Higher-Order Perturbation in Sequence Space: the Certainty Correspondence

Adrien Auclert* Rodolfo Rigato† Matthew Rognlie‡ Ludwig Straub§

March 2026

Abstract

We show that higher-order perturbation solutions to rational expectations models exhibit a *certainty correspondence* property: they can be obtained from the derivatives of the model's perfect foresight solution. This extends the well-known certainty equivalence result beyond first order. For heterogeneous-agent models with large idiosyncratic state spaces, our approach is significantly faster than existing state-space perturbation or global solution methods. We apply our method to solve a HANK model and a menu cost model to third order in aggregate risk, quantifying the time variation in the fiscal multiplier and the responsiveness of inflation to costs, as well as the effect of perceived aggregate risk on the steady state and impulse responses.

*Stanford University, CEPR and NBER. Email: aauclert@stanford.edu.
†European Central Bank. Email: rodolfo.dinis_rigato@ecb.europa.eu.
‡Northwestern University and NBER. Email: matthew.rognlie@northwestern.edu.
§Harvard University, CEPR and NBER. Email: ludwigstraub@g.harvard.edu.
This research is supported by the National Science Foundation grant number SES-234393 and the Harvard Chae Initiative. The views expressed are those of the authors and do not necessarily reflect those of the European Central Bank.

1 Introduction

Heterogeneous-agent models with aggregate risk are notoriously difficult to solve in general equilibrium. A recent literature (Boppart, Krusell and Mitman 2018, Auclert, Bardóczy, Rognlie and Straub 2021) has made progress by exploiting *certainty equivalence*: the fact that to first order in shock size, aggregate risk is neutral, and future variables can be replaced by their expectations. A consequence is that to first order, the impulse response to an "MIT shock"—a one-time, unexpected shock, with perfect foresight thereafter—is the same as in a stochastic model. Such impulse responses can be obtained with sequence-space methods, sidestepping the curse of dimensionality that often arises in the state space, and enabling faster and easier computation.

The restriction to first-order aggregate risk, however, is currently a major drawback of the sequence-space approach. In a first-order approximation, there is no role for size- or history-dependent effects, nor is there any precautionary response to risk. It is impossible to study how aggregate shocks, or policies responding to them, affect welfare.

In this paper, we overcome these limitations by extending sequence-space methods to higher order. We show that certainty equivalence can be generalized to what we call the *certainty correspondence*: a systematic relationship between MIT-shock impulse responses and the higher-order perturbation solution. We derive this explicitly for the second- and third-order cases, showing how each term can be obtained via the certainty correspondence, and apply these methods numerically to obtain the full third-order solution of a HANK model and a menu cost model. Our perturbation solutions coincide with the state-space approach with pruning (Andreasen, Fernández-Villaverde and Rubio-Ramírez 2018), and can be easily used to derive generalized impulse response functions and moments of the stochastic solution. The new approach is straightforward to implement for anyone who can calculate nonlinear impulses to MIT shocks, and we provide additional tools to speed up general equilibrium.

The foundation of our approach is the nonlinear sequence-space solution $X_t = \mathcal{X}(\sigma, \epsilon_t, \epsilon_{t-1}, \ldots)$, which expresses the path of an economy as a function of anticipated risk σ and the history $\epsilon_t, \epsilon_{t-1}, \ldots$ of realized shocks.[1] Whenever a locally stable state-space solution $X_t = g(\sigma, X_{t-1}, \epsilon_t)$ exists, a sequence-space solution also exists, and we can use perturbation to approximate it around the deterministic steady state. To first order, this is exactly the approach advocated by Boppart et al. (2018): since $\mathcal{X}_\sigma = 0$ (certainty equivalence), the first-order solution is determined by the derivatives $\mathcal{X}_{-j}$ with respect to past shocks ϵ_{t-j},

[1] Lan and Meyer-Gohde (2013b) also study such a solution, which they call a "nonlinear moving average", and Lombardo and Uhlig (2018) study a similar function where the history is finite—starting at date 0—rather than infinite. We discuss the relation to these papers further in the related literature section.

which can be identified with the first-order impulse response to an MIT shock.

Moving to second order, certainty equivalence no longer holds, since there is a constant risk term $\mathcal{X}_{\sigma\sigma}$. All other terms in the second-order expansion, however, are independent of risk σ and can be obtained by differentiating impulse responses to MIT shocks. We show that, remarkably, the risk term $\mathcal{X}_{\sigma\sigma}$ can *also* be obtained from the information in MIT-shock solutions: it is the effect from perturbing steady-state expectations by the size-dependence term $\mathcal{X}_{0,0}$. This is the second-order certainty correspondence.

With the second-order sequence-space solution in hand, we can read off various concepts of interest. For instance, $\mathcal{X}_{\sigma\sigma}$ gives the risky steady state: the steady state when risk is anticipated, but no shocks are ever realized. Summing this with the size-dependent terms, $\mathcal{X}_{0,0} + \mathcal{X}_{1,1} + \ldots$, gives the second-order ergodic mean of $X_t - X$. Second-order impulse responses are almost identical to their MIT shock counterparts, but with a correction reflecting expected size dependence.

We follow a similar approach for the third-order perturbation solution. Most terms in the third-order expansion of $\mathcal{X}$ are independent of risk and can be obtained from MIT shocks. The risk-dependent terms have the form $\mathcal{X}_{-j,\sigma,\sigma}$, and capture how risk alters the transmission from ϵ_{t-j} to X_t. In part, these terms reflect the economy starting at the risky steady state $\mathcal{X}_{\sigma\sigma}$; this contribution is easy to obtain by shifting the steady state at which first-order impulses are calculated. The risk-dependent terms also reflect interactions between size and history dependence—which can again be obtained from MIT shocks, completing the third-order certainty correspondence.

As previously shown by Lombardo and Uhlig (2018), our sequence-space perturbation solution agrees with the pruned state-space solution. To obtain approximate moments of the stochastic solution, the conventional approach is to directly evaluate these moments on the perturbation approximation, as in Andreasen et al. (2018). We propose a different approach: a formal perturbation of the nonlinear moments themselves. This yields parsimonious formulas for each moment: for instance, evaluating covariances at any horizon using the first-order solution as in Auclert et al. (2021) is accurate to third order, and augmenting this with some terms from the second- and third-order solutions increases accuracy to fifth order.

The main requirement of our approach is the ability to compute nonlinear responses to MIT shocks, and to repeatedly differentiate them around the steady state—possibly via either numerical or automatic differentiation. Although we do not require that these responses are obtained with any particular solution method, one can exploit analytical results to improve upon a naive approach. For heterogeneous-agent models, nonlinear impulse responses are typically obtained by reducing the problem to a limited number of

unknown sequences, then iterating until convergence with a variant of Newton's method. This iterative process, however, is unnecessary for a perturbation solution. Instead, we show that the nth-order solution can be obtained by evaluating equilibrium conditions on the $n-1$th-order solution, finding the nth-order error, and then applying the same inverse Jacobian used to solve the first-order system. This "one-shot principle" makes higher-order solutions much easier to obtain.

Throughout the paper, we apply our methods to a number of models. Our primary application is a standard though "smooth" heterogeneous-agent model, first in partial and then in general equilibrium.[2] In partial equilibrium, we study the response to income shocks—finding that, as expected from the concavity of the consumption function, the second-order term for contemporaneous consumption is negative. Interestingly, however, we find that the second-order terms for consumption in subsequent periods are *positive*—so that, as we scale up an income shock, the delayed effect on consumption increases more than proportionally. This result extends to a general equilibrium HANK setting where income fluctuates due to deficit-financed transfers. As argued in Auclert, Rognlie and Straub (2023), these intertemporal spillovers may be party responsible for sustained high demand in the US after the 2020–21 fiscal interventions; the fact that these intertemporal spillovers are even stronger for large shocks strengthens this argument, given the unprecedented size of these interventions. We also study how the impact transfer multiplier varies with the level of current economic activity, finding that the multiplier can be twice as large or more in times when the economy is depressed, balance sheets are weak, and MPCs are high.

Our second main application is a menu cost model with idiosyncratic shocks in the spirit of Golosov and Lucas (2007). For that model, we study how the responsiveness of inflation to nominal marginal cost is affected by the current level of inflation. In line with other recent studies that use similar models (e.g. Blanco, Boar, Jones and Midrigan 2024), we find a significantly higher responsiveness of inflation to costs in times of high inflation.[3]

Related literature. The closest existing work to ours is Lan and Meyer-Gohde (2013b), which studies a nonlinear sequence-space solution of the same form (a "nonlinear mov-

[2]The "smooth" model is an otherwise standard incomplete market (SIM) model with additional i.i.d. shocks to income drawn from a smooth density. As we discuss, this feature of the model is important to make the model differentiable to the second and higher order, overcoming the difficulties with the non-smooth model highlighted by Bhandari, Bourany, Evans and Golosov (2023).

[3]We also verify that our formulas and our implementation of them are correct by checking that they produce the same solution as state-space perturbation up to third order on a neoclassical growth model.

ing average") and its perturbation solution up to third order, which is known as a Volterra expansion. Beyond this similar core, however, the two papers diverge: Lan and Meyer-Gohde use techniques that parallel the state-space solution[4], while we introduce the certainty correspondence. As we discuss further in section 2.3, the certainty correspondence does not hold in their framework, because of a difference in where expectations are taken.

Another closely related paper is Lombardo and Uhlig (2018), which uses a nonlinear sequence-space solution to explain why the state-space perturbation solution should be pruned: namely, because this makes it equivalent to the perturbed sequence-space solution.[5] We re-derive this important result in our setting. Unlike our paper and Lan and Meyer-Gohde, however, Lombardo and Uhlig only track the finite history since time 0, which makes their sequence-space solution vary with t.

Our paper also connects to the broader state-space perturbation literature, including the seminal work of Schmitt-Grohé and Uribe (2004) and Kim, Kim, Schaumburg and Sims (2008), and the pruned state-space system of Andreasen et al. (2018). Several recent papers, including Bilal (2023), Reiter (2023), and Bayer, Luetticke, Weiss and Winkelmann (2026), have extended state-space perturbation to the heterogeneous-agent context. For rich enough heterogeneity, however, this can require significant "model reduction"—as the size of a second-order state-space solution is cubic in the dimension of the state, and becomes prohibitively large and costly when this dimension is high. The recent work of Bhandari et al. (2023) provides a hybrid of state- and sequence-space approaches, and derives a second-order Volterra expansion for heterogeneous-agent models. To our knowledge, our work is the first to obtain a third-order solution, without model reduction, in the presence of substantial heterogeneity.

Finally, our paper is connected to a literature that seeks "global" solutions to heterogeneous-agent models to study the effect of aggregate risk beyond first order. This literature uses a variety of approaches, such as neural networks (Kase, Melosi and Rottner 2022), deep learning (Han, Yang et al. 2021, Azinovic-Yang and Žemlička 2025, Druedahl, Huleux and Røpke 2026), structural reinforcement learning (Yang, Wang, Schaab and Moll 2025), and alternative simulation-based methods (Lee 2025). We view this literature as highly complementary. Global solutions may be useful for models with sufficiently extreme nonlinearities, but they often require additional assumptions and substantial computational power—while our third-order solution uses an unreduced HANK model, and can be obtained in a matter of minutes on a laptop. The extent to which global solutions

[4]As they note in footnote 21: "Thus, our nonlinear moving average solution parallels nonlinear state space solutions in a manner analogous to the linear case, where the recursion is in the coefficients as opposed to the variables themselves."

[5]See also related work in Lan and Meyer-Gohde (2013a).

differ from our third-order solution for typical heterogeneous-agent models remains an intriguing open question.

2 Certainty correspondence

2.1 General setup

We consider a model driven by an exogenous shock process ϵ_t with i.i.d. innovations.[6] In the model, the endogenous variable $X_t \in \mathbb{R}^n$ is determined at t to solve

$$F\left(X_{t-1}, X_t, \mathbb{E}_t\left[X_{t+1}\right], \epsilon_t\right) = 0, \tag{1}$$

with F mapping $\left(\mathbb{R}^n\right)^3 \times \mathbb{R}$ to $\mathbb{R}^n$.

The expectations in (1) are taken assuming that $\epsilon_t = \sigma\bar{\epsilon}_t$, where $\bar{\epsilon}_t$ are drawn i.i.d. from a distribution that is symmetric around zero and has unit variance, and σ is a scalar that indexes the standard deviation of shocks. The actual, realized ϵ_t are given by $\epsilon_t = \sigma_R\bar{\epsilon}_t$. Our objective is to solve the rational expectations case where $\sigma = \sigma_R$, but our "certainty correspondence" result will connect this case to perfect-foresight or "MIT shock" impulse responses, where $\sigma = 0$ but $\sigma_R \neq 0$.

