Solving Heterogeneous Agent Models with Dispersed Information in Sequence Space

Jonas Overhage*

Preliminary draft, September 2026

Abstract

This paper studies how to solve rational expectations models in which households do not directly observe the aggregate objects that matter for their decisions. I consider economies with rational expectations, heterogeneous agents, endogenous equilibrium aggregates, and signals that combine aggregate and idiosyncratic components. The main contribution is a method that embeds signal extraction into a sequence space solution approach. The key factor is aggregate linearity, an assumption that is standard in sequence space methods. I show that this assumption implies a mapping that can be used to apply a Kalman filter for signal extraction, making it possible to solve for transition paths. The method is illustrated in a small open economy model with incomplete markets and endogenous labor supply, where households infer the aggregate state from their own observed wages.

*Institute for International Economic Studies, Stockholm University.

1 Introduction

A common premise in rational expectations (RE) models is that agents observe the aggregate state directly. One reason for this full information assumption is tractability: Bayesian updating about endogenous aggregates, conditional on some information set, generally does not have a closed-form expression, which makes solving such models technically difficult. Instead, models incorporating heterogeneity in agents' information sets usually deviate from rational expectations.

This paper develops a method for solving this class of problems under rational expectations. The key setting is one in which households do not observe endogenous aggregate objects directly, but instead observe individual signals that combine these objects with idiosyncratic noise. Crucially, these signals may also influence the household problem directly, meaning they may correlate with other household state or choice variables. Solving RE models then requires a method that jointly handles signal extraction and forward-looking household choice while maintaining general equilibrium consistency.

The proposed approach is based on a sequence space representation of local equilibrium dynamics. Around a steady state, I characterize the response of endogenous aggregate objects to an aggregate innovation by an impulse-response function (IRF). Given this representation, individual observations can be interpreted as noisy signals about the underlying innovation. Households update expectations using individual signals and form expectations over future aggregate paths implied by the underlying IRF transition path. In essence, households fully understand how the IRF to a specific initial innovation would look, but never observe that innovation itself. Equilibrium is obtained as a fixed point over the aggregate response objects. As households use this exact mapping in their signal extraction, expectations updating is fully consistent with aggregate dynamics.

The paper is structured as follows. Section 2 reviews the related literature. Section 3 presents a general environment with heterogeneous agents, endogenous aggregates, and noisy individual signals. Section 4 introduces the IRF-based solution method in this general notation. Section 5 specializes the framework to a small open economy with incomplete markets and endogenous labor supply. Section 6 concludes.

2 Related literature

This paper relates first to the literature on linear solution methods for heterogeneous-agent models. Reiter (2009) as well as Bayer and Luetticke (2020) develop perturbation approaches for heterogeneous-agent economies with aggregate risk. Boppart et al. (2018) and Auclert et al. (2021) show how aggregate dynamics can be characterized by sequence space responses to MIT shocks and sequence space Jacobians respectively. I build on this sequence space approach literature, but apply it to an environment in which households do not directly observe the aggregate objects relevant for their decisions.

Second, the paper relates to the literature on dispersed information and signal extraction in macroeconomics. Lucas (1972) provides the classic formulation in which agents infer aggregate conditions from individual observations, and Sims (2003) studies imperfect information through rational inattention. More recent work such as Angeletos and La'O (2010) analyzes macroeconomic fluctuations under noisy information. Relative to this literature, the key feature here is that households infer endogenous equilibrium outcomes rather than exogenous aggregate states.

Third, the paper relates to work on sticky expectations. Mankiw and Reis (2002) study sluggish adjustment through infrequent updating. Auclert et al. (2023) consider a similar framework, and solve a model with sticky expectations using sequence space methods. Relative to this, the contribution of this paper is to allow for expectations updating that is (i) not orthogonal to other state variables and (ii) building on a set of heterogeneous signals.

The most closely related work is by Auclert et al. (2026), who attempt to solve the exact same computational challenge from a different starting point. As both this paper as well as theirs are works in progress, a comparison between the two approaches is for now left to future work.

3 General Environment

This section lays out the general class of heterogeneous-agent economies with dispersed information this paper considers. The key restriction maintained throughout is that aggregate uncertainty is one-dimensional: On a transition path, all endogenous aggregate objects are driven by a single aggregate innovation. While models may of course have multiple shocks, this paper builds on an aggregate linearity assumption similar to Boppart et al. (2018). Essentially, it is assumed that responses to multiple shocks are equal to the superimposition of corresponding single shocks. From now on, I denote aggregate variables with capital letters (e.g. X) and vector-valued objects in bold (e.g. $\mathbf{X}$).

