Abstract

Computational Methods for Economic Models with Function-Valued States

David Childers

2016

Underlying much modern economic theory and econometric methodology is a function space setting. Motivated by this perspective, this dissertation provides a set of tools to construct and solve dynamic economic models with functions as state variables and apply them in models of inequality over space and across individuals. I show how to characterize the solution to function-valued models by linearization in function space, provide a set of algorithms to compute this solution numerically in both regular and ill-posed models, and prove that the algorithms are consistent. The power and efficacy of the methods are illustrated in several examples including a dynamic stochastic model of trade, migration, and economic geography.

Chapter 1: Solution of Rational Expectation Models With Function-Valued States

Many variables of interest to economists take the form of time varying distributions or functions. This high-dimensional 'functional' data can be interpreted in the context of economic models with function valued endogenous variables, but deriving the implications of these models requires solving a nonlinear system for a potentially infinite-dimensional function of infinite-dimensional objects. To overcome this difficulty, I provide methods for characterizing and numerically approximating the equilibria of dynamic, stochastic, general equilibrium models with function-valued state variables by linearization in function space and representation using basis functions. These methods permit arbitrary infinite-dimensional variation in the state variables, do not impose exclusion restrictions on the relationship between variables or

limit their impact to a finite-dimensional sufficient statistic, and, most importantly, come with demonstrable guarantees of consistency and polynomial time computational complexity. Numerical evaluation of the approximation algorithms against a model with an exact benchmark demonstrates that they show speed and accuracy in line with the theoretical guarantees.

Chapter 2: A Dynamic Model of Economic Geography

To study the evolution over time of the spatial structure of economic activity, population, and welfare, I introduce a dynamic stochastic model of trade, migration, and economic geography. In this model, agents make forward-looking costly migration decisions in response to a spatial distribution of wages which is determined endogenously in spatial equilibrium as a function of the population distribution and patterns of persistent regional shocks to amenity values. A closed form characterization of the equilibrium is provided for a class of economies with idealized spatial structure and numerical methods and simulation results are provided for a general case allowing nonparametric spatial heterogeneity along a variety of attributes. The setting allows for a reevaluation of the relationship between spatial agglomeration externalities and population dynamics, suggesting that the sources of long-run spatial heterogeneity may differ substantially from those driving the response to temporary shocks.

Chapter 3: Solving Ill-posed Function-Valued Rational Expectations Models

In linear dynamic stochastic economic models with function space variables, a rational expectations solution taking function-valued inputs to function-valued outputs is defined by a decomposition, analogous to the generalized Schur decomposition of matrices, of a set of operator equations into components. A computationally feasible approximation of the solution may be constructed by projection on a set of basis functions. But when the operators defining the equilibrium conditions are not compact, the solution of the approximate system may fail to converge to the solution of the

true system. This failure arises from multiple breakdowns of continuity in the map from approximation to solution. A solution is devised that enables components to be constructed sequentially, applying regularization at each step. Due to the lack of compactness, standard regularization methods for linear ill-posed inverse problems do not suffice to ensure continuity, and inverses are instead constructed by uneven section methods based on the generalized sampling technique of Adcock *et al.* (2014a). Guidelines for tuning parameter selection are provided, and the performance of the algorithm is demonstrated on an example model.

Computational Methods for Economic Models with Function-Valued States

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
David Childers

Dissertation Director: Peter Phillips

May 2016

Contents

List of Figures

List of Tables

Acknowledgements

I would like to thank my advisors, Peter Phillips and Tony Smith for their guidance, encouragement and support. Without their willingness to let me take on an ambitious project and commitment to ensuring that I saw it through, this thesis would surely not have been possible.

I also owe substantial gratitude to Committee Member Costas Arkolakis, who provided tremendous assistance in completing the project and with the geographic applications.

John Geanakoplos acted as a mentor throughout my graduate school career, for which I am thankful; his constant encouragement of both analytical precision and real-world relevance has substantially strengthened my understanding of economics.

I wish to thank Xiaohong Chen for introducing me to many of the technical areas of inquiry which led to this thesis, as well as guidance along the way.

Other faculty members who provided advice or guidance include Don Andrews, Don Brown, Gerald Jaynes, Yuichi Kitamura, and David Pollard.

I would also like to thank many current and former fellow graduate students. Tim Christensen especially provided substantial advice and guidance. Valentina Assenova,

Brian Baisa, Laura Bakkensen, Claire Brennecke, Cristoph Breunig, Giovanni Compiani, Sabyasachi Das, Rahul Deb, Hao Deng, Yingni Guo, Wenxi Jiang, Lawrence Jin, Namrata Kala, Adam Kapor, Steve Karolyi, Roy Lederman, Margaret Leighton, Yu Luo, Sofia Moroni, Norma Padrón, Zhao Ren, Camilla Roncoroni, Jamin Speer, Cristina Tello-Trillo, Alexis Toda, Kieran Walsh, Ranran Wang, and Nicolas Werquin have all been greatly valued colleagues and friends.

Finally, I would like to thank my parents for all they have done for me.

Notation

$\mathcal{H}$, with any subscript, is assumed to be a complete separable Hilbert space. $\mathcal{B}$, with any subscript, is Banach space. The notation $\| \ \|$ is overloaded: if the object a is an element of $\mathcal{B}$, $\|a\|$ is the norm of a in $\mathcal{B}$, and if a is in a Hilbert space $\mathcal{H}$, $\|a\| = \langle a, a \rangle^{\frac{1}{2}}$ is the norm of a in $\mathcal{H}$, where $\langle \ , \ \rangle$ is the associated inner product. If $a \in \mathcal{H}_a$ but it is not clear from context the space on which a lives, the norm may be denoted $\|a\|_{\mathcal{H}_a}$. $\mathcal{L}(\mathcal{H}_a \to \mathcal{H}_b)$ is the space of bounded linear operators from $\mathcal{H}_a$ to $\mathcal{H}_b$, equipped with the operator norm: for $A \in \mathcal{L}(\mathcal{H}_a \to \mathcal{H}_b)$, $\|A\| = \sup_{\|x\|_{\mathcal{H}_a}=1} \|Ax\|_{\mathcal{H}_b}$. If clarification is required, this norm may be denoted $\|A\|_{op}$. A^* denotes the (Hermitian) adjoint of A: $\forall x \in \mathcal{H}_a, y \in \mathcal{H}_b$, $\langle Ax, y \rangle = \langle x, A^* y \rangle$. A sequence of operators $A_i \in \mathcal{L}(\mathcal{H}_a \to \mathcal{H}_b)$, $i \in \mathbb{N}$ is said to converge in operator norm topology, or 'in norm' to A if $\|A_i - A\| \to 0$. For Γ a Cauchy contour in the extended complex plane $\mathbb{C}_\infty$ (see Conway (1978, Ch. 1 S. 6)) and $f(\lambda) : \ \mathbb{C}_\infty \to \mathcal{L}(\mathcal{H}_a \to \mathcal{H}_b)$ a function from one complex variable to a linear operator, $\int_\Gamma f(\lambda) d\lambda$ is the path integral of $f(\lambda)$ over the curve Γ, as defined in Gohberg *et al.* (1990, Ch. I). I is the identity operator: if the space $\mathcal{B}_a$ on which it acts needs to be specified, it is written $I_{\mathcal{B}_a}$. For $A \in \mathcal{L}(\mathcal{H}_a \to \mathcal{H}_b)$, $\text{Im}(A)$ is the image of A and $\text{Ker}(A)$ is the kernel of A. For a pair of bounded operators (B, A) each in $\mathcal{L}(\mathcal{H}_a \to \mathcal{H}_b)$, following Gohberg *et al.* (1990), define the spectrum $\sigma(B, A)$

as those $\lambda \in \mathbb{C}$ such that $\lambda A - B$ is not invertible, accompanied by the point ∞ if and only if A does not have bounded inverse, and the resolvent set $\rho(B, A)$ as $\mathbb{C}_\infty \backslash \sigma(B, A)$. An operator pair is said to be Γ-regular if for some nonempty subset $\Gamma \subset \mathbb{C}_\infty$, $\Gamma \subset \rho(B, A)$. Brackets $A[h]$ may optionally be used to denote that h is an argument of linear operator A, parentheses $A(h)$ generally denote that h is an argument of (possibly) nonlinear operator A. For nonlinear functions and operators, $F(a, b)$, F_a and F_b are the partial derivatives with respect to arguments a and b respectively. For a variable x, which may be a function, x' denotes the variable in the next time period, not the derivative. The Fourier transform of a function $f(x)$ is denoted with the scale convention $\hat{f}(\omega) := \mathcal{F}[\mathrm{f(x)}](\omega) := \int \exp(-2\pi\iota\omega x) f(x) dx$.

Chapter 1

Solution of Rational Expectation Models With Function-Valued States

Many variables of interest to economists take the form of time varying distributions or functions. This high-dimensional 'functional' data can be interpreted in the context of economic models with function valued endogenous variables, but deriving the implications of these models requires solving a nonlinear system for a potentially infinite-dimensional function of infinite-dimensional objects. To overcome this difficulty, I provide methods for characterizing and numerically approximating the equilibria of dynamic, stochastic, general equilibrium models with function-valued state variables by linearization in function space and representation using basis functions. These methods permit arbitrary infinite-dimensional variation in the state variables, do not impose exclusion restrictions on the relationship between variables or limit their impact to a finite-dimensional sufficient statistic, and, most importantly, come with demonstrable guarantees of consistency and polynomial time computational complexity. Numerical evaluation of the approximation algorithms against a model with an exact benchmark demonstrates that they show speed and accuracy in line with the theoretical guarantees.

1.1 Introduction

In order to understand and evaluate the causes and consequences of economic heterogeneity, it is helpful to have an analytical framework in which the distribution of heterogeneity can change over time and can both affect and be affected by other variables. A perspective in which some of the state variables of an economic model are endogenous random functions allows distributions, as well as objects like demand and supply curves or policy or value functions, to be treated as data. While descriptive models and methods for function-valued time series are undergoing rapid development,[1] interpreting this data requires formulating economic models capable of generating the observed functional data and deriving their implications. For models featuring forward looking decision making and endogenous aggregate variables, this derivation typically requires solving a computationally intractable infinite-dimensional system of nonlinear expectational difference equations. Although heuristic or strongly model dependent methods have been proposed, to date there appears to be no general purpose algorithm which provides a formal guarantee of even an approximate solution to rational expectations models with stochastic function-valued states.

This paper provides such an algorithm. In particular, it demonstrates how the equilibrium conditions for a general class of function-valued rational expectations models, including but not limited to heterogeneous agent dynamic stochastic general equilibrium models, can be linearized directly in function space, with solutions characterized locally by a functional linear process, a tractable empirical model for function-valued time series (Bosq, 2000). Construction of a local solution requires introducing a novel infinite-dimensional extension of the generalized Schur decomposition used to solve finite-dimensional rational expectations models (Klein, 2000) and developing perturbation theory for this object, which may be contributions of

[1] See Horváth & Kokoszka (2012); Bosq (2000); Morris (2014); Ferraty & Romain (2011) for surveys of the rapidly expanding field of functional data analysis, which focuses on modeling, estimation, and inference for series of observed or estimated functions.

independent mathematical interest. The solution can be implemented numerically by a procedure based on finite-dimensional projection approximations which converges to the local solution under mild regularity conditions. I analyze in detail a particular approximation algorithm in this class, a wavelet transform based procedure which yields an approximate solution accurate to within any desired degree in polynomial time.

To evaluate the method, I apply it to a dynamic spatial model of trade, migration, and economic geography, developed and further analyzed in Chapter 2, which introduces forward looking migration decisions and spatial shocks into the economic geography model of Krugman (1996). In the model, the spatial distribution of population, wages, and welfare over a continuum of locations is allowed to vary nonparametrically in response to persistent spatially correlated shocks to the desirability of different locations. Due to the spatial structure of trade and production, the spatial distribution of economic activity is a determined by the distribution of population across locations, while the distribution of population is determined by forward looking migration decisions which take into account the expected distribution of economic activity. In this setting, the relationship between these two functions is not easily reduced to low-dimensional summaries or split into "local" and "global" components, but is well characterized by a functional linear model representation. By exploiting an analytical characterization of the solution to (certain parameterizations of) this model, the speed and numerical accuracy of the algorithm are evaluated in practice and shown to be in line with the strong theoretical guarantees.

The core idea behind the solution method is *functional linearization*. By taking the functional derivatives of the equations defining an equilibrium, it is possible to construct a system of equations which can be solved for the functional derivatives at a fixed point in function space of the policy operator, a map from function-valued states to function-valued endogenous variables. In this way, it is possible to recover local

information about the solutions, which can then be used to construct a functional Taylor expansion of the policy operator which provides an accurate solution for all functions not too far from the function around which the model is linearized.

Constructing this linear approximation of the policy operators from the functional derivatives of the model equations requires solving a system of quadratic equations in linear operators. In the case of linear or linearized finite-dimensional rational expectations models, the analogous quadratic equation can be solved using matrix decomposition. In particular, Klein (2000) demonstrated that a solution can be found using the generalized Schur (or QZ) decomposition of the matrices of derivatives. In infinite dimensions, an analogous decomposition appears to be absent from the literature, in part because the finite-dimensional version is constructed by induction using eigenvalues, which may fail to exist or have countable cardinality in infinite-dimensional space. Nevertheless, it is possible to construct an analogous decomposition by other methods, described in detail in Appendix A. Under the conditions required for such a decomposition to exist and under further conditions analogous to the well known criterion of Blanchard & Kahn (1980) ensuring that the model has a linear solution, it is possible to solve for the first order expansion of the policy operator.

Calculating this local solution numerically requires representing it in a form that can be evaluated on a computer. A standard procedure for reducing problems in function spaces to finite-dimensional objects is to approximate the functions by projecting the space onto the span of a set of basis functions, such as wavelets, splines, or trigonometric or Chebyshev polynomials, and representing operators on function space in terms of their behavior with respect to the basis functions. These approaches are referred to as spectral methods and are commonly applied to solve integral and differential equations: see Boyd (2000), Chatelin (2011). If any function we are interested in can be represented reasonably accurately by a finite set of basis functions, the loss from the use of a finite set of functions may be small. The caveat here is

that, unlike in classical function approximation problems where the class of functions is known, 'the set of functions we are interested in' is not explicitly assumed, but must be determined by the properties of the model.

The issue that projection methods must overcome is that the class of functions well approximated by finite projection is in fact small in the class of all possible functions which could conceivably arise endogenously as outcomes of an implicitly defined model with function valued variables. To handle this concern, conditions must be imposed on the model which ensure both that the solutions themselves are continuous with respect to projection approximations and that the solutions are operators which have the property that they map functions which are well approximated by basis functions to functions which are well approximated by basis functions. Continuity properties of the generalized Schur decomposition are derived in Appendix A, and a set of restrictions on the model which ensure that basis function approximation is valid is described in Section 4.2.

While the precise statements of the sufficient conditions on a model for projection to be valid are somewhat technical, the conditions themselves are rather mild. Essentially, they rule out certain kinds of maps which take well-behaved smooth functions as input and produce jagged, noisy, or discontinuous functions as output. Many economic models can be represented in forms which satisfy these conditions, and many of those that do not can be modified slightly so that they do, for example by smoothing discontinuous cost functions or adding a small amount of noise to ensure that a distribution remains smooth.

Provided that the regularity conditions hold, implementing the solution is simple and fast. The linearized equilibrium equations can be approximated by projection, either analytically or numerically by quadrature, to produce two pairs of matrices, to which one can apply the finite-dimensional QZ decomposition, solve and combine to form a matrix approximation to the infinite dimensional policy operator. The

accuracy of the approximation is then determined by the number of basis functions used and the smoothness of the functions that they are used to approximate. If all the equilibrium conditions are defined using Hölder continuous functions, wavelets provide the smallest and fastest feasible representation. Implementing approximate projection using the Discrete Wavelet Transform, the method converges in a number of operations polynomial in the degree of accuracy of the solution and in numerical experiments gives demonstrably accurate results at high speed. High level conditions are also provided for more general procedures, including for the case when parts of the model are estimated directly from data.

The dynamics of economic heterogeneity have been considered from a variety of perspectives. Surveys of heterogeneous agent models are available in Krusell & Smith (2006); Heathcote *et al.* (2009); Guvenen (2011); Ljungqvist & Sargent (2004). A canonical framework is the Bewley model (Bewley, 1986), sometimes referred to as the Bewley-Huggett-Aiyagari model after the models and algorithms of Huggett (1993) and Aiyagari (1994). These models produce a time invariant cross-sectional distribution of income and wealth given by the stationary distribution endogenously induced by individual decisions which are themselves determined by the distribution. While not permitting any stochastic variation over time in distributions, the algorithms introduced to solve these nonstochastic models can be used as the first step in the linearization procedure I will provide, to find the point in the space of distributions around which to construct a linearized solution to a model with a stochastic distribution of heterogeneity.

To accommodate the setting where the distribution may evolve stochastically over time, Krusell & Smith (1998) introduce aggregate uncertainty into the Bewley model and provide a procedure to calculate approximate decision rules and generate dynamics of distributions jointly by simulation and representation of the impact of the distribution on decisions through a small set of moments. This method is particularly

well suited to the model in that it takes advantage of a feature the authors refer to as 'approximate aggregation.' In the Krusell-Smith model, due to the use of a one-dimensional source of aggregate variation, an economic structure in which the impact of the distribution on the decision problem occurs only indirectly through its impact on prices in a centralized market, and a set of preferences and constraints that yields a decision rule which appears close to linear in individual states over most of the state space, a low-dimensional set of statistics of the wealth distribution suffices to describe its dynamics with a high level of accuracy. As a result, the decision problem can be reduced to a low-dimensional nonlinear decision problem in these statistics with apparently minimal loss of accuracy.

However, many of the features which make the Krusell-Smith method well suited to their model and similar models are far from universal. In particular, the finite-dimensional set of aggregate shocks limits variation in the shape of the functions of interest, and may create difficulty in matching estimates of the functions from cross-sectional data. Unless variation over time in the function lies exactly on the same low-dimensional space along which the model implies the functions move, observed functions may not be consistent with any possible values of the aggregate shocks, and so full information statistical methods will reject the specification completely. The low-dimensional set of aggregate shocks may also prevent consideration of economically important features, whether they affect the shape of the functions directly, such as shocks to uncertainty, skewness, and higher moments of a distribution, as documented by Guvenen *et al.* (2012) for income distributions, or enter the model via other variables, such as the variety of sources uncertainty included in medium scale DSGE models. While procedures like the Krusell-Smith method can handle some increase in the dimensionality of the space of aggregate shocks by adding more statistics, increasing the dimension of the state space in the intertemporal decision problem can be computationally costly, with naive approaches based on discretiza-

tion or tensor product function representations scaling exponentially in the number of state variables, and more sophisticated approaches, such as the Smolyak method of Gordon (2011), requiring some degree of difficult to verify regularity in the induced distribution.

Moreover, the restriction of the impact of the distribution to acting only on a finite set of market prices is less tenable in situations where interactions are decentralized, or the distribution enters the decision problem directly. This can be the case, for example, in spatial models, like the one considered in Section 3, where market outcomes differ across locations due to costs of trade or other economic interactions over distance, and in which decisions depend on the entire spatial distribution of economic activity due to both local and long distance interactions.

Many extensions and alternatives to the Krusell-Smith method are available: a *Journal of Economic Dynamics and Control* symposium (Den Haan, 2010) compares a variety of methods. Methods based on linearization or perturbation are not new, and are explored in, among others, Reiter (2009), Chung (2007), Winberry (2014) and Veracierto (2014). Perturbation approaches, which build on the class of linear rational expectations solution methods introduced by Blanchard & Kahn (1980) and extended by Klein (2000) and others, describe variability locally, and are much more amenable to including high-dimensional aggregate shocks than global, fully nonlinear approaches. Reiter (2009) developed the approach of linearizing models around nonlinear functions and distributions, and noted that by doing so, local methods can capture differences of large magnitude in the heterogeneous state between individuals and completely nonlinear responses to those differences while maintaining the tractability of linear methods for aggregate variables.

All of the considered perturbation approaches differ from the one advocated here because they do not consider linearization in function space. Instead, they replace functions by finite-dimensional approximations, either by projection or discretization,

and then linearize and construct a solution based on applying algorithms applicable to finite-dimensional expectations models. The difference between linearizing and solving in infinite-dimensional space before taking a finite-dimensional approximation and taking a finite-dimensional approximation before linearizing and solving may seem minor, but the first approach is key to ensuring that the resulting solution is well defined in terms of the true solution of the model and that the algorithm produces an answer which is provably close to this solution. These approaches also fail to consider that, except under certain regularity conditions, applying a finite-dimensional solution method to the approximated equilibrium conditions need not guarantee that the solution is accurate, even if the approximation of the conditions is, because the solution is not in general continuous with respect to the approximation error.

More generally, none of the methods described, including Krusell-Smith, provide or attempt to provide any formal proof that the approximation converges to a true solution. Although numerical demonstrations may be used to assess features consistent with the accuracy of the methods and so diagnose certain inaccuracies in an approximated solution, they cannot certify that the output of the algorithm is valid. Because it comes with formal guarantees, the functional linearization approach introduced here provides for the first time a benchmark which can be used to characterize a solution to dynamic models with heterogeneous agents which can be assured to be accurate.

Outline

The structure of this paper is as follows. I describe the setting of rational expectations models with function valued states in Section 1.2. Section 1.3 characterizes and gives necessary conditions for the existence of solution to the linearized model, while Section 1.4 introduces projection algorithms for calculating this solution and

describes conditions for their consistency. Section 1.5 evaluates the performance of the procedure on the model developed in detail in Chapter 2. Section 1.6 concludes. Several appendices contain technical results: Appendix A.1 describes conditions for existence of an infinite-dimensional version of the generalized Schur decomposition, and Appendix A.2 gives conditions under which it is continuous. Appendix C collects all proofs, and Appendix B provides high level sufficient conditions for the existence of recursive equilibria in function-valued dynamic models.

1.2 Function-Valued Models and Linearization in Function Space

The class of dynamic economic models which may be placed in a framework amenable to linearization in function space is large. Many economic models define objects of interest, explicitly or implicitly, as functions which solve a set of equations representing conditions such as optimization, market clearing, self-consistency, feasibility, or accounting identities. For example, a consumption function is often represented implicitly as the solution to an Euler equation, or a value function as the fixed point of a Bellman operator. Most trivially, when economic variables take values in Euclidean space, all of the theory developed in this paper will continue to apply. To see how random functions may naturally enter the description of an economic model, let us first consider a simple and illustrative case, before providing a general characterization.

It is common in microeconometric study of the dynamics of income and consumption by consumers or production by firms to model individual behavior by a linear dynamic panel model. A simplified version of this model is given by the assumption that, for each agent i, the variable of interest ζ_{it} follows the autoregressive process $\zeta_{it+1} = \rho\zeta_{it} + \epsilon_{it+1}$, where ϵ_{it+1} is independent of ζ_{it} and across agents and $|\rho| < 1$. While it is conventional to take an interest in the individual persistence parameter

ρ, for the purposes of analysis of aggregates and welfare we may also be interested in the cross sectional distribution of the attribute ζ_{it}, which may be represented by pdf $f_t(\zeta)$. Given a measure 1 continuum of agents following this rule, the evolution of this distribution can be determined from its past value and the distribution of the shock ϵ_{it+1}. To model time varying effects such as aggregate shocks, we may let $\epsilon_{it+1} \overset{i.i.d.}{\sim} p_{t+1}(\epsilon)$ across agents, where the density function $p_{t+1}(.)$ may be taken as a function-valued random variable for each t. This models not only mean shifts, as would be captured by time fixed or random effects, but also distributional changes such as the changes in polarization or tail behavior of income risk as documented, for example, in Guvenen *et al.* (2012). Under this assumption, we have a dynamic equation for the evolution of the distribution of ζ, given by the convolution of the past distribution and the shock distribution

$$f_{t+1}(\zeta) = \int p_{t+1}(\zeta - \rho u) f_t(u) du \tag{1.1}$$

which provides a recursive representation for a function-valued economic variable of interest, $f_t(.)$, in terms of current and past values of the state, an operator mapping between them, and an exogenous shock which is also function valued, $p_{t+1}(.)$.

To formalize the linearization procedure for this and other models, and to provide a framework which permits both variables which are predetermined and those determined by forward looking expectations, we provide a notational framework for a general class of models. The notation and structure to be used follows closely that of Schmitt-Grohe & Uribe (2004), who described perturbation procedures for finite-dimensional rational expectations models, with the difference that we now allow state variables to be elements of an infinite-dimensional space. We consider in particular models with a recursive representation described by a set of equilibrium conditions which may be expressed as differentiable operators between separable Banach spaces.

A solution to the model defines a recursive law of motion for the endogenous variables in the system in terms of the exogenous variables and past values of endogenous variables. The law is determined implicitly as the solution of a nonlinear expectational difference equation

$$\mathbb{E}F(x, y, x', y', \sigma) = 0_{\mathcal{B}_2} \tag{1.2}$$

where $x \in \mathcal{B}_x$ is a set of predetermined variables, $y \in \mathcal{B}_y$ is a set of endogenous or 'jump' variables, a superscript x', y' indicates the values of these elements in the next time period $t+1$ and the absence thereof indicates values of variables known at time t, $\sigma \in \mathbb{R}$ is a scalar scaling parameter determining the size of fluctuations. The function $F(x, y, x', y', \sigma) : \mathcal{B}_x \times \mathcal{B}_y \times \mathcal{B}_x \times \mathcal{B}_y \times \mathbb{R} \to \mathcal{B}_2$, which we refer to as the *equilibrium operator*, is a map taking the values of the state variables today and tomorrow and the scaling parameter to a space $\mathcal{B}_2$, and $\mathbb{E}$ is the (Bochner) expectation with respect to the law of motion induced by the solution of the model, to be made explicit shortly.[2]

Uncertainty in the model is incorporated solely via exogenous Banach random elements z' on probability space $(\mathcal{B}_z, \Sigma_z, \mu^z)$, which enter into the exogenous law of motion generating a subset of the predetermined variables x_2, with $(x_1, x_2) \in \mathcal{B}_{x_1} \times \mathcal{B}_{x_2} = \mathcal{B}_x$, by the equation $x_2' = h_2(x_2) + \sigma z'$ for $h_2 : \mathcal{B}_{x_2} \to \mathcal{B}_{x_2}$ a given function describing the dependence of future values of x_2 on current values. The shocks z' are normalized to have zero mean $\mathbb{E}[z'] = 0$. As a result, F contains as one subcomponent the formula $x_2' - h_2(x_2)$.

While this form may appear somewhat restrictive, many apparent limitations may be addressed through inclusion of appropriate auxiliary variables and equations. For example, while only variables in two time periods are included, by including lags and

[2]The Bochner integral of a $\mathcal{B}$-valued random variable g on probability space (Ω, Σ, μ) is given by an element $\mathbb{E}g \in \mathcal{B}$ defined for simple functions $g = \sum_{i=1}^n f_i\{\omega \in A_i\}$ for f_i in $\mathcal{B}$, $A_i \in \Sigma$ as $\mathbb{E}g = \sum_{i=1}^n f_i \mu[A_i]$ and for more general random variables g as the strong limit of the Bochner integral of a sequence of simple functions g_n such that $\mu \|g - g_n\|_{\mathcal{B}} \to 0$. A measurable random element is Bochner integrable if and only if $\mu\|g\|_{\mathcal{B}} < \infty$.

leads as separate variables, systems dependent on more time periods may be brought into this recursive form. Likewise, while function-valued uncertainty z' is restricted to enter additively in the model, nonlinear effects of shocks may be included by adding an additional predetermined variable which is a function of the shock: e.g., if z_k enters nonlinearly in F, replacing z_k with x_{2k} and incorporating the equation $x'_{2k} = \mathbb{E}z'_k + \sigma(z'_k - \mathbb{E}z'_k)$ can recover the nonlinear effects. Overall, beyond imposing a recursive structure, the form provides a consistent notation but imposes only modest restrictions on the form of the economic model.

A (recursive) solution is given by a set of policy operators which solve the equilibrium equation for any value of the initial predetermined state x and the exogenous shocks z. In each period, y is given by the endogenously determined map $g(x, \sigma)$ from predetermined state x to endogenous state y (or x' to y'), and x' is given by the transition operator $h(x, \sigma) + \sigma\eta z'$ mapping the current predetermined state and shocks to next period's predetermined state, where η denotes the imbedding $\mathcal{B}_{x_2} \to \mathcal{B}_x$, i.e. for $z \in \mathcal{B}_{x_2}$, $\eta[z] = (0, z) \in \mathcal{B}_{x_1} \times \mathcal{B}_{x_2}$, and $h(x, \sigma) = (h_1(x, \sigma), h_2(x_2))$ includes both an endogenously determined transition component h_1 and an exogenous component h_2.

Definition 1.1. A *recursive solution* is a set of maps $g(x, \sigma) : \mathcal{B}_x \times \mathbb{R} \to \mathcal{B}_y$, $h_1(x, \sigma) : \mathcal{B}_x \times \mathbb{R} \to \mathcal{B}_{x_1}$, $h_2(x_2) : \mathcal{B}_{x_2} \to \mathcal{B}_{x_2}$ such that the *equilibrium conditions* hold:

$$\mathbb{E}F(x, g(x, \sigma), h(x, \sigma) + \sigma\eta z', g(h(x, \sigma) + \sigma\eta z', \sigma), \sigma) = 0_{\mathcal{B}_2} \tag{1.3}$$

for all x, σ, where the expectation $\mathbb{E}$ may now be defined, for each x, σ as the expectation with respect to the pushforward measure of μ^z on $\mathcal{B}_2$ generated by the function $F(x, g(x, \sigma), h(x, \sigma) + \sigma\eta z', g(h(x, \sigma) + \sigma\eta z', \sigma), \sigma) : (x, z', \sigma) \in \mathcal{B}_x \times \mathcal{B}_{x_2} \times \mathbb{R} \to \mathcal{B}_2$ evaluated at fixed x, σ.

It can be shown that this pointwise in x definition of a solution generates a stochastic process for (x_t, y_t) under mild measurability conditions on the functions chosen.

(i) Let $\{z_t\}_{t=0}^{\infty}$ be an i.i.d. sequence on the infinite product of independent copies of $(\mathcal{B}_z, \Sigma_z, \mu^z)$ and initial value x_0 be defined on $(\mathcal{B}_x, \Sigma_x)$ with distribution μ_0^x, where Σ_x is a sigma field containing Σ_z. (ii) Fix $\sigma \in \mathbb{R}$. Suppose $h(x, \sigma)$ is $(\mathcal{B}_x, \Sigma_x) \to (\mathcal{B}_x, \Sigma_x)$ measurable, $g(x, \sigma)$ is $(\mathcal{B}_x, \Sigma_x) \to (\mathcal{B}_y, \Sigma_y)$ measurable for some Σ_y, and F is measurable with respect to the product sigma field $\Sigma_x \otimes \Sigma_y \otimes \Sigma_x \otimes \Sigma_y$ on $\mathcal{B}_x \times \mathcal{B}_y \times \mathcal{B}_x \times \mathcal{B}_y$

The measurability restrictions on h and g do impose some nontrivial limitations on the class of solutions to be considered by ruling out auxiliary randomness in the policy functions for aggregate variables beyond that included in z. For certain classes of models, randomization may be necessary to ensure existence of a solution, see Miao (2006). If this can be incorporated in z by expanding the state space, this poses no difficulty, but because the model will be solved by approximating near a point with no aggregate variability, the method cannot accommodate models which have no solution without aggregate randomness.

Proposition 1.1. *The series defined recursively by* $x_0 \sim \mu_0^x$, $x_{2,t+1} = h_2(x_{2,t}) + \sigma z_{t+1}$, $x_{1t+1} = h_1(x_t, \sigma)$, $y_t = g(x_t, \sigma)$ $\forall t \geq 0$, *where* h, g *are a recursive solution satisfying Condition 1, is measurable with respect to the infinite product sigma field and* $\mathbb{E}F(x, g(x, \sigma), h(x, \sigma) + \sigma\eta z', g(h(x, \sigma) + \sigma\eta z', \sigma), \sigma)$ *coincides with the conditional expectation of* $F(x_t, g(x_t, \sigma), h(x_t, \sigma) + \sigma\eta z_{t+1}, g(h(x_t, \sigma) + \sigma\eta z_{t+1}, \sigma), \sigma)$ *at time* t *given* $x_t = x$.

Proof. See Appendix. □

In order to ensure computation of a stationary solution, the point around which the model is linearized is a nonstochastic steady state, which allows construction of a solution which is both local and recursive, by ensuring that the point around which the rule is calculated is the same in all periods.

Definition 1.2. A *nonstochastic steady state* is a set of values $(x^*, y^*) \in \mathcal{B}_x \times \mathcal{B}_y$

such that when $\sigma = 0$ and so function-valued uncertainty disappears, F satisfies

$$F(x^*, y^*, x^*, y^*, 0) = 0$$

Many recursive models will have such a point, at which all aggregate variables are unchanging over time. This is the equilibrium concept used in Bewley-Huggett-Aiyagari models, in which the distribution of heterogeneity is given by an invariant distribution generated by individual decision rules, and its existence can often be guaranteed by fixed point theorem. It may also be calculated consistently by standard methods, such as the iterative algorithms proposed by Huggett (1993) and Aiyagari (1994). In general, determining the nonstochastic steady state of a model involves solving a functional equation, which will differ depending on the details of the model. However, the problem involves determining only a single set of functions rather than an operator valid for any function, and is often quite feasible using standard methods. For example, in models where the decision rule is a function-valued state variable, recursive solutions are often available by dynamic programming, for which there are many feasible approximation algorithms with exponential or similarly fast convergence rates. Calculation of invariant distributions of Markov processes is also often achievable by iterative methods with exponential convergence. More broadly, in the absence of infinite-dimensional uncertainty, the problem usually reduces to a set of integral equations, for which a broad variety of standard numerical integral equation methods may be used.

A linearized solution of the model is given by first order Taylor expansion of $g(.)$ and $h(.)$ with respect to their arguments at the steady state. In order to solve for this, $g(.)$ and $h(.)$ and the operator $F : \mathcal{B}_x \times \mathcal{B}_x \times \mathcal{B}_y \times \mathcal{B}_y \to \mathcal{B}_2$ must be differentiable with respect to their arguments. In Banach space, the appropriate notion of derivative for linearization is (usually) the Fréchet derivative, which is defined analogously to the

Fréchet derivative in Euclidean space.[3] If $F(x)$ is operator between Banach spaces $\mathcal{B}_1 \to \mathcal{B}_2$, the Fréchet derivative, if it exists, is the continuous linear operator DF satisfying

$$\lim_{\|h\|_1 \to 0} \frac{\|F(x+h) - F(x) - DF[h]\|_2}{\|h\|_1} = 0. \tag{1.4}$$

In practice, calculation of Fréchet derivatives of Banach space-valued operators is not difficult: they obey many of the standard rules of Euclidean-valued derivatives including linearity, additivity, and the product rule, and many standard operators have known derivatives: see e.g. Kesavan (2004). Most importantly, the Fréchet derivative follows a version of the chain rule: for two Fréchet differentiable operators F, G, $D(F \circ G)[h] = DF[DG[h]]$. Fréchet differentiability is strictly stronger than directional, or Gateaux differentiability, which requires only the existence of a limit in the direction of a fixed element h. As the Gateaux derivatives of F in any direction $h \in \mathcal{B}$ may be calculated as the scalar derivative $\frac{d}{d\tau}F(x + \tau h)$ at $\tau = 0$ and must coincide with the Fréchet derivative when the latter exists, the form of the Fréchet derivative is easily determined. The Fréchet derivative preserves linear operators, so integration, differentiation, multiplication by a function, and any composition thereof have derivatives equal to themselves. A special class of operators which arises frequently in economic models is the composition of one function with another, referred

[3]For operators defined or differentiable only on subsets (not necessarily subspaces) of infinite-dimensional Banach spaces, such as the positive cone of non-negative measures, it may sometimes be desirable to consider the Hadamard derivative (see Flett (1980)) tangential to a set, which requires a derivative to be defined uniformly only over compact sets and so is weaker than the Fréchet derivative, which requires uniformity over closed balls, which in infinite dimensions are not compact. While the chain rule and a version of the implicit function theorem also apply for this class of derivatives, and so a linear approximation may be defined by the same equations with the Hadamard in place of the Fréchet derivative (and the derivatives exactly coincide on finite-dimensional spaces), the Taylor expansion will in be defined only over the subset on which a Hadamard derivative exists, and have a remainder with size dependent on the direction of the path of the approach, rather than just the norm. The domain restriction may not be a problem in practice, as in most cases the derivatives may be extended by the Hahn-Banach theorem to a larger space in a canonical way, for example, by removing positivity restrictions. If a solution does exist and is Hadamard differentiable, the Hadamard derivatives of the operators of interest will coincide on their domain with the extended operators on the total space. When this is the case, the same first order approximation may be constructed and approximated by the algorithm provided, but will be valid only for directions in which Hadamard differentiability holds.

to as a Nemytskii operator. Under appropriate boundedness, differentiability, and integrability conditions on $f(s_1, s_2)$, the composition $f(g(x), s_2)$, viewed as a map from the function $g(.)$ of x to the function $f(g(x), s_2)$ of x, s_2 is a Fréchet differentiable function of $g(.)$ at the point $g^*(.)$ with derivative $f_{s_1}(g^*(x), s_2) \cdot [h(.)]$: that is, the derivative is given by multiplication of the direction in which g changes by the partial derivative of f with respect to the element with which it is composed (Kesavan, 2004). In Banach space, Taylor's theorem for Fréchet derivatives gives a linear approximation of a differentiable operator $F(x) : \mathcal{B}_1 \rightarrow \mathcal{B}_2$ as $F(x + h) = F(x) + DF[h] + o(\|h\|_1)$.

It is important to contrast a Taylor expansion in function space with a local solution for finite-dimensional models. The point at which the linearization is constructed is the stationary state of the model in the situation where the variance of function-valued shocks is taken to 0. This is not the same as shutting down all variability in the model. In most heterogeneous agent models, individuals face a distribution of idiosyncratic uncertainty which may be arbitrarily dispersed and induces a nondegenerate stationary distribution of heterogeneity, in which the state of each individual evolves stochastically over time. In these models, the steady state function is the stationary distribution of heterogeneity, in the absence of aggregate shocks. For example, the unemployment rate can be constant over time while each individual faces employment risk, with the number of people entering and leaving unemployment equal. Similarly, linearization of the policy operator does not imply all decision rules are linear: decisions with respect to individual characteristics may be arbitrarily nonlinear. Instead, the relationship between function-valued state variables is expressed in terms of linear operators. Loosely, the value of one nonlinear function at each point can be thought of as approximated by a linear function of the values of each other function at each point.[4] As a result, a first order functional Taylor expansion can describe rather

[4]This description is accurate for discrete functions: for general functions, the proper statement is that each linear functional of the output function is equal to a different linear functional of the input function.

complicated patterns of behavior.

To illustrate the process of linearization, consider the law of motion for the distribution provided by the panel model in (1.1): a linear approximation will describe the law of motion in the case of 'small' changes in the distribution. To consider the model in the case of small i.i.d. over time aggregate shocks to the cross sectional distribution $p_{t+1}(\epsilon)$ of idiosyncratic shocks, write the law deviations from the mean as $p_{t+1}(\epsilon) - p^*(\epsilon) = \sigma z_{t+1}(\epsilon)$, for $z_{t+1}(\epsilon)$ an i.i.d. over time Bochner mean 0 random function so that at $\sigma = 0$ the distribution of ϵ is constant over time at a fixed distribution $p^*(\epsilon)$. In the above notation $x_2' = p_{t+1}(\epsilon) - p^*(\epsilon)$, $z' = z_{t+1}(\epsilon)$ and $h_2(x_2) = 0$ because we have assumed that the exogenous aggregate shocks are not persistent.[5] To complete the description of the model, we may take as the endogenous predetermined variables $x_1 = f_t$, $x_1' = f_{t+1}$, and

$$F(x_1, x_2, x_1', x_2') = \begin{bmatrix} f_{t+1}(\zeta) - \int p_{t+1}(\zeta - \rho u) f_t(u) du \\ p_{t+1}(\epsilon) - p^*(\epsilon) \end{bmatrix}$$

as the equilibrium operator defining the model. In this case, all variables are predetermined or exogenous, so there is no y variable. A linear approximation with respect to f and p is given by taking the functional derivative of F with respect to p_t,f_t, p_{t+1}, f_{t+1} around a nonstochastic steady state f^*, p^* satisfying $p(\epsilon) = p^*(\epsilon)$, $f^*(\zeta) = \int p^*(\zeta - \rho u) f^*(u) du$, which exists whenever $|\rho| < 1$ under mild conditions on the density p^* of the error term: see Christensen (2014). Applying the chain rule and the product rule, a Taylor expansion of the law of motion for f_{t+1} in p_{t+1} and f_t

[5] Allowing $h_2(x_2)$ in the model to be nonzero would represent persistence in the aggregate shock to the distribution of error terms. After linearization, the cross sectional distribution of ζ, the observable individual characteristic, would then be approximated by a functional ARMA(1,1) process, instead of a functional AR(1).

is given by

$$f_{t+1}(\zeta) = \int p^*(\zeta - \rho u) f^*(u) du + \int p^*(\zeta - \rho u)[f_t(u) - f^*(u)] du$$
$$+ \int [p_{t+1}(\zeta - \rho u) - p^*(\zeta - \rho u)] f^*(u) du + o(\|(f_t - f^*, p_{t+1} - p^*)\|).$$

Substituting in the (already linear) law of motion $p_{t+1} = p^* + \sigma z_{t+1}$ and the steady state relation, obtain

$$f_{t+1}(\zeta) - f^*(\zeta) = \int p^*(\zeta - \rho u)[f_t(u) - f^*(u)] du + \sigma \int z_{t+1}(\zeta - \rho u) f^*(u) du + o(\|(f_t - f^*, p_{t+1} - p^*)\|)$$

which expresses the deviation from the steady state in time $t+1$ as given by a linear operator applied to the deviation from steady state in time t plus, by linearity of the expectation and of the integral operator applied to $z_{t+1}(.)$, a mean 0 exogenous Banach random element. That is, it may be written as

$$f' - f^* \approx B[f - f^*] + \sigma \varepsilon'$$

for some linear operator B and some mean zero noise ε', a linear functional autoregression as in Bosq (2000), so long as both the noise and the deviation from a steady state are small. As similar procedures may be applied to more general dynamic panel data models, one sees that a functional linear process may provide a local approximation to the law of motion for distributions of cross-sectional aggregates for a wide range of commonly used empirical models of individual and aggregate behavior.[6]

In what follows, I will show how to use the functional derivatives of a broad class of dynamic function-valued models to solve for the linearized dynamics and responses of the state variables of those models to endogenous and exogenous changes.

[6]For models as simple as in 1.1, it is possible to characterize the behavior without approximations: the random linear operator model generalizes the random coefficients model to infinite dimensions, and has been analyzed in Skorohod (1984).

1.3 Characterization of Equilibrium Solution

While simple models with a high degree of structure may have linear approximations to solutions which can be found easily by heuristic methods, for larger or more complicated systems, and especially those where components are mutually determined, a more systematic approach is required. Fortunately, the steps involved in finding a linear solution may be described explicitly and so reduced to an algorithm which automates construction. The idea behind the method is to use a decomposition of the equilibrium conditions into components which may be solved separately and recursively, by taking a component whose evolution may be expressed as a function of past variables and a component which is solved by iterating forward expectations of future variables. While in some models, the components which are solved by looking backwards and the components which are solved by looking forwards may be identified with separate variables in the system, this is not true in general. Instead, this separation must be determined endogenously in such a way that initial and end-point conditions of the system are satisfied. What this often consists of is the requirement that some choice variables or other endogenous variables must be chosen to affect the expected evolution of other variables so that they satisfy an endpoint condition.

This is the source of the logic behind the cross-equation restrictions implied by many classical rational expectations models: the path of the endogenous variables must be determined jointly, and so expectations regarding one variable possibly far in the future may cause another variable to move far in advance. For example, in the Krugman (1979) model of balance-of-payments crises (and more recent models of the same), foreign exchange investor behavior is tied down by expected future optimality conditions at the point when the sovereign runs out of reserves. In the Dornbusch (1976) overshooting model, the nominal exchange rate follows a nonmonotonic path to ensure consistency with both short run price rigidity and a long run purchasing power parity anchor. In infinite-horizon settings, the role of an end-point condition

which coordinates expectations is played by analogous long-run optimality or consistency conditions. For example, in the fiscal theory of the price level, the price adjusts to ensure consistency of long run expectations of the government budget deficit. Most commonly (at least in real models: see Cochrane (2011) for a discussion of complications in nominal models), long-run behavior is determined by a condition, such as transversality, which is satisfied when variables follow a dynamic path which is stationarity. While many types of long run restrictions are possible in models with function-valued state variables, we will provide an algorithm for this most common case, in which the specified model takes recursive form over an infinite horizon and endpoint conditions require (or permit) a stationary solution. Although some modification is possible, including requiring asymptotic convergence (or slow divergence) at a particular rate possibly above or below 1, due to the infinite-dimensional nature of the parameter space, arbitrary endpoint conditions introduce substantial complications and so these will not be discussed further.

The requirement that it is possible to separate into solvable components also imposes one more technical limitation: to ensure orthogonality of projections, in what follows, we specialize from the setting of arbitrary separable Banach spaces to require all variables to live on separable Hilbert spaces: $\mathcal{H}_1 = \mathcal{H}_x \times \mathcal{H}_y$ and $\mathcal{H}_2$ replace $\mathcal{B}_1 = \mathcal{B}_x \times \mathcal{B}_y$ and $\mathcal{B}_2$, respectively. This will also be helpful in the numerical implementation. For models defined on spaces which can be densely embedded into a Hilbert space, it is often possible to extend the derivatives to the full Hilbert space by completion. However, norm convergence results must then be taken with respect to the Hilbert space norm.

For an economic model with recursive solution which is differentiable and generates a stationary stochastic process, we describe conditions that the functional derivatives of the solution operators $g(.)$ and $h(.)$ must satisfy, which will allow these derivatives to be calculated numerically.

Let the equilibrium conditions for the model of interest be given by 1.3 on page 13

$$G(x,\sigma) := \mathbb{E}F(x, g(x,\sigma), h(x,\sigma) + \sigma\eta z', g(h(x,\sigma) + \sigma\eta z', \sigma), \sigma) = 0$$

for all x, σ and assume $G(x,\sigma)$ is Fréchet differentiable with respect to x, σ.

Take the derivative with respect to x (evaluated at $(x^*, x^*, y^*, y^*, 0)$) to obtain

$$F_x + F_{x'}h_x + F_y g_x + F_{y'} g_x h_x = 0$$

In matrix form

$$\begin{bmatrix} F_{x'} & F_{y'} \end{bmatrix} \begin{bmatrix} I & 0 \\ 0 & g_x \end{bmatrix} \begin{bmatrix} h_x \\ h_x \end{bmatrix} = -\begin{bmatrix} F_x & F_y \end{bmatrix} \begin{bmatrix} I \\ g_x \end{bmatrix} \tag{1.5}$$

Define $A = \begin{bmatrix} F_{x'} & F_{y'} \end{bmatrix}$, $B = -\begin{bmatrix} F_x & F_y \end{bmatrix}$ mapping $\mathcal{H}_1 := \mathcal{H}_x \times \mathcal{H}_y \to \mathcal{H}_2$.

We seek to partially characterize the policy operators $h(x,\sigma)$ and $g(x,\sigma)$ by solving for their first derivatives with respect to the 'predetermined' state variable x, h_x and g_x. Written as

$$A \begin{bmatrix} I & 0 \\ 0 & g_x \end{bmatrix} \begin{bmatrix} h_x \\ h_x \end{bmatrix} = B \begin{bmatrix} I \\ g_x \end{bmatrix}$$

this can be seen as an equation in terms of a pair of linear operators (B, A) which may be solved in terms of a joint decomposition of the pair. In general, multiple solutions to this system are possible: however, additional considerations provide some constraint as to the nature of acceptable solutions. In particular, conditions such as transversality conditions in optimization and No Ponzi Game conditions often rule out equilibria in which (some) state variables explode. More generally, a local solution method is attractive largely to the extent that the system remains with high

probability in a neighborhood of the state around which the linearization applies.

For these reasons, we seek a solution to these equilibrium conditions which also induces stable, or stationary, dynamics. For finite-dimensional deterministic dynamical systems, sufficient conditions for the local stability around the steady state may be characterized by the eigenvalues of the linearized transition rule: in discrete time, eigenvalues less than one in modulus imply stability, in continuous time, eigenvalues must have real part less than 0. For infinite-dimensional dynamical systems, analogous conditions apply (see Gohberg *et al.* (1990, Ch. IV.3)). For rational expectations models characterized in terms of expectations, dynamics of state variables may be characterized not only by past values, but also by expectations of future values, and, in particular, certain variables may be allowed to 'jump,' which is to say that in response to a stochastic change in the current state, some variables may change discontinuously in order to satisfy the equilibrium conditions. As a result, stability conditions for this class of models differ from those for deterministic dynamical systems. Most notably, they may exhibit 'saddle-path stability,' in which the system evolves toward the steady state only along a lower-dimensional manifold and so only a (possibly null) subset of eigenvalues satisfy the stability conditions. A stable solution exists in such a case if the jump variables may adjust to ensure that the system stays on this stable manifold.

In the finite-dimensional case, stable solutions to this matrix pair equation may be characterized in terms of the Jordan decomposition of the pair, as in the seminal work of Blanchard & Kahn (1980), or in the case where singularity may be possible or numerical stability is desired, in terms of the generalized Schur decomposition as in Klein (2000). In the infinite-dimensional case, one may, under certain regularity conditions, apply analogues of these decompositions. To provide robustness to singularity and ensure numerical stability, this paper applies an analogue of the generalized Schur decomposition. As such a decomposition appears to be absent from the litera-

ture, Appendix A provides a detailed characterization and a proof of existence under a mild set of regularity conditions. The key idea of the proof is to use the generalized resolvent operator to construct potentially non-orthogonal subspaces on which the operator pair acts corresponding to elements of the spectrum outside and inside the unit circle, and then show that orthogonalizing the subspaces to ensure unitarity of the transform preserves the spectrum.

For the purpose of characterizing the equilibrium conditions of the model, the following description suffices. If (B, A) satisfy the regularity conditions of Lemma(.9) in Appendix A, among which are that (B, A) are bounded operators and that (B, A) are Γ-regular on the unit circle: $\gamma A - B$ has bounded inverse for any complex γ satisfying $|\gamma| = 1$, i.e. the unit circle is in the resolvent set, there exists a decomposition

$$(B, A) = (Q^* T U, Q^* S U)$$

in which U and Q are unitary operators and S and T may be decomposed as

$$(T, S) = \left(\begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix}, \begin{bmatrix} S_{11} & S_{12} \\ 0 & S_{22} \end{bmatrix} \right)$$

conformable with the decomposition $Q = \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix}$ and $U = \begin{bmatrix} U_1 \\ U_2 \end{bmatrix}$ such that the images of U_1^* and U_2^* respectively decompose $\mathcal{H}_1$ into two orthogonal subspaces $\mathcal{H}_{11}$ and $\mathcal{H}_{12}$ and the spectrum of (T_{11}, S_{11}) lies inside the unit circle, so S_{11} has bounded inverse. We may further decompose $U_1,\ U_2$ by considering their actions on $\mathcal{H}_y$ and $\mathcal{H}_x$. Write $U_{11} := U_1 \varphi^X$, $U_{12} := U_1 \varphi^Y$, $U_{21} := U_2 \varphi^X$, $U_{22} := U_2 \varphi^Y$ where $\varphi^X : \mathcal{H}_x \to \mathcal{H}_x \times \{0\} \subseteq \mathcal{H}_1$ and $\varphi^Y : \mathcal{H}_y \to \{0\} \times \mathcal{H}_y \subseteq \mathcal{H}_1$ are imbeddings.

Remark. The assumption of boundedness of the operator pair is not fundamental. Rather, it reflects the choice of space on which the operators are defined. See Kur-

batova (2009) for a way in which to define the domain on which the pair acts so that boundedness holds and the above decomposition may be constructed for operator pairs unbounded with respect to the original choice of space $\mathcal{H}_X$ by restriction to a subspace. The use of potentially unbounded operators may be useful if equilibrium conditions are defined in terms of differential operators, as is common in continuous time versions of the models studied in this paper, as in Achdou *et al.* (2013). In discrete time, the conditions of interest are generally defined in terms of integral equations and so boundedness usually holds without restrictions.

In contrast, Γ-regularity imposes nontrivial restrictions. By requiring existence of a bounded operator with bounded inverse between the two spaces, it requires that $\mathcal{H}_X$ and $\mathcal{H}_Y$ be isomorphic, reflecting the traditional condition that to have a unique set of solutions, it is necessary that there be as many equations as unknowns. In addition to ruling out unit roots, invertibility on the unit circle also rules out long memory behavior. To see this, note that because the resolvent set of an operator pair is open, invertibility must also hold in an open neighborhood of the unit circle, and so it cannot be the case that the spectrum has a limit point in the unit circle, as occurs in certain processes with long memory.

In terms of classes of operators which this assumption excludes, it rules out the presence of a continuous spectrum in the neighborhood of the unit circle. A prominent example of an operator pair with a continuous spectrum is an identity paired with a multiplication operator (which can arise as the functional derivative of a composition operator), i.e. (F, I) with $F[g(x)] = f(x) \cdot g(x)$, which has continuous spectrum taking all values attained by $f(x)$. If $|f(x)|$ has a limit point equal to 1, this operator pair is not Γ-regular. In this case a spectral decomposition can be constructed analytically, and a solution will exist with long memory or unit root behavior (depending on the behavior of $f(x)$ as it approaches 1), but for general models which fail to be Γ-regular with no closed form spectral decomposition, numerical approximations of the

decomposition based on projection methods will be highly unstable.

