Full-Information Estimation of Heterogeneous Agent Models Using Macro and Micro Data

Preliminary and Incomplete

Laura Liu — *Federal Reserve Board*

Mikkel Plagborg-Møller* — *Princeton University*

December 19, 2018

Abstract: We develop a generally applicable full-information inference method for heterogeneous agent models, combining aggregate time series data and repeated cross sections of micro data. To deal with unobserved aggregate state variables that affect cross-sectional distributions, we compute a numerically unbiased estimate of the model-implied likelihood function. Employing the likelihood estimate in a Markov Chain Monte Carlo algorithm, we obtain fully efficient and valid Bayesian inference.

*`laura.liu@frb.gov` (Liu) and `mikkelpm@princeton.edu` (Plagborg-Møller). We are grateful for helpful comments from Yoosoon Chang, Simon Mongey, Ulrich Müller, Frank Schorfheide, Neil Shephard, Tom Winberry, and seminar participants at Princeton and the 2018 Midwest Econometrics Group Conference. This paper does not necessarily reflect the views of the Board of Governors or the Federal Reserve System.

1 Introduction

Macroeconomic models with heterogeneous agents have exploded in popularity in recent years.[1] New micro data sets – including firm and household surveys, social security and tax records, and censuses – have exposed the empirical failures of traditional representative agent approaches. The new models not only improve the fit to the data, but also make it possible to meaningfully investigate the causes and consequences of inequality among households or firms along several dimensions, including endowments, financial constraints, age, size, location, etc.

So far, however, empirical work in this area has only been able to exploit limited features of the micro data sources that motivated the development of the new models. As emphasized by Ahn, Kaplan, Moll, Winberry, and Wolf (2017), the burgeoning academic literature has mostly calibrated model parameters and performed over-identification tests by matching a few empirical moments that are deemed important *a priori*. This approach may be highly inefficient, as it ignores that the models' implied macro dynamics and cross-sectional properties often fully determine the entire *distribution* of the observed macro and micro data. The failure to exploit the joint information content of macro and micro data stands in stark contrast to the well-developed inference procedures for estimating *representative* agent models using only macro data (Herbst and Schorfheide, 2016).

To exploit the full information content of macro and micro data, we develop a general technique to perform Bayesian inference in heterogeneous agent models. We assume the availability of aggregate time series data as well as repeated cross sections of micro data. Evaluation of the joint macro and micro likelihood function is complicated by the fact that model-implied cross-sectional distributions typically depend on unobserved aggregate state

[1] For references and discussion, see Krueger, Mitman, and Perri (2016), Ahn, Kaplan, Moll, Winberry, and Wolf (2017), and Kaplan and Violante (2018).

variables. To overcome this problem, we devise a way to compute a numerically unbiased estimate of the model-implied likelihood function of the macro and micro data. As argued by Andrieu, Doucet, and Holenstein (2010) and Flury and Shephard (2011), such an unbiased likelihood estimate can be employed in standard Markov Chain Monte Carlo (MCMC) procedures to generate draws from the fully efficient Bayesian posterior distribution given all available data.

The starting point of our analysis is the insight that existing solution methods for heterogeneous agent models directly imply what the distribution of the macro and micro data is, given structural parameters. These models are typically solved by imposing a flexible functional form (e.g., a discrete histogram or continuous parametric family) on the relevant cross-sectional distributions. The distributions are governed by time-varying unobserved state variables (e.g., moments). To calculate the model-implied likelihood, we decompose it into two parts. First, the likelihood of the *macro data* can be evaluated using standard linear state space methods (Mongey and Williams, 2017; Winberry, 2018), provided that the model is solved using the method of Reiter (2009), which linearizes with respect to the macro shocks but not the micro shocks. Second, the likelihood of the repeated cross sections of micro data, *conditional on the macro states*, can be evaluated by simply plugging into the assumed cross-sectional density. The key challenge that our method overcomes is that the econometrician typically does not directly observe the macro state variables. Instead, the observed macro time series are imperfectly informative about the underlying states.

Our procedure can loosely be viewed as a Bayesian version of a two-step approach: First we estimate the latent macro states from macro data, and then we construct the model-implied cross-sectional likelihood conditional on the macro states. More precisely, we obtain a *numerically unbiased* estimate of the likelihood by averaging the cross-sectional likelihood across repeated draws from the smoothing distribution of the hidden states given the macro data. We emphasize that, despite being based on a likelihood estimate, our method is fully

Bayesian and automatically takes into account all sources of uncertainty about parameters and states. An attractive computational feature is that evaluation of the micro part of the likelihood lends itself naturally to parallel computing. Hence, computation time scales well with the size of the data set.

We perform finite-sample valid and fully efficient Bayesian inference by plugging the unbiased likelihood estimate into a standard MCMC algorithm. The generic arguments of Andrieu, Doucet, and Holenstein (2010) and Flury and Shephard (2011) imply that the ergodic distribution of the MCMC chain is the full-information posterior distribution that we would have obtained if we had known how to evaluate the exact likelihood function (not just an unbiased estimate of it). This is true no matter how many smoothing draws are used to compute the unbiased likelihood estimate. In principle, we may use any MCMC algorithm that relies only on evaluating (the unbiased estimate of) the posterior density, such as Random Walk Metropolis-Hastings or Sequential Monte Carlo.

We illustrate the joint inferential power of macro and micro data through two numerical examples: a heterogeneous household model (Krusell and Smith, 1998) and a heterogeneous firm model (Khan and Thomas, 2008). In both cases we assume that the econometrician observes certain standard macro time series as well as intermittent repeated cross sections of, respectively, (i) household employment and income and (ii) firm inputs. Using simulated data, and given flat priors, we show that our method accurately recovers the true structural model parameters. Importantly, for several structural parameters, the micro data reduces the length of posterior credible intervals by more than 40%, relative to inference that exploits only the macro data. In fact, we give examples of parameters that can only be identified if micro data is available.

Empirical application: TO BE COMPLETED.

Literature. We appear to be the first to exploit the full, combined information content available in macro and micro data for estimating heterogeneous agent models. We build on the idea of Mongey and Williams (2017) and Winberry (2018) to estimate heterogeneous agent models from the linear state space representation obtained from the Reiter (2009) model solution approach. Whereas Winberry (2018) exploits only macro data for estimation, Mongey and Williams (2017) also track a particular cross-sectional moment over time. In contrast, we exploit the entire model-implied likelihood function given repeated micro cross sections. Hahn, Kuersteiner, and Mazzocco (2018) develop asymptotics for estimation using interdependent micro and macro data sets, but their approach is not applicable to our setting since the exact micro likelihood in a heterogeneous agent model is usually not available in closed form, due to the need to integrate out unobserved state variables. Unlike the rich microeconometric literature on heterogeneous agent models (Arellano and Bonhomme, 2017), our work explicitly seeks to estimate the deep parameters of a macro model by also incorporating macro time series data. Chang, Chen, and Schorfheide (2018) propose a reduced-form approach to estimating the feedback loop between aggregate time series and heterogeneous micro data. In likelihood estimation of *representative* agent models, micro data has mainly been used to inform the prior, as in Chang, Gomes, and Schorfheide (2002).