3.1 Households, aggregates, and equilibrium objects

Time is discrete and indexed by t. The economy is populated by a unit mass of households indexed by $i \in [0, 1]$. Household i enters period t with individual state vector $\mathbf{a}_{it}$ and chooses a control vector $\mathbf{u}_{it}$. The aggregate economy is summarized by a vector of endogenous equilibrium objects,

$$\mathbf{X}_t = \begin{bmatrix} X_{1t} \\ \vdots \\ X_{Mt} \end{bmatrix}, \tag{1}$$

where the components of $\mathbf{X}_t$ may represent prices, wages, quantities, or other aggregate objects relevant for household decisions. The key feature that makes changes to existing solution methods necessary is that $\mathbf{X}_t$ is not exogenous. Instead, it is determined in equilibrium through aggregation over household decisions:

$$\mathbf{X}_t = \mathcal{G}_t\big(\{\mathbf{u}_{it}, \mathbf{a}_{it}\}_{i\in[0,1]}, \mu^t\big), \tag{2}$$

where $\mathcal{G}_t(\cdot)$ is the equilibrium mapping implied by technology, market clearing, and other aggregate constraints, and μ^t denotes the history of aggregate innovations up to date t. These innovations capture exogenous shocks to the economy, which are assumed to be scalar. In practice, these may correspond to productivity or monetary policy shocks.

Households solve dynamic problems in which current and future values of $\mathbf{X}_t$ have implications for optimal behavior. As a result, expectations about aggregate conditions are part of the household state. The central assumption for the method developed below is that endogenous sequences respond linearly to aggregate shocks, as is standard in sequence space methods (see Boppart et al., 2018, Auclert et al., 2021). I consider for simplicity here a case with a scalar aggregate innovation μ, but under this assumption the setup extends to cases with multiple shocks as well.

3.2 Information structure

Households do not observe $\mathbf{X}_t$ directly. Instead, household i at each time t observes a vector of individual signals,

$$\mathbf{x}_{it} = \begin{bmatrix} x_{it}^{(1)} \\ \vdots \\ x_{it}^{(K)} \end{bmatrix}, \tag{3}$$

where each component may load on one or more aggregate objects and also contains idiosyncratic noise. In levels, let the signal system be written as

$$\mathbf{x}_{it} = \mathcal{H}_t(\mathbf{X}_t) + \boldsymbol{\nu}_{it}, \tag{4}$$

where $\boldsymbol{\nu}_{it}$ is a $K \times 1$ vector of iid disturbances, and $\mathcal{H}$ is a function relating aggregate states to individual signals. Linearized around a steady state, the signal system can be written as:

$$\tilde{\mathbf{x}}_{it} = \mathbf{H}_t \tilde{\mathbf{X}}_t + \boldsymbol{\nu}_{it}. \tag{5}$$

Here, $\mathbf{H}_t$ is a $K \times M$ loading matrix and tildes denote percentage deviations from steady state approximated following $\log x_t - \log x_{\text{ss}}$. This formulation allows the observed signals to enter the household budget constraint or utility directly, but does not require them to do so. The setup thus admits signals correlated with other household states, such as information carried by wages or prices, as well as signals that matter exclusively through the information they convey about aggregate conditions.

The disturbance vector $\boldsymbol{\nu}_{it}$ is assumed to be mean-zero conditional on aggregate information and orthogonal to aggregate innovations:

$$E[\boldsymbol{\nu}_{it} \mid \mu^t] = \mathbf{0}. \tag{6}$$

Additionally, assume that $\boldsymbol{\nu}_{it}$ is conditionally independent across households and has covariance matrix

$$\boldsymbol{\Sigma}_{\nu,t} = \mathrm{Var}(\boldsymbol{\nu}_{it} \mid \mu^t), \tag{7}$$

known to the household. Correlation across signals within a household is permitted through the off-diagonal elements of $\boldsymbol{\Sigma}_{\nu,t}$.

The essential additional assumption is that all aggregate variation relevant for the signals is driven by the same aggregate shock. Equivalently, on a transition path, any aggregate object that enters the signal system must admit a representation in terms of that same one-dimensional aggregate disturbance. Under this restriction, the household faces a multi-dimensional signal-extraction problem about a scalar state.

3.3 General household problem and equilibrium

Let Ω_{it} denote household i's information set at date t, consisting of the history of individual states and realized signals. Households have preferences over stochastic paths of controls and outcomes of the form

$$E_0 \sum_{t=0}^{\infty} \beta^t \, u(\mathbf{s}_{it}, \mathbf{u}_{it}), \tag{8}$$

where $\mathbf{s}_{it}$ denotes the vector of individual payoff-relevant states, $\mathbf{u}_{it}$ the vector of controls, and $\mathbf{X}_t$ the vector of endogenous aggregate objects. Individual states evolve according to

$$\mathbf{a}_{i,t+1} = \mathcal{A}_t(\mathbf{a}_{it}, \mathbf{u}_{it}, \mathbf{X}_t, \boldsymbol{\varepsilon}_{i,t+1}), \tag{9}$$

where $\boldsymbol{\varepsilon}_{i,t+1}$ collects idiosyncratic shocks. Household choices are further restricted by a generic feasibility condition,

$$\mathcal{B}_t(\mathbf{a}_{it}, \mathbf{a}_{i,t+1}, \mathbf{u}_{it}, \mathbf{X}_t, \mathbf{x}_{it}) \geq \mathbf{0}, \tag{10}$$

which may represent a budget constraint, borrowing constraint or another period restriction. The observed signal vector $\mathbf{x}_{it}$ may enter this constraint directly. Likewise, signals may enter current utility directly, or matter only through the information they convey about aggregate conditions.

