Efficient Aggregation in Heterogeneous Agents Models with Bounded Rationality*

David Evans
University of Oregon

Giorgi Nikolaishvili
Wake Forest University

March 8, 2025

Abstract

Simulations of heterogeneous agents models with boundedly rational agents and multiple markets in the style of Krusell and Smith (1997, 1998) require solving for market-clearing conditions at each period to determine temporary equilibria. Existing solution procedures can be computationally intensive, especially in settings with a high-dimensional state space and a large number of agents. We suggest a new method for approximating temporary equilibria in heterogeneous agents models that offers a speedup of multiple orders of magnitude compared to a conventional benchmark. We demonstrate the effectiveness of our procedure by applying it to a model with heterogeneous boundedly rational agents featuring a large idiosyncratic state space.

JEL Classifications: C63; C68

Keywords: Heterogeneous agents; equilibrium approximation; policy function aggregation; market clearing

*This work benefited from access to the University of Oregon high performance computing cluster, Talapas. All errors are our own.

1 Introduction

When simulating heterogeneous agents (HA) economies in the style of Huggett (1993) and Aiyagari (1994) augmented with aggregate shocks, it is necessary to solve for the temporary equilibrium (TE) every period. In a TE, agents' decisions and prices are simultaneously pinned down by market clearing conditions – in turn, these outcomes determine aggregate dynamics. Practically, this involves finding agent-level policy rules as a function of individual and aggregate states, along with prices, then solving for equilibrium prices by applying a sequential root solver to a set of market clearing conditions that depend on the aggregate of individual policy rules as a function of aggregate states and prices (Bakota, 2022). As described in Den Haan et al. (2010), the repeated aggregation of individual policy rules during the computation of TE in a simulation may take a long time, even when using the standard projection methods described in Judd (1992, 1996) to approximate individual policy functions. This issue is most pertinent to HA models with multiple markets and boundedly rational agents, such as those with endogenous labor supply (Krusell and Smith, 1998; Chang and Kim, 2006, 2007; Krusell et al., 2010), as well as those with multiple assets (Krusell and Smith, 1997).

We suggest a new approach to approximating TE in HA models that bypasses a redundancy in conventional projection methods. Once the policy rules are approximated using a set of basis functions over individual states, aggregate states, and prices, it is possible to perform a single aggregation over the individual states each period. In doing so, we construct an approximation to the aggregate demand functions, conditional on the current aggregate states, as summarized by a vector of coefficients for the basis functions over prices. We find that our approach significantly outperforms existing aggregation methods in terms of execution speed.

2 Methodology

Suppose that we aim to find the TE of an HA economy featuring bounded rationality. We model this as a set of demand functions $\widetilde{x}(z, P, Y)$, which depend on individual states z (including beliefs) with distribution Ω, prices P, and a finite-dimensional vector of observable aggregates Y.[1] The aggregate state, $Z = (\Omega, Y)$, follows a law of motion $Z' = G(Z, \nu')$, where Z' and ν' denote the realization of the aggregate state and shocks in the next period, respectively. Our objective is to find the set of market-clearing prices P^* given Z, such that

$$\widetilde{X}(P^*, Z) = \overline{X}(P^*, Z)\,, \tag{1}$$

where

$$\widetilde{X}(P, Z) \equiv \int \tilde{x}(z, P, Y)\, d\Omega(z) \tag{2}$$

is aggregate demand and $\overline{X}$ is aggregate supply.

We use the model developed by Evans et al. (2023) with the notation adapted to the above setting to tangibly illustrate the challenge of solving the general class of TE shown in Eq. (1). We then present the conventional approach to solving for the TE of such multi-market HA economies, and introduce our proposed method as an extension of a basis function approach to preapproximating policy rules. Finally, we compare the execution time of our approach with that of the conventional TE aggregation method.

[1]For example, in the Krussell-Smith model, Y contains the first moment of the distribution of assets.

2.1 Example Environment

Let time be discrete. The economy is populated by a continuum of agents such that a given agent is endowed with a unit of labor per period and derives utility from consumption c and leisure l according to the instantaneous utility function $u(c, l)$. Each agent has a unique effective unit of labor for each unit of nominal labor supplied and receives a corresponding wage that can be separated into the following two components: (1) a common aggregate component w, and (2) an idiosyncratic efficiency component ε distributed across the population. We assume $\{\varepsilon\}$ to be a Markov process with time-invariant transition function Π distributed i.i.d. across agents. In each period, an agent can trade one-period claims to capital for net return r, limited by the exogenous borrowing constraint $\underline{a}$. The goods and factor markets are assumed to be competitive.