Outline. Section 2 shows that heterogeneous agent models imply a statistical model for the macro and micro data. We then present our method for computing an unbiased likelihood estimate that is used to perform efficient Bayesian inference. Section 3 illustrates the inferential power of combining macro and micro data using two simple numerical examples, a heterogeneous household model and a heterogeneous firm model. Section 4 contains the empirical application (TO BE COMPLETED). Section 5 concludes and discusses possible future research directions. Appendix A contains supplementary material.

2 Method

We first describe how heterogeneous agent models imply a statistical model for the macro and micro data, then we introduce our procedure for computing a numerically unbiased estimate of the likelihood, and finally we describe the posterior sampling procedure.

2.1 Set-up

We assume the availability of macro time series as well as repeated cross sections of micro data, as summarized in Figure 1. Let $\mathbf{x} \equiv \{x_t\}_{1 \leq t \leq T}$ denote the vector of observed time series data (e.g., real GDP growth), with sample size T. At a subset $\mathcal{T} \subset \{1, 2, \ldots, T\}$ of time points we additionally observe the micro data $\mathbf{y} \equiv \{y_{i,t}\}_{1 \leq i \leq N_t, t \in \mathcal{T}}$, where $y_{i,t}$ is a vector (e.g., the consumption choices of household i or the balance sheet of firm i). At each time t, $\{y_{i,t}\}_{1 \leq i \leq N_t}$ is sampled i.i.d. from the population distribution *conditional* on some macro state vector z_t. Since the micro data arises from repeated cross sections, we assume that, conditional on $\mathbf{z} \equiv \{z_t\}_{1 \leq t \leq T}$, the micro data is independent across time t. The state variable vector z_t is chosen sufficiently rich so that $y_{i,t}$ is independent of $(\mathbf{x}, \{z_\tau\}_{\tau \neq t})$ conditional on z_t. The states may be partially unobserved, i.e., $z_t \neq x_t$ in general. This last fact complicates the evaluation of the exact likelihood function, as shown below.

Given the structural parameter vector θ, the heterogeneous agent model implies functional forms for the macro observation density $p(x_t \mid z_t, \theta)$, the macro state transition density $p(z_t \mid z_{t-1}, \theta)$, and the micro sampling density $p(y_{i,t} \mid z_t, \theta)$. Depending on the model solution method, $p(y_{i,t} \mid z_t, \theta)$ is often given by a flexible parametric density indexed by certain time-varying parameters ψ_t, in which case ψ_t is included in z_t. A special case is when the model is solved by approximating cross-sectional distributions with discrete histograms; here ψ_t denotes the probability masses of the discrete distributions.

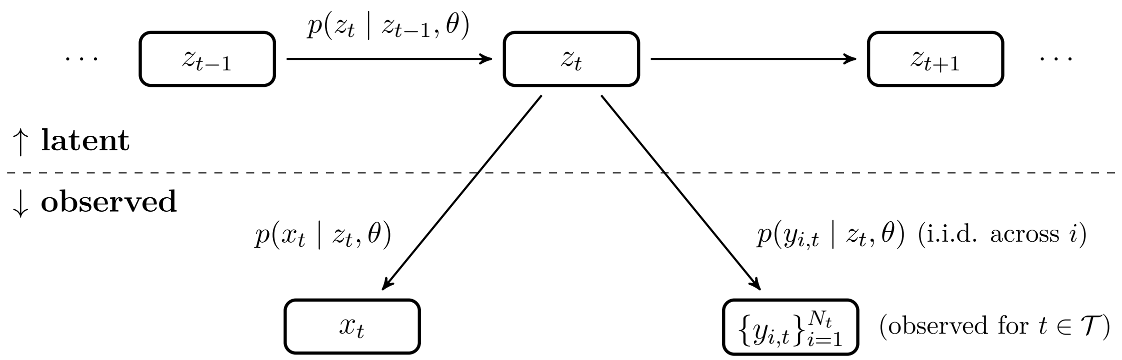

Figure 1: Diagram of the distribution of the macro and micro data implied by a heterogeneous agent model. The state vector z_t includes any time-varying parameters that govern the cross-sectional distribution $p(y_{i,t} \mid z_t, \theta)$.

2.2 Unbiased likelihood estimate

Our likelihood estimate is based on decomposing the joint likelihood into a macro part and a micro part (conditional on the macro data):

$$
\begin{aligned}
p(\mathbf{x}, \mathbf{y} \mid \theta) &= \overbrace{p(\mathbf{x} \mid \theta)}^{\text{macro}} \, \overbrace{p(\mathbf{y} \mid \mathbf{x}, \theta)}^{\text{micro}} \\
&= p(\mathbf{x} \mid \theta) \int p(\mathbf{y} \mid \mathbf{z}, \theta) p(\mathbf{z} \mid \mathbf{x}, \theta) \, d\mathbf{z} \\
&= p(\mathbf{x} \mid \theta) \int \prod_{t \in \mathcal{T}} \prod_{i=1}^{N_t} p(y_{i,t} \mid z_t, \theta) p(\mathbf{z} \mid \mathbf{x}, \theta) \, d\mathbf{z}. \qquad (1)
\end{aligned}
$$

The macro part of the likelihood is easily computable. If the representative agent model is solved using the popular linearization method of Reiter (2009), the (approximate) heterogeneous agent model can be written as a linear state space model in the macro observables and macro shocks:

$$
\begin{aligned}
x_t &= S(\theta) z_t, \\
z_t &= A(\theta) z_{t-1} + B(\theta) \varepsilon_t,
\end{aligned}
$$

where the matrices $S(\cdot)$, $A(\cdot)$, and $B(\cdot)$ are complicated functions of the structural parameters θ and of the model's micro heterogeneity. Hence, assuming i.i.d. Gaussian macro shocks ε_t, the macro part of the likelihood $p(\mathbf{x} \mid \theta)$ can be obtained in closed form from the Kalman filter (as discussed by Ahn, Kaplan, Moll, Winberry, and Wolf, 2017; Mongey and Williams, 2017; Winberry, 2018). This is computationally cheap even in models with many state variables and/or observables.