The generalized Schur decomposition allows us to rewrite our decomposition as

$$Q^* \begin{bmatrix} S_{11} & S_{12} \\ 0 & S_{22} \end{bmatrix} \begin{bmatrix} U_{11} & U_{12} \\ U_{21} & U_{22} \end{bmatrix} \begin{bmatrix} I & 0 \\ 0 & g_x \end{bmatrix} \begin{bmatrix} h_x \\ h_x \end{bmatrix} = Q^* \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix} \begin{bmatrix} U_{11} & U_{12} \\ U_{21} & U_{22} \end{bmatrix} \begin{bmatrix} I \\ g_x \end{bmatrix} \quad (1.6)$$

Unitarity of Q allows it to cancel on both sides, leaving, after simplification,

$$\begin{bmatrix} S_{11} & S_{12} \\ 0 & S_{22} \end{bmatrix} \begin{bmatrix} (U_{11} + U_{12}g_x)h_x \\ (U_{21} + U_{22}g_x)h_x \end{bmatrix} = \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix} \begin{bmatrix} U_{11} + U_{12}g_x \\ U_{21} + U_{22}g_x \end{bmatrix}.$$

To find a stable solution, first solve for g_x, which determines the jump variables in terms of the predetermined variables, and then use this to find the value of h_x. To ensure that the second line holds trivially, it is sufficient to find $g_x : \mathcal{H}_x \to \mathcal{H}_y$ such that

$$U_{21} + U_{22}g_x = 0 \quad (1.7)$$

always. In principle, there may be many solutions, one solution, or no solution to this problem. In the case that $U_{22}U_{22}^*$ has bounded inverse on the space $\mathrm{Im}(U_2)$, at least one solution exists, given by what is referred to in numerical analysis as the 'minimum norm solution' (Golub & van Loan, 1996, Ch. 4) to the linear equation (1.7),

$$g_x = -U_{22}^*(U_{22}U_{22}^*)^{-1}U_{21}. \quad (1.8)$$

It is worth noting how the condition that $U_{22}U_{22}^*$ has bounded inverse relates to the eigenvalue criteria in Blanchard & Kahn (1980) and subsequent rational expectations

solution procedures. The existence of a bounded inverse implies $U_{22}U_{22}^*$ is bijective, and so U_{22} is surjective onto $\text{Im}(U_2)$, which is mapped isometrically to $\mathcal{H}_{12}$ by the continuous and invertible linear transformation U_2^*. Therefore, there exists a linear surjection from $\mathcal{H}_y \to \mathcal{H}_{12}$. In finite dimensions, this requires that the dimension of the space of 'jump variables' y is at least as large as the dimension of the eigenspace corresponding to the 'unstable' generalized eigenvalues. Note however that in infinite dimensions, both of these spaces are infinite dimensional and the spectrum is generally uncountable, so this criterion cannot be expressed in terms of a relationship between the 'number of eigenvalues greater than one' and the 'number of jump variables'.

There is also an analogous condition characterizing uniqueness of the solution. Consider the case in which U_{22} has nontrivial null space. Then if g_x is a solution and $\tilde{g}$ is an operator whose range is a subset of Ker U_{22}, $g_x + \tilde{g}$ is also satisfies $U_{21} + U_{22}(g_x + \tilde{g}) = 0$. Thus, a solution is unique only if U_{22} has trivial null space. Formally, a solution is unique if and only if U_{22} is $\mathcal{H}_y$-complete: $\forall y \in \mathcal{H}_y$, $U_{22}y = 0$ implies $y = 0$. If U_{22} is complete and surjective, then it is bijective, and so, by the bounded inverse theorem has a bounded inverse and so

$$g_x = -{U_{22}}^{-1}U_{21} \tag{1.9}$$

is the unique solution.[7] In finite dimensions, a necessary condition for a linear operator to have trivial null space is that the domain and range spaces are of the same dimension. This therefore corresponds to the case in which the number of jump variables and unstable eigenvalues is exactly equal. Note that while there is a burgeoning literature on the characterization and implications of completeness in econometric models (see, e.g., Andrews (2011)), this is generally in the context of operators which are not surjective and do not have bounded inverse.

[7] If you don't like the axiom of choice, this will still be the unique solution but may or may not be a bounded operator.

If U_{22} is surjective but not complete, the system is said to be underdetermined, and there may be many solutions, of which $-U_{22}^*(U_{22}U_{22}^*)^{-1}U_{21}$, the minimum norm solution is one. In this case, any solution to the system must be equal to $-U_{22}^*(U_{22}U_{22}^*)^{-1}U_{21}$ plus an operator $\tilde{g}$ whose range is in the kernel of U_{22}, which is the complement of the range of U_{22}^*, and so for all $x \in \mathcal{H}_x$ $\| - U_{22}^*(U_{22}U_{22}^*)^{-1}U_{21}x + \tilde{g}x\| = \| - U_{22}^*(U_{22}U_{22}^*)^{-1}U_{21}x\| + \|\tilde{g}x\| \geq \| - U_{22}^*(U_{22}U_{22}^*)^{-1}U_{21}x\|$ hence the description 'minimum norm'. This corresponds to the case in finite dimensions in which there are more jump variables than unstable eigenvalues. In this case, one may calculate a canonical solution with minimum norm, but there are also a continuum of other solutions in which arbitrary terms may be added in the eigenspaces corresponding to the jump variables so long as these terms are sent to zero by the expectation operator. This situation corresponds to the partial identification result when completeness fails in nonparametric instrumental variables estimation described in Santos (2012). While indeterminacy in the finite-dimensional case has received extensive study, for brevity and to avoid technical complications, I will consider only cases in which the solution is unique, in which case U_{22}^{-1} is bounded and well defined.

Given a solution for g_x, the evolution equation for the predetermined variables may be expressed in terms of this solution. Imposing 1.7, the equilibrium conditions hold if $S_{11}(U_{11} + U_{12}g_x)h_x = T_{11}(U_{11} + U_{12}g_x)$. Since S_{11} has bounded inverse by construction, this gives

$$h_x = (U_{11} + U_{12}g_x)^{-1}S_{11}^{-1}T_{11}(U_{11} + U_{12}g_x) \tag{1.10}$$

is a solution so long as $U_{11} + U_{12}g_x$ has bounded inverse. Moreover, this operator is similar to $S_{11}^{-1}T_{11}$ and so has identical spectrum. In particular, by the construction of S_{11} and T_{11} the spectrum of this operator is inside the complex unit circle. So, by Gohberg *et al.* (1990, Thm IV.3.1), the difference equation $x_{t+1} = h_x x_t,\ x_0 = x \in \mathcal{H}_x$

has a unique solution for any given x, given by $x_t = (h_x)^t x$, which converges to 0. Thus, we say that h_x is a stable solution. Moreover, under these conditions, Bosq (2000, Thm 3.1) implies that the Hilbert AR(1) functional linear process given by $x_{t+1} = h_X x_t + \xi_t$, where ξ_t is a $\mathcal{H}_x$ random element uncorrelated over t has a unique covariance stationary solution, and so we are justified in referring to h_x as a stationary solution.

Note that existence of a solution to the operator equation (1.5) is a necessary condition for the existence of a differentiable solution consistent with the equilibrium conditions of the model but is not sufficient. For an overview of high level conditions that might be used to ensure existence of a solution, see Appendix B.

1.4 Algorithm

Having a formula for the functional derivatives of the policy operators in terms of the functional derivatives of the equilibrium conditions is not sufficient to implement the formula unless the components of that formula, defined in terms of the generalized Schur decomposition can be found. While there are some cases where this can be done analytically, these require a high degree of structure to be imposed, often requiring, for example, the model to take a partial equilibrium structure where aggregate variables are taken as exogenous, or the opposite, require individual decisions not to depend on the aggregate state. Beyond these and some other idiosyncratic cases, a numerical procedure is needed to construct the solution. This can be done using projection of the equilibrium conditions onto a finite set of basis functions, so long as the model takes a structure where the approximation error this introduces can be controlled. Conditions under which this holds can often be verified easily, and in particular hold for the economic geography model evaluated in Section 5.

1.4.1 Cases in Which a Known Decomposition Exists

To compute the functional derivatives of the equilibrium policy operators of a rational expectations model with Hilbert-valued states, it is generally necessary to separate the state space into forward looking and backward-looking, or 'unstable' and 'stable' components. In some special cases, these components correspond to known or analytically identifiable state variables. This generally requires that certain derivatives equal 0: a type of exclusion restriction which ensures that backward-looking variables are not influenced by forward looking variables or vice versa. Exclusion restrictions of these sort are prevalent in partial equilibrium models, in which a forward looking decision may be made given a persistent and purely exogenous state variable. For example, if the feedback between population and economic activity were to be removed from the geography model described in Chapter 2, the migration decision problem given an exogenous distribution of wages would fall into this class. Similar exclusion restrictions may also arise in cases where the equilibrium environment and decision problem are carefully tailored so that a persistent backward looking state has no impact on forward looking decisions which do affect the state. A special case of this structure is when decision making is purely myopic, either due to a carefully tailored incentive structure or due to behavioral constraints on the decision makers. Models with these kinds of exclusion restrictions may be described as *triangular*. While the restrictions required to ensure that such a condition holds are often stringent, the computational and analytical tractability that they allow makes them an important special case.

Let us consider two kinds of triangular models, roughly corresponding to the cases described above where forward-looking decisions are not influenced by a persistent state and where a persistent state is not influenced by a forward looking decision. I will call such cases *upper triangular* and *lower triangular*, respectively, for reasons that will become apparent. In the upper triangular case, the partial derivative of

the equilibrium conditions with respect to the predetermined state variables x is 0 in the equations describing the forward looking decision. This can occur in models with myopic decision making, either due to behavioral constraints or due to a structure of preferences, production or technology designed to produce the knife-edge condition that the optimal decision is independent of the state of the world. In the notation from above, in such cases, the derivatives may be decomposed as

$$(B, A) = [-F_x \ -F_y, F_{x'} \ F_{y'}] = \begin{bmatrix} T_{11} & T_{12} & S_{11} & S_{12} \\ 0 & T_{22} & 0 & S_{22} \end{bmatrix},$$

without any (additional) unitary transformation, so $Q = U = I$, and the forward and backward looking state variables may be identified with y and x, respectively. Applying previous results on the derivatives of the policy functions, obtain $g_x = 0$ and $h_x = S_{11}^{-1} T_{11}$. This says that, consistent with the intuition, the forward looking state has 0 derivative with respect to the persistent backward-looking one, and the backward looking state evolves autonomously. This is a locally stable solution if the spectrum of $S_{11}^{-1} T_{11}$ lies within the unit circle.

The lower triangular case occurs when the derivative of the equilibrium conditions with respect to the jump state variable y is 0 in the equations describing the backward looking persistent state variable. This can occur in partial equilibrium or small open economy type settings, in which aggregate states or distributions are determined completely exogenously. In this case, the derivatives are decomposed as

$$(B, A) = [-F_x \ -F_y, F_{x'} \ F_{y'}] = \begin{bmatrix} T_{11} & 0 & S_{11} & 0 \\ T_{21} & T_{22} & S_{21} & S_{22} \end{bmatrix},$$

without any unitary transform, so again y and x may be identified with forward and backward-looking state variables. In this case, a slightly different calculation may be

applied to obtain the derivatives of the policy operators. Applying the condition

$$\begin{bmatrix} S_{11} & 0 \\ S_{21} & S_{22} \end{bmatrix} \begin{bmatrix} h_x \\ g_x h_x \end{bmatrix} = \begin{bmatrix} T_{11} & 0 \\ T_{21} & T_{22} \end{bmatrix} \begin{bmatrix} I \\ g_x \end{bmatrix}$$

obtain the solution $h_x = S_{11}^{-1} T_{11}$, which is to say that x evolves according to a law of motion which does not depend on the forward looking jump state y. The derivative g_x of the policy function giving y in terms of x, satisfies the equation

$$T_{22} g_x = S_{21} h_x - T_{21} + S_{22} g_x h_x$$

which may either be expressed recursively as an infinite sum, or by treating g_x as an element of the Banach space of bounded linear operators, in terms the inverse of the linear operator $T_{22}[.] - S_{22}[.] S_{11}^{-1} T_{11}$, if it exists, as $g_x = (T_{22}[.] - S_{22}[.] S_{11}^{-1} T_{11})^{-1} (S_{21} S_{11}^{-1} T_{11} - T_{21})$. [8] While neither of these formulas is particularly straightforward to apply, often inverses may be computable in closed form, allowing simple evaluation of the effect of a state variable of interest on an intertemporal decision problem.

1.4.2 Numerical Evaluation by Projection

In practice, one often finds that the restrictions required to ensure that a system takes triangular form are not economically sensible. Especially in general equilibrium problems, forward looking decisions both influence and are influenced by the evolution of persistent states. In such cases, it becomes necessary to apply a method which can separate the forward and backward looking subspaces under general conditions. Unfortunately, closed form solutions are rarely available for the generalized Schur decomposition of systems of operator equations, and one must instead turn to

[8] If we may use the familiar notation $\mathrm{vec}(g_x)$ to denote the map from the space of operators to an isomorphic Banach space, viewed as a vector space, we may write this formula suggestively as $\mathrm{vec}(g_x) = (I \otimes T_{22} - (S_{11}^{-1} T_{11})^* \otimes S_{22})^{-1} \mathrm{vec}(S_{21} S_{11}^{-1} T_{11} - T_{21})$, which gives the solution in terms of the finite-dimensional vec operator and the Kronecker product when $\mathcal{H}_x$ and $\mathcal{H}_y$ are finite dimensional.

numerics. For an algorithm to be useful, it must take data which are computable from representations of the derivatives of the equilibrium conditions and output an approximate decomposition. This suggests application of methods based on sampling, where the derivative operators are approximated by finite-dimensional objects to which a decomposition may be applied numerically.

A particularly simple way to perform this approximation is to approximate any infinite-dimensional Hilbert space by an increasing sequence of subspaces, possibly spanned by a standard set of basis functions. On such spaces, the derivative operators of interest are finite-dimensional matrices. As the number of basis functions grows, representation of any function in the space becomes increasingly accurate, and one may hope that at a sufficient level of detail, the finite-dimensional system accurately approximates the infinite-dimensional one. If this is the case, it may be possible to simply apply solution algorithms for finite-dimensional linear rational expectations algorithms to produce finite-dimensional approximations of the policy functions.

While intuitively appealing, there is an important step missing in the above logic. In order for the finite-dimensional solutions to be accurate, at least asymptotically, it is necessary that when the input of the finite-dimensional rational expectation algorithm is sufficiently close to the truth, that the output also be close: the solution must be continuous with respect to some topology. It turns out that the difficult step to show here is the continuity of the generalized Schur decomposition. While the generalized Schur decomposition is known to be stable in finite dimensions (Golub & Van Loan, 1996), only limited results are available for the infinite-dimensional case. In infinite dimensions, different topologies are not equivalent, and so one must choose a topology with respect to which the approximation converges. The results of Stewart (1973) ensuring continuity of the generalized Schur subspaces and the Rayleigh components ((S_{ij}, T_{ij}) in previous notation) apply in infinite-dimensional spaces, demonstrating, under certain conditions, the continuity of these objects with

respect to the Hilbert-Schmidt norm. In infinite dimensions, the use of this norm imposes unduly strong summability conditions on most operators of interest, to the point that many operators used in practice, such as the identity, do not have finite norm. As a result, in the appendix, I have demonstrated a generalization of this result to the operator norm. If a sampling procedure converges to the true derivatives in operator norm, the generalized Schur decomposition will also converge in the same norm.

While reassuring, continuity in operator norm is in fact of limited applicability without some important auxiliary hypotheses. In particular, it is known that a finite-dimensional matrix may approximate an infinite-dimensional operator in operator norm only if that operator is compact. This presents something of a difficulty, as essentially no economic models with function-valued states have derivatives which are compact operators. However, there exists a limited but far from trivial subclass of models in which it is nevertheless possible to construct the generalized Schur decomposition of a set of operators which consistently approximates the true equilibrium derivatives in operator norm, and so to which the continuity result applies. I refer to models which satisfy this condition as *asymptotically diagonal*.

Definition 1.3. The operator pair (B, A) is *asymptotically diagonal* if there exists a known linear isometry such that $\mathcal{H}_1$ is isometrically isomorphic to $\mathcal{H}_2$, and the representation of the operator pair with respect to this isometry (which will also be denoted (B, A)) satisfies the decomposition $(B, A) = (B_I, A_I) + (B_C, A_C)$ such that B_C and A_C are compact and there exist known finite partitions of $\mathcal{H}_x, \mathcal{H}_y \subset \mathcal{H}_1 \cong \mathcal{H}_2$ into orthogonal subspaces $\{\mathcal{H}_j\}_{j=1}^J$ conformable with the partition into $\mathcal{H}_x$ and $\mathcal{H}_y$, usually corresponding to variables making up X and Y, such that for each pair $(i, j) \in \{1 \dots J\}^2$, $A_{I_{ij}} := \text{Proj}_{\mathcal{H}_i} A_I \text{Proj}_{\mathcal{H}_j}$ and $B_{I_{ij}} := \text{Proj}_{\mathcal{H}_i} B_I \text{Proj}_{\mathcal{H}_j}$ satisfy $A_{I_{ij}}$ and $B_{I_{ij}}$ are each either equal either to the zero operator or to a scalar multiple of the identity I_{ij}, where I_{ij} is defined for $i = j$ as the identity operator on $\mathcal{H}_i$ and for

$i \neq j$ is defined as the identity from $\mathcal{H}_j$ to $\mathcal{H}_i$ if $\mathcal{H}_i \cong \mathcal{H}_j$.

Informally, this statement says that asymptotically diagonal systems can be broken up into a compact part and a part for which all subcomponents are equal to the identity. The typical form for an asymptotically diagonal operator pair is a set of square block operators acting on a space of J functions, where each block contains an identity operator, a compact operator, or a sum of a compact operator and an identity operator. For example if $J = 2$, (B, A) may take the form

$$\left(\begin{bmatrix} c_1 I_{11} + C_1 & c_2 I_{12} + C_2 \\ c_3 I_{21} + C_3 & c_4 I_{22} + C_4 \end{bmatrix}, \begin{bmatrix} c_5 I_{11} + C_5 & c_6 I_{12} + C_6 \\ c_7 I_{21} + C_7 & c_8 I_{22} + C_8 \end{bmatrix} \right) \tag{1.11}$$

where c_1 through c_8 are real scalars (possibly 0) and C_1 through C_8 are compact operators, for example integral operators of the form $\int K(x, z)[f(z)]dz$ for some bounded smooth function $K(x, z) : \ [0, 1)^2 \to \mathbb{R}^1$ in the case where $\mathcal{H}_j$ is $L^2[0, 1)$. Here (B_C, A_C) collects the C components, and (B_I, A_I) collects the cI components.

Asymptotic diagonality ensures that the model has a tractable form 'up to a compact perturbation.' In particular, it can be seen that (B_I, A_I) is block diagonal with respect to any orthonormal basis of $\mathcal{H}_1$ conformable with the partition into subspaces $\{\mathcal{H}_j\}_{j=1}^J$ with blocks which are J-dimensional square pencils which are, importantly, all identical. For example, for a pair in the form of (1.11), for any orthonormal basis $\{\phi_{i1}\}_{i=1}^\infty$ of $\mathcal{H}_{j=1}$, which must have a corresponding basis $\{\phi_{i2}\}_{i=1}^\infty$ for $\mathcal{H}_{j=2}$ if $\mathcal{H}_{j=1} \cong \mathcal{H}_{j=2}$, (if not, the off-diagonal components c_2, c_3, c_6, and c_7 must all be 0, as no identity can be defined), the action of (B_I, A_I) on the coefficients corresponding to functions (ϕ_{i1}, ϕ_{i2}) is given by the pair of 2×2 matrices

$$\left(\begin{bmatrix} c_1 & c_2 \\ c_3 & c_4 \end{bmatrix}, \begin{bmatrix} c_5 & c_6 \\ c_7 & c_8 \end{bmatrix} \right)$$

for *any* $i = 1 \ldots \infty$. This provides a representation of (B_I, A_I) as block diagonal with respect to the the orthonormal basis $\{\{e_{ij}\}_{i=1}^{\infty}\}_{j=1}^{2}$ of $\mathcal{H} = \mathcal{H}_{j=1} \times \mathcal{H}_{j=2}$ with $e_{i1} = (\phi_{i1,}0)$, $e_{i2} = (0, \phi_{i2})$, with blocks corresponding to pairs identified by common index i. As a result, to construct the generalized Schur decomposition of (B_I, A_I), it suffices to calculate a single J-dimensional decomposition of the matrix pencil representing any particular block and to concatenate the identical and orthogonal blocks.

Generally speaking, the isometry condition will be fulfilled by any model which uniquely determines an equilibrium, as it generalizes the familiar requirement that a model have an identical number of equations and unknowns, so the space into which the equilibrium conditions map will generally have a canonical isomorphism to the space of unknown states. This holds similarly for the J subspaces, which usually correspond to interpretable variables in the context of the model, with isomorphisms between spaces of variables likewise defined canonically. For example, in the example model of economic geography in Chapter 2, the distribution of wages and the distribution of amenities may be defined as functions on the same space defined in the same units (such as dollars, or utils).

The use of a restriction of this kind is that identity components are common components of the derivative operators of many models, because many conditions take the form of defining a variable or assigning it a value, but are not compact, and so cannot be approximated directly by finite-dimensional approximations. The remainder of (B_I, A_I) after projection onto any subspace does not go to 0, but because it takes a tractable diagonal form, it is known. In contrast, for the compact component, the remainder when projecting onto an increasing sequence of subspaces does go to 0 and so is asymptotically negligible. By combining these two components, it is possible to use a finite-dimensional projection to approximate the operator pencil on a finite-dimensional subspace and leave a remainder on the orthogonal complement space a

which is known up to an asymptotically negligible perturbation. In this way, one can use a finite set of computations to compute a generalized Schur decomposition corresponding to an operator pencil which is close in norm to the true infinite-dimensional pencil, and so by the continuity in norm of the decomposition, yields a decomposition which is close to the true one.

The disadvantage of such a restriction is that requiring all components of the derivatives to either be compact or to be composed of identity operators restricts the functional forms of allowable models, potentially in ways which rule out economically meaningful effects. For example, in a model with a distribution of characteristics which evolve independently across individuals driven by a Markov process, the distribution is a state variable and its evolution is described by an adjoint Markov operator. When the conditional density of the process given any initial state is sufficiently smooth, the operator mapping the density today to the density tomorrow will be compact and the density tomorrow enters via an identity. However, when there is a point mass in the conditional density, the transition operator need not be compact. Point masses can describe inertia, such as that induced by fixed costs (Stokey, 2008) or indexation, or mass movement along a discontinuous path. Similarly, decision problems where the object chosen is a function, ubiquitous in economics in the form of best-response policies, can yield first order conditions in which the function which is a choice variable enters into a nonlinear utility function, resulting in a functional derivative which is a multiplication operator, which may be noncompact. In some cases, it may be possible to transform the condition into one where the noncompact operator is an identity by applying its inverse to the equation, but this can eliminate only one non-identity operator from the equation. In the case where this decision problem over functions faces a state variable which is also a function, as in games or contracting problems, or government choosing nonlinear policies over a continuum of agents, goods, or locations, there may be multiple nonlinear operators in the decision

problem which may prevent reducing to an asymptotically diagonal form.

When possible, applying an invertible transformation to both sides of an equilibrium equation can ensure that the asymptotically diagonal form holds without making changes to the model itself. In other cases, it may be possible to construct a modified model which is close to the original but which satisfies the condition that its derivatives are asymptotically diagonal. For example, if compactness fails due to a law of motion with discrete jumps to a fixed value, creating a discontinuity in the distribution at that value, the discontinuity may be removed if the discrete jump is accompanied by a small amount of continuously distributed noise, thereby smoothing out the conditional distribution, though the shape of the resulting distribution may be very close if the noise is small enough. Similarly, discrete actions induced by hard constraints can be made to vary continuously by replacing hard constraints with smooth but sharply growing penalties which induce similar but smooth behavior. These sorts of smoothing transformations are commonly used to employ numerical methods which rely on smoothness (see Den Haan (2010) for commentary), though it should be noted that these changes in the model may not be innocuous. While the resulting behavior at the individual level may be extremely similar, the resulting operator describing the law of motion for the distribution across individuals, which is now compact, as desired, must be far away from the true operator for some input functions. As a result, this approach does not guarantee that the resulting aggregate behavior will be close. Instead, it provides a solution to a different model with similar individual level behavior. However, if the additional noise or smoothing of the constraint is empirically justified, this is not necessarily a concern. For example, the extreme value heterogeneity in location preference in the model of migration decisions not only ensures a smooth law of motion for the population distribution, it also reflects the believable feature that there is idiosyncratic preference heterogeneity which ensures that individuals do not all move to the same place.

When the derivatives of the equilibrium conditions of a model are asymptotically diagonal (and a unique stable equilibrium exists local to the steady state), computation of the first order approximation of the policy operators is both straightforward and computationally fast, in the sense that a consistent approximation can be computed to any desired precision in time polynomial in the number of basis functions used in the approximation. The procedure consists of projecting the equilibrium derivative operators onto a finite-dimensional orthogonal subspace, computing the policy operators on that subspace by applying directly a standard first order rational expectations solution algorithm for finite-dimensional models, and computing the policy operator on the orthogonal complement of that subspace analytically using (B_I, A_I). The operator norm precision of the resulting approximation is then asymptotically of no higher order than the operator norm error in the projection approximation of (B_C, A_C). While compactness alone ensures only that this projection error goes to 0 as the number of basis functions increases, when the compact component takes the form of integral operators $\int K(x,z)[f(z)]dz$, mild smoothness conditions (or other limited complexity conditions) on the kernel, can be used to ensure a rate of convergence. Moreover, in the case where projections cannot be calculated analytically, for example because the kernel function can only be accessed by point evaluation and so integrals must be computed approximately by quadrature, similar smoothness conditions ensure that the additional error induced is controllable.

As a very wide variety of schemes for approximating such an operator may be applied, I first provide a general purpose bound in terms of operator norm error, then offer an example of a set of conditions and an approximation scheme which ensure that this bound goes to 0. Specifically, under a Hölder continuity and compact support condition on the kernel, I demonstrate convergence rates for a Coiflet representation and 1-point quadrature scheme derived from Beylkin *et al.* (1991) which justify the high speed and accuracy that this method has exhibited in numerical experimenta-

tion relative to a wide variety of other basis function choices and quadrature schemes. Results in Beylkin *et al.* (1991) may easily be used to extend to the case of singular kernels and other wavelet classes. One particular alternative which may be attractive in some cases is to use an operator calibrated to data. In this case, plugging in any matrix-valued operator norm consistent estimator (as in Guillas (2001) for functional autoregressions, or Park & Qian (2012) or Benatia *et al.* (2015) for functional regression) produces a consistent estimator of the policy functions.

Formally, sufficient conditions for consistent approximation are given by

(i) (B, A) $\mathcal{H}_1 \to \mathcal{H}_2$ is an asymptotically diagonal pair of bounded operators, Γ-regular with respect to closed Cauchy curve Γ (i.e., per Definition (.1), $\lambda A - B$ is invertible for all λ in a closed curve $\Gamma \subset \mathbb{C}_\infty$ separating the extended complex plane into an interior and exterior subsets), with generalized Schur decomposition with respect to Γ given by

$$(B, A) = [Q_1^*, Q_2^*] \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix}, \begin{bmatrix} S_{11} & S_{12} \\ 0 & S_{22} \end{bmatrix} \begin{bmatrix} U_1 \\ U_2 \end{bmatrix}$$

(ii) $\text{dif}(\begin{matrix} T_{11} & S_{11} \\ T_{22} & S_{22} \end{matrix}) > 0$, where the dif operator is defined in 8 on page 181 in Appendix B as a measure of continuity of the generalized Schur decomposition with respect to perturbations (iii) $U_{22} = U_2 \varphi^X$ is invertible

These conditions on the derivatives of the model are not entirely general but apply to fairly broad classes of models. As mentioned before, asymptotic diagonality rules out certain classes of models which display excessive 'frequency mixing'. The general property of operators which this rules out is a transfer of energy between frequencies which fails to dissipate as frequency increases to infinity. In these cases, input functions with a high degree of regularity are passed to outputs which may be irregular, impeding the ability to represent the system uniformly in time with

respect to classes of regular functions which can be well approximated by standard function approximation. This transfer of energy to higher and higher frequencies is commonly described in models of physical systems as an aspect of (weak) turbulence, and generally requires numerical methods different from those described here. Γ-regularity ensures that forward and backward looking components of the system can be distinguished. In the typical case where Γ is the complex unit circle, it rules out unit roots and continuous spectra around unity, and so imposes some restriction on the time series properties of the systems which can be analyzed by this method.

Condition (ii) on the dif operator of the pair similarly imposes that the forward and backward looking components are well-separated, ensuring their continuity with respect to small perturbations in the operators: see Appendix A.2 for an exact definition and further discussion. Heuristically speaking, the dif constant is a measure of the separation between the forward and backward subspaces which depends on the spectral gap between the subspaces and the degree of nonnormality (or deviation from a diagonalizable pair) of the operator pair. In the case where the operator pair (B, A) is diagonalizable, it is equal to the minimum distance between the spectra of (T_{11}, S_{11}) and (T_{22}, S_{22}) and so positivity is implied by Γ-regularity. Γ-regularity is also sufficient in the case that (B, A) is finite dimensional (see Stewart & Sun (1990) Thm VI.1.11) or in the case in which either B or A is invertible, in which case it follows from the Sylvester-Rosenblum theorem for operators (Bhatia & Rosenthal, 1997), though the exact size will depend on the degree of nonnormality.[9]

Condition (iii) is necessary for existence and uniqueness of derivatives of a policy function which are consistent with the equilibrium conditions: it ensures that there is a correct model to be approximated.

To provide a consistent approximation, it is necessary to choose a sequence of

[9]I conjecture that there may exist an analogue of the Sylvester-Rosenblum theorem for the operator defining the dif constant, in which case condition (ii) is entirely redundant. However, the method of proof does not straightforwardly generalize.

finite-dimensional orthogonal subspaces which converge to $\mathcal{H}$. Generally these will be defined as the closed linear span of an increasing sequence of functions in a set of complete orthonormal bases of $\{\mathcal{H}_j\}_{j=1}^J$, though orthonormality is mainly a computational and notational convenience. As one often does not have access to an exact projection, it is sufficient to request a consistent approximation to one instead. For consistency, we require approximations satisfying the following properties

(i) Let $\{\pi_j^{K_j}\}_{j=1}^J$ be J orthogonal projections onto K_j-dimensional orthogonal subspaces of $\{\mathcal{H}_j\}_{j=1}^J$ respectively such that Im $\pi_i^{K_i} \cong$ Im $\pi_j^{K_j}$ if $\mathcal{H}_i \cong \mathcal{H}_j$ (i.e., π_j^K and $\pi_i^{K_i}$ map to subspaces which are identified of elements of the partition which are themselves identified), and let $\pi^K = \sum_{j=1}^J \pi_j^{K_j}$ project onto the $K = \sum_{j=1}^J K_j$ -dimensional union of these subspaces. Define $(B^K, A^K) := \pi^K(B, A)\pi^K$, and $(B_C^K, A_C^K) := \pi^K(B_C, A_C)\pi^K$. Let

$$\max\{\left\|B_C^K - B_C\right\|_{op}, \left\|A_C^K - A_C\right\|_{op}\} \leq \eta_K$$

for some sequence η_K decreasing to 0 as $K \to \infty$.

(ii) Let $(\tilde{B}^K, \tilde{A}^K)$ be a sequence of matrix approximations of (B^K, A^K) on a Euclidean space isomorphic to Im π^K satisfying

$$\max\{\left\|\tilde{B}^K - B^K\right\|_{op}, \left\|\tilde{A}^K - A^K\right\|_{op}\} \leq \zeta_K$$

for some sequence ζ_K decreasing to 0 as $K \to \infty$.

In practice, as the J subspaces of $\mathcal{H}$ represent distinct functions used as state variables (for example, a value function and a distribution over agents), these approximations are given by first choosing an appropriate complete series basis for each function of interest and representing each function with respect to an increasing number of terms in that series. The numerical representation of the operators with respect to the series $(\tilde{B}^K, \tilde{A}^K)$ is then calculated by interpolation, quadrature, exact

sampling in some special cases, or estimation. Note that we require only consistency of the projections over (B_C, A_C). Both on and off the projected space (B_I, A_I) has exact representation as a set of scalar multiples of identity matrices on Im π^K and as identity operators on the orthogonal complement of that space.

Given a choice of spaces onto which to project and a consistent approximation of the projected operators, we can define our approximate operators by calculating policy operators $\tilde{g}_x^K$, $\tilde{h}_x^K$ and $g_x^{K\perp}, h_x^{K\perp}$ from (B, A) separately on Im π^K and Ker π^K, respectively and composing them. We may define these operators as follows.

Denote the generalized Schur decomposition with respect to Γ of the finite-dimensional matrix representation of $(\tilde{B}^K, \tilde{A}^K)$ as

$$[\tilde{Q}_1^{*K}, \tilde{Q}_2^{*K}] \begin{bmatrix} \tilde{T}_{11}^K & \tilde{T}_{12}^K \\ 0 & \tilde{T}_{22}^K \end{bmatrix}, \begin{bmatrix} \tilde{S}_{11}^K & \tilde{S}_{12}^K \\ 0 & \tilde{S}_{22}^K \end{bmatrix} \begin{bmatrix} \tilde{U}_{11}^K & \tilde{U}_{12}^K \\ \tilde{U}_{21}^K & \tilde{U}_{22}^K \end{bmatrix}.$$

Note that because this is a finite-dimensional matrix pair, this may be calculated in $O(K^3)$ time by the QZ algorithm: see Golub & Van Loan (1996). Applying the formulas for the policy operators to this restricted space, define $\tilde{g}_x^K = -(\tilde{U}_{22}^K)^{-1}\tilde{U}_{21}^K$, $\tilde{h}_x^K = (\tilde{U}_{11}^K + \tilde{U}_{12}^K\tilde{g}_X^K)^{-1}(\tilde{S}_{11}^K)^{-1}\tilde{T}_{11}^K(\tilde{U}_{11}^K + \tilde{U}_{12}^K\tilde{g}_X^K)$. These define an approximation of g_x and h_x respectively on the space Im π^K.

As the restriction of the policy function to this space need not, in general, consistently approximate the policy functions on $\mathcal{H}$ as a whole, we supplement by an approximation on the orthogonal complement space, $g_x^{K\perp}, h_x^{K\perp}$. To do this, we approximate by considering only (B_I, A_I) on this space: this is a reasonable approximation because for K large enough, the contribution of (B_C, A_C) on the remainder becomes negligible. Consider a set of complete orthonormal bases of $\mathcal{H}_j$, $\{e_{ij}\}_{j=1}^J$, $i = 1 \ldots \infty$, where e_{sj} and e_{tk} are identified if $s = t$ and $\mathcal{H}_j \cong \mathcal{H}_k$. Then, by construction, for all i, (B_I, A_I) maps the closure of their span $\overline{\text{Span}}\{e_{ij}\}_{j=1}^J$ to itself and moreover, the representation of this map is identical for all i. Informally, (B_I, A_I) is ($\mathcal{H}$-equivalent

to by Parseval's identity) a block diagonal matrix pair over this complete orthonormal basis with identical $J \times J$ blocks. Further, because an identity matrix has identity representation with respect to any choice of basis, we may choose a basis such that $\{e_{ij}\}_{j=1}^{J},\ i = K+1 \ldots \infty$ are a complete orthonormal basis of Ker π^K (these may or may not be the remaining elements of an orthonormal basis the projection onto the span of which defines π^K, though this representation is convenient). By the orthogonality of the blocks, it is sufficient to define the policy function separately on each block. This can be done by applying the solution formula to any J-dimensional block i, regarded as a pair of $J \times J$ matrices, (B_I^i, A_I^i). These have generalized Schur decomposition

$$(B_I^i, A_I^i) = [Q_1^{i*}, Q_2^{i*}] \begin{bmatrix} T_{11}^i & T_{12}^i \\ 0 & T_{22}^i \end{bmatrix}, \begin{bmatrix} S_{11}^i & S_{12}^i \\ 0 & S_{22}^i \end{bmatrix} \begin{bmatrix} U_{11}^i & U_{12}^i \\ U_{21}^i & U_{22}^i \end{bmatrix}$$

on each block i, where U_{11}^i and U_{21}^i acts on the J_1 elements contained in $\mathcal{H}_x$ and U_{12}^i and U_{22}^i act on the $J - J_1$ elements contained in $\mathcal{H}_y$. The corresponding block of the policy operators are given by $g_x^i = -(U_{22}^i)^{-1} U_{21}^i$, $h_x^i = (U_{11}^i + U_{12}^i g_x^i)^{-1} (S_{11}^i)^{-1} T_{11}^i (U_{11}^i + U_{12}^i g_x^i)$. To define our approximation on the orthogonal complement of Im π^K, we simply concatenate the blocks, giving sequential representations

$$g_x^{K\perp} = \sum_{i=K+1}^{\infty} \sum_{j=J_1+1}^{J} \sum_{k=1}^{J_1} (g_x^i)_{(j-J_1)k} \langle e_{ik}, [.] \rangle e_{ij}$$

and

$$h_x^{K\perp} = \sum_{i=K+1}^{\infty} \sum_{j=1}^{J_1} \sum_{k=1}^{J_1} (h_x^i)_{jk} \langle e_{ik}, [.] \rangle e_{ij}.$$

Note that since each block is identical, calculation of the policy function needs be performed for only one representative block, with running time dominated by the QZ algorithm, of order $O(J^3)$ typically negligible.

The approximation to the policy operators on $\mathcal{H}$ are given by $g_K := \tilde{g}_x^K + g_x^{K\perp}$ and $h_K := \tilde{h}_x^K + h_x^{K\perp}$. A summary of the steps leading to their construction is provided as Algorithm 1. Under the conditions given, these approximations are consistent in operator norm. We have the result

Theorem 1.1. *Let (B, A) and their approximations $(\tilde{B}^K, \tilde{A}^K)$ satisfy Conditions (1.4.2) and (1.4.2). Then $\|g_K - g_x\|_{op} \to 0$ and $\|h_K - h_x\|_{op} \to 0$ as $K \to \infty$. In particular, there exists some $\bar{K}$ and some constant C such that for $K > \bar{K}$, $\|g_K - g_x\|_{op} \leq C(\zeta_K + \eta_K)$ and $\|h_K - h_x\|_{op} \leq C(\zeta_K + \eta_K)$.*

Proof. See Appendix. □

The idea behind the consistency argument is to show that the generalized Schur decomposition of the combined approximation on and off Im π^K converges in operator norm and then apply perturbation theorems ensuring continuity in operator norm for the Schur projectors and Rayleigh components. Then by applying orthogonality, show that the policy functions corresponding to the generalized Schur decompositions on and off Im π^K are equivalent to the policy functions corresponding to the Schur decomposition of the approximate operator as a whole. The exact constant C and $\bar{K}$ are both decreasing functions of the dif constant of (B, A). While the rate of convergence is unaffected, for highly non-normal operators or those with a small gap between the spectrum of the forward and backward looking components, the constant on the rate may be large.

1.4.3 Implementation: Wavelet Transform

Overall, the computational effort needed to obtain ϵ-close approximations is driven by the rates η_K and ζ_K. If efficient (or exact) evaluation schemes are used, the projection error η_K tends to dominate: this may not be the case if the value of the projection coefficients is determined by estimation, in which case the accuracy of ζ_K is limited

Algorithm 1 Construction of g_K, h_K

Inputs: An equilibrium operator $F(x, y, x', y', \sigma)$ satisfying Condition 1.4.2, Im π^K a K-dimensional subspace satisfying Condition 1.4.2, and $\{\{e_{ij}\}_{j=1}^{J}\}_{i=K+1}^{\infty}$ a conformable orthonormal basis for the orthogonal complement of Im π^K

Output: g_K, h_K approximate functional derivatives of recursive solution with respect to x

1. Compute steady state (x^*, y^*) s.t. $F(x^*, y^*, x^*, y^*, 0) = 0$
2. $(B, A) \leftarrow \left(-\begin{bmatrix} F_x & F_y \end{bmatrix}, \begin{bmatrix} F_{x'} & F_{y'} \end{bmatrix}\right)$ Calculate functional derivatives at steady state
3. Decompose (B, A) into $(B_I, A_I) + (B_C, A_C)$ compact and identity components as per Definition 1.3
4. Construct $(\tilde{B}^K, \tilde{A}^K)$, a K-dimensional approximate projection of (B, A) onto Im π^K, satisfying Condition 1.4.2, using Algorithm 2 or other method
5. Build components of policy operator on Im π^K and Ker π^K
   (a) Build policy operators on Im π^K using approximate projections
      i. $[\tilde{Q}_1^{*K}, \tilde{Q}_2^{*K}] \begin{bmatrix} \tilde{T}_{11}^K & \tilde{T}_{12}^K \\ 0 & \tilde{T}_{22}^K \end{bmatrix}, \begin{bmatrix} \tilde{S}_{11}^K & \tilde{S}_{12}^K \\ 0 & \tilde{S}_{22}^K \end{bmatrix} \begin{bmatrix} \tilde{U}_{11}^K & \tilde{U}_{12}^K \\ \tilde{U}_{21}^K & \tilde{U}_{22}^K \end{bmatrix} \leftarrow QZ(\tilde{B}^K, \tilde{A}^K)$ Apply QZ algorithm to obtain generalized Schur decomposition of $(\tilde{B}^K, \tilde{A}^K)$
      ii. $\tilde{g}_x^K \leftarrow -(\tilde{U}_{22}^K)^{-1}\tilde{U}_{21}^K$, $\tilde{h}_x^K \leftarrow (\tilde{U}_{11}^K + \tilde{U}_{12}^K\tilde{g}_X^K)^{-1}(\tilde{S}_{11}^K)^{-1}\tilde{T}_{11}^K(\tilde{U}_{11}^K + \tilde{U}_{12}^K\tilde{g}_X^K)$
   (b) Build policy operators on Ker π^K by analytical decomposition of (B_I, A_I)
      i. $[(B_I^i, A_I^i)]_{jk} \leftarrow \langle B_I e_{ij}, e_{ik}\rangle, \langle A_I e_{ij}, e_{ik}\rangle \ \forall j, k = 1 \dots J$ Construct (B_I^i, A_I^i) (identical for all i) using $\{e_{ij}\}_{j=1}^J$ for some i
      ii. $[Q_1^{i*}, Q_2^{i*}] \begin{bmatrix} T_{11}^i & T_{12}^i \\ 0 & T_{22}^i \end{bmatrix}, \begin{bmatrix} S_{11}^i & S_{12}^i \\ 0 & S_{22}^i \end{bmatrix} \begin{bmatrix} U_{11}^i & U_{12}^i \\ U_{21}^i & U_{22}^i \end{bmatrix} \leftarrow QZ(B_I^i, A_I^i)$ Apply QZ algorithm to (B_I^i, A_I^i)
      iii. $g_x^i \leftarrow -(U_{22}^i)^{-1}U_{21}^i$, $h_x^i \leftarrow (U_{11}^i + U_{12}^i g_x^i)^{-1}(S_{11}^i)^{-1}T_{11}^i(U_{11}^i + U_{12}^i g_x^i)$ Build policy functions over $\overline{\text{Span}}\{e_{ij}\}_{j=1}^J$
      iv. Add identical components for all $i = K+1 \dots \infty$

$$g_x^{K\perp} = \sum_{i=K+1}^{\infty} \sum_{j=J_1+1}^{J} \sum_{k=1}^{J_1} (g_x^i)_{(j-J_1)k} \langle e_{ik}, [.]\rangle e_{ij}$$

$$h_x^{K\perp} = \sum_{i=K+1}^{\infty} \sum_{j=1}^{J_1} \sum_{k=1}^{J_1} (h_x^i)_{jk} \langle e_{ik}, [.]\rangle e_{ij}$$

   for J_1 blocks in $\mathcal{H}_x$, $J - J_1$ in $\mathcal{H}_y$.
6. $g_K \leftarrow \tilde{g}_x^K + g_x^{K\perp}$, $h_K \leftarrow \tilde{h}_x^K + h_x^{K\perp}$ Add components

by the quantity of data available. To more precisely quantify the size of these errors, we provide an example of a set of conditions on (B, A), the approximating subspace Im π^K, and the evaluation method for the projections which provides precise rates. In more general situations, where these conditions don't hold or other approximation methods are desired, similar approaches can guarantee the high level conditions of the preceding theorem.

In particular, I cover the example of Coiflet sampling and representation as in Beylkin *et al.* (1991) for Fredholm integral operators with Hölder-continuous kernel. Fredholm integral operators are a canonical example of operators which are given by a compact component and potentially a component given by an identity, and appear frequently in examples. Wavelet sampling methods provide a particularly fast and accurate method for approximating these operators even when the kernel can only be accessed by pointwise evaluation, perhaps because it is a complicated function which has itself been numerically approximated, such as a function of a steady state calculated numerically by fixed point iteration.

(i) Let $\{\mathcal{H}_j\}_{j=1}^J$ be given by the spaces of square integrable periodic functions of dimension d_j with domain normalizable to $[0,1)^{d_j}$, $\mathcal{H}_j = L^2_{per}[0,1)^{d_j}$. Let (B_C, A_C) consist on each block (i,j) of $r = B$ or $r = A$ of integral operators mapping $f(y) \in \mathcal{H}_j$ to $f(x) \in \mathcal{H}_i$ $f(x) = \int_{[0,1)^{d_j}} K_{r,ij}(x,y)[f(y)]dy$ such that for all r, i, j $\sup_{x,y \in [0,1)^{d_i} \times [0,1)^{d_j}} |K_{r,ij}(x,y)| < \infty$ and $K_{r,ij}(x,y) \in \Lambda^{\alpha_{r,ij}}([0,1)^{d_i} \times [0,1)^{d_j})$, the space of $\alpha_{r,ij}$-Hölder continuous periodic functions on $[0,1)^{d_i} \times [0,1)^{d_j}$ for some $\alpha_{r,ij} > 0$.[10]

(ii) Let Im π^{K_j} be the subspace spanned for each j by a tensor product of d_j one-dimensional orthonormal Coiflet wavelet multiresolution analyses with mother wavelet ψ and scaling function ϕ with each bounded, having support which is a compact interval, and a number of vanishing moments greater than or equal to $\min_{r,i,j} \alpha_{r,ij}$. Let

[10] A function f(x) is Hölder continuous on domain I of order $\alpha \in (0,1]$ if $\sup_{x,y \in I} |f(x) - f(y)| \leq K|x-y|^\alpha$ and is Hölder continuous of non-integer order $\alpha > 1$ if it is $\lfloor \alpha \rfloor$ times continuously differentiable with $\lfloor \alpha \rfloor^{th}$ derivatives Hölder continuous of order $\alpha - \lfloor \alpha \rfloor$.

the matrix representation of $(\tilde{B}^K, \tilde{A}^K)$ on this space be given by $\pi^K(B_I, A_I)\pi^K$ plus a matrix where the i, j block is given by the discrete wavelet transforms over rows then columns of the $K_i \times K_j$ matrix whose (s, t) entry is $\frac{1}{\sqrt{K_i K_j}} K_{r,ij}(x_s, y_t)$, where $\{x_s\}_{s=1}^{K_i}$ and $\{y_t\}_{t=1}^{K_j}$ are dyadic grids over $[0, 1)^{d_i}$ and $[0, 1)^{d_j}$ respectively.

Remark. On (i): These assumptions can be slightly relaxed through different choices of wavelet basis. Periodicity is convenient for proofs because it does not require any special treatment of boundaries: it also fits the example model presented. Depending on the problem, this may be relaxed by one of a number of boundary extension methods: see Mallat (2008). Compact support can be replaced by a tail condition by sampling an increasing spatial domain. Boundedness of the kernel can likewise be dispensed with provided the operator remains compact and some knowledge of the singularity is available: Beylkin *et al.* (1991) provides methods and convergence results for many singular integral operators. It is likely that Hölder regularity could be replaced with more general Besov classes which may exhibit less uniform regularity, at the expense of more difficult analysis of the quadrature approximation. In both of these cases, speed of the algorithm may be enhanced by pruning away those basis functions whose inner product with the kernel is below a small threshold. As the wavelet representation is often sparse, one may incur minimal error in the approximation of (B, A) while substantially reducing the size K of the matrix for which one calculates the generalized Schur decomposition, an order K^3 operation which dominates the quadratic time to evaluate and threshold a higher order wavelet representation.

On (ii): As described, the procedure represents each kernel in terms of a tensor product of multiresolution wavelet bases instead of a single multidimensional multiresolution analysis as advocated in Beylkin *et al.* (1991): while such a representation has desirable features for thresholding procedures, a tensor product representation of the operator ensures that functions in the domain and range space are represented in

terms of the original d_j-dimensional orthonormal wavelet basis. For j with $d_j > 1$, either a tensor product wavelet basis or a multidimensional wavelet multiresolution analysis may be used in calculating the basis functions: the space spanned by a finite representation is identical. In practice, the multidimensional MRA is preferred computationally. The moment condition is assumed to hold for the one-dimensional wavelets generating the tensor product or multiresolution basis.

The requirement that both wavelet and scaling function have compact support, α vanishing moments, and generate an orthonormal basis strongly restricts the choice of wavelet class. The use of Coiflets (Beylkin *et al.* , 1991) (or certain mild generalizations, as in Wei (1998), which also maintain these properties) is in fact required to achieve optimal rates via the procedure described. The purpose of this assumption is to ensure that the operator can be represented directly in terms of the discrete wavelet transform of its evaluations at a set of points, effecting a 'one-point quadrature' scheme for the calculation of the coefficients of the representation. For more general classes of wavelets, the use of the discrete wavelet transform of the evaluation points of a smooth function to substitute for the projection onto a wavelet basis results in an error which is of higher order than the error induced by restricting to a projection onto a finite basis.

Other classes of wavelets may be used if the projection is approximated by a multipoint quadrature scheme, as described in Beylkin *et al.* (1991) or Sweldens & Piessens (1994), at the cost of additional preprocessing before applying the discrete wavelet transform. If neither multipoint quadrature nor the use of Coiflets is acceptable, it is also possible to use interpolating wavelets, which do not form an orthogonal basis and result in a more complicated representation of $\pi^K(B_I, A_I)\pi^K$. General considerations regarding wavelet sampling are discussed in Mallat (2008). One case in which specialized classes of wavelets may be necessary is when the domain is not rectangular or is a subset of a non-Euclidean manifold, as may occur with geographic data restricted

to an irregularly shaped geographic unit or on the surface of the Earth. In this case, a variety of alternative bases and sampling methods are available.

The procedure for constructing approximate projections $(\tilde{B}^K, \tilde{A}^K)$ using the Coiflet basis is laid out in Algorithm 2. Under the above conditions, it can easily be shown that one obtains rapid convergence of the approximation algorithm:

Theorem 1.2. *Let (B,A) and $(\tilde{B}^K, \tilde{A}^K)$ satisfy (1.4.2),(1.4.2), and (1.4.3). If $\bar{\alpha} = \min_{r,i,j} \frac{2\alpha_{r,ij}}{d_i+d_j}$ and $\bar{d} = \max_j 2d_j$, there exists $C > 0$ such that $\eta_K = O(J \max_{r,i,j} (K_i K_j)^{-\alpha_{r,ij}/(d_i+d_j)})$ and $\zeta_K = O(C^{\bar{d}} J \max_{r,i,j} (K_i K_j)^{-\alpha_{r,ij}/(d_i+d_j)})$. As a result, operator norm ϵ-approximations of h_x and g_x such that $\|h_K - h_x\|_{op} \leq \epsilon$ and $\|g_K - g_x\|_{op} \leq \epsilon$ can be calculated using a basis of $K = O(J(\frac{JC^{\bar{d}}}{\epsilon})^{\frac{1}{\bar{\alpha}}})$ functions in $O(J^{3+\frac{3}{\bar{\alpha}}} C^{\frac{3\bar{d}}{\bar{\alpha}}} \epsilon^{-\frac{3}{\bar{\alpha}}})$ operations.*

This result shows that a polynomial time approximation scheme is feasible for this class of models. Due to the accurate quadrature properties of compactly supported wavelet multiresolution analyses, the error from projection and the error from quadrature are of the same order in K, up to constants. While a curse of dimensionality exists with respect to the number of variables entering as arguments of the functions used as state variables, the fact that the functions are themselves infinite-dimensional objects does not impede impede approximation. Further, when the operators are reasonably smooth, as measured by Hölder exponent of the integral kernels, the rate of convergence can be quite rapid. If one is interested in the policy operators as a whole, rather than just their derivatives, this approach only provides a first order Taylor expansion. As a result, it provides accurate approximations within a local neighborhood of the nonstochastic steady state. In the case where the policy operators are continuously Fréchet differentiable, operator norm approximation of the first derivative ensures that the approximated Taylor expansions of the operators are ϵ-close to the true operators uniformly over an open neighborhood of this steady state.

Remark. The dependence on J, which in most applications has the interpretation of the number of independent functions which constitute the equilibrium objects (e.g.,

Algorithm 2 Construction of $(\tilde{B}^K, \tilde{A}^K)$ using wavelet quadrature

Inputs: Block operators $(B, A) = (B_I, A_I) + (B_C, A_C)$ s.t. (B_C, A_C) is composed of integral operators $\int_{[0,1)^{d_j}} K_{r,ij}(x, y)[.]dy \; \forall i, j \in 1 \dots J, r \in \{B, A\}$ satisfying Condition 1.4.3(i), $\{K_j\}_{j=1}^J$ number of evaluation points for each block

Output: $(\tilde{B}^K, \tilde{A}^K)$ satisfying Condition 1.4.2

1. $[K_{r,ij}]_{s,t} \leftarrow \frac{1}{\sqrt{K_i K_j}} K_{r,ij}(x_s, y_t)$ for x_s, y_t on evenly spaced grids of size K_i, K_j over $[0,1)^{d_i}$, $[0,1)^{d_j}$ respectively, $\forall i, j, r$. Construct matrices to represent kernels of integral operators
2. $(\tilde{B}_C^K, \tilde{A}_C^K)_{r,ij} \leftarrow (\text{DWT}[(\text{DWT}[K_{r,ij}])^*])^* \; \forall i, j, r$ Construct approximate projection coefficients by discrete wavelet transform of rows then columns of $K_{r,ij}$, using Coiflet wavelets basis satisfying Condition 1.4.3(ii)
3. $(\tilde{B}_I^K, \tilde{A}_I^K) \leftarrow \pi^K (B_I, A_I) \pi^K$ Represent identity operators by $K_i \times K_j$ identity matrices
4. $(\tilde{B}^K, \tilde{A}^K) \leftarrow (\tilde{B}_I^K, \tilde{A}_I^K) + (\tilde{B}_C^K, \tilde{A}_C^K)$ Add components

a value function, a distribution of individual states, and so on) and is usually a fixed feature of the model, will in general be conservative, as it is based on the worst case that all blocks of (B_C, A_C) contain an integral operator and that the difficulty of approximation of each operator, measured by $\frac{2\alpha_{r,ij}}{d_i + d_j}$, is roughly equal. If the row and column corresponding to subspace j for all but a subset S of subspaces do not contain an integral operator or contain only operators which are substantially smoother and so require fewer basis functions to approximate to ϵ accuracy, and only $K = O(S(\frac{JC^{\bar{d}}}{\epsilon})^{\frac{1}{\bar{\alpha}}})$ basis functions will be needed. This may be the case, for example, if one block contains an operator which is substantially harder to approximate than others (due to being higher-dimensional, less smooth, or both), in which case $S = 1$. In most applications, J is fixed and very small, though it could grow, for example, if some components are represented by a functional autoregressive model of high order.