In period t, an agent holds claims a, experiences idiosyncratic efficiency ε, and faces factor prices $P = (r, w)$. Additionally, an agent has a vector of beliefs $\psi \in \mathbb{R}^n$, which comprises coefficients of the forecasting model used to form expectations of next period's shadow price λ', where the current shadow price is defined as

$$\lambda(a, \varepsilon, \psi) \equiv (1 + r)\, u_c(c(a, \varepsilon, \psi), l(a, \varepsilon, \psi))\,. \tag{3}$$

The set of individual states $z = (a, \varepsilon, \psi)$ has a distribution Ω.

The representative firm rents capital k at the real rental rate $r + \delta$, hires effective labor n at the real wage w, and produces output under perfect competition using CRTS technology $\theta f(k, n)$, where δ is the capital depreciation rate. θ is a stationary process that affects total factor productivity with dynamics given by $\theta' = \nu' \theta^{\rho}$, where $|\rho| < 1$ and ν' is i.i.d. with a log-normal distribution. There are no capital installation costs. Profit-maximizing behavior by the firm implies that factors earn their respective marginal products:

$$w = \theta\, f_n(k,n) \quad \text{and} \quad r + \delta = \theta\, f_k(k,n)\,. \tag{4}$$

The aggregate state is $Z = (\Omega, Y)$, where Y is a vector of observable aggregates, including θ, that all agents observe in the current period and could use to condition their forecasts of λ'. The transition dynamics of the aggregate state is governed by $Z' = G(Z, \nu')$, where ν' is the total factor productivity shock of the next period.

The individual consumption, labor, and claims demands of agents are functions of idiosyncratic states, prices, and observable aggregates. We may express the individual demand function as $\tilde{x}(z,P,Y) = (\tilde{c}(z,P,Y), \tilde{a}(z,P,Y), \tilde{l}(z,P,Y))$, which satisfies the following conditions:

$$u_c(\tilde{c}(z,P,Y), \tilde{l}(z,P,Y)) \geq \beta\, \tilde{\lambda}^e(\tilde{a}(z,P,Y), \varepsilon, \psi, Y) \tag{5}$$

$$\text{and } \tilde{a}(z,P,Y) \geq \underline{a}, \text{ with c.s.}$$

$$u_l(\tilde{c}(z,P,Y), \tilde{l}(z,P,Y)) = u_c(\tilde{c}(z,P,Y), \tilde{l}(z,P,Y))w \tag{6}$$

$$\tilde{a}(z,P,Y) = (1+r)a + w \cdot \varepsilon \cdot (1 - \tilde{l}(z,P,Y)) - \tilde{c}(z,P,Y)\,, \tag{7}$$

where $\tilde{\lambda}^e$ represents an agent's forecast of λ' based on Y. Specifically, each agent forms a forecast using the following perceived law of motion (PLM):

$$\log \tilde{\lambda}^e = \log \overline{\lambda}^e + \langle \psi, Y \rangle\ , \tag{8}$$

where $\langle \cdot, \cdot \rangle$ is the standard inner product on $\mathbb{R}^n$, and $\overline{\lambda}^e$ is a time-t forecast of λ' in a corresponding stationary recursive equilibrium without any aggregate risk. We may express

$$\overline{\lambda}^e(\bar{a}(a,\varepsilon), \varepsilon) = \int \overline{\lambda}(\bar{a}(a,\varepsilon), \varepsilon')\Pi(\varepsilon, d\varepsilon')\,, \tag{9}$$

where ε' is an individual efficiency shock in period $t+1$, $\bar{a}$ is the stationary equilibrium level of claims in the current period, and $\overline{\lambda}(a,\varepsilon) = (1+\overline{r})u_c(\overline{c}(a,\varepsilon), \overline{l}(a,\varepsilon))$ such that $\overline{c}$, $\overline{l}$, and $\overline{r}$ represent the stationary equilibrium levels of consumption, labor, and the capital rate of return, respectively.[23]

This brings us to the principal challenge addressed in this paper. Given agent-specific states $z = (a, \varepsilon, \psi)$ and observable aggregates Y, conditions (5)–(7) determine agents' decision rules as a function of prices $P = (r, w)$. However, the realized value of prices and other endogenous aggregates is determined by market clearing, i.e., TE. Mechanically, this determination requires tracking the evolving distribution of agent-specific states. TE imposes $r = \theta f_k(k,n) - \delta$ and $w = \theta f_n(k,n)$, where (k,n) is determined by the following market clearing conditions:

$$k = \int a \cdot \Omega(dz) \quad \text{and} \quad n = \int (1 - \tilde{l}(z,P,Y))\, \Omega(dz)\,. \tag{10}$$

Notice that n depends on the agent-level policy rules $\tilde{l}(z,P,Y)$, which in turn depend implicitly on current factor prices P. All must be determined jointly in the TE as a solution to a system of nonlinear equations. Solving the system necessitates aggregating over the individual decisions of all agents for each candidate solution (n,r,w), which may pose a significant computational hurdle in a corresponding simulation with a large number of agents.

