

Locally Robust Estimation of DSGE Models

Seung Hyun Kim

October 5, 2026

Abstract

Empirical DSGE applications often calibrate parameters governing steady-state values of macroeconomic variables and then treat those calibrations as fixed when conducting likelihood-based inference on other parameters or objects of interest. This produces first-order distortions in posterior inference when the calibration is locally misspecified, while full-information estimation of all parameters may be computationally infeasible. This paper proposes an augmentation of the likelihood of DSGE models that protects inference on parameters of interest against local calibration error and does not require estimation of the calibrated parameters. The method linearizes the model's unconditional mean and impulse responses around a baseline calibration and estimates the coefficients on the resulting derivative directions jointly with parameters that are not calibrated. Under local misspecification, the augmented likelihood delivers an asymptotically unbiased and efficient estimator of the target parameters. We apply the method to Krusell–Smith and one-asset HANK models. Simulations show that even modest calibration errors can severely distort credible-set coverage under a fixed-calibration likelihood, whereas the augmented model restores valid inference.

Keywords: DSGE models; heterogeneous-agent models; calibration; local misspecification; Bayesian inference.

JEL Classification: C11, C13, C32, E32.

1 Introduction

Consider a Dynamic stochastic general equilibrium (DSGE) model with parameters θ . In practice, many empirical DSGE applications divide the parameter vector, either explicitly or implicitly, into parameters estimated from the likelihood, β , and parameters fixed outside the estimation, η . This typically involves choosing η to match some cross-sectional moment, time-series moment, steady-state target, or counterfactual implication. In Smets and Wouters (2007), a leading example of an

estimated medium-size DSGE model, five parameters including the depreciation rate of capital are calibrated. More recently, Acharya et al. (2023) explicitly calibrate parameters that determine the steady state of their heterogeneous agent New Keynesian (HANK) model. Melosi et al. (2022), who estimate heterogeneous agent DSGE models using neural networks, also clearly delineate calibrated parameters from estimated ones. These calibrated parameters are then plugged into the model when estimating β , and are typically treated as fixed in the sense that the effect of this first-stage uncertainty is often ignored in subsequent inference.

Uncertainty in the calibrated values is a concern for many DSGE models of interest. As far back as Canova (1994), it was recognized that different sources of information used in calibrated exercises may produce incompatible calibrations for the same parameter. In the same vein, Ríos-Rull et al. (2012) show that the calibrated value of the Frisch elasticity can differ vastly depending on the type of external micro-level data used in the calibration. This suggests that calibration error is not just a numerical issue. If matching different external moments imply different values for the same parameter, we should suspect misspecification in the calibration of first order or perhaps larger. Under such first-order misspecification, standard inference on β breaks down.

The alternative to calibration—joint full-information estimation of θ —is no panacea. For one, it may simply be computationally infeasible to jointly estimate the parameters in modern DSGE models. For example, in heterogeneous agent models, estimation of parameters governing the steady state values of the aggregate macro variables can be prohibitively costly. Winberry (2018) notes that it takes a simple real business cycle (RBC) model with firm heterogeneity 35 seconds to compute the steady state given a set of parameter values—which would imply close to 200 hours to draw just 20,000 samples from the posterior distribution.

Moreover, macroeconomists may want to calibrate a subset of the parameters even when joint estimation is feasible in finite time. Ríos-Rull et al. (2012) propose a method for choosing a calibration of the Frisch elasticity by making use of both the second moments of aggregate macro variables and external micro-level data, instead of only utilizing the information in the former. Dridi et al. (2007) emphasize how calibration can be statistically motivated as a means of protecting DSGE models against possible model misspecification, to which likelihood-based methods are vulnerable. There is also the issue that, even if steady-state parameters are identified in the population, in finite samples the likelihood surface could be irregular or highly multimodal and complicate likelihood-based inference on these parameters. Indeed, this partly explains the preponderance of Bayesian methods when it comes to joint estimation of DSGE models, since prior information helps regularize the likelihood surface (Herbst and Schorfheide (2015))—but imposing priors that are too tight is indistinguishable from calibration.

This paper introduces a new Bayesian estimation paradigm for DSGE models serving as an alternative to the two extremes of full-information joint estimation and inference under fixed

calibrations. In recognition of the computational difficulties with joint estimation, the procedure starts from the fixed calibration baseline model. The key is that the model is augmented in a manner that gives it enough flexibility for inference to remain valid in the presence of local misspecification in the chosen calibration. As a result, the advantages of calibration are preserved while the statistical shortcomings arising from ignoring calibration uncertainty are rectified.

Mechanically, as proposed in Kim (2026), the baseline model is augmented in a manner that restores the score directions that are lost when we fix the steady-state parameters at their chosen calibrations. For many DSGE models this restoration is implemented through a simple linearization procedure. It is well known that DSGE models usually admit a Wold representation of the aggregate macro variables with Gaussian errors, so that the likelihood is determined by the long run mean of the macro variables and their impulse responses to structural shocks, which in turn depend on the parameters $\theta = (\beta, \eta)$. The basic idea behind the proposed augmentation is to linearize the long run mean and impulse responses around a baseline calibrated value η^* of the nuisance parameters. Rather than jointly re-estimating the calibrated parameter vector η , the procedure retains the baseline calibration η^* and augments the likelihood with the first-order changes in the model's unconditional mean and impulse responses induced by perturbations to η^* . The coefficients on these derivative directions, denoted by δ , are estimated jointly with the parameters of interest β . Under local calibration misspecification, this augmented likelihood yields an estimator of β that remains asymptotically unbiased.

The augmented model nests the baseline model with η fixed at η^* since the two models coincide when δ is equal to 0. The flexibility of the augmented model that allows it to remain robust to local calibration error arises then from varying δ away from 0; in effect, δ absorbs the effects of local misspecification that would have otherwise distorted inference for β . Joint estimation of (β, δ) in the augmented model is nowhere near as computationally intensive as jointly estimating (β, η) in the original model because δ enters linearly into the augmented impulse responses, as opposed to η of which the original impulse responses are highly non-linear functions. Furthermore, under full augmentation estimates of the auxiliary parameter δ is, to first order, equivalent to estimating the deviation $\eta - \eta^*$ of the nuisances from their calibrations. This provides us with a formal way to update the calibration of η using the information in aggregate macro data without having to jointly estimate (β, η) .

The proposed model augmentation is applied to two different heterogeneous agent models: the Krusell and Smith (1998) model and a one-asset HANK model in the spirit of McKay et al. (2016). For the Krusell and Smith (1998) model, simulation studies suggest even slight misspecification in the calibrated parameters is enough to severely distort inference for the parameters of interest, in particular the total factor productivity (TFP) shock standard deviation σ_z . For instance, we find that, when the capital share parameter α is misspecified by 0.03 and the sample size is $T = 250$, a

credible interval of nominal coverage rate 95% for σ_z does not cover the truth at all. Inference on impulse responses is also affected, with a 95% pointwise credible band for the impulse response of output to a TFP shock covering the true impulse response less than 50% of the time for short and long horizons. In contrast, when we estimate the augmented version of the model, correct coverage is restored for all the objects whose inference is distorted under the baseline model.

When applying the model augmentation procedure to the one-asset HANK model with Smets and Wouters (2007) data, we find evidence that the chosen baseline calibration is misspecified. This calibration error is shown to affect estimation of the NKPC slope parameter κ , with the baseline model underestimating the value of κ compared to the augmented model. Taken together, the simulation study and the real-data application demonstrate that heterogeneous agent DSGE models are sensitive even to errors in calibration that we normally consider small; and that the proposed method of model augmentation allows us to conduct valid inference in the presence of these errors.

Throughout the paper, we do not attempt to robustify inference against local misspecification in steady-state parameters that are weakly identified or unidentified from aggregate macro data. Theoretically, the augmentation principle is developed under the assumption that the model is locally identified at the chosen calibration η^* in the sense of Iskrev (2010), without which the Bernstein-von Mises (BvM) theorem does not hold. Also recall that the augmentation procedure is meant to complement the baseline calibration by incorporating information of the steady-state parameters contained in the aggregate macro data with which the model likelihood is constructed. The information in aggregate macro data about weakly identified or unidentified steady state parameters is simply too limited to affect the baseline calibration and any subsequent inference in a meaningful manner.

Literature Review. This paper is the latest in an extensive literature on the calibration and estimation of DSGE models. As in Canova (1994), this paper is interested in safeguarding against calibration error. Del Negro and Schorfheide (2008) present a useful classification of parameters into steady state parameters, parameters determining the propagation of shocks, and parameters related to the law of motion of exogenous shock processes. In our applications the first class of parameters are denoted η and calibrated, while the other two classes of parameters are collected into β and estimated, as aggregate macroeconomic time series often contain limited information about steady state parameters (see, for example, Liu and Plagborg-Møller (2023)).

Meanwhile, the likelihood-based framework of this paper follows that of An and Schorfheide (2007) and subsequent Bayesian estimation methods for DSGE models. Auclert et al. (2021) introduce a sequence space Jacobian approach to the estimation of heterogeneous agent models, which proves especially well-suited to the form of model augmentation proposed in the present paper. Consequently, the two heterogeneous agent model applications in this paper both build on

the likelihood construction in Auclert et al. (2021).

The approach taken in the present paper to the estimation of heterogeneous agent models complements many other modern approaches. Winberry (2018) proposes to estimate heterogeneous agent models by employing a projection-based approach to approximating the wealth distribution, but is silent on the effects of calibration error. Liu and Plagborg-Møller (2023) develop a joint full-information estimation method for heterogeneous agent models using both macro and micro data, but they remain silent on the computational roadblock of repeated steady-state computations.

A recent attempt to marry calibration with rigorous statistical analysis is the work of Cocco and Plagborg-Møller (2025), who provide conservative standard errors for calibrated parameters based on the estimated variance of the associated moments. In contrast to their work, the present paper focuses more on efficient likelihood-based inference in the presence of local calibration error instead of on conducting inference on the calibrated parameters per se.

This paper is also related to the literature that seeks to find a middle ground between maintaining the internal consistency of DSGE models and improving their fit to observed data. A leading example of such a compromise is the DSGE-VAR model of Del Negro, Schorfheide, et al. (2007), where one starts by estimating a VAR model implied by some DSGE model and then sequentially relaxes the cross-equation restrictions implied by the DSGE model to improve model fit. The augmented model proposed in this paper is similar to a DSGE-VAR in the sense that it is grounded in a calibrated DSGE model and allows for slight data-driven deviations from the baseline model to better match the dynamics of the data.

The rest of the paper is organized as follows. In section 2, we study a simplified version of the Brock and Mirman (1972) model that helps illustrate how the linearization works. Section 3 lays out a general roadmap for model augmentation. Section 4 deals with applications to the Krusell and Smith (1998) model and one-asset HANK model. Section 5 concludes.

2 A Simple DSGE Model

To better illustrate how the proposed model augmentation works, we focus on the case of a simple stochastic neoclassical growth model with log utility and full depreciation.

2.1 The Model and its Likelihood

The representative household solves the problem

$$\max_{\{c_t\}_{t \geq 0}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta_d^t \log(c_t) \right] \quad \text{subject to} \quad c_t + a_{t+1} = (1 + r_t)a_t + w_t,$$

where labor supply is normalized to 1 and β_d is the discount rate. The prices r_t, w_t are determined via the firm's problem of maximizing profits $Z_t K_t^\alpha L_t^{1-\alpha} - w_t L_t - (1 + r_t) K_t$ with respect to K_t, L_t , where Z_t is TFP, and $\log$ TFP $z_t = \log(Z_t)$ is assumed to follow a mean zero AR(1) process $z_t = \rho_z z_{t-1} + \epsilon_t$ with iid standard normal innovations. Finally, the market clearing conditions require that $L_t = 1$ and $K_t = a_t$ for any $t \geq 0$.