Because households do not observe $\mathbf{X}_t$ directly, optimal choices depend not only on current individual states, but also on expectations over current and future aggregate objects. The household problem can therefore be written recursively as

$$V(\mathbf{a}_{it}, \Omega_{it}) = \max_{\mathbf{u}_{it}} \left\{ u(\mathbf{s}_{it}, \mathbf{u}_{it}) + \beta E\left[V(\mathbf{a}_{i,t+1}, \Omega_{i,t+1}) \mid \Omega_{it} \right] \right\}, \tag{11}$$

subject to the feasibility condition and the household's expectations over future aggregate outcomes. Here the information set Ω_{it} is the full history of realized signals and the expectations induced by them. Note that aggregate $\mathbf{X}_t$ does affect household utility, but is not included as a state due to being unobserved.

A recursive equilibrium consists of household decision rules, a law of motion for expectations, and an aggregate process $\{\mathbf{X}_t\}_{t=0}^{T}$ such that: (i) household decisions are optimal given expectations, (ii) expectations are consistent with the signal structure, and (iii) aggregate outcomes are generated by the equilibrium mapping following from household behavior. The computational challenge is that expectations depend on the information extracted from the vector of observed signals $\mathbf{x}_{it}$, while the aggregate objects being inferred are themselves pinned down by household decisions. In essence, how households should optimally form expectations relates directly to equilibrium mappings, and vice-versa.

4 IRF-based solution method

This section presents the proposed solution method in general notation. The method exploits the IRF representation of aggregate dynamics and uses it to transform the filtering problem into inference about a scalar aggregate innovation, even when both the vector of endogenous aggregate objects and the vector of observed signals are multidimensional. This section relies on aggregate linearity: This assumption allows the following characterization of sequence dynamics around a steady state.

4.1 Linear representation of aggregate dynamics

Suppose the economy is linearized around a full information steady state. Consider the response to a one-time aggregate innovation μ_0 at date 0, which has variance σ_μ^2. Under aggregate linearity, the resulting path of the vector of endogenous aggregate objects can be represented as

$$\tilde{\mathbf{X}}_t = \mathbf{\Gamma}_t \mu_0, \tag{12}$$

where $\mathbf{\Gamma}_t$ is an $M \times 1$ vector of impulse-response coefficients,

$$\mathbf{\Gamma}_t = \begin{bmatrix} \Gamma_t^1 \\ \vdots \\ \Gamma_t^M \end{bmatrix}. \tag{13}$$

Equivalently, each aggregate object satisfies

$$\tilde{X}_t^m = \Gamma_t^m \mu_0, \qquad m = 1, \ldots, M. \tag{14}$$

An IRF of length T is then fully characterized by a sequence of vectors $\{\boldsymbol{\Gamma}_t\}_{t=0}^T$, capturing all movements in aggregate variables in response to a shock. Equation 12 is central to the numerical method. It summarizes how a scalar aggregate disturbance maps into the entire future path of the endogenous equilibrium objects that factor into the household problem. The crucial restriction is that, for an IRF, the aggregate shock is one-dimensional. This does not preclude models with multiple aggregate shocks, but it does require the aggregate linearity assumption. Following Boppart et al. (2018), responses to multiple innovations could simply be superimposed from multiple IRFs.

4.2 Expectations updating from individual signals

Combining the linear representation of aggregate dynamics with the signal structure yields

$$\tilde{\mathbf{x}}_{it} = \mathbf{H}_t \boldsymbol{\Gamma}_t \mu_0 + \boldsymbol{\nu}_{it}. \tag{15}$$

Define the $K \times 1$ vector of composite signal loadings

$$\boldsymbol{\Lambda}_t \equiv \mathbf{H}_t \boldsymbol{\Gamma}_t. \tag{16}$$

Following this, the signal system becomes:

$$\tilde{\mathbf{x}}_{it} = \boldsymbol{\Lambda}_t \mu_0 + \boldsymbol{\nu}_{it}. \tag{17}$$

Thus, each household observes multiple noisy signals about the same aggregate innovation μ_0. The role of the loading vector $\mathbf{\Lambda}_t$ is to summarize how the one-dimensional aggregate disturbance is transmitted into the full vector of observed signals once equilibrium responses of the aggregate objects have been taken into account.