The novelty of our approach is that we compute an unbiased estimate of the micro likelihood conditional on the macro data. Although the integral in expression (1) cannot be computed analytically in realistic models, we can obtain a *numerically unbiased* estimate of the integral by random sampling:

$$\int \prod_{t \in \mathcal{T}} \prod_{i=1}^{N_t} p(y_{i,t} \mid z_t, \theta) p(\mathbf{z} \mid \mathbf{x}, \theta)\, d\mathbf{z} \approx \frac{1}{J} \sum_{j=1}^{J} \prod_{t \in \mathcal{T}} \prod_{i=1}^{N_t} p(y_{i,t} \mid z_t = z_t^{(j)}, \theta), \tag{2}$$

where $\mathbf{z}^{(j)} \equiv \{z_t^{(j)}\}_{1 \le t \le T}$, $j = 1, \dots, J$, are draws from the joint smoothing density $p(\mathbf{z} \mid \mathbf{x}, \theta)$ of the latent states. Again using the Reiter (2009) linearized model solution, the Kalman smoother can be used to produce these state smoothing draws with little computational effort (e.g., Durbin and Koopman, 2002). As the number of smoothing draws $J \to \infty$, the likelihood estimate converges to the exact likelihood, but we show below that finite J is sufficient for our purposes, as we rely only on the numerical unbiasedness of the likelihood estimate, not its consistency.

Our likelihood estimate can loosely be interpreted as arising from a two-step approach: First we estimate the states from the macro data, and then we plug the state estimates $z_t^{(j)}$ into the micro sampling density. However, we will argue next that the unbiased likelihood estimate makes it possible to perform valid Bayesian inference that fully takes into account all sources of uncertainty about states and parameters.

The expression on the right-hand side of the likelihood estimate (2) is parallelizable over smoothing draws j, time t, and/or individuals i. Thus, given the right computing environment, the computation time of our method scales well with the dimensions of the micro data.[2] This is particularly helpful in models where evaluation of the micro sampling density involves numerical integration, as in the household model in Section 3.1 below.

2.3 Posterior sampling

Now that we have a numerically unbiased estimate of the likelihood, we can plug it into any generic MCMC procedure to obtain draws from the posterior distribution, given a choice of prior density. We may simply pretend that the likelihood estimate is exact and run the MCMC algorithm as we otherwise would, as explained by Andrieu, Doucet, and Holenstein (2010) and Flury and Shephard (2011). Despite the simulation error in estimating the likelihood, the ergodic distribution of the MCMC chain will equal the fully efficient posterior distribution $p(\theta \mid \mathbf{x}, \mathbf{y})$. This is true no matter how small the number J of smoothing draws is. Still, the MCMC chain will typically exhibit better mixing if J is moderately large so that proposal draws are not frequently rejected merely due to numerical noise. In principle, we can use any generic MCMC method that requires only the likelihood and prior density as inputs, including Metropolis-Hastings and Sequential Monte Carlo approaches.[3]

3 Numerical illustrations

We now present two proof-of-concept simulation exercises that demonstrate how the combination of macro and micro data can sharpen structural inference. We deliberately keep the dimensionality of these inference problems small in order to focus attention on the core

[2] If the micro data were panel data instead of repeated cross sections, the scope for parallelization would be significantly reduced due to serial dependence, as briefly discussed in Section 5.

[3] See Herbst and Schorfheide (2016) for a review of such generic MCMC methods.

workings of our procedure. The first exercise estimates a heterogeneous household model in the spirit of Krusell and Smith (1998). The second exercise considers a heterogeneous firm model as in Khan and Thomas (2008). In both cases, we apply our inference procedure to data simulated from the structural model in order to examine the inferential accuracy relative to the true parameters. Our main take-away is that micro data can dramatically reduce the length of posterior credible intervals relative to using only macro data, and in some cases micro data is essential for identification.

3.1 Heterogeneous household model

In our first example, we estimate a version of the Krusell and Smith (1998) model of heterogeneous households exposed to idiosyncratic employment risk and aggregate productivity shocks. Households can only save in a non-state-contingent bond and are subject to a borrowing constraint. Since the core of the model is entirely standard, we refer to the exposition in Winberry (2016).

We add two features to the baseline model. First, the observed macro variable equals log real GDP plus independent measurement error. Second, individual households are each endowed with a permanent labor productivity level λ_i, which is drawn at the beginning of time from a lognormal distribution with mean parameter $E[\log \lambda_i] = \mu_\lambda \leq 0$ and variance parameter chosen such that $E[\lambda_i] = 1$. We rig the model so that the household problem scales linearly with λ_i; thus, macro aggregates are unaffected by the dispersion of individual productivity levels. As discussed below, μ_λ is an example of a parameter that is only identifiable if micro data is available.

We aim to estimate the households' discount factor β, the standard deviation σ_e of the measurement error in log output, and the individual productivity heterogeneity parameter μ_λ. All other parameters are assumed known for simplicity.

Calibration, simulated data, prior. We assume that the econometrician observes aggregate data on log output with measurement error, as well as repeated cross sections of household employment status and after-tax income (after-tax labor income plus unemployment benefits plus interest on assets).

We adopt the annual parameter calibration in Winberry (2016). In particular, $\beta = 0.96$. We choose the true measurement error standard deviation σ_e so that about 20% of the variance of observed log output is due to measurement error, yielding $\sigma_e = 0.02$.[4] The individual heterogeneity parameter μ_λ is chosen to be -0.25, implying that the model's cross-sectional 20th to 90th percentile range of log after-tax income roughly matches the range in U.S. data (Piketty, Saez, and Zucman, 2018, Table I).

Using this calibration, we simulate a single data set with $T = 100$ periods of macro data, and micro data consisting of 1,000 households observed at each of the ten time points $t = 10, 20, 30, \ldots, 100$. The data is simulated using the same approximate model solution method as is used to compute the unbiased likelihood estimate, see below.

Finally, we choose the prior on $(\beta, \sigma_e, \mu_\lambda)$ to be flat in the natural parameter space.

Computation. Following Winberry (2016, 2018), we solve the model using a Dynare implementation of the Reiter (2009) method.[5] This allows us to use Dynare's Kalman filter/smoother procedures. Evaluation of the micro sampling density $p(y_{i,t} \mid z_t, \theta)$ requires evaluation of an integral, since household income is a convolution of individual productivity λ_i and assets. Appendix A.1.1 describes how we numerically compute this integral and how we deal with the point mass in the asset distribution at the borrowing constraint. We average the likelihood across $J = 500$ smoothing draws.