It is well known that we can analytically solve for consumption and savings in this model as $c_t = (1 - \alpha\beta_d)e^{z_t}K_t^\alpha$ and $K_{t+1} = \alpha\beta_d e^{z_t}K_t^\alpha$. Output Y_t then evolves according to the equation $Y_t = e^{z_t}(\alpha\beta_d)^\alpha Y_{t-1}^\alpha$, so that log output $y_t = \log(Y_t)$ is determined as

$$y_t = \alpha \log(\alpha\beta_d) + \alpha y_{t-1} + z_t.$$

Defining $s_t = (y_t, z_t)$, we can see that s_t is a VAR(2) process given as

$$s_t = \underbrace{\begin{pmatrix} \alpha \log(\alpha\beta_d) \\ 0 \end{pmatrix}}_{\Phi_0} + \underbrace{\begin{pmatrix} \alpha & \rho_z \\ 0 & \rho_z \end{pmatrix}}_{\Phi_1} s_{t-1} + \underbrace{\begin{pmatrix} 1 \\ 1 \end{pmatrix}}_R \epsilon_t,$$

and the model can be written in state space form as

$$y_t = \begin{pmatrix} 1 & 0 \end{pmatrix} s_t \quad \text{(Measurement)}$$

$$s_t = \Phi_0 + \Phi_1 s_{t-1} + R \epsilon_t. \quad \text{(Transition)}$$

Suppose that the value of β_d is known, so that the model parameters are (α, ρ_z) . The SVMA representation for log output y_t is

$$y_t = \mu(\alpha) + \sum_{j=0}^{\infty} \Psi_j(\rho_z, \alpha) \epsilon_{t-j}, \quad \text{where}$$

$$\mu(\alpha) = \frac{\alpha \log(\alpha\beta_d)}{1 - \alpha} \quad \text{and}$$

$$\Theta_j(\rho_z, \alpha) = \frac{\rho_z^{j+1} - \alpha^{j+1}}{\rho_z - \alpha} \quad \forall j \geq 0.$$

Capital share α affects both the steady state and dynamics/autocovariances of log output, whereas ρ_z only affects the latter. We are interested in robustifying ρ_z to calibration error in the steady state parameter α .

2.2 The Augmented Model

The augmented model is constructed by linearizing the long run mean and the impulse responses around the chosen calibration α^* :

$$y_t = \mu(\alpha^*) + \frac{\partial \mu(\alpha^*)}{\partial \alpha} \delta + \sum_{j=0}^{\infty} \left[\Psi_j(\rho_z, \alpha^*) + \delta \frac{\partial \Psi_j(\rho_z, \alpha^*)}{\partial \alpha} \right] \epsilon_{t-j}.$$

The auxiliary parameter δ indexes the first-order departure from the baseline calibration α^* that is relevant for matching the unconditional mean and dynamics of output. Theoretically, it is a local nuisance parameter that allows the augmented model to move in the tangent direction generated by perturbations of α^* , thereby absorbing the first-order effects of calibration error on the likelihood. Practically, it can be used to estimate the structural discrepancy $\alpha - \alpha^*$, as we show below.

It turns out that the augmented model in this simple case is equivalent to the following state-space model:

$$y_t = \mu^{\mathcal{A}}(\delta) + \begin{pmatrix} 1 & \delta - \alpha^* & 0 \end{pmatrix} s_t \quad (\text{Measurement})$$


$$s_t = \begin{pmatrix} 2\alpha^* + \rho_z & -(2\alpha^* \rho_z + (\alpha^*)^2) & \rho_z (\alpha^*)^2 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} s_{t-1} + \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \epsilon_t, \quad (\text{Transition})$$

where


$$\mu^{\mathcal{A}}(\delta) = \frac{\alpha^* \log(\alpha^* \beta_d)}{1 - \alpha^*} + \frac{\log(\alpha^* \beta_d) + 1 - \alpha^*}{(1 - \alpha^*)^2} \delta.$$

The state-space formulation makes clear exactly where the augmented model adds flexibility to the baseline model with fixed α . Note how δ appears only in the mean and factor loadings of the measurement equation. Its presence in the mean indicates that δ gives the augmented model the flexibility to match the long run mean of the data if the baseline calibration does not. Moreover, the fact that δ appears in the second element of the loading vector means that the augmented model essentially extends the baseline model with fixed α , which is AR(1), into an AR(2) process. By adding an additional lag the augmented model becomes better equipped to match the dynamics of the data compared to the baseline model.


Figure 1: **Locally robust inference in the simple model.** The data-generating process sets $\alpha_0 = \alpha^* + 1/\sqrt{T}$ with $T = 1000$. Panel (a) compares the prior and posterior distributions of the implied calibrated parameter $\alpha^* + \delta$. Panel (b) compares posterior inference for ρ under the baseline, augmented, and full models. Panel (c) reports posterior mean impulse responses and 95% uniform credible bands for the baseline and augmented models. Panel (d) reports results from 100 Monte Carlo replications.



(a) Calibration update



(b) Posterior distribution of ρ



(c) Impulse-response inference

	Baseline	Augmented
Mean $\hat{\rho}$	0.7125	0.6923
Coverage of ρ_0	0.89	0.98
Mean $\hat{\alpha}$	—	0.3613
Coverage of α_0	—	1.00
IRF RMSE	0.0320	0.0322
Uniform IRF coverage	0.79	0.96

(d) Monte Carlo performance

2.3 Three Robustness Results

The augmented model above represents a full augmentation of the baseline model since the nuisance parameter is scalar-valued. There are three different ways in which we can make use of a fully augmented model: to estimate the parameter ρ_z , to update the calibration α^* based on information in y_t , and to estimate an economically meaningful object such as impulse responses that depend on both ρ_z and α . To get a brief glimpse of how the augmented model helps robustify these three procedures, we generate a single time series based on a preliminary calibration of $\alpha_0 = 0.33$, $\rho_{z,0} = 0.7$ and $\beta_d = \frac{1}{1+r_{ss}}$ for target steady state interest rate $r_{ss} = 0.01$. We also present Monte Carlo results obtained across multiple different realized time series.

2.3.1 ρ_z Estimation

In figure 1b we compare the posterior densities of ρ_z obtained from the baseline model with α fixed at α^* , the full model where ρ_z is estimated jointly with α , and the augmented model where ρ_z is estimated jointly with the auxiliary parameter δ . $\rho_{z,0}$ comfortably lies within the symmetric 95% credible interval around the posterior mode when ρ_z is estimated jointly with α or through the augmented model, whereas the baseline posterior density at $\rho_{z,0}$ is very small when α^* is positively locally misspecified. The augmented model both recenters the posterior distribution closer to $\rho_{z,0}$ and introduces additional uncertainty, so that the resulting posterior looks much more like that of the full model.

2.3.2 Calibration Updating

It is also possible to conduct locally robust inference on the steady-state parameter α itself. The construction is deceptively simple: given a posterior sample $\{\delta^{(i)}\}_{1 \leq i \leq N}$ of the auxiliary parameter δ , a posterior sample of α is constructed as $\{\alpha^* + \delta^{(i)}\}_{1 \leq i \leq N}$. Figure 1a demonstrates how information about the capital share contained in the data is used to update its prior distribution. Reflecting how the data is generated under a value of α greater than its calibration α^* , posterior mass is predominantly assigned to the region above α^* .

2.3.3 Impulse Response Estimation

Now suppose we are interested in obtaining the impulse response of output to productivity shocks. At each horizon $h \geq 0$, we can form a posterior sample of size N of the h -period ahead impulse response from the baseline model by first obtaining a posterior sample $\{\rho_z^{(i)}\}_{1 \leq i \leq N}$ of ρ_z and then plugging these posterior values into the h -lag MA coefficient:

$$\left\{ \Psi_h \left(\rho_z^{(i)}, \alpha^* \right) \right\}_{1 \leq i \leq N}.$$

Local calibration error in α^* will contaminate this posterior sample of impulse responses—not only is the posterior sample of ρ_z centered incorrectly, the local misspecification of α will directly shift the resulting h -period ahead impulse responses.

Meanwhile, suppose we have a posterior sample $\{\rho_z^{(i)}, \delta^{(i)}\}_{1 \leq i \leq N}$ obtained from the augmented model. The resulting impulse responses can be constructed as

$$\left\{ \Psi_h \left(\rho_z^{(i)}, \alpha^* \right) + \delta^{(i)} \frac{\partial \Psi_h \left(\rho_z^{(i)}, \alpha^* \right)}{\partial \alpha} \right\}_{1 \leq i \leq N}.$$

The resulting posterior sample is robust to local calibration error because it corrects for the two

sources of misspecification when estimating impulse responses in the baseline model. First, it uses a locally robust sample of ρ_z instead of one that may be centered at the incorrect mode in the limit. Meanwhile, the added term involving posterior draws of δ serves as a linear correction to the baseline impulse response $\Psi_h(\rho_z^{(i)}, \alpha^*)$ that serves to offset any biases that arise from using α^* in the impulse response formula instead of α_0 .

In figure 1c we plot the impulse responses of output to a productivity shock for horizons 0 to 10 that are implied by each model given a single simulation of the data, where the construction of the 95% uniform bands follows Montiel Olea and Plagborg-Møller (2019). Although the difference between the augmented and baseline bands are not too large, the augmented model just about manages to cover the true IRF; the same cannot be said for the baseline credible bands. The table in figure 1d reveals that this is not an artefact of the particular time series with which the impulse responses in figure 1c were constructed. Despite similar RMSEs, the IRFs in the augmented model manage to achieve nominal coverage 95% while those produced by the baseline model only covers the true IRF 79% of the time.

2.4 Takeaways

The simple model elucidates the central mechanism of the paper. The calibrated parameter α affects both the unconditional mean and the propagation of shocks. Fixing α therefore makes inference on ρ_z sensitive to calibration error. The augmented model resolves this tension by keeping α^* as the expansion point while adding the local nuisance direction generated by perturbations in α . This allows the likelihood to absorb local calibration error without requiring full nonlinear estimation of α .

3 The General Model

In this section we introduce the general augmentation scheme for Gaussian SVMA models, followed by comments on the details of implementation and a heuristic explanation of why the augmentation delivers robustness against local misspecification. Further details about the econometrics behind the procedure are found in Kim (2026).

3.1 Setup and Notation

Let $\{X_t\}_{t \in \mathbb{Z}}$ be a doubly infinite n -dimensional time series generated as

$$X_t = \mu(\theta) + \sum_{j=0}^{\infty} \Psi_j(\theta) \epsilon_{t-j}, \quad \epsilon_t \stackrel{iid}{\sim} \mathcal{N}[O_{n_\epsilon \times 1}, I_{n_\epsilon}].$$

Here ϵ_t can be viewed as n_ϵ -dimensional structural shocks, $\Psi_j(\theta)$ the j -period ahead impulse response of the aggregate macro variables X_t to the structural shocks, and $\mu(\theta)$ the long-run mean of the aggregate macro variables.

The k -dimensional parameter vector θ can be decomposed as $\theta = (\beta, \eta)$, where the blocks are k_β and k_η -dimensional, respectively. β collects the estimable parameters in the model—such as shock parameters and, in HANK models, Taylor rule parameters. η is a nuisance parameter block that cannot be jointly estimated with β through likelihood-based methods in finite time, or for which there exist external micro-level datasets that are more informative than aggregate macro variables. These include what we have up to now referred to as steady-state parameters. Due to reasons discussed above, we wish to avoid estimating η jointly with β . Instead, η is fixed at some predetermined value η^* . In what follows, we will be interested in the estimation of some function $H(\theta)$ of the parameters, such as impulse responses, multipliers, MPCs, or some subset of the estimable parameters β .

The asymptotic theory treats η^* as a fixed non-random quantity and focuses on the effect of local misspecification in η^* . In other words, we are interested in how inference on $H(\theta)$ is affected when the data is generated under $\eta^* + \frac{h}{\sqrt{T}}$ instead of η^* for some $h \in \mathbb{R}^{k_\eta}$ ¹. Errors in η^* of order $T^{-1/2}$ are our focus because errors of a smaller order are asymptotically negligible, while errors of a higher order can be detected with probability approaching 1 on the basis of the likelihood.

3.2 How the Augmentation Works

Let

$$\dot{\ell}_T(\theta) = \begin{pmatrix} \dot{\ell}_{T,\beta}(\theta) \\ \dot{\ell}_{T,\eta}(\theta) \end{pmatrix}$$

be the T -sample score function of the full model evaluated at the parameter vector θ , where $\dot{\ell}_{T,\beta}$ is the β -score and $\dot{\ell}_{T,\eta}$ the η -score.