A *deterministic steady state* of the model is a constant vector $X \in \mathbb{R}^n$ that solves

$$F\left(X, X, X, 0\right) = 0.$$

For a given σ, a *nonlinear state-space solution* of the model is a function g such that

$$X_t = g\left(X_{t-1}, \epsilon_t; \sigma\right) \tag{2}$$

solves equation (1) for all X_{t-1} and all ϵ_t. This is a standard solution concept. Classic state-space perturbation (e.g., Schmitt-Grohé and Uribe 2004, Fernández-Villaverde, Rubio-Ramírez and Schorfheide 2016) performs a Taylor expansion of the function g in (2) around the deterministic steady state $(X, 0; 0)$.[7]

In this paper, we consider instead a *nonlinear sequence-space solution* of the model, by which we mean a function $\mathcal{X}$ such that, given σ,

$$X_t = \mathcal{X}\left(\sigma; \epsilon_t, \epsilon_{t-1}, \epsilon_{t-2}, \ldots\right) \tag{3}$$

[6] We assume a single exogenous shock for notational convenience. Results generalize to multiple exogenous shocks, see the appendix.

[7] We provide conditions for the local existence and differentiability of (2) in the appendix.

solves equation (1) for all sequences $\{\epsilon_t, \epsilon_{t-1}, \epsilon_{t-2}, \ldots\}$.[8] Our sequence-space perturbation approach will be to perform a Taylor expansion of the $\mathcal{X}$ function in (3) around the deterministic steady state $(0; 0, 0, 0, \ldots)$.

The state-space solution has a nonlinear vector autoregressive (VAR) form while the sequence-space solution has a nonlinear vector moving average (VMA) form. It is a well-known result from time series (eg Hamilton 1994, chapter 10) that, under a stability condition, linear VAR and VMA are equivalent representations of the same process. The following result shows that this equivalence carries over to nonlinear solutions of the model in (1).

Proposition 1. *Assume that the state-space solution g exists and is stable in a neighborhood of the steady state. Then the sequence-space solution $\mathcal{X}$ exists locally as well, and is equivalent to the state-space solution in the sense that*

$$\mathcal{X}\left(\sigma; \epsilon_t, \epsilon_{t-1}, \epsilon_{t-2}, \ldots\right) = g\left(\mathcal{X}\left(\sigma; \epsilon_{t-1}, \epsilon_{t-2}, \epsilon_{t-3}, \ldots\right), \epsilon_t; \sigma\right) \quad \forall \sigma, \forall\left(\epsilon_t, \epsilon_{t-1}, \epsilon_{t-2}, \ldots\right) \tag{4}$$

In addition, if g is $\mathcal{C}^n$, then so is $\mathcal{X}$.

Given proposition 1, state-space and sequence-space perturbation solutions are closely connected. We review this connection, which has previously been established by Lan and Meyer-Gohde (2013b) and Lombardo and Uhlig (2018), in section 2.3 below. Our main contribution is to further connect the sequence-space perturbation solution to the (repeated-surprise) perfect-foresight solution, $\mathcal{X}(0; \epsilon_t, \epsilon_{t-1}, \epsilon_{t-2}, \ldots)$.

Generalized impulse response functions. A common way to summarize the behavior of the model is in terms of its impulse responses. We define the generalized impulse response of X as the revision of the forecast of X, j periods ahead due to an innovation to ϵ at date t, $IRF_j^X \equiv \mathbb{E}_t[X_{t+j}] - \mathbb{E}_{t-1}[X_{t+j}]$. This revision generally depends on the size of perceived shocks σ, on the size of the date-t innovation ν, and on the state entering period t, though it does not further depend on the calendar time t beyond these states. If we consider a state-space solution (2), then the state entering date t is X_{t-1},

$$IRF_j^X\left(\nu, X_{t-1}; \sigma\right) \equiv \mathbb{E}\left[X_{t+j} | \epsilon_t = \nu, X_{t-1}; \sigma\right] - \mathbb{E}\left[X_{t+j} | X_{t-1}; \sigma\right] \quad j \geq 0 \tag{5}$$

[8] Note that σ_R does not directly appear in either (2) or (3); given σ, only the realized sequence of ϵ_t matters for X_t.

while if we consider a sequence-space solution (3), then the state is $\{\epsilon_{t-k}\}_{k=1}^{\infty}$,

$$IRF_j^X\left(\nu, \{\epsilon_{t-k}\}_{k=1}^{\infty}; \sigma\right) \equiv \mathbb{E}\left[X_{t+j}|\epsilon_t = \nu, \{\epsilon_{t-k}\}_{k=1}^{\infty}\right] - \mathbb{E}\left[X_{t+j}|\{\epsilon_{t-k}\}_{k=1}^{\infty}\right] \quad j \geq 0 \quad (6)$$

Symmetry in σ. We note that since $\bar{\epsilon}_t$ is distributed symmetrically around zero, the distributions of $\sigma\bar{\epsilon}_t$ and $-\sigma\bar{\epsilon}_t$ are identical. This implies:

Lemma 1. Both g and $\mathcal{X}$ are even functions of σ. Assuming they exist, the derivatives of both with respect to σ around $\sigma = 0$ are identically zero.

We will illustrate our results using a number of examples.

Example 1: The Smooth Standard Incomplete Markets (SIM) Model. As our main application, we consider a heterogeneous-agent model where consumers smooth idiosyncratic income risk subject to a borrowing constraint. The Bellman equation is

$$\begin{aligned} V_t(a,e) &= \mathbb{E}_\xi\left[\max_{c,a'} u(c) + \beta\mathbb{E}_t[V_{t+1}(a',e')|e]\right] \\ \text{s.t.} \qquad & a' + c = (1+r)a + Y_t\xi e; \qquad a' \geq \underline{a} \end{aligned} \quad (7)$$

where e follows a finite Markov chain with a fixed transition matrix Π, ξ is a smoothly distributed i.i.d. shock with density $f(\cdot)$ realized at the beginning of the period, $u(\cdot)$ is CRRA utility, and Y_t is aggregate income, which follows an AR(1) process, $Y_t - 1 = \rho(Y_{t-1} - 1) + \epsilon_t$.

Optimal consumption and asset policy conditional on (a, e, ξ) are characterized by a first-order condition $u'(c) \geq \beta\mathbb{E}_t[V_{a,t+1}(a', e')|e]$, which holds with equality unless the borrowing constraint binds. The marginal value function $V_{a,t}$ is then given by the envelope condition $V_{a,t}(a,e) = (1+r)\mathbb{E}_\xi u'(c(a,e,\xi))$.

For each e, let $\Phi_t(a', e)$ denote the end-of-period CDF of assets a' saved by households, conditional on current income state e. Given the Φ_{t-1} from the end of the previous period, the asset policy at t, together with the fixed Π and distribution of ξ, determines Φ_t. Finally, aggregate consumption at t is given by $C_t = \sum_{e,e_-} \Pi_{e_-e} \int\int c_t(a,e,\xi) f(\xi) d\xi d\Phi_{t-1}(a, e_-)$.

In practice, we represent $V_{a,t}(a,e)$ using its value at n_a asset gridpoints a, for each of the n_e income states e, and let $\boldsymbol{v}_t$ be the vector of these $n_a \times n_e$ values. We represent $\Phi_t(a', e)$ with a cubic spline interpolating over the same gridpoints for each e, and let $\boldsymbol{\phi}_t$ be the vector of the $n_a \times n_e$ spline coefficients.

To put the system in the notation of (1), we stack these vectors, along with C_t and Y_t,

in the vector $X_t \equiv (\boldsymbol{v}_t, \boldsymbol{\phi}_t, C_t, Y_t)$. We then stack the $2(n_a n_e + 1)$ equations into

$$F\left(X_{t-1}, X_t, \mathbb{E}_t\left[X_{t+1}\right], \epsilon_t\right) \equiv \begin{bmatrix} \boldsymbol{v}_t - \mathcal{V}(Y_t, \mathbb{E}_t \boldsymbol{v}_{t+1}) \\ \boldsymbol{\phi}_t - \Lambda(Y_t, \mathbb{E}_t \boldsymbol{v}_{t+1}, \boldsymbol{\phi}_{t-1}) \\ C_t - \mathcal{C}(Y_t, \mathbb{E}_t \boldsymbol{v}_{t+1}, \boldsymbol{\phi}_{t-1}) \\ Y_t - 1 - \rho(Y_{t-1} - 1) - \epsilon_t \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \tag{8}$$

where $\mathcal{V}$, Λ, and $\mathcal{C}$ are constructed as described above, acting on the discretized representation.

The appendix provides more details for both the underlying model and its computation. One crucial feature of this model is that it is "smooth": due to the smooth ξ shock, $V_{a,t}$ and Φ_t are both infinitely differentiable functions. This is important since, as emphasized by Bhandari et al. (2023), the standard incomplete market model without a smooth ξ is not second-order differentiable, which complicates the implementation of perturbation.

In section 3.2, we further develop a general equilibrium version of this model.

Example 2: The Neoclassical Growth Model. To facilitate comparison with existing state-space methods, we also consider the neoclassical growth model with shocks, as in e.g. Schmitt-Grohé and Uribe (2004). In this model, a planner maximizes the utility $\mathbb{E}\sum \beta^t C_t^{1-\sigma}/(1-\sigma)$ of a representative agent, subject to the feasibility constraint $C_t + K_t - (1-\delta)K_{t-1} = e^{a_t} K_{t-1}^{\alpha}$, and an AR(1) for log TFP $a_t = \rho a_{t-1} + \epsilon_t$. The optimum in this model is characterized by an Euler equation $C_t^{-\sigma} = \beta \mathbb{E}_t \left(1 - \delta + \alpha e^{a_{t+1}} K_t^{\alpha-1}\right) C_{t+1}^{-\sigma}$.

To put this in the notation of (1), we define $\lambda_t \equiv \left(1 - \delta + \alpha e^{a_t} K_{t-1}^{\alpha-1}\right) C_t^{-\sigma}$ to be the shadow value of capital at t, and define $X_t \equiv (a_t, \lambda_t, K_t, C_t)$. Then we can stack

$$F\left(X_{t-1}, X_t, \mathbb{E}_t\left[X_{t+1}\right], \epsilon_t\right) \equiv \begin{bmatrix} \lambda_t - \left(1 - \delta + \alpha e^{a_{t+1}} K_t^{\alpha-1}\right) C_{t+1}^{-\sigma} \\ C_t^{-\sigma} - \beta \mathbb{E}_t\left[\lambda_{t+1}\right] \\ C_t + K_t - (1-\delta)K_{t-1} - e^{a_t} K_{t-1}^{\alpha} \\ a_t - \rho a_{t-1} - \epsilon_t \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \tag{9}$$

It is important to note that in both (8) and (9), all expectations at time t show up only through $\mathbb{E}_t[X_{t+1}]$, in line with our form for F: in (8), through $\mathbb{E}_t \boldsymbol{v}_{t+1}$, and in (9), through $\mathbb{E}_t \lambda_{t+1}$. It is usually possible to cast models in this form, which emerges naturally in dynamic programming—where the future appears through the expectation of a value function or its derivative. This feature of F will be crucial to deriving the higher-order certainty correspondence.

2.2 First-order perturbation and certainty equivalence

Assuming differentiability, and noting that $\mathcal{X}_\sigma = 0$ by lemma 1, a first-order Taylor approximation of X_t around the deterministic steady state can be written as

$$X_t \approx X + \sum_{j=0}^{\infty} \frac{\partial \mathcal{X}}{\partial \epsilon_{-j}} \epsilon_{t-j} \tag{10}$$

where all derivatives are evaluated at $(\sigma, \epsilon_t, \epsilon_{t-1} \ldots) = (0, 0, 0, \ldots)$. Hence, to first order, $X_t - X$ is $MA(\infty)$ in ϵ_t. An important feature of (10) is that σ does not appear, so that anticipated risk does not affect the first-order solution. Further, (10) implies

$$IRF_j^X \left(\nu, \{\epsilon_{t-k}\}_{k=1}^{\infty}; \sigma\right) = \frac{\partial \mathcal{X}}{\partial \epsilon_{-j}} \nu$$

so the impulse response is linear in ν and independent of the state.