1.5 Evaluation

The above procedures may be applied to construct an approximate linearized solution to the model of trade, migration, and economic geography of Chapter 2. Details of the model and its solution are described fully in that chapter: here I provide only an assessment of the performance of the above algorithms with respect to the linearized solution of a calibrated version of that model. The model consists of a system of nonlinear operator equations in three function-valued variables, $x_1 = \lambda(.),\ x_2 = \nu(.),\ y = V(.)$ each in $L^2_{per}[0,1)$ and the linearized equilibrium conditions may be represented as block operators acting on (x_1, x_2, y) by

$$(B, A) = \left(- \begin{bmatrix} 0 & 0 & I \\ P & 0 & \beta P - \beta P P \\ 0 & \Gamma & 0 \end{bmatrix}, \begin{bmatrix} \frac{d\omega}{d\lambda} & I & \beta P \\ I & 0 & 0 \\ 0 & I & 0 \end{bmatrix} \right)$$

in which P, $\frac{d\omega}{d\lambda}$, and Γ are compact linear integral operators and β is a scalar parameter, all defined explicitly in Chapter 2. It is shown there that there exists a parameterization in which each integral operator is translation invariant, and so has a representation in the class of convolution operators. The representation in terms of compact and identity operators clearly demonstrates that the model is asymptotically diagonal, and under this parameterization, sufficient conditions on structural parameters may be derived such that the model satisfies Condition (1.4.2), and so provides a suitable case for application of the method.

To evaluate the approximation algorithm, several numerical comparisons are performed for the translation invariant parameterization of this model defined in Chapter 2 using two separate choices of subspace for projection π^K. One choice is Coiflet wavelets, implemented via Algorithm (2). Another choice, particularly well suited to this model, is the Fourier basis of trigonometric polynomials. It is demonstrated

in Chapter 2 that not only can the projections be calculated exactly without additional quadrature approximation in this case, the functional derivatives of the policy operators g_x and h_x are also available in semi-closed form,[11] for any bandlimited input function. This generates a near-exact benchmark for the error in the operator approximated by wavelet quadrature, which I compute at different levels of K. Accuracy can be compared for impulse responses to function-valued shocks, as well as for simulations. The Fourier and wavelet methods appear to exhibit a high degree of agreement, whether expressed in squared error norm over the grid points (a proxy for L^2 norm, controlled by the theory) or in maximum norm over grid points (not controlled by the theory). Error is largest for components and at parameter values at which the component of the Fourier coefficients which must be approximated numerically has substantial impact, suggesting that this quadrature error may be a non-negligible factor contributing to the discrepancy between wavelet and semi-closed form representations, but overall the discrepancy primarily measures the effect of the wavelet quadrature and projection as controlled by Theorem (1.2). For the Fourier representations, integer frequencies $\frac{-K}{2}$ to $\frac{K}{2}$ are used for each of $J = 3$ functions $\nu(.)$, $\lambda(.)$, and $V(.)$, giving $3 \times (K + 1)$ basis functions, for symmetry, while for wavelets K grid points are used to represent the scaling function coefficients for each function, with K given by a power of 2.

I evaluate the policy operator by constructing an impulse response to a smooth but spatially localized shock $\varepsilon(.)$ to the exogenously evolving component of the model $\nu(.)$, a scaled Gaussian spike centered at location 0.5, with functional form $\exp(50000(. - 0.5)^2)$. Relative accuracy of the Fourier and wavelet representations of the model for the above shock are measured in Table (1.5.1), for $K = 256$, 512, and 1024, for the maximum error at any grid point over 80 periods of the impulse response. Figures (1.5.1) and (1.5.2) represent the Euclidean norm difference (over an evenly

[11] Minor numerical error is introduced into the computation by numerical calculation of an integral term in the coefficients, see the complete model description for full details.

Table 1.5.1: Numerical IRF Discrepancy, Fourier vs. Wavelet Representations

K	max pointwise, ν_t	max pointwise, λ_t	max pointwise, V_t	Running Time (seconds)
256	0.0107	3.9549e-07	1.9362e-06	11.607549
512	3.4594e-07	5.0737e-08	7.6597e-08	96.187571
1024	8.9301e-11	1.2976e-08	1.9643e-08	376.833220

spaced grid) for each $j = \{1, 2, 3\}$ at each time point between the wavelet and Fourier representations at the different values of K, a proxy for the L^2 norm. Note that even for $K = 512$, the errors are already extremely small, with maximum pointwise error on the order of 10^{-7} or smaller for $\nu_t(x)$ a function with values ranging from 0 to 1, and 10^{-8} for $\lambda_t(x)$ and $V_t(x)$, functions with range of about 0.1. The order of this error decreases significantly for $K = 1024$, both for maximum and squared average error.

The clock time to compute the wavelet solutions, also displayed in Table (1.5.1), is relatively fast and increases roughly in cubic proportion to K, taking under two minutes for $K = 512$, including producing all figures and evaluation metrics, coded in Matlab using the default QZ function on a 2011 Macbook Pro with 2.8 GHz Intel i7 processor and 2 GB RAM. This level of speed and accuracy on a far from state of the art setup suggests that the procedure may be useful in applications where it is applied repeatedly, for example to estimate parameters. The Fourier representation takes only a few seconds for any K, which should be expected as it allows calculating solutions for each frequency separately and so takes time linear in K. This feature is only a result of the special structure of this model and is not likely to generalize.

1.6 Conclusion

The idea that heterogeneity matters for economic outcomes, not only at the individual level but through the set of interdependencies linking behavior at the individual level to the environment faced by others, is a core principle in economics. Function-valued

Figure 1.5.1: Euclidean Discrepancy, K=512

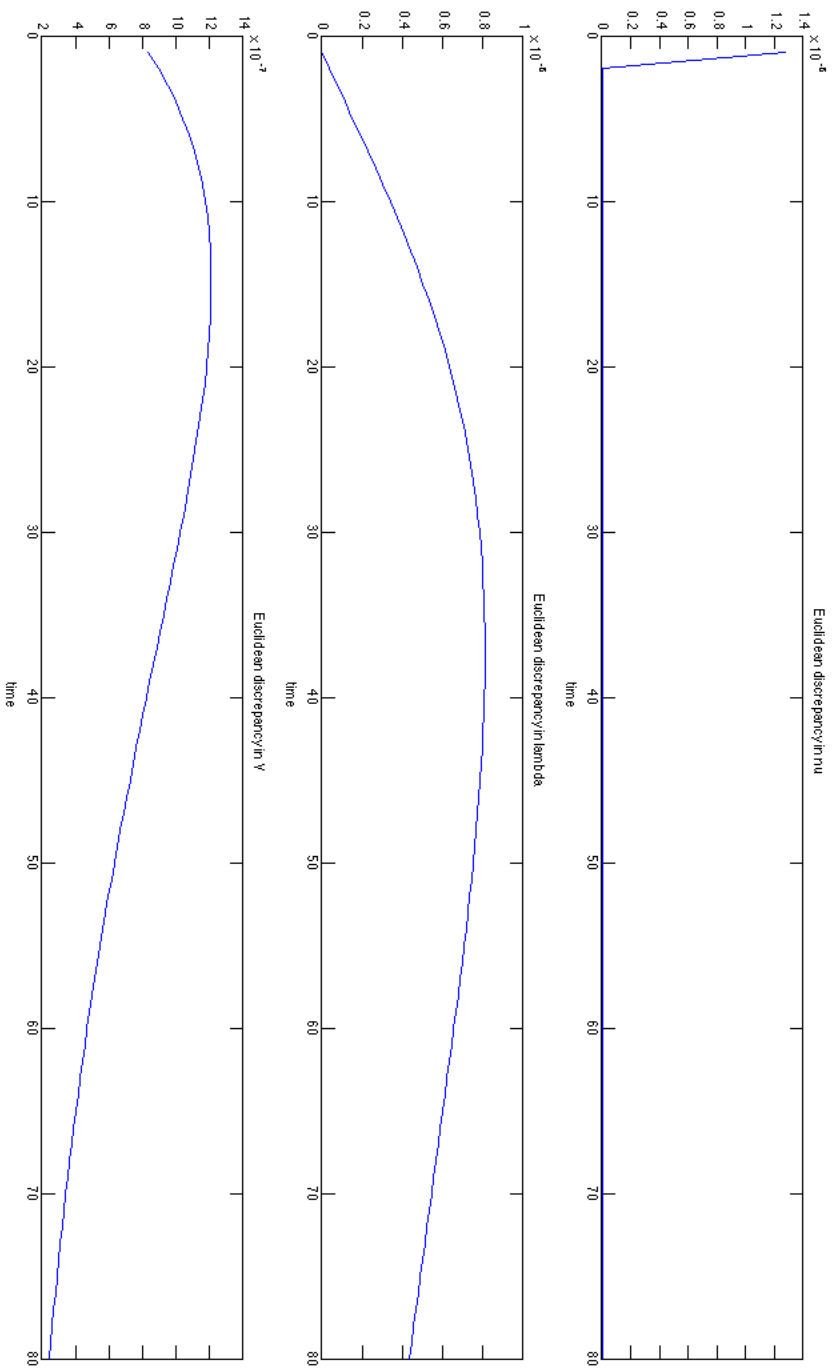

Figure 1.5.2: Euclidean Discrepancy, K=1024

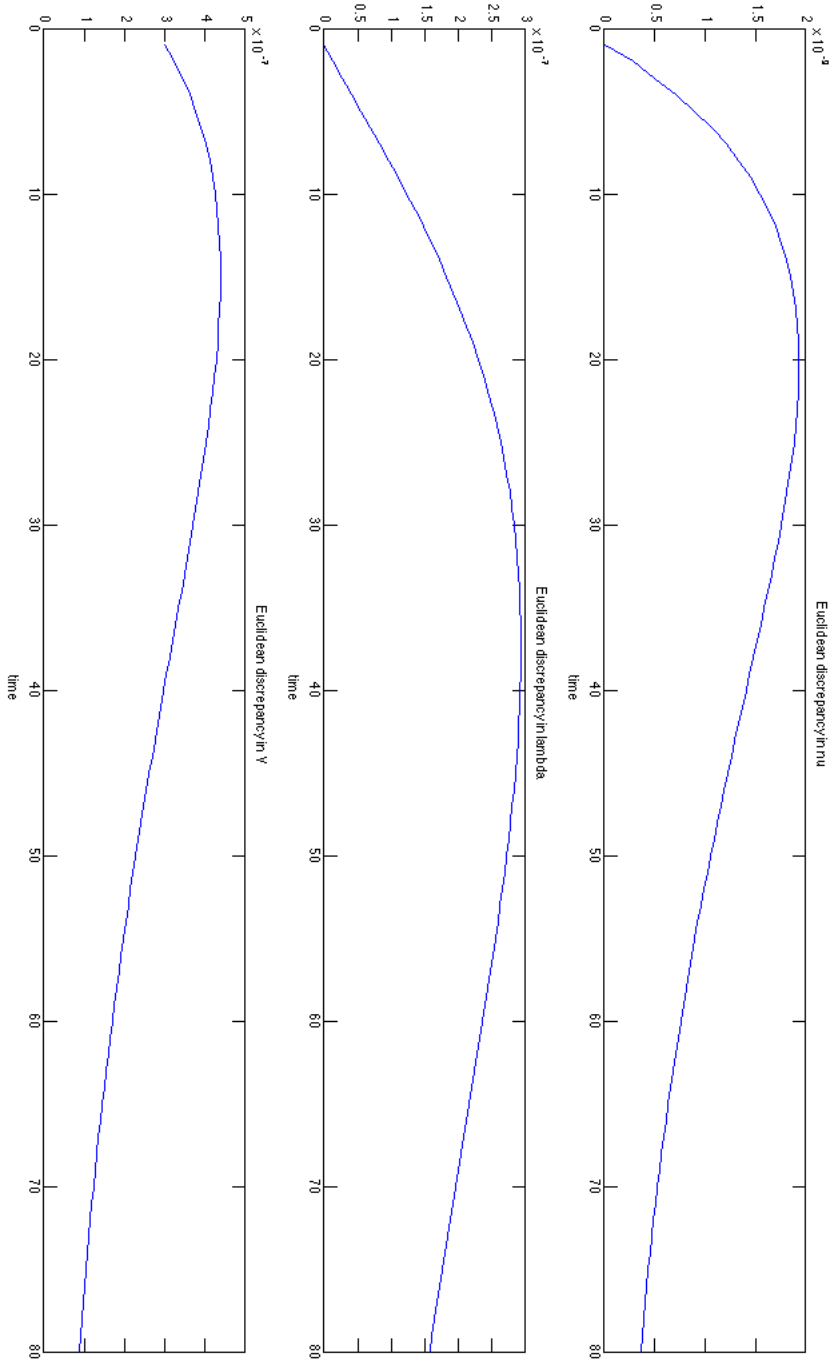

stochastic processes, by describing how patterns of heterogeneity change over time and relate to other variables, provide an analytical framework in which these interdependencies can be modeled and evaluated directly rather than considering only aggregate variables. While describing economic decision making in these environments can be challenging due to the high dimension of the relevant variables, a substantial amount of information can be recovered by describing the problem locally near a point where infinite-dimensional uncertainty disappears. A linearized solution allows consideration of responses to any possible pattern or shape that can be considered, accurately representing the behavior of the system in an infinite-dimensional set of possible inputs. Moreover, for many systems, this response can be calculated quickly and accurately, uniformly over all possible directions by projection representations of the functional derivatives of the system.

The dynamics of economic interactions over space, typically challenging to describe due to the fact that people in different locations must respond differently to the geographic patterns of economic activity induced by trade and spatially inhomogeneous regional disturbances, provide a demonstration of the rich patterns of relationships that can be captured by allowing decisions and distributions to respond to the precise geographic pattern of shocks. Responses can differ substantially based on distance, but also based on expectations of perceived future spatial distributions. Although spatial interactions provide a case which illustrates the full importance of allowing for response to potentially arbitrarily shaped patterns of heterogeneity, the function-valued approach seems promising for a wide variety of applications. These include understanding the mechanisms behind the dynamics of income and wealth inequality over business cycles, analyzing both through the relationship with capital markets, as has been explored in existing studies of incomplete markets models with aggregate shocks, as well as other potential economic mechanisms and policies. They may also be useful for studying a variety of patterns of interaction which depend on

the entire shape of the distribution of heterogeneity, such as matching markets in labor or other contexts or interactions through a social or economic network.

While for some applications, existing methods may be used to characterize the dynamics of economic heterogeneity, albeit without explicit guarantees of accuracy, the function-valued approach may still be desirable as a framework for data analysis. By explicitly allowing the model to incorporate uncertainty of arbitrary shape, the models described allow a complete characterization of the variation in micro and macroeconomic data and open the possibility of comparing the model directly to cross-sectional micro data. Because linearized function-valued models generate dynamics consistent with functional linear processes, estimation and inference methods from functional data analysis may be applied to evaluate them empirically. Given the speed and accuracy of the solution methods, they may also open up the possibility of using functional data methods to perform full information structural estimation of models with heterogeneous agents.

Chapter 2

A Dynamic Model of Economic Geography

To study the evolution over time of the spatial structure of economic activity, population, and welfare, I introduce a dynamic stochastic model of trade, migration, and economic geography. In this model, agents make forward-looking costly migration decisions in response to a spatial distribution of wages which is determined endogenously in spatial equilibrium as a function of the population distribution and patterns of persistent regional shocks to amenity values. A closed form characterization of the equilibrium is provided for a class of economies with idealized spatial structure and numerical methods and simulation results are provided for a general case allowing nonparametric spatial heterogeneity along a variety of attributes. The setting allows for a reevaluation of the relationship between spatial agglomeration externalities and population dynamics, suggesting that the sources of long-run spatial heterogeneity may differ substantially from those driving the response to temporary shocks.

2.1 Introduction

The spatial structure of economic activity may vary over time due to changes in economic variables which differentially affect certain locations and the propagation of these effects through channels which link economic activity across space. Trade creates spatially complex patterns of linkages in production and consumption which induce nontrivial relationships between local changes and the global distribution of economic activity and welfare. On somewhat longer timescales, the economic structure of different locations is transformed by migration, induced in part by differences in economic wellbeing across regions. To maintain tractability in the study of economies with complicated spatial structure, it is common to evaluate welfare effects of changes in economic variables in *trade equilibrium*, where trade and production adjust but the distribution of factors remains fixed, and to consider patterns of population across cities or locations in *spatial equilibrium*, in which population has adjusted to equate welfare across regions. While convenient for characterizing cross-sectional variation, these frameworks make it difficult to study the effects of spatially distributed uncertainty and the dynamics of adjustment to spatial shocks in which the spatial distributions of population, production, and consumption evolve gradually and interdependently.

A particular difficulty faced by these static frameworks is reconciling the high degree of spatial concentration of economic activity and population corresponding to urban agglomerations with dynamic processes of migration which must lead to these agglomerations. In long run equilibrium, migration across regions is often assumed to act as an equilibrating force which attenuates or eliminates differences in welfare (for the marginal resident) between locations (Roback, 1982; Piyapromdee, 2015). At the same time, substantial evidence in economic geography suggests that pervasive agglomeration externalities can create complementarities which induce or magnify regional heterogeneity, causing population and economic activity to concentrate in

space (see Glaeser & Gottlieb (2009) for a survey). If these forces are sufficiently powerful, movements in population may in some sense be self-reinforcing, so that transitory population changes may induce changes in local welfare which induce further migration (Krugman, 1991; Allen & Arkolakis, 2014). Under an exogenously specified behavioral equation for population movements chosen for tractability (see Krugman (1996, p. 109 Footnote 1)), Krugman (1996) shows that a model with these features exhibits local dynamic instability in the spatial structure, with small perturbations from a geographically homogeneous steady state diverging towards spatially variegated, history-dependent regional patterns. This class of processes provides a potential explanation for persistent economic differences across locations and the origin of cities, but implies that the joint dynamics of migration, wages, and economic activity play a destabilizing rather than stabilizing role.

These contrasting views of the role of migration in the dynamic evolution of regional activity and welfare call for a more exacting characterization of the interaction between population movements and spatial structure. To investigate these issues, I introduce fully forward-looking costly dynamic migration decisions and stochastic aggregate shocks perturbing the desirability of population flows between locations into a model subsuming the spatial structure of trade and production of Krugman (1996) but also permitting extremely rich (nonparametric) quantitative ex ante spatial heterogeneity as expressed in Allen & Arkolakis (2014). To capture the temporal structure of the dynamics, the model includes dynamic idiosyncratic heterogeneity inducing migration flows in equilibrium, subject to adjustment costs, so that adjustment to regional shocks must take place in the short run by movements of prices and quantities and only gradually by shifts in population. This feature, along with explicit modeling of preferences, also has the advantage of allowing evaluation of the welfare implications of regional shocks. This fully dynamic structure replaces the more common assumptions of immediate adjustment of populations to equalize welfare across

regions, complete markets in insurance for regional uncertainty, or complete labor immobility commonly imposed to ensure tractability in spatial models. An exception is provided by Caliendo *et al.* (2015), which is descriptively rich, and from which the model incorporates the structure of the intertemporal migration problem, as introduced originally by Artuç *et al.* (2010).

The proposed model differs from Caliendo *et al.* (2015) in market structure and in allowing a continuum of locations, which is feasible in a high-dimensional dynamic model using the functional linearization techniques introduced in Chapter 1. As will be shown, the use of a continuum provides several advantages. Beyond the greater generality and ability to fit to arbitrary resolution, the continuum allows substantially greater tractability and analytical insight under certain parameterizations of the model. Most importantly, in the special case, also considered by Krugman (1996), of a completely spatially homogeneous geography with a continuum of locations where no location differs ex ante from any other, the linearized solution to the model can be described analytically. This permits detailed characterization of the effects of different structural parameters on the responses to shocks and to the local dynamic stability of the model. In particular, it permits reevaluation of the hypothesis of history dependence, in which small temporary shocks to fundamentals can induce permanent changes in the spatial structure of the economy, in a self-reinforcing process in which population growth induces local productivity growth through agglomeration effects, which in turn raises wages and attracts more population growth. It will be shown that in a forward-looking economy, the conditions under which such a process induces permanent changes may be substantially more stringent, in essence because perception that growth will be temporary dampens the migration response, reducing the agglomeration effect and thereby slowing down each stage of the cycle so that the total effect is finite and transient.

Beyond the tractable spatially homogeneous special case, the generality and preci-

sion of functional linearization also permit numerical characterization of the dynamics for substantially more realistic spatial patterns. In order to implement this procedure, I introduce a simple and fast iterative method for computing a steady state equilibrium for nonparametric spatial patterns of trade and migration costs or evolution of shocks, and apply it in combination with the wavelet-based approximate solution algorithm in Chapter 1. I then construct simulations and impulse responses to shocks under highly rugged spatial patterns to demonstrate how economic variables evolve when locations are not ex ante identical. The results highlight the role of migration and trade costs in differences between regions and in their responses to shocks.

2.2 Model: Trade, Migration, and Economic Geography

We begin with the intertemporal decision problem, which can be analyzed independently of the static equilibrium structure. Notation follows Krugman (1996). Individuals working in the tradeables sector at location x in geography G, a set of locations with a distance metric which for now we take to be a subset of Euclidean space, receive in each period t a real wage $\omega_t(x)$ and a value of regional amenities $\nu_t(x)$, both taken as given by the worker. A worker in location x at time t may decide to move to location x' in period $t+1$ at a cost $c(x', x)$ which is a convex function of distance traveled. Workers are risk neutral with time-separable additive utility and discount the future at rate β. In each period they also receive independent and identically distributed across time and worker shocks $\epsilon_t(x')$ to their utility for each potential choice of location x', distributed according to a Gumbel process (Maddison *et al.* , 2014), whose finite-dimensional marginal distributions are independent Type I extreme value random variables. The Bellman equation for the decision problem is

therefore given by

$$\tilde{V}_t(x) = \max_{x'}\{\omega_t(x) + \nu_t(x) + c(x', x) + \epsilon_t(x') + \beta E_t \tilde{V}_{t+1}(x')\}$$

For reasons of tractability, it is easier to work with the conditional expectation of this equation: denoting $V_t(x) := E_t\tilde{V}_{t+1}(x)$, we obtain

$$V_t(x) = E_t \max_{x'}\{\omega_{t+1}(x) + \nu_{t+1}(x) + c(x', x) + \epsilon_{t+1}(x') + \beta V_{t+1}(x')\}$$

As a result, the location decision satisfies a continuous analogue of a multinomial logit decision rule: the conditional density of choices at location x' given current location x is given by

$$p(x'|x, V) = \exp(c(x', x) + \beta V(x'))/\int \exp(c(x', x) + \beta V(x'))dx'.$$

The use of extreme value shocks to generate a logit formulation for the policy function is similar to that used in Caliendo *et al.* (2015), with the difference that here the decision rule is defined over a continuum. To simplify notation, we will write the partition function of this conditional density as

$$f(x, V) := \int \exp(c(x', x) + \beta V(x'))dx'.$$

Using the closed form characterization for the expectation of the maximum of a Gumbel process, it is possible to write the expectation over the maximum in terms of the partition function, allowing the Bellman equation to be simplified to

$$V_t(x) = E_t\omega_{t+1}(x) + \nu_{t+1}(x) + \log f(x, V_{t+1}) + \gamma \tag{2.1}$$

where γ is the Euler-Mascheroni constant (≈ 0.577). Due to this explicit form, no numerical optimization is needed to compute the value function. Since it can be shown that Blackwell's conditions hold, the Bellman operator is a contraction and the steady state value can be found by iteration of the contraction mapping.

The above constitutes the forward looking component of the model. To determine the implications of the chosen policy for dynamics of the equilibrium, assume that *at each location* there is a continuum of workers, who each receive independent and identically distributed preference shocks, and that the total mass of workers has measure 1 and is distributed across locations at time t with density at location x given by $\lambda_t(x)$. Since the conditional density over locations given an initial state x is given by $p(x'|x, V)$, the time evolution of the density of workers across regions is given by the (adjoint) Markov transition operator

$$\lambda_{t+1}(x') = \int_G p(x'|x, V_t)\lambda_t(x)dx \tag{2.2}$$

taking the current distribution of population $\lambda_t(x)$ to the next period distribution $\lambda_{t+1}(x)$.

Together, λ_t and $\tilde{V}_t$ constitute the endogenous function-valued state variables of the model. To complete the model, one computes a static spatial equilibrium which generates a value of real wages at each location $\omega_t(x)$ given a distribution of population across places. A number of assumptions on market structure, trade, and geographical spillovers are possible here, with many models of trade and geography taking similar functional forms as discussed by Allen & Arkolakis (2014). A simple benchmark choice is the model of increasing returns, monopolistic competition, and iceberg trade costs of Krugman (1996), whose static structure can be borrowed without change. Specifically, we copy the block of equations (A.24)-(A.27) of that model to determine wages given population. See Krugman (1996) for derivation and more detailed explanation.

Variables included in these equations are $Y(x)$, output at location x, $T(x)$, the price index at location x, and $w(x)$, the nominal wage in terms of the nontraded good. Parameters used are σ, the elasticity of substitution of the CES aggregator across varieties, μ, the Cobb-Douglas expenditure share on traded goods, $\tilde{\tau}$, a scale factor, and $\tau(x,z)$, the effective economic distance in the proportional iceberg trade costs $1-\mathrm{e}^{-\tilde{\tau}\tau(x,z)}$ of shipping a good from point x to point z, and a normalizing constant $\tilde{c}$. Given a predetermined distribution of population $\lambda_t(x)$, a static equilibrium of the model is given by functions $\{Y_t(x), T_t(x), w_t(x), \omega_t(x)\}$ satisfying the system of nonlinear integral equations

$$
\begin{aligned}
Y_t(x) &= 1-\mu+\mu\lambda_t(x)w_t(x) && (2.3)\\
T_t(x) &= \left[\tilde{c}\int_G \lambda_t(z)w_t(z)^{1-\sigma}e^{\tilde{\tau}(1-\sigma)\tau(x,z)}dz\right]^{\frac{1}{1-\sigma}} && (2.4)\\
w_t(x) &= \left[\tilde{c}\int_G Y_t(z)T_t(z)^{\sigma-1}e^{-\tilde{\tau}(\sigma-1)\tau(x,z)}dz\right]^{\frac{1}{\sigma}} && (2.5)\\
\omega_t(x) &= w_t(x)T_t(x)^{-\mu} && (2.6)
\end{aligned}
$$

This system of equations is not analytically tractable, and has no explicit solution for ω_t in terms of λ_t. However, a solution in general exists and under certain conditions on parameters one may be able to calculate an implicit solution.

The dynamic specification of the model is completed by the inclusion of aggregate uncertainty. As described above, while agents take into consideration the expectation of real wages, in the absence of additional inputs, these evolve deterministically. A number of potential sources of aggregate uncertainty can arise which affect the evolution of population across regions. For the purpose of the decision problem over locations, however, any source of uncertainty which affects the static equilibrium of the model exerts its effect only through its impact on the real living standards at different locations, $\omega_t(x)+\nu_t(x)$. Amenity value is exogenous in this specification of the model, and shocks to amenities across locations can reflect natural mechanisms

like patterns of weather or natural disasters, or outcomes of (exogenous) local policies. Disturbances to variables determined within the static equilibrium of the model, such as changes in productivity (which may vary by location) in traded or nontraded sectors, changes in trade costs, or relative preferences for different varieties of good, will all show up in real wages. Further, because these are determined as the outcome of a purely static process, any persistence in these deviations (aside from that transmitted through the dynamics of population, described above), must come from outside the model. As a result, for the purposes of deriving the dynamics of economic activity and population, it is equivalent to model all shocks as changes to the exogenous value of amenities $\nu_t(x)$ at time t, and to provide exogenously specified dynamics for these shocks.

While many forms are possible, because the model will end up being linearized, it is sufficient to consider a linear specification for the dynamics of $\nu_t(x)$. For simplicity of illustration and, later, computation, I consider a first order functional autoregression specification with a kernel representation of the transition operator.

$$\nu_{t+1}(x) = \int_G \Gamma(x,z)\nu_t(z)dz + \varepsilon_{t+1}(x) \tag{2.7}$$

In the above, $\Gamma(.,.)$ is some bounded, smooth, square-integrable function parameterizing the degree of spatial diffusion of shocks, and $\varepsilon_t(x)$ is an i.i.d. function-valued Banach random element with covariance operator Σ. Note that the additive formulation of the shock ν_t is without loss of generality even when interpreted as shocks to the trade component of the model, as subsequent to linearization, up to appropriate reparameterization of Γ and Σ, all specifications lead to a representation in the linearized Bellman equation as an additive shock to $\omega_t(x)$. While this is without loss of generality for the purpose of determining dynamic properties of the model, specification of the particular form in which shocks enter could be used to aid identification

of the effects of these particular shocks, by specifying Γ and Σ as results of the composite effects of multiple shocks. However, since all static variables are determined jointly and contemporaneously, identification requires the aid of functional form assumptions on the covariance of shocks or the validity of external instruments, and so no conclusions can be drawn without additional assumptions. The additive structure does have implications for identification of the model in the case where amenities are not directly observed, as the impact of shocks to real wages and amenities must be disentangled to identify the effects of each, but given the dynamics of the sum, the division has no effect on the dynamics of the other model variables, at least to first order. Similarly, as in Caliendo *et al.* (2015), the idiosyncratic valuation shocks which induce diffusion of population may be divided between real income and amenities, affecting interpretation but not the dynamics of the model.

While we will later consider parameterizations under which it is necessary to solve numerically for many components of linearized model, for this model, it is possible to construct a particularly tractable special case in which the steady state and projections of derivatives can be computed exactly. In this case, we set the effective distance for trade $\tau(x, z) = |x - z|$, and likewise set the distance cost for migration to $c(x', x) = c(x' - x)$, so that the cost between any two locations is not changed if we translate the origin and the destination by an identical distance. If we assume that the geography is spatially homogeneous, such as the case of a circle, a sphere, an infinite line or plane, or higher-dimensional analogues of the preceding, the steady state of the system has a closed form solution. In particular, set $\nu_t(x)$ to 0 in all periods and conjecture that the initial distribution of population is uniform over the real line, in the sense that population measure over any interval is given by Lebesgue measure over the interval. Then it can be seen that a solution of the static equilibrium component of the model is given by $\{Y_t(x), T_t(x), w_t(x), \omega_t(x)\}$ which are all constant over x. To ensure a simple normalization, let $\tilde{c} = \frac{\tau(\sigma-1)}{2}$, which ensures that all of

these functions are equal to 1. Plugging this into the Bellman equation under the assumption that $\omega_t = 1$ is constant over time, shows that, because $c(x' - x)$ is translation invariant, $V(x) = \bar{V}$ constant is the unique solution of the Bellman equation. Placing this in $p(x'|x, V)$, we obtain that $p(x'|x, \bar{V}) \propto \exp(c(x' - x))$ and so is also translation invariant, and in particular if $c(x' - x) = \log g(x' - x)$ for any nonnegative function $g(.)$, the transition equation for λ_t is given by a convolution with a density proportional to $g(.)$. For example, if $c(x' - x) = -\frac{1}{2c}(x' - x)^2$, quadratic adjustment costs, equation (2.2) is given by convolution with a Gaussian with standard deviation c, and if $c(x' - x) = -\frac{1}{c}|x' - x|$, equation (2.2) is given by convolution with a Laplace distribution with dispersion parameter c. Because convolution is spatially invariant, the unique steady state of this transition equation on a translation-invariant domain is the uniform distribution, thus verifying the initial conjecture. For convenience, note that in steady state the partition function $f(x, \bar{V})$ is a constant, $\bar{f}$. To generate a specification in which the dynamics share the translation invariance property of the steady state, $\Gamma(x, z)$ can be chosen to be translation invariant, equal to $\Gamma(x - z)$ for a bounded square integrable univariate function, thereby restricting to shocks which do not diffuse differently from ex-ante identical locations.

Given the existence of a steady state, the dynamics of the model local to this point can be expressed by taking functional derivatives of the operators. To express this model in format appropriate for solution by a functional linear rational expectations algorithm, express the model recursively in terms of jump variable $V(x)$ and predetermined variables $\lambda(x)$ and $\nu(x)$ and their next period values V', λ', ν', solving out the static variables, which may be expressed at each time as a deterministic function of these three states, which are completely sufficient to solve for the dynamics of the model. We consider perturbations of V and ν as elements of $L^2(\mathbb{R})$ and perturbations of λ, a probability distribution, as an element of $L^2_0(\mathbb{R})$, the space of square integrable functions on $\mathbb{R}$ integrating to 0, ensuring that densities integrate to 1.

First, the transition equation is linear with respect to λ with derivative given by a convolution of the argument with density proportional to $\exp(c(x' - x))$, an operator which we can denote as $P[.] := \int \frac{1}{\bar{f}} \exp(c(x' - x) + \beta\bar{V})[.]dx'$. This can be interpreted as convolution with Gibbs distribution with potential given by the cost of moving: in the absence of disturbances to the value of a different locations, given a current population at each location, next period population spreads out by an amount proportional to the cost of distance. The Bellman equation is linear in V with derivative equal to the identity and has functional derivative with respect to V' given by $\frac{\beta}{\bar{f}} \int \exp(c(x'-x)+\beta\bar{V})[.]dx' = \beta P[.]$. The transition equation has derivative equal to the identity with respect to λ' and has derivative with respect to V given by $\beta \int_G \frac{1}{\bar{f}} \exp(c(x'-x)+\beta\bar{V})[.] - \frac{1}{\bar{f}} \exp(c(x'-x)+\beta\bar{V}) \frac{\beta}{\bar{f}} \int \exp(c(z'-x)+\beta\bar{V})[.]dz'dx$, which equals $\beta P - \beta PP$. The transition equation for ν is linear in ν and ν', with derivative with respect to ν given by $\Gamma[.] := \int \Gamma(x, z)[.]dz$ and ν' by the identity. Finally, although no closed form expression exists for $\omega(x)$ in terms of $\lambda(x)$, its functional derivative $\frac{d\omega}{d\lambda}$ with respect to $\lambda(x)$, which is all that is needed, can be determined by implicit differentiation of equations (2.3),(2.4), (2.5), and (2.6): the exact formula is derived in Appendix C as equation (13).

Together these calculations fully characterize the derivatives of the model's equilibrium conditions with respect to the state variables. Arranging these derivatives into blocks with elements given by linear operators, the linearization of the equilibrium conditions of this model can be expressed in a form suitable for application of our solution methods, as a pair of linear operators representing the derivatives of the equilibrium conditions of the model with respect to today's state variables (λ, ν, V)

and tomorrow's state variables (λ', ν', V').

$$\left(\begin{bmatrix} 0 & 0 & I \\ P & 0 & \beta P - \beta PP \\ 0 & \Gamma & 0 \end{bmatrix}, \begin{bmatrix} \frac{d\omega}{d\lambda} & I & \beta P \\ I & 0 & 0 \\ 0 & I & 0 \end{bmatrix} \right) \tag{2.8}$$

In this pair of operators, the columns correspond to function-valued state variables, while the rows correspond to the linearized equations defining the equilibrium. In order, these are the Bellman equation, the transition law for the population distribution, and the law of motion for the function-valued shock to the distribution of amenities. Derivatives expressed in this form may be used to solve for the linearized dynamics and responses of the state variables of the model to endogenous and exogenous changes.

2.3 Implementation and Evaluation: Spatially Homogeneous Case

In the case of an economy with a geographic structure which is spatially homogeneous, in the sense that all locations are completely identical, with the relationship between any two locations dependent only on the distance between them, several analytical or partially analytical results can be derived regarding the dynamics of the economy. In particular, using the functional linearization techniques of Chapter 1, for which the sufficient conditions can be easily verified in this model, explicit formulas for the local dynamics in a neighborhood of the nonstochastic steady state can be constructed, and the dynamic stability of this process can be evaluated. The law of motion can also be calculated with negligible loss of precision using the numerical approximation techniques introduced in that chapter in order to provide a unified methodology with the case, explored in the subsequent section, where a less stylized but also less

analytically tractable geographic structure is introduced.

2.3.1 Steady State and Exact Projections

To go from the functional derivatives of the equilibrium conditions to a linear solution, it suffices to find projections onto a complete set of basis functions. The structure of the model makes that task particularly simple, because when the basis used is the standard Fourier basis of trigonometric polynomials, the projections can be calculated exactly without numerical integration. The structure can also be used to verify the conditions which ensure that projection approximations are valid. In fact, the structure allows even more to be said about the solution than can be inferred from (1.1). Because the model is block diagonal with respect to the Fourier basis, the solution operator can be calculated exactly for any input given by a Fourier basis function, and so for any bandlimited function.

The linearized equilibrium conditions in this model are given by equation (2.8), in which $P[.] = \frac{1}{f}\int \exp(c(x'-x)+\beta\bar{V})[.]dx'$, Γ is likewise an integral operator, and $\frac{d\omega}{d\lambda}$ can be shown to be defined in terms of the composition of a number of convolution operators with respect to a Laplace distribution and their inverses. As can therefore be seen, the model is composed of identity and integral operators, exactly the structure needed for the projection approaches developed in the first chapter of this thesis to be valid. Moreover, examining the expressions for the derivatives of the economic geography model it can be seen that all of the integral operators are expressed in terms of convolution operators. By the convolution theorem, all convolution operators (and their inverses, as well as the identity) are diagonal in a Fourier basis, and so all operators can be expressed as a convolution with distributions, or equivalently, as multiplication of the Fourier transform of the input by a known function.

Because each functional derivative in the model is diagonal with respect to the Fourier transform, the model can be broken down into blocks corresponding to in-

dividual frequencies: there is no interaction across frequencies. Within a frequency, the linearized model can be written in terms of 3×3 matrices of derivatives of each component with respect to perturbations at that frequency. The exception is at frequency 0, where only derivatives with respect to V and ν are taken, as, by Parseval's theorem, functions $L_0^2(\mathbb{R})$ can be represented in the Fourier domain as sequences of Fourier coefficients with the coefficient at frequency 0 equal to 0.

Among other things, this block diagonal structure implies that Condition 1.(ii) regarding the modulus of continuity of the Schur decomposition holds so long as Condition 1.(i) holds. Conditions 1.(i) and 1.(iii), requiring existence and uniqueness of a Schur decomposition into components inside and outside the unit circle with unstable subspace isomorphic to the space spanned by the jump variable (in this case V), may also be verified for any given set of parameters by ensuring the conditions hold for each finite-dimensional subsystem. In order for the system to have a locally stable rational expectations equilibrium, it must be the case that at each frequency, the system has two generalized eigenvalues inside the unit circle, corresponding to the predetermined variables ν and λ, and one generalized eigenvalue outside, corresponding to the jump variable V. Such a condition is not general: it requires restrictions on the parameter values to ensure that such a solution exists.

Impressionistically, because the value of a location is a weighted average of future wages (a function of population), and because the current population is a weighted average of past values, the system remains stable only if this mutual reinforcement is not too strong. Otherwise, at certain frequencies, at which more than one eigenvalue is unstable, the linearized model implies that value grows without bound and population does as well: this is the conclusion of Krugman (1996), which does not derive dynamics from forward-looking decisions. However, the stability condition on the eigenvalues is substantially weaker than the condition imposed by Krugman, that the impact of population on wages be negative for all frequencies. Positive feedback is consistent

with stability of a rational expectations equilibrium so long as the effect on wages is expected to be temporary. Moreover, if the feedback is temporary, the population response is damped, and so the degree of mutual reinforcement is even lower. As a result, only frequencies where the parameterization implies that the feedback from population to wages is so large that no policy rule which eventually returns to steady state can be constructed are a problem for calculating a forward looking solution.

To consider which frequencies might be problematic, note that at extremely high frequencies, because convolution with a smooth density dampens high frequency fluctuations, the mutual reinforcement phenomenon is dampened and eventually disappears, so these frequencies are stable. Similarly, due to the dispersive forces in the geographic equilibrium model, at extremely low frequencies, increasing population actually reduces wages, ensuring stability. It is at intermediate frequencies where population growth and real wage growth are complementary, and parameters must be chosen so that at these frequencies the degree of complementarity is not so great as to prevent the mean-reversion induced by the dispersion of population due to idiosyncratic tastes from ensuring eventual return to uniformity after a temporary shock. This suggests that a parameterization of adjustment costs which ensures that medium to high frequency fluctuations are rapidly smoothed out is needed. However, degree of smoothing and size of adjustment costs have a nontrivial relationship. For quadratic costs, a higher scale is equivalent to a smaller variance of the Gaussian flows and so results in less smoothing. However, while changing from quadratic to linear (in absolute distance) costs results in Laplace flows with substantially more movement to long distances as it lowers costs of moving large distances, it raises costs of moving short distances and so decreases mean reversion at medium to high frequencies. In practice, stability holds for a very broad range of parameter values.

Formally, at each frequency not equal to 0, the model is represented by a 3×3 block of the Bellman equation, the population transition, and the shock transition at

that frequency. At a representative frequency ϕ, the model can be taken as a set of matrix equations in Fourier transform of the vector of endogenous functions at that frequency. The matrix of derivatives with respect to $(\hat{\lambda}_\phi, \hat{\nu}_\phi, \hat{V}_\phi)$ is

$$B_\phi = \begin{bmatrix} 0 & 0 & 1 \\ \hat{P}_\phi & 0 & \beta\hat{P}_\phi - \beta\hat{P}_\phi^2 \\ 0 & \hat{\Gamma}_\phi & 0 \end{bmatrix}$$

where $\hat{P}_\phi$ is the Fourier transform of $\frac{1}{\bar{f}}\exp(c(x)+\beta\bar{V})$ evaluated at frequency ϕ and $\hat{\Gamma}_\phi$ is the Fourier transform of $\Gamma(x)$ evaluated at frequency ϕ. The matrix of derivatives with respect to $(\hat{\lambda}'_\phi, \hat{\nu}'_\phi, \hat{V}'_\phi)$ is

$$A_\phi = \begin{bmatrix} \frac{\hat{d\omega}}{d\lambda}_\phi & 1 & \beta\hat{P}_\phi \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

where $\frac{\hat{d\omega}}{d\lambda}_\phi$ the Fourier transform of $\frac{d\omega}{d\lambda}$ at frequency ϕ, is derived in Appendix C.

Finally, at frequency 0, by dropping the transition equation which does not act over this frequency because perturbations of λ are restricted to lie in L_0^2, the space of functions integrating to 0, to ensure that the density λ integrates to 1, the system is represented by 2×2 blocks of derivatives with respect to $(\hat{\nu}_\phi, \hat{V}_\phi)$ and $(\hat{\nu}'_\phi, \hat{V}'_\phi)$ given by

$$(B_0, A_0) = \left(\begin{bmatrix} 0 & 1 \\ \hat{\Gamma}_0 & 0 \end{bmatrix}, \begin{bmatrix} 1 & \beta\hat{P}_0 \\ 1 & 0 \end{bmatrix} \right)$$

To construct an approximate solution from these projections, note that because the operator pairs are block diagonal, a fully upper triangular infinite-dimensional system can be constructed so long as each block can be placed in upper triangular form. Together, each pair of matrices forms a finite-dimensional linear rational

expectations system which can be evaluated by standard algorithms for calculating perturbative expansions of such systems, such as the Schmitt-Grohe & Uribe (2004) procedure, based on the algorithm of Klein (2000). Here, no changes need to be made to the finite-dimensional procedure: it is simply applied independently at each integer frequency ϕ. The derivatives of the policy functions are then given by the collection of derivatives at each frequency. For each $\phi \neq 0$, the policy functions $\hat{h}_\phi : (\hat{\lambda}_\phi, \hat{\nu}_\phi) \to (\hat{\lambda}'_\phi, \hat{\nu}'_\phi)$ and $\hat{g}_\phi : (\hat{\lambda}_\phi, \hat{\nu}_\phi) \to \hat{V}_\phi$ are given by 2×2 and 1×2 matrices. The first order approximate policy operators are then represented with respect to the Fourier basis as block-diagonal infinite matrices $\hat{h}$ and $\hat{g}$, with $\hat{h}_\phi$ and $\hat{g}_\phi$ on the diagonals, respectively, so that for general inputs in $L^2_0(\mathbb{R}) \times L^2(\mathbb{R})$, they may be represented as $h = \mathcal{F}^{-1}\hat{h}\mathcal{F}$ and $g = \mathcal{F}^{-1}\hat{g}\mathcal{F}$ where $\mathcal{F}$ is the Fourier transform and $\mathcal{F}^{-1}$ is the inverse Fourier transform.

For bandlimited perturbations, a finite representation h^K_x, g^K_x given by concatenating the first K frequencies is exact. More generally, the functional derivatives generated by taking an increasing finite collection of frequencies converge in the strong operator topology, and for any components which are compact, in the operator norm topology. Operator norm convergence follows from application of (1.1). To see that the conditions are met, note that for smooth adjustment costs and transition functions, $\hat{P}_\phi$, $\hat{\Gamma}_\phi$, and $\frac{\hat{d\omega}}{d\lambda}_\phi$ converge to 0, and so compactness and convergence of the projected derivatives in operator norm follows. Moreover, as $\phi \to \infty$, (B_ϕ, A_ϕ) converges to

$$(B^i_I, A^i_I) = \left(\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \right),$$

and by the continuity of the generalized Schur decomposition with respect to perturbations, so do the policy functions at each frequency. It can be shown that the first derivatives of the policy functions $g^i_x = \hat{g}_\infty : (\hat{\lambda}_\infty, \hat{\nu}_\infty) \to \hat{V}_\infty$ and $h^i_x = \hat{h}_\infty :$

$(\hat{\lambda}_\infty, \hat{\nu}_\infty) \to (\hat{\lambda}'_\infty, \hat{\nu}'_\infty)$ generated by calculating the finite-dimensional linear rational expectations solution for this pair are given by matrices which are identically 0, and so $h_K = h_x^K$ and $g_K = g_x^K$. As a result, by taking an increasing set of frequencies, the finite representation can be used to compute a response which is accurate uniformly over all input functions, and not just bandlimited ones.

It is possible to determine the rate of convergence directly from the exact representations rather than by applying the rate results from (1.1). Note that perturbation results for generalized eigenvectors and eigenvalues imply a linear rate of convergence in the Frobenius norm of the perturbation (see Stewart & Sun (1990)), while sufficiently smooth functional forms for adjustment costs and for the transition operator for the exogenous shocks, and the exponential form chosen for trade costs, generate rates of convergence for the entries which are faster than linear in ϕ. As a result, given sufficiently smooth parameterizations, the blocks of the policy function corresponding to each frequency converge at a rate comparable to the slowest rate of each of the components. So long as this converges to 0, this implies that the policy operators are compact (and if this rate is faster than linear, they are Hilbert-Schmidt), and so the policy operators given by taking an increasing finite sequence of blocks converge to the true policy functions in operator norm. One note regarding the form of this convergence is that the perturbation theorem for the Schur subspaces applies only under a separation condition on the generalized eigenvalues, while (B_I^i, A_I^i) has the generalized eigenvalues $(0, \infty, \infty)$. This implies that the blocks corresponding to forward and backward looking components are well separated, while within the block of backward looking components the eigenvalues are not asymptotically well separated and the generalized Schur vectors are not stable. However, the block itself is stable in the sense that the span of the Schur vectors converges, and so the policy functions, which are determined only by the sub-blocks of the Schur matrices, also converge.

Formally, this may be stated as

Lemma 2.1. *(i)* $g_x^i = \hat{g}_\infty = (0,0)$, $h_x^i = \hat{h}_\infty = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$ *Suppose*

$$\gamma_\phi := \left\|(B_\phi, A_\phi) - (B_I^i, A_I^i)\right\|_F \to 0$$

as $|\phi| \to \infty$. *Then* $\|\hat{g}_\phi - \hat{g}_\infty\|_F = O(\gamma_\phi^{\frac{1}{2}})$ *and* $\left\|\hat{h}_\phi - \hat{h}_\infty\right\|_F = O(\gamma_\phi^{\frac{1}{2}})$ *for large* $|\phi|$, *and so converge to 0 and*

$$h[\lambda(x), \nu(x)] = \hat{h}_0[\int \nu(x)dx] + \sum_{\phi \in \mathbb{Z}\setminus\{0\}} (\hat{h}_\phi \begin{bmatrix} \int \exp(-2\pi i\phi x)\lambda(x)dx \\ \int \exp(-2\pi i\phi x)\nu(x)dx \end{bmatrix}) \circ \begin{bmatrix} \exp(-2\pi i\phi x) \\ \exp(-2\pi i\phi x) \end{bmatrix}$$

and

$$g[\lambda(x), \nu(x)] = \hat{g}_0[\int \nu(x)dx] + \sum_{\phi \in \mathbb{Z}\setminus\{0\}} (\hat{g}_\phi \begin{bmatrix} \int \exp(-2\pi i\phi x)\lambda(x)dx \\ \int \exp(-2\pi i\phi x)\nu(x)dx \end{bmatrix}) \cdot \exp(-2\pi i\phi x)$$

are compact. (ii) Suppose in addition that $\gamma_\phi = O(|\phi|^{-(1+\epsilon)})$ *for some* $\epsilon > 0$. *Then* $h[\lambda(x), \nu(x)]$ *and* $g[\lambda(x), \nu(x)]$ *are Hilbert-Schmidt.*

Proof. See appendix. □

This result not only gives compactness and rates of convergence, it also implies that the approximated policy operators converge in a stronger norm, the Hilbert Schmidt norm. The demonstration that these operators are compact and Hilbert-Schmidt implies that in principle, the policy function in this model could be consistently estimated from a time series of observations of (λ, ν) by procedures such as those described in Bosq (2000); Guillas (2001).

2.3.2 Calibration and Numerical Evaluation

To characterize the dynamics of population and values in the model, I calculate the first derivatives of the policy operators for a fixed set of parameters. For the adjustment cost function $c(x)$, in order to generate a Gaussian conditional distribution of population movements in steady state, I adopt a quadratic specification $c(x) = \frac{1}{2\sigma_c}x^2$, where σ_c parameterizes the cost of moving and is also the standard deviation of the conditional Gaussian distribution. For the kernel describing the persistence of the exogenous shocks $\Gamma(x)$, in order to ensure both stationarity and decay of coefficients to represent smooth diffusion of shocks from their initial locations, I choose a rescaled Gaussian pdf, $\Gamma(x) = \frac{k_\Gamma}{\sqrt{2\pi}\sigma_\Gamma}\exp(-\frac{1}{2\sigma_\Gamma^2}x^2)$, where $|k_\Gamma| < 1$ ensures stationarity at all frequencies and σ_Γ measures the speed at which shocks spread, or, more directly, how rapidly the autoregressive coefficient on each frequency goes to 0 as the frequency increases.

For the static equilibrium of the model, I borrow parameterizations from Krugman (1996), who considers the ranges $\sigma \in \{4, 5, 6\}$, $\mu \in \{0.2, 0.3, 0.4\}$. As within this range the qualitative behavior of the model is similar, all experiments reported are carried out with $\sigma = 4$, $\mu = 0.4$. While the trade cost parameter τ is left unspecified in the parameterization as it merely normalizes the unit of distance in the model, the relative values of τ, σ_Γ, and σ_c determine the characteristic length scales at which trade, productivity (or other shock) diffusion, and migration operate. Note however that because trade costs are specified as exponential, while migration and productivity diffusion follow a Gaussian and so squared exponential rate of increase in distance, that the numbers are not directly comparable. This specification implies that trade at long distances is relatively less costly than migration or diffusion of changes in the economic environment. While difficult to place on a comparable scale, this seems to be qualitatively reasonable for a global or national scale, with long-distance exchange relatively common while long distance migration is comparatively

rare. For the purposes of simulations, and without any claim to represent empirically reasonable values, simulations set $\tau = 0.2$, $\sigma_\Gamma = 0.04$, and $\sigma_c = 0.05$, representing again fairly small trade costs and fairly slow diffusion of population and amenity value from an initial location. Along with a value of $k_\Gamma = 0.98$ and discount rate $\beta = 0.96$, these are designed to ensure that fluctuations in the spatial distribution of population and amenity values are persistent and that the model generates substantial variation in the expectations of future distributions.

While in principle a closed form is available for the policy functions at each frequency for arbitrary parameter values, it is an unintuitive nonlinear function of the roots of a cubic polynomial, so instead we verify the stability conditions at each frequency numerically. By the stability of the system at (B_I^i, A_I^i) and the convergence of (B_ϕ, A_ϕ) to (B_I^i, A_I^i), it is sufficient to verify the eigenvalue condition for the finite set of frequencies where the derivatives differ by more than some small constant from (B_I^i, A_I^i).

In practice, and in contrast to the generically explosive limit generated by the ad-hoc dynamics imposed in Krugman (1996), only for relatively extreme parameterizations does the model with forward-looking decision making lack an equilibrium which is locally stable. The complementarity between wages and population at intermediate frequencies generated by agglomerative forces in the model and the substitutability at low frequencies generated by the dispersive forces are reflected in the cross-derivatives of the transition operator $\hat{h}$ mapping shocks to living standards and population this period to those next period. The complementarity and substitutability manifest as a positive coefficient in the map from the shock $\hat{\nu}_\phi$ to amenity value today to population tomorrow at intermediate frequencies and a negative coefficient at low frequencies, respectively. However, the presence of a positive coefficient does not generate explosive behavior if the shock itself is mean-reverting, as assumed, and the autonomous dynamics of population are also stable. Here, except when the elasticity of substitu-

tion across varieties σ is extremely low so the benefits of agglomeration in a region with large population and a large variety of goods is high, the natural smoothing of population across regions generated by heterogeneous idiosyncratic preferences is the dominant determinant of the speed of adjustment of population at a given frequency. As a result, even for very strong agglomerative forces, it is also necessary for adjustment costs of moving to be quite large before complementarities at some frequency dominate and generate dynamics which are locally unstable.

In part, this expresses an important difference between the myopic and forward-looking models. In the myopic case, even small complementarities result in a cumulative process which continues without bound, while in a forward looking setup, if the effects of such complementarities are transient, their impact on value and so on decisions is bounded and so is attenuated. From an economic perspective, forward-looking decisions respond less strongly to changes perceived as temporary, and so even in the presence of complementarities, regional shocks need not be destabilizing. To be fair, however, some of the difference also reflects the additional dispersive force provided by idiosyncratic preference shocks, though it's not clear how one would generate a smooth transition law as in Krugman (1996) even with myopic decision making without some other smoothing force.

2.3.2.1 Finite Domain

While solving the model on an infinite domain ensures a great deal of tractability, it has some disadvantages, of which lack of realism is a minor but nontrivial one. From the perspective of demonstrating existence, compactness of the domain permits the use of standard existence theorems which are unavailable on unbounded space. Further, when approximating the integrals via an expansion in basis functions other than trigonometric polynomials, as may be needed for nonperiodic variations of the model, it permits use of compactly-supported basis functions, such as B-splines or

(certain classes of) wavelets, without requiring an unbounded number to cover the entire domain.