Let $m_{i,t|t-1}$ denote household i's prior mean for μ_0 before observing $\mathbf{x}_{it}$, and let $P_{t|t-1}$ denote the corresponding prior variance. At time 0, priors $m_{i,0|\text{pre}} = 0, P_{i,0|\text{pre}} = \sigma_\mu^2$ are used, reflecting the true, full information steady state hit with a one-off shock with standard deviation σ_μ. As the noise vector $\boldsymbol{\nu}_{it}$ is Gaussian with covariance matrix $\mathbf{\Sigma}_{\nu,t}$, posterior expectations are given by a standard multivariate Kalman filter. In particular, the posterior mean is

$$m_{i,t|t} = m_{i,t|t-1} + \mathbf{K}_t \left(\tilde{\mathbf{x}}_{it} - \mathbf{\Lambda}_t m_{i,t|t-1} \right), \tag{18}$$

where the Kalman gain is

$$\mathbf{K}_t = P_{t|t-1} \mathbf{\Lambda}_t' \left(\mathbf{\Lambda}_t P_{t|t-1} \mathbf{\Lambda}_t' + \mathbf{\Sigma}_{\nu,t} \right)^{-1}. \tag{19}$$

The posterior variance is

$$P_{t|t} = P_{t|t-1} - P_{t|t-1} \mathbf{\Lambda}_t' \left(\mathbf{\Lambda}_t P_{t|t-1} \mathbf{\Lambda}_t' + \mathbf{\Sigma}_{\nu,t} \right)^{-1} \mathbf{\Lambda}_t P_{t|t-1}. \tag{20}$$

Here, the household's inference problem remains scalar: all signals help identify the same aggregate shock at time 0. Given posterior expectations about μ_0, the household forms expectations over future aggregate conditions through the same impulse responses:

$$E_{it}[\tilde{\mathbf{X}}_{t+s}] = \mathbf{\Gamma}_{t+s} m_{i,t|t}, \qquad s \geq 0. \tag{21}$$

Or, componentwise,

$$E_{it}[\tilde{X}^m_{t+s}] = \Gamma^m_{t+s} m_{i,t|t}, \qquad m = 1, \dots, M. \tag{22}$$

These expectations enter current decision rules and thereby affect equilibrium outcomes.

4.3 Equilibrium fixed point

The remaining step is to impose consistency between the conjectured impulse-response sequence and the one implied by household behavior. The computational procedure is as follows.

First, guess a sequence $\{\boldsymbol{\Gamma}_t\}_{t=0}^T$ for the response of the endogenous aggregate vector to the innovation μ. Second, given this sequence and the signal-loading vectors $\{\mathbf{H}_t\}_{t=0}^T$, construct the implied composite loadings $\boldsymbol{\Lambda}_t = \mathbf{H}_t\boldsymbol{\Gamma}_t$. Third, backwards iterate to solve the households' dynamic problem, updating expectations about μ_0 according to equation 18. Fourth, iterate the distribution forward and aggregate the resulting transitions across households to recover the implied path of the aggregate vector $\tilde{\mathbf{X}}_t$. Finally, update the conjectured sequence and iterate until convergence.

At a fixed point, the sequence $\{\boldsymbol{\Gamma}_t\}_{t=0}^T$ used by households to interpret their signals coincides with the sequence following from aggregating their choice variables. The method therefore preserves rational expectations under the aggregate linearity assumption. Because expectations updating follows directly from the true IRF coefficients, no separate loop is needed to find the "correct" expectations.

The main limitation is also clear from this formulation. The method requires that all aggregate movements along the IRF relevant for households must originate linearly from the same scalar disturbance μ_0. If this relation is nonlinear, or interacts in a nonlinear way with a secondary exogenous shock, signal extraction this way would no longer be valid. Note that all tests of the aggregate linearity assumption from described by Boppart et al. (2018) can be applied to this method as well.

5 Application: Small Open Economy

This section makes the general framework concrete: The numerical method is applied to a small open economy with incomplete markets, endogenous labor supply, and noisy information about aggregate wages. Relative to Sections 3 and 4, the application is intentionally kept simple. In particular, it corresponds to the scalar special case

$$M = 1, \qquad K = 1. \tag{23}$$

5.1 Model environment

The economy is populated by a unit mass of households who trade a one-period foreign bond and supply labor. The gross world interest rate $1 + r$ is exogenous and satisfies $1 + r < 1/\beta$. Aggregate productivity follows

$$\log z_{t+1} = \rho \log z_t + \mu_{t+1}, \quad \mu_{t+1} \sim N(0, \sigma_\mu^2) \tag{24}$$

where μ_{t+1} is an aggregate productivity innovation. Households have preferences

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_{it}^{1-\sigma}}{1-\sigma} - \frac{\chi}{1+\phi} h_{it}^{1+\phi} \right]. \tag{25}$$