We next define the sequence

$$X_j^{MIT}(\nu) \equiv \mathcal{X}(0; 0, \ldots, 0, \underbrace{\nu}_{j}, 0 \ldots)$$

where the argument at a lag of j is ν, as the impulse response to an *MIT shock*: a one-time, unanticipated shock of size ν at date 0, starting from the steady state and with perfect foresight thereafter. We then clearly have:

$$\frac{\partial \mathcal{X}}{\partial \epsilon_{-j}} = (IRF_j^X)'(0) = (X_j^{MIT})'(0) \tag{11}$$

Hence, the $MA(\infty)$ coefficients in (10) can be obtained as the local impulse response to an MIT shock (see, e.g. Boppart et al. 2018). This is an especially helpful result for heterogeneous-agent models, since there are well-known procedures to calculate this MIT shock impulse response in these models, both in partial equilibrium where a simple backward-forward procedure suffices, or in general equilibrium where one can use the sequence-space Jacobian method (Auclert et al. 2021). We summarize this result as:

Proposition 2 (Certainty equivalence)**.** *The first-order perturbation of X_t in the sequence space is given by the linear $MA(\infty)$ in equation (10). The constant term is the deterministic steady state X, and the coefficients on past innovations are given by (11), so they can be obtained from linearized perfect-foresight impulse responses.*

Given the $MA(\infty)$ representation of X_t in (10), we can immediately calculate the sec-

ond moments of X_t analytically. We show in section 2.6 that these moments are third-order accurate from the perspective of the aggregate risk solution.

2.3 Second-order perturbation and certainty correspondence

Assuming second differentiability, and noting that lemma 1 also implies that any mixed derivatives $\mathcal{X}_{\sigma j}$ with respect to σ and ϵ_{t-j} are zero, a second-order Taylor approximation of X_t, also known as the second-order Volterra expansion, can be written as

$$X_t \approx X + \sum_{j=0}^{\infty} \frac{\partial \mathcal{X}}{\partial \epsilon_{-j}} \epsilon_{t-j} + \frac{1}{2} \left(\mathcal{X}_{\sigma\sigma}\sigma^2 + \sum_{j,k\geq 0} \frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-j} \partial \epsilon_{-k}} \epsilon_{t-j}\epsilon_{t-k} \right) \tag{12}$$

where, again, all derivatives are evaluated at $(\sigma, \epsilon_t, \epsilon_{t-1} \ldots) = (0, 0, 0, \ldots)$. Hence, to second-order, up to constants, X_t is a quadratic infinite moving-average process in ϵ_t.

Using symmetry of mixed partials, the new second-order terms in (12) can be usefully grouped as follows:

$$\underbrace{\frac{1}{2}\mathcal{X}_{\sigma\sigma}\sigma^2}_{\text{Risk}} + \underbrace{\frac{1}{2}\sum_{j=0}^{\infty} \frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-j}^2} \epsilon_{t-j}^2}_{\text{Size}} + \underbrace{\sum_{j=0}^{\infty}\sum_{l=1}^{\infty} \frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-j} \partial \epsilon_{-(j+l)}} \epsilon_{t-j}\epsilon_{t-(j+l)}}_{\text{History}} \tag{13}$$

Now, anticipation of future risk shows up as a constant "risk" term.

Notions of steady states. In the presence of aggregate risk, there are two natural notions of "steady state". The *risky steady state* (Coeurdacier, Rey and Winant 2011) arises when agents continually anticipate aggregate risk, even though realized risk is $\sigma_R = 0$, so that $\epsilon_t = 0$ at all dates. Given (12), the system is then perpetually at $X + \frac{1}{2}\mathcal{X}_{\sigma\sigma}\sigma^2$. The *ergodic mean* is the average value of the system when risk realizes at its expected value $\sigma_R = \sigma$. This mean can be immediately derived from (12) as $\mathbb{E}[X] = X + \frac{1}{2}\mathcal{X}_{\sigma\sigma}\sigma^2 + \frac{1}{2}\sum_{j=0}^{\infty} \frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-j}^2}\sigma^2$. While these notions all collapse to the deterministic steady state X in the first-order solution (10), they generally differ from it in the second-order solution.

Generalized impulse response functions. Given (12), we calculate the generalized impulse response function as :

$$IRF_j\left(\nu, \{\epsilon_{t-k}\}_{k=1}^{\infty}; \sigma\right) = \frac{\partial \mathcal{X}}{\partial \epsilon_{-j}}\nu + \frac{1}{2}\frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-j}^2}\left(\nu^2 - \sigma^2\right) + \sum_{k=1}^{\infty} \frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-j} \partial \epsilon_{-(j+k)}} \nu \epsilon_{t-k} \tag{14}$$

In addition to the linear term from the first order impulse response, the generalized impulse response is now nonlinear in ν due to the quadratic term in (14), as well as state-dependent. The impulse response at date j is concave in size ν if $\frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-j}^2} < 0$ and convex otherwise; when past innovations are 0, it coincides with the linear impulse response when $\nu = \pm\sigma$, with this dependence on σ reflecting expected size dependence.[9] Further, past innovations change the linear impulse response around 0.

Certainty correspondence. For the size and history terms, we can use the same connection to perfect foresight that we used for the first order term. First, we have:

$$\frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-j}^2} = (X_j^{MIT})''(0) \tag{15}$$

which connects the size effect in the quadratic MA expression for X_t to the size effect in the MIT shock impulse response. Equations (11) and (15) then imply that

$$X_j^{MIT}(\nu) - X \approx \frac{\partial \mathcal{X}}{\partial \epsilon_{-j}}\nu + \frac{1}{2}\frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-j}^2}\nu^2 \tag{16}$$

so that the MIT shock impulse response now differs from the generalized impulse response (14) since the latter features expected size dependence while the former does not. Further, for $l \geq 1$ we can define the sequence

$$X_j^{MIT,l}(\epsilon, \nu) \equiv \mathcal{X}(0; 0, \dots 0, \underbrace{\epsilon}_{j-l}, 0, \dots, 0, \underbrace{\nu}_{j}, 0 \dots)$$

as the impulse response to a sequence of two consecutive MIT shocks: a first unexpected shock of size δ occurring at date $-l$, and a second unexpected shock of size ϵ occurring at date 0. Then we have that

$$\frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-j} \partial \epsilon_{-(j+l)}} = \frac{\partial^2 X_j^{MIT,l}}{\partial \nu \partial \epsilon} \tag{17}$$

which connects the history effect in the second-order Volterra expansion of X_t to the local interaction term of the impulse response to two consecutive MIT shocks. Indeed, (11), (15) and (17) imply that

$$X_j^{MIT,l}(\epsilon, \nu) - X \approx \frac{\partial \mathcal{X}}{\partial \epsilon_{-j}}\nu + \frac{1}{2}\frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-j}^2}\nu^2 + \frac{\partial \mathcal{X}}{\partial \epsilon_{-(j+l)}}\epsilon + \frac{1}{2}\frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-(j+l)}^2}\epsilon^2 + \frac{\partial^2 \mathcal{X}}{\partial \epsilon_{-j} \partial \epsilon_{-(j+l)}}\nu\epsilon \tag{18}$$

[9] This result was previously noted by Andreasen et al. (2018).

so the MIT shock impulse response is the sum of the two nonlinear impulse responses to each shock independently, plus the same interaction term that enters the generalized impulse response in (14).

Equations (15) and (17) show that the new size and history coefficients in the second-order sequence space perturbation can be obtained from the perfect foresight solution, just like in the certainty equivalence result of proposition 2. To complete the certainty correspondence result, we now show how to obtain the risk term using the same information. To do this, we introduce the concept of a *perturbed-expectations steady state*, which for any δ solves the equation

$$F\left(X^{pss}(\delta), X^{pss}(\delta), X^{pss}(\delta) + \frac{\partial^2 \mathcal{X}}{\partial \epsilon_0^2}\delta, 0\right) = 0, \tag{19}$$

Observe that, by (15), $X^{pss}(\delta)$ can be solved just like the deterministic steady state, provided one replaces, everywhere they appear, occurrences of $\mathbb{E}_t[X_{t+1}]$ by the (yet unknown) steady state plus the immediate size dependence $(X_0^{MIT})''(0)$, scaled by δ. Then, we have that:

$$\mathcal{X}_{\sigma\sigma} = (X^{pss})'(0) \tag{20}$$

The intuition is that in the presence of aggregate risk, expectations $\mathbb{E}_t X_{t+1}$ pick up $\frac{1}{2}\frac{\partial^2 \mathcal{X}}{\partial \epsilon_0^2}\sigma^2$ in every period, from the expected size dependence term. This is a constant perturbation to expectations, and the risky steady state reequilibrates in response to this perturbation, following (19) and (20).

This leads us to our core result:

Proposition 3 (Second-order certainty correspondence)**.** *The second-order perturbation of X_t in the sequence space is given by the quadratic MA(∞) in equation (12). The coefficients on the new terms in (13) are given by equations (15), (17), and (20), so they can be obtained from perfect-foresight impulse responses.*

Relationship to literature. Two earlier papers, Lan and Meyer-Gohde (2013b) and Lombardo and Uhlig (2018), also pursue perturbation in the sequence space, with Lan and Meyer-Gohde introducing the Volterra expansion in (12). Lombardo and Uhlig connect this to perturbation in the state space and provide a theory of pruning, which we further discuss below. Lan and Meyer-Gohde obtain expressions for the coefficients in this expansion in terms of the derivatives of the F function, paralleling the usual approach to perturbation in the state space. By contrast, our result in proposition 3 shows that we

can solve for these coefficients using perfect-foresight impulse responses, bypassing the need to explicitly obtain and manipulate derivatives of F. This is a significant improvement because obtaining these derivatives can be very complicated, especially in the case of heterogeneous-agent models.

Finally, it is worth noting that the system that we start from, (1), is different from that used in Lan and Meyer-Gohde (2013b), and much of the literature. Lan and Meyer-Gohde write their model in the form

$$\mathbb{E}_t \left[F \left(X_{t-1}, X_t, X_{t+1}, \epsilon_t \right) \right] = 0, \tag{21}$$

with the expectation operator on the outside of the F function. While equation (21) may appear to be strictly more general, dynamic programming usually enables us to rewrite the variables in X_t so that expectations only appear through the mean $\mathbb{E}_t[X_{t+1}]$, as argued in section 2.1. This reformulation is key to enabling us to obtain the $\mathcal{X}_{\sigma\sigma}$ term solely via perfect foresight, as in equation (20). Otherwise, as the appendix shows, we would have another term, reflecting the interaction of risk in X_{t+1} with the nonlinearity of F.[10]

Equivalence to second-order pruned state space. Consider perturbation in the state space rather than the sequence space. Given lemma 1, a second-order expansion of (2) around $\sigma = 0$ yields

$$X_t \approx g_X X_{t-1} + g_\epsilon \epsilon_t + \frac{1}{2} \left(X_{t-1}' g_{XX} X_{t-1} + g_{X\epsilon} X_{t-1} \epsilon_t + g_{\epsilon\epsilon} \epsilon_t^2 + g_{\sigma\sigma} \sigma^2 \right) \tag{22}$$

As Kim et al. (2008) have pointed out, simulations of X_t starting an initial X_0 using equation (22) can be explosive even when simulations of the nonlinear law of motion (2) are well behaved, due to the behavior of the quadratic term $X_{t-1}' g_{XX} X_{t-1}$. The typical solution in the literature is to "prune" the solutions (e.g. Andreasen et al. 2018). Suppose that we simulate separately $X_t^{(1)}$ and $X_t^{(2)}$ using:

$$\begin{aligned} X_t^{(1)} &= g_X X_{t-1}^{(1)} + g_\epsilon \epsilon_t \\ X_t^{(2)} &= g_X X_{t-1}^{(2)} + g_\epsilon \epsilon_t + \frac{1}{2} \left(X_{t-1}^{(1)\prime} g_{XX} X_{t-1}^{(1)} + g_{X\epsilon} X_{t-1}^{(1)} \epsilon_t + g_{\epsilon\epsilon} \epsilon_t^2 + g_{\sigma\sigma} \sigma^2 \right) \end{aligned}$$

Then the behavior of the resulting $X_t = X_t^{(1)} + X_t^{(2)}$ remains well behaved provided that the stability condition in proposition 1 is satisfied. Lombardo and Uhlig (2018)'s result

[10] Indeed, Lan and Meyer-Gohde (2013b) have precisely this additional term in their equation (36); in their notation, it appears as $f_{y^{+2}}(y_0 \otimes y_0)\mathbb{E}_t(\varepsilon_{t+1} \otimes \varepsilon_{t+1})$. Without this term, their equation (36) would parallel our derivation of $\mathcal{X}_{\sigma\sigma}$.

is that the X_t obtained from this procedure is identical to the X_t obtained in the second-order sequence-space solution (12).[11] Therefore, a benefit of sequence-space perturbation over state-space perturbation is that it automatically yields stable, "pruned" solutions.