In the symmetric case, the loss of tractability is rather minor: by setting G to be a circle of circumference 1 with coordinates $x \in [0, 1)$ parameterizing the location[1] and changing the normalizing constant $\frac{\tau(\sigma-1)}{2}$ to $\frac{\tau(\sigma-1)}{2-2e^{-\tau(\sigma-1)/2}}$ in formulas (2.4) and (2.5), it can be easily seen that the steady state equilibrium remains uniformly distributed with $\bar{\lambda}(x) = \bar{\omega}(x) = \bar{W}(x) = \bar{T}(x) = 1 \ \forall x \in G$, and $\bar{V}(x)$ constant. The only material difference to the dynamics is that now instead of convolution with a Laplace or Gaussian distribution as the representation of the effect of population on wages or the dynamics of $\nu(x)$ or $\lambda(x)$ respectively, these operators are replaced by convolution with truncated (and recentered and renormalized) Laplace or Gaussian distributions, e.g. $\Gamma(x) = \frac{1}{1-2\Phi(\frac{1}{2})} \frac{k_\Gamma}{\sqrt{2\pi}\sigma_\Gamma} \exp(-\frac{1}{2\sigma_\Gamma^2}(x - \frac{1}{2})^2) 1[0 \leq x < 1]$. This reflects the economic structure of the problem: in a finite space, there is a finite maximum trade cost and finite maximum migration cost, and so a minimum impact of one location on another. Truncation does not change the ability to represent the operators as diagonal with respect to a Fourier basis, though now the result holds by the circular convolution theorem. The Fourier transform of a product is given by the convolution of the Fourier transforms, and so, by a change of variables, in the derivation of $\frac{d\omega}{d\lambda}$, $H(\phi)$ is replaced by $H(\phi) * \text{Sinc}(\phi)$, where $\text{Sinc}(\phi) = \frac{\sin \pi\phi}{\pi\phi}$ is the Fourier transform of $1[-\frac{1}{2} \leq x < \frac{1}{2}]$. While this convolution has no simple closed form expression, it is easily calculated numerically by quadrature.

For parameterizations with rapid increase in trade or migration costs over distance, this transformation has minimal effect, as the truncation only affects the far tails. For small trade or migration costs, it increases impact at some frequencies and decreases it at others, reflecting the periodicity induced by the circular shape. Numerical experiments suggest that even for relatively small costs, the impact of this

[1] By symmetry, the initial point 0 can be assigned to any arbitrary location.

change is limited. As a result, the main impact is on ensuring proper scaling and allowing testing approximate equilibrium computation using a wavelet basis.

Due to the relatively high cost of constructing the numerical integrals needed to evaluate the convolutions with the sinc function, and for comparison with the inhomogeneous case, it is also convenient to construct a representation using the wavelet based numerical approximation techniques of Chapter 1. To represent the circular convolutions with respect to a wavelet basis, the operators are first written in terms of the distance on a circle with $x' - x$ replaced with arc length along the diameter of the circle: $d(x', x) = \text{mod}(x' - x + \frac{1}{2}, 1) - \frac{1}{2}$ is the distance between points x', $x \in [0, 1)$ on the circumference. For example $\Gamma[\nu](x') = \int_0^1 \Gamma(d(x', x))[\nu(x)]dx$ describes the value of the amenity value $\nu'(x')$ next period at each point $x' \in [0, 1)$ given an initial distribution $\nu(x)$. Construction of wavelet approximations consists of sampling the kernels (e.g. $\Gamma(d(x', x))$) at an evenly spaced grid of $K \times K$ points on $[0, 1) \times [0, 1)$ and applying the discrete wavelet transform to the rows and columns of the resulting matrix. The kernels used in this model are infinitely differentiable at most values of x, y but nondifferentiable at $d(x, y) = 0.5$ due to the finite domain creating a maximal possible level of trade or migration costs at the antipodal location on the circle where counterclockwise or clockwise movements meet. For the exponential trade costs, there is also a point of nondifferentiability at $d(x, y) = 0$. For wavelet representations, there is a tradeoff between vanishing moments to represent the smooth parts parsimoniously and width of the scaling function which creates distortions at nonsmooth points. Although higher order Coiflets will achieve faster rates asymptotically, for finite values of K, lower order Coiflets may yield better performance, which is borne out in numerical experiments. As a compromise, level 3 Coiflets are used in all simulations and evaluations.

Two additional sets of approximations are made beyond those described in Chapter 1 Equation (1.2). To ensure that perturbations to the population distribution $\lambda(x)$

remain in the space of mean 0 functions, the wavelet representations of operators acting on this space are orthogonalized with respect to the the wavelet representation of the constant function. While for Haar wavelets this demeaning is exact, for other bases it yields a representation which is approximately orthogonal to constants. Rather than defining the exact kernel for $\frac{d\omega}{d\lambda}$ and applying the wavelet transform to it directly, because it is composed of convolutions with a Laplace distribution and identity operators, it may be constructed by applying the products and inverses of the wavelet representations of these operators. Because all applications are continuous (note that the inverse is applied to a sum of identity and convolution operators and so is well posed), so long as the convolution operators themselves are consistently approximated, $\frac{d\omega}{d\lambda}$ is as well.

2.3.3 Results

To evaluate the approximation algorithm, several numerical comparisons are performed using both the Fourier and the wavelet representations of the model, at different levels of K. Accuracy can be compared for impulse responses to function-valued shocks, as well as for simulations. Although the Fourier representation is exact in principle, at least for bandlimited functions, when restricted to a finite domain, the value of the Fourier coefficient at each frequency must be computed by quadrature due to the convolution with the sinc function. Nevertheless, the Fourier and wavelet methods appear to exhibit a high degree of agreement, whether expressed in squared error norm over the grid points (a proxy for L^2 norm, controlled by the theory) or in maximum norm over grid points (not controlled by the theory). Error is largest for components which are strongly impacted by the finite diameter of the geography, and declines for parameter values which ensure that the cutoff has limited effect on the representation, suggesting that the numerical error induced by quadrature may be a non-negligible factor contributing to the discrepancy between Fourier and wavelet

representations. For the Fourier representations, integer frequencies $\frac{-K}{2}$ to $\frac{K}{2}$ are used for each of $J = 3$ functions $\nu(x)$, $\lambda(x)$, and $V(x)$, giving $3 \times (K + 1)$ basis functions, for symmetry, while for wavelets K grid points are used to represent the scaling function coefficients for each function, with K given by a power of 2.

To describe the behavior of the model, first consider the impulse response to a smooth but spatially localized shock $\varepsilon(x)$ to the amenity value of locations, a scaled Gaussian spike centered at location 0.5, with functional form $\exp(50000(x - 0.5)^2)$. This may represent a nearly exactly localized improvement, as might occur in response to a local policy initiative or favorable productivity shock. As can be seen in Figure (2.3.1), the response of amenity value over time and space, calculated from $K = 1024$ using the Fourier representation, in spite of the high persistence parameter k_Γ and the relatively small standard deviation of the diffusion kernel σ_Γ, this shock spreads out rapidly from the initial location and diffuses from a local region to an eventually larger and larger area. Note that while the space coordinate is represented on a line segment, the model is defined over a circle, so the edges are connected.

The population response, displayed in Figure (2.3.2), follows the amenity shock but is much more dispersed, and responds slowly, peaking over 10 periods later and then declining gradually. The population in regions far from the center declines, as people move towards the more desirable area, with a nadir over 20 periods later. Despite the slow speed of adjustment, movements begin the first period after the shock, as individuals anticipate the spread of the amenity over space and the possibility of moving in the future to more desirable areas, which are desirable in part because they provide the option value of moving even close to the center in future at lower cost and so taking advantage of the improved amenity there. This is displayed clearly in the plot of welfare, $V_t(x)$ in Figure (2.3.2), which jumps immediately, with peak at the location of the shock but high values substantially more broadly dispersed, with a nontrivial jump in welfare over the entire domain, as even regions for which the value

Figure 2.3.1: Impulse Response of $\nu_t(x)$ to $\varepsilon_0(x) = \exp(50000(x - 0.5)^2)$

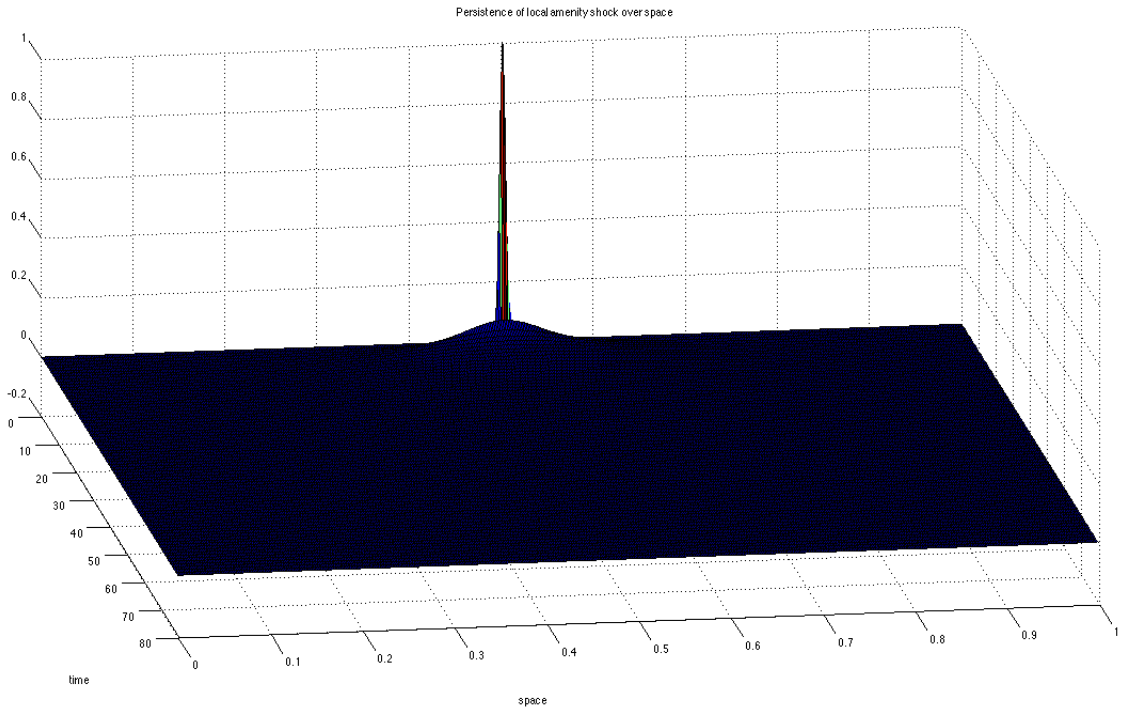

of the shock immediately and in the first few periods is essentially negligible face the prospect of higher welfare in the future as the amenity spreads out and population moves to regions positively affected by the shock.

To consider the behavior of the model in response to more complex patterns of input, I use it to produce simulated time paths. The shocks $\varepsilon_t(x)$ are drawn from a spatially correlated Gaussian process, a simulated fractional Brownian motion (started at 0) with Hurst parameter 0.7 and so a degree of Hölder regularity no greater than 0.7. Wavelet quadrature is easily capable of representing functions with this degree of regularity and so the simulations are drawn from the representation of the model with respect to a wavelet basis, with $K = 512$. The dynamic law of motion is defined in terms of the dynamics of $\lambda_t(x)$ and $\nu_t(x)$. In order to simulate also the implied paths of real and nominal wages and prices, $\omega_t(x)$, $w_t(x)$, and $T_t(x)$, which are deterministic functions of $\lambda_t(x)$, it is possible to either use the linearized version of the deterministic function, or simply plug the simulated value of $\lambda_t(x)$ into Equations

Figure 2.3.2: Impulse Responses of $\lambda_t(x)$ & $V_t(x)$ to $\varepsilon_0(x) = \exp(50000(x-0.5)^2)$

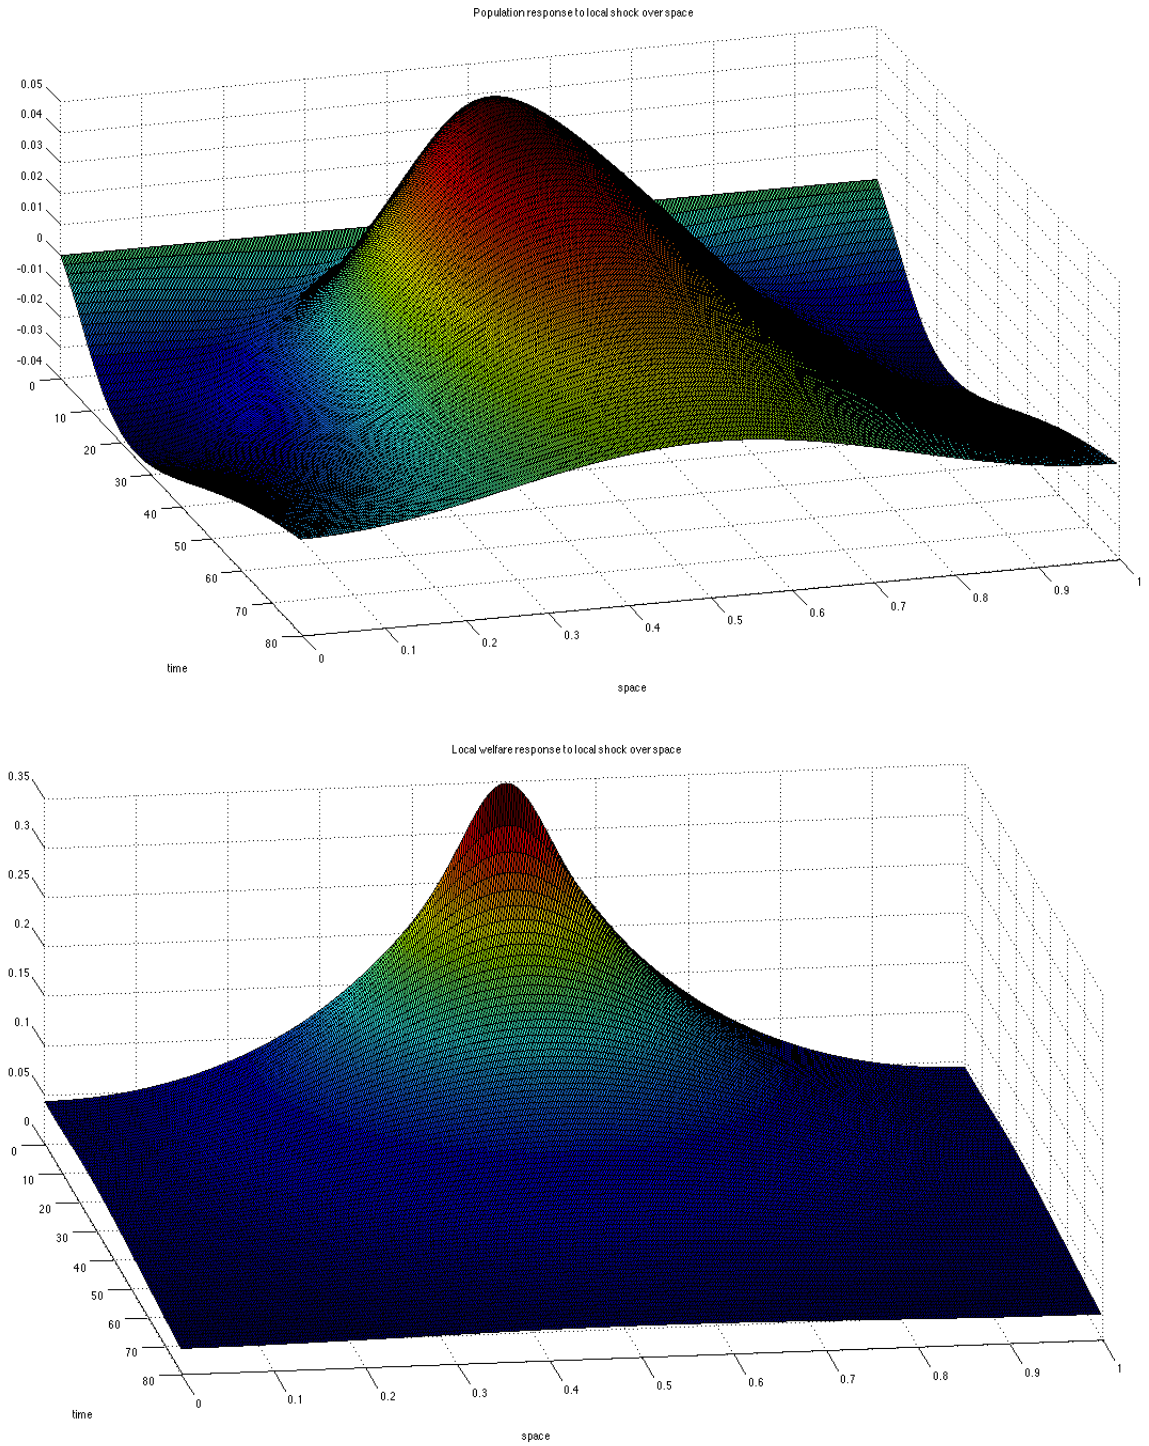

(2.3)-(2.6): to the extent that the simulated values of $\lambda_t(x)$ are only accurate to first order due to linearization, the latter procedure does not reduce the order of the error, but may eliminate a source of higher-order bias (and also does not require computing implicit functional derivatives) and so is used for the simulations shown. Time paths are displayed in Figures (2.3.3), (2.3.4), and (2.3.5).

One feature which stands out is the low degree of smoothness of $\nu_t(x)$, the persistent shock process, and $V_t(x)$, the welfare of residents at each location x, in contrast to the fairly high degree of smoothness of population movements $\lambda_t(x)$. This contrast is as should be expected, because $V_t(x)$ is a jump variable, and so adjusts immediately to reflect changes in the state, while population is a predetermined variable, and so changes only in response to expected future changes in welfare, which, because shocks to amenity value are expected to be smoothed out over time, substantially discounts the high frequency variations which impart roughness to the spatial distribution of current welfare. This is in line with standard reasoning for rational expectations decision problems: because moving is costly, transitory variation, expressed by the rough local movements in amenity values, has minimal effect on forward looking decisions. In contrast, low frequency changes, which are expected to be more persistent, do induce population movements, and the simulation does show periods of time where there are large population movements between regions. The simulation also exemplifies the expressive power of functional methods, as it allows description of the welfare and behavioral consequences of extremely finely detailed patterns of aggregate shocks, which would be difficult to express even with smooth nonparametric function representations, let alone low-dimensional parametric approximations.

The path of real wages displayed in the simulations tracks the population distribution quite closely: there does not appear to be a substantial degree of inter-regional interaction which could in principle result in wage patterns that differ in shape from the population patterns. This suggests that at the parameter values chosen, the ef-

fects emphasized by Krugman (1996) in which trade costs can cause global patterns by inducing a strong response at a particular range of sinusoidal frequencies, do not appear to be a major influence. Breaking the real wage down into a nominal wage and a price index, as displayed in Figure (2.3.5), shows that both components contribute essentially similar response shapes, with high population areas facing both higher nominal wages due to increasing returns production technology and lower costs of living as these areas are also located closer to more productive regions and so face a lower shipping cost for the bundle of goods that they consume. The combination of these effects yields what appears to be an extremely strong positive impact of population shocks on local wages.

As overviewed in Piyapromdee (2015), the empirical literature which has examined the local wage response to plausibly exogenous migration flows finds a range of responses from small positive effects on wages to moderate negative effects on wages for some population subgroups depending on a variety of specification and data issues which the current model is not rich enough to capture, but large positive effects do not seem to be empirically plausible. As this is the main economic mechanism inducing agglomeration in Krugman-style geography models, this suggests a potentially important empirical shortcoming. A number of plausible changes may help to reconcile the agglomeration effects with this counterfactual prediction regarding wages. First, the wage evidence mostly describes short run impacts, while geography models have been used to describe changes over decades or even longer, and so may omit certain features of the adjustment process. The assumption of increasing returns, which induces a positive wage response to a local increase in labor supply, derives from free entry of firms producing new varieties. If entry response takes time, or if production requires capital accumulated locally, increasing returns may be a better description of long run than immediate behavior. The agglomeration externalities may also in part come through endogenous response of non-wage amenities, as in

Diamond (2016), which would still induce agglomeration but would not be measured in wage response. The precise effect of such modifications on the dynamics may be nontrivial, as any change in the speed of wage response will also induce a qualitative as well as quantitative change in the response of migration flows.

2.4 Spatially Inhomogeneous Case

In order to examine the behavior of population, wages, and welfare in situations where locations are not ex ante identical, it is necessary to move from closed form to numerical characterizations of dynamics. In realistic geographic settings, there are a variety of sources of persistent heterogeneity which introduce asymmetries in the response of economic activity and population across regions in addition to the possibility of spatially inhomogeneous shocks. I will consider in particular three sources of heterogeneity, though others are possible even in the context of the model presented above. These are differences in trade costs between location pairs, differences in migration costs between location pairs, and patterns of diffusion across space of innovations affecting desirability of different locations (either through wages or amenities) which differ across location pairs. Differences in trade costs may reflect policy, as in the presence of tariffs or other trade barriers, geographic features which influence transportation costs, either natural such as rivers or mountains, or manmade such as highway or railroad networks, or information frictions. Migration costs may arise for similar reasons. Sources of heterogeneity in diffusion vary depending on the source of variation in desirability of different locations. For policy shocks it represents patterns of policy spillovers between locations, for weather or natural variation it reflects the physical processes underlying these dynamics, and for productivity it comes from patterns of productivity spillovers. From the perspective of the model, these sources of heterogeneity can be thought of as anything which affects the shape of the map

Figure 2.3.3: Simulated Geographic Equilibrium: Amenities and Welfare

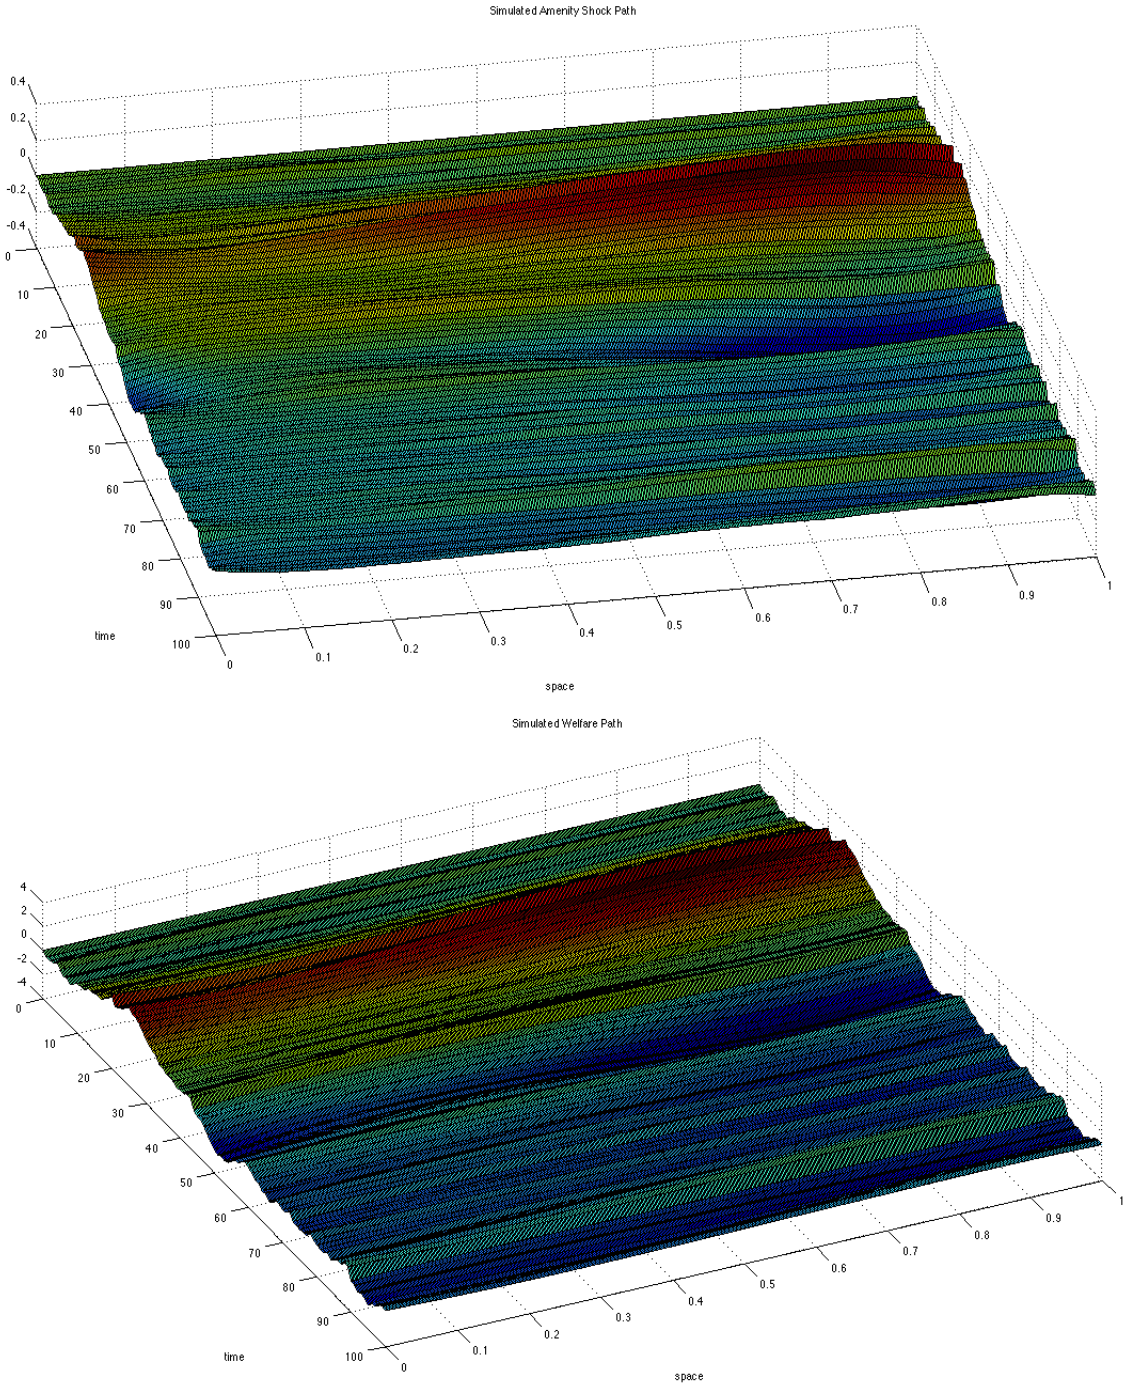

Figure 2.3.4: Simulated Geographic Equilibrium: Population and Real Wages

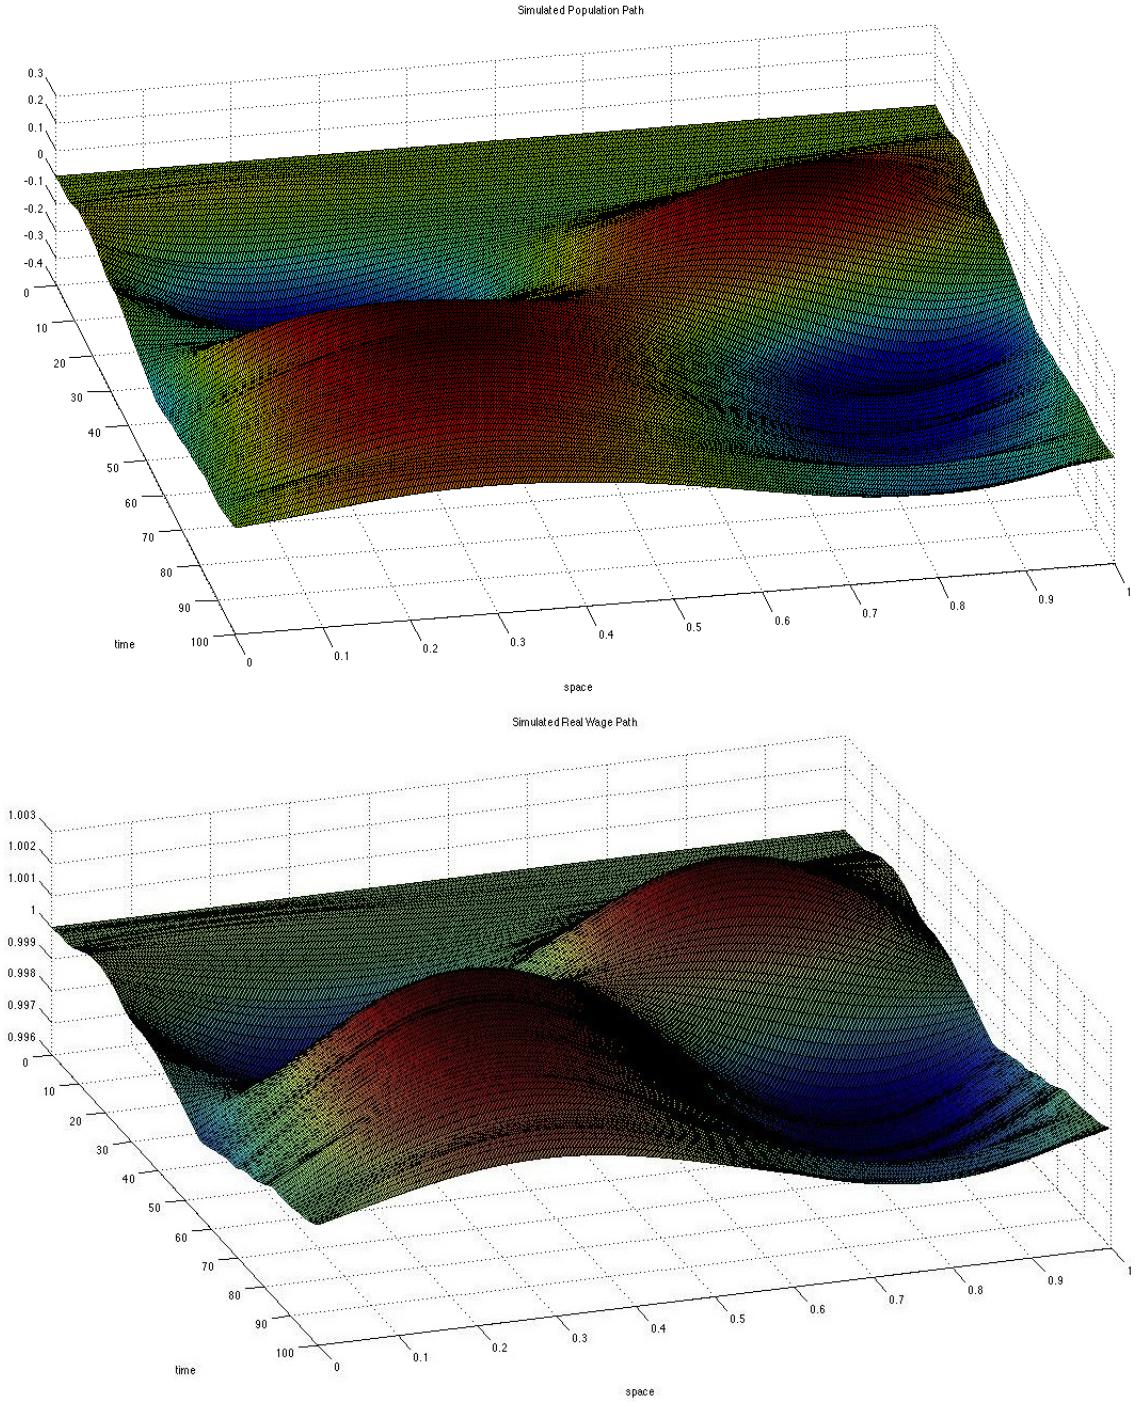

Figure 2.3.5: Simulated Geographic Equilibrium: Nominal Wages and Local Price Index

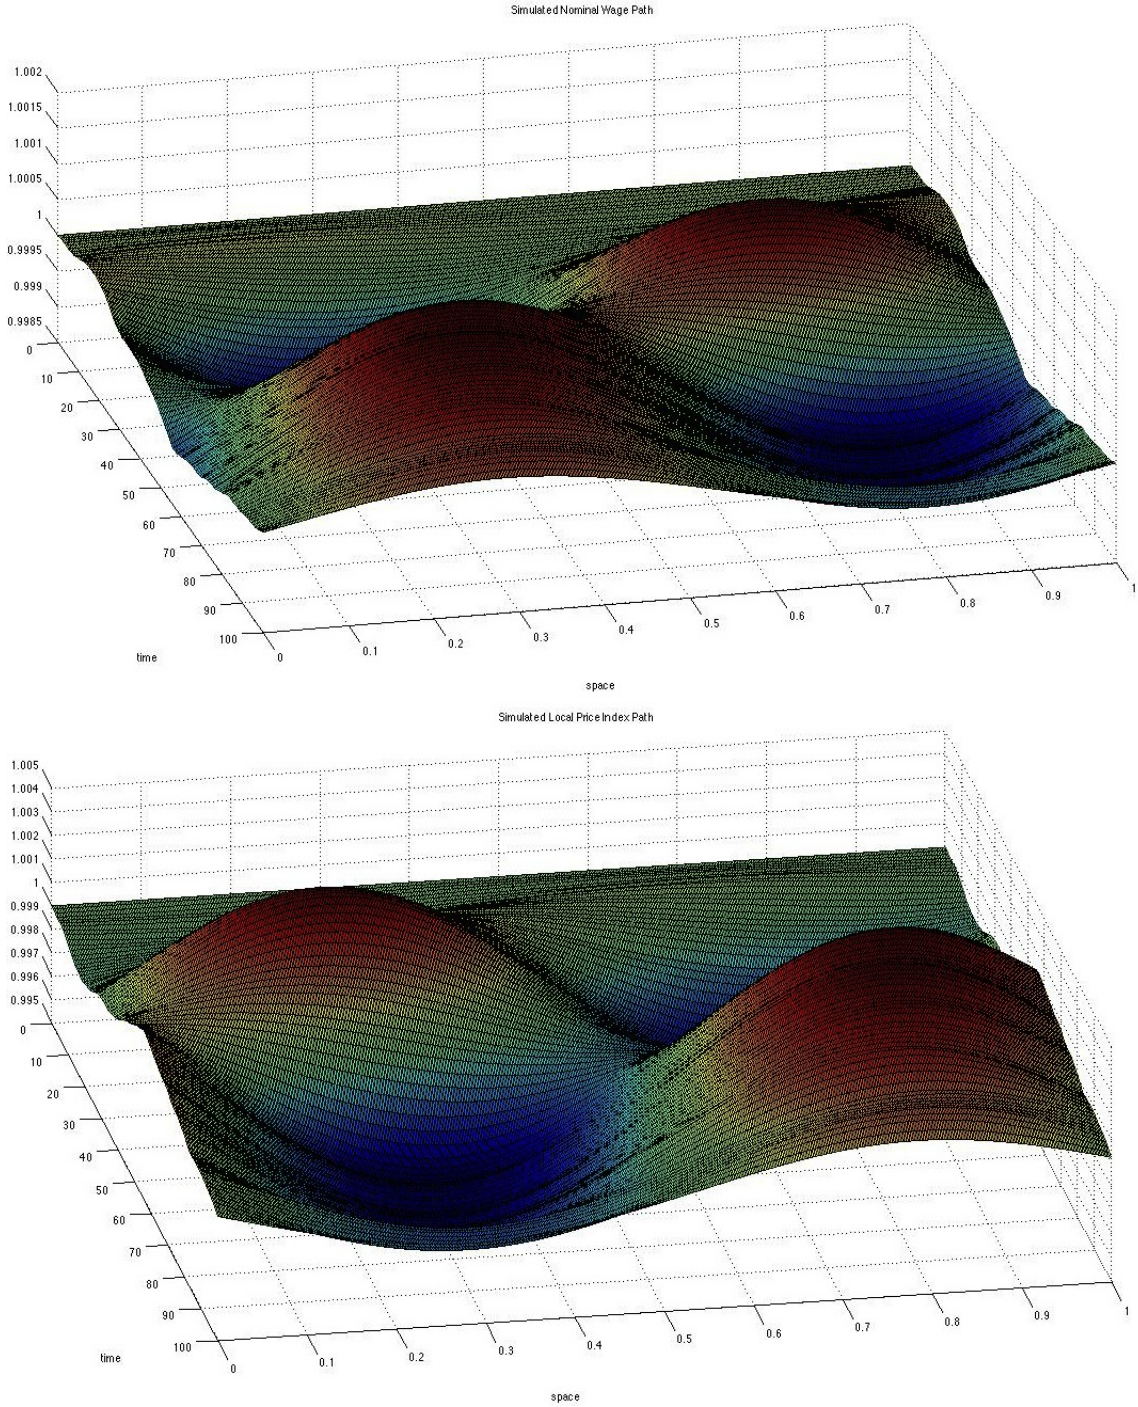

$\omega(\lambda(x))$ between population and welfare, anything which changes the shape of the law of motion of population given differences in welfare, and anything which affects the shape of the shock transition process $\Gamma[\nu]$. As such, these variations may also reflect many other sources of spatial inhomogeneities, such as preference or technology differences across regions.

One notable implication of heterogeneity across regions is that one may no longer assume that a spatially homogeneous steady state exists, and the shape of the steady state distributions is determined by the nonlinear features of the model. This affects both the point around which the model is linearized and also the values of the functional derivatives at that point. While equation (2.8) still characterizes the structure of the derivatives around the steady state, the operators P, $\frac{d\omega}{d\lambda}$, and Γ composing this solution differ. P is now given by $\int \frac{1}{f(x,\bar{V}(.))} \exp(c(x',x) + \beta\bar{V}(x')[.]dx'$, where $\bar{V}(.)$ is no longer a constant, and so heterogeneity in the value induces an additional source of asymmetry in the evolution of population. Effectively, this introduces a drift towards persistently high value locations and away from persistently low value locations which ensures that in addition to the gradual diffusion introduced by idiosyncratic population movements which smooths out temporary fluctuations in population, there is also systematic movement towards high value locations whenever the population distribution is insufficiently concentrated in these high value regions. The term $\frac{d\omega}{d\lambda}$ takes a slightly more complicated form, now no longer separable into frequency bands reflecting the systematic influences of locations at fixed distances, as the effective distances are made irregular and origin specific both by the presence of variable trade costs and also by differences in steady state population and wages, which affect the response of local labor and product markets.

To calculate steady states numerically, I propose a heuristic iterative procedure, "Algorithm" (3), based on contraction mapping principles, to solve the system of nonlinear operator equations generating the steady states. In general, such iterative

methods require that the operators be contractions, or satisfy similar conditions, to be guaranteed to converge, in addition to the obvious requirement that a steady state exists to converge to. These properties are in general parameter dependent: Allen & Arkolakis (2015) provide a characterization of existence and uniqueness of a steady state in a model which is essentially identical to the one provided here, up to a change of variables. It should also be noted that even in the absence of guarantees, a fixed point induced by the method will approximately satisfy all the equations for a steady state, up to a user-determined tolerance ϵ. In practice, for all parameter values tested, the method converges and does so rapidly (in no more than 5-6 seconds, with a tolerance of 10^{-15}) so long as the trade and migration costs are strictly nonnegative.

While a realistic calibration is outside the scope of this chapter, it is possible to construct a broad variety of reasonably complex spatial patterns, albeit still on the circle, by using randomly generated cost functions drawn from a stochastic process. To generate patterns of costs which might reflect certain features of trade and transportation costs, I consider cost functions constructed by the methods of Allen & Arkolakis (2014), who consider bilateral trade costs functions given by choosing a path between locations which minimizes the weighted effective distance computed as the integral over the path of location-specific transportation costs along that path. This results in a symmetric cost function reflecting optimal choice of transportation route. To construct bilateral cost functions, I therefore draw location-specific transportation costs for each location on the circle from a strictly nonnegative stochastic process, then for each location pair find the cost reflecting the optimal path. I choose for the stochastic process a small positive constant (0.2 in the simulated examples) plus the square of a Gaussian process with exponential covariance kernel $Cov(x, z) = \exp(-d(x, z))$, as this induces a path which is bounded away from 0 and with probability 1 is Hölder smooth of any finite order, by Kolmogorov's continuity criterion (see, e.g. Hairer (2009) Theorem 3.17).

Algorithm 3 Iterative Method to Construct Steady State

Input: Model Parameters, Trade costs $\tau(x, z)$, Migration Costs $c(x, z)$, K number of grid points, $\epsilon > 0$ tolerance
Output: Approximate discretized steady state functions $\bar{\lambda}(x)$, $\bar{V}(x)$, $\bar{\omega}(x)$

1. Discretize operators (2.1)-(2.7) along K-point dyadic grid
2. Conjecture strictly nonnegative $\lambda_a(x)$
3. Iteratively solve Hammerstein equations (2.3), (2.4), (2.5), (2.6)
   (a) Guess initial $w_0(x)$
   (b) Apply (2.4), then (2.3), then (2.5) with $\lambda(x) = \lambda_a(x)$ to $w_n(x)$ obtain $w_{n+1}(x)$
   (c) Iterate until $\max |w_{n+1}(x) - w_n(x)| < \epsilon$
   (d) Apply (2.6) to $w_{n+1}(x)$ to obtain $\omega_b(x)$
4. Iterate discretized Bellman equation $V_{n+1}(x) = T(V_n(x)) := \omega_b(x) + \log f(x, V_n(x)) + \gamma$ until $\max |V_{n+1}(x) - V_n(x)| < \epsilon$ to obtain $V_b(x) = V_{n+1}(x)$
5. Iterate discretized transition equation (or calculate Frobenius eigenvector) $\lambda_{n+1}(x) = \int \frac{\exp(c(x,z)+\beta V_b(z))}{\int \exp(c(x,z)+\beta V_b(z))dz} \lambda_n(x)dx$ until $\max |\lambda_{n+1}(x) - \lambda_n(x)| < \epsilon$ to obtain $\lambda_b(x) = \lambda_{n+1}(x)$
6. If $\max |\lambda_a(x) - \lambda_b(x)| > \epsilon$ $\lambda_a(x) \leftarrow \lambda_b(x)$, Return to 3, else go to 7
7. $\bar{\lambda}(x) \leftarrow \lambda_a(x)$, $\bar{V}(x) \leftarrow V_b(x)$, $\bar{\omega}(x) \leftarrow \omega_b(x)$

Figure (2.4.1) displays the resulting steady state function values of $\bar{\lambda}(.)$, $\bar{V}(.)$, and $\bar{\omega}(.)$ for independent random draws of cost functions $c(x', x)$, $\tau(x, z)$ from the process described above at 512 grid points, scaled by $4/512$ for $\tau(x, z)$ and $4/(3 \times 512)$ for $c(x', x)$, and with all other parameters identical to those used in the numerical simulations for the spatially homogeneous case. As can be seen, trade and migration costs induce concentration of population in steady state. Population is generally denser in areas with high steady state welfare and lower in areas with low steady state welfare, though the correlation is imperfect due to migration costs which differ by location. Also notable is that the distribution of wages and the distribution of welfare are far from perfectly correlated, even though in steady state welfare is determined entirely by the expected present discounted value of wages, which are constant over time. The reason for this is that, due to idiosyncratic shocks which induce movement in steady state, individuals do not stay in a single location. As a result, the value of a location includes not only its wages, but also the relative likelihood of moving to a different area and the migration costs of doing so. This is one example of how idiosyncratic volatility may make aggregate statistics misleading as indicators of welfare or individual decisions, because individuals also consider the option value of responding to idiosyncratic risks.

The spatially inhomogeneous case also produces quite different implications for responses to aggregate shocks. To introduce richer dynamics, it is possible to utilize a parameterization of the law of motion for the amenity shock process $\Gamma(x, z)$ which is also spatially inhomogeneous. For the purpose of constructing example IRFs and simulations, I choose the function $\Gamma(x', x) = \frac{k_\Gamma}{\int \exp(-\frac{1}{2\sigma_\Gamma^2}(q(x',x))^2)dx} \exp(-\frac{1}{2\sigma_\Gamma^2}(q(x', x))^2)$, with $q(x', x)$ drawn independently from the same process generating $\tau(x, z)$. This parameterization preserves the general shape and magnitude, but induces spatial inhomogeneities in the degree of diffusion. It may be interpreted as amenities also following a least cost diffusion path with respect to a randomly generated measure

Figure 2.4.1: Simulated Spatially Inhomogeneous Steady State $\bar{\lambda}(.)$, $\bar{V}(.)$, and $\bar{\omega}(.)$

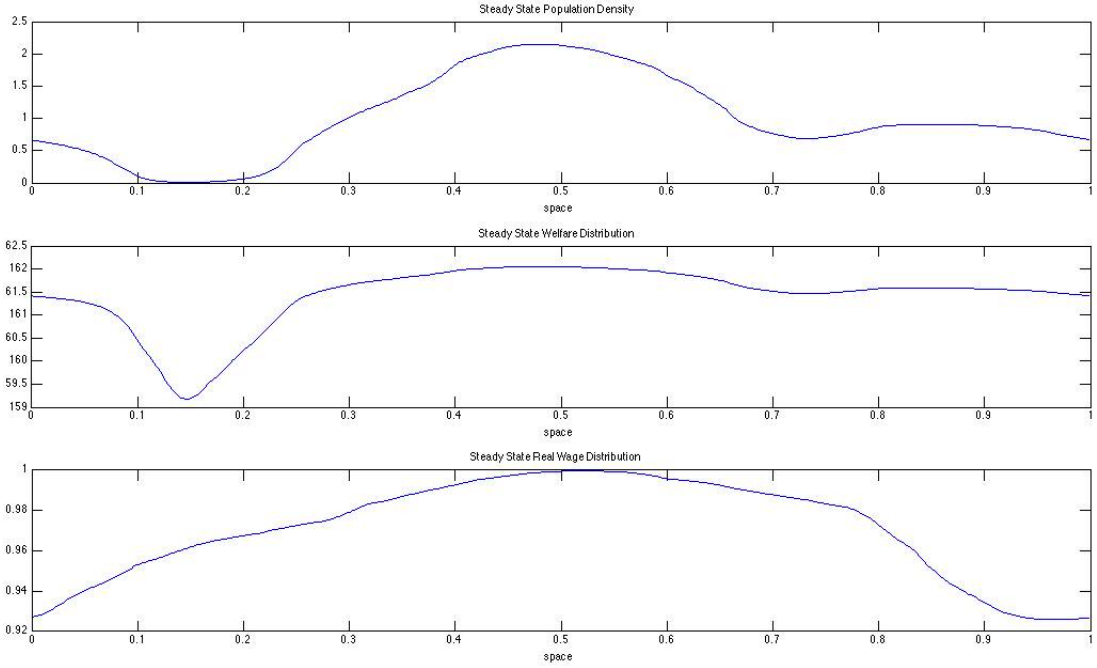

of local obstacles. In order to evaluate the behavior of the model in this case, it is no longer sufficient to construct IRFs to shocks localized at a single point. As a simple exhibition of the effects of inhomogeneity, I instead construct IRFs to the same spatially Gaussian shock as used in the homogeneous case, but now centered at 0.1, 0.5, and 0.9. Graphs of Impulse response functions are displayed in figures (2.4.2), (2.4.3),(2.4.4),(2.4.5),(2.4.6), and (2.4.7), respectively.

Notable features of the responses of $\nu_t(.)$ include that while behavior begins localized and spreads out in each case, the spread is biased in one direction at 0.1 and 0.9, and spreads out more rapidly at 0.5. This may be part of the reason for the difference in welfare responses. While the initial welfare response contains a highly localized spike in the 0.1 and 0.9 cases, superimposed on a more diffuse bump around the edges which does not differ substantially between the two cases, it is substantially more diffuse in the 0.5 case. More interestingly, despite the differences in the shocks and welfare responses, the population response is qualitatively similar in all cases: population grows near the edges of the map and shrinks at the center. This

is particularly remarkable in the 0.5 case since the area which loses population is also the area which gains the most in amenity value. This seems to reflect a global structure induced by the costs to trade, migration, and diffusion. The structure can also be seen to some extent in simulations of the process, drawn with the same pattern of shocks as in the spatially homogeneous case, displayed in Figures (2.4.8) and (2.4.9), which display the simulated deviations from steady state. While the pattern of amenity shocks is quite close to that from the spatially homogeneous case, as is to be expected from the relatively similar diffusion law, the welfare distribution displays greater variability in the center and at the edges than in between, a pattern also observed in the impulse responses of welfare to shocks, and the population distribution also displays this pattern, substantially different from that in the homogeneous case, with the large decreases and increases in population in response to shocks shifted to the right, reflecting the greater unity of response of the regions near the edges in this case.

Another notable feature of the impulse responses in this case is that while responses to aggregate shocks can show substantial differences based on the location of the effects of the shock in the short run, in the medium run, many types of shocks display variation along a similar characteristic shape. The reason for this is that the long run behavior is dominated by the eigenvectors corresponding to the largest eigenvalues of the transition operator. In the homogeneous case, this usually corresponds to very low frequencies. In the case with heterogeneity, this shape may be less regular. It is also worth noting that this low dimensional variation is not necessarily aligned with the functional principle components of the data generated by this model, which are a function of both the shock covariance and the transition law. This suggests that in high dimensional settings such as spatial models the information needed to determine policy or causal effects may differ substantially from the data most useful for prediction.

Figure 2.4.2: Impulse Response of $\nu_t(x)$ to $\varepsilon_0(x) = \exp(50000(x-0.1)^2)$

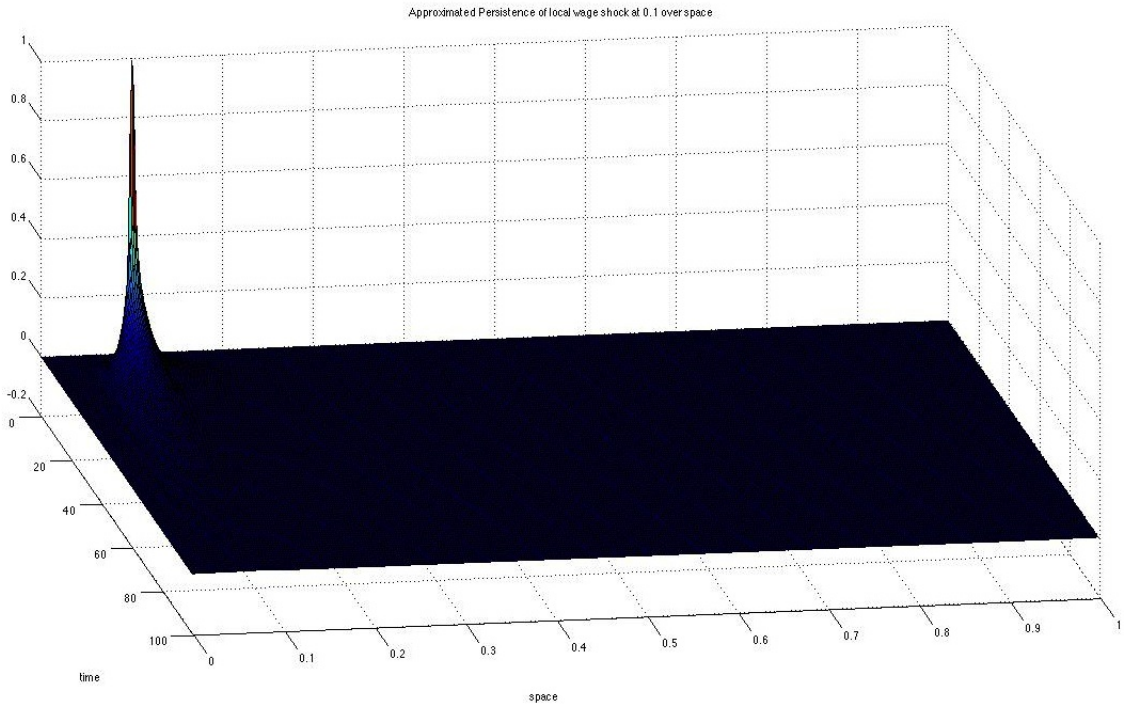

2.5 Conclusion and Future Work

By incorporating a forward-looking migration decision into a quantitative model of economic geography, it becomes possible to formulate hypotheses regarding the spatial dynamics of population and economic activity in response to spatially heterogeneous disturbances and evaluate the welfare implications of these movements at local and global levels. In this chapter, we have provided a first attempt at investigating these dynamics. The results from the dynamic model demonstrate the potential role of anticipation of future changes in regional economic status in modulating both the magnitude and direction of population movements, and also the effect of idiosyncratic heterogeneity and steady state migration in spreading the welfare impact of local changes over space, with migration acting as a form of insurance against otherwise non-insurable regional shocks. They also show how the immediate and delayed effects of a shock can differ substantially in a model with rich spatial structure, with the immediate effects driven primarily by the properties of the shock itself, but the impact

Figure 2.4.3: Impulse Responses of $\lambda_t(x)$ & $V_t(x)$ to $\varepsilon_0(x) = \exp(50000(x-0.1)^2)$

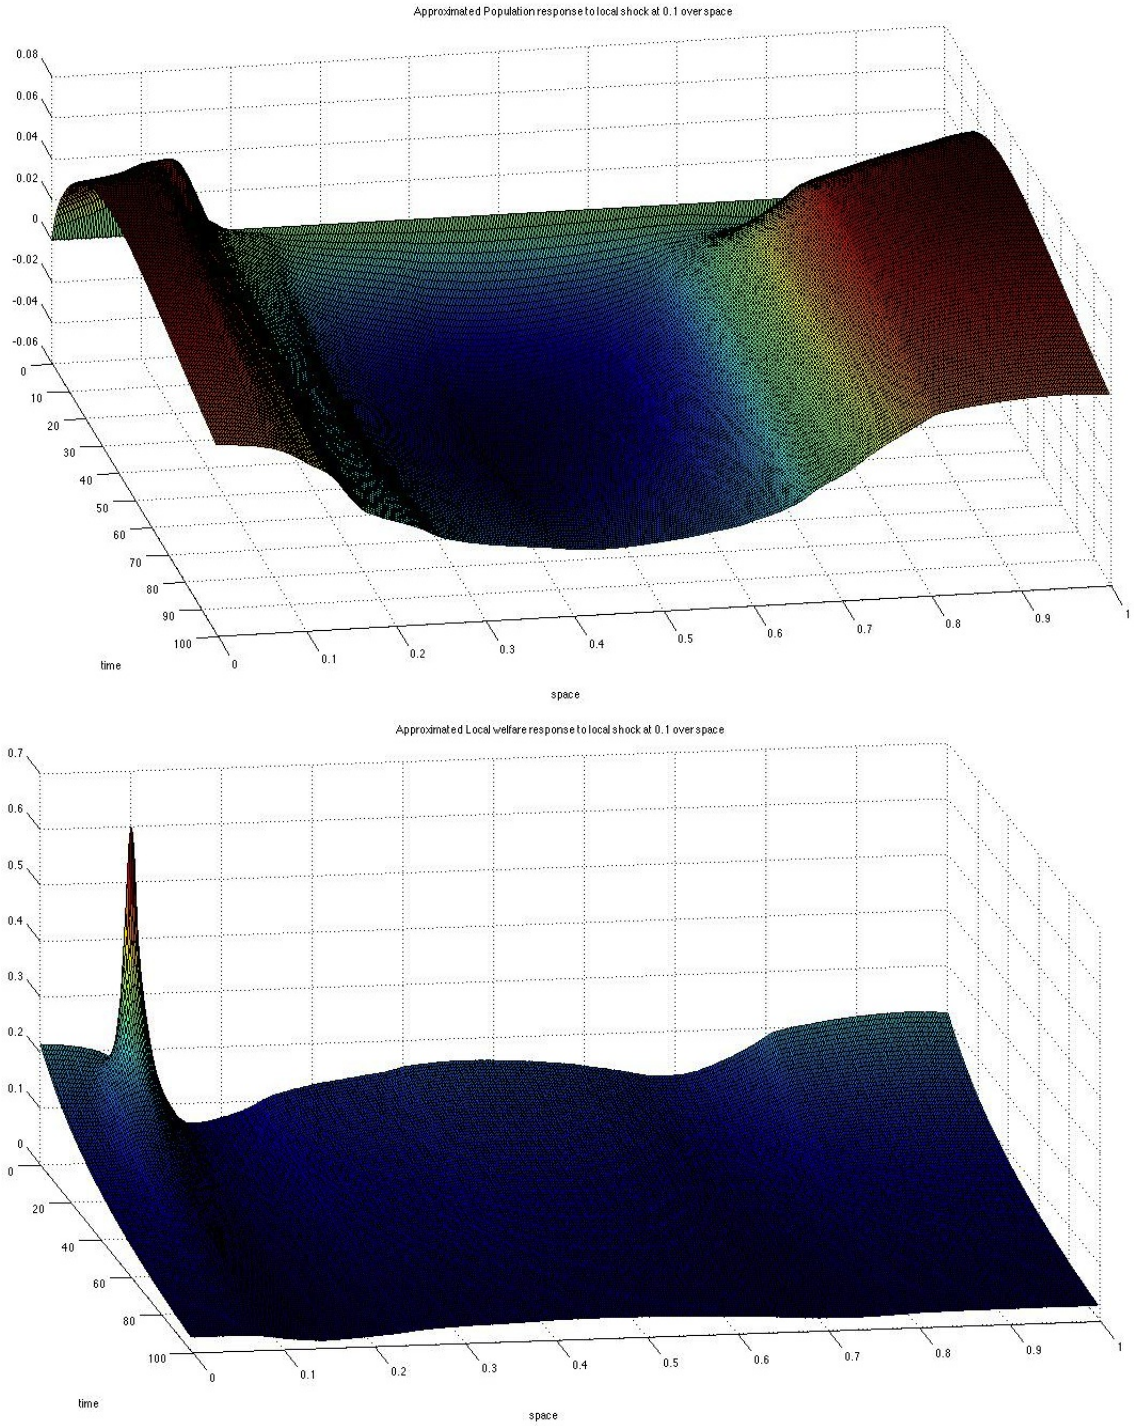

Figure 2.4.4: Impulse Response of $\nu_t(x)$ to $\varepsilon_0(x) = \exp(50000(x-0.5)^2)$

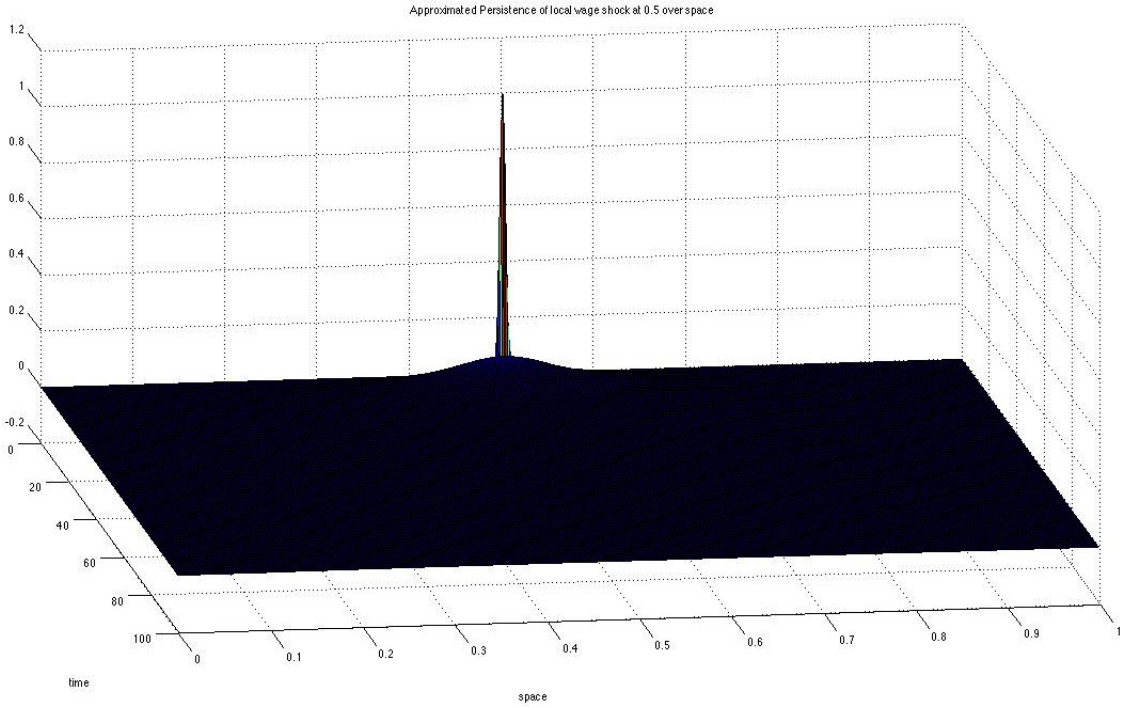

over the long term shaped substantially by the global features of the economy which channel variation along certain spatial patterns, determined by the eigenfunctions of the laws of motion.