[4] One possible real-world interpretation of the measurement error is that it represents the statistical uncertainty in estimating the natural rate of output (recall that the model abstracts from nominal rigidities).

[5] We thank Tom Winberry for making the code available on his website.

Heterogeneous household model: Posterior density

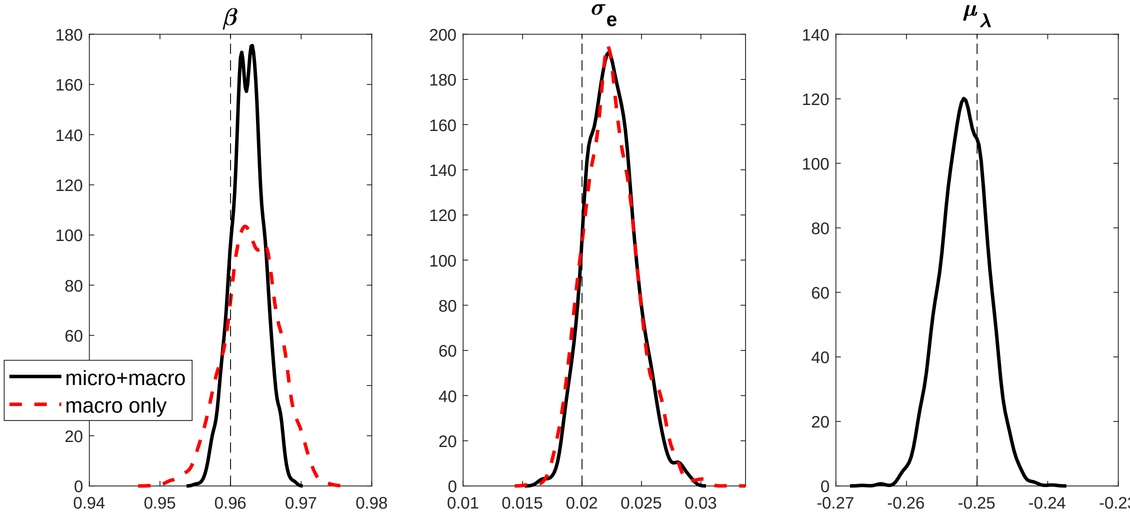

Figure 2: Posterior densities with (solid curve) and without (dashed curve) conditioning on the micro data. Both sets of results use the same simulated data set. Vertical dashed lines indicate true parameter values. Posterior density estimates from the 8,000 retained MCMC draws using Matlab's `ksdensity` function with default settings. The third display omits the macro-only results, since μ_λ is not identified in that case.

For clarity, our MCMC algorithm is a basic Random Walk Metropolis-Hastings algorithm with tuned proposal covariance matrix and adaptive step size (Atchadé and Rosenthal, 2005). We start the algorithm at parameter values far from the truth.[6] Appendix A.1.2 shows that the MCMC chain exhibits reasonably fast mixing. Using parallel computing on 28 cores, likelihood evaluation takes about as long as Winberry's (2016) procedure for computing the model's steady state.

Results. Figure 2 depicts the posterior densities of the three parameters after discarding a burn-in sample. The posterior is concentrated close to the true values of the three parameters. The figure also shows the posterior density *without* conditioning on the micro data. Ignoring the micro data leads to substantially less accurate inference about β, with the 90% equal-

[6]Specifically, $\beta = 0.9$, $\sigma_e = 0.04$, and $\mu_\lambda = -1$.

tailed credible interval for this parameter being 68% wider. Even more starkly, μ_λ can only be identified from the cross section (by construction). In contrast, almost all the information about σ_e comes from the macro data.

3.2 Heterogeneous firm model

As our second example, we estimate a version of the heterogeneous firm model of Khan and Thomas (2008). Firms are subject to idiosyncratic and aggregate productivity shocks, and they face non-convex investment adjustment costs. In addition to the productivity shock, there is a second aggregate shock that affects investment efficiency. We refer to the exposition of this standard model in Winberry (2018).

We aim to estimate the AR(1) parameter ρ_ε and innovation standard deviation σ_ε of the firms' idiosyncratic productivity process. Khan and Thomas (2008) showed that these parameters have little impact on the aggregate macro implications of the model; hence, micro data is needed. All other structural parameters are assumed known for simplicity.

Calibration, simulated data, prior. We adopt the annual calibration of Winberry (2018), which in turn follows Khan and Thomas (2008). However, we make an exception in setting the true idiosyncratic AR(1) parameter $\rho_\varepsilon = 0.53$, following footnote 5 in Khan and Thomas (2008).[7] We then set $\sigma_\varepsilon = 0.0364$, so that the variance of the idiosyncratic (log) productivity process is unchanged from the baseline calibration in Khan and Thomas (2008) and Winberry (2018). The macro implications of our calibration are virtually identical to the baseline in Khan and Thomas (2008), as those authors note.

We assume that the econometrician observes time series on aggregate output and investment, as well as repeated cross sections of micro data on firms' capital and labor inputs.

[7] This avoids numerical issues that arise when solving the model for high degrees of persistence.

We simulate a single data set with sample size $T = 50$, while micro cross sections of size $N = 1000$ are observed at each of the five time points $t = 10, 20, \ldots, 50$. Unlike in Section 3.1, we do not add measurement error to the macro observables.

The prior on $(\rho_\varepsilon, \sigma_\varepsilon)$ is chosen to be flat in the natural parameter space.

Computation. As in the previous numerical example, we solve and simulate the model using the Winberry (2018) Dynare solution method.[8] Computation of the micro sampling density is simple in this example, since – conditional on macro states – the micro observables (capital and labor) are log-linear transformations of the micro state variables (capital and idiosyncratic productivity).[9] We use $J = 500$ smoothing draws to compute the unbiased likelihood estimate. The MCMC routine is the same as in Section 3.1. We start the algorithm at parameter values far away from the truth.[10] Appendix A.1.3 shows that the MCMC chain mixes well. Likelihood evaluation using 12 parallel cores is several times faster than computing the model's steady state.

Results. Despite the finding in Khan and Thomas (2008) that macro data is essentially uninformative about the idiosyncratic productivity parameters, these are accurately estimated when the micro data is used also. Figure 3 shows the posterior densities of ρ_ε and σ_ε. The posterior distribution of each parameter is concentrated close to the true parameter values. We cannot meaningfully compare these results with inference that relies only on macro data, since the macro likelihood is almost entirely flat as a function of $(\rho_\varepsilon, \sigma_\varepsilon)$,

[8] We again used the code kindly made available on Tom Winberry's website.