[2]In essence, we assume that agents' forecasts of the following period's shadow price consists of two components: (1) a rational forecast of next period's shadow price in a corresponding stationary equilibrium without aggregate risk; and (2) a boundedly rational forecast of the expected deviation of next period's shadow price from its stationary level attributed to variation in aggregate risk, as predicted by variation in the observed aggregate variables. For an in-depth discussion of this PLM and how agents update their beliefs, refer to Evans and McGough (2021) and Evans et al. (2023). In this paper – to approximate a TE – we treat the distribution of ψ over the population of agents as given.

[3]Giusto (2014) offers an alternative approach to incorporating agent-level bounded rationality in an HA model.

2.2 Approximation Methods

We present three distinct approaches to approximating TE in HA models and apply all of them in our example setting. The first method involves explicitly solving the policy rules for each individual agent in the simulation and aggregating their respective decisions in each iteration of the root solver at each period. The second method involves interpolating agent-level policy rules via basis function approximation at the start of a simulation and using them in the market-clearing root solver. While pre-computing equilibrium policy rules could improve performance, we demonstrate that this approach is inefficient in our boundedly rational multi-market setting as individual policy functions must be aggregated at each step of the root solver. However, we show that the structure of the interpolated policy functions allows us to aggregate them only once at the beginning of the period. This allows us to construct an approximation of aggregate demand conditional on the aggregate state using the basis functions over prices. In doing so, we eliminate the computational cost of aggregating the policy functions in each iteration of the root solver. This approach significantly outperforms the first method in terms of execution time, without suffering from a notable loss in accuracy.

Method #1 (Naive Global Approximation): Let Ω be approximated with N points $\{\overline{z}_i\}_{i\in\mathbb{N}_N}$ with corresponding weights $\{\omega_i\}_{i\in\mathbb{N}_N}$. The approximation for the demand schedule $\widetilde{X}(P,Z)\in\mathbb{R}^J$ for a given P and Z may be expressed as the following weighted sum:

$$\widehat{X}^j(P,Z) \equiv \sum_i \omega_i\, \tilde{x}^j(\overline{z}_i,P,Y)\,, \tag{11}$$

where $\tilde{x}$ is derived individually for each $(\overline{z}_i,P,Y)$ tuple, and $j=1,\ldots,J$ is an index for the set of goods in the economy. Repeating the above approximation for all J goods yields $\widetilde{X}(P,Z)$, which can then be used to solve the equilibrium vector of prices P^* using the market clearing condition in Eq. (1).

Solving the policy functions N times for a given P can be computationally taxing, especially if N is large. If an M number of steps is required for a root solver to converge to an approximation of P^*, then each $\tilde{x}^j$ has to be recomputed a total of $M \cdot N$ times. In our application, for a given wage level w and distribution $\widehat{\Omega}$ of individual states (a, ε, ψ), we analytically solve the level of capital k and the return on capital r. Using the above objects, we then solve the quantity of labor supplied individually by each agent, $1 - l_i$, according to conditions (5)–(7). Averaging these individual labor supply decisions across the set of all agents, as in Eq. (11), yields the aggregate labor supply n_t, as shown in Eq. (10).

Method #2 (Interpolated Global Approximation): This method adds a layer of approximation to Method 1. Suppose that $\tilde{x}$ is approximated by $\hat{x}$ via projection, such that the demand for the j-th good is represented by

$$\hat{x}^j(z, P, Y) = \sum_k \sum_l \sum_m c^j_{klm} \, \Phi^z_k(z) \, \Phi^P_l(P) \, \Phi^Y_m(Y) \,, \tag{12}$$

where Φ^P_l, Φ^z_k and Φ^Y_m are the basis functions that each depend on prices, idiosyncratic states, and the aggregate observables, respectively. The coefficients for this approximation, c^j_{klm}, need to be computed only once prior to simulation. The approximation for the demand schedule $\widetilde{X}(P, Z)$ may be expressed as the following weighted sum:

$$\widehat{X}^j(P, Z) \equiv \sum_i \omega_i \, \hat{x}^j(\overline{z}_i, P; Y) \,, \tag{13}$$

where $\overline{z}_i$ and ω_i are defined as before. Once again, Eq. (13) may be used to solve for P^* in Eq. (1).