A fully dynamic formulation also permits evaluation of hypotheses regarding the origins of spatial inequality which rely on cumulative processes of divergence. In the popular Krugman model of trade featuring trade costs, differentiated goods, increasing returns, and traded and non-traded sectors, the interplay of agglomeration and congestion externalities, steady state equilibrium features spatial concentrations of population and economic activity whose location is determined in a nontrivial manner by the global interplay of these forces with even small ex ante geographic differences. This global equilibrium furthermore exhibits comparative statics which suggest that this distribution may be sensitive to small changes in fundamentals, as the presence of increasing returns induces a positive wage response to increases in population which reinforces rather than dampens the initial population change. While particular form

Figure 2.4.5: Impulse Responses of $\lambda_t(x)$ & $V_t(x)$ to $\varepsilon_0(x) = \exp(50000(x-0.5)^2)$

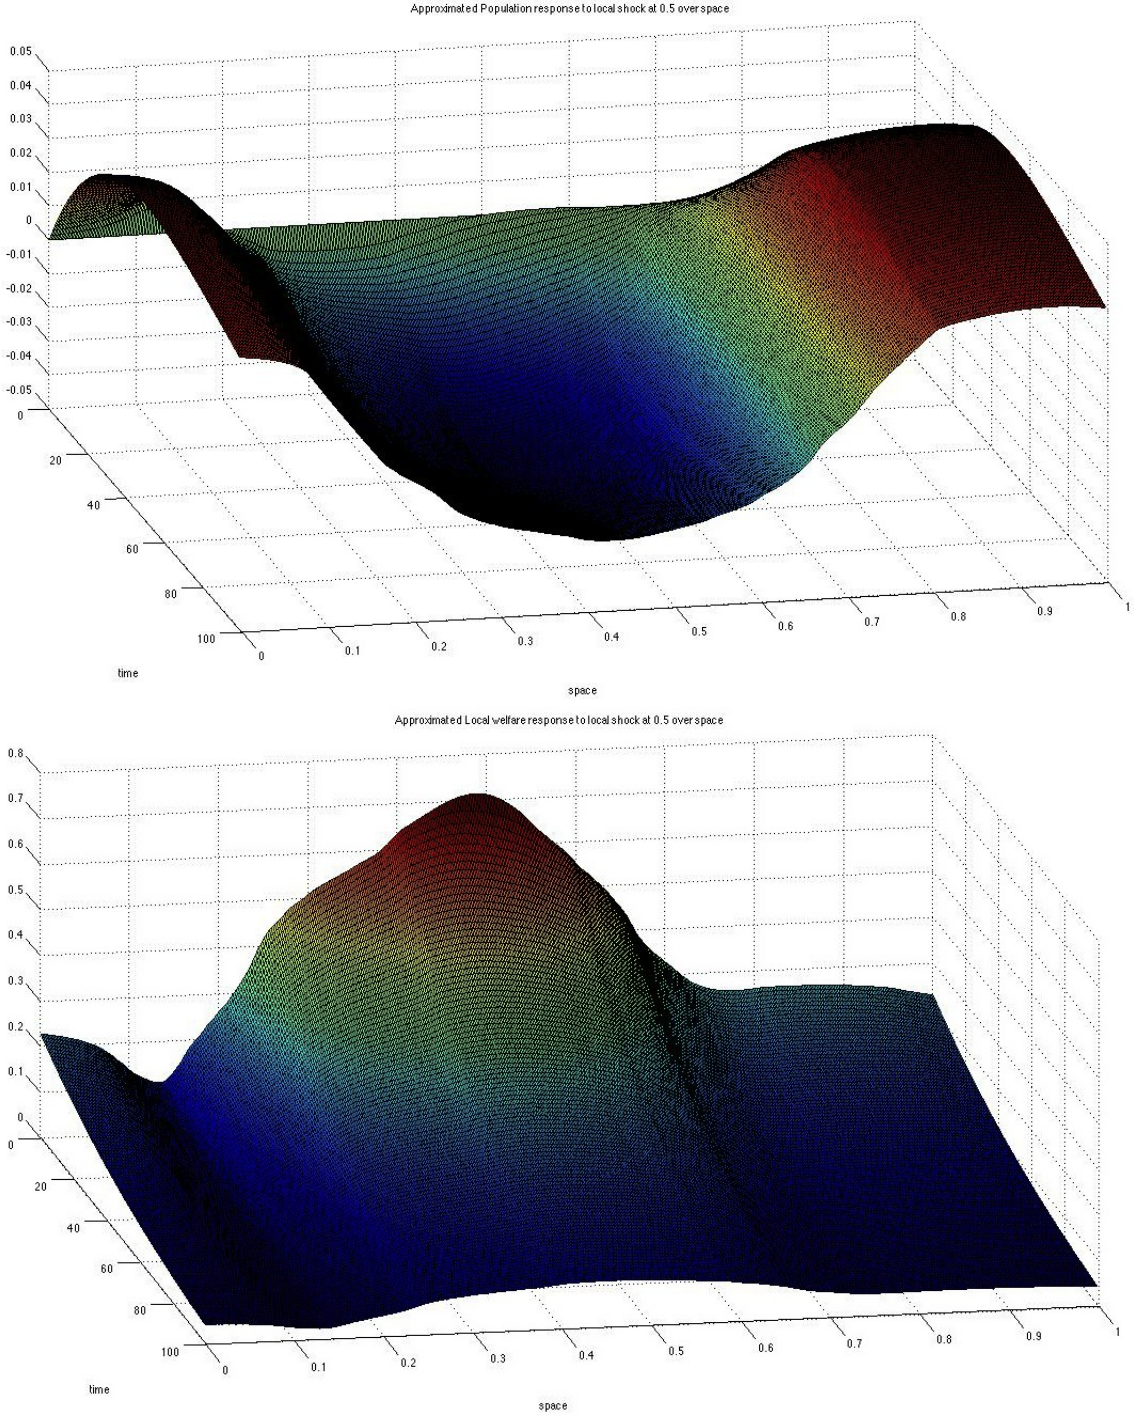

Figure 2.4.6: Impulse Response of $\nu_t(x)$ to $\varepsilon_0(x) = \exp(50000(x-0.9)^2)$

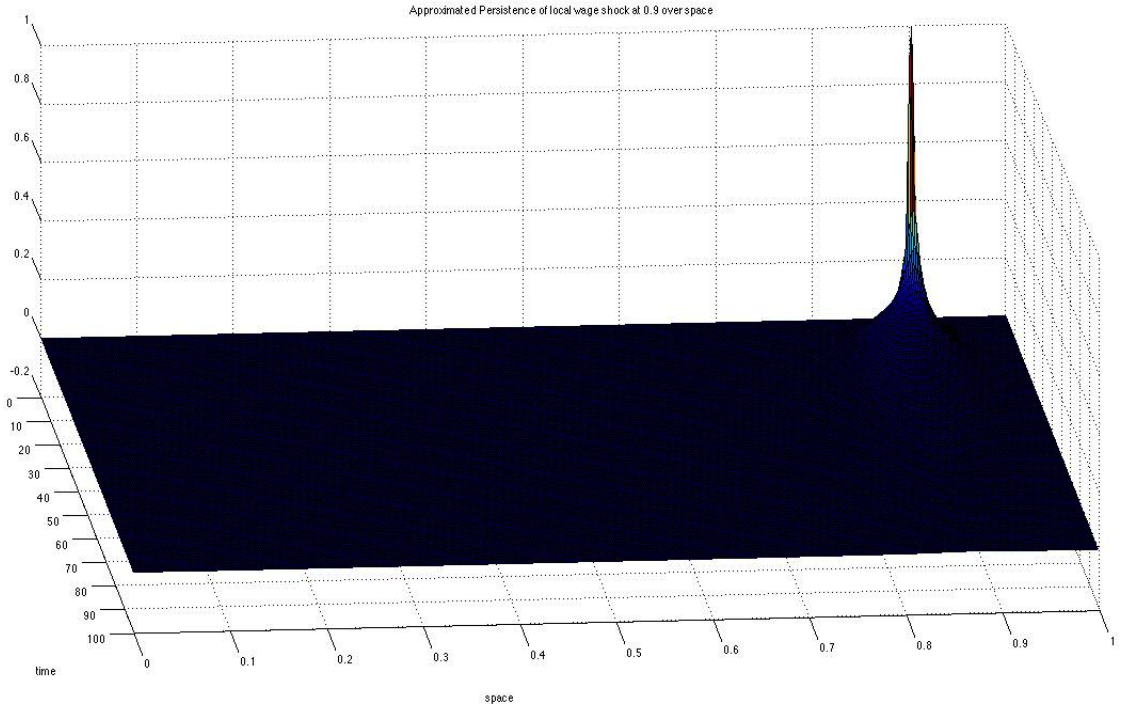

of the response is governed by the structure of production and trade imposed by this model, a wide variety of sources of pecuniary and nonpecuniary spatial agglomeration externalities have been documented since Marshall (1890) which likewise suggest that population changes may be self-reinforcing.

A dynamic perspective adds a number of caveats to this class of static accounts of regional heterogeneity. First, migration costs in the presence of steady state migration flows tilt the response of population to persistent wage differences both directly by keeping people in some regions longer and indirectly by changing the desirability of different locations. Moreover, as the comparative static effect is derived essentially under the assumption of a permanent change, it need not reflect the impact of a temporary shock, in which case, under costly migration, population may only respond partially, and also temporarily. While this does not quite refute the 'history-dependence' hypothesis, that small temporary shocks induced a self-reinforcing cycle leading to long term regional differences, it does suggest that it is not quite as ro-

Figure 2.4.7: Impulse Responses of $\lambda_t(x)$ & $V_t(x)$ to $\varepsilon_0(x) = \exp(50000(x-0.9)^2)$

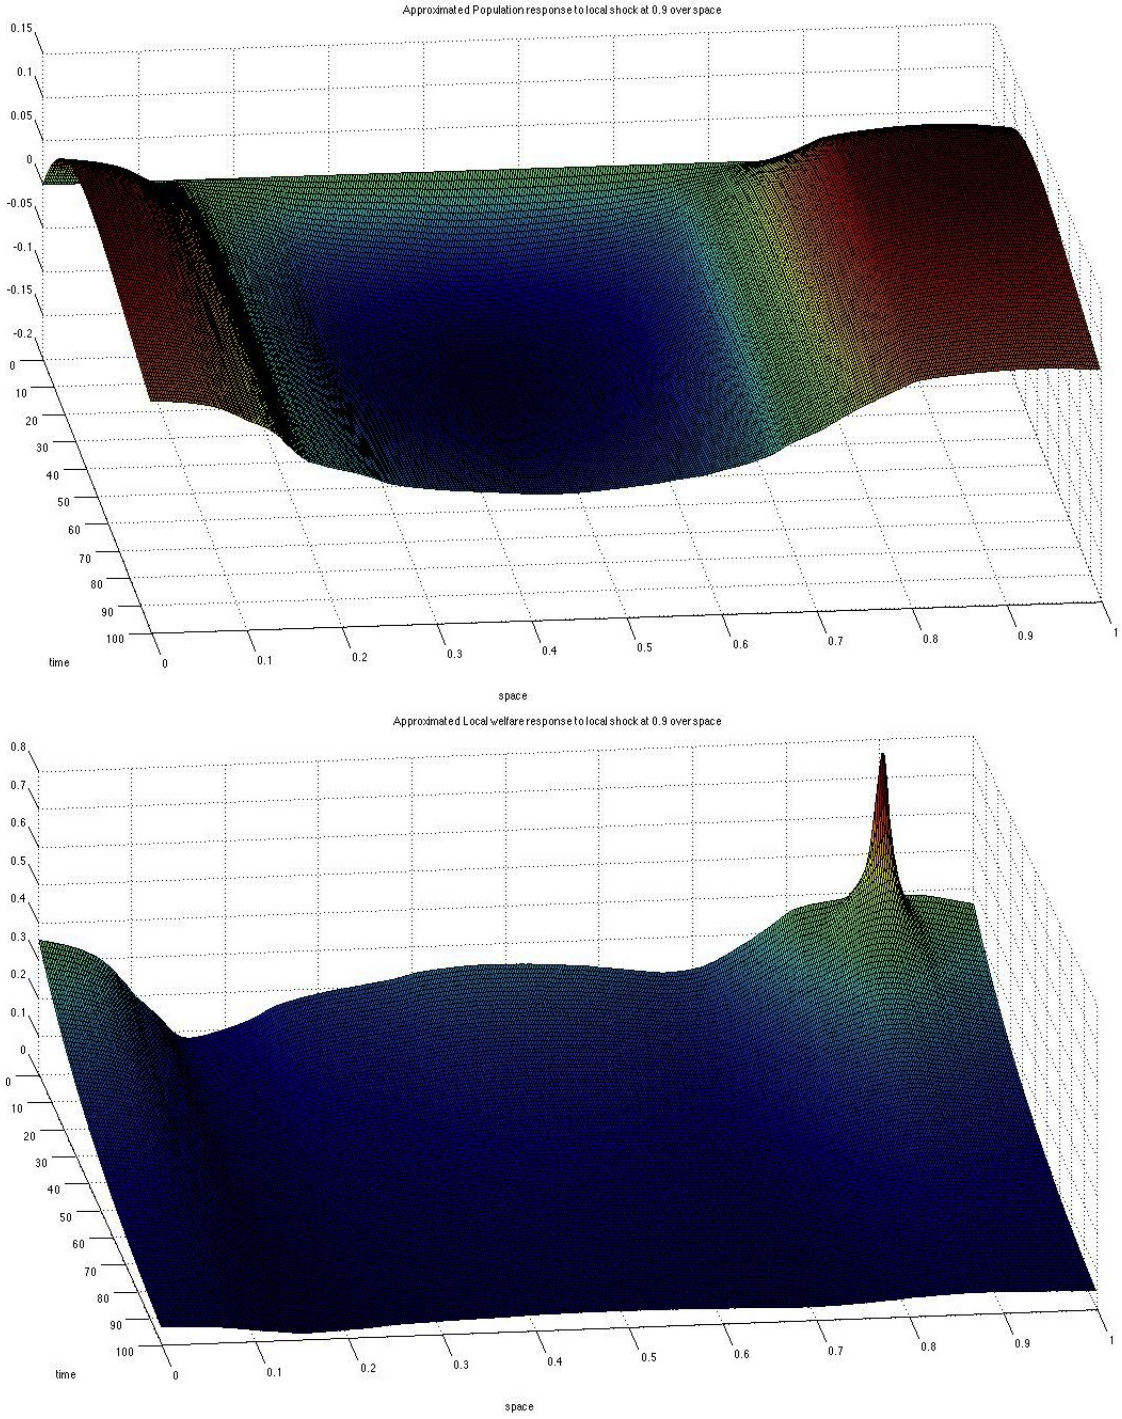

Figure 2.4.8: Simulated Inhomogeneous Geographic Equilibrium: Amenities and Welfare

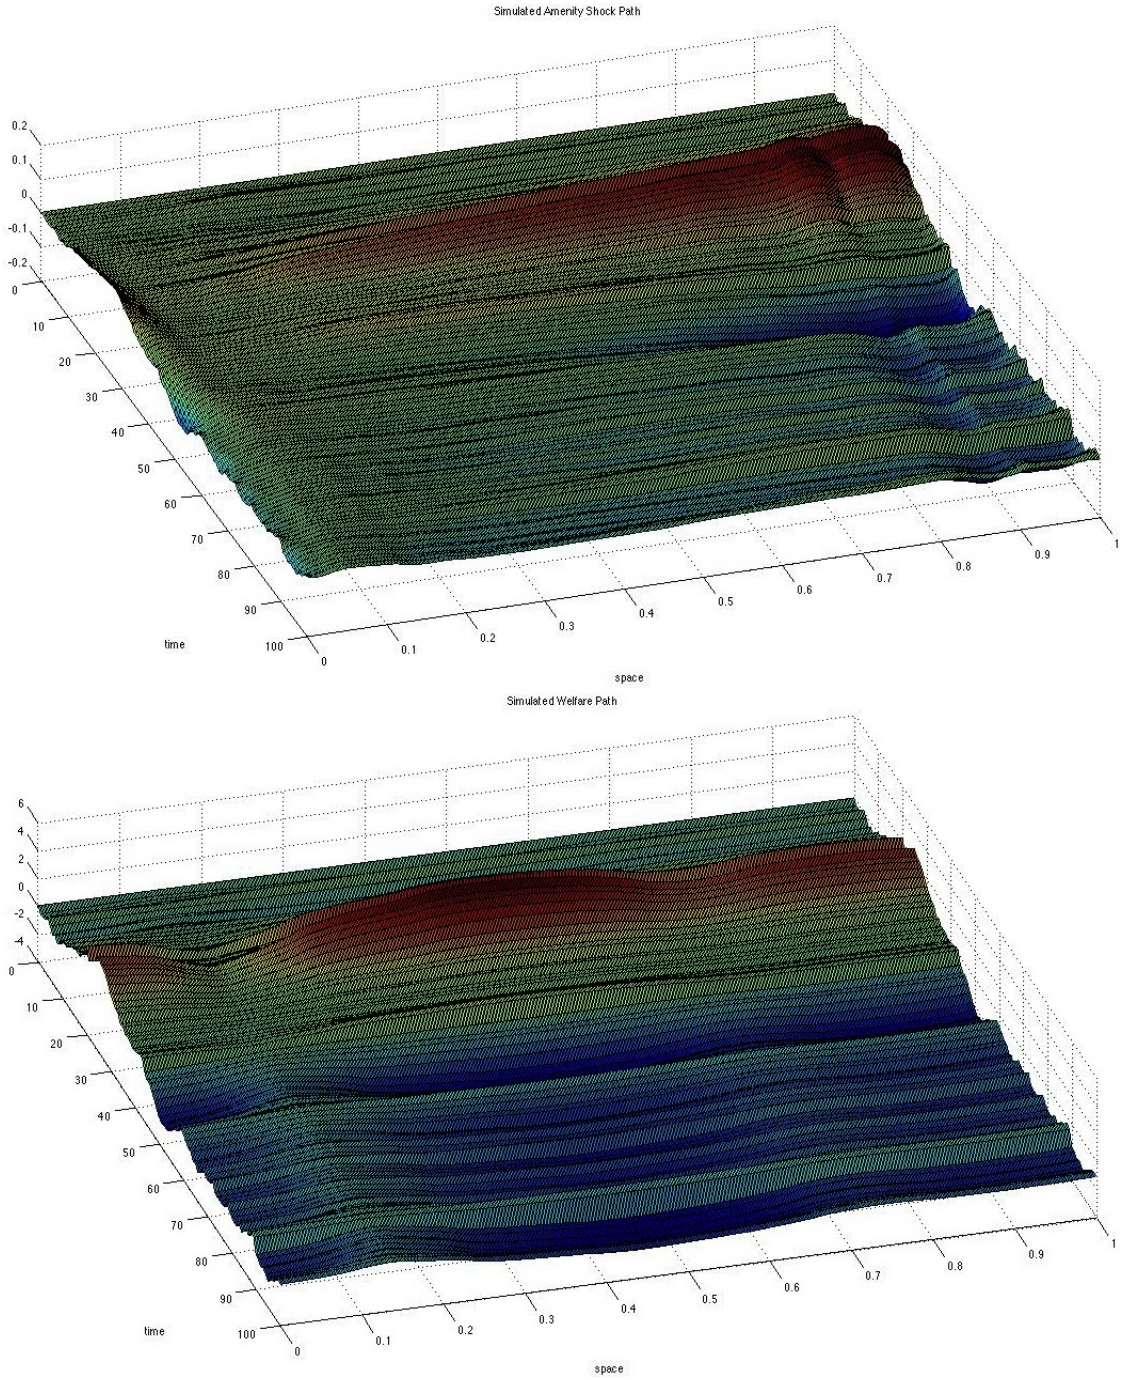

Figure 2.4.9: Simulated Inhomogeneous Geographic Equilibrium: Population

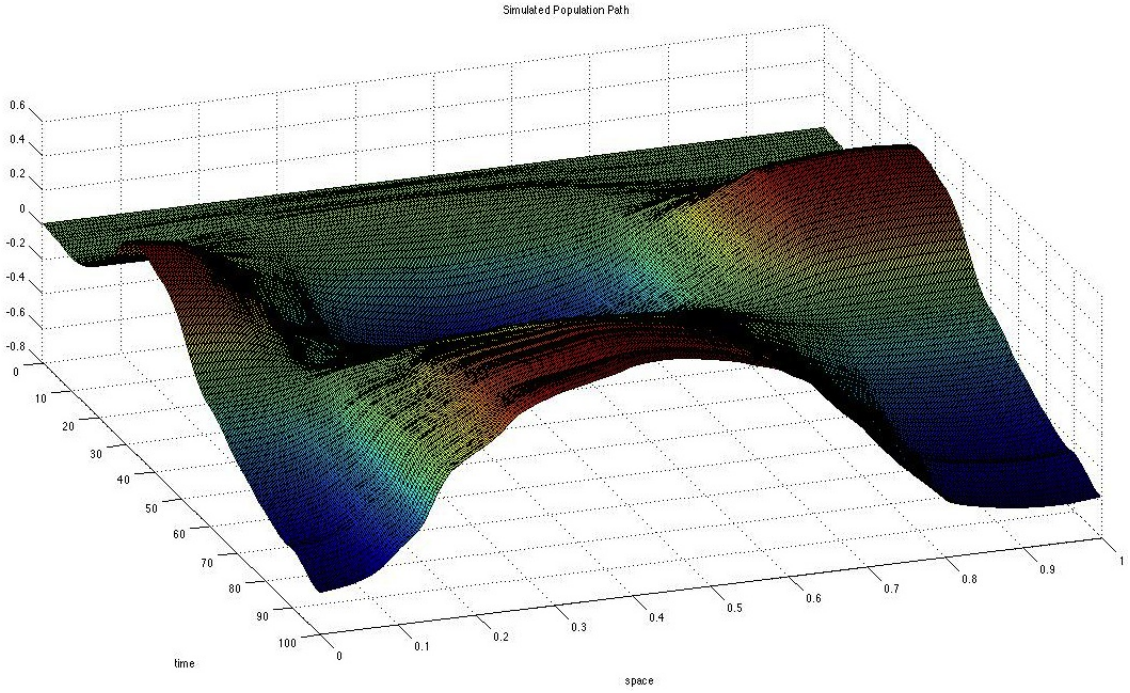

bust an implication of this class of models as previously hypothesized. In order to obtain permanent differences from temporary shocks, one must either assume that the self-reinforcing externalities are substantially larger than those needed previously to induce such an explosive process, impose a substantially different decision process which may not be forward looking, or to rely on the possibility of multiple steady state equilibria, in which case there may exist dynamic rational expectations equilibria which differ from the local description of dynamics offered in this chapter and may involve movement not confined to the neighborhood of any one steady state.

The model does provide a clear and testable empirical implication of local, if not global instability, that can be examined without reference to the sources of this instability. As the equilibrium can locally be described as a functional linear process, instability takes the form of a nonstationary or explosive subspace in the time series of population density functions, analogous to the unit root subspace described in Chang *et al.* (2014). In the homogeneous case of Section 2.3 this corresponds to nonstationarity at a particular frequency: in more general cases as described in Section 2.4, it

need not correspond to a particular basis function. Explicit tests of this hypothesis may require extension of tests for explosive roots to the functional time series setting.

While the model and procedures introduced here represent a useful foundation for studying the spatial dynamics of population and economic activity, precise answers will require substantial further work. From a theoretical perspective, evaluation of the relationship between history dependence and increasing returns may require extending from the local approach described here to a global approach which can be used to evaluate dynamics in the case when multiple steady states may exist. From an empirical standpoint, the model should be evaluated using empirically plausible parameter values estimated using geographic data over different time periods to assess the plausibility and magnitude of implied responses to regional shocks. It is likely that matching these empirical responses may require reevaluation of some of the dynamic features of the model. In particular, additional sources of individual and local heterogeneity might be incorporated, and the structure of production and trade, which is essentially static, may require the incorporation of dynamic elements, such as capital or persistent relationships, to capture the time-series structure. In all cases, the representation of the solution of dynamic models of economic geography derived here as functional stochastic processes provides a reduced form framework for evaluating and testing the empirical implications.

Chapter 3

Solving Ill-posed Function-Valued Rational Expectations Models

In linear dynamic stochastic economic models with function space variables, a rational expectations solution taking function-valued inputs to function-valued outputs is defined by a decomposition, analogous to the generalized Schur decomposition of matrices, of a set of operator equations into components. A computationally feasible approximation of the solution may be constructed by projection on a set of basis functions. But when the operators defining the equilibrium conditions are not compact, the solution of the approximate system may fail to converge to the solution of the true system. This failure arises from multiple breakdowns of continuity in the map from approximation to solution. A solution is devised that enables components to be constructed sequentially, applying regularization at each step. Due to the lack of compactness, standard regularization methods for linear ill-posed inverse problems do not suffice to ensure continuity, and inverses are instead constructed by uneven section methods based on the generalized sampling technique of Adcock *et al.* (2014a). Guidelines for tuning parameter selection are provided, and the performance of the algorithm is demonstrated on an example model.

3.1 Introduction: Motivation and Overview

Dynamic economic models with function-valued states, such as heterogenous agents models where the distribution of heterogeneity evolves endogenously over time, are an increasingly popular means of studying high-dimensional economic phenomena. In order to make these classes of models tractable to analyze and compute, it is common to make simplifying assumptions which require that the dynamics have an "essentially" finite-dimensional character, so that a finite-dimensional approximate model describes the dynamics reasonably well. This is the assumption, either implicit or explicit, behind approaches based on approximate sufficient statistics or smooth function approximation. However, there is a substantial range of possible dynamic behavior in infinite dimensional spaces which simply cannot be described uniformly well in a finite-dimensional context, and so may be considered as "essentially infinite-dimensional." In these situations, any finite dimensional description of the system, no matter how large, may be substantially affected by some components which have been left out of the approximation. At best this results in potential sources of fluctuations which do not enter the model description. At worst, the dynamics even of the finite dimensional component may be completely changed through their interactions with the missing elements. Moreover, the class of models which exhibit this behavior is in some sense generic, arising any time the fundamental structure describing the dynamics is non-smooth.

In this paper, I ask what can be said about these kinds of models, and show that while this high-dimensional structure makes some descriptive tasks difficult, a substantial amount can still be learned about their solutions. In particular, while the non-compactness makes uniform approximation of the model by a finite-dimensional representation impossible, the model can still be approximated pointwise. The solution which describes the dynamics implied by the model is not a continuous function

of the pointwise approximation of the model, so the solution problem is ill-posed.[1] However, this only means that a pure "plug-in" type estimator of the true solution will not converge. Instead, by analyzing the properties of the map from model to solution, which consists of multiple interdependent steps, it is possible to construct a more careful approximation, which uses regularization to extract as much information as possible from each step and pass it on to the next step, which again must use regularization to handle the limited information coming from the previous step. The final result is a procedure which eventually recovers the solution, which sometimes may itself be a well-behaved object, even when the model that defines it is not.

Formally, this paper considers a generalization of the class of recursive linearized rational expectations models with function-valued state variables described in Chapter 1, which provides a coherent framework for characterizing the local behavior of many economic models with continuous heterogeneity. The approach in that chapter requires a set of conditions which limit consideration to models which are well-behaved, in the sense that finite-dimensional approximations can be used to build estimates of the operators defining the model which are consistent in a uniform sense, i.e. in operator norm. With this strong property, plug-in estimates of the solution with no further regularization can be shown to converge. There are, however, situations where the condition which ensures that this is possible, referred to as "asymptotic diagonality," need not hold. These cases are intimately connected with the presence of non-smooth or non-compact components of a model, which take input functions which are smooth or well-behaved to output functions which are non-smooth or poorly behaved. Qualitatively, such operators are associated with the phenomena of "scattering" or "turbulence," in which high frequency fluctuations, instead of dissipating or at least remaining isolated at high frequencies, influence behavior at lower frequencies, so that large scale fluctuations cannot be controlled without looking to higher and

[1] See Horowitz (2013) for a survey of ill-posed problems in economics.

higher scales.

3.2 Literature Review

The technical contributions of this chapter draw from a number of literatures. Linearization in function space is an old technique for examining the properties of high dimensional dynamical systems. In statistical mechanics, it is referred to as linear response theory (Sethna, 2006, Ch. 10). The methodology has seen some application in economics, most notably in the economic geography model of Krugman (1996) which simplifies the problem substantially by using adaptive expectations. The extension of this methodology to the case of rational expectations is described in Chapter 1, where qualitatively similar antecedents are also reviewed. The mathematical theory underlying this approach is based on separation into subspaces by a generalized Schur decomposition and numerical approximation thereof, described in Appendix A. The problem of estimating subspaces corresponding to an unknown infinite-dimensional operator is well-studied in mathematics and increasingly so in statistics, with the literature on high-dimensional or functional PCA and CCA. The study of the properties of subspaces corresponding to eigenfunctions or generalized eigenfunctions is the field of mathematics known as perturbation theory. Most of the statistical literature deals with highly structured operators, such as covariance operators, which exhibit desirable properties such as symmetry and compactness, and so permit consistency to be achieved by quantitative bounds on distance such as the Davis-Kahan theorem or analogues. The procedure in Chapter 1 essentially follows this approach, by imposing structure sufficient to apply a set of perturbation inequalities derived in Appendix A.2 and building off of the finite-dimensional results of Stewart (1973), which take advantage of the assumed compact structure. An older literature mainly applied in

physics (see Kato (1976) for a textbook overview and Krieger (2001) for an economic application) applies to a broader class of operators, but investigates parametric perturbations, which is useful when the solution is known up to a scalar parameter, but not for arbitrary infinite-dimensional deviations. While this makes the results not directly applicable to the case of infinite-dimensional function-valued perturbations, the characterizations of subspaces used in this literature apply, and weaker results may be derived directly from these characterizations.

The main inspiration for the approach in this chapter to analyzing perturbations in the absence of the high degree of structure needed for previously applied perturbation methods to work is the 'Generalized Sampling' approach developed by Adcock and Hansen, et. al. in a series of papers mostly about compressed sensing of images viewed as functions rather than a vector of pixels (Adcock *et al.* , 2014b), (Adcock *et al.* , 2014a). The comparison is illustrative: they point to a number of issues which may be generated by ignoring the propagation of bias induced by discretization into pixels, and provide a series of methods for alleviating it.

The idea is as follows: imagine you have an image which you believe to be sparse in a wavelet basis, such as the Haar basis, but instead of a subset of wavelet coefficients, you observe a sample of Fourier coefficients instead (this is the norm in tomographic imaging due to the projection-slice theorem) and wish to invert it. The Gram operator which maps a function represented in wavelet coefficients to a function represented in Fourier coefficients is an infinite-dimensional unitary operator, meaning it preserves angles between vectors. By sampling only a subset of Fourier coefficients, we get a subset of the rows, but because of sparsity in the wavelet basis, we are interested in the columns. Since Haar wavelets and Fourier series are incoherent, it is not the case that most of the energy in a given set of columns is concentrated in the subset of rows sampled. As a result, if one tries to recover as many wavelet coefficients as Fourier coefficients sampled, then even if the number of Fourier coefficients sampled grows, the

matrix will become arbitrarily ill-conditioned, with minimal singular value decaying to zero as the number of samples grows. As a result, as more samples are taken, reconstruction error may (potentially) blow up to infinity. This is remarkable because a unitary operator always has condition number one and so the true inverse which is being approximated is very well conditioned. The error arises precisely because for a non-compact operator, a finite dimensional subsample cannot consistently estimate the singular values. The solution however, is quite simple: instead of sampling rows and columns evenly, sample many more rows than columns. It is known for finite-dimensional matrices that the minimal singular value of a subset of columns is at least as large as that of the matrix as a whole (this is Cauchy's interlacing theorem: see e.g. Tao (2011, Ch 1.3)), and this applies also for infinite-dimensional operators (with minimal singular value replaced by inverse of the norm of the inverse). If the columns are well approximated, the minimal singular value can be kept well above zero and the inverse is well-posed, and so one can eventually recover any finite set of wavelet coefficients.

In my application, the unitary transformation of interest is the unitary operator which rotates the axes so that the components to be solved forward and the components to be solved backwards are in known position. We wish to recover elements of the forward looking subspace from a sample of the basis functions we use to approximate the functional derivatives, which are chosen mainly for computational convenience since we rarely know the true eigenfunctions, nor can we easily compute them to high accuracy in the absence of special structure. Without the structure imposed by knowledge of the true Gram matrix in Adcock *et al.* (2014b), we will not obtain rates of convergence with or without uneven sections, but we can guarantee consistency for large enough number of rows sampled.

3.3 Consistent Approximation of the Policy Functions

3.3.1 Overview of Solutions to Rational Expectations Models in Function Space

The class of problems to which this procedure is directed is a slight generalization of the class of recursive functional rational expectations models described in Chapter 1. The object to be solved for is still the stable recursive solution of a nonlinear expectational difference equation over a Hilbert space, in the sense of Definition 1 in Chapter 1, a set of maps $g(x,\sigma) : \mathcal{H}_x \times \mathbb{R} \to \mathcal{H}_y$, $h(x,\sigma) : \mathcal{H}_x \times \mathbb{R} \to \mathcal{H}_x$ which satisfy $\mathbb{E}F(x, g(x,\sigma), h(x,\sigma) + \sigma\eta z', g(h(x,\sigma) + \sigma\eta z', \sigma), \sigma) = 0$ for a model defined by given set of equilibrium conditions F, which induces a stationary law of motion for x. The goal is still to construct the first order Taylor expansion of these solution operators around a nonstochastic steady state (Chapter 1 Definition 2), which is assumed to exist. Given this framework, under Γ-regularity of the operators (B, A), $A = \begin{bmatrix} F_{x'} & F_{y'} \end{bmatrix}, B = -\begin{bmatrix} F_x & F_y \end{bmatrix}$ composed of the functional derivatives of the equilibrium conditions F at the steady state, the functional derivatives of these operators take precisely the same form, defined in terms of the generalized Schur decomposition of these operators, $(B, A) = (Q^*TU, Q^*SU)$, $(T, S) = \left(\begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix}, \begin{bmatrix} S_{11} & S_{12} \\ 0 & S_{22} \end{bmatrix} \right)$ conformable with the decomposition $Q = \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix}$ and $U = \begin{bmatrix} U_1 \\ U_2 \end{bmatrix}$ such that the images of U_1^* and U_2^* respectively decompose $\mathcal{H}_1 = \mathcal{H}_x \times \mathcal{H}_y$ into two orthogonal subspaces $\mathcal{H}_{11}$ and $\mathcal{H}_{12}$ and the spectrum of (T_{11}, S_{11}) lies inside the unit circle, so S_{11} has bounded inverse. U_1 and U_2 are further decomposed as $U_{11} := U_1\varphi^X$, $U_{12} := U_1\varphi^Y$, $U_{21} := U_2\varphi^X$, $U_{22} := U_2\varphi^Y$ where

$\varphi^X : \mathcal{H}_x \to \mathcal{H}_x \times \{0\} \subseteq \mathcal{H}_1$ and $\varphi^Y : \mathcal{H}_y \to \{0\} \times \mathcal{H}_y \subseteq \mathcal{H}_1$ are imbeddings. Given this representation, the first order approximate solutions satisfy Chapter 1 Formulas 3.3 and 3.6 as

$$U_{21} + U_{22} g_x = 0$$

$$h_x = (U_{11} + U_{12} g_x)^{-1} S_{11}^{-1} T_{11} (U_{11} + U_{12} g_x)$$

The difficulty arises in the assumptions needed to ensure that these local solutions, defined in terms of infinite-dimensional objects, can be consistently approximated by computable, finite-dimensional methods. For a broad variety of models, there does not exist an 'asymptotically diagonal' representation of the functional derivatives, defined in terms only of compact operators and scalar multiples of the identity, and so the procedures in Chapter 1, which require this assumption even to construct the approximation are not applicable. We will instead offer an alternative procedure which does not require this structure.

Moreover, because the procedure is not based on the operator norm perturbation theorem used to demonstrate the validity of that method, the additional assumptions required to ensure that the approximation yields a continuous perturbation of the true operators in operator norm are not needed. As a result, Chapter 1 Condition 2, which assures both existence of functional derivatives of the solution and the regularity needed to use the method provided, can be weakened to just the conditions ensuring existence, viz. Γ-regularity to ensure decomposability and invertibility of U_{22} to ensure that a solution exists which is stable. In practice, even invertibility of U_{22} is slightly stronger than needed, if all one wishes is to ensure existence of a solution, as opposed to existence and uniqueness. In the presence of multiple solutions, as occurs in the case of models with indeterminacy, the procedure provided will converge to one solution, the 'minimum norm solution' so long as $U_{22} U_{22}^*$ has bounded inverse. The weakness of the conditions imposed is what permits the applicability of the described procedure

to ill-posed models: it is also what occasions the substantially more intricate analysis required to ensure validity of approximations for a much broader class of objects.

3.3.2 Models with Non-compact Functional Derivatives

Relaxing the assumption that the functional derivatives of the model of interest consist only of compact and identity operators allows consideration of models with substantially different economic structure, and potentially, equilibrium economic behavior. Non-compact operators can derive from a variety of different economic situations: the commonality in all cases is simply that the output, when faced with an input which is function-valued, cannot be summarized to arbitrarily high accuracy using only a finite set of sufficient statistics chosen before the input is known. The identity operator presents the clearest example: in order to exactly reproduce any input over an infinite-dimensional space, it must allow a completely unrestricted class of possible outputs. The same logic applies to other operators which are *unitary*, preserving the norm of any possible input. This class includes any transformation that can be seen as a change of basis, including rotations, permutations, and common transformations such as the Fourier or wavelet transform. Operators of this sort can appear in any case where evolution over time preserves certain quantities or symmetries of the system, though admittedly this is less common in economic than physical applications. One class of non-compact operator which is ubiquitous in economic applications is the *multiplication operator* $g(s) \to m(s) \cdot g(s)$ for some non-constant $m(s)$. This arises most often as the functional derivative of a composition operator $g(s) \to f(g(s))$ with respect to $g(.)$, for $m(s) = \frac{d}{dx} \mid_{g(s)} f(.)$. These arise, for example, in Euler equations for intertemporal optimization problems with respect to any state-dependent variables. Relatedly, it may also occur that one wants to differentiate $f(g(s))$ with respect to the function $f(.)$, rather than $g(.)$. In this case, the functional derivative is given by the *pointwise evaluation operator* $\text{eval}_{g(s)}[f(.)] := f(g(s))$. In general, this

operator is not compact, and, depending on the Hilbert space in which $f(.)$ lives, may also fail to be bounded (see Berlinet & Thomas-Agnan (2004)). These can occur in any case where the state variable is a function giving the response to another function, as is ubiquitous in game theoretic settings or in dynamic settings.

One intuitive way to look at an evaluation operator is to see it as a degenerate form of the broader class of kernel integral operators $f(s) \rightarrow \int k(s,t)f(t)dt$ where the kernel $k(s,t)$ is, for every s given by the generalized function, or Schwarz distribution, $\delta_{g(s)}$, rather than a function. While the case in which the kernel is not even a well-defined function is extreme, non-compactness can occur in many cases where it is a function, but is not smooth or bounded. Typical examples include *singular integral operators*, in which $k(t,s)$ asymptotes to infinity at some point in t, see Beylkin *et al.* (1991) for discussion. These may arise in the case of a rapidly-decaying potential function, for example, in spatial models where trade costs (or other costs of interactions) rise rapidly with distance. More generally, integral operators arise in any case defined by a conditional expectation, and may fail to be compact in many cases where the conditional density in this expectation is non-smooth. This can occur in models where some variables evolve deterministically or contain components which jump or flow in a non-smooth manner, as frequently arises in models with fixed costs (Stokey, 2008).

It should be noted that the presence of these operators in an economic model need not indicate that the behavior of the model is necessarily ill-posed. This requires also that the non-compactness cannot be eliminated by some transformation which leaves the behavior of the solution unchanged. The simplest way to perform such a transformation is to multiply the row of the block operators (B, A) by the inverse of the non-compact operator in question. This transforms the non-compact operator into an identity. For a unitary transformation, this is simply given by inverting the change of basis. Because the composition of any bounded operator with a compact operator is also compact, so long as this inverse is bounded and the operator being

inverted is the only non-compact operator in the row, this transforms a row containing a non-compact operator into one without, without changing the solution. However, if there is more than one non-compact operator in the row, this procedure fails. It is precisely this class of equations in which ill-posedness is a fundamental issue. For example, the equation $f'(s) = \int k(s,t)f(t)dt$ for a singular kernel $k(s,t)$, which could represent the transition equation for a population of agents evolving according to a law of motion inducing a non-smooth conditional density, has functional derivatives $(B, A) = (\int k(s,t)[.]dt,\ I)$, both of which are non-compact.

Equilibrium relationships of this non-compact type typically reflect a fundamental non-smoothness or high-dimensionality in the dynamics of the economy of interest. The previous evolution equation example provides useful intuition. The set of behaviors inducing the singular conditional density may be relatively simple, but they can induce dynamics at the aggregate level which are highly nontrivial. Extremely fine features of the distribution at time t may have nontrivial impact on the distribution at $t+1$ (and any aggregate variables which depend on it), making description of the dynamics using low-dimensional statistics difficult or impossible. Some care is warranted here; in the case where the space of shocks facing the economy is essentially low-dimensional, in the sense of having a compact covariance operator, it may be highly feasible to construct an estimator which is consistent with respect to, for example, mean squared Hilbert norm error of forecasts, even if no estimator is consistent uniformly over all possible inputs. However, when considering potential policy options, which need not be aligned with the historical distribution of shocks, the effect of these higher order components cannot be ignored.

One further way to avoid non-compactness is to assume a model in which the non-smoothness does not appear. Such a model may in many ways be very similar to one in which it does, if the smoothing is induced, for example, by inducing a small amount of smooth noise or heterogeneity into the deterministic or non-smooth

transition, resulting in a transition equation which for any individual is, with high probability, close to the deterministic value, but which induces a transition equation which is smooth, if rapidly changing at some points. For example, in the fixed cost case, one may switch a deterministic fixed cost for a stochastic one, as used in generalized SS models (Caballero & Engel, 1999), or otherwise introduce risk into deterministic actions. While such changes are in some cases economically plausible, they are not innocuous with respect to predictions for aggregate variables. For an infinitely large class of potential inputs, the operators corresponding to the smoothed and unsmoothed kernels differ by an amount bounded away from 0. To the extent that components of this infinitely large set of variables interact with other variables, the implied predictions for aggregate variables may differ substantially, no matter how small the variance of the noise term added. As such, the choice to use a well-posed rather than ill-posed model should not be made merely based on convenience, but also based on understanding of the effects on both individual and aggregate level dynamics.

3.3.3 Discussion of Methods

In order to calculate the first order derivatives of the policy function which solves the rational expectations equilibrium, it is necessary to evaluate and compose certain functions of the Fréchet derivatives of the equilibrium conditions, in particular their generalized Schur decompositions and inverses thereof. Except in certain exceptional cases, often requiring both partial equilibrium structure and specialized form for the equilibrium conditions such as isotropy with respect to the Haar measure of a known finite-dimensional group, finding a closed form for these functions is intractable and often impossible. Numerical approximation suggests itself as an alternative approach, but standard techniques encounter a number of pitfalls. In order for a technique based on 'plugging in' approximations to the primitive objects of the model to yield consis-

tent approximation of the model outcomes, two conditions are required: first, these approximations must be consistent in some topology, and second, the model outcomes must be continuous with respect to this same topology. This raises immediate problems for any approach based on finite-dimensional approximations of the equilibrium operators when compactness does not hold, as the fact that these operators are not compact implies that no finite dimensional approximation of any sort can consistently approximate them in operator norm.[2] As a result, to use finite-dimensional methods in this non-compact case, it is necessary to use a weaker sense of consistency for the convergence of the primitives, such as convergence with respect to the strong operator topology of pointwise convergence in Hilbert space instead of the operator norm topology which requires uniformity over the Hilbert space domain. However, when the primitives no longer converge in the stronger sense, it is no longer necessary that the function of the primitives remains continuous. Indeed, I will show that with respect to the strong operator topology, the policy function is no longer a continuous function of the primitives without certain auxiliary conditions, and so a direct plug-in approximation cannot be shown to be consistent, even in this considerably weaker sense.

Instead, I propose a slightly modified procedure, which overcomes the discontinuities in three ways.

The first is to impose sufficient conditions on the approximation such that the

[2]Those familiar with numerical analysis may object that there do exist frequently used numerical methods which are able to handle certain classes of non-compact operators. Many classes of differential equations characterized by non-compact differential operators are solved routinely by discretization and finite-element methods. In general, however, such methods find weak solutions characterized by functions of the operator of interest which are compact, such as an inverse or resolvent. In the models specified here, where neither the operators nor their inverses are compact, these results are not directly applicable. More relevant, certain models may be characterized by parametric perturbation theory, in which one works directly with infinite-dimensional objects which are more tractable than the operators of interest and differ only by a finite dimensional parameter. While such an approach may be useful in certain restricted classes of model with function-valued state space, and indeed has been used for heterogeneous agent models (e.g. Krieger (2001)), defining an appropriate parametric deviation must rely on the structure of the model, and may be impractical for deviations which are infinite-dimensional in nature, such as arbitrary shocks to cross-sectional distributions.

components of the policy function which require taking an inverse, which is not continuous with respect to the strong operator topology, converge. This is guaranteed by what Chatelin (2011, Ch. 3) refers to as stable convergence, sufficient conditions for which include strong operator topology convergence and a uniform bound on minimum singular values. For sieve-type approximations, I show that such a bound is guaranteed by an infinite-dimensional variant of Cauchy's interlacing inequalities. Use of this result is facilitated by the characterization of (generalized) Riesz projections in terms of a complex path integral over a (generalized) resolvent operator, which reduces the issue of convergence of spectral projections to application of dominated convergence to the more easily characterized resolvent operator.

The second technique is to leverage pointwise convergence into uniform convergence over finite sets of basis functions, which can yield uniform (operator norm) convergence for finite-dimensional operators constructed from the outputs, allowing characterization of the spectral properties of finite-dimensional intermediate representations and also ensuring that the policy function, which is compact under reasonable conditions, converges in operator norm even though the primitives characterizing it do not. The cost of the additional power afforded by this technique is the necessity each time it is used to have substantially fewer elements in the approximation of the subsequent step, resulting in both slower rates of convergence and the imposition of the number of elements as an additional tuning parameter in the application of the method. This represents a particular difficulty as the strong operator topology is not metrizable and so no guidance may be provided on optimal rates for these tuning parameters.[3]

[3]It is possible, under fairly mild conditions, to show convergence of the primitives in slightly stronger senses which are metrizable, and indeed induce a Banach space, but are still substantially weaker than operator norm convergence. It is not yet clear if these metrics may be passed through all steps of the proof in order to yield rates of convergence in terms of the number of basis of functions for each step. Investigation of this avenue, in order to provide guidance on tuning parameter selection, remains a priority. The simplest examples of such a space are the weighted L^p spaces of $p-$(Bochner) integrable measurable linear operators between fixed Hilbert spaces, when the measure on the space is tight and so sufficiently concentrated on a small set, as are traditional stochastic process priors used

Third, certain components of the approximation for which (procedures analogous to psuedo-)inverses are required cannot be guaranteed to be nonsingular and so may not converge even under a stronger norm. For these components, consistency can be achieved by spectral cutoff-type regularization of the finite-dimensional approximations. Note that this is a separate issue from the first concern, as unlike in traditional ill-posed inverse problems, cutoff or Tikhonov-type regularization will not, in general, eliminate the continuity problem for inverses of operators converging in strong operator topology to non-compact operators. The reason for this is that even if the inverse is bounded and so the problem is 'well-posed' according to the traditional Tikhonov classification, if the operator and its inverse are not compact the singular functions cannot be consistently estimated and so one cannot confine the bias induced by the regularization to any particular eigenspace.

3.3.4 Step-by-Step Construction

The approximation of the equilibrium policy function proceeds in a sequence of steps, each component constructed from the previous component. Beginning with an approximation of the Fréchet derivatives of the operators, one constructs the Riesz projector, then the unitary operators inducing the Schur decomposition, then the triangular form, and finally assembles the components into approximations of the policy operators g_x and h_x.

We begin with some preliminary lemmas.

For purposes of numerical approximation, we would like to write h_x and g_x in forms which are amenable to numerical approximation. In particular, in order to ensure stability of approximations of inverses, it is helpful to use constructions which ensure that any inverse which is approximated is of a self-adjoint positive definite operator. For-

in Bayesian nonparametrics such as Gaussian processes and Bayesian sieves. Aside from requiring an additional uniform integrability assumption, by putting almost all weight on compact and so "nearly" finite dimensional spaces, these appear to produce behavior very similar to pointwise convergence.

tunately, in the context of constructing the solution to a linear rational expectations model, representations in this form are easily constructed. For g_x, this is achievable by using the analogue of the minimum norm solution $g_x = -U_{22}^*(U_{22}U_{22}^*)^{-1}U_{21}$ of the linear equation $U_{21} + U_{22}g_x = 0$: this will permit the use of approximations of U_{22} and U_{21} for which the analog equation need not have a unique solution, even if the true equation does. A related transformation can be shown to yield a representation of h_x for which the analog constructed with approximated components yields a stable inverse.