Verification on the neoclassical growth model. Given the equivalence between the pruned state-space and sequence-space solutions, we can use a standard state-space perturbation solution of the neoclassical model as an alternative means to construct the sequence-space solution, and validate that this gives the same results as using the certainty correspondence. We have done so, finding that the two agree up to a small numerical error from differentiation.

2.4 Application to the smooth SIM model

We now apply our results to the smooth SIM model described in section 2.1. Our calibration, which is at an annual frequency, mostly follows Auclert, Rognlie and Straub (2024a), except for the addition of the smooth i.i.d. income shocks. We assume that the discrete component of idiosyncratic income follows a Markov chain in logs, $\log e_t = \rho_e \log e_{t-1} + \sigma_e \eta_t$, where η_t is i.i.d. with mean zero and variance 1, with income persistence $\rho_e = 0.91$, and $\sigma_e = 0.5\sqrt{1-\rho_e^2}$, discretized to a $n_e = 7$ point Markov chain using Rouwenhorst's method. We further assume that the distribution for the smooth i.i.d. shock ξ is the sum of a constant and a lognormal distribution with standard deviation $\sigma = 0.05$, constructed such that $\mathbb{E}[\xi] = 1$ overall.[12] We set $\underline{a} = 0$, and set the asset gridpoints on which to store the value function (which are also the knots for the cubic spline representing the distribution) to be $n_a = 200$ double-exponentially spaced points. In the deterministic steady state, we set the annual interest rate to $r = 4\%$, and pick β, the level of income Y and the aggregate amount of assets A so that aggregate consumption is $C = 1$ and the first-order aggregate marginal propensity to consume $\frac{\partial C_0}{\partial Y_0}$ is equal to 0.51 as in Auclert et al. (2024a). This procedure yields $\beta = 0.9117$, $Y = 0.9927$ and $A = 0.18$. Finally, for our aggregate shock process, we assume that aggregate income follows an i.i.d. income process $Y_t - Y = \sigma \bar{\epsilon}_t$, with annual standard deviation $\sigma = 0.1$.

We compute the terms necessary to obtain the second-order perturbation of the model as follows. First, we numerically construct the functions $C_j^{MIT}(\nu)$ mapping an innovation to income of size ν to the perfect-foresight response of consumption. This can be done by solving backward for all policies, then solving forward for the distribution, and aggre-

[11] The same observation is made by Bhandari et al. (2023).

[12] The constant ensures that the natural borrowing limit is positive.

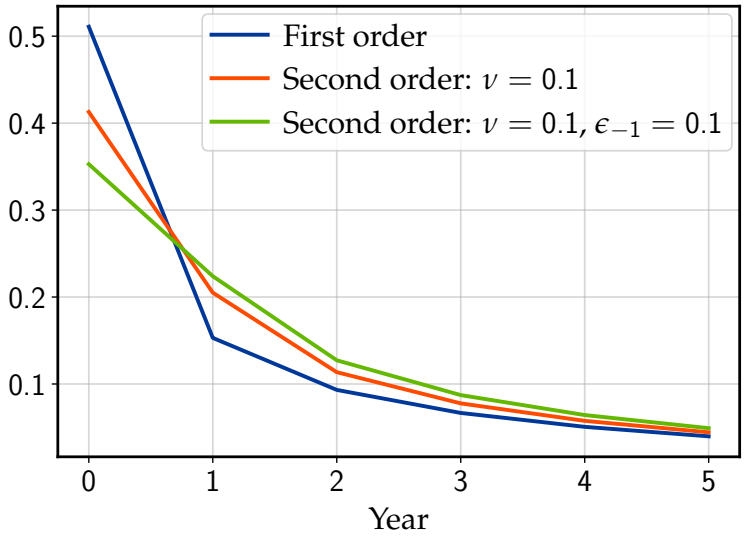


Figure 1: Intertemporal MPCs: first and second order

gating.[13] We similarly construct the sequence $C_j^{MIT,l}(\epsilon, \nu)$ that results from a sequence of two shocks. We then obtain the derivatives of these functions using numerical differentiation with central finite-difference schemes. In particular, to obtain the second derivative required for (15) we use

$$(C_j^{MIT})''(0) \simeq \frac{C_j^{MIT}(\nu) + C_j^{MIT}(-\nu) - 2C}{\nu^2}$$

while for (17) we use

$$\frac{\partial^2 C^{MIT,l}}{\partial \nu \partial \epsilon} \simeq \frac{C_j^{MIT,l}(\epsilon, \nu) - C_j^{MIT,l}(\epsilon, -\nu) - C_j^{MIT,l}(-\epsilon, \nu) + C_j^{MIT,l}(-\epsilon, -\nu)}{4\nu\epsilon}$$

Observe that to compute all the size and history terms up to lag L, this requires a total evaluation of $2 + 4L$ partial equilibrium nonlinear impulse responses, which is not very costly especially if L is not too large.[14]

We first explore the resulting size and history dependence. In figure 1 we report the first order impulse response $(C_j^{MIT})'(0)$ against the second-order impulse response normalized by size, $(C_j^{MIT}(\nu) - C)/\nu$, computed using the quadratic approximation (16) for a shock of size $\nu = \sigma$. These both have the interpretation of aggregate intertemporal marginal and average propensities to consume, respectively. The blue line reports

[13] Note that since this is a partial equilibrium model, here there is no fixed point to solve in the sequence space. We discuss how to efficiently implement perturbation in general equilibrium with heterogeneous agents in section 3.

[14] In practice, $L \simeq 20$ is enough that the history terms tend to have died out.

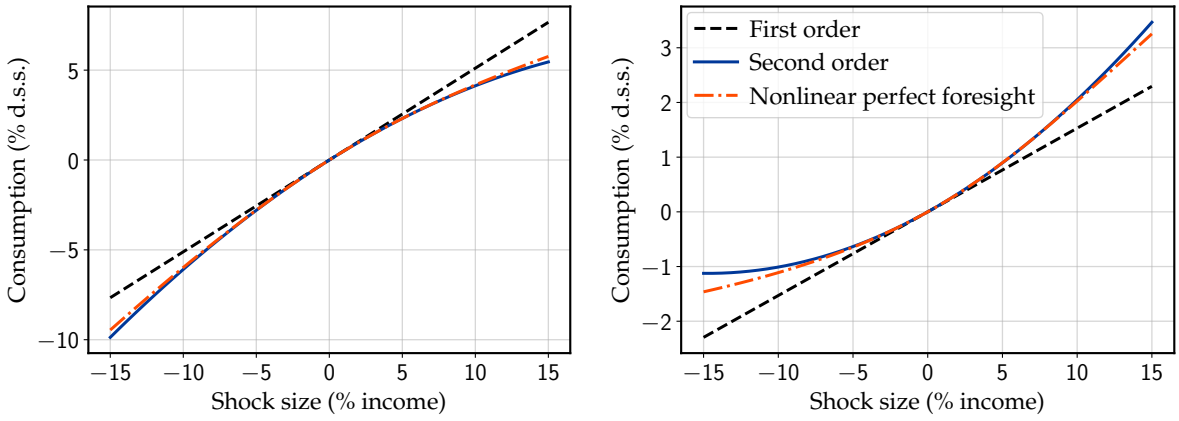
the well-known pattern from models with idiosyncratic income risk and borrowing constraints: after a large impact effect of income on consumption (targeted to 0.51 in our calibration), the consumption response remains elevated for several years after the shock. The orange line shows that the effect, once scaled by size, is attenuated on impact for a larger shock. This is due to the well-known concavity of the consumption function, manifesting in a negative $\frac{\partial^2 \mathcal{X}}{\partial \epsilon_0^2}$ term. However, an interesting feature of this problem is that, while the impact effect from a larger income shock is lower, the effect is also significantly more persistent. This is a manifestation of the intertemporal budget constraint of agents: the additional savings from period 0 must later be spent down as agents return to their buffer stock of assets. In this calibration, this effect appears immediately in year 1.

The green line shows the way in which the impulse response to $\nu = \sigma$ is modified by the presence of a shock of size $\epsilon = \sigma$ at date -1. Qualitatively, this history effect is similar to that of a larger-sized shocks, acting via the same mechanism: the past income shock is saved, boosting agent balance sheets as they come into period 0, which in turn further dampens the impact effect and magnifies the persistence of the effect on consumption of the income shock at date 0.

Figures 2 and 3 visualize these results in a different way, by considering how the level of the consumption response $C_j^{MIT}(\nu) - C$ varies with shock size on impact ($j = 0$, left panels) and at date 1 ($j = 1$, right panels). The dashed line in these figures represents the linear impulse response. The blue line visualizes the second order approximation: this has a concave shape at date 0, and a convex shape at date 1, for the reason discussed above. In figure 2, we visualize this against the fully nonlinear response $C_j^{MIT}(\nu) - C$. For the range of shocks considered here (aggregate income shocks of up to 15%), the second order approximation does an excellent job at capturing the nonlinear solution, although the fit is slightly worse for the date 1 impulse response. Finally, figure 3 shows the additional effect from the presence of a positive shock of size 0.1 at date -1: this has the effect of rotating the impulse response linearly around 0 clockwise at date 0 (a negative interaction term as balance sheets are stronger immediately) and counterclockwise at date 1 (a positive interaction term due to the intertemporal budget constraint.)

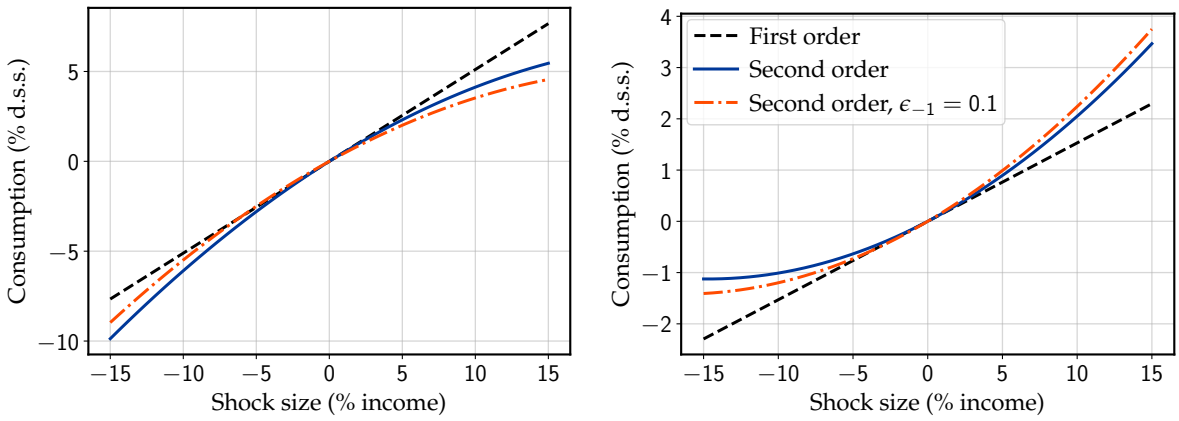
We next turn to the effect of income uncertainty σ on aggregate asset accumulation in the risky steady state[15] To evaluate this, we implement equation (20). We first obtain all the terms in $\frac{\partial^2 \mathcal{X}}{\partial \epsilon_0^2}$ that are necessary to calculate the perturbed expectations steady state in (19). Inspecting the F equation for the SIM model in (8), we note that the only element of X_t that enters with an $\mathbb{E}_t[X_{t+1}]$ is the marginal value function v_t. We therefore obtain the

[15]Note that this is closely related to the effect on aggregate consumption, since the country budget constraint implies $C = Y + rA$.