Household i enters period t with bond holdings b_{it} and idiosyncratic labor efficiency η_{it}. This idiosyncratic efficiency follows

$$\log \eta_{it} \sim \mathcal{N}\left(-\frac{\sigma_\eta^2}{2}, \sigma_\eta^2\right). \tag{26}$$

Effective labor supplied by household i is

$$n_{it} = \eta_{it} h_{it}, \tag{27}$$

and aggregate effective labor is

$$N_t = \int_0^1 \eta_{it} h_{it}\, di. \tag{28}$$

Production is carried out by a representative competitive firm with technology

$$Y_t = z_t N_t^{1-\alpha}. \tag{29}$$

Profit maximization implies that the wage per efficiency unit is

$$W_t = (1-\alpha) z_t N_t^{-\alpha}. \tag{30}$$

The realized wage of household i is therefore

$$w_{it} = W_t \eta_{it}. \tag{31}$$

Profits $\Pi_t = \alpha z_t N_t^{1-\alpha}$ are rebated to foreign firm owners for the sake of simplicity. The household budget constraint is

$$c_{it} + b_{i,t+1} = (1 + r)b_{it} + w_{it}h_{it}, \tag{32}$$

subject to a no-borrowing constraint

$$b_{i,t+1} \geq 0. \tag{33}$$

Under full information, households observe both the aggregate wage W_t and their idiosyncratic efficiency η_{it} before choosing labor supply and saving. The first-order condition for labor supply is

$$\chi h_{it}^{\phi} = c_{it}^{-\sigma} W_t \eta_{it}, \tag{34}$$

and the Euler condition for bond holdings is

$$c_{it}^{-\sigma} = \beta(1 + r)E_t[c_{i,t+1}^{-\sigma}] + \lambda_{it}, \tag{35}$$

where $\lambda_{it} \geq 0$ is the multiplier on the borrowing constraint, together with complementary slackness.

5.2 Benchmark results under full information

The model is calibrated externally to standard values, as the results discussed here are qualitative in nature. Parameters can be found in table 1. Note that the interest rate in this model must, by necessity, be calibrated below $1/\beta - 1$. Due to the presence of a borrowing constraint, savings provide an insurance value in addition to the mere rolling over of consumption to the next period. With $r = 1/\beta - 1$, there would thus be runaway savings, with asset positions increasing over time.

Table 1: Model parameters

Parameter	Value	Description
σ	2.0	Utility curvature (CRRA)
β	0.985	Discount factor
r	$(1/\beta - 1)/2$	Interest rate
α	1/3	Capital share
χ	0.1	Weight on leisure in utility
ϕ	1.5	inverse Frisch elasticity
ρ	0.95	Persistence of productivity shocks
σ_μ	0.007	Std. dev. of productivity shocks
σ_η	0.15	Std. dev. of idiosyncratic efficiency shocks

Under full information, the model can be solved using standard techniques following Boppart et al. (2018). Full information IRFs for this model, in response to a one standard deviation μ shock are depicted in figure 1. As is standard in such models, these IRFs are fully front-loaded, with maximum impact at the time of the shock.