Lemma 3.1. *Let $(U_{22}U_{22}^*)^{-1}$ be bounded and let $g_x = -U_{22}^*(U_{22}U_{22}^*)^{-1}U_{21}$, the minimum norm solution. Then*

$$h_x = (U_{11}+U_{12}g_x)^{-1}S_{11}^{-1}T_{11}(U_{11}+U_{12}g_x) = (\varphi^{X*}\varphi^X+g_x^*g_x)^{-1}(g_x^*U_{12}^*+U_{11}^*)S_{11}^{-1}T_{11}(U_{11}+U_{12}g_x) \tag{3.1}$$

Proof. See Appendix C. □

As a corollary, note that $\varphi^{X*}\varphi^X + g_x^*g_x$ is a quadratic form satisfying

$$\inf_{\|x\|_{\mathcal{H}_x}=1}\|(\varphi^{X*}\varphi^X + g_x^*g_x)x\| \geq \inf_{\|x\|_{\mathcal{H}_x}=1}\|\varphi^{X*}\varphi^X x\| = 1$$

and so this inverse always exists and is bounded. Therefore, if $(U_{22}U_{22}^*)^{-1}$ is bounded, a stable solution exists for h_x.

The above representations of the objects to be approximated ensure the ability to pass strong operator topology convergence through inverses, by ensuring uniform upper bounds on the inverse of the approximating sequence, or, equivalently in the finite-dimensional case, uniform lower bounds on the minimum singular value. For approximations which are constructed by projecting the operator of interest onto a lower dimensional space to form a submatrix, a useful tool is the Cauchy interlacing inequalities (see, e.g. Tao (2011, Ch 1.3)), which provide bounds on the eigenvalues

of a compression, or symmetric submatrix, of a Hermitian matrix in terms of the eigenvalues of the original matrix, and in particular ensure that the minimum eigenvalue, considered on the lower-dimensional space, is no smaller than the minimum eigenvalue of the original matrix. This then implies that for non-Hermitian matrices, a submatrix consisting of a subset of the rows has smallest singular value no smaller than that of the original matrix. Since the objects to be approximated are not in general finite-dimensional or even compact, the standard version of the theorem does not apply. However, in this situation, a bound on the norm of the inverse can be recovered. The following result is likely classical, but lacking a reference and because the proof is simple, is included for completeness.

Lemma 3.2. *Let $A \in \mathcal{L}(\mathcal{H} \to \mathcal{H})$ be a bounded positive self-adjoint operator with bounded inverse. Let P be an orthogonal projection from $\mathcal{H}$ to some finite-dimensional subspace $Im(P)$. Then, considering PAP as an operator on $Im(P)$, $\lambda_{\min}(PAP) \geq \frac{1}{\|A^{-1}\|}$*

Remark. The previous lemma may be extended to non-self-adjoint operators S by considering the necessarily positive and self-adjoint operator SS^*, as $\sigma_{\min}(PS) = \lambda_{\min}^{\frac{1}{2}}(PSS^*P) \geq \frac{1}{\|(SS^*)^{-1}\|^{\frac{1}{2}}}$. Note that this provides bounds on one-sided projections only: for projections of domain and range for non-self-adjoint operators, one obtains only $\sigma_{\min}(PSP) \geq \frac{1}{\|(SPS^*)^{-1}\|^{\frac{1}{2}}}$, where the latter quantity may not be well controlled even if $\|(SS^*)^{-1}\|$ is.

To construct approximations of the generalized Schur decomposition of the Fréchet derivatives of the equilibrium conditions, rather than simply finding an approximation of the derivative operators and taking its generalized Schur decomposition, instead we follow the construction in the existence proof for this decomposition by first approximating the Riesz projector onto the stable generalized eigenspace and then passing through it a set of vectors which are then orthonormalized to construct a unitary operator representing synthesis with respect to an orthonormal basis of the span of

the Riesz projection. In contrast to the full calculation of the generalized Schur decomposition, this is different in two ways.

First, we do not presume that the vectors passed through the Riesz projector are restricted to the eigenspaces associated to isolated eigenvalues of the true operator pencil, and as a result, within the stable subspace, the operator generated by this approximations will not in general be upper triangular with isolated eigenvalues along the diagonal. As the construction of the equilibrium policy function only requires a block upper-triangular decomposition with respect to the stable and unstable subspaces, this does not impede construction of the policy function. Second, for the finite-dimensional approximation of the unitary operator, the number of vectors used to construct an orthonormal basis need not and, in general should not, equal the dimension of the space spanned by the approximated Riesz projector, as the dimension of the approximation represents a tuning parameter controlling relative approximation error rather than any fundamental aspect of the system itself.

Such an approach involves some loss of information relative to methods used for finite-dimensional problems. By eliminating explicit calculation of eigenvectors and eigenvalues it does preclude the ability to assess the existence and uniqueness of a stable solution by counting stable versus unstable eigenvalues as first proposed by Blanchard & Kahn (1980). Existence of a stable solution must instead be presumed or verified at the outset. This should not be considered particularly surprising: while the non-existence of a linear isomorphism between two spaces may be assessed in finite dimensions simply by comparing the dimension of the spaces, when both are infinite-dimensional more sophisticated assessments are needed. This is particularly the case when the spaces are being represented by finite-dimensional approximations, for which a mismatch in dimensionality may be a result of the approximation rather than of the spaces themselves. Construction of approximations from subsets of the space must therefore be undertaken with care, as the method may not provide the useful service

of finite dimensional rational expectations algorithms of failing to produce output when the assumptions for existence of a solution do not hold. As the Blanchard & Kahn (1980) criterion is analogous to testing the order condition as a necessary condition for a rank condition, and the conditions analogous to the rank condition are supplied by the completeness of U_{22} and its adjoint, this difficulty should be seen in the light of the difficulty of testing for completeness: Canay *et al.* (2013) provide impossibility results in a distinct but related setting.

To define the forward and backward looking subspaces properly, we make the following assumption

Assumption 1: Let the pair of Fréchet derivatives of the equilibrium conditions (B, A) be a pair of bounded operators between separable Hilbert spaces $B \in \mathcal{L}(\mathcal{H}_1 \to \mathcal{H}_2)$ $A \in \mathcal{L}(\mathcal{H}_1 \to \mathcal{H}_2)$ Γ-regular with respect to Cauchy curve $\Gamma = \{\exp(\iota\lambda) : \lambda \in [0, 2\pi)\}$.

The above condition specializes the assumptions of Lemma .9 for constructing a block-wise Schur decomposition to the setting of the derivatives of the equilibrium conditions and ensures that the blocks separate the stable and unstable subspaces. To see that using the positively-oriented unit circle as the separating Cauchy curve satisfies the assumptions of Lemma .9, note that the inner domain Δ_+ is the interior of the unit circle, so $0 \in \Delta_+$ and outer domain Δ_- is the exterior, so $\infty \in \Delta_-$. The 'forward looking' or 'exterior' subspace is then given by the Schur subspace corresponding to the exterior of Γ, while the 'backward looking' or 'interior' subspace is the subspace corresponding to the interior of Γ.

Formally, the conditions for strong operator topology convergence of the approximation of the Riesz projector onto the stable eigenspace are as follows.

Lemma 3.3. *Let (B, A) satisfy Assumption 1. Let π^{k_1} be an orthonormal projection onto a k_1 dimensional set of basis vectors spanning a subspace monotonically increasing to $\mathcal{H}_2$ and π^{k_2} be an orthonormal projection onto a k_2 dimensional set of basis*

vectors similarly increasing to $\mathcal{H}_1$*, and define* $(\tilde{B}, \tilde{A}) = (\pi^{k_1} B \pi^{k_2}, \pi^{k_1} A \pi^{k_2})$*. Define the Riesz projector onto the stable subspace* $P_1 := \frac{1}{2\pi\iota} \int_\Gamma (\zeta A - B)^{-1} A d\zeta$ *and its approximation* $\tilde{P}_1^{k_1,k_2} := \frac{1}{2\pi\iota} \int_\Gamma ((\zeta \tilde{A}^* - \tilde{B}^*)(\zeta \tilde{A} - \tilde{B}))^{-1} (\zeta \tilde{A}^* - \tilde{B}^*) \tilde{A} d\zeta$*. Then,* $\forall x \in \mathcal{H}_1$ *there exists a sequence* $k_1(k_2)$ *increasing in* k_2 *such that* $\forall \epsilon > 0$*,* $\exists K_2$ *such that for all* $k_2 \geq K_2$*,* $\|(P_1 - \tilde{P}_1^{k_1(k_2),k_2})x\| < \epsilon$*.*

In order to construct the unitary operators $U = [U_1, U_2]$ and $Q = [Q_1, Q_2]$ which induce a block upper triangular decomposition of the pair (B, A) using only the pointwise convergence of the Riesz projector P_1 which maps onto the space spanned by U_1, the space spanned by U_1 and that spanned by U_2 will be treated asymmetrically. The reason is that U_1 may be constructed from points in the span of P_1 directly, but because the space on which U_2 resides is orthogonal to the entire space spanned by U_1, pointwise convergence of P_1 only gives access to constructions which are orthogonal to a finite dimensional subspace. The issue may be resolved by first constructing a projector onto the orthogonal complement of a (growing) finite dimensional subspace, which will converge pointwise, and only then passing through this a smaller number of points to construct an approximation of U_2.

For both U_1 and U_2 (and the components of Q), which in the proof of existence of the upper triangular decomposition are constructed by orthonormalizing bases of their respective subspaces, construction proceeds by first taking the map of a fixed set of points in $\mathcal{H}_1$ onto the subspace and then using the outputs to construct a set of orthonormal vectors within the space. While in principle Gram-Schmidt orthonormalization, or more numerically robust procedures such as QR decomposition which also orthonormalize a given set of vectors could be used, these procedures are discontinuous and so not robust to approximation error in the presence of singularity induced when the projection onto the subspace renders outputs collinear. Although it might be possible to simply assume that the vectors used to construct U_1 and U_2 induce for any finite dimension a full-rank (quasi-)matrix (in the sense of Townsend

& Trefethen (2014)), primitive conditions for this would require at least some knowledge of the characteristics of the stable and unstable subspaces which is generally not available, as otherwise there would be no need to use numerical approximation. Instead, the issue may be handled by judicious regularization which eliminates eliminates components which appear to be close to singular. In particular, we use a truncated singular value decomposition, which uses as the set of orthonormal vectors in the stable subspace the left singular vectors of the projection onto that subspace of a set of vectors in $\mathcal{H}_1$ corresponding to singular values larger than a small threshold, eliminating those corresponding to singular values numerically close to 0.[4]For Q_1, which is constructed from an orthonormal basis of the range space of A and B on the stable subspace, a similar technique may employed by passing the orthonormalized vectors approximating U_1 through A and orthonormalizing the outputs, though in this case orthonormality of U_1 and the assumed nonsingularity of A on the stable subspace ensure that no singularity problem arises and so the QR decomposition may be used instead of SVD.

This construction, for any finite dimensional approximation, induces a slightly different version of the generalized Schur decomposition at each level of approximation, analogous to different ways of sorting the eigenvalues within the stable and unstable blocks in the finite-dimensional QZ algorithm as used in Klein (2000). It may also be compared to the difference between the real and the complex QZ algorithms, as, similarly to the real QZ algorithm, this method preserves the block triangularity structure, but not necessarily triangular structure within a block. As the set of singu-

[4]Another way to construct an orthonormal basis from a possibly singular set of vectors, while also preserving a fixed ordering of the vectors, is the rank revealing QR decomposition, (Golub & van Loan, 1996, Ch 5.4). While this decomposition in principle is also robust to singularity (Hong & Pan, 1992), actual construction of a column pivoting which is guaranteed to achieve these guarantees appears not to be achievable with any known polynomial time algorithm (Çivril, 2014). One disadvantage of the ordering generated by the singular value decomposition is that, while it generates a monotonically increasing span, this span does not have a natural basis in terms of the generating vectors which is invariant as the dimension grows. It is this feature which necessitates the baroque construction in which Schur components are defined on sequences of distinct but isometrically isomorphic Hilbert spaces.

lar vectors span a monotonically increasing sequence of subspaces as k_3 grows, with only the partition of those subspaces changing, the approximate Schur operators induce a triangular array of sequences converging to equivalent (formally, isometrically isomorphic) Schur operators which each generate identical equilibrium solutions, this does not impede convergence to a solution.

Formally, to approximate U_1 we take a fixed set of (column) basis vectors $\{z_i\}_{i=1}^{k_3}$ in $\mathcal{H}_1$ and denote by Z^{k_3} the (quasi-)matrix with z_i in the i^{th} column. We may then take $U_1^{k_3,\varepsilon_{k_3}}$ defined by $[U_1^{*k_3,\varepsilon_{k_3}} U_{extra}^{*k_3,\varepsilon_{k_3}}] D^{k_3} V^{k_3} = svd(\tilde{P}_1^{k_1(k_2),k_2} Z^{k_3})$ as an approximation of U_1, where $U_1^{*k_3,\varepsilon_{k_3}}$ is the set of left singular vectors of $\tilde{P}_1^{k_1(k_2),k_2} Z^{k_3}$ corresponding to singular values greater than a threshold $\varepsilon_{k_3} > 0$, which should be set to a value below the smallest non-zero singular value of $\tilde{P}_1^{k_1(k_2),k_2} Z^{k_3}$, and $U_{extra}^{*k_3,\varepsilon_{k_3}}$ is the set of left singular vectors corresponding to smaller singular values. Specifically, we may say this generates convergence in the following sense.

Lemma 3.4. *Let the conditions of 3.3 hold and let Z^{k_3} be a sequence of nonsingular quasimatrices mapping $\mathbb{R}^{k_3} \to \mathcal{H}_1$ constructed by selecting as columns the first k_3 elements of a fixed sequence $\{z_i\}_{i=1}^{\infty}$ whose closed linear span converges to $\mathcal{H}_1$ (i.e. $Z^{k_3} = \sum_{i=1}^{k_3} \langle e_i, . \rangle z_i$ where e_i is the standard basis of $\mathbb{R}^{k_3}$). Then, there exist increasing sequences $k_2(k_3)$ and $k_1(k_2(k_3))$, decreasing sequence $\varepsilon_{k_3} > 0$ such that the sequence of approximations $U_1^{k_3,\varepsilon_{k_3}}$ defined by $[U_1^{*k_3,\varepsilon_{k_3}} U_{extra}^{*k_3,\varepsilon_{k_3}}] D^{k_3} V^{k_3} = svd(\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)} Z^{k_3})$, satisfies, $\forall x \in \mathcal{H}_1$, for any $\epsilon > 0$, that there exists K_3 such that $\forall k_3 > K_3$ $\|(U_1^{k_3,\varepsilon_{k_3}} - U_1^{(k_3)})x\| < \epsilon$, where $\forall k_3$, $U_1^{(k_3)}$ is an analysis operator generating a representation in ℓ_2 of a complete orthonormal basis of $Im(P_1)$.*

Remark. The particular unitary operator $U_1^{(k_3)}$ which is approximated by each element of the sequence is not constant, in the sense that it maps the same element of $\mathcal{H}_1$ to a different representation. This is, however, immaterial to solution of the equilibrium, as the input to the policy operator is mapped to a space by U_1 and then mapped back to $\mathcal{H}_1$ by its inverse after a suitable transformation. As the range space of $U_1^{(k_3)}$ is

isometrically isomorphic for all k_3 and the transformation on that range space changes precisely so as to offset this difference (see below for elaboration), this lack of a single limit does not present a challenge. The important property of the limit which is to be shown in this lemma is the property that $U_1^{(k_3)*}U_1^{(k_3)}$ is an orthogonal projection onto $Im(P_1)$, and so separates the stable subspace from the unstable subspace.

Construction of an estimator of U_2 which is pointwise consistent proceeds in a substantially identical fashion to construction of the estimator for U_1. The main difference is that the subspace onto which the test functions are projected is not $Im(P_1)$ but its orthogonal complement. As a result, the estimator of P_1 may not be used directly. Moreover, because $\tilde{P}_1^{k_1(k_2),k_2}$ is only pointwise consistent, the projector onto the orthogonal complement of the span of $\tilde{P}_1^{k_1(k_2),k_2}$,

$$I - \tilde{P}_1^{k_1(k_2),k_2}(\tilde{P}_1^{k_1(k_2),k_2*}\tilde{P}_1^{k_1(k_2),k_2})^+\tilde{P}_1^{k_1(k_2),k_2*},$$

is also not consistent, even if the pseudo-inverse is regularized by cutting off numerically small singular values. The reason for this is simple: as P_1 is a projection onto a subspace, it does not have bounded inverse on $\mathcal{H}_1$ as a whole, and so no 'minimal singular value' condition of the type used in 3.2 applies, and so its inverse and as a result, its resolvent and so its singular vectors, are not guaranteed to converge, and so regularization cannot be applied to distinguish zero and nonzero singular values. The issue is resolved by instead passing to a sequence of finite dimensional subspaces on which convergence can be guaranteed. In fact, this has been achieved by the previous lemma, which guarantees convergence of an operator spanning a finite-dimensional subspace of $Im(P_1)$, in the form of $U_1^{k_3,\varepsilon_{k_3}}$. Projection onto the orthogonal complement of this smaller subspace can then be calculated by $I - U_1^{k_3,\varepsilon_{k_3}*}U_1^{k_3,\varepsilon_{k_3}}$, which is pointwise consistent. From there, the proof is identical.

Lemma 3.5. *Let the conditions of 3.3 hold and let Z'^{k_4} be a sequence of nonsingular*

quasimatrices mapping $\mathbb{R}^{k_4} \rightarrow \mathcal{H}_1$ *constructed by selecting as columns the first* k_4 *elements of a fixed sequence* $\{z'_i\}_{i=1}^{\infty}$ *whose closed linear span converges to* $\mathcal{H}_1$ *(i.e.* $Z'^{k_4} = \sum_{i=1}^{k_4} \langle e_i, . \rangle z'_i$ *where* e_i *is the standard basis of* $\mathbb{R}^{k_4}$*). Then, there exist increasing sequences* $k_3(k_4)$, $k_2(k_3(k_4))$ *and* $k_1(k_2(k_3(k_4)))$*, decreasing sequence* $\varepsilon_{k_4} > 0$ *such that the sequence of approximations* $U_2^{k_4,\varepsilon_{k_4}}$ *defined by* $[U_2^{*k_4,\varepsilon_{k_4}} U_{extra}^{*k_4,\varepsilon_{k_4}}] D^{k_4} V^{k_4} = svd((I - U_1^{k_3,\varepsilon_{k_3}*} U_1^{k_3,\varepsilon_{k_3}}) Z'^{k_4})$*, satisfies,* $\forall x \in \mathcal{H}_1$*, for any* $\epsilon > 0$*, that there exists* K_4 *such that* $\forall k_4 > K_4 \ \|(U_2^{k_4,\varepsilon_{k_4}} - U_2^{(k_4)})x\| < \epsilon$*, where* $\forall k_4$*,* $U_2^{(k_4)}$ *is an analysis operator generating a representation in* ℓ_2 *of a complete orthonormal basis of* $\mathcal{H}_1 \backslash Im(P_1)$*.*

Remark. Although there need not be any relation between $\{z'_i\}_{i=1}^{\infty}$ and $\{z_i\}_{i=1}^{\infty}$, there is nothing prohibiting using the same sequence, and doing so may improve the performance of the approximation as the component of z'_i lying inside $Im(P_1)$ may be approximated more accurately if z'_i was itself considered in the approximation of $Im(P_1)$. For the same reason, a reasonable choice of Z^{k_3} might include elements lying inside the span of π^{k_2}. For example, if an ordered orthonormal basis such as orthogonal or trigonometric polynomials or wavelets is used for approximation, π^{k_2} may project onto the first k_2 elements, Z^{k_3} may include the first k_3 elements, and Z'^{k_4} may include the first k_4.

Construction of an estimate of the left Schur operator Q follows the same general pattern, with a few differences. First, as only the components of the generalized Schur decomposition corresponding to the stable subspace are needed for construction of a stable equilibrium solution, only Q_1, which acts on the stable subspace, needs to be constructed, though one could produce an estimate of Q_2 in a similar way if desired. Second, the space on which Q_1 is constructed is the image of the projector referred to as π_2 in the proof of existence of the Schur decomposition. Rather than constructing an additional estimator of π_2, points in this subspace are constructed using an estimate of an operator with equivalent range, AP_1, using as inputs the orthonormal vectors constructed to span $\mathrm{Im}(P_1)$. This allows construction of a basis

to proceed by the QR algorithm instead of the singular value decomposition.

Lemma 3.6. *Let the conditions of 3.3 hold. Define an approximation $Q_1^{k_3}$ of Q_1 by $Q_1^{k_3 *} R^{k_3} = qr(\tilde{A}\tilde{P}_1^{k_1(k_2),k_2} U_1^{k_3,\varepsilon_{k_3}*})$. Then, $\forall x \in \mathcal{H}_2$, there exist increasing sequences $k_2(k_3)$ and $k_1(k_2(k_3))$ such that for all $\epsilon > 0$, $\exists K_3$ such that $\forall k_3 > K_3$, $\|(Q_1^{k_3} - Q_1^{(k_3)})x\| < \epsilon$ where $\forall k_3$, $Q_1^{(k_3)}$ is an analysis operator generating a representation in ℓ_2 of a complete orthonormal basis of $Im(\pi_2)$. $Q_1^{k_3 *}$ also converges in strong operator topology to $Q_1^{(k_3)*}$.*

Given uniformly bounded estimates of Q_1, U_1, A, and B, one may easily construct pointwise consistent approximations of the Schur matrices corresponding to the $(1,1)$ block of the block upper triangular decomposition of (B, A), denoted $(T_{11}, S_{11}) = (Q_1 B U_1^*, Q_1 A U_1^*)$ by plugging in the approximated analogues. However, aside from consistency, the analogue estimators do not generate an approximation of the policy function with desirable properties. First, because S_{11} enters through an inverse in $S_{11}^{-1} T_{11}$, so it is also necessary to achieve stable convergence, which may be achieved by 3.2 provided that the approximation of S_{11} used approximates a subset of columns and not just a submatrix. As with the Riesz projector, this is achieved by asymmetry in the size of the row and column dimensions, here using a lower dimensional approximation of U_1 than of Q_1. One minor complication in this procedure is that the spaces on which the elements of the sequences $U_1^{(k_3)}$ and $Q_1^{(k_3)}$ live are not constant, which is problematic as one needs a fixed vector to which to apply a pointwise convergent operator. To stabilize the spaces, instead of a subset of an arbitrary Hilbert space represented by ℓ_2, we construct representations of $U_1^{(k_3)}$ and $Q_1^{(k_3)}$ as projections in $\mathcal{H}_1$ so that S_{11} and T_{11} may be viewed as operators on $\mathcal{H}_1$ with a single fixed limit rather than a sequence of limits on different but isomorphic Hilbert spaces. More interestingly, under the stronger compactness conditions imposed for .10, T_{11} and $S_{11}^{-1} T_{11}$ can be shown to be compact, opening the possibility that stronger results than pointwise convergence might be achieved. By applying pointwise convergence on an increasing

set of singular vectors, I show that under this additional condition, one may indeed construct an estimator which is consistent in the operator norm topology. This ability to 'leverage' a weaker result into a stronger one is key to ensuring that the approximation of the policy function constructed converges uniformly over inputs. As the compactness condition may not hold in all cases, I define two estimators of h_X: one which achieves uniform convergence but requires the compactness condition, and another which achieves only strong operator topology convergence, but does not require the compactness condition to hold.

Lemma 3.7. *(i) Let the conditions of 3.3 hold, and define*

$$(T_{11}^{(k_3)}, S_{11}^{(k_5)}) = (Q_1^{(k_3)} B U_1^{(k_3)*}, Q_1^{(k_3)} A U_1^{(k_5)*}).$$

Define $T_{11}^{k_3} = Q_1^{k_3} \tilde{B} U_1^{k_3*}$ *and* $S_{11}^{k_5} = Q_1^{k_3} \tilde{A} U_1^{k_5,\varepsilon_{k_5}*}$ *where* $U_1^{k_5,\varepsilon_{k_5}*}$ *is constructed as in 3.4 using the first* k_5 *elements of* $\{z_i\}_{i=1}^{\infty}$. *Then* $\forall x \in \mathcal{H}_1$ *there exist sequences* k_5, $k_3(k_5)$, $k_2(k_3(k_5))$, $k_1(k_2(k_3(k_5)))$, ε_{k_3}, ε_{k_5} *such that for any* $\epsilon > 0$ $\exists K_5$ *such that for all* $k_5 > K_5$

$$\|(U_1^{k_5*}(S_{11}^{k_5*} S_{11}^{k_5})^{-1} S_{11}^{k_5*} Q_1^{k_3} (Q_1^{k_3*} T_{11}^{k_3} U_1^{k_3}) - U_1^{(k_5)*} S_{11}^{(k_5)-1} T_{11}^{(k_3)} U_1^{(k_3)}) x\| < \epsilon.$$

(ii) Assume in addition to the previous conditions that Ω_1 *and* Ω_2 *are compact operators. Then* $T_{11}^{(k_3)}$ *and* $S_{11}^{(k_5)-1} T_{11}^{(k_3)}$ *are compact operators. Define* $thresh_{k_6}(T_{11}^{k_3})$ *as the optimal rank* k_6 *approximation of* $T_{11}^{k_3}$ *constructed by thresholding all but the largest* k_6 *singular values. Precisely, for any operator* M, $thresh_{k_6}(M) = U^{k_6} D^{k_6} V^{k_6*}$ *where* $[U^{k_6} U^{extra}] \begin{bmatrix} D^{k_6} & 0 \\ 0 & D^{extra} \end{bmatrix} \begin{bmatrix} V^{k_6*} \\ V^{extra*} \end{bmatrix} = svd(M)$ *is the division of the SVD of* M *into the component corresponding to the first* k_6 *singular vectors, ordered from largest to smallest, and the remainder. Then for fixed* k_6 *such that* $\sigma_{k_6}(Q_1^{(k_3)*} T_{11}^{(k_3)} U_1^{(k_3)}) > \sigma_{k_6+1}(Q_1^{(k_3)*} T_{11}^{(k_3)} U_1^{(k_3)})$ *there exist sequences* $k_3(k_5)$, $k_2(k_3(k_5))$, $k_1(k_2(k_3(k_5)))$, ε_{k_3}, ε_{k_5}

such that

$$\|U_1^{k_5*}(S_{11}^{k_5*}S_{11}^{k_5})^{-1}S_{11}^{k_5*}Q_1^{k_3}thresh_{k_6}(Q_1^{k_3*}T_{11}^{k_3}U_1^{k_3})-U_1^{(k_5)*}S_{11}^{(k_5)-1}Q_1^{(k_3)}thresh_{k_6}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)})\| \to 0$$

as $k_5 \to \infty$. *Similarly there exist sequences* k_6, $k_5(k_6)$, $k_3(k_5(k_6))$, $k_2(k_3(k_5(k_6)))$, $k_1(k_2(k_3(k_5(k_6))))$, ε_{k_3}, ε_{k_5} *such that*

$$\|U_1^{(k_5)*}S_{11}^{(k_5)-1}Q_1^{(k_3)}thresh_{k_6}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}) - U_1^{(k_5)*}S_{11}^{(k_5)-1}T_{11}^{(k_3)}U_1^{(k_3)}\| \to 0$$

as $k_6 \to \infty$.

Remark. $U_1^{(k_5)*}S_{11}^{(k_5)-1}T_{11}^{(k_3)}U_1^{(k_3)}$ is invariant to the set of basis vectors used to construct $U_1^{(k_3)}$ and $U_1^{(k_5)}$, and so may be written as $U_1^*S_{11}^{-1}T_{11}U_1$. As S and T enter into the policy operator only through this quantity, we thus have that the that the sequence of limits of the unitary Schur operators chosen does not affect the calculated equilibrium outcome.

Given these preliminary results, we are now ready to construct an approximation of the (first derivative of the) equilibrium policy function. Recall that the minimum norm equilibrium map from the predetermined variables X to the jump variables Y, if a stable equilibrium exists, is given by $g_X = -\varphi^{Y*}U_2^*(U_2\varphi^Y\varphi^{Y*}U_2^*)^{-1}U_2\varphi^X$ and the map from the predetermined variables to these variables next period is given by $h_X = (U_1(\varphi^X + \varphi^Y g_X))^{-1}S_{11}^{-1}T_{11}U_1(\varphi^X + \varphi^Y g_X)$ which is equal to

$$h_X = (\varphi^{X*}\varphi^X + g_X^*g_X)^{-1}(g_X^*\varphi^{Y*} + \varphi^{X*})U_1^*S_{11}^{-1}T_{11}U_1(\varphi^X + \varphi^Y g_X)$$

by 3.1. Given these formulations of the equilibrium functions, we may construct approximations by simply plugging in the estimates of the components that have already been constructed. For an appropriate sequence of tuning parameters, this produces an estimator of g_X which is consistent in strong operator topology and two

possible estimators of h_X: one which is consistent in strong operator topology, using the estimate and assumptions from Part (i) of 3.7, and one which is consistent in operator norm, using the estimate and assumptions from Part (ii) of 3.7.

Theorem 3.1. *Let the conditions of 3.3 hold, and also assume $U_2\varphi^Y\varphi^{Y*}U_2^*$ has bounded inverse. Let*

$$g_X^{k_4} = -\varphi^{Y*}U_2^{k_4,\varepsilon_{k_4}*}(U_2^{k_4,\varepsilon_{k_4}}\varphi^Y\varphi^{Y*}U_2^{k_4,\varepsilon_{k_4}*})^{-1}U_2^{k_4,\varepsilon_{k_4}}\varphi^X$$

and

$$h_X^{k_4,k_5} = (\varphi^{X*}\varphi^X+g_X^{k_4*}g_X^{k_4})^{-1}(g_X^{k_4*}\varphi^{Y*}+\varphi^{X*})U_1^{k_5*}(S_{11}^{k_5*}S_{11}^{k_5})^{-1}S_{11}^{k_5*}Q_1^{k_3}(Q_1^{k_3*}T_{11}^{k_3}U_1^{k_3})(\varphi^X+\varphi^Y g_X^{k_4}).$$

Then, for any $x \in \mathcal{H}_X$, there exist sequences k_3, k_2, k_1, ε_{k_3}, and ε_{k_4}, all depending on k_4, such that for any $\epsilon > 0$, $\exists K_4$ such that $\forall k_4 > K_4$, $\|(g_X^{k_4} - g_X)x\| < \epsilon$. Likewise, $\forall x \in \mathcal{H}_X$, there exist sequences k_5, k_4, k_3, k_2, k_1, ε_{k_3}, ε_{k_4}, and ε_{k_5}, depending on k_4 and k_5 such that for any $\epsilon > 0$ $\exists K_4, K_5$ such that $\forall k_4 > K_4$ and $\forall k_5 > K_5$ $\|(h_X^{k_4,k_5} - h_X)x\| < \epsilon$. (ii) Suppose also that Ω_1 and Ω_2 are compact operators. Let $h_X^{k_6} =$

$$(\varphi^{X*}\varphi^X+g_X^{k_4*}g_X^{k_4})^{-1}(g_X^{k_4*}\varphi^{Y*}+\varphi^{X*})U_1^{k_5*}(S_{11}^{k_5*}S_{11}^{k_5})^{-1}S_{11}^{k_5*}Q_1^{k_3*}thresh_{k_6}(Q_1^{k_3*}T_{11}^{k_3}U_1^{k_3})(\varphi^X+\varphi^Y g_X^{k_4}).$$

Then there exist sequences k_6, k_5, k_4, k_3, k_2, k_1, ε_{k_3}, ε_{k_4}, and ε_{k_5}, all depending on k_6, such that $\|h_X^{k_6} - h_X\| \to 0$ as $k_6 \to \infty$.

3.3.5 Advice on Tuning Parameter Selection

While the above results provide guarantees that for some values of tuning parameters k_1 to k_5, ε_{k_3}, ε_{k_4}, ε_{k_5} and possibly k_6 the policy operators produced by the algorithm will accurately characterize the true equilibrium operators, the guarantee

is qualitative rather than quantitative. Largely this is an unavoidable result of the nature of the problem. In order to permit the endogenous state variables to vary in a way which is restricted only by the economic characteristics of the decision problem rather than ex ante restriction to a smaller space not justified by economic theory, one must permit unrestricted dependence of equilibrium decisions on a topologically 'large' space of potential inputs. This fact, along with the requirements of spectral theory that the separation of inputs into 'forward' and 'backward' looking subspaces be mediated by norm-preserving (unitary) transformations, result in the necessity of estimating operators which are not compact. Operators which are not compact may not be consistently approximated in uniform norm. As a result, components which depend on these operators must be constructed from procedures which approximate different parts of the space at different rates, and so for particularly poor choices of input may require arbitrarily large approximating subspaces. In consequence, no quantitative guidance may be provided for choices of k_2, k_3, and k_4 which applies uniformly over all inputs. Nevertheless, one may hope that for some reasonable choices of input, and for the other tuning parameters, that more can be said, and indeed this is the case.

A common feature of the structure of economic models in this class which permits somewhat more to be said about choice of approximation is that even when asymptotic diagonality fails, there may exist a 'privileged' subset of $\mathcal{H}_1$ on which one expects most of the variation in endogenous objects to reside. This is easiest to see when the conditions for operator norm consistent estimation of h_X are met: the equilibrium operator may be split into two compact components, Ω_1 and Ω_2. By compactness, for any desired level of approximation ϵ, there exists a finite-dimensional subspace, given by the span of the right singular vectors of Ω_1 and Ω_2 in decreasing order by singular value, to which the restriction of these operators yields error no more than ϵ uniformly over inputs in $\mathcal{H}_1$. By choosing π_2 so that the subspace it spans contains

this subspace, and the span of π_1 to contain the span of the corresponding right singular vectors, the action of the equilibrium conditions may be well-approximated within the forward and backward subspaces. In particular, the component of h_X given by action on the subspace of $\mathcal{H}_X$ corresponding to eigenvectors inside the unit circle is given by a compression of Ω_1 onto the subspace spanned by this eigenspace[5], and so must also be well approximated by the first left and right singular vectors of Ω_1. This suggests also choosing these singular vectors to generate Z^{k_3} (and Z^{k_5}). Similarly, the components of U_2 which generate those components of g_X which map onto the principal singular vectors of h_X may likewise be chosen for Z^{k_4}. The number of such functions used to achieve a given degree of accuracy is determined by the rate of decay of the singular values of Ω_1 and h_X, with faster decay indicating that a smaller number of functions is needed to achieve a given level of uniform accuracy.

In practice, choosing singular vectors or eigenvectors of the equilibrium conditions is infeasible, as they are unknown, which is the entire point of approximating them numerically. However, the rates achievable with perfect knowledge of these components represent a lower bound on the number of basis functions needed to achieve a given level of accuracy. While the exact number needed is not knowable in general, one may use qualitative properties of the equilibrium conditions to provide estimates. In particular, the rate of decay of the singular values determines the approximation number of the compact components for which bounds exist (Carl & Stephani, 1990, Ch. 2). One may use Jackson-type inequalities to convert knowledge of the properties of the range of a compact operator into estimates of singular values. In general, these inequalities say that if an operator maps the unit ball in Hilbert space to a set which is 'small', in the sense of having a small metric entropy or covering number, then then the operator has correspondingly rapid decay of singular values. The archetypal example is the case in which the output of an operator maps ar-

[5] i.e. $U_1^* S_{11}^{-1} T_{11} U_1 = U_1^* U_1 \Omega_1 U_1^* U_1$

bitrary functions to smooth functions, in the sense of having derivatives which are bounded or square integrable. Standard choices of basis functions, such as trigonometric or orthogonal polynomials, yield an approximation error which is uniformly small over all smooth functions, and so an operator whose output is known to be smooth must have range well-approximated by a finite-dimensional subspace, and in particular must have singular values which decay at least as fast as the approximation error using standard basis functions. The dependence of the approximation error on the degree of smoothness is well known for many choices of basis function. For trigonometric polynomials, the error in approximating a function (in L^2 on a bounded domain) decays exponentially in number of basis functions used for functions which are infinitely differentiable, and polynomially of order k for functions with discontinuous $k-1^{st}$ derivative: i.e., functions with a discontinuity have approximation error which decays linearly, while continuous functions with a break in the first derivative have error which decays quadratically. The same general rates apply to Chebyshev polynomials: see Boyd (2000, Ch. 2) for precise results and a useful discussion of heuristics and rules of thumb for choosing an appropriate number of basis functions given only qualitative information about the function to be approximated. Note that while information about the range of compact components does not imply that the right singular vectors satisfy similar smoothness properties, the same techniques may be applied to the adjoint operators to provide similar bounds. Take as an example a canonical choice of compact operator, a Fredholm integral operator of the first kind $f(x) = \int k(x,y)[f(y)]dy$. Under reasonable conditions, if the kernel $k(x,y)$ has square integrable k^{th} derivative with respect to x, it will map $f(y) \in L^2$ to $f(x)$ with square integrable k^{th} derivative. Similarly, the adjoint operator $f(y) = \int k(x,y)[f(x)]dx$ will have smooth range if the kernel $k(x,y)$ has square integrable k^{th} derivative with respect to y.

While smoothness principles are useful for characterizing both the rate at which

singular values may be expected to decay and providing a guess at which sets of basis functions may be used to accurately approximate the compact components, one must be aware that smoothness conditions do not imply that singular vectors or eigenvectors may be approximated at standard rates for smooth functions. The reason for this is simple: the full set of singular vectors must span the entire space on which the operator acts, but the set of smooth functions is only a small subset of this larger space. While the first few singular vectors of a smooth operator will themselves be smooth, higher order singular vectors must be orthogonal to the preceding set of smooth singular vectors, and so cannot themselves be smooth. As the size of the singular value decreases, the degree of smoothness, and so approximability by a fixed set of basis functions, of the corresponding singular vector decreases linearly. For example, the eigenfunctions of a convolution operator are the trigonometric polynomials, and so to approximate the n^{th} singular vector to error of norm less than one, one must use exactly n trigonometric polynomials, and no other set of basis functions can do better. It is for this reason that no uniform approximation rates can be given for the unitary components of the generalized Schur or eigendecomposition.

In the case where compactness on subspaces does not hold, similar principles apply. While the domain of (B, A) consists of the entirety of $\mathcal{H}_1$ and due to the presence of non-compact components the restriction of (B, A) to any finite-dimensional subspace will always be uniformly bounded away from the truth, it is often the case that 'many' or 'most' components of (B, A) are compact. As a result, these components may be uniformly approximated on a finite-dimensional subspace, outside of which they are approximately equal to zero. Outside of this space, even if (B, A) is not asymptotically diagonal, it may take a relatively simple structure, approximately, given block-wise by a mix of identity operators and multiplication by smooth functions, which preserve the smoothness properties of the input. As a result, one may hope that after a sufficient number of basis functions are used to well-approximate the compact component,

the equilibrium operators map inputs of higher order basis functions to outputs do not require substantially higher order basis functions to approximate, and so the equilibrium policy operators are well approximated on the space of inputs of lower order. Note that this requirement is far from trivial. For a smooth multiplication operator, such as a low order polynomial, a polynomial input may be mapped to an output which requires only a slightly higher order polynomial to approximate well. Because this applies no matter how high the order of the input, no square approximation of the operator is uniformly valid, but the degree of uneven-ness may be bounded. In contrast, such a pattern could continue indefinitely if one component of the equilibrium conditions is a unitary operator which rotates inputs well represented in one basis into outputs represented in a completely unknown basis: this is the worst case scenario for the algorithm described in terms of rates, as one could potentially need an arbitrarily large number of additional basis functions to approximate the output of one additional basis function provided as input. However, such a component does not seem to be a natural element of most economic models.

Slightly more realistically, an economic model may have components with singular functions which are mutually incoherent, so that a set of functions which well represents one operator must poorly represent the other. In this case, so long as at least one of the components is compact, one may include an additional set of basis functions which well-represent the finite-dimensional space on which the bulk of the energy of the offending component resides, increasing the required number of functions for a given approximation error but not increasing the required rate of growth for each component. The typical example of this is when one component is well represented in a Fourier basis (i.e., it is smooth), and the other is well-represented by spatially localized functions (i.e., it is spiky). In this case one may use either the span of two sets of basis functions, trigonometric polynomials and functions well-representing spikes, or use a basis which may approximate both to reasonable accuracy. When this occurs,

the number of basis functions needed for a given accuracy will be at least the worse of the two rates required for accurate representation of each, with additional overhead or an increased constant factor induced by the additional component.

For the tuning parameters used to ensure that inverses are stable through use of uneven section methods, viz. k_1 relative to k_2 and k_3 relative to k_5, aside from needing to be at least as large as their counterpart, the Cauchy interlacing argument provided does not provide quantitative guidance on reasonable ratios.[6] As usual, this is because in large spaces, it is in principle possible for the choice of right approximating subspace and left approximating subspace to be so mismatched relative to the operator that an arbitrarily large number of left bases is needed for any given right input to ensure the minimum singular values is bounded away from 0. But we may again consider 'typical' behavior by examining classes of operators believed to be qualitatively similar to those which will appear in standard models. In particular, given choices of basis functions which are sensible in they provide accurate approximations to the ordered set of singular functions of the model with not too many additional functions, the output generated by expanding the input space by one dimension should be well-approximated with a small number of extra basis functions. This can be and has been quantified for several important cases. For the unitary transform (equivalent to an infinite-dimensional Gram matrix) between trigonometric polynomials and Daubechies wavelets, both reasonable choices for representing smooth functions and so likely to be typical of the discrepancy between the basis functions typically chosen and the unknown singular functions of the operators involved in the equilibrium conditions, Adcock *et al.* (2014b) show that a constant ratio of trigonometric polynomials to wavelets is sufficient to preserve the size of the minimum singular value as the number of wavelets used goes to infinity. Depending on the order of the wavelet basis and exact degree of approximation, the required constant is small, around a

[6] k_3 must also be large relative to k_4 for these purposes, but the dominant factor driving their relationship is the necessity of the pointwise convergence of projections in 3.5.

factor of 2 to 4. They note however that oversampling is necessary: for ratios below a computed constant, the minimum singular value decays exponentially quickly as the number of wavelets grows. This suggests that for components of the proposed algorithm which rely on similar oversampling and for reasonable basis function choice, while some constant factor of oversampling is needed, the degree of oversampling requires need not grow with the number of basis functions, at least if similar results hold: k_1 may be chosen as Ck_2 and k_3 as Ck_5, for some C on the order of 2 to 4, or slightly larger to be safe. For this ratio one also has a potential diagnostic: one can simply look at the minimum singular value of the operators $\lambda\tilde{A}$-$\tilde{B}$, $S_{11}^{k_5}$, and $U_2^{k_4,\varepsilon_{k_4}}\varphi^Y$ for different values of k_1, k_3, and k_3, respectively. While a small value is not necessarily indicative of a need for a higher degree of oversampling, as nothing requires the true minimum singular values of these operators to be of a particular scale, a relatively large value (on the order of 0.1 would certainly count as large) which does not substantially change with k_1 or k_3 respectively will ensure stability of the inverse.

Another set of tuning parameters used are the values ε_{k_3}, ε_{k_4}, and ε_{k_5}, used to ensure that the orthonormalization process does not blow up numerical noise into fictitious singular vectors. In general, the procedure should not be strongly sensitive to choice of these parameters. To induce a singularity, a vector in Z^{k_3}, Z^{k_4}, and Z^{k_5} must not only be near the boundary of the subspace onto which it is being projected, but the proportion of this vector lying in the subspace must be arbitrarily close to 0. As such, it may often be the case that all vectors are far enough from this edge case that a relatively wide range of choices for ε will lead to exactly the same numerical output. This is the case, for example, if Z^{k_3} and Z^{k_5} are chosen optimally, as vectors completely inside the interior eigenspace, and will continue to be true for values chosen anywhere approximately close to this choice, where the approximation need not increase in accuracy with sample size. More generally, if ε is chosen slightly on the large side and so excludes a singular vector, this has the effect of reducing the

dimension of the approximating subspace by 1. Unless the target of interest has substantial mass within the one-dimensional subspace spanned by this singular vector, the pointwise approximation should remain nearly as accurate, equivalent to using just one fewer basis function. An ε which is slightly too small will in contrast end up magnifying numerical noise and so including projection in a direction which is close to random. A potential diagnostic is to examine the smallest singular values of the projected matrices to see what changes in ε end up excluding, and to check if the singular vectors it excludes are stable as the dimension of the projector changes, which one hopes would characterize true singular vectors, or vary substantially, which is what would be expected for numerical error. While again this may not perfectly distinguish, as true singular vectors associated to small singular values will also generally be quite variable at moderate degrees of accuracy, as suggested by Wedin's Sin Θ theorem, this is arguably a positive contribution, as a poorly estimated contribution from a true singular vector may itself result in substantial approximation error relative to the small reduction in bias of adding one more basis function.

In conclusion, basis functions should be chosen in accordance with the known smoothness properties of the model, in numbers large enough to ensure numerical stability of the outputs. One wants enough basis functions chosen to yield numerical accuracy on typical inputs, which should have polynomial size representations if moderately smooth or logarithmic size representations if supersmooth, and enough basis functions must be chosen in each step to ensure stability of the previous step. Typically, one expects for a model with well-chosen basis functions each k_1 to be of the same order of magnitude as k_2 and k_3 to be of the same order of magnitude as k_5, but a constant multiple larger. For k_2 relative to k_3, k_3 relative to k_4, and k_4 and k_5 relative to k_6, convergence of projections pointwise should follow rates for nonparametric projection estimation of smooth functions, with the improvement that the bases used for projection are redundant across the different functions to be estimated, so that

rather than compounding, the error is primarily driven by the approximation in the hardest function to approximate, which, since eigenfunctions and similar bases are generally decreasing in smoothness as order increases, is usually the last, in order. Although not foolproof, these principles should allow selection of tuning parameters of reasonable magnitudes: further refinement may then be achieved by numerical experimentation and examination of singular values and singular vectors.

3.4 Application and Evaluation

3.4.1 Example Model: Target Tracking with Adjustment Costs

To fix ideas regarding the framework provided, we provide a simple function-valued model to which the decomposition and approximation procedures may be applied. This model of decision-making under adjustment costs provides an illustration of the concepts involved in solving a high-dimensional rational expectations model, while maintaining a high enough degree of tractability that an explicit (partially) closed form solution is available. This may be used both to describe the structure of a stationary rational expectations equilibrium in Hilbert space and to assess the properties of a numerical solution method. The concept is that each period, the agent must choose a function so as to minimize the distance of this function from an independent and identically distributed function valued target, subject to an adjustment cost quadratic in the distance of the current choice from a shifted version of last period's decision. The agent's problem is given as

$$\max_{\{a_t\}_{t=0}^{\infty}} \mathbb{E} \sum_{t=0}^{\infty} \beta^t (\|a_t - \eta_t\|_{\mathcal{H}}^2 + \|B[a_t - a_{t-1}]\|_{\mathcal{H}}^2)$$

where $a_t \in \mathcal{H}$ an infinite-dimensional separable Hilbert space, B is a bounded compact linear operator $\mathcal{H} \to \mathcal{H}$ parameterizing the adjustment cost, and η_t is an$\mathcal{H}$-

valued exogenous target satisfying $\eta_{t+1} = \sigma z_{t+1}$, where z_{t+1} is a noise term independent and identically distributed over time with mean 0 and covariance operator $\Sigma[.] := \mathbb{E}\langle z_t, [.]\rangle z_t$.

Although this problem is primarily designed as a purely formal exercise to investigate the properties of a solution method, discussion is in order about its potential economic interpretation. In form, this problem is most similar to the Rotemberg (1982) model of optimal pricing subject to quadratic adjustment costs. In such a model, we may interpret the choice variable a_t as representing the prices charged by a monopolist for a continuum of goods and η_t as the period-t static profit-maximizing price profile given product-specific but correlated demand shocks z_t. In this case, B might represents a distance and weighting function indicating which goods or, potentially, combinations of goods, have prices which are more costly to adjust. Assuming B is self-adjoint and has a discrete spectral decomposition $B[.] = \sum_{i=1}^{\infty} \lambda_i \langle \phi_i, [.]\rangle \phi_i$, we may interpret each eigenfunction ϕ_i as a bundle of goods, and the corresponding eigenvalue as the cost of adjusting that particular bundle: a large λ_i means changing the price of bundle ϕ_i is difficult, while a small λ_i means that changing the price of the corresponding bundle is easy.

This model may admit other interpretations as well, depending on the space $\mathcal{H}$ chosen and the adjustment cost operator B. For example, this may represent a communications problem. Each period, a monetary policy maker observes a set of economic signals, consisting of quantitative economic indicators and verbal descriptive reports. To influence market sentiment in a desired direction, the policy maker chooses an announcement, which may naturally be thought of as an element of the high-dimensional vector space of documents composed of a set (not necessarily finite) number of words, to be as close as possible to the optimal statement given market information. However, statements which deviate strongly from past statements may result in undesirable confusion, and so the policy maker must also take into account

the costs of increasing 'semantic' distance from previous statements. In this case, a_t is the policy-maker's public statement and B represents a distance function on the space of documents, possibly weighted by economic importance of the words. See, for example, Coifman & Maggioni (2006) or Mohri *et al.* (2012) for examples of vector-space representations of documents and distance functions characterizing semantic similarity.

Other interpretations may include personal style or fashion. Each period, some appearance or social behavior η_t may be considered fashionable, but an individual faces costs, potentially both monetary and in terms of desire to maintain a coherent personal identity, as in Akerlof & Kranton (2000), of changing their personal style. Given a vector space representing components of personal identity, the operator B parameterizes which are the most psychologically fundamental and which may be adjusted freely without damaging one's self-image.

The feature which makes this model require a nontrivial decomposition into stable and unstable subspaces is the adjustment cost. The presence of the adjustment cost induces backward looking behavior because past decisions influence the current state and so current decisions, but also induces forward looking behavior, since an agent must take into account the effect of their current decision on all future decisions. The transversality condition induces selection of a stationary path. Economically, the agent recognizes that backward looking selection of a component which corresponds to an unstable eigenvalue may end up, due to adjustment costs, causing the agent to select a value further and further from the expected optimal state (here normalized to 0), as adjustment back would be too costly, and thus incurs a large expected loss with chosen value far away from expected η_t for a considerable duration.[7] Decision

[7]Technically, due to discounting, the transversality condition only requires that discounted utility not explode, which may be satisfied for any decision path inside the complex circle around the origin of radius $\frac{1}{\beta}$. While the stationary solution remains a solution to the optimization problem, the solution algorithm may easily be set to integrate over the appropriate path in order to incorporate these additional components.

rules generating sample paths in violation of the transversality condition correspond to suboptimal 'bubbles'. In the communication example, committing to a decision rule which violates this condition might correspond to making a statement which it would be very costly to walk back, requiring increasingly strong statements each period in order to clarify the previous statement, eventually resulting in need to say arbitrarily absurd things (from the perspective of the static optimal choice) simply in order to cover for previous missteps. In the identity example, this would correspond to cultivating an increasingly extreme and unpopular idiosyncratic style to maintain one's personal self-concept. Note though that an agent who is not committed to a suboptimal rule would revert to the optimal policy rule immediately if it perceived such a spiral out of control to be occurring, as there is no time consistency problem in this model. In the opposite direction, forward looking assignment of components corresponding to stable paths incurs a loss in deviating from current period optimal policy, for bounded future gain, and so the agent prefers to set these components based on past decisions and current shocks. Precisely which components are more costly to solve forward and which are more costly to solve back depends strongly on the characteristics of the adjustment cost parameterization B.

Taking first order conditions, we have

$$\langle a_t - \eta_t, . \rangle + \langle B(a_t - a_{t-1}), B. \rangle = \beta \mathbb{E} \langle B(a_{t+1} - a_t), B. \rangle$$

simplifying to the condition

$$(I + (1+\beta)B^*B)a_t - \eta_t - B^*Ba_{t-1} = \beta \mathbb{E} B^*Ba_{t+1}$$

To represent the model in the format described, add the equations $\eta_{t+1} = \sigma z_{t+1}$ and the auxiliary variable $u_{t+1} = a_t$ so that only time t and $t+1$ terms are included.

With this, the system may be represented as

$$\begin{bmatrix} \beta B^*B & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{bmatrix} \begin{bmatrix} \mathbb{E}a_{t+1} \\ u_{t+1} \\ \eta_{t+1} \end{bmatrix} = \begin{bmatrix} I + (1+\beta)B^*B & -B^*B & -I \\ I & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a_t \\ u_t \\ \eta_t \end{bmatrix} + \sigma \begin{bmatrix} 0 \\ 0 \\ z_{t+1} \end{bmatrix}$$

Note that this model is asymptotically diagonal, and so may be analyzed by the procedures of Chapter 1. However, as it again has an exact solution, it does provide a useful test case for the algorithm introduced here, which may be used even when the problem is not ill-posed, if this characteristic is not known ex ante.

3.4.2 Construction of Exact Solution

Due to the quadratic form of the utility function, this problem results in a characterization of the solution which is linear without approximations. In this representation, a_t is y, the jump variable, and u_t, η_t are components of x, the predetermined variable. To simplify notation, let $L := B^*B = \sum_{i=1}^{\infty} \lambda_i \langle \phi_i, . \rangle \phi_i$, where $\{\phi_i\}_{i=1}^{\infty}$ is assumed to be a complete orthonormal eigenbasis of $\mathcal{H}$ and $\{\lambda_i\}_{i=1}^{\infty}$ are non-negative eigenvalues. Using this spectral decomposition, it is possible to solve the generalized eigenvalue problem analytically to construct a generalized Schur decomposition of the matrix pair. By taking the (unbounded) inverse of the matrix of coefficients on time $t+1$ variables, the problem may be converted into a standard eigenvalue problem for the unbounded operator given by

$$C = \begin{bmatrix} \frac{1}{\beta}L^{-1} + \frac{1+\beta}{\beta}I & -\frac{1}{\beta}I & -\frac{1}{\beta}L^{-1} \\ I & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

As all elements are functions of L, I and 0, the spectral subspaces may be defined in terms of the eigendecomposition of L. First, note that this operator has a non-

trivial null space on $\mathcal{H} \times \mathcal{H} \times \mathcal{H}$, which, since the null space of I is $\{0\}$ and so the first component of each element must be 0, contains elements of the form $[0', a', b']'$ such that $-\frac{1}{\beta} Ia - \frac{1}{\beta} L^{-1} b = 0$ or $b = -La$. As a result, the null space is given by the closed span of the set $\{[0', a', (-La)']' : \ a \in \mathcal{H}\}$. This set is the spectral subspace corresponding to generalized (generalized) eigenvalue 0. Note that the closure condition is important here: $\{La : \ a \in \mathcal{H}\}$ is dense in $\mathcal{H}$ but there exist elements of $\mathcal{H}$ outside of it. Those elements in the closure of this set correspond to points at which an inverse of C may be defined but is not continuous, and so properly correspond to 0 as an element of the continuous rather than point spectrum. Economically, this may be interpreted as saying that the system treats identically a current state $x = \{u_t, \eta_t\} = \{0, 0\}$ and a state $x = \{u_t, \eta_t\} = \{a, -La\}$ for any $a \in \mathcal{H}$: that is, there exists, for any value of the shock, a value of the past decision which would "exactly cancel it out," leaving the decision problem the same, or, in the cases in the closure, a sequence of past values which come arbitrarily close to canceling out the shock. Due to the fact that this continuous component of the spectrum is restricted to the null space, this construction does not induce any difficulty in distinguishing the interior and exterior components of the spectrum, which remain discrete.