If η is fixed at η^* and the model is treated as one whose only parameter is β , then standard M-estimation theory tells us that the maximum likelihood estimator $\hat{\beta}_T^{nv}$ of β admits the first-order expansion $\sqrt{T}(\hat{\beta}_T^{nv} - \beta_0) = \mathbb{I}_{*,\beta\beta}^{-1} \frac{1}{\sqrt{T}} \dot{\ell}_{T,\beta}(\beta_0, \eta^*) + o_p(1)$, where $\mathbb{I}_{*,\beta\beta}$ is the covariance matrix of the β -score $\dot{\ell}_{T,\beta}(\beta_0, \eta^*)$. If η^* is the true value of η , then $\hat{\beta}_T^{nv}$ is asymptotically normal with mean zero and covariance matrix given by $\mathbb{I}_{*,\beta\beta}^{-1}$. However, if η^* is locally misspecified for the truth, so that

¹Sometimes calibrations are viewed as first-stage estimators rather than fixed plug-in values. If these estimators are constructed from micro-level data independent of the aggregate variables entering the likelihood, they can be treated as fixed for the likelihood-based analysis considered here. When the same data are used in both stages, additional complications arise from conditioning on the first-stage estimator; see Chernozhukov et al. (2022). Sample splitting provides one way to avoid this dependence.

the true value of η is given as $\eta_0 = \eta^* + \frac{h}{\sqrt{T}}$ for some $h \in \mathbb{R}^{k_\eta}$, then $\hat{\beta}_T^{nv}$ will be asymptotically biased. The bias term depends on the unknown direction and magnitude of misspecification h , so that standard inference based on $\hat{\beta}_T^{nv}$ is invalidated.

The asymptotic bias of the maximum likelihood estimator $\hat{\beta}_T^{nv}$ is linear in the covariance between the β -score and the η -score. In other words, the source of the bias is the fact that this covariance is not accounted for when we fix η at η^* ; mathematically, this can be seen from the absence of the η -score in the first-order expansion of $\sqrt{T}(\hat{\beta}_T^{nv} - \beta_0)$. What the proposed augmentation procedure does is restore this missing η -score.

Concretely, we introduce a new model with parameters (β, δ) , where δ is a vector of auxiliary parameters. The new model is chosen so that, when δ is equal to the zero vector, the likelihood of the model coincides exactly with the original model with η fixed at η^* —this is the likelihood-matching condition. This ensures that the new model nests the baseline model, and therefore can truly be viewed as an augmentation of the baseline. Furthermore, the augmented model is designed so that the score with respect to (β, δ) at $(\beta_0, 0)$ is exactly the score with respect to $\theta = (\beta, \eta)$ at $\theta^* = (\beta_0, \eta^*)$ —this is the score-matching condition. This means that the maximum likelihood estimator of β in the augmented model has exactly the same first-order expansion as the joint maximum likelihood estimator of β in the original model, and thus is robust against local misspecification in η^* . Under regularity conditions that ensure asymptotic equivalence between Bayesian and frequentist methods, this implies that jointly sampling from the posterior of (β, δ) in the augmented model delivers a posterior sample of β that is also robust to local calibration error.

For stationary Gaussian SVMA models, it turns out that we can construct an augmented model that satisfies the required likelihood and score matching conditions by simply linearizing the long run mean and impulse response functions around the calibration η^* . Below we delve into the details of how this linearization is implemented for two cases of interest.

3.3 The Augmentation Principle

Given a fixed calibration η^* , a posterior sample for $H(\theta)$ is usually formed through the following two-step procedure that treats η as fixed and equal to η^* :

Step 1 Posterior Sample for β

We obtain a posterior sample $\{\beta^{(i)}\}_{1 \leq i \leq N}$ of some size N from

$$X_t = \mu(\beta, \eta^*) + \sum_{j=0}^{\infty} \Psi_j(\beta, \eta^*) \epsilon_{t-j}, \quad \epsilon_t \stackrel{iid}{\sim} \mathcal{N}[O_{n_\epsilon \times 1}, I_{n_\epsilon}] \quad (3.1)$$

viewed as a model with parameters β .

Step 2 Posterior Sample for $H(\theta)$

Form a posterior sample for $H(\theta)$ of the same size N as $\{H(\beta^{(i)}, \eta^*)\}_{1 \leq i \leq N}$.

Under appropriate differentiability conditions on $H(\cdot)$, a delta-method type result for total variation distance implies that, if the BvM theorem holds for the posterior sample $\{\beta^{(i)}\}_{1 \leq i \leq N}$, then it also holds for $\{H(\beta^{(i)}, \eta^*)\}_{1 \leq i \leq N}$. The resulting asymptotic equivalence between Bayesian and frequentist inference means that, if the maximum likelihood estimator for $H(\theta)$ under fixed calibration is asymptotically biased, which it generally is under local calibration error, then so is the posterior sample $\{H(\beta^{(i)}, \eta^*)\}_{1 \leq i \leq N}$. Intuitively, the failure of posterior inference under fixed calibration stems from the fact that we are only varying β when forming the posterior sample and ignoring any potential variation, or uncertainty, in η^* .

The corresponding augmented model is formed in a manner that linearly takes into account potential variation in η alongside β . The desired posterior sample for $H(\theta)$ is then obtained as follows:

Step 1 Forming the Augmented Model

Specify the following model for X_t that depends on the parameters (β, δ) where δ is k_η -dimensional:

$$X_t = \mu^{\mathcal{A}}(\beta, \delta) + \sum_{j=0}^{\infty} \Psi_j^{\mathcal{A}}(\beta, \delta) \epsilon_{t-j}, \quad \epsilon_t \stackrel{iid}{\sim} \mathcal{N}[O_{n_\epsilon \times 1}, I_{n_\epsilon}] \quad (3.2)$$

where

$$\mu^{\mathcal{A}}(\beta, \delta) = \mu(\beta, \eta^*) + \frac{\partial \mu(\beta, \eta^*)}{\partial \eta'} \delta \quad (3.3)$$

$$\Psi_j^{\mathcal{A}}(\beta, \delta) = \Psi_j(\beta, \eta^*) + \sum_{\ell=1}^{k_\eta} \frac{\partial \Psi_j(\beta, \eta^*)}{\partial \eta_\ell} \delta_\ell \quad \forall j \geq 0 \quad (3.4)$$

Step 2 Posterior Sample for ν

Form a posterior sample $\{\beta^{(i)}, \delta^{(i)}\}_{1 \leq i \leq N}$ of (β, δ) of some size N .

Step 3 Posterior Sample for $H(\theta)$

Form a posterior sample for $H(\theta)$ of the same size N as

$$\left\{ H(\beta^{(i)}, \eta^*) + \frac{\partial H(\beta^{(i)}, \eta^*)}{\partial \eta'} \delta^{(i)} \right\}_{1 \leq i \leq N}$$

The posterior sample for $H(\theta)$ constructed as above will deliver valid inference even when η^* is locally misspecified for the truth; that is, in the limit, its credible intervals will have actual coverage

equivalent to the nominal coverage rate and point estimators formed from the posterior sample will be asymptotically unbiased. The robustness of the posterior sample obtained from the model augmentation procedure can be attributed to the fact that, at least locally, the variation in η required for inference to remain valid is captured by the variation in the auxiliary parameter δ .

3.3.1 Optional Parsimonious Augmentation

Here we consider a special case where $H(\theta)$ selects a subset of the estimable parameters β , so that it does not depend on the value of the nuisance parameters η . Denote the k_1 -dimensional subset of β that we want to estimate as β_f , and the remaining k_{β_f} elements of β as β_N , so that β can be further decomposed as (β_f, β_N) . For mnemonic purposes, the subscript f stands for “focus” and N for “nuisance”. In the following, we use the notation

$$\theta = (\beta_f, \beta_N, \eta), \quad \beta = (\beta_f, \beta_N), \quad \bar{\eta} = (\beta_N, \eta), \quad (3.5)$$

where $\bar{\eta}$ stands for the collection of all nuisance parameters regardless of whether they are estimable.

If we want to form a posterior sample of $H(\theta)$ in general, then because this functional depends on the nuisance parameter η as well as the estimable parameters β we must robustify the model against every possible direction in which η^* might be locally misspecified. Therefore the auxiliary parameter we introduce is the same dimension as η . However, if we want to form a posterior sample of just β_f , then we need only robustify the model against the directions of misspecification that are relevant to β_f —and it is well known from semiparametric efficiency theory that there are only as many of these directions as there are parameters in β_f . As such, in the second case we need only k_{β_f} auxiliary parameters in the augmented model, leading to a much more parsimonious augmentation if $k_{\beta_f} \ll k_{\eta}$, that is, if there are many more nuisance parameters than parameters of interest.

Under the fixed calibration η^* , a posterior sample for β_f is formed by obtaining a posterior sample $\{\beta_f^{(i)}, \beta_N^{(i)}\}_{1 \leq i \leq N}$ from the full model with η fixed at η^* . The draws $\{\beta_f^{(i)}\}_{1 \leq i \leq N}$ of β_f then serves as our desired posterior sample. Since this case corresponds to $H(\theta) = \beta_f$, it follows from the discussion above that inference based on $\{\beta_f^{(i)}\}_{1 \leq i \leq N}$ will fail under local calibration error in η^* .

To robustify inference using a parsimonious augmentation of the baseline model we first let $\theta^* = (\beta_0, \eta^*)$, where $\beta_0 = (\beta_{f,0}, \beta_{N,0})$ are the true values of the estimable parameters, and define $\mathbb{I}_*$ as the $k \times k$ Fisher information matrix evaluated at the point θ^* . We assume that $\mathbb{I}_*$ is non-singular² and partition it as

$$\mathbb{I}_* = \begin{pmatrix} \mathbb{I}_{*,\beta_f\beta_f} & \mathbb{I}_{*,\beta_f\bar{\eta}} \\ \mathbb{I}_{*,\bar{\eta}\beta_f} & \mathbb{I}_{*,\bar{\eta}\bar{\eta}} \end{pmatrix}$$

²We are assuming that the model is locally identified at θ^* .

conformably along the dimensions of β_f and $\bar{\eta} = (\beta_N, \eta)$. The non-singularity of $\mathbb{I}_*$ implies that $\mathbb{I}_{*,\bar{\eta}\bar{\eta}}$ is also non-singular, so that the $k_{\bar{\eta}} \times k_{\beta_f}$ matrix $\mathbf{B}_0 = \mathbb{I}_{*,\bar{\eta}\bar{\eta}}^{-1} \mathbb{I}_{*,\bar{\eta}\beta_f}$ is well-defined. Now we form a locally robust posterior sample of β_f as follows:

Step 1 Forming the Augmented Model

Specify the following model for X_t that depends on the parameters (β, δ) where δ is k_{β_f} -dimensional:

$$X_t = \mu^{\mathcal{A}}(\beta, \delta) + \sum_{j=0}^{\infty} \Psi_j^{\mathcal{A}}(\beta, \delta) \epsilon_{t-j}, \quad \epsilon_t \stackrel{iid}{\sim} \mathcal{N} [O_{n_\epsilon \times 1}, I_{n_\epsilon}] \quad (3.6)$$

where

$$\mu^{\mathcal{A}}(\beta, \delta) = \mu(\beta, \eta^*) + \frac{\partial \mu(\beta, \eta^*)}{\partial \bar{\eta}'} \mathbf{B}_0 \delta \quad (3.7)$$

$$\Psi_j^{\mathcal{A}}(\beta, \delta) = \Psi_j(\beta, \eta^*) + \sum_{\ell=1}^{k_\beta} \left(\sum_{r=1}^{k_{\bar{\eta}}} \frac{\partial \Psi_j(\beta, \eta^*)}{\partial \bar{\eta}_r} \mathbf{B}_{0,r\ell} \right) \delta_\ell \quad \forall j \geq 0 \quad (3.8)$$

Step 2 Posterior Sample for β

Form a posterior sample $\{\beta_f^{(i)}, \beta_N^{(i)}, \delta^{(i)}\}_{1 \leq i \leq N}$ of (β, δ) of some size N . The collection of the draws $\{\beta_f^{(i)}\}_{1 \leq i \leq N}$ for β_f is our posterior sample for β_f .

Notice how the matrix $\mathbf{B}_0$ acts as a matrix that selects the directions of the nuisance parameter $\bar{\eta}$ that are relevant for estimation of β_f . Mathematically, $\mathbf{B}_0$ is defined as the population regression coefficient from regressing the β_f -score onto the $\bar{\eta}$ -score, so that it effectively summarizes the parts of the $\bar{\eta}$ -score that affect estimation of β_f .

3.4 Practical Considerations

Here we address details related to the practical implementation of the augmentation procedure. There are three aspects of the implementation that require some care: computing the derivatives of $\mu(\cdot)$ and $\Psi_j(\cdot)$ involved in the construction of the augmented model; computing the selection matrix $\mathbf{B}_0$; and choosing the prior over δ .