This approach requires the policy function for the j-th good to be approximated only once, after which it is inputted into Eq. (13) to compute the sum. However, given a large N, the sum in Eq. (13) may still take a substantial amount of time to compute. This is especially true if the vector of idiosyncratic states z is high-dimensional, in which case each instance of a precomputed $\hat{x}$ may take a long time to execute – causing Method 1 to outpace Method 2.

Method #3 (Fast Approximation): Finally, we present our proposed method. Notice that substituting Eq. (12) into Eq. (13) yields

$$\widehat{X}^j(P, Z) = \sum_i \omega_i \sum_l \sum_k \sum_m c^j_{klm}\, \Phi^z_k(z)\, \Phi^P_l(P)\, \Phi^Y_m(Y)\,, \tag{14}$$

which can be rearranged by distributing ω_i, switching the order of summation, and factoring Φ^P_l in the following manner:

$$\widehat{X}^j(P, Z) = \sum_l \Phi^P_l(P) \sum_i \omega_i \sum_k \sum_m c^j_{klm} \Phi^z_k(z) \Phi^Y_m(Y)\,. \tag{15}$$

Finally, defining $C^j_l(Z) \equiv \sum_i \omega_i \sum_k \sum_m c^j_{klm}\, \Phi^z_k(z)\, \Phi^Y_m(Y)$ allows us to express Eq. (15) as

$$\widehat{X}^j(P, Z) = \sum_l C^j_l(Z)\, \Phi^P_l(P)\,. \tag{16}$$

With this formulation of $\widehat{X}$, since C^j_l is independent of P, it suffices to compute C^j_l only once before initializing the root solver to find P^*. This approach strictly dominates Method 2 by relying on the same initial projections but fewer mappings in the aggregation stage, so that execution times are lower but the accuracy is identical. More importantly, this method outpaces Method 1 without much loss in accuracy, as discussed in Section 3.

The execution speed of a TE solver becomes crucial in practice when the dynamic recursive equilibrium of an HA economy is simulated over multiple periods of time, in other words, when the TE solver is used repeatedly (Bakota, 2022). For such applications, the speedup offered by our method is particularly attractive. The case becomes even stronger in the context of stationary recursive equilibria, in which the idiosyncratic distribution Ω

is time-invariant – it becomes sufficient to compute C_l^j only once in the initial period of a simulation, beyond which it can be used in all future periods.

3 Application

In our application, we simplify certain elements of the model for convenience. We fix the TFP shock to $\theta = 1$, and evaluate the TE given an ergodic distribution of individual states $\widehat{\Omega}$.[4]

We compare the performance of the described TE solution methods across three dimensions: (1) the execution time of the aggregation procedure – in other words, the time it takes to find the sum of individual policy rules given all necessary inputs; (2) the execution time of the multidimensional mapping characterizing the TE, which needs to be computed at each step of the nonlinear solver; (3) the execution time of the nonlinear solver that approximates the set of equilibrium prices. We generate samples of these execution times and present the corresponding summary statistics in Table 1. Our proposed method (Method 3) offers significant speed-ups across all stages of the TE solution procedure.

In addition to comparing execution speed, we compare the accuracy with which our method estimates the aggregate labor supply schedule. In Figure 1, we plot the percentage deviation of the labor supply schedule approximated using Methods 2 and 3 from that obtained using Method 1 over a large interval surrounding the steady state wage level.[5] We find that our approach offers accuracy similar to Method 1, since the two labor supply schedules appear to be practically identical.

[4] In principle, the time and accuracy can depend on θ and Ω, but we found this not to be the case in our application.

[5] Note that Methods 2 and 3 must yield the same approximation.

	Method #1	Method #2	Method #3
Aggregation	3.30e-2	4.20e+1	4.00e-6
	(2.40e-2)	(3.94e+1)	(3.20e-6)
Temporary Eq. Mapping	6.11e-2	4.33e+1	3.99e-6
	(5.01e-2)	(4.32e+1)	(3.44e-6)
Nonlinear Solver	6.82e-1	4.86e+2	3.36e-4
	(6.57e-1)	(4.39e+2)	(3.22e-4)
Sample Size			1000

Table 1: Temporary equilibrium solution methods execution times. *Note:* Mean execution time in seconds (minimum execution time in parentheses).