[9] Moreover, it is easy to simulate the micro data, since Winberry (2018) shows that a multivariate normal distribution (with parameters pinned down by the model's equilibrium conditions) provides an accurate approximation to the distribution of idiosyncratic state variables.

[10] Specifically, $\rho_\varepsilon = 0.7$ and $\sigma_\varepsilon = 0.02$.

Heterogeneous firm model: Posterior density

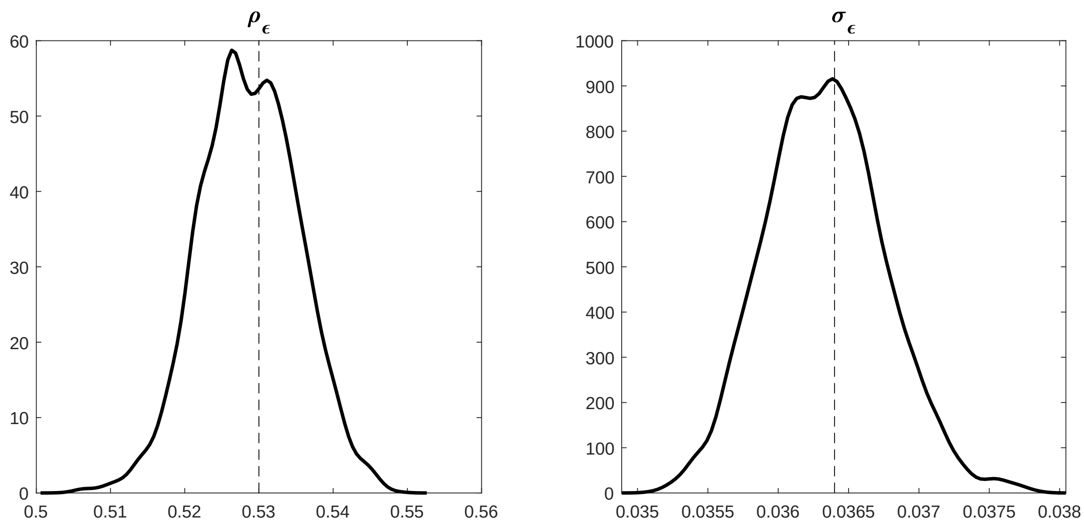

Figure 3: Posterior densities. Vertical dashed lines indicate true parameter values. Posterior density estimates from the 8,000 retained MCMC draws using Matlab's `ksdensity` function with default settings.

consistent with Khan and Thomas (2008).[11] Thus, micro data is essential to inference about these parameters.

3.3 Take-aways

Our two examples demonstrate the versatility of our method and the utility of efficient use of the micro data. When combined with standard macro time series, micro data can substantially sharpen inference about certain parameters, such as the discount factor in the heterogeneous household model. Indeed, some parameters may only be identified if micro data is available in addition to the macro data, such as the idiosyncratic productivity parameters in both the household and firm models.[12] On the other hand, there do exist

[11] In our MCMC chain, the standard deviation (after burn-in) of the log macro likelihood across all Metropolis-Hastings proposals of the parameters is only 0.10, while it is 10.3 for the log micro likelihood.

[12] Adding additional macro observables to our examples does not change this conclusion, since all macro variables in the two models are almost perfectly collinear with the baseline observables (modulo measurement

parameters that are mostly identified off macro data, such as the measurement error variance in the heterogeneous household model. Although the results reported here are derived from a single simulated data set in each example, we have confirmed that additional simulations yield qualitatively identical conclusions.

4 Empirical application

TO BE COMPLETED.

5 Conclusion

We develop a method to exploit the full information content in macro and micro data when estimating heterogeneous agent models. As we demonstrate through economic examples, the *joint* information content available in micro and macro data is often much larger than in either of the two separate data sets. Our inference procedure can loosely be interpreted as a two-step method: First we estimate the underlying macro states from macro data, and then we evaluate the likelihood by plugging into the cross-sectional sampling densities given the estimated states. However, our method delivers finite-sample valid and fully efficient Bayesian inference that takes into account all sources of uncertainty about parameters and states. The procedure is most computationally attractive when the model is solved using the popular linearization technique of Reiter (2009). The computation time of our procedure scales well with the size of the data set, as the method lends itself to parallel computing.

Our method is generally applicable whenever macro data and repeated cross sections of micro data are available. Our technique may serve as a formal and efficient alternative to the currently popular but *ad hoc* calibration (moment matching) approach. This

error). This is because we use as many macro observables as there are macro shocks.

would be analogous to the now widespread use of full-information likelihood methods for estimating *representative* agent models (Herbst and Schorfheide, 2016). From a policy perspective, Dynamic Stochastic General Equilibrium models have long played an important role in central bank decision making (Del Negro, Eusepi, Giannoni, Sbordone, Tambalotti, Cocci, Hasegawa, and Linder, 2013). Our technique represents a first step towards extending existing *representative* agent policy models to incorporate empirically relevant sources of microeconomic heterogeneity.

Our research suggests several possible avenues for future research. First, it would be interesting to extend our method to allow for panel data. Unlike in the case of repeated cross sections, panel data complicates the evaluation of the micro likelihood due to the intricate serial dependence of individual decisions. Second, since our method works for a wide range of generic MCMC posterior sampling procedures, it would be interesting to investigate the scope for improving on the Random Walk Metropolis-Hastings algorithm that we use for conceptual clarity in our examples. Third, more empirical work is needed to gauge the usefulness of micro data in large-scale structural models.

A Appendix

A.1 Numerical illustrations: Details

A.1.1 Heterogeneous household model: Micro sampling density

Here we describe how we evaluate the micro sampling density of household income, given employment status, in the simulations in Section 3.1. In the Winberry (2016) model, the after-tax income of household i at time t is given by

$$s_{i,t} = \lambda_i(c_{i,t} + R_t a_{i,t}),$$

where R_t is the gross real interest rate, $a_{i,t}$ is the household's assets relative to its permanent productivity λ_i, and $c_{i,t}$ depends on the wage rate w_t, the unemployment benefit level, the tax rate, and the employment status of the individual (see the budget constraint on p. 3 of Winberry, 2016). Conditional on the macro states r_t and w_t and on the individual employment status, the randomness in $s_{i,t}$ is coming from λ_i and $a_{i,t}$. The latter has a point mass at zero (the borrowing constraint in the calibration). One can show that the density of $s_{i,t}$ (conditional on macro states and employment status $e_{i,t}$) is given by

$$\frac{d}{ds}P(s_{i,t} \leq s \mid e_{i,t}, z_t, \theta) = \pi_t \frac{f(\frac{s}{c_{i,t}} \mid \mu_\lambda)}{c_{i,t}} + (1 - \pi_t) \int_0^\infty g(a \mid \psi_t) \frac{f(\frac{s}{c_{i,t}+R_t a} \mid \mu_\lambda)}{c_{i,t} + R_t a} \, da. \tag{3}$$

Here π_t is the point mass at 0 for the asset distribution, $f(\cdot \mid \mu_\lambda)$ is the lognormal density function for λ_i, and $g(\cdot \mid \psi_t)$ is the model-implied cross-sectional density for relative assets $a_{i,t}$ conditional on them being positive. The latter density is governed by certain time-varying aggregate states ψ_t, which are pinned down by the model's equilibrium conditions. π_t and ψ_t are part of the state vector z_t, along with w_t (which determines $c_{i,t}$) and R_t. In

Winberry (2016) the functional form for $g(\cdot \mid \psi_t)$ is chosen to be a flexible exponential family of densities, see his equation 1 on p. 5.