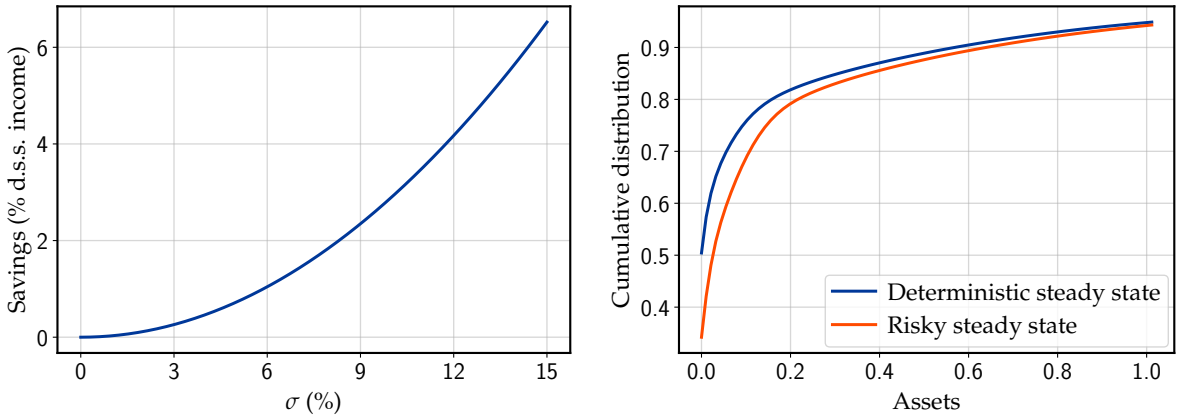
(a) Response at $t = 0$ (b) Response at $t = 1$

Figure 2: Response of consumption to income shocks: second order vs perfect foresight



(a) Response at $t = 0$ (b) Response at $t = 1$

Figure 3: Response of consumption to two consecutive income shocks: second order



(a) Aggregate assets

(b) Wealth distribution

Figure 4: The risky steady state: aggregate savings and aggregate income uncertainty

impact size-dependence in that marginal value function at each of the $n_a n_e$ points on the grid using numerical second differentiation. We then replace the marginal value function by its steady state value plus δ times this derivative, and solve for the pseudo steady state given δ using the usual steady state routine. We can then use equation (20) to recover $\mathcal{X}_{\sigma\sigma}$.

Figure 4 visualizes the outcome of this procedure. The left panel plots the effect of aggregate risk on risky steady-state aggregate assets as a function of σ, in other words the term $\frac{1}{2}\mathcal{X}_{\sigma\sigma}\sigma^2$ corresponding to total assets. The right panel plots the cumulative distribution function of assets, both at the deterministic steady state and at the risky steady state with $\sigma = 0.1$. Perceived aggregate risk makes agents save more at every point in the distribution, resulting in a rightward shift of the cdf and a significant increase in aggregate precautionary savings. Notably, this lowers the mass of agents at the borrowing constraint by a significant amount. In turn, this implies that MPCs are significantly lower in the risky steady state, compared to the deterministic steady state. This has important implications for the third order perturbation.

2.5 Third-order perturbation

Assuming third differentiability, and using lemma 1 again to observe that all mixed derivatives $\mathcal{X}_{\sigma jk}$ with respect to σ, ϵ_{t-j} and ϵ_{t-k} are zero, a third-order Taylor approximation (Volterra expansion) of X_t can now be written

$$X_t \approx X_t^{(2)} + \frac{1}{6}\sum_{j=0}^{\infty}\frac{\partial^3 \mathcal{X}}{\partial \epsilon_{-j}\partial\sigma^2}\epsilon_{t-j}\sigma^2 + \frac{1}{6}\sum_{j=0}^{\infty}\sum_{k=0}^{\infty}\sum_{l=0}^{\infty}\frac{\partial^3 \mathcal{X}}{\partial \epsilon_{-j}\partial\epsilon_{-k}\partial\epsilon_{-l}}\epsilon_{t-j}\epsilon_{t-k}\epsilon_{t-l} \tag{23}$$

where $X_t^{(2)}$ is the second-order Volterra expansion in equation (12) and, again, all derivatives are evaluated at $(\sigma, \epsilon_t, \epsilon_{t-1} \ldots) = (0, 0, 0, \ldots)$. Hence, to third order, up to constants, X_t is a cubic infinite moving-average process in ϵ_t.

The new term on the right in (23) corresponds to all triple interactions of shocks, and can further be rewritten as

$$\begin{aligned}\frac{1}{6}\sum_{j=0}^{\infty}\frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}^3}\epsilon_{t-j}^3 \quad &+ \quad \frac{1}{2}\sum_{j=0}^{\infty}\sum_{l=1}^{\infty}\frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}^2\partial\epsilon_{-(j+l)}}\epsilon_{t-j}^2\epsilon_{t-(j+l)} + \frac{1}{2}\sum_{j=0}^{\infty}\sum_{l=1}^{\infty}\frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}\partial\epsilon_{-(j+l)}^2}\epsilon_{t-j}\epsilon_{t-(j+l)}^2 \\ &+ \quad \sum_{j=0}^{\infty}\sum_{l=1}^{\infty}\sum_{l'=1}^{\infty}\frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}\partial\epsilon_{-(j+l)}\partial\epsilon_{-(j+l+l')}}\epsilon_{t-j}\epsilon_{t-(j+l)}\epsilon_{t-(j+l+l')}\end{aligned}$$

These correspond, respectively, to cubic size dependence, the interaction of quadratic size dependence with history, the quadratic size history effect, and the history effect from the interaction of two past shocks. Just like with second-order certainty correspondence, these terms are all straightforwardly related to the corresponding MIT shock solution. Indeed, we have:

$$\frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}^3} = (X_j^{MIT})'''(0) \tag{24}$$

and for $l \geq 1$ we have

$$\frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}^2\partial\epsilon_{-(j+l)}} = \frac{\partial^3 X_j^{MIT,l}}{\partial^2\nu\partial\epsilon} \qquad \frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}\partial^2\epsilon_{-(j+l)}} = \frac{\partial^3 X_j^{MIT,l}}{\partial\nu\partial^2\epsilon} \tag{25}$$

while, for $l \geq 1$ and $l' \geq 1$ we can define the sequence

$$X_j^{MIT,l,l'}(\epsilon', \epsilon, \nu) \equiv \mathcal{X}(0; 0, \ldots 0, \underbrace{\epsilon'}_{j-(l+l')}, 0, \ldots 0, \underbrace{\epsilon}_{j-l}, 0, \ldots, 0, \underbrace{\nu}_{j}, 0 \ldots)$$

and then we have

$$\frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}\partial\epsilon_{-(j+l)}\partial\epsilon_{-(j+l+l')}} = \frac{\partial^3 X_j^{MIT,l,l'}}{\partial\nu\partial\epsilon\partial\epsilon'} \tag{26}$$

Equations (24), (25) and (26) show that all triple shock interaction terms in the third order Volterra expansion can be obtained via perfect-foresight impulse responses.

The additional new term in the third order sequence space perturbation (23) is the term from the $\epsilon_{t-j}\sigma^2$ interactions. This term quantifies how the presence of aggregate risk affects the impulse response to shocks, and how the occurrence of past shocks affects

the precautionary motive (time-varying precautionary savings). We separate these effects into two:

$$\frac{\partial^3 \mathcal{X}}{\partial \epsilon_{-j} \partial \sigma^2} = \left(\frac{\partial^3 \mathcal{X}}{\partial \epsilon_{-j} \partial \sigma^2} \right)^{rss} + \left(\frac{\partial^3 \mathcal{X}}{\partial \epsilon_{-j} \partial \sigma^2} \right)^{tvps} \tag{27}$$

where *rss* stands for the risky steady state, and *tvps* for time-varying precautionary savings.

Given our routines to calculate the risky steady state, the *rss* term in (27) is simple to calculate using a perfect-foresight impulse response: obtain the first-order impulse response to a shock starting from the risky steady state distribution rather than the deterministic steady state, with all expectations modified using expected size dependence as in the risky steady state calculation. The *tvps* term, on the other hand, requires us to modify our concept of perturbed expectations to accommodate time variation in these perturbations. Concretely, we consider the deterministic *perturbed-expectations transition*, which for any δ solves the equation

$$F\left(X_{t-1}^{pt}(\delta), X_t^{pt}(\delta), X_{t+1}^{pt}(\delta) + \frac{\partial^3 \mathcal{X}}{\partial \epsilon_{-t} \partial \epsilon_0^2} \delta, 0 \right) = 0,$$

where the $\frac{\partial^3 \mathcal{X}}{\partial \epsilon_{-t} \partial \epsilon_0^2}$ term is calculated using (25). Then we have that:

$$\left(\frac{\partial^3 \mathcal{X}}{\partial \epsilon_{-j} \partial \sigma^2} \right)^{tvps} = (X_j^{pt})'(0) \tag{28}$$

This completes our argument that we can obtain all the terms in the third-order perturbation using perfect-foresight sequences.

Proposition 4 (Third-order certainty correspondence)**.** *The third-order perturbation of X_t in the sequence space is given by the cubic MA(∞) in equation (23). The coefficients on the new terms are given by equations (24), (25), (26), (27), and (28), so they can be obtained from perfect-foresight impulse responses.*

Again, the argument in Lombardo and Uhlig (2018) can be used to show that the third order perturbation we obtain is equivalent to the third-order pruned state space solution. As in the second-order case, we use this pruned solution for the neoclassical growth model to validate our certainty correspondence results.

Finally, analogously to (14), we can derive the third order generalized impulse re-

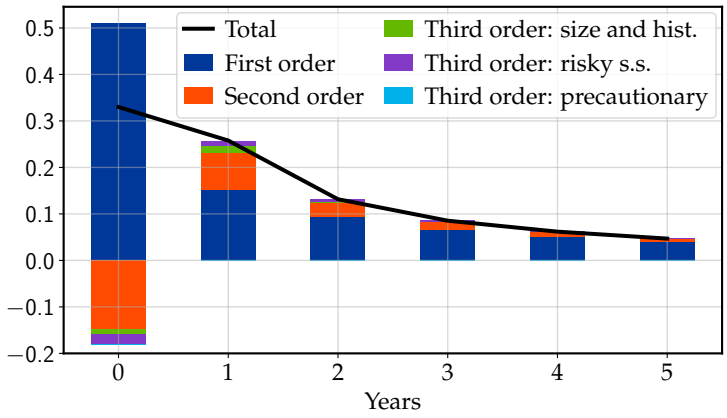


Figure 5: Intertemporal MPCs up to third order for a shock of size $\nu = 0.15$ and $\sigma = 0.1$

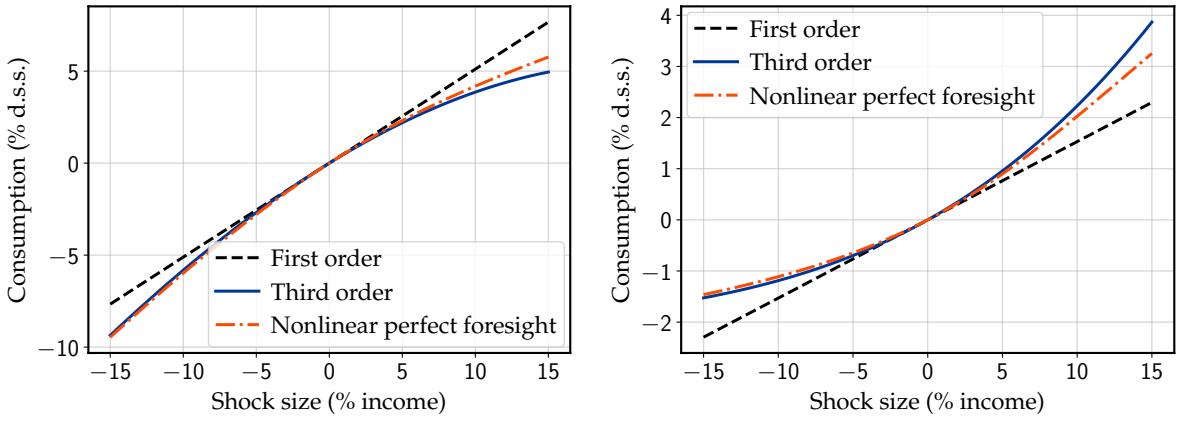
sponse. This reads:

$$\begin{aligned} IRF_j^{(3)}(\nu,\sigma) &= IRF_j^{(2)}(\nu,\sigma) + \frac{1}{6}\frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}\partial\sigma^2}\sigma^2\nu + \frac{1}{6}\frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}^3}\nu^3 + \frac{1}{2}\left(\sum_{l=1}^{\infty}\frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}^2\partial\epsilon_{-(j+l)}}\epsilon_{t-l}\right)(\nu^2-\sigma^2) \\ &+ \sum_{l=1}^{\infty}\frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}\partial\epsilon_{-(j+l)}^2}\epsilon_{t-l}^2\nu + \sum_{l=1}^{\infty}\sum_{l'=1}^{\infty}\frac{\partial^3\mathcal{X}}{\partial\epsilon_{-j}\partial\epsilon_{-(j+l)}\partial\epsilon_{-(j+l+l')}}\epsilon_{t-l}\epsilon_{t-(l+l')}\nu \end{aligned} \tag{29}$$

where $IRF_j^{(2)}(\nu,\sigma)$ denotes the second-order impulse response in (14).