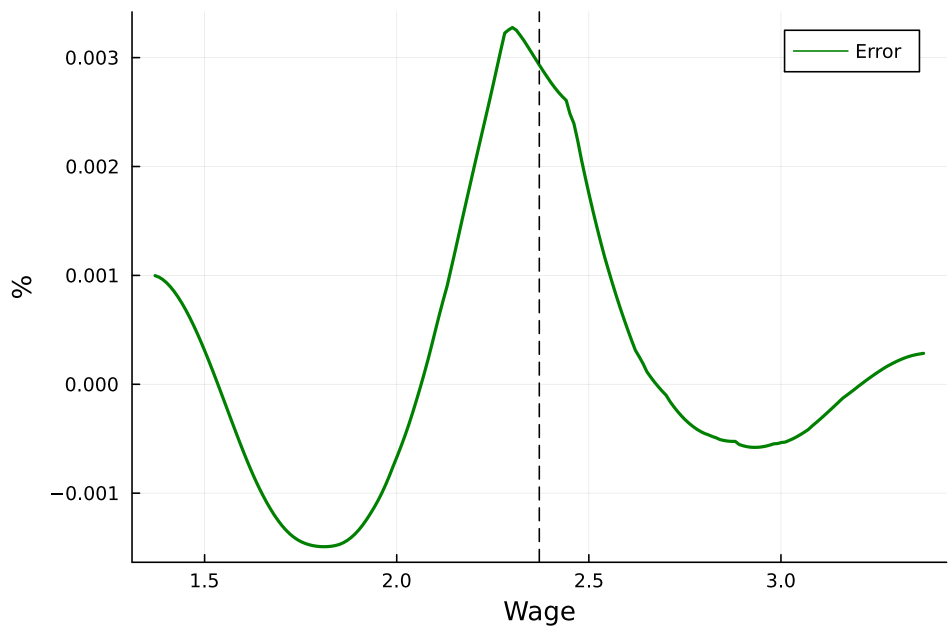

Figure 1: Percentage deviation of aggregate labor supply as a function of wage approximated using Method 3 from that obtained using Method 1. *Note:* The steady state wage level is represented by the dashed vertical line.

References

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684. https://doi.org/10.2307/2118417.

Bakota, I. (2022). Market clearing and krusell-smith algorithm in an economy with multiple assets. *Computational Economics*, 62:1007–1045. https://doi.org/10.1007/s10614-022-10290-2.

Chang, Y. and Kim, S.-B. (2006). From individual to aggregate labor supply: A quantitative analysis based on a heterogeneous agent macroeconomy. *International Economic Review*, 47(1):1–27. https://doi.org/10.1111/j.1468-2354.2006.00370.x.

Chang, Y. and Kim, S.-B. (2007). Heterogeneity and aggregation: Implications for labor-market fluctuations. *American Economic Review*, 97(5):1939–1956. https://doi.org/10.1257/aer.97.5.1939.

Den Haan, W. J., Judd, K. L., and Juillard, M. (2010). Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty. *Journal of Economic Dynamics and Control*, 34(1):1–3. https://doi.org/10.1016/j.jedc.2009.07.001.

Evans, D., Li, J., and McGough, B. (2023). Local rationality. *Journal of Economic Behavior & Organization*, 205:216–236. https://doi.org/10.1016/j.jebo.2022.11.005.

Evans, G. W. and McGough, B. (2021). Agent-level adaptive learning. In *Oxford Research Encyclopedia of Economics and Finance*. Oxford University Press. https://doi.org/10.1093/acrefore/9780190625979.013.620.

Giusto, A. (2014). Adaptive learning and distributional dynamics in an incomplete markets model. *Journal of Economic Dynamics and Control*, 40:317–333. https://doi.org/10.1016/j.jedc.2014.01.007.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5):953–969. https://doi.org/10.1016/0165-1889(93)90024-M.

Judd, K. L. (1992). Projection methods for solving aggregate growth models. *Journal of Economic Theory*, 58(2):410–452. https://doi.org/10.1016/0022-0531(92)90061-L.

Judd, K. L. (1996). Chapter 12 approximation, perturbation, and projection methods in economic analysis. In *Handbook of Computational Economics*, volume 1, pages 509–585. Elsevier. https://doi.org/10.1016/S1574-0021(96)01014-3.

Krusell, P., Mukoyama, T., and Şahin, A. (2010). Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations. *The Review of Economic Studies*, 77(4):1477–1507. https://doi.org/10.1111/j.1467-937X.2010.00700.x.

Krusell, P. and Smith, A. A. (1997). Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns. *Macroeconomic Dynamics*, 1(2):387–422. https://doi.org/10.1017/S1365100597003052.

Krusell, P. and Smith, A. A. (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy*, 106(5):867–896. https://doi.org/10.1086/250034.