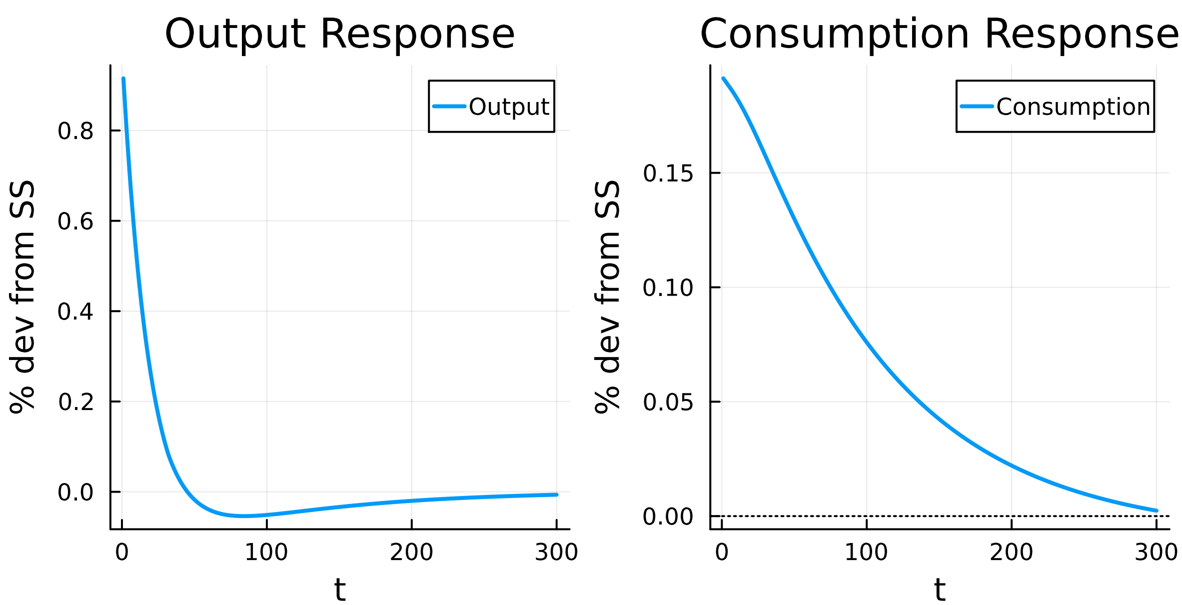

Figure 1: Full information impulse response functions

5.3 Information Friction, Solution Method

Let there now be an information friction in the model: Households no longer observe W_t and η_{it} separately. Instead, after aggregate and idiosyncratic shocks are realized, household i observes only its realized wage w_{it} and then chooses labor supply and saving. The realized wage is therefore both correlated with other states and choices, because it enters the household budget constraint directly, as well as informative, because it contains information about aggregate conditions. In the notation of Sections 3 and 4, this application is the special case

$$\mathbf{X}_t = W_t, \qquad \mathbf{x}_{it} = w_{it}. \tag{36}$$

Since

$$w_{it} = W_t \eta_{it}, \tag{37}$$

a log-linear approximation around the steady state gives

$$\tilde{w}_{it} = \tilde{W}_t + \tilde{\eta}_{it}. \tag{38}$$

Thus the observed wage admits the additive signal structure required by the general method. In this application, the single observed signal is both part of the household's budget constraint and a noisy measurement of the aggregate object to be inferred. This friction creates the key general equilibrium problem. The aggregate wage satisfies

$$\tilde{W}_t = \tilde{z}_t - \alpha \tilde{N}_t, \tag{39}$$

so the aggregate object households seek to learn about is endogenous in equilibrium. Aggregate labor supply depends on household decisions, but household decisions depend on beliefs about current and future aggregate wages. The signal-extraction problem and the equilibrium problem must therefore be solved jointly.

To apply the solution method, consider a one-off aggregate productivity innovation μ_0 at date 0, followed by no further aggregate innovations. Since productivity follows an AR(1) in logs, the resulting path of aggregate productivity is

$$\tilde{z}_t = \rho^t \mu_0. \tag{40}$$

Under linearization, aggregate labor and the aggregate wage respond according to

$$\tilde{N}_t = \Gamma_t^N \mu_0, \qquad \tilde{W}_t = \Gamma_t^W \mu_0, \tag{41}$$

where $\{\Gamma_t^N\}$ and $\{\Gamma_t^W\}$ denote the corresponding IRF coefficients. Combining these expressions with equation (39) implies

$$\Gamma_t^W = \rho^t - \alpha \Gamma_t^N. \tag{42}$$

The SOE therefore fits the general IRF representation with a scalar disturbance μ_0 and a scalar endogenous aggregate object $X_t = W_t$. The observed wage is then a noisy signal about that same disturbance:

$$\tilde{w}_{it} = \Gamma_t^W \mu_0 + \tilde{\eta}_{it}. \tag{43}$$

Operationally, households learn about the underlying aggregate innovation μ_0, and use the equilibrium IRF coefficients $\{\Gamma_t^W\}$ to map posterior beliefs about μ_0 into expectations over future aggregate wages. In this sense, households do not directly filter the aggregate wage itself. Rather, they infer the initial shock that generates the sequence $\{W_t\}_{t=0}^T$.

Given a conjectured wage IRF $\{\Gamma_t^W\}$, households update beliefs about μ_0 from observed wages, form expectations over future wage paths, solve their labor-savings problem, and thereby generate an implied aggregate labor response. Equilibrium is obtained by iterating on the wage IRF until the conjectured and implied sequences coincide.