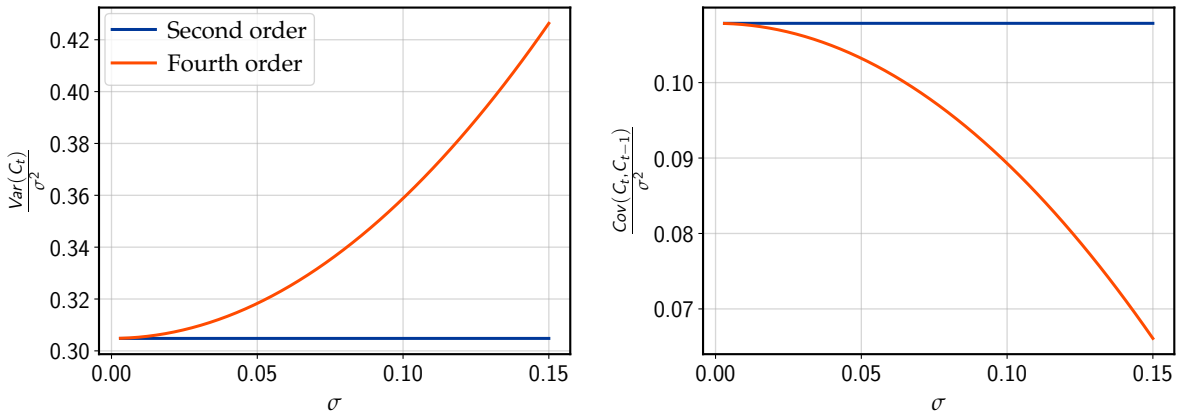
Application to the smooth SIM model. We continue to illustrate these results using the smooth SIM model. Figure 5 shows non-constant terms from the impulse response of consumption to income in (29) for a shock of size $\nu = 0.15$ and standard deviation of shocks $\sigma = 0.10$, with zero past innovations $\epsilon_{t-l} = 0$. The first- and second-order term are the same as in figure 1: the second order term comes in to lower iMPCs at date 0 due to the concavity of the consumption function, and raise them at date 1 and later, an overshoot due to the intertemporal budget constraint. The third order terms, in this application, are collectively relatively small, and the precautionary savings effect with i.i.d. aggregate shocks is actually zero. We will find more significant differences when considering a general equilibrium version of this model in section 3.2.

Figure 6 further zooms in on the third order MIT shock impulse response, comparing it to the nonlinear perfect-foresight solution. Relative to figure 2, we see that the third-order approximation provides only a slight improvement towards the nonlinear perfect-foresight solution, consistent with figure 5.



(a) Response at $t = 0$ (b) Response at $t = 1$

Figure 6: Response of consumption to income shocks: third order vs perfect foresight



(a) Consumption variance (b) Consumption first autocovariance

Figure 7: Ergodic moments: higher order approximations

2.6 Model moments

The ergodic moments of the model are an object of broad interest. A typical approach in the literature is to use the moments from the perturbation approximation–for instance, second moments calculated from the $MA(\infty)$ implied by the first-order approximation are routinely used to estimate models via method of moments or likelihood methods (e.g. Auclert et al. 2021). Here, we ask how these calculations relate to the perturbation of the primitive ergodic model moments as a function of σ, and show that we can calculate fourth-order accurate moments using the third order approximation derived above.

In particular, for vector-valued X_t, we are interested in the autocovariance matrix

$$\Sigma_h(\sigma) = \text{Cov}[\mathcal{X}(\sigma, \sigma\bar{\epsilon}_t, \sigma\bar{\epsilon}_{t-1}, \ldots), \mathcal{X}(\sigma, \sigma\bar{\epsilon}_{t-h}, \sigma\bar{\epsilon}_{t-h-1}, \ldots)] \tag{30}$$

We then have the following proposition, which further assumes that $\bar{\epsilon}$ has the same moments, up to at least fourth order, as a standard normal distribution.[16]

Proposition 5 (Ergodic moments). *A fifth-order accurate solution to the autocovariance matrix in (30) is given by*

$$\Sigma_h(\sigma) = \sigma^2\Sigma_{h,2} + \sigma^4\Sigma_{h,4} + O(\sigma^6)$$

where $\Sigma_{h,2} = \sum_{i=0}^{\infty} \mathcal{X}_{h+i}\mathcal{X}_i^\top$ *is the autocovariance matrix obtained from the first-order perturbation solution, and*

$$\Sigma_{h,4} = \frac{1}{2}\sum_{i=0}^{\infty}\sum_{j=0}^{\infty} \mathcal{X}_{h+i,h+j}\mathcal{X}_{ij}^\top + \frac{1}{6}\sum_{i=0}^{\infty}\left(\mathcal{X}_{h+i}\mathcal{X}_{\sigma\sigma i}^\top + \mathcal{X}_{\sigma\sigma,h+i}\mathcal{X}_i^\top\right) + \frac{1}{2}\sum_{i=0}^{\infty}\sum_{j=0}^{\infty}\left(\mathcal{X}_{h+i}\mathcal{X}_{ijj}^\top + \mathcal{X}_{h+i,jj}\mathcal{X}_i^\top\right)$$

Proposition 5 shows, first, that the second-moments calculated from the first-order solution are accurate up to third order in σ. This provides a justification for the typical procedure in the literature. The proposition further offers a correction to the moments that is accurate up to fifth order. This correction involves the second and third derivatives of $\mathcal{X}$, which we have shown how to obtain in sections 2.3 and 2.5 respectively. Notably, however, it does not require *all terms* from the second and third-order solutions: for instance, the triple interactions $\mathcal{X}_{ijk}$, which are the most numerous and therefore computationally intensive part of the third-order solution, do not appear.[17]

Figure 7 illustrates these results by visualizing the ergodic variance of consumption $\text{Var}(C_t)$ and the first order autocovariance of consumption $\text{Cov}(C_t, C_{t-1})$, normalized by income shock variance σ^2, for different values of σ. The horizontal blue line shows the second-order approximation to variance–the usual term obtained in the literature, including from typical Dynare output–which is constant once normalized by σ^2. The orange line shows how this term is modified in a fifth-order accurate approximation of the moments. For extremely large aggregate shocks ($\sigma \approx 0.15$), the scaled ergodic variance is around 40% larger, and the ergodic autocovariance around 30% lower. This suggests that this correction from the first-order solution might make a significant difference in certain applications.

[16] This additional assumption is only required for the fourth-order term $\sigma^4\Sigma_{h,4}$. Our original assumption of symmetric $\bar{\epsilon}$ was sufficient for the $\sigma^2\Sigma_{h,2}$ to be accurate to third order.

[17] This is in contrast to the more common approach in, for instance, Andreasen et al. (2018), where the entire second- or third-order pruned perturbation solution is used when approximating autocovariances. Our derivation shows that other components of the solution are not needed for fifth-order accuracy—and while they might be needed to higher order, an even higher-order approximation to moments would involve terms from the fourth- and fifth-order solutions.

3 Implementation in the sequence space

The certainty correspondence result we derived in section 2 is a general result that relates the Volterra expansion of the nonlinear sequence-space solution up to third order to nonlinear perfect-foresight impulse responses computed up to the same order. In practical applications, however, the nonlinear perfect-foresight impulse response may be slow to obtain. This is in particular the case in general equilibrium heterogeneous-agent models, which require solving for certain fixed points in sequences (for instance solving for endogenous paths of prices or quantities.) In this section, we describe a general approach to speed up the implementation of the formulas from section 2, bypassing nonlinear fixed points. We then provide two applications to general equilibrium heterogeneous agent models.

3.1 General equilibrium in sequence space

When solving for nonlinear perfect-foresight impulse responses, a common approach, applicable in many general equilibrium models, is to directly solve for the path of aggregate "unknown" variables $\mathbf{U} = (\mathbf{U}_0, \mathbf{U}_1, \ldots)$ given exogenous "shocks" $\mathbf{Z} = (\mathbf{Z}_0, \mathbf{Z}_1, \ldots)$, reducing the system to a certain function

$$\mathbf{H}(\mathbf{U}(\sigma), \mathbf{Z}(\sigma)) = 0 \tag{31}$$

where σ indexes the standard deviation of primitive shocks (see, for instance, Auclert et al. (2021), equation 29). This formulation differs from the primitive formulation in (1) in several important ways. First, $\mathbf{Z}$ and $\mathbf{U}$ are only perfect foresight impulse responses. Second, they are more compact than the full perfect foresight solution tracing innovations to ϵ to the full set of endogenous variables X. For instance, with an AR(1) driving stochastic process $\tilde{Z}_t - Z = \rho(\tilde{Z}_{t-1} - Z) + \sigma\bar{\epsilon}_t$, then it is sufficient to write the shock $\mathbf{Z}_t(\sigma) = Z + \sigma\rho^t$. Similarly, $\mathbf{U}$ only contains the aggregate sequences that require a meaningful fixed point, rather than the full set of value functions and distribution points underlying the F system in (1).

For the case of systems like (31), we now provide a way to obtain the MIT shock perturbation solutions needed for equations (11), (15), (17) and (24)-(26) without the need to solve for a general equilibrium fixed point.

To see how this works, first differentiate (31) with respect to σ to obtain

$$\mathbf{H}_U(\mathbf{U}(\sigma), \mathbf{Z}(\sigma))\mathbf{U}'(\sigma) + \mathbf{H}_Z(\mathbf{U}(\sigma), \mathbf{Z}(\sigma))\mathbf{Z}'(\sigma) = 0 \tag{32}$$

Evaluating (32) at $\sigma = 0$, letting $\mathbf{H}_U$ and $\mathbf{H}_Z$ denote derivatives around the steady state, and using the notational conventions gives $\mathbf{H}_U\mathbf{U}'(0) + \mathbf{H}_Z\mathbf{Z}'(0) = 0$. This can be solved to obtain the usual first-order solution

$$\mathbf{U}'(0) = -\mathbf{H}_U^{-1}\mathbf{H}_Z\mathbf{Z}'(0)$$

Calculating this solution only requires the sequence-space Jacobians $\mathbf{H}_U$ and $\mathbf{H}_Z$, which can be obtained rapidly via the methods in Auclert et al. (2021). This delivers the first-order MIT shock solution in (11).

Next, assume that the exogenous $\mathbf{Z}$ is defined such that $\mathbf{Z}''(0) = 0$, such as in our AR(1) example. Differentiating (32) again with respect to σ, evaluating at $\sigma = 0$, now gives

$$\mathbf{H}_U\mathbf{U}''(0) + \mathbf{H}_{UU}(\mathbf{U}'(0))^2 + 2\mathbf{H}_{UZ}\mathbf{U}'(0)\mathbf{Z}'(0) + \mathbf{H}_{ZZ}(\mathbf{Z}'(0))^2 = 0 \tag{33}$$

where, e.g., $\mathbf{H}_{UU}d\mathbf{U}^2$ denotes the application of the vector-valued Hessian $\mathbf{H}_{UU}$ to $d\mathbf{U}$ on both input dimensions. (33) can be solved to obtain the second-order term in a solution for $\mathbf{U}$:

$$\mathbf{U}''(0) = -\mathbf{H}_U^{-1}\left(\mathbf{H}_{UU}(\mathbf{U}'(0))^2 + 2\mathbf{H}_{UZ}\mathbf{U}'(0)\mathbf{Z}'(0) + \mathbf{H}_{ZZ}(\mathbf{Z}'(0))^2\right) \tag{34}$$

This equation for $\mathbf{U}''(0)$ has a simple intuition. First, calculate the first-order impulse response $\mathbf{U}'(0)$, which satisfies $\mathbf{H}$ to first order since $\mathbf{H}_U\mathbf{U}'(0) + \mathbf{H}_Z\mathbf{Z}'(0) = 0$. Now, evaluate $\mathbf{H}$ to *second order* on these first-order impulse responses, obtaining the term in parentheses in (34). This will generally not be zero, implying a second-order error in $\mathbf{H}$ that must be offset by some $\mathbf{U}''(0)$, which we obtain by applying $-\mathbf{H}_U^{-1}$.