We evaluate the numerical integral in (3) by combining numerical integration and interpolation. First, we use a univariate numerical integration routine to evaluate the integral on an equi-spaced grid of values for $\log s$. To evaluate the expression at values of s outside this grid, we use cubic spline interpolation. In practice, a small number of grid points is sufficient in this application, since the density (3) is a very smooth function of s.

To simulate micro data from the cross-sectional distribution, we use the inverse probability transform. For this purpose, we compute the cumulative distribution function using numerical integration.

A.1.2 Heterogeneous household model: MCMC diagnostics

We now document that the MCMC chain in Section 3.1 exhibits reasonably fast mixing. Figures 4 and 5 depict the trace plot and autocorrelation function of the MCMC draws, respectively. The trace plot shows that the chain quickly locates the posterior mode. The autocorrelation plot shows that the dependence of the chain dies out after about 150 lags for β and σ_e, and much quicker for μ_λ. The average acceptance rate after the 1,000 burn-in draws is 29%, which is close to the target of 30% that we set in the Atchadé and Rosenthal (2005) adaptive step size algorithm.

A.1.3 Heterogeneous firm model: MCMC diagnostics

The MCMC procedure also works well in the heterogeneous firm example in Section 3.2. Figures 6 and 7 depict the trace plot and autocorrelation function of the MCMC draws, respectively. As in the household model, the chain quickly locates the posterior mode and achieves moderately fast mixing afterwards. The average acceptance rate after the 2,000 burn-in draws is 31%, which is close to the target of 30% for the adaptive step size algorithm.

Heterogeneous household model: MCMC chain

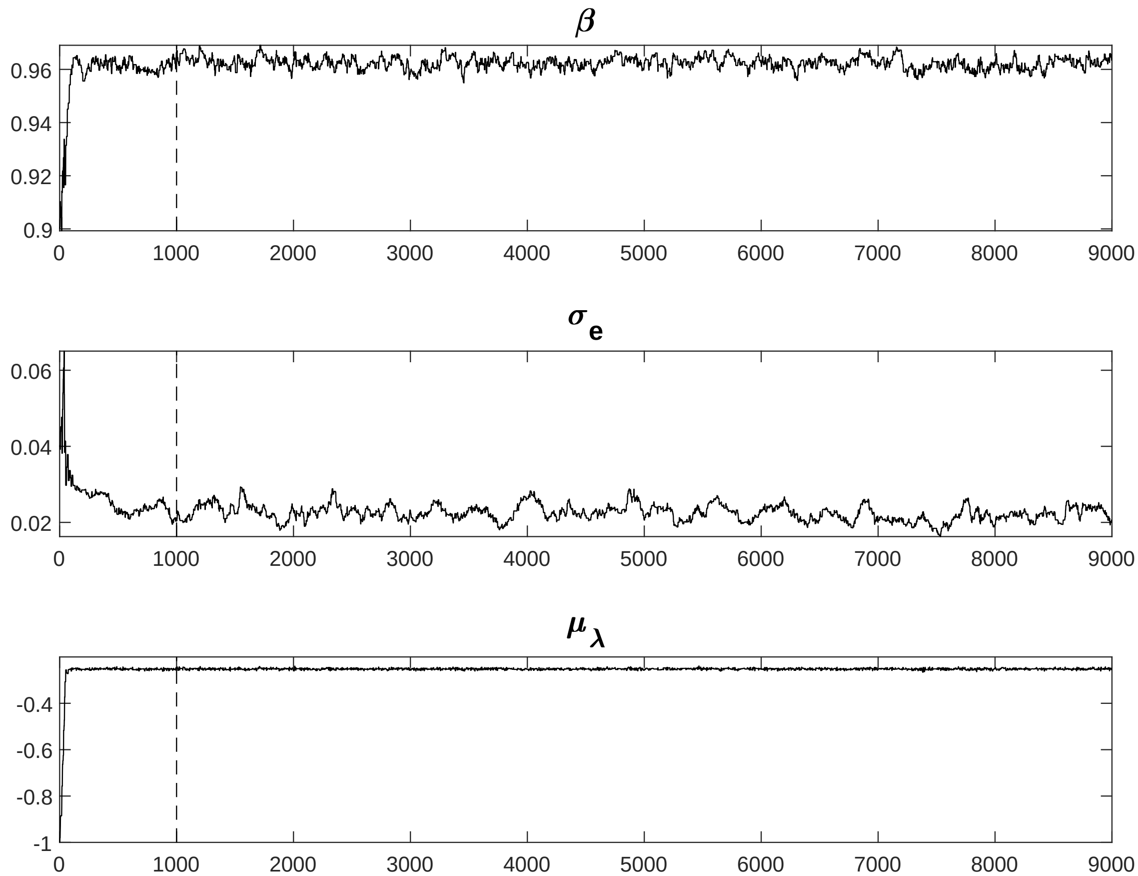

Figure 4: Trace plot of the MCMC chain. Vertical dashed line indicates the end of the burn-in sample.