3.4.1 Numerical Derivatives for Model Augmentation

The proposed model augmentation involves the derivatives $\frac{\partial \mu(\beta, \eta^*)}{\partial \eta'}$ and $\frac{\partial \Psi_j(\beta, \eta^*)}{\partial \eta'}$ of the long run mean and impulse responses with respect to the steady-state parameters η . Consider HA models

where the functions $\mu(\cdot)$ and $\Psi_j(\cdot)$ can be expressed as functions of the type

$$f(\beta, \eta, X_{ss}(\eta)) \quad (3.9)$$

where $\eta \mapsto X_{ss}(\eta)$ maps steady-state parameters into the corresponding steady-state values of the aggregate macro variables and the mapping $(\beta, \eta, x_{ss}) \mapsto f(\beta, \eta, x_{ss})$ is a function with known analytical derivatives. The complexity of the derivatives with respect to η arise solely from the mapping $X_{ss}(\cdot)$, which in HA models is found as the solution to some implicit function characterized by steady-state market clearing conditions. Notably, $\mu(\cdot)$ and $\Psi_j(\cdot)$ admit representations as in equation (3.9) when the likelihood is obtained via the SSJ approach in Auclert et al. (2021), and as such the subsequent applications opt for the representation of the likelihood in terms of SSJs instead of a linear state-space representation.

By the chain rule, the total derivative of the mapping $\eta \mapsto f(\beta, \eta, X_{ss}(\eta))$ with respect to the j th steady state parameter η_j is given as

$$\frac{d}{d\eta_j} f(\beta, \eta, X_{ss}(\eta)) = \frac{\partial f(\beta, \eta, X_{ss}(\eta))}{\partial \eta_j} + \frac{\partial f(\beta, \eta, X_{ss}(\eta))}{\partial X_{ss}} \times \frac{\partial X_{ss}(\eta)}{\partial \eta_j}.$$

While analytical expressions are normally available for $\frac{\partial f(\beta, \eta, X_{ss}(\eta))}{\partial \eta_j}$ and $\frac{\partial f(\beta, \eta, X_{ss}(\eta))}{\partial X_{ss}}$, the partial derivative $\frac{\partial X_{ss}(\eta)}{\partial \eta_j}$ is usually computed numerically in a way that preserves the market clearing conditions that define $X_{ss}(\eta)$. Specifically, suppose that the steady state values $X_{ss}(\eta)$ are defined implicitly via the market clearing condition

$$M(X_{ss}(\eta), \eta) = 0.$$

By the implicit function theorem,

$$\frac{\partial X_{ss}(\eta)}{\partial \eta_j} = - \left(\frac{\partial M(X_{ss}(\eta), \eta)}{\partial X'_{ss}} \right)^{-1} \frac{\partial M(X_{ss}(\eta), \eta)}{\partial \eta_j}.$$

The partial derivatives of the mapping $(x, \eta) \mapsto M(x, \eta)$ are generally unknown, so we plug in the two-sided numerical approximations to $\frac{\partial M(X_{ss}(\eta), \eta)}{\partial X'_{ss}}$ and $\frac{\partial M(X_{ss}(\eta), \eta)}{\partial \eta_j}$ into the above formula to compute $\frac{\partial X_{ss}(\eta)}{\partial \eta_j}$. Letting the mapping $X_{ss}(\cdot)$ be n_x -dimensional, we therefore require only $2 \times (n_x + k_\eta)$ evaluations of the steady state to obtain the desired derivatives. Since we can reuse $\frac{\partial X_{ss}(\eta^*)}{\partial \eta_j}$ under different values of β when computing $\frac{d}{d\eta_j} f(\beta, \eta^*, X_{ss}(\eta^*))$, model augmentation is much less computationally costly compared to recomputing the steady state values for every candidate value of θ , which is what would happen under full-information joint estimation of the parameters.

3.4.2 An Estimator for $\mathbf{B}_0$

If we are interested in conducting inference on the subset β_f of β , $\mathbf{B}_0$ is given as the population regression coefficient $\mathbb{I}_{*,\bar{\eta}\bar{\eta}}^{-1}\mathbb{I}_{*,\bar{\eta}\beta_f}$ from regressing the β -score onto the $\bar{\eta}$ -score at the baseline $\theta^* = (\beta_0, \eta^*)$. This quantity is not known in general, and must be replaced in practice with a consistent estimator $\mathbf{B}_T^*$.

While the formula for the population Fisher information $\mathbb{I}(\theta)$ evaluated at the parameter vector θ is known for stationary Gaussian SVMA models, computing $\mathbb{I}(\theta)$ involves numerically integrating over the frequency domain. A cleaner solution is to work with the finite-sample Fisher information $\hat{\mathbb{I}}_T(\theta_T^*)$ where θ_T^* is a consistent estimator for θ^* , which we know will be itself consistent for $\mathbb{I}_*$.

We construct θ_T^* as (β_T^*, η^*) , where β_T^* is the point estimator for $\beta_f = (\beta, \beta_N)$ obtained from the baseline model (3.1) with η fixed at η^* . Because we consider only local calibration error in η^* , while β_T^* may be asymptotically biased it is still consistent for the true value β_0 , making θ_T^* consistent for θ^* . $\mathbf{B}_T^*$ is then constructed by plugging in θ_T^* into the function $\hat{\mathbb{I}}_T(\cdot)$, whose (i, j) th element is given by

$$\hat{\mathbb{I}}_{T,ij}(\theta) = \frac{1}{2T} \text{tr} \left(\mathbf{V}_{T,i}(\theta) \Sigma_T(\theta)^{-1} \mathbf{V}_{T,j}(\theta) \Sigma_T(\theta)^{-1} \right) + \frac{1}{T} (\iota_T \otimes \dot{\mu}_i(\theta))' \Sigma_T(\theta)^{-1} (\iota_T \otimes \dot{\mu}_j(\theta)), \quad (3.10)$$

where ι_T is the T -dimensional vector of ones, $\Sigma_T(\theta)$ is the $nT \times nT$ covariance matrix of $(X_1, \dots, X_T)$ evaluated at the parameter vector θ , and for $1 \leq i \leq k$,

$$\begin{aligned} \dot{\mu}_i(\theta) &= \frac{\partial \mu(\theta)}{\partial \theta_i}, \\ \Sigma_T^{\sigma(i)}(\theta) &= \text{Covariance matrix of } (Y_1, \dots, Y_T) \text{ where } Y_t = \sum_{j=0}^{\infty} \left(\Psi_j(\theta) + \frac{\partial \Psi_j(\theta)}{\partial \theta_i} \right) \epsilon_{t-j} \\ \Sigma_T^{D(i)}(\theta) &= \text{Covariance matrix of } (Y_1, \dots, Y_T) \text{ where } Y_t = \sum_{j=0}^{\infty} \left(\frac{\partial \Psi_j(\theta)}{\partial \theta_i} \right) \epsilon_{t-j} \\ \mathbf{V}_{T,i} &= \Sigma_T^{\sigma(i)}(\theta) - \Sigma_T^{D(i)}(\theta) - \Sigma_T(\theta). \end{aligned}$$

It is shown in the appendix that the $\mathbf{V}_{T,i}$ constructed as such is equivalent to the derivative of $\Sigma_T(\theta)$ with respect to θ_i . The above definition reuses the derivatives of $\Psi_j(\cdot)$ obtained during the construction of the augmented model and therefore is computationally faster than having to additionally compute the numerical derivative of $\Sigma_T(\theta)$.

3.4.3 Prior Selection for δ

The economic meaning of the auxiliary parameters δ in the augmented model may not be immediately clear, complicating the choice of a prior distribution for δ . To this end, we leverage the asymptotic equivalence between δ and η . In the case where $H(\theta)$ is the object of interest, the maximum likelihood estimator for δ in the augmented model is asymptotically equivalent to that of η in the original model. Therefore, given a reasonable prior covariance matrix Σ_η^* for η , we can form a prior for δ simply as $\mathcal{N} [O_{k_\eta \times 1}, \Sigma_\eta^*]$ ³.

Meanwhile, if the subset β_f of β is the object of interest, the maximum likelihood estimator for δ is now asymptotically equivalent to that of $(\mathbf{B}'_0 \mathbb{I}_{*, \bar{\eta} \bar{\eta}} \mathbf{B}_0)^{-1} \mathbf{B}'_0 \mathbb{I}_{*, \bar{\eta} \bar{\eta}} \bar{\eta}$. As such, in this case a natural information-weighted prior covariance matrix for δ would assume the sandwich form

$$(\mathbf{B}'_0 \mathbb{I}_{*, \bar{\eta} \bar{\eta}} \mathbf{B}_0)^{-1} \mathbf{B}'_0 \mathbb{I}_{*, \bar{\eta} \bar{\eta}} \Sigma_\eta^* \mathbb{I}_{*, \bar{\eta} \bar{\eta}} \mathbf{B}_0 (\mathbf{B}'_0 \mathbb{I}_{*, \bar{\eta} \bar{\eta}} \mathbf{B}_0)^{-1},$$

where Σ_η^* now reflects not only the prior variability in η but also in β_N .

This mapping shifts the burden of specifying a prior covariance for δ to specifying a prior covariance for η , for which there are many statistical and economic approaches one could take. An ad-hoc approach would be to specify Σ_η^* on the basis of the variability of popular calibrations in the literature (e.g. the discount rate β_d is usually calibrated between 0.985 and 0.99, so a reasonable prior standard deviation would be 0.0025). Although intuitive, this approach is silent on the prior correlation between the elements in η .

One alternative is to specify Σ_η^* as the conservative asymptotic variance of corresponding GMM estimator if η^* is formed through moment-matching (e.g. Cocci and Plagborg-Møller (2025)). Another is to follow Del Negro and Schorfheide (2008) and elicit Σ_η^* from the steady-state relationships they imply. A simple version of this elicitation proceeds as follows. Let $\eta \mapsto X_{ss}(\eta)$ be the mapping from steady-state parameters to the steady-state values of endogenous variables, where $X_{ss}(\cdot)$ and η have the same dimension, and Σ_x the diagonal matrix specifying how much prior variation we are willing to allow for each steady-state value. If $X_{ss}(\cdot)$ has a non-singular derivative at η^* , then denoting the derivative by $DX_{ss}(\eta^*)$, the induced prior covariance for η is⁴

$$\Sigma_\eta^* = DX_{ss}(\eta^*)^{-1} \Sigma_x DX_{ss}(\eta^*)^{-1'}.$$

³In practice, due to the local nature of the augmentation we truncate the Gaussian prior around the origin so that it preserves 90% of its mass; see Kim (2026) for details.

⁴This can be understood heuristically as a version of the measurement equation specification $X_{ss}^o = X_{ss}(\eta) + \varepsilon$ in Del Negro and Schorfheide (2008) with ε given as a random vector with diagonal covariance matrix Σ_x .

3.5 Implementation Algorithm

Consider a heterogeneous agent DSGE model involving n aggregate macro variables collected in X_t that depends on the parameters θ . We conduct locally robust Bayesian inference on the parameter $H(\theta)$ as follows.

Algorithm 1: Locally Robust Bayesian Estimation

Input: DSGE model with estimated parameters $\beta = (\beta_f, \beta_N)$, steady-state parameters η , and baseline calibration η^* .

1 Baseline model. Solve the model at η^* and construct the unconditional mean and impulse responses entering into the baseline SVMA representation in (3.1). Ensure that the mean and impulse response functions admit an expression of the form in (3.9).

2 Nuisance derivatives. Compute the derivatives of the unconditional mean and impulse responses with respect to η at η^* , using the chain rule and a combination of analytical and numerical partial derivatives.

3 Augmentation directions. For inference on general $H(\theta)$, this step is skipped. For inference on β_f , obtain a preliminary consistent estimate θ_T^* of $\theta^* = (\beta_0, \eta^*)$ using the baseline model, and then construct a consistent estimator B_T^* of B_0 using the finite-sample Fisher information formula laid out in equation (3.10).

4 Prior specification. Construct a prior for δ centered at zero whose scale reflects uncertainty about η , and furnish priors for $\beta = (\beta_f, \beta_N)$ as well.

5 Augmented estimation. Construct the augmented model (3.2) (equivalently for inference on β_f , (3.6) with mean and impulse responses given according to the formulas in equations (3.3) and (3.4) (equivalently, equations (3.7) and (3.8)).

6 Posterior inference. Construct posterior draws of (β, δ) from the augmented model under the given prior distributions. Using these draws construct a posterior sample of β_f or the composite parameter $H(\theta)$.

Output: Locally robust posterior sample for object of interest.

4 Application to HA Models

We apply the machinery developed in the preceding section to robustify heterogeneous agent models against local calibration error. These types of models naturally lend themselves to model augmentation due to the computational difficulty of estimating them via full-information methods.