Equation (34) delivers the terms necessary to the second-order sequence space solution in (15) and (17). Observe, critically, that this bypasses the need for solving a nonlinear fixed point, since it requires a single nonlinear evaluation of equilibrium conditions, and then applying the inverse Jacobian $\mathbf{H}_U^{-1}$ to re-equilibrate the system.

Continuing this argument, the appendix shows that this "one-shot principle" can be applied to obtain the third order $\mathbf{U}'''(0)$ necessary for equations (24)-(26) as well: evaluate the $\mathbf{H}$ function on the second-order solution to find the third-order error, and apply $\mathbf{H}_U^{-1}$ to find $\mathbf{U}'''(0)$.

3.2 Application 1: state-dependent fiscal multipliers

Our first general equilibrium application is a New Keynesian version of the smooth SIM model from section 2, along the lines of the "intertemporal Keynesian cross" model in

Auclert et al. (2024a). Households now solve the Bellman equation

$$V_t(a,e) = \mathbb{E}_\xi \left[\max_{c,a'} u(c) + \beta \mathbb{E}_t[V_{t+1}(a',e')|e] \right]$$

$$\text{s.t.} \qquad a' + c = (1+r)a + (Y_t - T_t)\xi e; \qquad a' \geq \underline{a}$$

which is identical to (7), except that aggregate household income is now $Y_t - T_t$, equal to production net of taxes. Production is from labor, $Y_t = N_t$, nominal prices are flexible, nominal wages are sticky, and monetary policy maintains the ex-ante real interest rate equal to r at all times. The government's budget constraint is

$$B_t + T_t = (1+r)B_{t-1}$$

The level of bonds B_t is stochastic, and follows an AR(1) process $B_t - B = \rho(B_{t-1} - B) + \sigma\bar{\epsilon}_t$. In equilibrium, the goods market must clear, so $\int c_t dD_t = Y_t$, and the asset market clears, $\int a_t dD_t = B_t$

In this model, positive shocks to B_t enable the government to temporarily lower taxes T_t; since households are not Ricardian, this in turn affects consumption, which in equilibrium affects aggregate income Y_t and therefore consumption C_t. Letting $H_t(Y_s, B_s) = A_t(Y_s, B_s) - B_t$, the model can be solved under perfect foresight by finding the sequence $\mathbf{Y}$ that solves $\mathbf{H}(\mathbf{Y}, \mathbf{B}) = 0$ given a sequence for bonds $\mathbf{B}$.

Risky steady state. We first calculate the risky steady state of this model to the second-order. Here, given that r is held constant at the deterministic steady state by monetary policy, the effect of aggregate risk is to reduce the average level of income. Figure 8 computes the magnitude of this effect, which is substantial. A more involved version of this model would have monetary policy react to the deflation created here by lowering the real interest rate in the risky steady state.

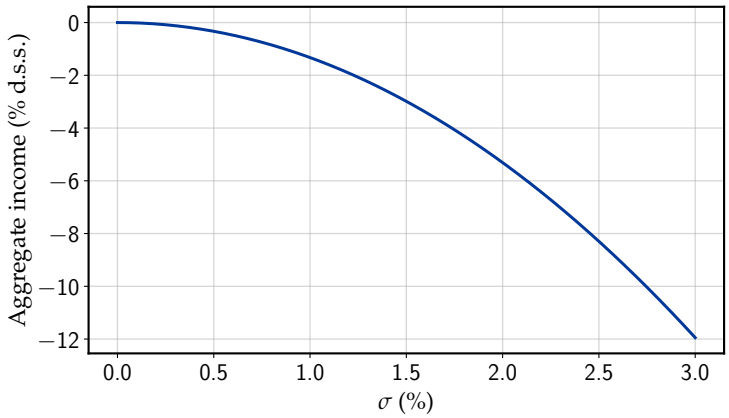


Figure 8: Disposable income in the risky steady state

Generalized impulse responses. We next turn to the generalized impulse responses of output to bond shocks, that is, the dynamic transfer multiplier. Figures 9 and 10 show that the basic nonlinear behavior of this model in partial equilibrium translates over exactly to general equilibrium: larger transfers boost output less on impact, but more later on. The third order terms remain relatively small, though the time-varying precautionary savings term now starts to play more of a role.

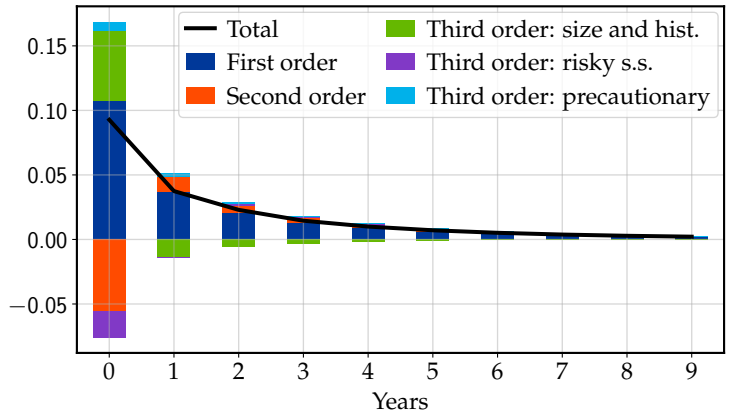
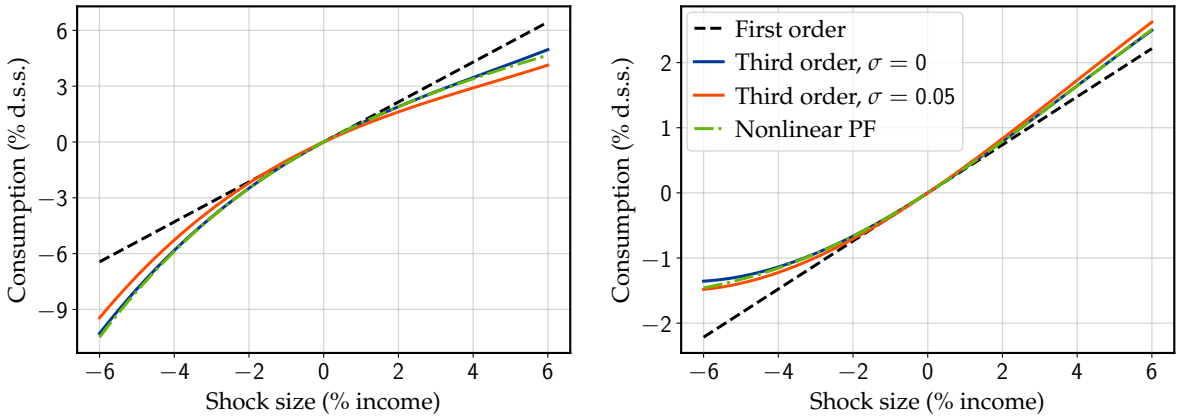


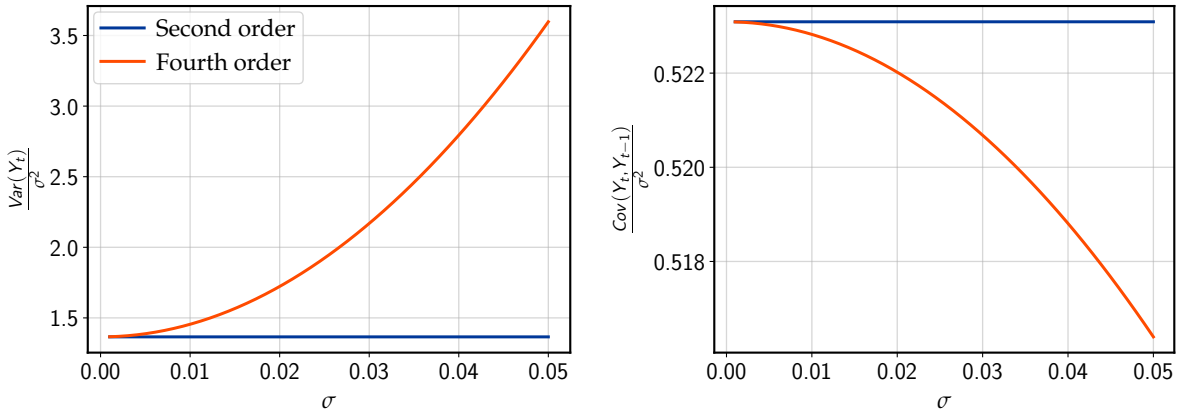
Figure 9: Decomposition of the GIRF for $\sigma = 0.05$ and $\nu = 0.1$



(a) Response at $t = 0$ (b) Response at $t = 1$

Figure 10: Response of output to a fiscal shock

Moments. Next, we evaluate how the second moments of output are affected by the computation of a fifth-order accurate solution. Just like in the partial equilibrium case, figure 11 finds that the fourth-order correction increases output (and consumption) variance and reduces output (and consumption) autocovariance.



(a) Output variance (b) Output first autocovariance

Figure 11: Higher order approximation of moments in the IKC model

Nonlinear multipliers. Finally, we study how much history dependence there is in the transfer multiplier dY_t/dB_t. We simulate the model twice, once using the second-order Volterra expansion and once using the third-order Volterra expansion, and calculate at any point in the simulation the first-order responsiveness of output to an innovation to B_t using our formulas for (12) and (23). Figure 12 finds that the transfer multiplier is

about twice as large when a sequence of negative shocks has pushed output 10% below the deterministic steady state, compared to when it is at the risky steady state.

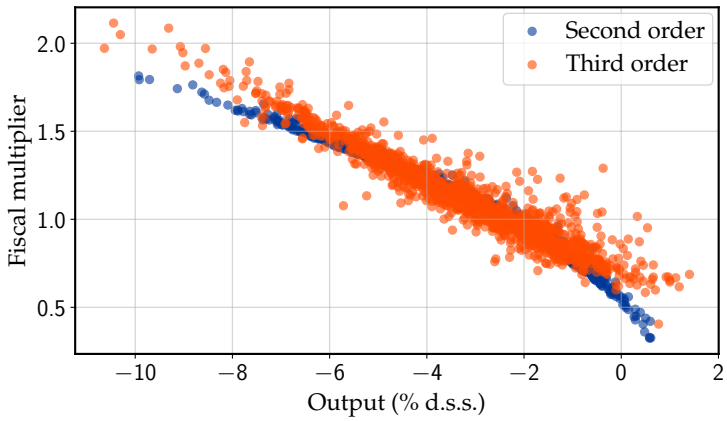


Figure 12: Time-varying fiscal multipliers for $\sigma = 0.015$

3.3 Application 2: nonlinear Phillips curves

Our second general application is a menu cost model with idiosyncratic shocks in the spirit of Golosov and Lucas (2007) and Alvarez, Lippi and Souganidis (2023). The model features a random menu cost, trend inflation, and strategic complementarity, as in Auclert, Rigato, Rognlie and Straub (2024b) and Alvarez et al. (2023). Firms have a quadratic objective, they try to keep their price gap x as close as possible to aggregate log nominal marginal cost MC_t. They can reset their price either by paying a fixed menu cost ζ, or by waiting for a free opportunity to adjust their cost, which arrives with Calvo probability λ. The only state is the price gap x, and their Bellman equation at date t is given by:

$$V_t(x) = \min \frac{1}{2}(x - MC_t)^2 + \beta(1-\lambda)\mathbb{E}_\epsilon[\min\{V_{t+1}(x+\epsilon), \zeta + \min_{x^*} V_{t+1}(x^*)\}] + \beta\lambda \min_{x^*} V_{t+1}(x^*)$$

We assume linear aggregation, so that the log price index is $P_t = \int x dD_t(x)$, where D_t is the distribution of gaps. Finally, we model strategic complementarity as roundabout production, with firms producing from an aggregate of labor and an input bundle made up of production from other firms; this implies that log nominal marginal cost is

$$MC_t = \chi W_t + (1-\chi)P_t$$

The log nominal wage trends at rate μ, $W_t = \mu t + \tilde{W}_t$, where $\tilde{W}_t = \rho\tilde{W}_{t-1} + \sigma\bar{\epsilon}_t$ follows an AR(1) process and is the source of aggregate risk in the model. We calibrate the model at a quarterly frequency so that, in the deterministic steady state, the price adjustment

frequency is freq $= 0.25$, the average size of price changes is $\mathbb{E}[\Delta p] = 0.1$, the share of free adjustments in total adjustments if 0.75, strategic complementarity is $1 - \chi = 0.3$ and trend inflation is $\mu = 5\%$ annually. General equilibrium requires solving for the nonlinear fixed point $MC_t = \chi W_t + (1 - \chi) P_t(\{MC_s\})$, which is our **H** function. Further details on the methods used to solve the model are provided in Auclert et al. (2024b)

Risky steady state. We first consider effect on the risky steady state aggregate price level (detrended by wage inflation). Figure 13 shows that cost uncertainty reduces the risky steady-state price level, an effect due to a widening of the inaction region because of the presence of cost uncertainty, which in turn erodes average real prices.