Heterogeneous household model: MCMC autocorrelations

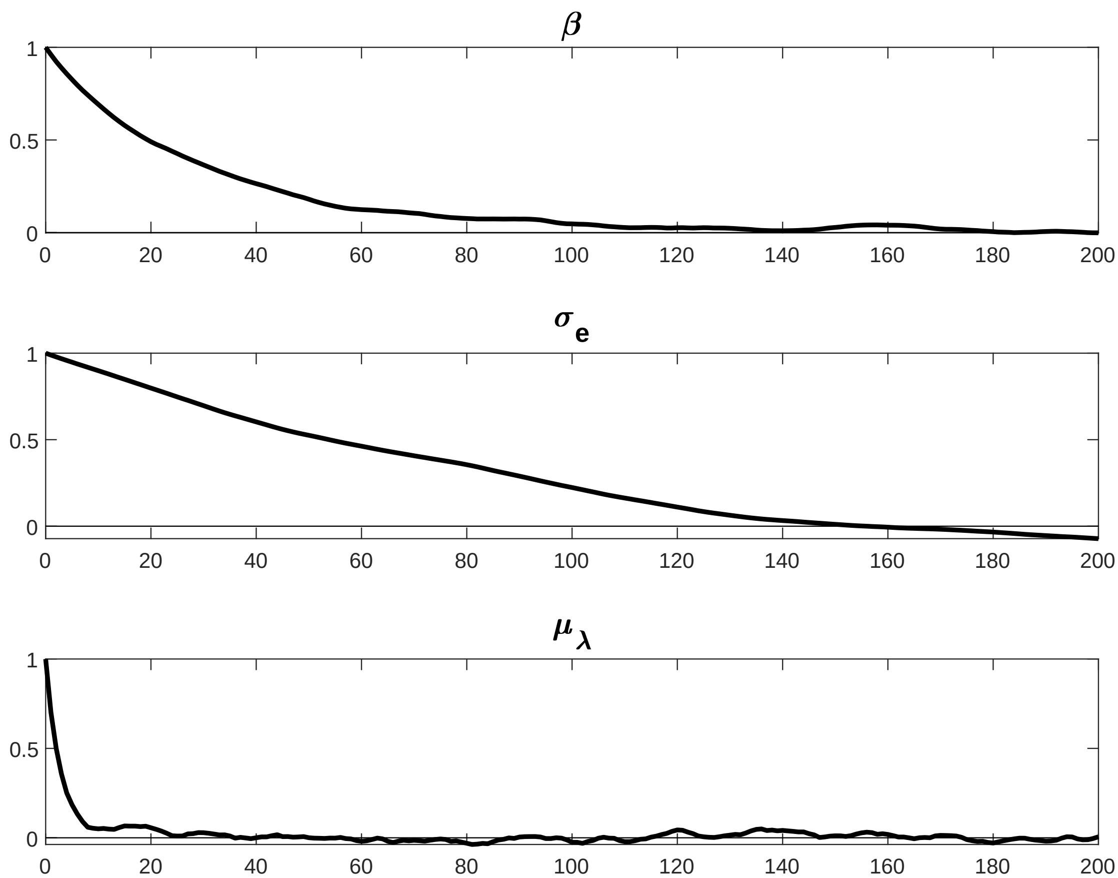

Figure 5: Autocorrelation function of MCMC draws, discarding burn-in sample.

Heterogeneous firm model: MCMC chain

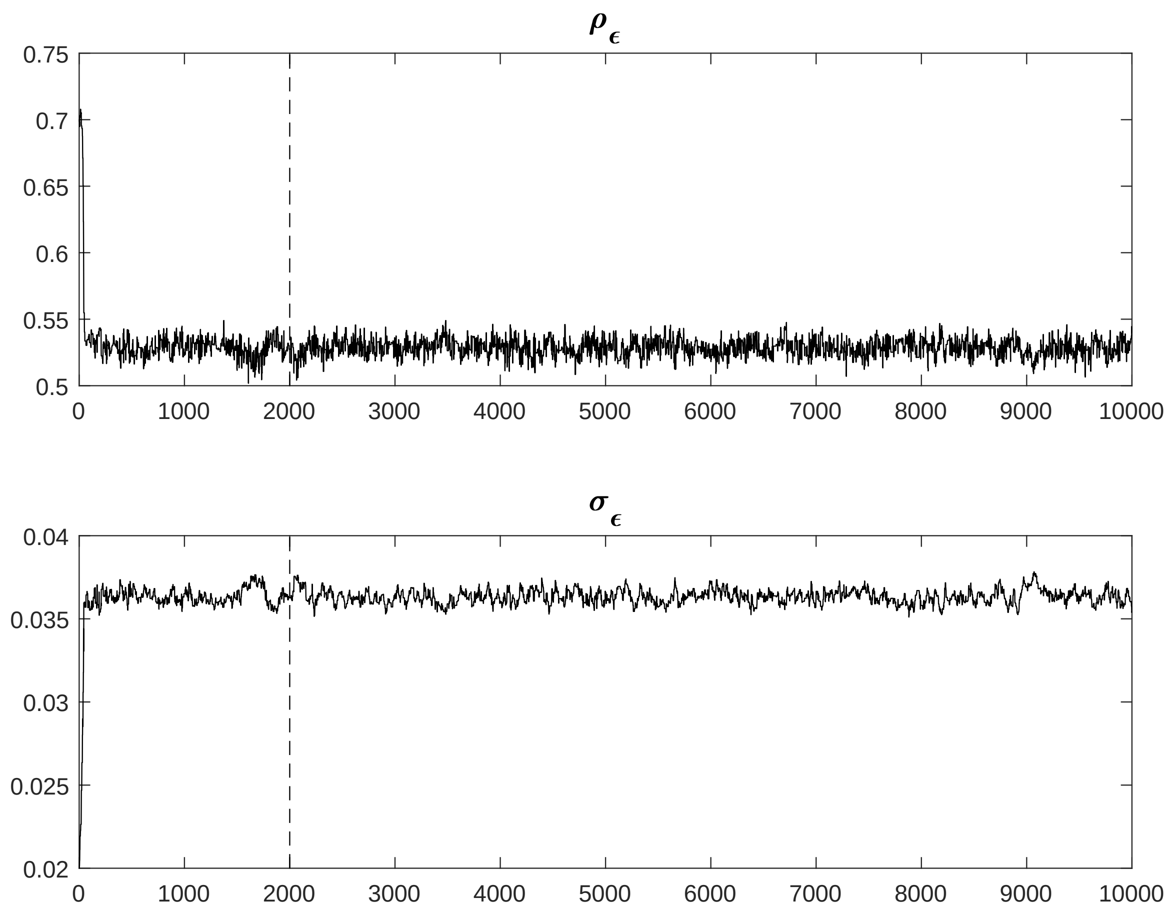

Figure 6: Trace plot of the MCMC chain. Vertical dashed line indicates the end of the burn-in sample.