5.4 Results, noisy information

5.4.1 Impulse responses

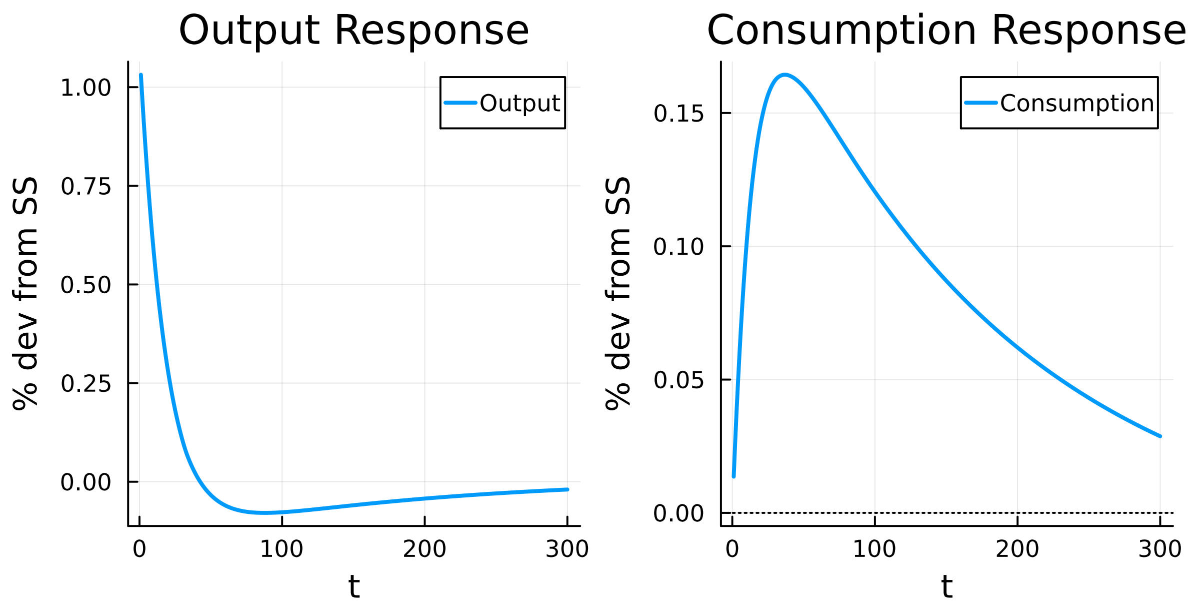

Figure 2: Noisy information impulse response functions

Having solved the model, noisy information IRFs can be studied. The responses to a one standard deviation productivity shock are depicted in figure 2. The largest difference here is in consumption: Due to the signal extraction, this now exhibits a hump-shaped form, with the shock only reaching its full impact on consumption after several periods.

5.4.2 Simulation and accuracy

The solution in this paper gained considerable simplicity and computational speed by considering only linear dynamics around a steady state. To evaluate the cost of these assumptions, I compare the results to the global solution method from Overhage (2026). Following the approximate aggregation method established in that paper, the model is solved for a perceived law of motion

$$\log W_{t+1} = a_0 + a_1 \log W_t + a_2 \mu_t, \quad \mu_t \sim \mathcal{N}(0, \sigma_\mu^2), \tag{44}$$

which implies signal extraction following from Kalman updates given by:

$$\hat{W}_{it|t} = (1 - k_t)\hat{W}_{it|t-1} + k_t \hat{w}_{it}^s, \quad k_t = \frac{P_{t|t-1}}{P_{t|t-1} + \sigma_\eta^2}. \tag{45}$$

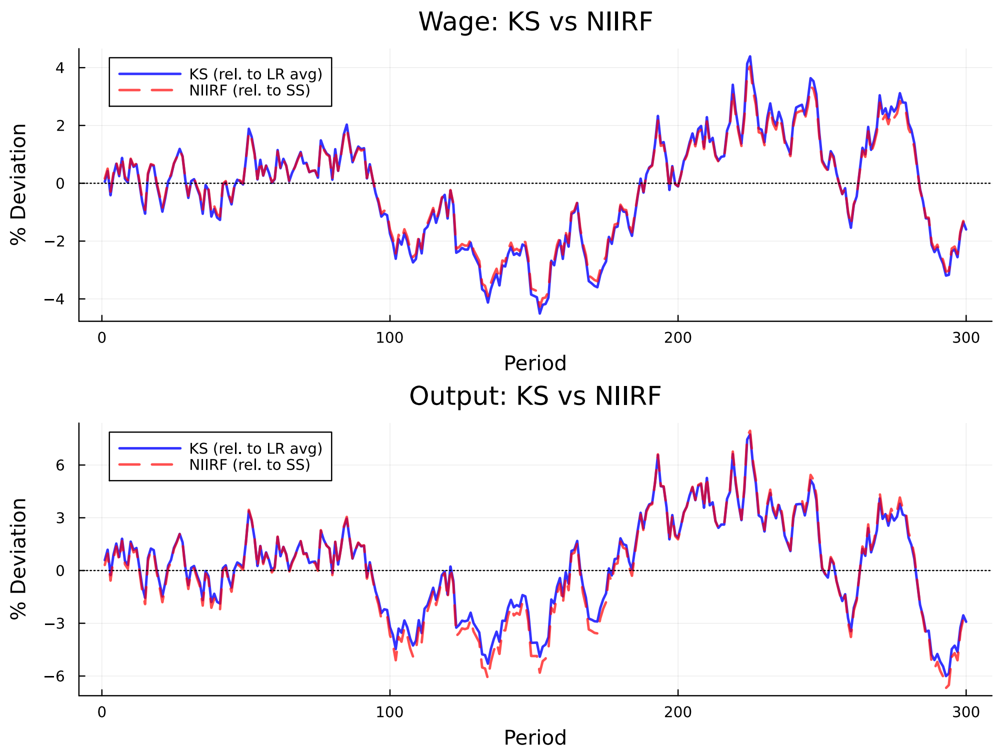

Figure 3: Simulation results

Simulating the model for the same series of shocks, using both the nonlinear method as well as the linearization around the steady state presented in this paper allows a comparison of the results. Figure 3 shows the results of this simulation. Note that the largest disagreements in absolute terms occur after especially large shocks. This is unsurprising, as the IRF-based approach relies on aggregate linearity, accuracy should be expected to decrease as we move further from the steady state.