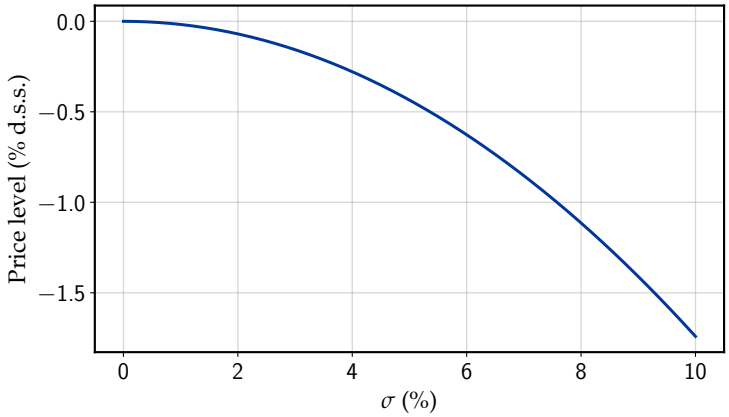
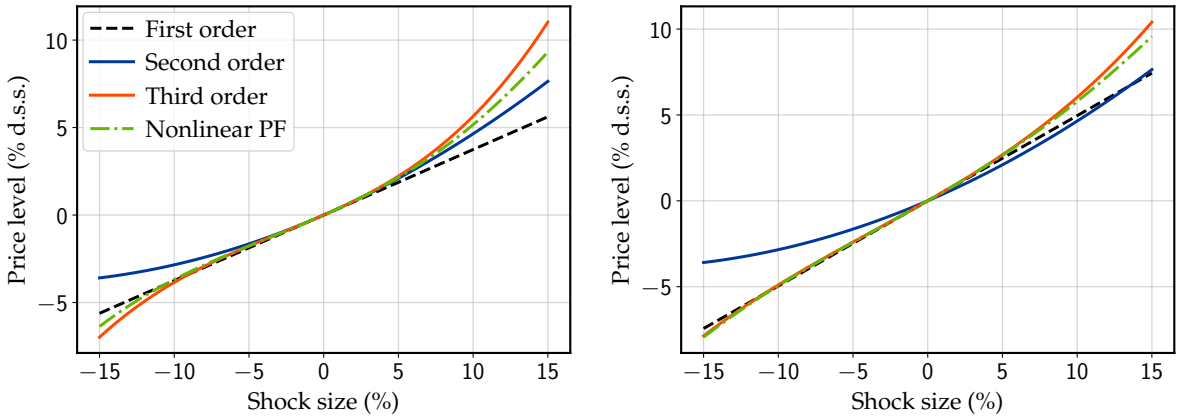


Figure 13: Price level in the risky steady state

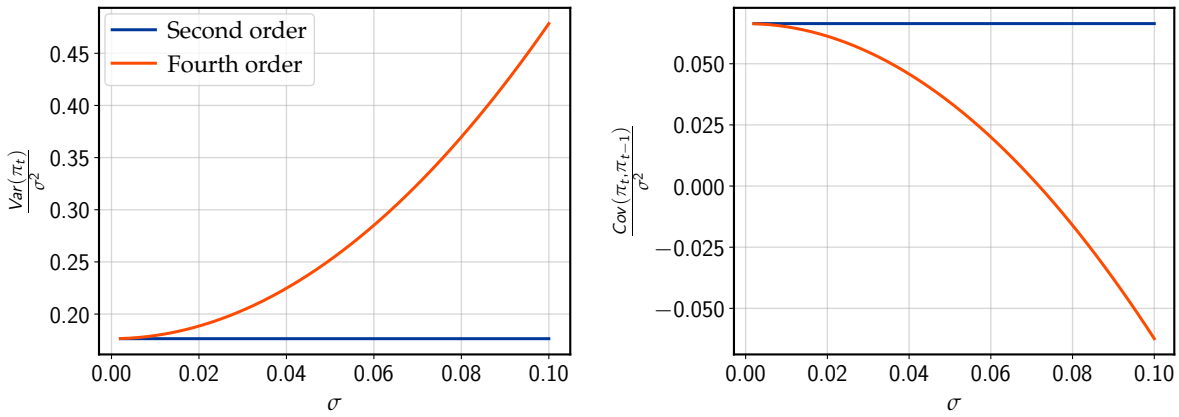
Impulse responses. Next, we turn to the impulse response of the price level to cost shocks of different sizes. As is well known of menu cost models, cost shocks tend to have nonlinear effects on prices, as larger shocks in either direction trigger more repricing, this is a persistent effect so is visible in both the immediate and the delayed response of prices (see figure 14).



(a) Response at $t = 0$ (b) Response at $t = 1$

Figure 14: Response of the price level to marginal cost shocks

Moments. We then turn to two ergodic moments of the model: the variance of inflation and the autocovariance of inflation. Relative to the typical second-order accurate moment, with large cost shocks a fourth-order accurate solution doubles the variance of inflation, and it can flip the sign of the autocovariance of inflation–which starts from a positive but low level, as is typical of menu cost models.



(a) Inflation variance (b) Inflation first autocovariance

Figure 15: Higher order approximation of moments in the menu cost model

State-dependent responsiveness of inflation. Finally, we evaluate the state dependence in the responsiveness of inflation to marginal cost shocks, $d\pi_t/d\epsilon_t$, conducting the same types of second and third order simulations as in our fiscal multiplier application. As pointed out by other studies on nonlinear Phillips curves (eg Blanco et al. (2024)), at

higher rates of current inflation, the responsiveness of inflation steepens; this is particularly pronounced in the third order solution since then the relationship between average inflation and inflation responsiveness becomes convex.

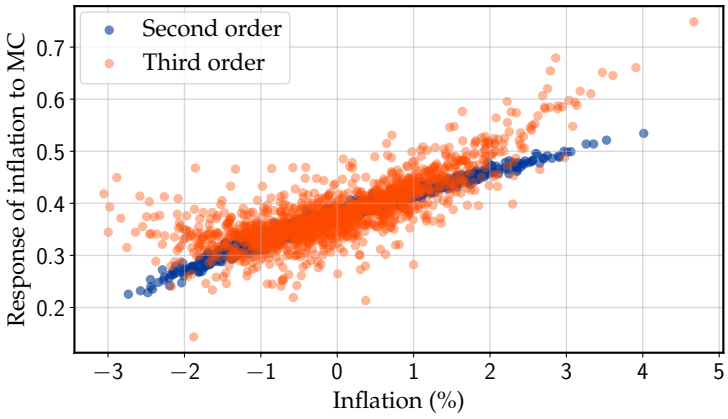


Figure 16: Time-varying response of inflation to marginal costs for $\sigma = 0.025$

4 Conclusion

This paper shows how to obtain higher-order perturbation solutions to heterogeneous-agent models using information derived from well-known routines to solve the model under perfect foresight. This allows us to solve a HANK model and a menu cost model to third order in a matter of minutes on a laptop computer. An open question is the extent to which these solutions differ from those obtained by global solution methods, which represent another active area of research.

References

Alvarez, Fernando, Francesco Lippi, and Panagiotis Souganidis, "Price setting with strategic complementarities as a mean field game," *Econometrica*, 2023, *91* (6), 2005–2039.

Andreasen, Martin M, Jesús Fernández-Villaverde, and Juan F Rubio-Ramírez, "The Pruned State-Space System for Non-Linear DSGE Models: Theory and Empirical Applications," *Review of Economic Studies*, 2018, *85* (1), 1–49.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 2021, *89* (5), 2375–2408.

_ , Matthew Rognlie, and Ludwig Straub, "The trickling up of excess savings," in "AEA Papers and Proceedings," Vol. 113 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2023, pp. 70–75.

_ , _ , and _ , "The Intertemporal Keynesian Cross," *Journal of Political Economy*, 2024, *132* (12), 4068–4121.

_ , Rodolfo Rigato, Matthew Rognlie, and Ludwig Straub, "New pricing models, same old Phillips curves?," *The Quarterly Journal of Economics*, 2024, *139* (1), 121–186.

Azinovic-Yang, Marlon and Jan Žemlička, "Deep Learning in the Sequence Space," *arXiv preprint arXiv:2509.13623*, 2025.

Bayer, Christian, Ralph Luetticke, Maximilian Weiss, and Yannik Winkelmann, "An endogenous gridpoint method for distributional dynamics," *Journal of Monetary Economics*, 2026, p. 103895.

Bhandari, Anmol, Thomas Bourany, David Evans, and Mikhail Golosov, "A Perturbational Approach for Approximating Heterogeneous-Agent Models," *Working Paper*, 2023.

Bilal, Adrien, "Solving Heterogeneous Agent Models with the Master Equation," *Working Paper*, 2023.

Blanco, Andres, Corina Boar, Callum Jones, and Virgiliu Midrigan, *Non-linear inflation dynamics in menu cost economies* number w32094, National Bureau of Economic Research, 2024.

Boppart, Timo, Per Krusell, and Kurt Mitman, "Exploiting MIT Shocks in Heterogeneous-Agent Economies: The Impulse Response as a Numerical Derivative," *Journal of Economic Dynamics and Control*, 2018, *89*, 68–92.

Coeurdacier, Nicolas, Hélène Rey, and Pablo Winant, "The Risky Steady State," *American Economic Review*, 2011, *101* (3), 398–401.

Druedahl, Jeppe, Raphael Huleux, and Inge Røpke, "Fiscal multipliers in a globally solved HANK model with aggregate risk," *Manuscript*, 2026.

Fernández-Villaverde, Jesús, Juan Rubio-Ramírez, and Frank Schorfheide, "Chapter 9 - Solution and Estimation Methods for DSGE Models," in John B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics*, Vol. 2, Elsevier, 2016, pp. 527–724.

Golosov, Mikhail and Robert E Lucas, "Menu costs and Phillips curves," *Journal of Political Economy*, 2007, *115* (2), 171–199.

Hamilton, James D., *Time Series Analysis*, Princeton University Press, 1994.

Han, Jiequn, Yucheng Yang et al., "Deepham: A global solution method for heterogeneous agent models with aggregate shocks," *arXiv preprint arXiv:2112.14377*, 2021.

Kase, Hanno, Leonardo Melosi, and Matthias Rottner, *Estimating nonlinear heterogeneous agents models with neural networks*, Centre for Economic Policy Research, 2022.

Kim, Jinill, Sunghyun Kim, Ernst Schaumburg, and Christopher A. Sims, "Calculating and Using Second-Order Accurate Solutions of Discrete Time Dynamic Equilibrium Models," *Journal of Economic Dynamics and Control*, 2008, *32* (11), 3397–3414.

Lan, Hong and Alexander Meyer-Gohde, "Pruning in Perturbation DSGE Models: Guidance from Nonlinear Moving Average Approximations," *Working Paper*, 2013.

_ **and** _ , "Solving DSGE Models with a Nonlinear Moving Average," *Journal of Economic Dynamics and Control*, 2013, *37* (12), 2643–2667.

Lee, Hanbaek, "Global nonlinear solutions in sequence space and the generalized transition function," Technical Report, Working paper 2025.

Lombardo, Giovanni and Harald Uhlig, "A Theory of Pruning," *International Economic Review*, 2018, *59* (4), 1825–1836.

Reiter, Michael, "State reduction and second-order perturbations of heterogeneous agent models," *IHS Working Paper*, 2023.

Schmitt-Grohé, Stephanie and Martín Uribe, "Solving Dynamic General Equilibrium Models Using a Second-Order Approximation to the Policy Function," *Journal of Economic Dynamics and Control*, 2004, *28* (4), 755–775.

Yang, Yucheng, Chiyuan Wang, Andreas Schaab, and Benjamin Moll, "Structural Reinforcement Learning for Heterogeneous Agent Macroeconomics," *arXiv preprint arXiv:2512.18892*, 2025.