In HA models the agents' steady state wealth distribution, along with steady state parameters such as the discount rate or relative risk aversion, determines the steady state values of aggregate macro variables. A simple example is the steady state real interest rate: in the stochastic neoclassical growth model of section 2, the steady state value of r_t is given as $r_{ss} = \frac{1}{\beta_d} - 1$. In contrast, in the Krusell and Smith (1998) model the same value is defined implicitly via the capital market clearing condition, where the interest rate determines capital supply both through the optimal saving policy and through the distribution of wealth across households.

The presence of the wealth distribution as an additional infinite-dimensional state variable destroys the simple mapping between the steady state parameters and the steady state values of aggregate macro variables, which now have to be recovered via costly bisection loops. Estimating these steady state parameters using maximum likelihood or Bayesian methods now requires one to recompute the steady state values for every new combination of parameters, rendering estimation extremely costly.

Below we estimate and augment two HA models: a heterogeneous agent version of the RBC model a la Krusell and Smith (1998), and the one-asset HANK Model found in works such as McKay et al. (2016). We solve both HA applications with the sequence-space Jacobian method of Auclert et al. (2021). Consider a model with only steady state parameters η and parameters β_N related to the evolution of the exogenous variables. The SSJ method decomposes the impulse-response vector as $\Psi(\beta_N, \eta) = \mathbf{G}(\eta)d\mathbf{Z}(\beta_N)$, where $\mathbf{G}(\eta)$ is the general equilibrium Jacobian describing the response of the endogenous macro variables to changes in the exogenous variables, and $d\mathbf{Z}(\beta_N)$ is a vector encoding how the exogenous variables respond to a time 0 MIT shock. This decomposition is particularly useful here because the nuisance derivatives $\frac{\partial \mathbf{G}(\eta)}{\partial \eta'}$ are computed once at the baseline calibration and then reused at every posterior draw. Thus, the augmentation retains a low-dimensional sampling problem without repeatedly resolving the HA steady state. In the HANK application below, the parameters of interest also enter the GE Jacobian. Their derivatives are nevertheless inexpensive to evaluate because they do not require recomputing the household steady state or its distribution.

4.1 The HA RBC Model

4.1.1 Model Setup

A continuum of households on the unit interval $[0, 1]$ have CRRA preferences, and choose consumption and supply fixed hours of labor, normalized to 1. These households own shares of capital from which they receive financial income. Households are subject to uninsurable idiosyncratic

labor productivity shocks e_t as well as a borrowing constraint. Household i thus faces the problem

$$\max_{\{c_t^i\}_{t \geq 0}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta_d^t \frac{c_t^{1-\sigma} - 1}{1-\sigma} \right] \quad \text{s.t.} \quad c_t^i + a_{t+1}^i = (1+r_t)a_t^i + w_t e_t^i, \quad a_{t+1}^i \geq 0$$

where β_d is the discount rate and $\sigma > 0$ is the relative risk aversion coefficient. We assume that $\log(e_t^i) = m_e + \rho_e \log(e_{t-1}^i) + \sigma_e v_t^i$ for $v_t^i \sim \mathcal{N}[0, 1]$ that are iid across i and t . We denote the corresponding transition kernel for e_t^i by $P(\cdot, \cdot)$. Finally, w_t, r_t are the time t wage level and interest rate on capital respectively.

There is a representative firm using a Cobb-Douglas production function with capital share $\alpha \in (0, 1)$ to produce output using two factors, labor and capital. Production depends on TFP Z_t , so that the firm's problem at time t is maximizing profits $Z_t K_t^\alpha L_t^{1-\alpha} - w_t L_t - (r_t + \delta_d) K_t$ with respect to the factors K_t, L_t , where $\delta_d \in (0, 1)$ is the depreciation rate of capital. We assume that $\log(Z_t) = \rho_z \log(Z_{t-1}) + \sigma_z \epsilon_t$ for $\epsilon_t \sim \mathcal{N}[0, 1]$ that is iid across time. The market clearing conditions require that $K_t = \int_0^1 a_t^i di$ and $L_t = \int_0^1 e_t^i di$.

A recursive equilibrium for this model is standard and its definition can be found in, say, Winberry (2018). It is characterized by a collection of value, policy and pricing functions satisfying the household and firm optimality conditions, the market clearing conditions for capital and labor, and finally a specification for how the wealth distribution evolves alongside a condition that ensures household beliefs are consistent. A full definition is provided in the appendix.

4.1.2 Steady State

When the economy is at the steady state, so that $Z_t = 1$ for $t \geq 0$, the model reduces to that of Aiyagari (1994). The steady state wage w_{ss} is determined as $w_{ss} = (1 - \alpha) \left(\frac{\alpha}{r_{ss} + \delta_d} \right)^{\frac{\alpha}{1-\alpha}}$ through the firm's first order conditions. The steady state interest rate r_{ss} solves the capital market clearing condition $K^d(r) = K^s(r)$, where capital demand $K^d(r)$ is derived once again from the firm's first order conditions and capital supply $K^s(r)$ is computed from the steady-state savings policy function and the stationary wealth distribution.

Obtaining the steady-state savings policy function requires iterating through either the value function or the policy function until convergence is reached; and computing the stationary distribution requires iterating over the transition kernel induced by the policy function found in the first step until convergence. The combination of this backward and forward iteration steps accounts for the bulk of the computational cost associated with solving the KS model.

The parameters involved in the search for r_{ss} are collected into $\eta = \{\alpha, \beta_d, \delta_d\}$ and we refer to them as steady state parameters. Despite playing an important role in the determination of the steady state, the household heterogeneity parameters m_e, ρ_e, σ_e^2 are not included in η . This is because

the proposed model augmentation procedure robustifies inference against local misspecification in calibrated parameters that are locally identified—and it has been noted in the literature (e.g. Liu and Plagborg-Møller (2023)) that parameters governing household heterogeneity are weakly identified or unidentified using aggregate macro data. The relative risk aversion coefficient σ is likewise excluded, since evaluations of the Fisher information at the chosen calibration reveals that σ is quite weakly identified by output data.

4.1.3 The Wold Representation and Computing General Equilibrium Jacobians

In this model, the parameters of interest are the parameters governing the evolution of TFP, collected into $\beta = \{\rho_z, \sigma_z\}$. Suppose that an economy that is at the steady state at time 0 will, in response to any shock at time 0, have returned to the steady state by time T . Since Z_t follows an AR(1) process in logs, its deviation dZ_t from its steady state value 1 admits the Wold representation $dZ_t \approx \log(Z_t) = \sum_{j=0}^{\infty} \rho_z^j \sigma_z \epsilon_{t-j}$. Let $d\mathbf{Z}(\beta)$ be the vector collecting the MA coefficients of dZ_t , and $\mathbf{G}(\eta) = \frac{\partial \mathbf{Y}}{\partial \mathbf{Z}}(\eta)$ the general equilibrium Jacobian for output whose $(t+1, s+1)$ th element is $\frac{\partial Y_t}{\partial Z_s}(\eta)$, or the impulse response of time t output to a deviation at time s of TFP from its steady state value, it is shown in Auclert et al. (2021) that the DGP for Y_t can be approximated via the finite-order MA process

$$Y_t \approx Y_{ss}(\eta) + \sum_{j=0}^{T-1} \Psi_j(\beta, \eta) \epsilon_{t-j},$$

where $\Psi_j(\beta, \eta) = (\mathbf{G}(\eta) d\mathbf{Z}(\beta))_j$ and $Y_{ss}(\eta)$ is the steady state value of output that depends on the steady state parameters η . The truncation at lag $T-1$ reflects our assumption that any transition path in response to an MIT shock will have length at most T .

Define $\frac{\partial \mathbf{K}}{\partial \mathbf{Z}}(\eta)$ so that its $(t+1, s+1)$ th element is $\frac{\partial K_{t+1}}{\partial Z_s}(\eta)$, the impulse response of time t saving to a deviation in time s TFP from its steady state value. Since time t output Y_t is given by $Y_t = Z_t K_t^\alpha (L^*)^{1-\alpha}$, it follows from the chain rule that

$$\mathbf{G}(\eta) = Y_{ss}(\eta) I_T + (r_{ss}(\eta) + \delta_d) \frac{\partial \mathbf{K}}{\partial \mathbf{Z}}(\eta) I_T(-1)$$

where $I_T(-1)$ is the $T \times T$ matrix with its 1s on its first subdiagonal and zeros elsewhere. Intuitively, the first term in $\mathbf{G}(\eta)$ represents the direct effect of changes in TFP on output; while the second term is the indirect effect through capital that arrives with a one-period lag. The only missing piece required to compute the impulse responses is the sequence space Jacobian $\frac{\partial \mathbf{K}}{\partial \mathbf{Z}}(\eta)$, for which we use the Fake news algorithm detailed in Auclert et al. (2021).

4.1.4 Augmentation of the Likelihood

Let the baseline calibration of η be denoted by η^* . Then, the general augmented model is constructed as

$$Y_t = Y_{ss}(\eta^*) + \frac{\partial Y_{ss}(\eta^*)}{\partial \eta'} \delta + \sum_{j=0}^{T-1} \Psi_j^{\mathcal{A}}(\beta, \delta) \epsilon_{t-j}$$

where the expressions for $\Psi_j^{\mathcal{A}}(\beta, \delta)$ are given in equation (3.4)⁵. Collecting the augmented MA coefficients into the vector $\Psi^{\mathcal{A}}(\beta, \delta)$, since $\Psi(\beta, \eta^*) = \mathbf{G}(\eta^*) d\mathbf{Z}(\beta)$, we can express

$$\Psi^{\mathcal{A}}(\beta, \delta) = \left\{ \mathbf{G}(\eta^*) + \sum_{\ell=1}^{k_\eta} \frac{\partial \mathbf{G}(\eta^*)}{\partial \eta_\ell} \delta_\ell \right\} d\mathbf{Z}(\beta),$$

As such, the MA coefficients for the augmented model are also separable in δ and β . Computationally, this decomposition indicates that we need only compute the derivatives $\frac{\partial \mathbf{G}(\eta^*)}{\partial \eta_\ell}$ once for $1 \leq \ell \leq k_\eta$ and reuse them to construct $\Psi^{\mathcal{A}}(\beta, \delta)$ for different pairs (β, δ) . The same goes for the mean augmentation, where we reuse the derivative $\frac{\partial Y_{ss}(\eta^*)}{\partial \eta'}$ that is computed once before sampling (β, δ) from the augmented model. These derivatives are computed using two-sided numerical differentiation as noted in the preceding section. Sampling from the posterior distribution now follows algorithm 1.

4.1.5 Prior Selection and Posterior Sampler

For ρ_z we follow Auclert et al. (2021) and impose a beta prior with mean 0.5 and standard deviation 0.2. Likewise, we impose an inverse gamma prior on σ_z with mean 0.4 and standard deviation 4. These are diffuse priors, and posterior results are relatively insensitive to these choices.

The prior variance of δ is informed by the prior information on the nuisance parameters η . As detailed in section 3.4.3, the prior covariance Σ_η^* of the calibrated parameters is chosen based on steady-state target matching. Given that $\eta = \{\alpha, \beta_d, \delta_d\}$, we choose to target the steady state relationships

$$\frac{K_{ss}}{Y_{ss}} = \frac{\alpha}{r_{ss} + \delta_d}, \quad r_{ss} = r^*(\alpha, \beta_d, \delta_d), \quad w_{ss} = (1 - \alpha) \left(\frac{\alpha}{r_{ss} + \delta_d} \right)^{\frac{\alpha}{1-\alpha}}.$$

We choose to target the steady state capital-output ratio alongside the steady state prices because α is literally the capital share of production. We implement the procedure so that one prior standard

⁵Here we opt for the full augmentation since we are interested in both the shock parameters and impulse responses, which are composite objects.

deviation in the calibrated parameters corresponds, to first order, to a 15% deviation in each steady state target. At our present calibration

$$\Sigma_{\eta}^* = \begin{pmatrix} 0.0016 & -6.76 \times 10^{-7} & 0.0003 \\ -6.76 \times 10^{-7} & 2.30 \times 10^{-6} & 1.15 \times 10^{-6} \\ 0.0003 & 1.15 \times 10^{-6} & 6.57 \times 10^{-5} \end{pmatrix},$$

so that the prior standard deviations of the calibrated parameters are roughly 0.04, 0.0015 and 0.008 for α , β_d and δ_d .