Despite this, broadly, the results look quite similar, which can also be verified by estimating a law of motion following equation 44 on the simulated series from the IRF-based method from this paper. These coefficients are included in Table 2. As both solutions rely on approximations, in one case this is local aggregate linearity as common in sequence space methods (Auclert et al. (2021), Boppart et al. (2018)), whereas in the other it is a behavioral approximation (Krusell and Smith (1998)), neither solution solves the full rational expectations equilibrium.

Table 2: Krusell-Smith versus IRF-based coefficients

	a_0	a_1	a_2	R^2
KS	-0.023751	0.965398	0.821026	0.999915
IRF	-0.024999	0.964058	0.771015	0.999936

It is also possible to compare how closely both simulations align in terms of various endogenous equilibrium objects. These results are contained in Table 3. Results overall are very close, with the correlation between $\{W_t\}_{t=0}^T$ from both methods greater than 0.999 and correlation for $\{Y_t\}_{t=0}^T$ greater than 0.996.

Table 3: Simulation Statistics: KS vs. NIIRF

Variable	KS	NIIRF	Difference
Mean W	0.502135	0.498783	-0.003352
Std W	0.010830	0.010064	-0.000767
Mean N	2.338953	2.384324	0.045371
Std N	0.043492	0.044560	0.001067
Mean Y	1.761719	1.783896	0.022176

6 Conclusion

This paper develops a numerical approach for solving rational expectations models with heterogeneous agents and dispersed information when the aggregate objects of interest are endogenous equilibrium outcomes. The method uses a local sequence space representation to convert the informational problem into inference about a low-dimensional aggregate innovation. Households update expectations from individual signals, while equilibrium is recovered as a fixed point over aggregate impulse responses.

A small open economy is considered here as an example. More broadly, the method is designed to be portable across a wider class of heterogeneous-agent models in which expectations, individual signals, and endogenous aggregates are tightly intertwined. This portability is especially useful in environments where a fully global treatment of filtering and equilibrium would be computationally prohibitive: Adding an information friction and solving this way is fast: Relative to full information, only one additional state variable, namely the expectation over an initial shock μ_0, is required to solve an IRF – irrespective of the complexity of the model under consideration.

References

Angeletos, G.-M., & La'O, J. (2010). Noisy business cycles. In D. Acemoglu, K. Rogoff, & M. Woodford (Eds.), *Nber macroeconomics annual 2009, volume 24* (pp. 319–378). University of Chicago Press. https://doi.org/10.1086/648301

Auclert, A., Bardóczy, B., Rognlie, M., & Straub, L. (2021). Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. *Econometrica*, *89*(5), 2375–2408. https://doi.org/10.3982/ECTA17434

Auclert, A., Monnery, H., Rognlie, M., & Straub, L. (2026, July). *Spend What You (L)earn* [Presentation slides, NBER Summer Institute, Monetary Economics]. https://conference.nber.org/conf_papers/f246767/f246767.slides.pdf

Auclert, A., Rognlie, M., & Straub, L. (2023). *Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model* [R&R, American Economic Review].

Bayer, C., & Luetticke, R. (2020). Solving discrete time heterogeneous agent models with aggregate risk and many idiosyncratic states by perturbation. *Quantitative Economics*, *11*(4), 1253–1288. https://doi.org/10.3982/QE1243

Boppart, T., Krusell, P., & Mitman, K. (2018). Exploiting MIT shocks in heterogeneous-agent economies: The impulse response as a numerical derivative. *Journal of Economic Dynamics and Control*, *89*, 68–92. https://doi.org/10.1016/j.jedc.2018.01.002

Krusell, P., & Smith, A. A., Jr. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, *106*(5), 867–896. https://doi.org/10.1086/250034

Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory*, *4*(2), 103–124. https://doi.org/10.1016/0022-0531(72)90142-1

Mankiw, N. G., & Reis, R. (2002). Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve*. *The Quarterly Journal of Economics*, *117*(4), 1295–1328. https://doi.org/10.1162/003355302320935034

Overhage, J. (2026). Inferring the aggregate: Information frictions and macroeconomic dynamics. *Working Paper*.

Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, *33*(3), 649–665. https://doi.org/10.1016/j.jedc.2008.08.010

Sims, C. (2003). Implications of Rational Inattention. *Journal of Monetary Economics*, *50*, 665–690. https://doi.org/10.1016/S0304-3932(03)00029-1