Next we may consider the non-zero elements of the spectrum. As the third row maps all elements to 0, non-zero eigenvectors must have 0 as the final element. Further, since L has an orthonormal eigenbasis, they must be combinations of eigenvectors of L. So, eigenfunctions of C are of the form $\varphi_i \propto [a_i \phi_i', \phi_i', 0']'$ for scalar a_i, where the proportionally reflects setting the coefficient on the second component to 1, as eigenfunctions are identified only up to scale. Setting up the problem $\gamma_i \varphi_i = C\varphi_i$, this reduces to the quadratic equation $\gamma_i = a_i = [(\frac{1}{\beta\lambda_i} + \frac{1+\beta}{\beta})a_i - \frac{1}{\beta}]\frac{1}{a_i}$. Setting $z_i = (\frac{1}{\beta\lambda_i} + \frac{1+\beta}{\beta})$ and applying the quadratic formula, we find two roots, $\gamma_i^{\pm} = \frac{1}{2}(z_i \pm \sqrt{z_i^2 - \frac{4}{\beta}})$ corresponding to eigenvectors proportional to $\varphi_i^+ \propto [\gamma_i^+ \phi_i', \phi_i', 0']'$ and $\varphi_i^- \propto [\gamma_i^- \phi_i', \phi_i', 0']'$, respectively. Note that by compactness and self-adjointness, λ_i

are real and non-negative and $\lambda_i \searrow 0$ as $i \to \infty$, $\gamma_i^+ \nearrow \infty$ and $\gamma_i^- \searrow 0$ as $i \to \infty$. By the discreteness of the spectrum of L and the fact that 0 is its unique accumulation point, at most a finite number of γ_i^+ will be inside the unit circle and at most a finite number of γ_i^- will be outside. In fact we can say more. Economically interesting parameters require that $\beta \in (0,1)$, and since λ_i are real and non-negative, $z_i \geq \frac{1+\beta}{\beta}$ for all i. Plugging this restriction into $\gamma_i^{\pm}$ yields that $\gamma_i^+ > \frac{1}{\beta} > 1$ for all i and $\gamma_i^- < 1$ for all i, so the spectrum is well-separated and never intersects the unit circle. We see then that the discrete spectrum of the operator pair contains an infinite number of components inside the unit circle, corresponding to eigenvalues γ_i^- and eigenvectors φ_i^- and an infinite number outside, corresponding to eigenvalues γ_i^+ and eigenvectors φ_i^+. By the completeness and orthogonality of $\{\phi_i\}_{i=1}^{\infty}$, it is apparent that the union of the closed span of $\{\varphi_i^{\pm}\}_{i=1}^{\infty}$ and the 0 spectral subspace defined above span $\mathcal{H} \times \mathcal{H} \times \mathcal{H}$ and so provide a full characterization of the spectrum of C and so of the generalized spectrum of the equilibrium operators. To decompose into forward and backward subspaces, the forward subspace corresponds to all eigenvalues γ_i^+, while the backward subspace corresponds to the 0 spectrum and all eigenvalues γ_i^-. By the orthogonality of $\{\phi_i\}_{i=1}^{\infty}$, $\{\varphi_i^+\}_{i=1}^{\infty}$ and $\{\varphi_i^-\}_{i=1}^{\infty}$ are each orthogonal sequences: however, they are not mutually orthogonal: if eigenvectors are normalized to have norm one, $\left\langle \varphi_i^+, \varphi_i^- \right\rangle = (\gamma_i^+ \gamma_i^- + 1)/(\sqrt{1+\gamma_i^{+2}}\sqrt{1+\gamma_i^{-2}})$. However, orthogonality does imply $\left\langle \varphi_i^+, \varphi_j^- \right\rangle = 0$ for $i \neq j$. To construct the Schur vectors, it is sufficient to orthogonalize the eigenfunctions corresponding to the interior and exterior subspaces. We may also note that for any vector v in the 0-spectrum subspace, it may be represented as $v = \sum_{i=1}^{\infty} \frac{c_i}{\sqrt{1+\lambda_i^2}}[0', \phi_i', -\lambda_i \phi_i']'$ for some sequence c_i, and so we may take the components of this sum $v_i = \frac{1}{\sqrt{1+\lambda_i^2}}[0', \phi_i', -\lambda_i \phi_i']'$ to be a basis for the zero spectral subspace, with the property that $\left\langle \varphi_i^{\pm}, v_i \right\rangle = \frac{1}{\sqrt{1+\lambda_i^2}\sqrt{1+\gamma_i^{\pm 2}}}$ and $\left\langle \varphi_i^{\pm}, v_j \right\rangle = 0$ for $i \neq j$.

As a result, we may construct the Schur matrices as follows. Take as the first

set of rows of U_1 the eigenvectors φ_i^-. Next, orthogonalize the 0 spectral subspace to the previous space by projecting, for each i, the vector $\frac{1}{\sqrt{1+\lambda_i^2}}[0', \phi_i', -\lambda_i\phi_i']'$ onto the orthogonal complement of the span of φ_i^-. For each i, this generates the vector $\frac{1}{\sqrt{(\frac{-\gamma_i^-}{1+\gamma_i^{-2}})^2+(\frac{\gamma_i^{-2}}{1+\gamma_i^{-2}})^2+\lambda_i^2}}[\frac{-\gamma_i^-}{1+\gamma_i^{-2}}\phi_i', \frac{\gamma_i^{-2}}{1+\gamma_i^{-2}}\phi_i', -\lambda_i\phi_i']'$. We may then append the complete set of these rows to the analysis operator generating coefficients with respect to φ_i^- to generate U_1. To calculate U_2, we take the eigenvectors φ_i^+ and, for each i, project them onto the orthogonal complement of the span of the interior eigenvector with the same index and $\frac{1}{\sqrt{1+\lambda_i^2}}[0', \phi_i', -\lambda_i\phi_i']'$. The analysis operator onto this basis generates U_2. The utility of this procedure is that to calculate any single Schur vector, it requires access to at most a finite number of basis vectors (in fact, no more than two). As such, while the Schur decomposition in theory requires knowledge of the entire space to construct, and particularly so for U_2, which must be orthogonal to an infinite-dimensional space, the orthogonality properties characterizing this particular operator ensure that the components can be separated and so individual vectors can be constructed without accessing the full space. Q_1 may be constructed, as before, by orthonormalizing the operator AU_1, which again has a closed form solution whose calculation requires access to only a finite set of eigenvectors. Columns are given by $\frac{1}{\sqrt{1+(\beta\gamma_i^-\lambda_i)^2}}[-\beta\gamma_i^-\lambda_i\phi_i', \phi_i', 0']'$ corresponding to φ_i^- and $\frac{1}{\sqrt{(\frac{-\beta\gamma_i^-\lambda_i}{1+(\beta\gamma_i^-\lambda_i)^2})^2+(\frac{(\beta\gamma_i^-\lambda_i)^2}{1+(\beta\gamma_i^-\lambda_i)^2})^2+\lambda_i^2}}[\frac{-\beta\gamma_i^-\lambda_i}{1+(\beta\gamma_i^-\lambda_i)^2}\phi_i', \frac{(\beta\gamma_i^-\lambda_i)^2}{1+(\beta\gamma_i^-\lambda_i)^2}\phi_i', -\lambda_i\phi_i']'$ corresponding to the image of the null space, both for $i = 1 \ldots \infty$. T_{11} and S_{11} may then be calculated as $Q_1^*BU_1$ and $Q_1^*AU_1$, respectively, and h_X and g_X can be constructed from the formulas $g_X = -(U_2\varphi^Y)^{-1}U_2\varphi^X$ and $h_X = (I + g_X^*g_X)^{-1}(U_1\varphi^X + U_1\varphi^Y g_X)^* S_{11}^{-1}T_{11}(U_1\varphi^X + U_1\varphi^Y g_X)$, where all elements are truncated to contain only those components acting on $\{\phi_i\}_{i=1}^n$.

By the orthogonality of these components with $i \leq n$ and those with $i > n$, this provides an exact characterization of h_X and g_X on inputs in the closed span of (the

tensor product over $\mathcal{H} \times \mathcal{H} \times \mathcal{H}$ of) $\{\phi_i\}_{i=1}^n$. While the complexity of the formulas impedes precise discussion of the qualitative form of the optimal policy, it does permit exact numerical calculations. One may note that the lack of interaction between eigenfunctions indicates that the policy takes a form in which the decision regarding a particular component is determined only by shocks to a particular component and past values of that component. Numerical results in Section 2 demonstrate that the policy functions have intuitive properties. As is apparent from the utility functions, high values of η and u in a component result in a higher contemporaneous choice of the corresponding component of a. Components corresponding to eigenfunctions of L with larger eigenvalues are highly persistent, reflecting the higher costs of adjusting them from past values, while components corresponding to low eigenvalues have little persistence. In part, this reflects the i.i.d. assumption for the exogenous target, which results in little amplification of the persistence induced by adjustment costs.

3.4.3 Solution Accuracy in the Adjustment Cost Model

The quadratic adjustment cost model described earlier provides a useful test case for the approximation algorithm, as it requires a nontrivial generalized Schur decomposition, but this decomposition has a known form. In particular, due to the explicit construction of the generalized Schur vectors, which do not interact directly, when restricted to inputs in the span of a finite subset of Schur vectors, the solution may be calculated exactly from the finite-dimensional submatrices. One implication of this fact is that, were A and B to be approximated by their finite-dimensional projections onto this span, standard linear rational expectations solution algorithms would produce exact answers for band-limited inputs, making this a case in which new techniques are unnecessary. When the eigenfunctions of the equilibrium condition operator are not known, approximation error once again becomes a concern, though again due to the asymptotic diagonality, this approximation error could also

be handled by the methods in Chapter 1. In order to provide at least some challenge, rather than use the eigenfunctions as a basis, I evaluate the solution algorithm by approximating the equilibrium conditions in a basis other than the eigenbasis before applying the solution algorithm. One may see by inspection that for large i, due to the presence of the identity components, the elements of T_{11} do not converge to 0 and as a result, h_X is not a compact operator. As a result, to estimate h_X, we apply only the estimator $h_X^{k_4,k_5}$, which does not require compactness, and not $h_X^{k_6}$, which does.

To assess the quality of the approximation, we provide a calibration expected to be moderately challenging. We set $\lambda_i = \frac{15}{i^3}$, a cubic decay rate being characteristic of a kernel operator whose output has a second order singularity: that is, it outputs functions which are differentiable, but whose second derivative has a discontinuity, a situation which occurs frequently in economic models when decision makers expect a cost function which contains a kink at an unknown point, for example due to phase-in of a policy or to an endogenous constraint, as is common in models with collateral or limited liability. We set $\beta = 0.96$, within the usual range for discounting in macroeconomic models: at this value, $\gamma_i^+ > \frac{1}{\beta}$ for all i, and $\gamma_i^- < 1$, so the transversality condition holds so long as only eigenvectors corresponding to eigenvalues inside any complex circle of radius in the range $(\max \gamma_i^-, \min \gamma_i^+)$ are selected for the backward-looking component. In practical terms, this means that the complex unit circle is a consistent choice for the integration path when calculating the generalized Schur decomposition. Since the exact solution is available analytically, the particular choice of functions for the eigenvectors of L matters only in relation to the choice of basis functions used to approximate them. To that end, we consider a pair of basis choices which are simple to work with and likely to represent the degree of coherence between eigenfunctions and chosen representations likely to occur in practice. Following Adcock *et al.* (2014b), we choose trigonometric polynomials and Haar wavelets on the unit interval, which have the advantage that the change

of basis may be calculated explicitly in terms of sinc functions. The trigonometric polynomials $f_l(t) = e^{2\pi\imath lt}\mathbf{1}[t \in [0,1]]$ for $l = -\infty \dots \infty$ are the eigenfunctions of convolution, integration, and differentiation operators and so are common in applications. When the basis functions chosen are Haar wavelets, given by the functions $\phi_{0,0}(t) = \mathbf{1}[t \in [0,1]]$, $\psi_{j,k}(t) = 2^{j/2}\mathbf{1}[2^j t - k \in [0, \frac{1}{2})] - 2^{j/2}\mathbf{1}[2^j t - k \in [\frac{1}{2}, 1]]$, for $j = 0, \dots, \infty$, $k = 0, \dots, 2^j - 1$, trigonometric polynomials are a moderately difficult set of functions to represent: since these wavelets are discontinuous functions and have bounded and, for larger j, quite small, spatial support, not only do they have unbounded support in the frequency domain, the rate at which approximation error decreases as the number of trigonometric polynomials used to approximate one wavelet grows is relatively slow. Still, in spite of the discontinuity, both bases represent smooth functions well, and if the eigenfunctions are ordered from low to high frequency (we set $\phi_1 = f_0$, $\phi_i = f_{-\frac{i}{2}}$ i even, $\phi_i = f_{\frac{i-1}{2}}$ i odd), a Haar wavelet representation may achieve a high degree of accuracy with a moderate number of functions relative even to smoother choices such as orthogonal polynomials. Further, a Haar wavelet representation encodes exactly the same information as the popular histogram representation.

To evaluate the estimator, we use the same wavelet basis and number of basis functions on each of the three copies of $\mathcal{H}$, ordered hierarchically by scale and then shift, setting $k_1 = 3 \times 1024$, $k_2 = 3 \times 512$, $k_3 = 3 \times 256$, $k_4 = k_5 = 3 \times 128$, where for Z^{k_3}, Z^{k_4}, and Z^{k_5}, the first $k_3/3$, $k_4/3$, and $k_5/3$ wavelet basis functions are used, and as the exact wavelet coefficients of L are unavailable, these are calculated on its truncation to the first 2048 eigenfunctions, incurring an operator norm error of up to $\lambda_{2049} \approx 1.7 \times 10^{-9}$. Threshold parameters for truncating singularities are set at $\varepsilon_{k_3} = k_3\epsilon_m$, $\varepsilon_{k_4} = 2k_3\epsilon_m$, and $\varepsilon_{k_5} = 4k_3\epsilon_m$, for $\epsilon_m = 2^{-52}$: these are set to essentially exclude only singular values which are zero up to floating point error. Path integrals are computed numerically by parameterizing the unit circle and applying adaptive

Figure 3.4.1: IRF L^2 approximation error: Adjustment cost model via Haar basis

	$\eta_0 = u_0 = \sum_{i=1}^{64} \phi_i / \| \sum_{i=1}^{64} \phi_i \|$		$\eta_0 = u_0 = \sum_{i=1}^{1024} \frac{1}{i} \phi_i / \| \sum_{i=1}^{1024} \frac{1}{i} \phi_i \|$	
t	$\|u_t - \tilde{u}_t\| \times 10^4$	$\|a_t - \tilde{a}_t\| \times 10^4$	$\|u_t - \tilde{u}_t\| \times 10^4$	$\|a_t - \tilde{a}_t\| \times 10^4$
0	3.662	5.182	3.464	4.128
1	4.611	0.753	4.436	1.370
2	0.753	0.236	1.369	0.569
3	0.236	0.091	0.569	0.299
4	0.091	0.045	0.299	0.178
5	0.045	0.027	0.178	0.105
6	0.027	0.016	0.105	0.060
7	0.016	0.009	0.060	0.033

quadrature. Upon calculation of $g_X^{k_4}$ and $h_X^{k_4,k_5}$, the estimates are tested by calculating impulse response functions to some initial conditions representing different challenges. For the first, we take as initial impact $\eta_0 = u_0 = \sum_{i=1}^{64} \phi_i / \| \sum_{i=1}^{64} \phi_i \|$ and for the second, $\eta_0 = u_0 = \sum_{i=1}^{1024} \frac{1}{i} \phi_i / \| \sum_{i=1}^{1024} \frac{1}{i} \phi_i \|$. The first is strongly band-limited, but attaches high weight to middle frequencies resulting in a highly multimodal function. The second attaches some weight to very high frequencies which may not be well-approximated by the chosen basis, but the weight declines rapidly with frequency. Impulse response functions for predetermined variables are calculated as $(\tilde{u}_t, \tilde{\eta}_t) = (h_X^{k_4,k_5})^t (u_0, \eta_0)$ and $\tilde{a}_t = g_X^{k_4}(\tilde{u}_t, \tilde{\eta}_t)$, with $(\tilde{u}_0, \tilde{\eta}_0)$ simply denoting the projection π^{k_2} of the initial conditions onto the first k_2 elements of the wavelet basis. Figure 3.4.1 tabulates L^2 error in $\tilde{u}_t$ and $\tilde{a}_t$ for each $t = 0, \ldots, 7$. Subsequent figures display the true and approximated IRFs. [8]

Approximation error given the selected tuning parameters is modest, of order 10^{-4} for inputs of norm 1. Notably, the error is approximately as large for $\tilde{u}_0$ as for subsequent components in the IRF, in spite of the fact that the only source of bias in

[8] For the figures, only the real component of the approximated IRF is displayed. The presence of an imaginary component (counted in the L^2 error), should not be considered a weak point of the algorithm: due to the choice, for convenience of analysis, of complex trigonometric polynomials as eigenfunctions of L, the true solution does in fact call for state variables that are functions with range in $\mathbb{C}$, for any initial state whose Fourier coefficients are not real, which is true for any function not symmetric around $\frac{1}{2}$. While this feature is unlikely to be present in most economic applications, it would not be present given strictly real operators, and does not cause any practical difficulties.

Figure 3.4.2: True and estimated impulse responses to $\eta_0 = u_0 = \sum_{i=1}^{1024} \frac{1}{i}\phi_i / \| \sum_{i=1}^{1024} \frac{1}{i}\phi_i \|$

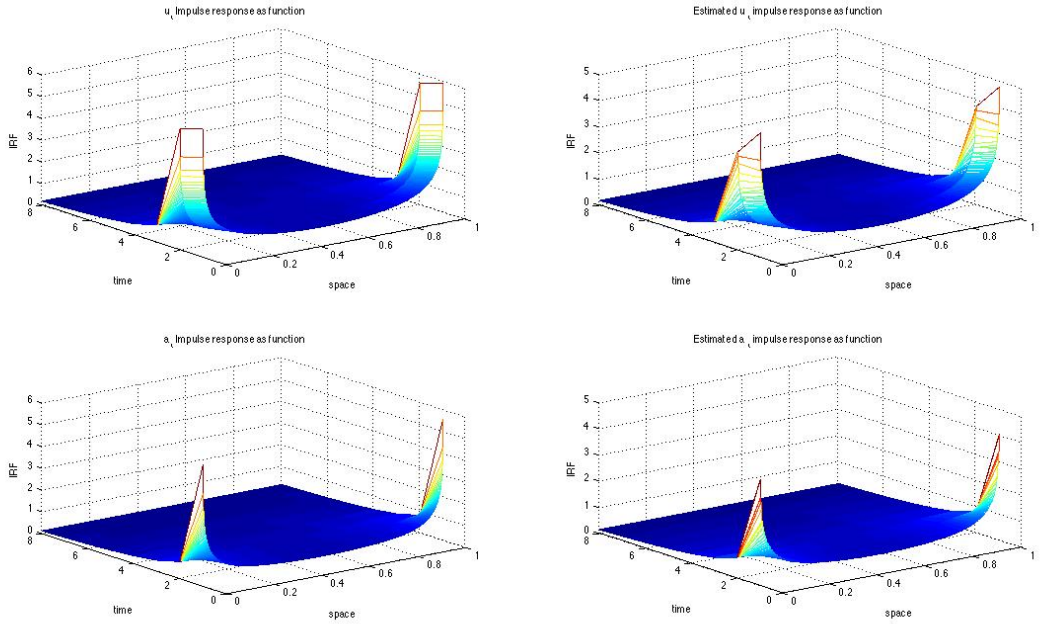

the estimate of the initial condition is that due to the basis function approximation. This indicates that the approximation error due to the subsequent components of the approximation including projection onto approximated eigenfunctions, need not substantially magnify the discretization bias, at least for sufficiently large values of tuning parameters. Repeated application of the approximated operator yields error which is well controlled, in part reflecting that the true norm declines over time for an IRF for a stationary solution and in part reflecting that $h_X^{k_4,k_5}$ yields errors which are small enough to be controlled even after passed through multiple iterations.

The reasonably good performance of the estimators $g_X^{k_4}$ and $h_X^{k_4,k_5}$ for this example suggests that it is possible to accurately approximate the solutions to a dynamic economic model with function-valued states using this method even when the solution is a highly complicated function, which requires a large number of basis functions to accurately represent and so would be difficult to parameterize in any parsimonious fashion.

Figure 3.4.3: True and estimated impulse responses to $\eta_0 = u_0 = \sum_{i=1}^{64} \phi_i / \| \sum_{i=1}^{64} \phi_i \|$

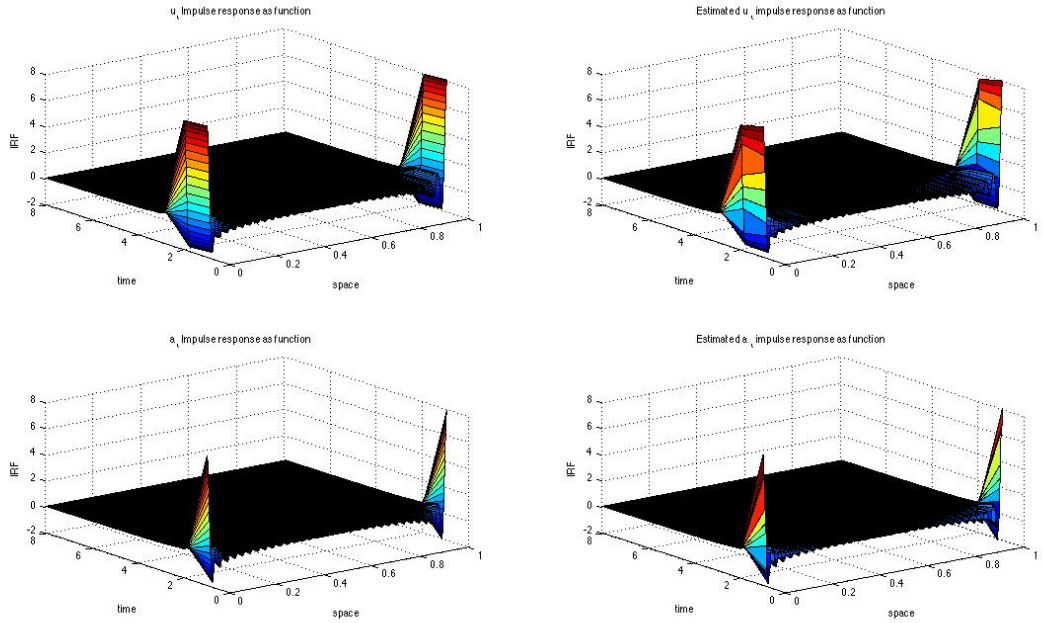

3.5 Conclusion

Many phenomena that economists want to describe are complex: they are the product of a high degree of heterogeneity across individuals, changing over time in response to and in anticipation of an uncertain environment which is high dimensional and may itself be determined endogenously by the decisions of individuals. As data on individual decisions and outcomes and over more complicated choices becomes available, models based on economic principles which directly examine the implications of this diversity become particularly desirable. At the same time, the computational challenges inherent in determining how these objects behave become increasingly daunting. While it is often possible and desirable to find simplifications which exploit structural features of certain problems which compress the decision space, either explicitly by ignoring all but a small set of features or implicitly by imposing symmetries or exclusion properties which restrict the possible combinations of interactions, such simplifications should be made carefully, with awareness of the tradeoff imposed

between tractability and bias. Ideally, such a tradeoff would be explicit and quantifiable. By constructing models which explicitly take into account the high-dimensional nature of the economic environment and providing an approximation scheme which provides consistent solutions over a set of economic states with, in theory, infinitely many dimensions of uncertainty, this procedure provides a step in that direction.

For many classes of economic problem, linearization of dynamics in function space provides a particularly tractable method for incorporating high dimensional uncertainty. By working directly in an infinite-dimensional space, it provides a construction which can be matched to arbitrary degree of accuracy to arbitrarily high dimensional data. This is particularly important for understanding changes in distributions over time, or in the case of counterfactual exercises, as a model which fits well to the current shape of a distribution may not describe the ways in which it may change in the future, especially under possibly complicated interventions which may themselves take shapes which have not previously been observed. In the case that such innovations are of interest, and are of moderate size (or the economic environment can be shown to exhibit a high degree of linearity), functional linearization provides one of the few procedures which can describe their dynamic impact. This is particularly the case in situations in which individuals anticipate and make plans based on these changes, in which case basing an estimate of the impact on a simple rule of thumb or atheoretical description of decision-making may miss the results. Fortunately, the functional linear model as a purely descriptive characterization of high-dimensional dynamic environments can also be used to approximate, in a precise sense, the behavior of forward looking agents in a fully-specified economic environment in response to such changes.

This can be achieved by extending standard procedures for determining the local behavior of dynamic, stochastic economic models with forward looking agents to the case of function-valued state variables. As functional linearization reduces exactly

to these standard procedures in the finite dimensional case, it becomes possible to extend the extremely wide array of such models already existing to incorporate heterogeneity and function-valued choice sets. This is illustrated in Section 4.1, in which a forward looking decision making problem with adjustment costs was extended to take arbitrarily high dimensional choice set and interaction between components of choices.

The main challenge in extension to infinite dimensional space is instead generalizing the decomposition made endogenously by forward looking decision makers between components which must be adjusted to achieve an outcome in the future and components which are adjusted based on the persistent influences of the past. Since many variables can interact, and in problems with function-valued states, an infinite number of components (in the sense of a basis function representation) of these functions may interact, determining how to make such a decomposition is a challenging process. Technically, the process involved is separating the linear operators which generate an implicit description of the process into a triangular form on orthogonal subspaces of the Hilbert space determined by the spectrum of the operators. I show under mild conditions that even when there is no way to reduce the space of these interactions to a fixed, manageable dimension, such a decomposition can be consistently approximated by increasing sequences of basis functions. As the unitary operators defining the decomposition are permitted to be truly high dimensional, in the sense that no finite representation can come close to describing their action uniformly over the space of inputs, and further they must be passed through a series of transformations to generate a solution, naïve representations of their behavior cannot be shown to accurately represent the behavior of the plans that are made. Intuitively, the economic reasoning behind this is that in the presence of interactions between components, restriction to a subset does not imply that approximation is accurate on that subset.

By applying the principles of generalized sampling (Adcock *et al.* , 2013) to construct accurate representations of the action of the decomposition applied at a point and applying to an increasing set of points which may then, with appropriate regularization, be passed through the transformations necessary to construct the linearized policy operators consistent with an equilibrium, one can construct finite-dimensional representations which consistently approximate the response of the endogenous variables to any given state. In the case where the true solution may be approximated uniformly well over all potential states (even though the decomposition used to generate cannot be), a simple variation of the procedure can ensure that the generated approximation converges uniformly. Given a nonstochastic steady state solution, which may be found by standard numerical methods, and a basis function representation of the derivatives of the equilibrium conditions, the algorithm to construct a consistent approximation of the linearized equilibrium response of function valued states may be implemented using standardized code over any model which fits the framework described. The order of approximations used and the choice of regularization parameters can be driven by simple diagnostics and rules of thumbs based on the type of functions being approximated. Application to a nontrivial example, with the high-dimensional adjustment cost model and the function-valued shocks fed into it parameterized in such a way as to require a particularly high order of approximation to accurately characterize solutions, shows that the algorithm can produce solutions with reasonably high accuracy even when conditions on the structure are not known.

Given the expressive power of the functional linearization method and the general applicability of the proposed equilibrium solution algorithm, they may be able to provide insights for a wide variety of high-dimensional dynamic economic phenomena heretofore difficult to analyze. These include models with nonlinear and non-smooth relationships between states, high-dimensional non-smooth dynamic decision problems, and other models with high-dimensional interactions. In each of these cases,

numerical characterization of the dynamics which takes into account the nontrivial impact on system behavior of high-dimensional interactions will require careful consideration of the propagation of errors and appropriate regularization to ensure that this behavior is not misrepresented by inappropriately ignoring these function-valued nature of the problem. Going forward, uses of these methods in realistic economic models, as well as assessment of the extent to which the phenomena uncovered here exhibit empirical relevance, are important future avenues to pursue.

Appendix A: Results on the Generalized Schur Decomposition

Existence of a Generalized Schur Decomposition for Pairs of Bounded Operators

The construction of a solution for the linear expectational difference equation defined by a linear or linearized rational expectations model in finite dimensions relies on the ability to partition the state space and the equilibrium equations into 'stable' and 'unstable' components which may be treated separately. This is generally achieved by either a Jordan decomposition, generating block-diagonal matrices, as in Blanchard & Kahn (1980) or by a generalized Schur decomposition, generating upper-triangular matrices,[9] as in Klein (2000). In practice, the latter has become preferred, as the Jordan decomposition of a matrix is not in general continuous while the generalized Schur decomposition, which is generated by unitary matrices, exhibits numerical stability in theory and practice. Such stability is particularly desirable in the infinite-dimensional case, as closed form solutions for the eigenfunctions are not in general feasible and finite-dimensional numerical procedures must by necessity induce some error into the representation of the operator pair of interest.

[9] This decomposition is often referred to as the QZ decomposition, in reference to the QZ algorithm often used to compute it. See Golub & van Loan (1996).

While generalization of the Jordan decomposition to infinite-dimensional operator pairs is well established (Kato, 1976; Gohberg *et al.* , 1990, Ch IV) and the Schur decomposition for a single infinite-dimensional operator has also been defined (Gohberg *et al.* , 1990, Ch II.3), an analogue of the generalized Schur decomposition for pairs of infinite-dimensional linear operators has not, to the best of my knowledge, been described. As in the case of the Schur decomposition of a single operator, extension to the infinite-dimensional case is slightly delicate, as the existence of the Schur or generalized Schur decomposition is based on an iterative construction which extends only in certain cases to an uncountable state space. In particular, the Schur decomposition may be extended to compact operators but not to arbitrary bounded or closed operators, for which a Jordan decomposition exists but a Schur decomposition may not. For the purposes of constructing an analogy of the generalized Schur decomposition which permits extension of rational expectation solution procedures, there are at least two ways around this difficulty. The first, and simplest, is to note that while solution requires splitting the domain into 'forward' and 'backward' subspaces, for a stationary solution there is no requirement that the restriction of the operator to these subspaces itself take upper triangular form. Instead, one can construct a block upper triangular decomposition which preserves the desirable feature of being generated by unitary transformation while eschewing the necessity to make restrictive compactness assumptions. Alternately, one may construct a generalized Schur decomposition analogously to the infinite-dimensional Schur decomposition, which does preserve an upper-triangular structure within blocks, under a modified and so slightly less onerous compactness condition than in the single operator case. In the following, I show existence of a blockwise decomposition under general conditions, and also decomposition which is upper triangular within blocks under a condition on compactness of certain transformations of the operator pair which does not imply that both operators are compact, and in particular allows the pertinent example of the

standard eigenvalue problem in which one of the operators in the pair is the identity operator, which is not compact on an infinite-dimensional space. This construction also has the advantage that it implies compactness of certain Schur components and so generates a solution for the law of motion which is itself compact.

Formally, let (M, G) be a pair of bounded linear operators acting between complex Hilbert spaces $\mathcal{H}_X$ and $\mathcal{H}_Y$, i.e. $M \in \mathcal{L}(\mathcal{H}_X \to \mathcal{H}_Y)$ $G \in \mathcal{L}(\mathcal{H}_X \to \mathcal{H}_Y)$. Following Gohberg *et al.* (1990), define the spectrum $\sigma(M, G)$ as those $\lambda \in \mathbb{C}$ such that $\lambda G - M$ is not invertible, accompanied by the point ∞ if and only if G does not have bounded inverse, and the resolvent set $\rho(M, G)$ as $\mathbb{C}_\infty \backslash \sigma(M, G)$, where $\mathbb{C}_\infty$ is the extended complex plane with the standard topology (see Conway (1978, Ch. 1 S. 6)).

Definition .1. An operator pair is said to be Γ-regular (with respect to a set Γ) if for some nonempty subset $\Gamma \subset \mathbb{C}_\infty$, $\Gamma \subset \rho(M, G)$.

Assume Γ is a Cauchy contour (c.f. Gohberg *et al.* (1990, p.6)) with inner domain Δ_+ and outer domain Δ_-, and that (M, G) is Γ-regular. For concreteness, we will often take Γ to be the positively oriented complex unit circle, in which case Γ-regularity means that the spectrum does not contain λ such that $|\lambda| = 1$. From a modeling perspective, this ensures stationarity by ruling out unit roots; this particular choice is not required to ensure existence of a generalized Schur decomposition. By Gohberg *et al.* (1990) Theorem IV.1.1, the above assumptions ensure the existence of (possibly oblique) projection operators π_1: $\mathcal{H}_X \to \mathcal{H}_X$ and π_2: $\mathcal{H}_Y \to \mathcal{H}_Y$ which partition $\mathcal{H}_X$ and $\mathcal{H}_Y$ into $\text{Im } \pi_1 \oplus \text{Ker } \pi_1$ and $\text{Im } \pi_2 \oplus \text{Ker } \pi_2$ respectively, and the operator pair (M, G) into components

$$(M, G) = \left(\begin{bmatrix} M_1 & 0 \\ 0 & M_2 \end{bmatrix} \begin{bmatrix} G_1 & 0 \\ 0 & G_2 \end{bmatrix} \right) : \text{Im } \pi_1 \oplus \text{Ker } \pi_1 \to \text{Im } \pi_2 \oplus \text{Ker } \pi_2 \quad (2)$$

such that (M_1, G_1) and (M_2, G_2) are Γ-regular, $\sigma(M_1, G_1) = \sigma(M, G) \cap \Delta_+$ and

$\sigma(M_2, G_2) = \sigma(M, G) \cap \Delta_-$. In words, this says one can separate the pair into a component with spectrum inside some domain and a component with spectrum outside.

Assume in addition that $0 \in \Delta_+$ and $\infty \in \Delta_-$. By the above result and the definition of the resolvent, this implies that G_1 and M_2 are invertible on their respective domains. In particular, $E = \begin{pmatrix} G_1^{-1} & 0 \\ 0 & M_2^{-1} \end{pmatrix}$: Im $\pi_2 \oplus$ Ker $\pi_2 \to$ Im $\pi_1 \oplus$ Ker π_1 is a bounded invertible operator and we may define the partition

$$(EM, EG) = \left(\begin{bmatrix} \Omega_1 & 0 \\ 0 & I_2 \end{bmatrix} \begin{bmatrix} I_1 & 0 \\ 0 & \Omega_2 \end{bmatrix} \right) : \text{ Im } \pi_1 \oplus \text{Ker } \pi_1 \to \text{Im } \pi_1 \oplus \text{Ker } \pi_1 \quad (3)$$

where $\Omega_1 = G_1^{-1} M_1$ and $\Omega_2 = M_2^{-1} G_2$. These operators have the following relationship with (M_1, G_1) and (M_2, G_2):

Lemma .8. *$\sigma(M_1, G_1) = \sigma(\Omega_1)$, and $\frac{1}{\lambda} \in \sigma(\Omega_2)$ if and only if $\lambda \in \sigma(M_2, G_2)$ (where $\frac{1}{\infty}$ may be defined to equal 0)*

Proof. Suppose λ is in the resolvent set of Ω_1. Then $\Omega_1 - \lambda I_1$ has some bounded inverse Z. Then $-ZG_1^{-1}$ satisfies $-ZG_1^{-1}(\lambda G_1 - M_1) = Z(\Omega_1 - \lambda I_1) = I_1$ and $-(\lambda G_1 - M_1)ZG_1^{-1} = -G_1 G_1^{-1}(\lambda G_1 - M_1)ZG_1^{-1} = G_1(\Omega - \lambda I_1)ZG_1^{-1} = G_1 G_1^{-1} = I_1$, so $\lambda \in \rho(M_1, G_1)$. That is, $\rho(\Omega_1) \subset \rho(M_1, G_1)$. Next, suppose $\lambda \in \rho(M_1, G_1)$. Then $\lambda G_1 - M_1$ has a bounded inverse Z, and $-ZG_1$ satisfies $-ZG_1(\Omega - \lambda I_1) = Z(\lambda G_1 - M_1) = I_1$ and $-(\Omega - \lambda I_1)ZG_1 = -G_1^{-1} G_1(\Omega - \lambda I_1)ZG_1 = G_1^{-1}(\lambda G_1 - M_1)ZG_1 = G_1^{-1} G_1 = I_1$, and so $\rho(M_1, G_1) \subset \rho(\Omega_1)$. Combining, $\rho(M_1, G_1) = \rho(\Omega_1)$ and so $\sigma(M_1, G_1) = \sigma(\Omega_1)$. Similar calculations show $\frac{1}{\lambda} \in \sigma(\Omega_2)$ if and only if $\lambda \in \sigma(M_2, G_2)$. If $\infty \in \sigma(M_2, G_2)$, G_2 is not invertible and so $M_2^{-1} G_2 - \frac{1}{\infty} I_2 = M_2^{-1} G_2$ must also have nontrivial kernel, and so be noninvertible. □

With this notation, it is possible to characterize conditions under which the oper-

ator pair (M, G) has a generalized Schur decomposition. As our construction makes use of complete orthonormal bases, we assume now that (M, G) are operators between *separable* Hilbert spaces $\mathcal{H}_X$ and $\mathcal{H}_Y$.

Lemma .9. *Let (M, G) be a pair of bounded operators $M \in \mathcal{L}(\mathcal{H}_X \to \mathcal{H}_Y)$ $G \in \mathcal{L}(\mathcal{H}_X \to \mathcal{H}_Y)$ Γ-regular with respect to a Cauchy curve with inner domain Δ_+* such that $0 \in \Delta_+$ *and outer domain Δ_- such that $\infty \in \Delta_-$. Define projectors π_1 and π_2 as in 2 with respect to Γ. Then, there exist unitary operators $Q = [Q^1, Q^2] : Im\pi_2 \oplus \mathcal{H}_Y / Im\pi_2 \to F_1 \oplus F_2$ and $P = [P^1, P^2] : Im\pi_1 \oplus \mathcal{H}_X / Im\pi_1 \to E_1 \oplus E_2$ such that (M, G) has the following block-wise generalized Schur decomposition*

$$(QMU^*, QGU^*) = \left(\begin{bmatrix} M_{11} & M_{12} \\ 0 & M_{22} \end{bmatrix}, \begin{bmatrix} G_{11} & G_{12} \\ 0 & G_{22} \end{bmatrix} \right)$$

from $E_1 \oplus E_2 \to F_1 \oplus F_2$

where E_1, E_2, F_1, and F_2 are spaces such that there exist linear isometric isomorphisms from $\mathrm{Im}\pi_1 \to E_1$, $\mathcal{H}_X/\mathrm{Im}\pi_1 \to E_2$, $\mathrm{Im}\pi_2 \to F_1$, and $\mathcal{H}_X/\mathrm{Im}\pi_2 \to F_2$, *respectively. Further, $\sigma(M_{11}, G_{11}) = \sigma(M_1, G_1) = \sigma(M, G) \cap \Delta_+$ and $\sigma(M_{22}, G_{22}) = \sigma(M_2, G_2) = \sigma(M, G) \cap \Delta_-$.*

Remark. The precise identity of the spaces E_1, E_2, F_1, and F_2 need not be considered for this result. However, a canonical choice of spaces would be to allow $E_1 = \mathrm{Im}\pi_1$, $E_2 = \mathcal{H}_X/\mathrm{Im}\pi_1$, $F_1 = \mathrm{Im}\pi_2$, $F_2 = \mathcal{H}_Y/\mathrm{Im}\pi_2$, in which case the Schur decomposition acts on the same space as (M, G).

Proof. We generate Q and U constructively, then verify their properties. Choose a complete orthonormal basis on $\mathrm{Im}\pi_1$, denoted $\{u_{1i}\}_{i=1}^{\infty}$ and then a complete orthonormal basis on the orthogonal complement of $\mathrm{Im}\pi_1$ in $\mathcal{H}_1$, denoted $\{u_{2i}\}_{i=1}^{\infty}$. The

eigenvectors are not, in general, such a basis, because Ω_1 and Ω_2 are not assumed self-adjoint and so nothing requires their eigenvectors to be orthogonal vectors. Then, U_1 is the operator $\sum_{i=1}^{\infty} \langle u_{1i}, . \rangle \, e_i^1$ where $\{e_i^1\}_{i=1}^{\infty}$ are an arbitrary orthonormal basis on E_1, a space isometrically isomorphic to $\mathrm{Im}\pi_1$, U_2 is the operator $\sum_{i=1}^{\infty} \langle u_{2i}, . \rangle \, e_i^2$ where $\{e_i^2\}_{i=1}^{\infty}$ are an arbitrary orthonormal basis on E_2, a space isometrically isomorphic to $\mathcal{H}_1/\mathrm{Im}\pi_1$. Likewise, choose a complete orthonormal basis $\{q_{1i}\}_{i=1}^{\infty}$ for the image of (M_1, G_1) and a complete orthonormal basis for the orthogonal complement of this space in $\mathcal{H}_Y$, $\{q_{2i}\}_{i=1}^{\infty}$. We define $Q_1 = \sum_{i=1}^{\infty} \langle q_{1i}, . \rangle \, f_i^1$ and $Q_2 = \sum_{i=1}^{\infty} \langle q_{2i}, . \rangle \, f_i^2$, for $\{f_i^1\}_{i=1}^{\infty}$ and $\{f_i^2\}_{i=1}^{\infty}$ orthonormal bases of F_1 and F_2, spaces isometrically isomorphic to the domains of Q_1 and Q_2 respectively.

Next, we show that these induce an upper triangular decomposition. We define

$$\left(\begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix}, \begin{bmatrix} G_{11} & G_{12} \\ G_{21} & G_{22} \end{bmatrix} \right) = \left(\begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} M \begin{bmatrix} U_1^* & U_2^* \end{bmatrix}, \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} G \begin{bmatrix} U_1^* & U_2^* \end{bmatrix} \right)$$

Using 2, we have that $(BU_1^*, AU_1^*) = (B_1U_1^*, A_1U_1^*)$ since the range of U_1^* is $\mathrm{Im}\pi_1$, and the restriction of (M, G) to this space is (M_1, G_1). Then, since the domain of Q_2 is orthogonal to $\mathrm{Im}(M_1, G_1)$, we have $(M_{21}, G_{21}) = (0, 0)$, so this is a triangular decomposition.

To characterize the spectrum of the decomposition, first note that $\sigma(M_1, G_1) = \sigma(M, G) \cap \Delta_+$ and $\sigma(M_2, G_2) = \sigma(M, G) \cap \Delta_-$ by Gohberg *et al.* (1990) Theorem IV.1.1. (M_{11}, G_{11}) may be written as $(Q_1MU_1^*, Q_1GU_1^*) = (Q_1M_1U_1^*, Q_1G_1U_1^*)$. Consider $\gamma \in \rho(M_1, G_1)$. Then $\gamma M_{11} - G_{11} = \gamma Q_1 M_1 U_1^* - Q_1 G_1 U_1^* = Q_1(\gamma M_1 - G_1)U_1^*$, which has inverse $U_1(\gamma M_1 - G_1)^{-1}Q_1^*$ which is bounded since $(\gamma M_1 - G_1)^{-1}$ is bounded, by definition of the resolvent set, and U_1 and Q_1^* are since they are unitary by construction. So $\sigma(M_{11}, G_{11}) \subset \sigma(M_1, G_1) = \sigma(M, G) \cap \Delta_+$.

Characterization of the spectrum of (M_{22}, G_{22}) requires a bit more care. (M_{22}, G_{22})

may be written as

$$\begin{aligned}(Q_2MU_2^*, Q_2GU_2^*) &= (Q_2M(\pi_1 + (I-\pi_1))U_2^*, Q_2G(\pi_1 + (I-\pi_1))U_2^*) \\ &= (Q_2M_1\pi_1U_2^*, Q_2G_1\pi_1U_2^*) + (Q_2M_2(I-\pi_1)U_2^*, Q_2G_2(I-\pi_1)U_2^*) \\ &= (Q_2M_2(I-\pi_1)U_2^*, Q_2G_2(I-\pi_1)U_2^*)\end{aligned}$$

where the second line follows from 2 and the final line follows from the fact that the domain of Q_2 is orthogonal to the range of (M_1, G_1). Consider $\gamma \in \rho(M_2, G_2)$. By definition of the resolvent, $T(\gamma) := (\gamma M_2 - G_2)^{-1}$ is a bounded operator for all such γ. Then $\gamma M_{22} - G_{22} = Q_2(\gamma M_2 - G_2)(I-\pi_1)U_2^*$. I claim that $U_2T(\gamma)Q_2^*$ is a bounded inverse of $\gamma M_{22} - G_{22}$. To see this, note that $Q_2^*Q_2$ is equal to $I_{\mathcal{H}_Y/\mathrm{Im}\pi_2}$ and $U_2^*U_2 = I_{\mathcal{H}_X/\mathrm{Im}\pi_1}$. As a result, we have $U_2T(\gamma)Q_2^*Q_2(\gamma M_2 - G_2)(I-\pi_1)U_2^* = U_2(I-\pi_1)U_2^* = U_2U_2^* = I_{\mathcal{H}_X/\mathrm{Im}\pi_1}$, where we use the fact that $U_2\pi_1 = 0$ since U_2 has domain orthogonal to the image of π_1. By the identity $(M_2, G_2)(I-\pi_1) = (I-\pi_2)(M_2, G_2)$, $Q_2(\gamma M_2 - G_2)(I-\pi_1)U_2^*U_2T(\gamma)Q_2^* = Q_2(I-\pi_2)(\gamma M_2 - G_2)U_2^*U_2T(\gamma)Q_2^* = Q_2(I-\pi_2)Q_2^* = I_{\mathcal{H}_Y/\mathrm{Im}\pi_2}$, since $Q_2\pi_2 = 0$. As a result, $\sigma(M_{22}, G_{22}) \subset \sigma(M_2, G_2) = \sigma(M, G) \cap \Delta_-$.

To show the reverse inclusion, note that $\sigma(M, G) = \sigma(QMU^*, QGU^*)$ by unitarity of Q and U. Next, we show that $\sigma(QMU^*, QGU^*) = \sigma(M_{11}, G_{11}) \cup \sigma(M_{22}, G_{22})$. Since Δ_+ and Δ_- are disjoint, $\sigma(M_{11}, G_{11}) \subset \sigma(M, G) \cap \Delta_+$, and $\sigma(M_{22}, G_{22}) \subset \sigma(M, G) \cap \Delta_-$, this implies that $\sigma(M_{11}, G_{11}) = \sigma(M, G) \cap \Delta_+$ and $\sigma(M_{22}, G_{22}) = \sigma(M, G) \cap \Delta_-$, as claimed. To show this, consider $\gamma \in \rho(M_{11}, G_{11}) \cap \rho(M_{22}, G_{22})$. Then $\gamma QMU^* - QGU^* = \begin{bmatrix} \gamma M_{11} - G_{11} & \gamma M_{12} - G_{12} \\ 0 & \gamma M_{22} - G_{22} \end{bmatrix}$ has bounded inverse given by

$$\begin{bmatrix} (\gamma M_{11} - G_{11})^{-1} & -(\gamma M_{11} - G_{11})^{-1}(\gamma M_{12} - G_{12})(\gamma M_{22} - G_{22})^{-1} \\ 0 & (\gamma M_{22} - G_{22})^{-1} \end{bmatrix}$$

and so $\sigma(QMU^*, QGU^*) \subset \sigma(M_{11}, G_{11}) \cup \sigma(M_{22}, G_{22})$. Next, suppose $\gamma \in \sigma(M_{11}, G_{11})$

and assume for contradiction that $\gamma \in \rho(QMU^*, QGU^*)$, and so $\gamma QMU^* - QGU^*$ has some bounded inverse $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$. Then $\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} \gamma M_{11} - G_{11} & \gamma M_{12} - G_{12} \\ 0 & \gamma M_{22} - G_{22} \end{bmatrix} = \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix}$ and so $a(\gamma M_{11} - G_{11}) = I$, implying that $\gamma M_{11} - G_{11}$ has bounded inverse a, a contradiction. Similarly, if $\gamma \in \sigma(M_{22}, G_{22})$, then if $\gamma QMU^* - QGU^*$ had some bounded inverse $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ then $\begin{bmatrix} \gamma M_{11} - G_{11} & \gamma M_{12} - G_{12} \\ 0 & \gamma M_{22} - G_{22} \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ would equal $\begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix}$, implying $(\gamma M_{22} - G_{22})d = I$, which is assumed false. As a result, $\sigma(QMU^*, QGU^*) \supset \sigma(M_{11}, G_{11}) \cup \sigma(M_{22}, G_{22})$, and the claim is shown. □

Slightly stronger assumptions than used in the above can yield stronger results. In particular, the assumption of compactness of Ω_1 and Ω_2 may permit the block triangular decomposition to be extended to a triangular decomposition within each block, as in the infinite-dimensional Schur decomposition in Gohberg *et al.* (1990). This provides a link to the finite-dimensional method, but is nowhere necessary for the application of the decomposition considered. However, compactness of the components does provide a useful sufficient condition for the necessary conditions, and also ensures the compactness of the solution operators, which is a condition commonly imposed for the validity of estimators of infinite-dimensional operators: see, e.g., Bosq (2000).

If this refinement is not needed, we may instead operate under a strictly weaker assumption: viz. that the spectrum of (M, G) is bounded away from Γ. To see that this is weaker, note that compactness implies that the unique accumulation point of the spectrum is at 0, and so by .8, the spectra of Ω_1 and Ω_2 and, as a result, of (M, G) must neither be inside of Γ or have limit point in Γ. Formally, we define a block triangular decomposition as follows. For notational convenience and analogy

to the finite-dimensional case, we take the decomposition to be defined as a pair on $\mathcal{L}(\mathcal{H}_X \to \mathcal{H}_Y)$ rather than over isometrically isomorphic spaces.

Lemma .10. *Let (M, G) be a pair of bounded operators $M \in \mathcal{L}(\mathcal{H}_X \to \mathcal{H}_Y)$ $G \in \mathcal{L}(\mathcal{H}_X \to \mathcal{H}_Y)$ Γ-regular with respect to a Cauchy curve with inner domain Δ_+* such that $0 \in \Delta_+$ *and outer domain Δ_- such that $\infty \in \Delta_-$. Define projectors π_1 and π_2 as in 2 with respect to Γ and Ω_1 and Ω_2 as in 3. Suppose in addition that Ω_1 and Ω_2 are compact operators. Then, there exist unitary operators $Q = [Q^1, Q^{1\perp}, Q^2, Q^{2\perp}] : F_1 \oplus F_1^\perp \oplus F_2 \oplus F_2^\perp \to F_1 \oplus F_1^\perp \oplus F_2 \oplus F_2^\perp$ and $P = [P^1, P^{1\perp}, P^2, P^{2\perp}] : E_1 \oplus E_1^\perp \oplus E_2 \oplus E_2^\perp \to E_1 \oplus E_1^\perp \oplus E_2 \oplus E_2^\perp$ such that (M, G) has the following (generalized Schur) decomposition*

$$(M,G) = \left(\begin{bmatrix} M_{11} & M_{11}^{off} & M_{12} & . \\ 0 & M_{11}^\perp & . & . \\ 0 & 0 & M_{22} & M_{22}^{off} \\ 0 & 0 & 0 & M_{22}^\perp \end{bmatrix}, \begin{bmatrix} G_{11} & G_{11}^{off} & G_{12} & . \\ 0 & G_{11}^\perp & . & . \\ 0 & 0 & G_{22} & G_{22}^{off} \\ 0 & 0 & 0 & G_{22}^\perp \end{bmatrix} \right)$$

from $E_1 \oplus E_1^\perp \oplus E_2 \oplus E_2^\perp \to F_1 \oplus F_1^\perp \oplus F_2 \oplus F_2^\perp$

where $E_1, E_1^\perp, E_2, E_2^\perp$ and $F_1, F_1^\perp, F_2, F_2^\perp$ are closed linear subspaces of $\mathcal{H}_X$ and $\mathcal{H}_Y$, respectively. Further, with respect to the orthonormal bases $\{\tilde{p}_i^1\}_{i=1}^\infty$ of E_1 and $\{\tilde{q}_i^1\}_{i=1}^\infty$ of F_1 generating the rows of P^1 and Q^1, respectively, (M_{11}, G_{11}) are upper triangular with $(M_{11})_{jj}/(G_{11})_{jj} = \lambda_j$ where λ_j is the j^{th} nonzero generalized eigenvalue (in some arbitrary fixed order) repeated a number of times equal to its multiplicity in $\sigma(M_1, G_1)$, and similarly with respect to the orthonormal bases $\{\tilde{p}_i^2\}_{i=1}^\infty$ of E_2 and $\{\tilde{q}_i^2\}_{i=1}^\infty$ of F_2 generating the rows of P^2 and Q^2, respectively, (M_{22}, G_{22}) are upper triangular with $(M_{22})_{jj}/(G_{22})_{jj} = \lambda_j$ where λ_j is the j^{th} finite generalized eigenvalue repeated a number of times equal to its multiplicity in $\sigma(M_2, G_2)$. In addition, $\sigma(M_{11}^\perp, G_{11}^\perp) \subset \{0\}$

and $\sigma(M_{22}^{\perp}, G_{22}^{\perp}) \subset \{\infty\}$.

Remark. $(G_{11}^{\perp})^{-1}M_{11}^{\perp}$ and $(M_{22}^{\perp})^{-1}G_{22}^{\perp}$ are examples of Volterra operators, as they are compact and quasinilpotent (with spectrum equal to zero only). As a result, they may be shown to be unitarily equivalent to a particular continuous analogue of an upper-triangular operator with respect to a (not necessarily countable) increasing chain of projections on subspaces of $\mathcal{H}_X$ (Gohberg *et al.* , 1993, Thm. XXI.1.5). In principle, a fully triangular representation of (M, G) in which $(M_{11}^{\perp}, G_{11}^{\perp})$ and $(M_{22}^{\perp}, G_{22}^{\perp})$ are also upper-triangular with respect to some chain of subspaces could be generated via an analogue for operator pairs of Gohberg *et al.* (1993, Thm. XXI.1.2). Such a decomposition is unnecessary for our purposes, as block-triangular structure is sufficient for representing a solution of the equilibrium conditions and the approximation techniques to be used do not take advantage of the continuous structure provided by the more intricate decomposition.