Heterogeneous firm model: MCMC autocorrelations

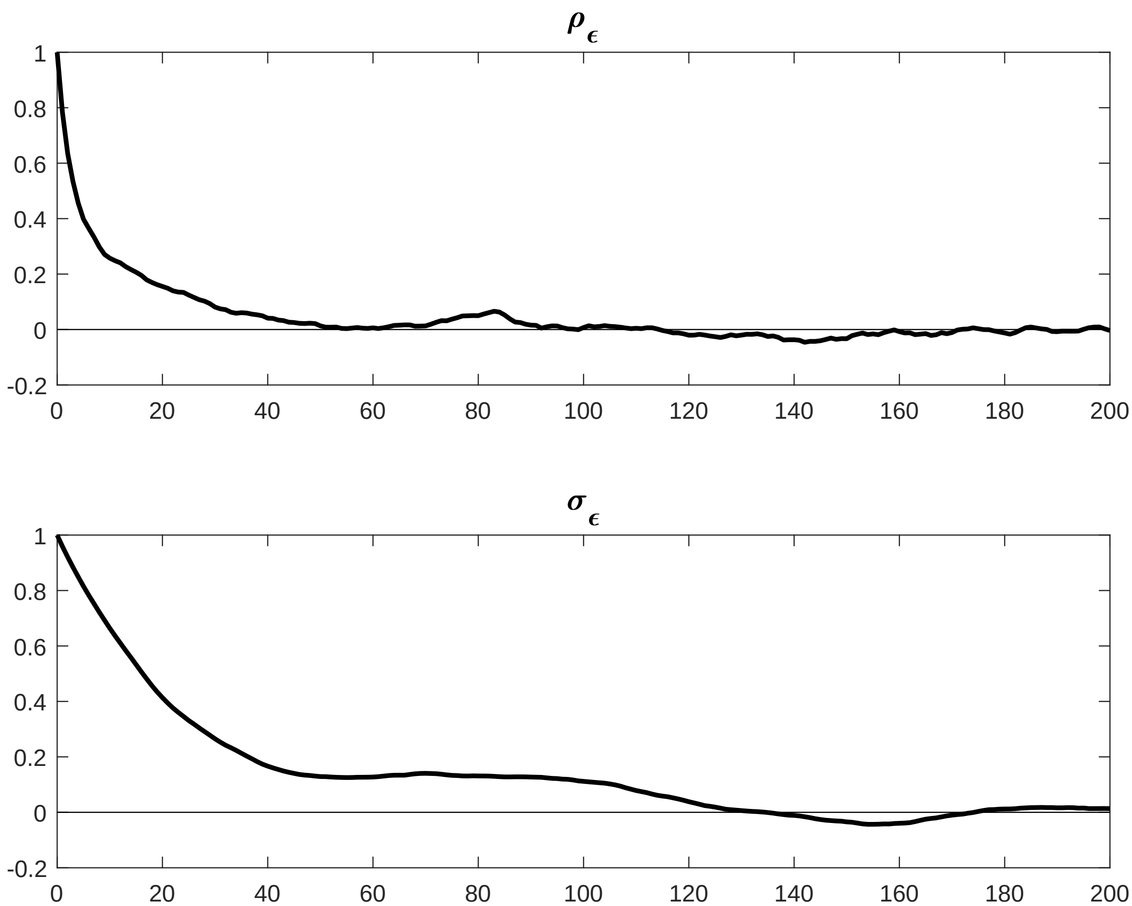

Figure 7: Autocorrelation function of MCMC draws, discarding burn-in sample.

References

Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf (2017): "When Inequality Matters for Macro and Macro Matters for Inequality," in *NBER Macroeconomics Annual, Volume 32*, ed. by J. A. Parker, and M. Woodford, pp. 1–75. University of Chicago Press.

Andrieu, C., A. Doucet, and R. Holenstein (2010): "Particle Markov chain Monte Carlo methods," *Journal of the Royal Statistical Society: Series B*, 72(3), 269–342.

Arellano, M., and S. Bonhomme (2017): "Nonlinear Panel Data Methods for Dynamic Heterogeneous Agent Models," *Annual Review of Economics*, 9(1), 471–496.

Atchadé, Y. F., and J. S. Rosenthal (2005): "On adaptive Markov chain Monte Carlo algorithms," *Bernoulli*, 11(5), 815–828.

Chang, M., X. Chen, and F. Schorfheide (2018): "Heterogeneity and Aggregate Fluctuations," Manuscript, University of Pennsylvania.

Chang, Y., J. F. Gomes, and F. Schorfheide (2002): "Learning-by-Doing as a Propagation Mechanism," *American Economic Review*, 92(5), 1498–1520.

Del Negro, M., S. Eusepi, M. Giannoni, A. Sbordone, A. Tambalotti, M. Cocci, R. Hasegawa, and M. H. Linder (2013): "The FRBNY DSGE Model," Federal Reserve Bank of New York Staff Reports 647.

Durbin, J., and S. J. Koopman (2002): "A simple and efficient simulation smoother for state space time series analysis," *Biometrika*, 89(3), 603–616.

Flury, T., and N. Shephard (2011): "Bayesian Inference Based Only on Simulated Likelihood: Particle Filter Analysis of Dynamic Economic Models," *Econometric Theory*, 27(05), 933–956.

Hahn, J., G. Kuersteiner, and M. Mazzocco (2018): "Central Limit Theory for Combined Cross-Section and Time Series," Manuscript, arXiv:1610.01697.

Herbst, E., and F. Schorfheide (2016): *Bayesian Estimation of DSGE Models*, Econometric and Tinbergen Institutes Lectures. Princeton University Press.

Kaplan, G., and G. L. Violante (2018): "Microeconomic Heterogeneity and Macroeconomic Shocks," *Journal of Economic Perspectives*, 32(3), 167–194.

Khan, A., and J. K. Thomas (2008): "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics," *Econometrica*, 76(2), 395–436.

Krueger, D., K. Mitman, and F. Perri (2016): "Macroeconomics and Household Heterogeneity," in *Handbook of Macroeconomics, Volume 2*, ed. by J. B. Taylor, and H. Uhlig, chap. 11, pp. 843–921. Elsevier.

Krusell, P., and A. A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106(5), 867–896.

Mongey, S., and J. Williams (2017): "Firm dispersion and business cycles: Estimating aggregate shocks using panel data," Manuscript, New York University.

Piketty, T., E. Saez, and G. Zucman (2018): "Distributional National Accounts: Methods and Estimates for the United States," *Quarterly Journal of Economics*, 133(2), 553–609.

Reiter, M. (2009): "Solving Heterogeneous-Agent Models by Projection and Perturbation," *Journal of Economic Dynamics and Control*, 33(3), 649–665.

Winberry, T. (2016): "User Guide for "A Toolbox for Solving and Estimating Heterogeneous Agent Macro Models"," Manuscript, University of Chicago Booth School of Business.

——— (2018): "A Method for Solving and Estimating Heterogeneous Agent Macro Models," *Quantitative Economics*, Forthcoming.