For sampling we use the random-walk Metropolis-Hastings (RWMH) algorithm as described in Herbst and Schorfheide (2015). The proposal variance is computed as the inverse negative Hessian of the posterior density evaluated at the posterior mode, where we use automatic differentiation to obtain the Hessian up to machine precision. The proposal variance is multiplied by a scaling parameter $\tau > 0$ that we tune over the burn-in phase to achieve the desired acceptance rate.

Table 1: **Calibration in the HA RBC model** The discount rate β_d and the location m_e of the idiosyncratic income shocks are chosen to match the steady state targets of output and the interest rate. σ_e^2 is given as $\Delta_e^2(1 - \rho_e^2)$.

<i>Steady State Parameters</i>			
σ	Relative Risk Aversion		1
α	Capital Share in Production		0.36
δ_d	Depreciation Rate of Capital		0.025
ρ_e	Persistence of Idiosyncratic Shock		0.966
Δ_e	Cross-sectional Standard Deviation of Idiosyncratic Shock		0.5
<i>Steady State Targets</i>			
Y_{ss}	Steady State Output		1
r_{ss}	Steady State Real Interest Rate		0.01

4.1.6 A Simulation Study

We conduct a simulation study to investigate the effect of local misspecification and evaluate how the augmented model performs under such forms of misspecification. Given a sample size $T = 250$, the calibration η^* of the steady state parameters is chosen as in table 1, which are almost exactly those used in Auclert et al. (2021). We then generate the artificial output data under $\rho_z = 0.7$, $\sigma_z^2 = 0.1$ and the exact same values as in table 1 for the steady-state parameters except for the capital share α —the data is generated under $\alpha_0 = \alpha^* - \frac{h}{\sqrt{T}}$ for $h = 0.3$ and 0.5 . This allows us to isolate the effects of local misspecification in the capital share. Under $T = 250$, α_0 is approximately 0.33 for $h = 0.5$ and 0.34 for $h = 0.3$, which are both calibrations that are found commonly in the

Table 2: **Monte Carlo study under local misspecification.** We report the coverage of credible intervals for ρ_z and σ_z , the absolute deviations of their posterior modes from the true values, and their posterior standard deviations.

h	T	Model	Coverage		Absolute Bias		Posterior SD	
			ρ_z	σ_z	ρ_z	σ_z	ρ_z	σ_z
0.5	250	Baseline	0.950	0.000	0.0358	0.0607	0.0432	0.0115
		Augmented	0.960	0.930	0.0393	0.0274	0.0462	0.0345
0.3	250	Baseline	0.950	0.180	0.0339	0.0402	0.0444	0.0124
		Augmented	0.960	0.960	0.0389	0.0269	0.0461	0.0349

Notes: Coverage is based on nominal 95% credible intervals over 500 simulations. Data are generated under $\alpha_0 = \alpha^* - h/\sqrt{T}$ with $\alpha^* = 0.36$. Absolute bias is computed as the Monte Carlo average of $|\theta_0 - \hat{\theta}_{\text{mode}}|$. Posterior standard deviations are averaged across Monte Carlo replications.

literature—choosing any one of these calibrations is usually taken to be innocuous when estimating macroeconomic models.


The simulation results in table 2 tells a consistent story across the magnitude of local misspecification h . Throughout the simulations, the actual coverage rate of the credible interval around ρ_z is close to its nominal rate of 95% for both the baseline and augmented models, indicating that local misspecification in α does not significantly distort inference on shock persistence. However, the 95% credible interval around the shock variance σ_z , when constructed from the baseline model, fails to cover the true value almost all of the time when the data is generated under $\alpha_0 = 0.33$ and severely undercovers even when $\alpha_0 = 0.34$. This suggests that even a local calibration error of the magnitude $0.5/\sqrt{T}$ for α , is enough to completely shift the posterior of σ_z away from the truth.

The failure of the baseline model stands in stark contrast to the performance of the augmented model, which attains a coverage rate close to the nominal rate for σ_z . The absolute bias of the posteriors for σ_z in the augmented model are noticeably smaller than those of the baseline model, demonstrating how the augmentation recenters the posterior distribution of parameters that were highly biased under the baseline model. In the case of ρ_z , both the absolute bias and posterior standard deviation are similar across the baseline and augmented models.

The poor coverage of the posterior of σ_z under the baseline model can be traced back to the general equilibrium consequences of local misspecification in α . When α changes from $\alpha^* = 0.36$ to $\alpha_0 = 0.33$, the steady-state interest rate falls only slightly from 0.01 to 0.0099 while drastically lowering capital demand

$$K^d(r) = \left(\frac{\alpha}{r + \delta_d} \right)^{\frac{1}{1-\alpha}} L^*.$$

Figure 2: **Monte Carlo study for impulse response coverage.** Below we display the coverage of pointwise credible bands for the impulse response of output and capital to TFP shocks up to 40 horizons. The impulse responses are evaluated for the baseline and augmented models, using data generated under the truth $\alpha_0 = 0.33$ and calibration $\alpha^* = 0.36$ with sample size $T = 250$.



This leads to a reduction in steady-state capital supply and, by implication, steady-state output. Numerically, steady-state capital falls from approximately 10.29 to 7.63, and steady-state output falls from 1 to approximately 0.81, a huge 20% decrease, in response to a change of α from 0.36 to 0.33. It is precisely this discrepancy that the baseline model cannot address under a fixed calibration, but which the augmented model is able to accommodate. Broadly speaking, this finding illustrates how a seemingly modest difference in calibrations can create substantial general equilibrium effects that ultimately result in large distortions in posterior inference.

Table 3: **Average 95% credible-band widths under local misspecification.** Average band widths are reported for the output and capital impulse responses under $T = 250$ and the calibration $\alpha^* = 0.36$. The final row reports the ratio of augmented to baseline average width.

	Output			Capital		
	1–4	5–12	13–40	1–4	5–12	13–40
Baseline	0.2369	0.1778	0.0469	0.3912	1.1694	0.9805
Augmented	0.2783	0.1479	0.0415	0.8351	1.4102	1.1456
Augmented / Baseline	1.17	0.83	0.89	2.13	1.21	1.17

Notes: Each entry is the average width of the nominal 95% pointwise credible band over the indicated horizon block. Horizon 0 is excluded. A ratio above one indicates that the augmented credible bands are wider on average.

Local misspecification in the steady state parameters also shifts the impulse responses of the endogenous macro variables to TFP shocks by altering their GE Jacobians. Under the same local misspecification setting for α as above, we compare the actual coverage of nominal 95% pointwise credible bands for the impulse response of output and capital to TFP shocks across the baseline and augmented models. When obtaining the posterior sample of the impulse response of output, we use the composition function $H(\beta, \eta) = \mathbf{G}(\eta)dZ(\beta)$; for the impulse response of capital, the composition function is $H(\beta, \eta) = \frac{\partial \mathbf{K}}{\partial \mathbf{Z}}(\eta)dZ(\beta)$.

Figure 2 shows us that the baseline model once again fails to achieve nominal coverage across horizons and macro variable, with undercoverage being much more severe for shorter and longer horizons compared to medium horizons, and for capital relative to output. The augmented model, meanwhile, restores coverage and has a tendency to slightly overcover at longer horizons. This demonstrates how, in the presence of even a small degree of calibration error, the usual fixed calibration approach to inference can be invalidated not only for the shock variance parameter σ_z , but also for more economically meaningful objects like impulse responses. Since the choice of whether to calibrate the capital share as 0.33 or 0.36 often comes down to the vintage of the data, model augmentation is a way to protect subsequent inference against the possibility of making the incorrect choice.

We now study what drives the undercoverage of the baseline impulse responses. The baseline impulse response is given as the product of a fixed GE Jacobian G and the vector $(1, \rho_z, \dots, \rho_z^{T-1})$. Broadly speaking, the off-diagonal elements of G grow smaller in magnitude the farther away they are from the diagonal. As such, the constant 1 loads heavily on impulse responses of shorter horizons, while ρ_z^j terms with large j that are extremely close to 0 dominate those of longer horizons. This means baseline impulse responses at the shortest and longest horizons do not vary enough to correct for the misspecification in G , leading to the poor coverage results in figure 2.

In contrast to the baseline impulse responses, the augmented model accounts for the uncertainty in the GE Jacobian through the auxiliary parameter δ , preventing it from undercovering the true impulse response at any horizon. Table 3 displays the average widths of the impulse response credible intervals for short (1 to 4 periods ahead), medium (5 to 12 periods ahead), and long (13 to 40 periods ahead) horizons, computed across Monte Carlo iterations. The pattern is clear: across most horizons, the credible intervals produced by the augmented model are wider than their baseline model counterparts.

For output, the interval widths are largely similar across models and horizons, implying that the augmented model restores coverage by primarily shifting the impulse response for output relative to the baseline model with little loss of efficiency. For medium horizons the augmented model actually produces credible bands that are narrower than those of the baseline model. In the case of capital, the same pattern persists across horizons, although the ratio of the widths of the augmented

bands to those of the baseline bands is higher in general. Taken together, we can infer that we would be working with a severely biased impulse response estimate for output, and an estimate that underreports the underlying uncertainty for capital.

4.2 The One-asset HANK Model

4.2.1 Model Setup

This time around, households choose both consumption and labor supply, while facing the same idiosyncratic labor productivity shocks and borrowing constraints as before. Their preferences are separable in consumption and labor supply. Instead of owning capital, households instead own equal shares of the representative firm, from which they receive dividends each period. There is a risk-free bond that matures the next period, and households pay lump-sum taxes that are adjusted according to their current labor productivity. Following McKay et al. (2016), we adopt a regressive taxation rule that taxes only those households with the highest labor productivity. Household i faces the problem

$$\max_{\{c_t^i\}_{t \geq 0}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta_d^t \left(\frac{c_t^{1-\sigma} - 1}{1-\sigma} - \psi \frac{(\ell_t^i)^{1+\nu}}{1+\nu} \right) \right] \quad \text{s.t.} \quad c_t^i + \frac{a_{t+1}^i}{1+r_t} = a_t^i + w_t \ell_t^i e_t^i + d_t - \bar{\tau}(e_t^i) \tau_t, \quad a_{t+1}^i \geq 0$$

where e_t^i follows the same DGP as in the HA RBC model with m_e fixed at 0 and $\bar{\tau}(e)$ is our regressive taxation rule. w_t, r_t are the time t wage level and real bond yield, respectively.

There is a final goods producer that produces the final goods using a CES aggregator of a continuum of intermediate inputs indexed by $j \in [0, 1]$. The CES elasticity is given by $\epsilon > 1$, so that the demand for the j th intermediate good is given as usual by $\tilde{P}_{jt}^{-\epsilon} Y_t$, where Y_t is the final good output and $\tilde{P}_{jt}$ is the price of the j th intermediate good relative to the final good. The j th intermediate good is produced as $Y_{jt} = Z_t n_{jt}^d$, where Z_t is TFP and n_{jt} is labor demand. Intermediate good producers are monopolistically competitive and face a price adjustment cost given as $\frac{\epsilon}{2\kappa} \left(\log \frac{\tilde{P}_{jt}}{\tilde{P}_{j,t-1}} \right)^2 Y_t$, where $\kappa > 0$. Under symmetric equilibrium, where each intermediate good producer sets the same price, inflation π_t evolves according to the exact NKPC

$$\log(1 + \pi_t) = \kappa \left(\frac{w_t}{Z_t} - \frac{1}{\mu} \right) + \frac{1}{1+r_t} \mathbb{E}_t \left[\log(1 + \pi_{t+1}) \frac{Y_{t+1}}{Y_t} \right]$$

for constant markup $\mu = \frac{\epsilon}{\epsilon-1} > 1$. The parameter $\kappa > 0$ thus serves as the slope of the Phillips curve. Dividends d_t are then calculated as the profits of the firm net adjustment costs. The derivations of the NKPC and the dividend formula are given in the appendix.

The nominal interest rate i_t is determined via a Taylor rule set by the monetary policy authority:

$i_t = r_{ss} + \phi\pi_t + \phi_y(Y_t - Y_{ss}) + \epsilon_{R,t}$, where Y_{ss} is the steady state level of output, r_{ss} the steady state real interest rate and $\epsilon_{R,t}$ an exogenous monetary policy shifter. The Fisher relation $(1 + i_t) = (1 + r_t)(1 + \mathbb{E}_t[\pi_{t+1}])$ connects the nominal and real interest rates. The government supplies a fixed supply $B > 0$ of bonds each period, and chooses taxes τ_t in order to balance the budget each period: $\tau_t = r_t B + G_t$, where G_t is government spending and treated as an exogenous variable in this model. Finally, market clearing requires that $B = \int_0^1 a_t^i di$ and $\frac{Y_t}{Z_t} = \int_0^1 e_t^i \ell_t^i di$.

