 2021	Natascha Hinterlang Josef Hollmayr	Classification of Monetary and Fiscal Dominance Regimes using Machine Learning Techniques
159 / 2021	Volker Wieland	The decline in euro area inflation and the choice of policy strategy
158 / 2021	Matthew Agarwala Matt Burke Patrycja Klusak Moritz Kraemer Kamiar Mohaddes	Rising Temperatures, Falling Ratings: The Effect of Climate Change on Sovereign Creditworthiness
157 / 2021	Yvan Lengwiler Athanasios Orphanides	Collateral Framework: Liquidity Premia and Multiple Equilibria
156 / 2021	Gregor Boehl Cars Hommes	Rational vs. Irrational Beliefs in a Complex World
155 / 2021	Michael D. Bauer Eric T. Swanson	The Fed's Response to Economic News Explains the "Fed Information Effect"
154 / 2021	Alexander Meyer-Gohde	On the Accuracy of Linear DSGE Solution Methods and the Consequences for Log-Normal Asset Pricing
153 / 2021	Gregor Boehl Philipp Lieberknecht	The Hockey Stick Phillips Curve and the Zero Lower Bound
152 / 2021	Lazar Milivojevic Balint Tatar	Fixed exchange rate - a friend or foe of labor cost adjustments?
151 / 2021	Thomas Jost Franz Seitz	Designing a European Monetary Fund: What role for the IMF?
150 / 2021	Gerhard Rösl Franz Seitz	Cash and Crises: No surprises by the virus
149 / 2021	Wolfgang Lechthaler Mewael F. Tesfaselassie	Endogenous Growth, Skill Obsolescence and Output Hysteresis in a New Keynesian Model with Unemployment