4.2.2 Steady State

The exogenous variables of this model are TFP Z_t , government spending G_t , and the monetary policy shifter $\epsilon_{R,t}$. We consider a zero-inflation steady state where $Z_{ss} = 1$, $G_{ss} = \epsilon_{R,ss} = 0$. There are ten endogenous macro variables in this model: output Y_t , consumption C_t , wage w_t , real interest rate r_t , inflation π_t , dividends d_t , taxation τ_t , labor demand and supply N_t^d, N_t^s , and aggregate savings A_t . We already know that $\pi_{ss} = 0$ by assumption, and since there is no capital in the model with steady state government spending equal to 0, $Y_{ss} = C_{ss} = N_{ss}$. By the bond market clearing condition, steady state savings are given by the bond supply $A_{ss} = B$, and by labor market clearing the steady state levels of N_t^d, N_t^s are both equal to N_{ss} . The NKPC, meanwhile, requires that $w_{ss} = \frac{1}{\mu}$ and $d_{ss} = (1 - w_{ss})N_{ss}$.

Steady state values for the interest rate and output, r_{ss} and N_{ss} , are the only objects left to pin down. Analogously to how the bond market clearing condition pinned down the steady state interest rate in the HA RBC model, the bond and labor market clearing conditions jointly pin down the pair (r_{ss}, N_{ss}) . Once again, for a given pair (r, N) of candidate steady state values, computing aggregate savings $A^s(r, N)$ involves first solving for the steady-state savings policy function and then deriving the stationary wealth distribution given (r, N) . Likewise, labor supply $N^s(r, N)$ requires the steady-state labor supply policy function integrated over individual savings and labor productivity with respect to the stationary wealth distribution. Due to the presence of two implicitly defined steady state values instead of just one, as in the HA RBC model, the search for the steady state is even more costly in the HANK model.

We collect the parameters involved in this step into $\eta = \{\beta_d, \mu, B\}$. The household heterogeneity parameters ρ_e, σ_e^2 are once again excluded because they are normally weakly identified or unidentified by aggregate macro data. Similarly, ψ is not independently calibrated but chosen to match the normalization $N_{ss} = 1$, and for this reason is also excluded from η . Finally, the risk aversion coefficient σ and the Frisch elasticity ν are again weakly identified from aggregate macro data and thus excluded.

4.2.3 The Wold Representation and Computing General Equilibrium Jacobians

Let $\log \text{TFP } \log(Z_t)$ follow a mean-zero AR(1) process with persistence ρ_z and error standard deviation σ_z . We assume that G_t and $\epsilon_{R,t}$ also similarly evolve according to AR(1) processes with persistence parameters ρ_G, ρ_R and error standard deviation σ_G, σ_R . These shock parameters are collected into the vector $\beta_N = \{\rho_z, \rho_G, \rho_R, \sigma_z, \sigma_G, \sigma_R\}$. The parameters of interest in the present exercise are $\beta_f = \{\phi, \phi_y, \kappa\}$, which affect the GE Jacobian but do not govern the evolution of the shock processes. Assuming that the economy at the steady state today will have returned to the steady state T periods after the arrival of some MIT shock today, we let $d\mathbf{Z}(\beta_N)$ be a $3T$ -dimensional vector describing how the exogenous variables $Z_t, G_t, \epsilon_{R,t}$ respond to a time 0 MIT shock, in that order. Following the convention established for the one-asset HANK model in Auclert et al. (2021), we let $\mathbf{G}(\beta_f, \eta)$ denote the $3T \times 3T$ GE Jacobian that collects the responses of the endogenous variables Y_t, w_t, π_t to changes in the exogenous variables, in this order.

In the appendix we explain how the GE Jacobian $\mathbf{G}(\beta_f, \eta)$ is recovered through the chain rule and implicit function theorem as

$$\mathbf{G}(\beta_f, \eta) = - \left[\mathbf{H}_U(\beta_f, \eta) + \mathbf{H}_{U_1} M_{1,U}(\beta_f, \eta) + \mathbf{H}_{U_2} (M_{2,U}(\eta) + M_{2,U_1}(\eta) M_{1,U}(\beta_f, \eta)) \right]^{-1} \\ \times \left[\mathbf{H}_Z(\beta_f, \eta) + \mathbf{H}_{U_1} M_{1,Z}(\eta) + \mathbf{H}_{U_2} M_{2,U_1}(\eta) M_{1,U}(\eta) \right].$$

$\mathbf{H}_{U_1}$ and $\mathbf{H}_{U_2}$ are sparse constant matrices, and $\mathbf{H}_U(\beta_f, \eta)$, $M_{1,U}(\beta_f, \eta)$, $\mathbf{H}_Z(\beta_f, \eta)$ and $M_{1,Z}(\eta)$ are sparse matrices whose non-zero elements can be expressed as very simple functions of the parameters β_f, η and the steady state values $r_{ss}(\eta), N_{ss}(\eta)$. The matrices $M_{2,U}(\eta)$ and $M_{2,U_1}(\eta)$ correspond to the household block of the problem, and for a given calibration of η they can be precomputed using the Fake news algorithm. As such, letting η^* be the chosen calibration of η , reconstructing $\mathbf{G}(\beta_f, \eta^*)$ for different values of β_f can be sped up considerably.

From the GE Jacobian $\mathbf{G}(\beta_f, \eta)$ for Y_t, w_t, π_t we can easily reconstruct the GE Jacobian $\mathbf{G}^o(\beta_f, \eta)$ for the observed macro variables Y_t, π_t, i_t . The DGP for the vector $X_t = (Y_t, \pi_t, i_t)$ is now given via the finite-order VMA process

$$X_t = X_{ss}(\eta) + \sum_{j=0}^{T-1} \Psi_j(\beta_f, \beta_N, \eta) \epsilon_{t-j}$$

where $X_{ss}(\eta)$ are the steady state values of X_t as a function of η and $\Psi_j(\beta_f, \beta_N, \eta) = (\mathbf{G}^o(\beta_f, \eta) d\mathbf{Z}(\beta_N))_j$ for $j \geq 0$.

Table 4: **Baseline calibration for one-asset HANK model** The discount rate β_d and the disutility of labor ψ are chosen to match the steady state targets of labor supply and the interest rate. σ_e^2 is given as $\Delta_e^2(1 - \rho_e^2)$.

<i>Steady State Parameters</i>			
σ	Relative Risk Aversion		2
ν	Frisch Elasticity of Labor		2
ρ_e	Persistence of Idiosyncratic Shock		0.966
Δ_e	Cross-sectional Standard Deviation of Idiosyncratic Shock		0.5
<i>Firm and Policy Parameters</i>			
μ	Markup		1.2
B	Bond Supply		5.6
<i>Steady State Targets</i>			
Y_{ss}	Steady State Output		1
r_{ss}	Steady State Real Interest Rate		0.01

4.2.4 Augmentation of the Likelihood

Let the baseline calibration of η be denoted by η^* . Then, the general augmented model is constructed as

$$X_t = X_{ss}(\eta^*) + \left(O_{3 \times k_{\beta_N}} \quad \frac{\partial X_{ss}(\eta^*)}{\partial \eta'} \right) \mathbf{B}_0 \delta + \sum_{j=0}^{T-1} \Psi_j^{\mathcal{A}}(\beta_f, \beta_N, \delta) \epsilon_{t-j}$$

where the expressions for $\Psi_j^{\mathcal{A}}(\beta_f, \beta_N, \delta)$ are given in equation (3.8) and $\mathbf{B}_0$ is the $k_{\bar{\eta}} \times k_{\delta}$ directional matrix. Once again, the augmented MA coefficients can be expressed as functions of $\mathbf{G}(\beta_f, \eta^*)$ and $d\mathbf{Z}(\beta_N)$ alongside the η -derivatives of $\mathbf{G}(\beta_f, \eta^*)$ and the derivatives of $d\mathbf{Z}(\beta_N)$. Differentiating $d\mathbf{Z}(\beta_N)$ with respect to β_N follows easily from the analytical formula for $d\mathbf{Z}(\beta_N)$. On the other hand, we can make use of the analytical formulas for $\mathbf{H}_U(\beta_f, \eta)$, $M_{1,U}(\beta_f, \eta)$, $\mathbf{H}_Z(\beta_f, \eta)$, $M_{1,Z}(\eta)$ alongside pre-computed numerical derivatives of the steady state values and household GE Jacobians $M_{2,U}(\eta)$ and $M_{2,U_1}(\eta)$ to quickly recompute $\frac{\partial \mathbf{G}(\beta_f, \eta^*)}{\partial \eta'}$ for different values of β_f . Details can be found in the appendix.

Inference for the augmented model then follows the exact same steps as in the HA RBC model. Since there are as many parameters of interest $\beta_f = \{\phi, \phi_y, \kappa\}$ as there are steady-state parameters in η , we directly reintroduce all three missing η -score directions.

4.2.5 Data, Calibration and Prior Selection

We estimate ϕ , ϕ_y and κ from the one-asset HANK model using the Smets and Wouters (2007) data from 1966:Q1 to 2004:Q4. Log output is linearly detrended and its steady state value is normalized to 1, so that the residual can be identified with the output gap. Inflation is obtained as the quarterly log difference of the GDP deflator, while we take 3-month moving averages of the monthly federal funds rate to transform it into a quarterly variable. All variables are expressed in terms of quarterly percentages, a scaling that is absorbed into the shock error standard deviations $\sigma_z, \sigma_G, \sigma_R$ and leaves the GE Jacobian unchanged. Our baseline calibration roughly follows those of Auclert et al. (2021), and are provided in table 4.

We use a beta prior with mean 0.5 and standard deviation 0.2 for all persistence parameters ρ_z, ρ_G, ρ_R , and an inverse gamma prior with mean 0.04 and standard deviation 0.4 for the error standard deviation parameters $\sigma_z, \sigma_G, \sigma_R$. For ϕ, ϕ_y and κ we use Gamma distributed priors with means and standard deviations (1.5, 0.25), (0.5, 0.25) and (0.1, 0.1) respectively. The prior on δ for the augmented model is chosen in the same manner as detailed in section 3.4.3, based on the implied prior covariance matrix Σ_η^* for η . Here we set Σ_η^* as the prior covariance matrix induced by the steady-state relationships

$$\tau_{ss} = r_{ss}B, \quad w_{ss} = \frac{1}{\mu}, \quad r_{ss} = r^*(\beta_d, \mu, B)$$


where $r^*(\cdot)$ represents the mapping from η to the steady state real interest rate. Steady state taxation τ_{ss} is chosen as a steady state target alongside the steady state prices since it is determined directly by the bond supply B . We once again construct Σ_η^* so that one prior standard deviation in the calibrated parameters corresponds, to first order, to a 15% deviation in each steady state target. The resulting prior covariance matrix under the present calibration is

$$\Sigma_\eta^* = \begin{pmatrix} 1.06 \times 10^{-5} & 3.62 \times 10^{-4} & 2.7 \times 10^{-3} \\ 3.62 \times 10^{-4} & 0.0032 & 0 \\ 2.7 \times 10^{-3} & 0 & 1.4112 \end{pmatrix},$$

implying prior standard deviations for β_d, μ and B equal to 0.0033, 0.18 and 1.19, respectively.

For sampling we once again use the RWMH algorithm with proposal variance given proportionally to the inverse negative Hessian of the posterior evaluated at the posterior mode.

Figure 3: **Posterior Densities of Model Parameters and Auxiliary Parameters** Below are the posterior densities of the model parameters ϕ, ϕ_y, κ obtained from the baseline and augmented models; and the posterior densities of the auxiliary parameters δ .



4.2.6 Estimating ϕ, ϕ_y, κ


The first row of figure 3 collects the posterior densities of ϕ, ϕ_y, κ obtained from the baseline and augmented models⁶. The two models roughly coincide when it comes to the posterior for the Taylor rule parameters ϕ, ϕ_y , with the augmented posterior for ϕ_y having a slightly fatter right tail compared to the baseline model. Meanwhile, the augmented posterior for κ is noticeably shifted to the right compared to its baseline posterior, with a much thicker right tail. Under the baseline model, there is little evidence supporting the hypothesis that $\kappa = 0.007$, while the hypothesis cannot be rejected outright under the augmented model. This suggests that the data prefers a calibration

⁶Compared to Auclert et al. (2021), the present model yields different estimates of ϕ_y and κ . In particular, κ is centered around a much smaller value than in Auclert et al. (2021). This discrepancy can be attributed to the differences in the specification of the exogenous variables in the one-asset HANK model. Auclert et al. (2021) include a markup shock in their model, whereas the present model instead contains a TFP shock that enters into the NKPC through real marginal cost $\frac{w_t}{Z_t}$. As such, in the one-asset HANK model with TFP shocks the data prefers a smaller value of the slope κ that prevents marginal cost fluctuations from exerting an overwhelming influence on inflation.

that implies lower nominal rigidity compared to the baseline calibration η^* .

The second row of figure 3 reveals that there is considerable evidence of misspecification in the calibration for the discount rate β_d , with the augmented model exhibiting a preference for smaller values of the discount rate compared to the baseline calibration. This implies a larger steady state real interest rate compared to the steady-state target 0.01, which is unsurprising since the sample mean of the federal funds rate used for estimation is larger than 0.01 and the augmented model uses its additional flexibility to better match the first and second moments of the aggregate macro variables. A smaller discount rate also explains why the posterior distribution of κ in the augmented model prefers larger values of κ compared to the baseline model—economic agents are less patient and therefore firms exhibit (slightly) less nominal rigidity.

Figure 4: Credible Bands and Posterior Medians for Impulse Responses Below are 95% credible bands around the posterior medians for the impulse response of each aggregate macro variable with respect to each of the exogenous shocks, from horizons 0 to 20.



4.2.7 Impulse Responses of Macro Aggregates

Given that we restore all three η -scores in the augmented model, it is possible to conduct inference on composite objects $H(\theta)$ that depend on all the model parameters by following the steps for case 1 in section 3. We focus on two objects: impulse responses of output, inflation and the nominal interest rate with respect to unit shocks in TFP, government spending and monetary policy, as well as the cumulative fiscal multiplier implied by the impulse response of output to unit government spending shocks.

Turning our attention first to the impulse responses, recall that the impulse response of the observed macro aggregates with respect to one standard deviation shocks in the exogenous variables is summarized in the $3T \times 3$ matrix $\mathbf{G}^o(\beta_f, \eta) d\mathbf{Z}^u(\beta_N)$, where $\mathbf{G}^o(\beta_f, \eta) \in \mathbb{R}^{3T \times 3}$ is the GE Jacobian for the observed macro variables and $d\mathbf{Z}^u(\beta_N) \in \mathbb{R}^{3T \times 3}$ traces the paths of the exogenous variables to unit shocks at time 0. Let $\{\beta_f^{(i)}, \beta_N^{(i)}, \delta^{(i)}\}_{1 \leq i \leq N}$ be a posterior sample of the parameters in the augmented model. Then,

$$\left\{ \left(\mathbf{G}^o(\beta_f^{(i)}, \eta^*) + \sum_{\ell=1}^{k_\eta} \frac{\partial \mathbf{G}^o(\beta_f^{(i)}, \eta^*)}{\partial \eta_\ell} \delta_\ell^{(i)} \right) d\mathbf{Z}^u(\beta_N^{(i)}) \right\}_{1 \leq i \leq N}$$

is a sample drawn from the posterior distribution of the impulse response of each macro variable for horizons 0 to T^7 .

The resulting median impulse responses and 95% credible intervals for horizons 0 to 20 are shown in figure 4. There is considerable overlap between the credible bands implied by the two models, indicating that the impulse responses they produce are roughly similar. The largest discrepancy can be found in the impulse responses of the macro variables with respect to a unit government spending shock—across all three variables the augmented model reports greater magnitude and uncertainty for the contemporaneous responses. Similarly, there is much more uncertainty associated with an output response to a unit TFP shock in the augmented model relative to the baseline model. This finding shows us that, even if

4.2.8 Cumulative Fiscal Multipliers

The unit impulse responses obtained in the preceding section can also be used to compute cumulative fiscal multipliers in response to a one-period shock to government spending. Following Ramey and Zubairy (2018), we define the H -period ahead cumulative fiscal multiplier as the ratio $\frac{\sum_{h=0}^H \Delta Y_{t+h}}{\sum_{h=0}^H \Delta G_{t+h}}$ of the cumulative change in output to the corresponding change in government spending. In the one-

⁷The derivatives of the sequence space Jacobians needed to obtain the derivatives of the GE Jacobians are pre-computed before running the augmented model, so obtaining a posterior sample of the impulse responses given a posterior sample of the parameters $(\beta_f, \beta_N, \delta)$ is computationally cheap.

Figure 5: **Posterior distributions of cumulative government spending multipliers** The left panel reports the posterior distribution of the cumulative multiplier through horizon $H = 12$, with vertical lines indicating posterior medians. The right panel reports posterior medians and 90% credible intervals at selected horizons.

Figure 6: Posterior Densities of $FM(12)$




Figure 7: Medians and Credible Intervals for $FM(H)$

	Baseline	Augmented
$H = 4$	0.310 [0.226, 0.394]	0.354 [0.265, 0.442]
$H = 8$	0.310 [0.226, 0.395]	0.353 [0.264, 0.442]
$H = 12$	0.311 [0.226, 0.396]	0.354 [0.263, 0.443]
$H = 20$	0.312 [0.226, 0.398]	0.355 [0.264, 0.447]

Notes: Entries report posterior medians, with 95% credible intervals in brackets.

asset HANK model, a unit government spending shock at time 0 shifts h -period ahead government spending by ρ_G^h , while the change in h -period ahead output is equal to the h -period ahead impulse response of output with respect to a unit government spending shock, denoted $\frac{\Delta Y_{t+h}}{\Delta G_t}$. Therefore, the cumulative fiscal multiplier is computed as

$$FM(H) = \frac{\sum_{h=0}^H \frac{\Delta Y_{t+h}}{\Delta G_t}}{\sum_{h=0}^H \rho_G^h}.$$

The panel on the right of figure 5 compares the 95% credible intervals and the posterior medians of the cumulative fiscal multipliers across different horizons obtained from the baseline and augmented models. In both models, the cumulative multiplier is relatively stable across horizons, with the augmented model reporting a higher value for $FM(H)$ uniformly across H . Interestingly, the augmented model seems to shift the posterior upward without drastically widening the posterior standard deviation, suggesting that in this particular case the augmented model mainly acts as a bias correction device. This is especially apparent in figure 5 that compares the posterior densities of $FM(H)$ for $H = 12$. This comparison suggests that local misspecification in η has a tangible effect on how we would conduct inference for $FM(H)$. For example, 0.45 is a value for $FM(12)$ that would be ruled out under the baseline model but seen as quite plausible if inference were to be based on the augmented model.

5 Conclusion

In this paper we introduce a method to robustify inference in DSGE models against local misspecification in steady-state parameters by linearizing the impulse responses around a baseline calibration.

We show, through an application to the Krusell and Smith (1998) model, that the augmented model restores correct coverage for both shock parameters under local calibration error whereas the baseline model exhibits drastically low coverage for the shock variance parameter. The robustness of the augmented model extends to the case where impulse responses are the quantities of interest. Once again, the baseline model fails to cover the true impulse response at short and long horizons, but the augmented model achieves close to the nominal coverage rate, albeit with slight overcoverage.

An application of model augmentation to a one-asset HANK model with TFP, monetary policy and government spending shocks provides evidence of local misspecification in the baseline calibration. This misspecification contaminates inference for the NKPC slope parameter κ , as seen through the much lower value of κ estimated from the baseline model compared to its augmented counterpart. The reported values for the auxiliary parameters δ imply that the data favors smaller values for all three steady-state parameters, the discount rate β_d , steady state markup μ and bond supply B , compared to their baseline calibrations. A failure to address potential misspecification also tangibly changes inference results for economic objects of interest such as fiscal multipliers.

The augmented model can therefore be positioned as a reasonable alternative to estimating the full DSGE model if we wish to either update the values of calibrated parameters based on the information contained in aggregate macro data without jointly estimating every parameter, or conduct likelihood-based inference in a manner that takes into account calibration uncertainty. This makes model augmentation especially well-suited for heterogeneous agent models, since aggregate macro data inherently contains little information on parameters governing household heterogeneity, and the presence of the wealth distribution as an additional state variable complicates the search for the steady state values of aggregate macro variables.

References

- Acharya, Sushant, William Chen, Marco Del Negro, Keshav Dogra, Aidan Gleich, Shlok Goyal, Ethan Matlin, Donggyu Lee, Reza Sarfaty, and Sikata Sengupta (2023). “Estimating HANK for central banks”. In: *FRB of New York Staff Report* 1071.
- Aiyagari, S Rao (1994). “Uninsured idiosyncratic risk and aggregate saving”. In: *The Quarterly Journal of Economics* 109.3, pp. 659–684.

- An, Sungbae and Frank Schorfheide (2007). “Bayesian analysis of DSGE models”. In: *Econometric reviews* 26.2-4, pp. 113–172.
- Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub (2021). “Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models”. In: *Econometrica* 89.5, pp. 2375–2408.
- Brock, William A and Leonard J Mirman (1972). “Optimal economic growth and uncertainty: the discounted case”. In: *Journal of Economic Theory* 4.3, pp. 479–513.
- Canova, Fabio (1994). “Statistical inference in calibrated models”. In: *Journal of Applied Econometrics* 9.S1, S123–S144.
- Chernozhukov, Victor, Juan Carlos Escanciano, Hidehiko Ichimura, Whitney K Newey, and James M Robins (2022). “Locally robust semiparametric estimation”. In: *Econometrica* 90.4, pp. 1501–1535.
- Cocci, Matthew D and Mikkel Plagborg-Møller (2025). “Standard errors for calibrated parameters”. In: *Review of Economic Studies* 92.5, pp. 2952–2978.
- Del Negro, Marco and Frank Schorfheide (2008). “Forming priors for DSGE models (and how it affects the assessment of nominal rigidities)”. In: *Journal of Monetary Economics* 55.7, pp. 1191–1208.
- Del Negro, Marco, Frank Schorfheide, Frank Smets, and Rafael Wouters (2007). “On the fit of new Keynesian models”. In: *Journal of Business & Economic Statistics* 25.2, pp. 123–143.
- Dridi, Ramdan, Alain Guay, and Eric Renault (2007). “Indirect inference and calibration of dynamic stochastic general equilibrium models”. In: *Journal of Econometrics* 136.2, pp. 397–430.
- Herbst, Edward P and Frank Schorfheide (2015). “Bayesian estimation of DSGE models”. In: *Bayesian Estimation of DSGE Models*. Princeton University Press.
- Iskrev, Nikolay (2010). “Local identification in DSGE models”. In: *Journal of Monetary Economics* 57.2, pp. 189–202.
- Kim, Seung H. (2026). *Model Augmentation for Locally Robust Bayesian Inference*. Working paper.
- Krusell, Per and Anthony A Smith Jr (1998). “Income and wealth heterogeneity in the macroeconomy”. In: *Journal of political Economy* 106.5, pp. 867–896.
- Liu, Laura and Mikkel Plagborg-Møller (2023). “Full-information estimation of heterogeneous agent models using macro and micro data”. In: *Quantitative Economics* 14.1, pp. 1–35.
- McKay, Alisdair, Emi Nakamura, and Jón Steinsson (2016). “The power of forward guidance revisited”. In: *American Economic Review* 106.10, pp. 3133–3158.
- Melosi, Leonardo, Hanno Kase, and Matthias Rottner (2022). “Estimating Nonlinear Heterogeneous Agents Models with Neural Networks”. In.

- Montiel Olea, José Luis and Mikkel Plagborg-Møller (2019). “Simultaneous confidence bands: Theory, implementation, and an application to SVARs”. In: *Journal of Applied Econometrics* 34.1, pp. 1–17.
- Ramey, Valerie A and Sarah Zubairy (2018). “Government spending multipliers in good times and in bad: evidence from US historical data”. In: *Journal of political economy* 126.2, pp. 850–901.
- Ríos-Rull, José-Víctor, Frank Schorfheide, Cristina Fuentes-Albero, Maxym Kryshko, and Raül Santaaulàlia-Llopis (2012). “Methods versus substance: Measuring the effects of technology shocks”. In: *Journal of Monetary Economics* 59.8, pp. 826–846.
- Smets, Frank and Rafael Wouters (2007). “Shocks and frictions in US business cycles: A Bayesian DSGE approach”. In: *American economic review* 97.3, pp. 586–606.
- Winberry, Thomas (2018). “A method for solving and estimating heterogeneous agent macro models”. In: *Quantitative Economics* 9.3, pp. 1123–1151.