Proof. Begin by noting that if $P_{\{\lambda_i\}}$ is a projector onto an eigenspace of Ω_1 corresponding to nonzero eigenvalue λ_i (sorted in arbitrary but fixed order), it is also a projector onto an eigenspace of (M_1, G_1) corresponding to the same eigenvalue. By compactness, any nonzero element of the spectrum of Ω_1 is isolated and an eigenvalue, and by equality of spectra corresponds to an isolated point in the spectrum $\sigma(M_1, G_1)$. As a result, one may write the projector onto the eigenspace associated with λ_i of Ω_1 as $P_{\{\lambda_i\}}^{\Omega_1} := \frac{1}{2\pi\iota}\int_{\Gamma_{\lambda_i}}(\zeta I_1 - \Omega_1)^{-1}d\zeta$, where Γ_{λ_i} is a closed Cauchy curve enclosing λ_i, and the projector onto the space associated with element λ_i of the spectrum of pair (M_1, G_1) as $P_{\{\lambda_i\}}^{(M_1,G_1)} := \frac{1}{2\pi\iota}\int_{\Gamma_{\lambda_i}}(\zeta G_1 - M_1)^{-1}G_1 d\zeta$ (See Gohberg *et al.* (1990, Ch.

I.2 and IV.1)). Since $\Omega_1 = G_1^{-1}M_1$,

$$P_{\{\lambda_i\}}^{(M_1,G_1)} = \frac{1}{2\pi\iota}\int_{\Gamma_{\lambda_i}} (\zeta G_1 - M_1)^{-1}G_1 d\zeta = \frac{1}{2\pi\iota}\int_{\Gamma_{\lambda_i}} ((G_1G_1^{-1})(\zeta G_1 - M_1))^{-1}G_1 d\zeta \quad (4)$$
$$= \frac{1}{2\pi\iota}\int_{\Gamma_{\lambda_i}} ((G_1)(\zeta I_1 - \Omega_1))^{-1}G_1 d\zeta$$
$$= \frac{1}{2\pi\iota}\int_{\Gamma_{\lambda_i}} (\zeta I_1 - \Omega_1)^{-1}G_1^{-1}G_1 d\zeta$$
$$= P_{\{\lambda_i\}}^{\Omega_1}$$

Compactness also guarantees that the dimension of the image of $P_{\{\lambda_i\}}$ is finite (Gohberg *et al.* , 1990, Thm II.3.2), and so by equality of spectra, the subspaces associated with points not equal to zero in the spectrum of (M_1, G_1) are all finite-dimensional. As a result, we may choose for each i, a finite set of basis vectors, of cardinality k_i, for the space $\mathrm{Im}P_{\{\lambda_i\}}$ and a basis for the image of the pair $(MP_{\{\lambda_i\}}, GP_{\{\lambda_i\}})$ which must be of dimension k_i as $GP_{\{\lambda_i\}}$ must be of full rank since G_1 is. In particular, as on this space the operator pair has a representation as a pair of $k_i \times k_i$-dimensional matrices, we may without loss of generality use orthonormal basis vectors $\{q_{i1}^1, \ldots, q_{ik_i}^1\}$ for the image of $(MP_{\{\lambda_i\}}, GP_{\{\lambda_i\}})$ and $\{p_{i1}^1, \ldots, p_{ik_i}^1\}$ for $\mathrm{Im}P_{\{\lambda_i\}}$ such that with respect to these bases, M and G are upper triangular with diagonal elements of M and G identically equal to σ_i and τ_i, respectively, where $\frac{\sigma_i}{\tau_i} = \lambda_i$. Such a representation exists by the generalized Schur decomposition for finite-dimensional matrix pairs (Stewart & Sun, 1990, Th. VI.1.9). Note that while these basis vectors are orthogonal within each block, in general $\mathrm{Im}P_{\{\lambda_i\}}$ is not necessarily orthogonal to $\mathrm{Im}P_{\{\lambda_j\}}$ for $i \neq j$ as these are oblique, not orthogonal projections.

For Ker π_1, compactness of Ω_2 permits an analogous construction of a countable sequence of finite-dimensional eigenprojections associated to isolated points of the spectrum, with the difference that the projection onto the space associated with point λ_i in the spectrum $\sigma(M_2, G_2)$ is equal to projection associated with nonzero

eigenvalue $\frac{1}{\lambda_i} \in \sigma(\Omega_2)$. That is to say, in the notation above, $P^{(M_2,G_2)}_{\{\lambda_i\}} = P^{\Omega_2}_{\{\frac{1}{\lambda_i}\}}$ for $\lambda_i \neq \infty$. Since Im $P^{\Omega_2}_{\{\frac{1}{\lambda_i}\}}$ is a finite-dimensional subspace of dimension k_i, and since $MP^{(M_2,G_2)}_{\{\lambda_i\}}$ is full rank since M_2 is, we may define sets of orthonormal basis vectors $\{q^2_{i1}, \ldots, q^2_{ik_i}\}$ on the image of $(MP^{(M_2,G_2)}_{\{\lambda_i\}}, GP^{(M_2,G_2)}_{\{\lambda_i\}})$ and $\{p^2_{i1}, \ldots, p^2_{ik_i}\}$ on $\mathrm{Im}P^{(M_2,G_2)}_{\{\lambda_i\}}$ such that with respect to these basis vectors, (M, G) has a representation as a pair of $k_i \times k_i$ upper-triangular Schur matrices with diagonal elements identically equal to the corresponding eigenvalue pair $\{\sigma_i, \tau_i\}$ where $\frac{\sigma_i}{\tau_i} = \lambda_i$.

For the space Im $\pi_1 \backslash \overline{\mathrm{Span}}\{p^1_{11}, \ldots, p^1_{1k_1}, p^1_{21,\ldots}\}$, choose choose an arbitrary complete orthonormal basis, say $\{p^{1\perp}_1, p^{1\perp}_2, \ldots\}$ and for Im $\pi_2 \backslash \overline{\mathrm{Span}}\{q^1_{11}, \ldots, q^1_{1k_1}, q^1_{21,\ldots}\}$, choose a basis $\{q^{1\perp}_1, q^{1\perp}_2, \ldots\}$. Likewise, for the space Ker $\pi_1 \backslash \overline{\mathrm{Span}}\{p^2_{11}, \ldots, p^2_{1k_1}, p^2_{21,\ldots}\}$, choose choose an arbitrary complete orthonormal basis, say $\{p^{2\perp}_1, p^{2\perp}_2, \ldots\}$ and for Ker $\pi_2 \backslash \overline{\mathrm{Span}}\{q^2_{11}, \ldots, q^2_{1k_1}, q^2_{21,\ldots}\}$ choose a basis $\{q^{2\perp}_1, q^{2\perp}_2, \ldots\}$. These bases may in general be infinite dimensional and are not necessarily orthogonal to the bases defined for other spaces. To produce the stated decomposition, these bases will be used to construct an orthogonal basis with the desired properties.

To produce the desired decomposition, order the sets of vectors as $(\{p^1_{11}, \ldots, p^1_{1k_1}\}$, $\{p^1_{21}, \ldots, p^1_{2k_2}\}, \ldots, \{p^{1\perp}_1, p^{1\perp}_2, \ldots\}, \ldots, \{p^2_{11}, \ldots, p^2_{1k_1}\}, \{p^2_{21}, \ldots, p^2_{2k_2}\}, \ldots, \{p^{2\perp}_1, p^{2\perp}_2, \ldots\}, \ldots)$ and $(\{q^1_{11}, \ldots, q^1_{1k_1}\}, \{q^1_{21}, \ldots, q^1_{2k_2}\}, \ldots, \{q^{1\perp}_1, q^{1\perp}_2, \ldots\}, \ldots, \{q^2_{11}, \ldots, q^2_{1k_1}\}, \{q^2_{21}, \ldots, q^2_{2k_2}\}$, $\ldots, \{q^{2\perp}_1, q^{2\perp}_2, \ldots\}, \ldots)$ and apply Gram-Schmidt orthonormalization to the countable sequences to produce a pair of orthonormal bases $\{\tilde{p}^1_{11}, \ldots, \tilde{p}^1_{1k_1}, \tilde{p}^1_{21,} \ldots, \tilde{p}^{1\perp}_1, \ldots,$ $\tilde{p}^2_{11}, \ldots, \tilde{p}^2_{1k_1}, \tilde{p}^2_{21,} \ldots, \tilde{p}^{2\perp}_1, \ldots\}$ and $\{\tilde{q}^1_{11}, \ldots, \tilde{q}^1_{1k_1}, \tilde{q}^1_{21,} \ldots, \tilde{q}^{1\perp}_1, \ldots, \tilde{q}^2_{11}, \ldots, \tilde{q}^2_{1k_1}, \tilde{q}^2_{21,} \ldots, \tilde{q}^{2\perp}_1, \ldots\}$ of $\mathcal{H}_X$ and $\mathcal{H}_Y$ respectively. We may then define $E_1 = \overline{\mathrm{Span}}\{\tilde{p}^1_{11}, \ldots, \tilde{p}^1_{1k_1}, \tilde{p}^1_{21,} \ldots\}$, $F_1 = \overline{\mathrm{Span}}\{\tilde{q}^1_{11}, \ldots, \tilde{q}^1_{1k_1}, \tilde{q}^1_{21,} \ldots\}$, $E^\perp_1 = \overline{\mathrm{Span}}\{\tilde{p}^{1\perp}_1, \tilde{p}^{1\perp}_2, \ldots\}$, $F^\perp_1 = \overline{\mathrm{Span}}\{\tilde{q}^{1\perp}_1, \tilde{q}^{1\perp}_2, \ldots\}$, $E_2 = \overline{\mathrm{Span}}\{\tilde{p}^2_{11}, \ldots, \tilde{p}^2_{1k_1}, \tilde{p}^2_{21,} \ldots\}$, $F_2 = \overline{\mathrm{Span}}\{\tilde{q}^2_{11}, \ldots, \tilde{q}^2_{1k_1}, \tilde{q}^2_{21,} \ldots\}$ $E^\perp_2 = \overline{\mathrm{Span}}\{\tilde{p}^{2\perp}_1, \tilde{p}^{2\perp}_2, \ldots\}$, and $F^\perp_2 = \overline{\mathrm{Span}}\{\tilde{q}^{2\perp}_1, \tilde{q}^{2\perp}_2, \ldots\}$, and decompose (M, G) into its restrictions to these spaces. We may define P and Q as the unitary operators whose rows are given by the basis vectors. That is, let $P^1 = \sum_{i,j=1}^{\infty} \langle \tilde{p}^1_{ij}, . \rangle \tilde{p}^1_{ij}$, $Q^1 = \sum_{i,j=1}^{\infty} \langle \tilde{q}^1_{ij}, . \rangle \tilde{q}^1_{ij}$,

$P^{1\perp} = \sum_{i=1}^{\infty} \langle \tilde{p}_i^{1\perp}, . \rangle \, \tilde{p}_i^{1\perp}$, $Q^{1\perp} = \sum_{i=1}^{\infty} \langle \tilde{q}_i^{1\perp}, . \rangle \, q_i^{1\perp}$, $P^2 = \sum_{i,j=1}^{\infty} \langle \tilde{p}_{ij}^2, . \rangle \, \tilde{p}_{ij}^2$, $Q^2 = \sum_{i,j=1}^{\infty} \langle \tilde{q}_{ij}^2, . \rangle \, \tilde{q}_{ij}^2$, $P^{2\perp} = \sum_{i=1}^{\infty} \langle \tilde{p}_i^{2\perp}, . \rangle \, \tilde{p}_i^{2\perp}$, and $Q^{2\perp} = \sum_{i=1}^{\infty} \langle \tilde{q}_i^{2\perp}, . \rangle \, q_i^{2\perp}$.

I claim that with respect to these bases, (M_{11}, G_{11}), has the desired properties. The proof of this fact follows by induction. Denote $P_m = \sum_{i,j=1}^{m} \langle \tilde{p}_{ij}^1, . \rangle \, \tilde{p}_{ij}^1$ and $Q_m = \sum_{i,j=1}^{m} \langle \tilde{q}_{ij}^1, . \rangle \, \tilde{q}_{ij}^1$. To show (M_{11}, G_{11}) are upper-triangular with respect to this basis, it suffices to show $(I - Q_s)MP_s = (I - Q_s)GP_s = 0$ for all $s \in \mathbb{N}$. It then also follows that $(I - Q^1)MP^1 = (I - Q^1)GP^1 = 0$, and so the $(2,1)$, $(3,1)$ and $(4,1)$ elements of M and G are indeed 0 as claimed. To see this, note that by definition of a closed span, for any $x \in E_1$, for all $\delta > 0$, $\exists s$ such that $\|P_s x - x\| < \delta$. Since M and G are continuous, for any $\epsilon > 0$ there exists $\delta > 0$ such that $\|z\| < \delta$ implies $\|Mz\| < \epsilon$, $\|Gz\| < \epsilon$, and so for any $x \in \mathcal{H}_X$, $\exists s \in \mathbb{N}$ s.t. $\|(I - Q^1)MP^1 x\| = \|(I - Q^1)MP_s x\| + \|(I - Q^1)M(P^1 - P_s)x\| < \epsilon$ and similarly $\|(I - Q^1)GP^1 x\| < \epsilon$.

Begin by showing that the first step of the induction chain holds. By construction of the generalized Schur decomposition for the finite-dimensional matrix pair, $\tilde{q}_{11}^1 = q_{11}^1 = \frac{1}{\|G\tilde{p}_{11}^1\|} G\tilde{p}_{11}^1$ and so $(I - Q_1)GP_1 = 0$ and likewise, since $\tilde{p}_{11}^1$ satisfies $M\tilde{p}_{11}^1 = \lambda_1 G\tilde{p}_{11}^1 = \lambda_1 \|G\tilde{p}_{11}^1\| \, \tilde{q}_{11}^1$, $(I - Q_1)MP_1 = 0$. Next, for arbitrary index $s = k \times \ell$ assume the inductive hypothesis $(I - Q_{s-1})MP_{s-1} = (I - Q_{s-1})GP_{s-1} = 0$. By the Gram-Schmidt process, $\tilde{p}_s^1 = \frac{1}{\|(I - P_{s-1})p_s^1\|}(I - P_{s-1})p_s^1$. Since p_s^1 is a generalized Schur vector of a finite-dimensional matrix pair,

$$q_s^1 = \frac{1}{\|(I - \sum_{j=1}^{\ell-1} \langle q_{k,j}^1, . \rangle \, q_{k,j}^1)Gp_s^1\|}(I - \sum_{j=1}^{\ell-1} \langle q_{k,j}^1, . \rangle \, q_{k,j}^1)Gp_s^1 \; (*)$$

and

$$(I - \sum_{j=1}^{\ell-1} \langle q_{k,j}^1, . \rangle \, q_{k,j}^1)Mp_s^1 = \lambda_k (I - \sum_{j=1}^{\ell-1} \langle q_{k,j}^1, . \rangle \, q_{k,j}^1)Gp_s^1 \; (**)$$

, or, in words, p_s^1 is a generalized eigenvector of the matrix pair on the space orthogonal to previous generalized Schur vectors within the block. Now consider $(I - Q_s)G\tilde{p}_s^1 = \frac{1}{\|(I - P_{s-1})p_s^1\|}(I - Q_s)G(I - P_{s-1})p_s^1 = \frac{1}{\|(I - P_{s-1})p_s^1\|}(I - Q_s)Gp_s^1$ by the inductive hypoth-

esis. By $(*)$, $Gp_s^1 \in \text{span}\{q_{k,1}^1, \ldots, q_s^1\} \subset \text{span}\{\tilde{q}_1^1, \ldots, \tilde{q}_s^1\}$ so $(I - Q_s)G\tilde{p}_s^1 = 0$, and since by the inductive hypothesis $(I - Q_s)G\tilde{p}_m^1 = 0$ for $m < s$, $(I - Q_s)GP_s = 0$. Similarly, by $(**)$ and the inductive hypothesis, $(I - Q_s)M\tilde{p}_s^1 = 0$, so it is also the case that $(I - Q_s)MP_s = 0$. By induction, $(I - Q_s)MP_s = (I - Q_s)GP_s = 0$ for all $s \in \mathbb{N}$.

To show that diagonals of (M_{11}, G_{11}) are the generalized eigenvalues, note that the s^{th} diagonal elements with respect to this basis are given by $\langle M\tilde{p}_s^1, \tilde{q}_s^1\rangle$ and $\langle G\tilde{p}_s^1, \tilde{q}_s^1\rangle$. Since $\tilde{p}_s^1 = \frac{1}{\|(I-P_{s-1})p_s^1\|}(I - P_{s-1})p_s^1$, $\tilde{q}_s^1 = \frac{1}{\|(I-Q_{s-1})q_s^1\|}(I - Q_{s-1})q_s^1$, $(I - Q_{s-1})M(I - P_{s-1}) = (I - Q_{s-1})M$ by triangularity, and Q_{s-1} is idempotent and self-adjoint since it is an orthogonal projection, $\langle M\tilde{p}_s^1, \tilde{q}_s^1\rangle = \frac{1}{\|(I-P_{s-1})p_s^1\|}\langle Mp_s^1, \tilde{q}_s^1\rangle$, and similarly $\langle G\tilde{p}_s^1, \tilde{q}_s^1\rangle = \frac{1}{\|(I-P_{s-1})p_s^1\|}\langle Gp_s^1, \tilde{q}_s^1\rangle$. By the finite-dimensional generalized Schur decomposition, $(**)$ holds, and so $\langle M\tilde{p}_s^1, \tilde{q}_s^1\rangle / \langle G\tilde{p}_s^1, \tilde{q}_s^1\rangle = \lambda_k$, and so (M_{11}, G_{11}) has the generalized eigenvalues along the diagonals as desired.

To demonstrate that the $(3,2)$ and $(4,2)$ blocks of (M, G) are equal to 0 is equivalent to requiring that $(I - [Q^1, Q^{1\perp}])M[P^1, P^{1\perp}] = (I - [Q^1, Q^{1\perp}])G[P^1, P^{1\perp}] = 0$. Because $(\{p_{11}^1, \ldots, p_{1k_1}^1\}, \{p_{21}^1, \ldots, p_{2k_2}^1\}, \ldots, \{p_1^{1\perp}, p_2^{1\perp}, \ldots\}, \ldots)$ span Im π_1 and $(\{q_{11}^1, \ldots, q_{1k_1}^1\}, \{q_{21}^1, \ldots, q_{2k_2}^1\}, \ldots, \{q_1^{1\perp}, q_2^{1\perp}, \ldots\}, \ldots)$ span Im π_2, we have by 2 that Im $M[P^1, P^{1\perp}] \subset$ Im $\pi_2 =$ Im $[Q^1, Q^{1\perp}]$ and Im $G[P^1, P^{1\perp}] \subset$ Im $\pi_2 =$ Im $[Q^1, Q^{1\perp}]$ so $M[P^1, P^{1\perp}] = [Q^1, Q^{1\perp}]M[P^1, P^{1\perp}]$ and $G[P^1, P^{1\perp}] = [Q^1, Q^{1\perp}]G[P^1, P^{1\perp}]$ so orthogonality holds.

The proof of the upper-triangular structure of (M_{22}, G_{22}) proceeds similarly to the above, by induction. Denote $P_m^2 = \sum_{i,j=1}^m \langle \tilde{p}_{ij}^2, . \rangle \tilde{p}_{ij}^2$ and $Q_m^2 = \sum_{i,j=1}^m \langle \tilde{q}_{ij}^2, . \rangle \tilde{q}_{ij}^2$. Further, denote $Q_m^u = [Q^1, Q^{1\perp}, Q_m^2]$ the projection onto the set of basis vectors of $\mathcal{H}_Y$ up to $\tilde{q}_m^2$ and similarly $P_m^u = [P^1, P^{1\perp}, P_m^2]$. To show (A_{22}, B_{22}) are upper-triangular with respect to this basis, it suffices to show $(I - Q_s^u)MP_s^2 = (I - Q_s^u)GP_s^2 = 0$ for all $s \in \mathbb{N}$. It then also follows by analogous $\delta - \epsilon$ argument that $(I - [Q^1, Q^{1\perp}, Q^2])MP^2 = (I - [Q^1, Q^{1\perp}, Q^2])GP^2 = 0$, and so the $(4,3)$ elements of M and G are 0 as claimed.

The proof is essentially identical to that for (M_{11}, G_{11}) except that all vectors are orthogonalized with respect to previous basis vectors, and the generalized Schur form of each matrix pair constructs q_{ij}^2 from M instead of G, as on Ker π_1 the spectrum excludes 0 and so M_2 is guaranteed to be invertible while G_2 is not.

Begin by showing the first step of the induction for (M_{22}, G_{22}). By construction of the generalized Schur decomposition for the finite-dimensional matrix pair, $\tilde{q}_{11}^2 = (I - [Q^1, Q^{1\perp}])q_{11}^2 = \frac{(I-[Q^1,Q^{1\perp}])}{\left\|Mp_{11}^2\right\|}Mp_{11}^2$ while $\tilde{p}_{11}^2 = (I - [P^1, P^{1\perp}])p_{11}^2$. As shown above, $(I - [Q^1, Q^{1\perp}])M[P^1, P^{1\perp}] = 0$ and so $\tilde{q}_{11}^2 = \frac{(I-[Q^1,Q^{1\perp}])}{\left\|Mp_{11}^2\right\|}M\tilde{p}_{11}^2$ and so $(I - Q_1^u)MP_1^u = 0$. Likewise, since $\tilde{p}_{11}^2$ satisfies $(I - [Q^1, Q^{1\perp}])Gp_{11}^2 = \frac{1}{\lambda}(I - [Q^1, Q^{1\perp}])Mp_{11}^2 = \frac{1}{\lambda}\left\|Mp_{11}^2\right\|\tilde{q}_{11}^2$, $(I-Q_1^u)BP_1^u = 0$ also. This shows that the first step of the induction holds: the continuation proceeds as for (M_{11}, G_{11}) except switching the order of M and G. Similarly, the presence of the eigenvalues along the diagonals is shown in a completely analogous manner.

It remains to show that $(M_{11}^\perp, G_{11}^\perp)$ satisfies $\sigma(M_{11}^\perp, G_{11}^\perp) \subset \{0\}$. In this, I follow Gohberg *et al.* (1990, Lemma II.3.4) closely. By construction, $(M_{11}^\perp z, G_{11}^\perp z) = (Q_1^\perp M_1 z, Q_1^\perp G_1 z)$ for all $z \in E_1^\perp$. By assumption, G is a bounded operator, so G_1 must be also and so $G_1^{-1}Q_1^\perp G_1$ must be as well. Since the compact operators are a closed ideal within the algebra of bounded operators on a Banach space (see, e.g. (Carl & Stephani, 1990)) and $G_1^{-1}M_1$ is compact by assumption, $G_1^{-1}Q_1^\perp M_1 = G_1^{-1}Q_1^\perp G_1 G_1^{-1}M_1$ is compact also, as is $\Omega_1^\perp := P_1^\perp G_1^{-1}Q_1^\perp M_1 P_1^\perp$, its restriction to $E_1^\perp$. Suppose for contradiction that μ is a nonzero element of $\sigma(M_{11}^\perp, G_{11}^\perp)$. Then by reasoning entirely analogous to .8, $\sigma(M_{11}^\perp, G_{11}^\perp) = \sigma(\Omega_1^\perp)$ and so by compactness μ is an isolated point in the spectrum of $\Omega_1^\perp$. Further, $\Omega_1^{\perp *}$ must have $\bar{\mu} \in \sigma(\Omega_1^{\perp *})$ as a nonzero point in the spectrum, and so by compactness, it must be an isolated point in the spectrum associated with (at least one) nonzero eigenvector, which we will call $x_0 \in E_1^\perp$. The upper triangular decomposition of (M_1, G_1) may be used to show $\Omega_1^{\perp *} = P_1^\perp (G_1^{-1}M_1)^* P_1^\perp$. To see this, note that multiplication of the the upper

triangular decomposition of M_1 by the inverse of the upper triangular decomposition of G_1 yields

$$(G_1^{-1}M_1)^* = \begin{pmatrix} G_{11}^{-1}M_{11} & -G_{11}^{-1}G_{11}^{off}G_{11}^{\perp -1}M_{11}^{off} \\ 0 & G_{11}^{\perp -1}M_{11}^{\perp} \end{pmatrix}^* = \begin{pmatrix} (G_{11}^{-1}M_{11})^* & 0 \\ (-G_{11}^{-1}G_{11}^{off}G_{11}^{\perp -1}M_{11}^{off})^* & \Omega_1^{\perp *} \end{pmatrix}$$

and so $\Omega_1^{\perp *} = P_1^{\perp}(G_1^{-1}M_1)^* P_1^{\perp}$ as claimed. As a result, x_0 is also an eigenvector of compact operator $(G_1^{-1}M_1)^*$ associated with eigenvalue $\bar{\mu}$, and so $x_0 \in E_1^{\perp} \cap$ Im $P_{\{\bar{\mu}\}}^{(G_1^{-1}M_1)^*}$.

However, we know also by 4 that Im $P_{\{\mu\}}^{G_1^{-1}M_1} =$ Im $P_{\{\mu\}}^{(M_1,G_1)} \subset E_1$, and by orthogonality of the decompositions, $E_1^{\perp}$ is orthogonal to Im $P_{\{\mu\}}^{G_1^{-1}M_1}$, and so must be a subset of Ker $(P_{\{\mu\}}^{G_1^{-1}M_1})^*$. Since this is an isolated eigenvalue of an operator on a Hilbert space, Gohberg *et al.* (1990, Prop I.2.5) gives that $(P_{\{\mu\}}^{G_1^{-1}M_1})^* = P_{\{\bar{\mu}\}}^{(G_1^{-1}M_1)^*}$, and so $E_1^{\perp} \subset$ Ker $P_{\{\bar{\mu}\}}^{(G_1^{-1}M_1)^*}$. This contradicts the previous assertion that there is a nonzero element x_0 in $E_1^{\perp} \cap$ Im $P_{\{\bar{\mu}\}}^{(G_1^{-1}M_1)^*}$ and so the original assertion that there is some $\mu \neq 0$ in $\sigma(M_{11}^{\perp}, G_{11}^{\perp})$.

The proof that $(M_{22}^{\perp}, G_{22}^{\perp})$ satisfies $\sigma(M_{22}^{\perp}, G_{22}^{\perp}) \subset \{\infty\}$ is essentially similar to the above, except using $((I-[Q^1, Q^{1\perp}])M(I-[P^1, P^{1\perp}]), (I-[Q^1, Q^{1\perp}])G(I-[P^1, P^{1\perp}]))$ in place of (M_1, G_1) and reversing the order of M and G. □

Perturbation Theory for the Generalized Schur Decomposition

Perturbation for generalized Schur subspaces associated with a subset of the spectrum is covered in Stewart (1973) for perturbations measured in Frobenius norm. In this section, I extend the results to perturbation in operator norm. In addition to bounds on the error in terms of the subspace angle between the approximate and true deflating

subspaces, this section will also consider approximation of the Rayleigh components of the operator pair corresponding to these subspaces. First, set up the generalized Schur subspace approximation problem exactly as in Stewart (1973).

Let $(A, B) \in \mathcal{L}(\mathcal{H}_1 \to \mathcal{H}_2, \mathcal{H}_1 \to \mathcal{H}_2)$ and unitary operators $X = (X_1, X_2)$ $\mathcal{H}_1 \to \mathcal{H}_1$ and $Y = (Y_1, Y_2)$ $\mathcal{H}_2 \to \mathcal{H}_2$ decompose (A, B) as

$$(Y^*AX, Y^*BX) = \left(\begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}, \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix} \right)$$

To find a perturbation bound, we search for the minimal rotations

$$U_X = \begin{pmatrix} I & -P^* \\ P & I \end{pmatrix} \begin{pmatrix} (I + P^*P)^{-1/2} & 0 \\ 0 & (I + PP^*)^{-1/2} \end{pmatrix}$$

$$U_Y = \begin{pmatrix} I & -Q^* \\ Q & I \end{pmatrix} \begin{pmatrix} (I + Q^*Q)^{-1/2} & 0 \\ 0 & (I + QQ^*)^{-1/2} \end{pmatrix}$$

such that $X' = (X'_1, X'_2) = XU_X$ and $Y' = (Y'_1, Y'_2) = YU_Y$ generate subspaces $\mathscr{R}(X'_1) = \mathscr{X} \subset \mathcal{H}_1$ and $\mathscr{R}(Y'_1) = \mathscr{Y} \subset \mathcal{H}_2$ which form a deflating pair of (A, B). A pair of subspaces $\mathscr{X}, \mathscr{Y}$ form a deflating pair if and only if $(A'_{21}, B'_{21}) = (0, 0)$. This is equivalent to

$$\begin{aligned} QA_{11} - A_{22}P &= A_{21} - QA_{12}P \\ QB_{11} - B_{22}P &= B_{21} - QB_{12}P \end{aligned} \tag{5}$$

In order to find (Q, P) which satisfy the above condition and are small relative to perturbations in operator norm, define a norm over the space of operator pairs over subspaces conformable to the pair (Q, P) as the largest operator norm of an operator

in the pair, i.e.

$$\|(Q,P)\|_{\mathcal{B}} = \max(\|Q\|, \|P\|)$$

If we can show that the conditions of Stewart (1973) Theorem 3.1 are satisfied for 5 using this norm, then this theorem will provide a bound on the operator norm of the rotation needed to generate such a decomposition. Define

$$T(Q,P) = \begin{pmatrix} QA_{11} - A_{22}P & QB_{11} - B_{22}P \end{pmatrix}$$

$$g = \begin{pmatrix} A_{21} & B_{21} \end{pmatrix}$$

$$\varphi(Q,P) = \begin{pmatrix} QA_{12}P & QB_{12}P \end{pmatrix}$$

To show a quadratic bound for $\varphi(Q,P)$, begin with the first term:

$$\begin{aligned} \|\varphi_1(Q,P)\| &\leq \|Q\| \, \|P\| \, \|A_{12}\| \\ &\leq \|(Q,P)\|_{\mathcal{B}}^2 \, \|A_{12}\| \end{aligned}$$

Combining with identical calculations for the second term yields quadratic bound

$$\|\varphi(Q,P)\|_{\mathcal{B}} \leq \eta \|(Q,P)\|_{\mathcal{B}}^2 \tag{6}$$

where

$$\eta = \|(A_{12}, B_{12})\|_{\mathcal{B}}$$

To demonstrate the Lipschitz property for this operator, again note

$$\begin{aligned}\|\varphi_1(Q,P)-\varphi_1(\tilde{Q},\tilde{P})\| &\leq \left\|Q-\tilde{Q}\right\|\|P\|\,\|A_{12}\|+\left\|\tilde{Q}\right\|\left\|P-\tilde{P}\right\|\|A_{12}\| \\ &\leq 2\max(\|(Q,P)\|_{\mathcal{B}},\|(\tilde{Q},\tilde{P})\|_{\mathcal{B}})\|(Q-\tilde{Q},P-\tilde{P})\|_{\mathcal{B}}\,\|A_{12}\|\end{aligned}$$

Combining with identical calculations for the second term gives Lipschitz condition

$$\|\varphi(Q,P)-\varphi(\tilde{Q},\tilde{P})\|_B \leq 2\eta\max(\|(Q,P)\|_{\mathcal{B}},\|(\tilde{Q},\tilde{P})\|_{\mathcal{B}})\|(Q-\tilde{Q},P-\tilde{P})\|_{\mathcal{B}} \qquad (7)$$

These demonstrate that conditions (i) and (ii) of Theorem 3.1 in Stewart (1973) continue to hold for the norm $\|.\|_{\mathcal{B}}$

Again defining

$$\gamma = \|g\|_{\mathcal{B}}$$

$$\delta = \|T^{-1}\|_{\mathcal{B}}^{-1}$$

one obtains

Lemma .11. *Suppose $T(Q,P) = g-\varphi(Q,P)$ with T, g, and φ defined as above, where φ satisfies the quadratic bound and Lipschitz conditions. Let $\delta > 0$ and $\gamma\eta/\delta^2 < 1/4$. Then*

$$\|\begin{pmatrix} Q, & P \end{pmatrix}\|_{\mathcal{B}} < 2\frac{\gamma}{\delta}$$

To determine precisely how the above theorem imposes bounds on errors in Schur subspaces, it is necessary to examine the stability properties of the term δ. Define

$$\mathrm{dif}(A,B) = \mathrm{dif}(\begin{matrix} A_{11} & B_{11} \\ A_{22} & B_{22} \end{matrix}) = \|T^{-1}\|_{\mathcal{B}}^{-1} \qquad (8)$$

Note that this operator depends on only the block diagonal terms of the pair (A, B). Define the perturbation $(E, F) \in \mathcal{L}(\mathcal{H}_1 \to \mathcal{H}_2, \mathcal{H}_1 \to \mathcal{H}_2)$ and define a partition of the operator conformable with that of (A, B) by $(E_{ij}, F_{ij}) = (Y_i^H E X_j, Y_i^H F X_j)$. We would like to define a bound on the term

$$\mathrm{dif}(A + E, B + F) = \mathrm{dif}(\begin{array}{cc} A_{11} + E_{11} & B_{11} + F_{11} \\ A_{22} + F_{22} & B_{22} + F_{22} \end{array})$$

Using the alternate characterization $\mathrm{dif}(A, B) = \inf_{\|Z\|_{\mathcal{B}}=1} \|T(Z)\|_{\mathcal{B}}$ where $Z \in \mathcal{B}$, one can derive lower and upper bounds

$$\mathrm{dif}(A, B) + \nu(E, F) \geq \mathrm{dif}(A + E, B + F) \geq \mathrm{dif}(A, B) - \nu(E, F)$$

where

$$\nu(E, F) = \max(\|E_{11}\| + \|E_{22}\|, \|F_{11}\| + \|F_{22}\|)$$

Combing this bound with the previous lemma, obtain

Theorem .2. *Let (A, B) and $(E, F) \in \mathcal{L}(\mathcal{H}_1 \to \mathcal{H}_2, \mathcal{H}_1 \to \mathcal{H}_2)$ and $X = (X_1, X_2)$ $\mathcal{H}_1 \to \mathcal{H}_1$ and $Y = (Y_1, Y_2)$ $\mathcal{H}_2 \to \mathcal{H}_2$ be unitary operators such that $\mathscr{R}(X_1)$ and $\mathscr{R}(Y_1)$ form a deflating pair of subspaces for the operator pair (A, B). Suppose these operators partition the pairs such that*

$$(Y^H A X, Y^H B X) = \left(\begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix}, \begin{bmatrix} B_{11} & B_{12} \\ 0 & B_{22} \end{bmatrix} \right)$$

$$(Y^H E X, Y^H F X) = \left(\begin{bmatrix} E_{11} & E_{12} \\ E_{21} & E_{22} \end{bmatrix}, \begin{bmatrix} F_{11} & F_{12} \\ F_{21} & F_{22} \end{bmatrix} \right)$$

Define

$$\delta = \mathit{dif}(A, B) - \nu(E, F)$$

along with

$$\gamma = \|(\ E_{21} \quad F_{21})\ \|_{\mathcal{B}}$$

and

$$\eta = \|(A_{12} + E_{12}, B_{12} + F_{12})\|_{\mathcal{B}}$$

Suppose $\delta > 0$ *and* $\gamma\eta/\delta^2 < 1/4$. *Then*

Then there is a pair of operators (Q, P) *with*

$$\left\| \begin{pmatrix} Q, & P \end{pmatrix} \right\|_{\mathcal{B}} \leq \frac{2\gamma}{\delta}$$

such that

$X_1' = (X_1 + X_2P)(I + P^*P)^{-1/2}$

$Y_1' = (Y_1 + Y_2Q)(I + Q^*Q)^{-1/2}$

and $\mathscr{R}(X_1')$ *and* $\mathscr{R}(Y_1')$ *form a pair of deflating subspaces for* $(A + E, B + F)$.

This is essentially identical to Theorem 5.7 of Stewart (1973) aside from the definition of the norms via which the terms are defined and the resulting difference in the lower bound on δ.

Via Theorem 2.7 in Stewart (1973), we know that

$$\|\sin\Theta(\mathscr{R}(X_1), \mathscr{R}(X_1'))\| \leq \|\tan\Theta(\mathscr{R}(X_1), \mathscr{R}(X_1'))\| = \|P\|$$

$$\|\sin\Theta(\mathscr{R}(Y_1), \mathscr{R}(Y_1'))\| \leq \|\tan\Theta(\mathscr{R}(Y_1), \mathscr{R}(Y_1'))\| = \|Q\|$$

both of which are less than $\left\| \begin{pmatrix} Q, & P \end{pmatrix} \right\|_{\mathcal{B}}$. As a result, we have the following corollary

Corollary .1. *Suppose* (A, B), (E, F), X *and* Y *satisfy the conditions of the theorem above. Then the operator pair* $(A + E, B + F)$ *has a right generalized Schur subspace*

$\mathscr{R}(X_1')$ such that $\|Proj_{X_1'} - Proj_{X_1}\|_2 \leq \frac{2\gamma}{\delta}$ and associated left generalized Schur subspace $\mathscr{R}(Y_1')$ such that $\|Proj_{Y_1'} - Proj_{Y_1}\|_2 \leq \frac{2\gamma}{\delta}$

As a result, for appropriately small approximation error in the operator pair of interest, a fixed, well-separated, primary generalized Schur subspace (and associated generalized Schur functions or vectors whose range spans it) of the perturbed pair differs by an amount which is on the order of the operator norm of the perturbation from the corresponding true subspace (and associated functions). This dependence on the order of the operator norm of the error may be particularly useful in the case of large or infinite-dimensional subspaces, for which the Frobenius norm of the error may increase as the square root of the dimension of the subspace. One loses, however, the set of sharp characterizations of the difference term δ in terms of spectral properties of the operator to be approximated which may be obtained when it is defined via the Frobenius norm. This seems necessary in general, however, as the Frobenius or Hilbert-Schmidt norm may fail to be finite in the infinite-dimensional case for otherwise well-behaved operators.

To bound the approximation error in the components (A_{11}, B_{11}) induced by an approximation, it is helpful to introduce an additional pair of subspaces to correspond to the right deflating pair $\mathscr{R}(X_1)$ and $\mathscr{R}(Y_1)$. Defining (X_1, X_2) and (Y_1, Y_2) as above so $\mathscr{R}(X_1)$ and $\mathscr{R}(Y_1)$ form a deflating pair, we look for operators V_1 and U_2 and R and S with $V_1 = Y_1 + Y_2 R^*$ and $U_2 = X_2 - X_1 S$ to solve

$$(V_1, Y_2)^* A (X_1, U_2) = \begin{pmatrix} I & R \\ 0 & I \end{pmatrix} \begin{pmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{pmatrix} \begin{pmatrix} I & -S \\ 0 & I \end{pmatrix} = \begin{pmatrix} A_{11} & 0 \\ 0 & A_{22} \end{pmatrix}$$

$$(V_1, Y_2)^* B (X_1, U_2) = \begin{pmatrix} I & R \\ 0 & I \end{pmatrix} \begin{pmatrix} B_{11} & B_{12} \\ 0 & B_{22} \end{pmatrix} \begin{pmatrix} I & -S \\ 0 & I \end{pmatrix} = \begin{pmatrix} B_{11} & 0 \\ 0 & B_{22} \end{pmatrix} \quad (9)$$

This holds if there exist S, R such that

$$\begin{aligned} A_{11}S - RA_{22} &= A_{12} \\ B_{11}S - RB_{22} &= B_{12} \end{aligned}$$

Theorem 5.9 in Stewart (1973) notes that if T is nonsingular, there exist S and R which solve this equation, and so (X_1, U_2) and (V_1, Y_2), which are not in general unitary, though are nonsingular, block diagonalize (A, B). Further, by the definitions of V_1 and U_2, one has $\|V_1\| = \|\sec\Theta(\mathscr{R}(V_1), \mathscr{R}(Y_1)\| = \|\sec\Theta_1\|$ and $\|U_2\| = \|\sec\Theta(\mathscr{R}(U_2), \mathscr{R}(X_2)\| = \|\sec\Theta_2\|$.

This block diagonalization can be used along with the perturbation formula to construct bounds on the approximation error in (A_{11}, B_{11}). Consider a perturbation (E, F) of (A, B) and define

$$((V_1, Y_2)^* E(X_1, U_2), (V_1, Y_2)^* F(X_1, U_2)) = \left(\begin{pmatrix} E_{11} & E_{12} \\ E_{21} & E_{22} \end{pmatrix}, \begin{pmatrix} F_{11} & F_{12} \\ F_{21} & F_{22} \end{pmatrix} \right)$$

so that perturbed operator pair satisfies

$$((V_1, Y_2)^*(A+E)(X_1, U_2), (V_1, Y_2)^*(B+F)(X_1, U_2)) = \left(\begin{pmatrix} A_{11}+E_{11} & E_{12} \\ E_{21} & A_{22}+E_{22} \end{pmatrix}, \begin{pmatrix} B_{11}+F_{11} & F_{12} \\ F_{21} & B_{22}+F_{22} \end{pmatrix} \right) \quad (10)$$

then, following Stewart & Sun (1990) VI.2.15, we have

Theorem .3. *Define*

$$\delta = \mathit{dif}(A, B) - \max(\|E_{11}\| + \|E_{22}\|, \|F_{11}\| + \|F_{22}\|)$$

along with

$$\gamma = \|(\ E_{21} \quad F_{21})\ \|_{\mathcal{B}}$$

and

$$\eta = \|(E_{12}, F_{12})\|_{\mathcal{B}}$$

Suppose $\delta > 0$ *and* $\gamma\eta/\delta^2 < 1/4$.

Then there is a pair of operators (Q, P) *with*

$$\left\| \begin{pmatrix} Q, & P \end{pmatrix} \right\|_{\mathcal{B}} \leq \frac{2\gamma}{\delta}$$

such that

$X_1' = X_1 + U_2P$

$Y_2' = Y_2 + V_1Q^*$

satisfy

$$((V_1, Y_2')^*(A+E)(X_1', U_2), (V_1, Y_2')^*(B+F)(X_1', U_2)) =$$
$$\left(\begin{pmatrix} A_{11} + E_{11} + E_{12}P & E_{12} \\ 0 & A_{22} + E_{22} + QE_{12} \end{pmatrix}, \begin{pmatrix} B_{11} + F_{11} + F_{12}P & F_{12} \\ 0 & B_{22} + F_{22} + QF_{12} \end{pmatrix} \right) \tag{11}$$

and so (A_{11}', B_{11}') $=(A_{11} + E_{11} + E_{12}P, B_{11} + F_{11} + F_{12}P)$ *form the generalized Rayleigh quotients of the perturbed operator pair, and as a result, we have*

$$\|A_{11} - A_{11}'\| \leq \|E_{11} + E_{12}P\| \leq \|E_{11}\| + \|E_{12}\|\frac{2\gamma}{\delta}$$

$$\|B_{11} - B_{11}'\| \leq \|F_{11} + F_{12}P\| \leq \|F_{11}\| + \|F_{12}\|\frac{2\gamma}{\delta}$$

Proof. Existence of a unique solution (Q, P) with the specified properties follows if there exist (Q, P) such that left multiplying (10) by $\begin{pmatrix} I & 0 \\ Q & I \end{pmatrix}$ and right multiplying by $\begin{pmatrix} I & 0 \\ P & I \end{pmatrix}$ sets the lower left elements in (11) to 0. This holds if there is unique solution to

$$\begin{pmatrix} Q(A_{11} + E_{11}) + (A_{22} + E_{22})P \\ Q(B_{11} + F_{11}) + (B_{22} + F_{22})P \end{pmatrix} = \begin{pmatrix} E_{21} \\ F_{21} \end{pmatrix} + \begin{pmatrix} QE_{12}P \\ QF_{12}P \end{pmatrix}$$

Existence of a unique solution here follows from application of Theorem 3.1 in Stewart (1973), the Lipschitz and norm bound shown for the quadratic component above, and the lower bound on δ which lower bounds the minimum singular value of the lefthand side. □

Appendix B: Existence of Solutions To Rational Expectations Models with Function-Valued States

In this appendix, I provide a set of sufficient conditions for the existence of a differentiable and stable solution to a recursive model with function valued states, which is the object for which an approximation algorithm is described in the main paper.

Local Existence of an Equilibrium

The requirement that the derivatives of the equilibrium policy operators satisfy the formulas (4.3) and (4.6) in the main paper represents a necessary condition that any differentiable recursive equilibrium solution must satisfy at the steady state. To ensure that an equilibrium characterized by this condition in fact exists and that it is locally stable, conditions beyond those needed to justify the existence of stable derivatives at this point may be needed. Existence of equilibria has previously been demonstrated in dynamic heterogeneous agent models of various types: for models with no aggregate disturbances, see Acemoglu & Jensen (2012), for models with aggregate disturbances see Miao (2006) or Bergin & Bernhardt (1995), and for models with finite numbers of agents, see Mertens & Judd (2012). In the case that an equilibrium of the model

may be characterized by the methods used in those papers, our method provides a numerical procedure to characterize a set of local properties. However, it may be desirable, in the case of a wide class of models which may be characterized by the methods in this paper, to provide a general existence argument based on the linearization. As such, we demonstrate that the local existence argument of Jin & Judd (2002) may be extended to the case of infinite-dimensional state variables.

In particular, under some additional continuity conditions, an implicit function theorem in Banach spaces (see e.g. Kesavan (2004, Ch. 1)) may be used to show that the conditions provided under which a linearization exist are also sufficient for the existence of an equilibrium with no aggregate shocks in a neighborhood of a steady state. This is a separate task from demonstration of the existence of a steady state equilibrium, which must be shown by other methods, as in Acemoglu & Jensen (2012) or by some other fixed point argument valid in infinite-dimensional spaces. In some cases, especially if individual agent decision rules are correspondence valued, the existence of a steady state may require the introduction of auxiliary state variables to generate a Markov structure; this occurs in Acemoglu & Jensen (2012) and Miao (2006), along with, more generally, much of the literature on implicit recursive contracts. As characterizing a steady state is a model-specific task, we will simply require that existence has been verified for the model in question and that the equilibrium conditions which are provided to the linearization procedure are sufficient to characterize an equilibrium.

Given the existence of an equilibrium away from the steady state but with no aggregate shocks, one may prove existence of an equilibrium with "small" aggregate shocks by applying the implicit function theorem to the deterministic system around $\sigma = 0$ as in Mertens & Judd (2012). Unfortunately, the continuation argument they use to demonstrate existence in the presence of large noise does not directly apply in this case, as it requires compactness of the equilibrium operators, which cannot

generally be guaranteed when the state-space is infinite-dimensional. As such, the approximations computed in this paper should be considered as valid in the "small noise" regime.

Formally, we have the following two theorems: proofs are provided in the next section.

Theorem .4. *Suppose there exist x^*, y^* such that $F(x^*, y^*, x^*, y^*, 0) = 0$, F is continuous with uniformly continuous Fréchet derivatives with respect to x, y, x', and y' in a neighborhood $\Omega \times \Omega \subset \mathcal{H}_1 \times \mathcal{H}_1$ of x^*, y^*, x^*, y^*, with Fréchet derivatives at this point given by $B = -F_x, -F_y$, $A = F_{x'}, F_{y'}$ such that (B, A) is a Γ-regular operator pair satisfying the conditions of Lemma (5) on existence of Schur decomposition for Γ the complex unit circle and so having generalized Schur decomposition $(B, A) = (Q^*TU, Q^*SU)$. Assume U_{22} is complete and has bounded inverse on Im U_2 and $M(y, x) := F_y y + (F_{x'} - F_{y'} U_{22}^{-1} U_{21})x : \mathcal{H}_y, \mathcal{H}_x \to \mathcal{H}_2$ has bounded inverse. Then, there exists a neighborhood $\mathcal{N}$ of x^* in $\mathcal{H}_x$ and continuous, Fréchet differentiable operators $g(x)$, $h(x)$ mapping $\mathcal{N}$ to $\mathcal{H}_y$ and $\mathcal{H}_x$, respectively, such that for any $x_0 \in \mathcal{N}$, the sequence $\{x_t, y_t\}_{t=0}^{\infty}$ defined recursively by $y_t = g(x_t)$, $x_{t+1} = h(x_t)$ satisfies $F(x_t, y_t, x_{t+1}, y_{t+1}, 0) = 0 \ \forall t \geq 0$ and converges in norm to x^*, y^*. Further, $g(x)$ and $h(x)$ are themselves Fréchet differentiable with first derivatives $g_X = -U_{22}^{-1} U_{21}$ and $h_X = (U_{11} + U_{12} g_X)^{-1} S_{11}^{-1} T_{11} (U_{11} + U_{12} g_X)$.*

Remark. Uniform continuity of the derivatives is needed to ensure convergence uniformly over time: a sufficient condition for this would be that F is twice continuously differentiable in $\Omega \times \Omega$. The assumption of continuous Fréchet derivatives could be replaced by strong Hadamard differentiability, with the corresponding weaker result that the policy operators are Hadamard differentiable, using the implicit function theorem for Hadamard derivatives of Craven & Nashed (1982) with only minor changes in the proof. Such a replacement may be necessary for certain classes of equilibrium operators: Craven & Nashed (1982) provide examples of operators which are strongly

Hadamard differentiable but not Fréchet differentiable. The same result may also be used to relax the assumption that the inverse of M is bounded, at cost of a weaker norm, which may be useful for example if M is compact. However, the argument for existence of a stochastic equilibrium in the next theorem does rely on the assumption that M has bounded inverse in a way that cannot be alleviated by using this weaker version of the implicit function theorem. Note that this theorem requires completeness of U_{22}, ruling out cases analogous to those in which there are more stable eigenvalues than predetermined variables, in which there may be indeterminacy. While in these cases at least one equilibrium may possibly still exist, this method of proof does not apply directly.

Proof. See next section. □

Next, one may use the deterministic recursive equilibrium constructed above, along with another application of the implicit function theorem, to demonstrate existence of an equilibrium for σ in a neighborhood of 0, and so with stochastic shocks. The demonstration of existence follows closely the argument (but not the notation) leading to Jin & Judd (2002, Theorem 6).

Theorem .5. *Suppose the conditions of the previous theorem hold, that* $\mathbb{E}F(x, y, x', y', \sigma)$ *is continuous and three times continuously differentiable with respect to* $x,\ y,\ x',\ y'$ *in a neighborhood of* $(x^*, y^*, x^*, y^*, 0)$ *and differentiable with respect to* σ *at* $(x^*, y^*, x^*, y^*, 0)$*, and suppose further that* $[F_y, F_{x'}]$ *has a bounded inverse from* $\mathcal{H}_2$ *to* $\mathcal{H}_y \times \mathcal{H}_x$ *and the operator* $[F_y, F_{x'}]^{-1}[F_{y'}, F_{y'}g_X] : \mathcal{H}_y \times \mathcal{H}_x \to \mathcal{H}_y \times \mathcal{H}_x$ *has spectrum inside the complex unit circle. Let* z *satisfy* $\mathbb{E}z' = 0$ *and have bounded support in* $\mathcal{H}_x$*. Then, there exists a neighborhood of* $\sigma = 0$ *and a neighborhood* $\mathcal{U}$ *of* x^* *on which there exist bounded functions with bounded derivatives* $g(x, \sigma),\ h(x, \sigma)$ *satisfying* $g(x, 0) = g(x),\ h(x, 0) = h(x)$ *as defined in the previous theorem and* $\mathbb{E}F(x, g(x, \sigma), h(x, \sigma) + \sigma\eta z', g(h(x, \sigma) + \sigma\eta z', \sigma), \sigma) = 0$ *for all* σ *and* x *in this neighborhood.*

Remark. The increase from once to three times continuous differentiability here is imposed in order to ensure uniform continuity of the second derivatives, which enter through a chain rule condition due to the recursive construction. This theorem demonstrates conditions which guarantee the existence of an implicit solution to a set of equilibrium equations. This condition is still weaker than those of Jin & Judd (2002), who require analyticity, albeit mainly for the ability to perform perturbations to arbitrary order. The requirement of bounded support is somewhat restrictive, but without other strong assumptions, is very hard to relax: see Jin & Judd (2002) for discussion. Note that for this result to guarantee existence of an equilibrium, a stationary solution to $\mathbb{E}F(x, g(x,\sigma), h(x,\sigma) + \sigma\eta z', g(h(x,\sigma) + \sigma\eta z', \sigma), \sigma) = 0$ must be a sufficient and not merely a necessary condition for equilibrium. That is, if desiderata for an equilibrium such as transversality or second order conditions are excluded from F, the solution is not guaranteed to satisfy them. Often, auxiliary equations may need to be added to the 'natural' characterization of a solution to fully characterize an equilibrium, as in Miao (2006). However, the result does guarantee existence of a candidate solution for which sufficient conditions may then be verified.

Proof. See next section. □

Remark. In the case where the equilibrium conditions are differentiable also in σ, the implicit function characterizes the derivatives with respect to σ.

Corollary. *Let $\mathbb{E}F(x, y, x', y', \sigma)$ be Fréchet differentiable in all its arguments (including σ) at $(x^*, y^*, x^*, y^*, 0)$. Further, suppose σ does not enter into F directly (as opposed to entering via the transition equation $x_2 = h_2(x_2) + \sigma z'$). Then $g(x,\sigma),\ h(x,\sigma)$ are differentiable in σ and have derivative at $\sigma = 0$ given by $g_\sigma(x, 0) = 0$, $h_\sigma(x, 0) = 0$.*

Proof. This follows from the implicit function theorem used to construct $g(x,\sigma),\ h(x,\sigma)$

and the fact that, evaluated at $\sigma = 0$

$$\frac{d}{d\sigma}\mathbb{E}F(x, g(x), h(x) + \sigma\eta z', g(h(x) + \sigma\eta z'), \sigma) = \mathbb{E}(F_{x'}(x, g(x), h(x), g(h(x)), 0)[\eta z'] + F_{y'}(x, g(x), h(x), g(h(x)), 0)g_x(h(x))[\eta z'] + F_\sigma(x, g(x), h(x), g(h(x)), 0)).$$

By the assumption that $\mathbb{E}z' = 0$ and the linearity of the Fréchet derivative for all x, the first two terms are 0. By assumption $F_\sigma = 0$, so the final term disappears as well. □

Remark. This result is a direct extension to this setting of the result of Schmitt-Grohe & Uribe (2004) that the first order impact of the standard deviation parameter is 0. However, the result is slightly stronger, as the implicit function theorem used here (as in Jin & Judd (2002)) takes as argument the policy operators as functions of x, and so the implicit function theorem characterizes the partial derivative of the operators $g(.)$ and $h(.)$ in the space of functions with respect to σ. The implication is that the zero first order effect of adding aggregate noise holds not only at the steady state x^* but also at any other initial condition x, a result which may be useful for extending models with transition dynamics. The assumption that $F_\sigma = 0$ is generally not restrictive, as σ is not a structural parameter but an auxiliary parameter scaling the deviation of the equilibrium away from the nonstochastic steady state. To embed structural assumptions regarding the variance of the shocks, the random element z may be taken to have arbitrary (trace class) covariance operator, which is then scaled by σ.

Proofs

Theorem .4

Proof. Define the space of potential sequences of (aggregate) states $\ell^1(\mathcal{H}_1)$ as the

space of sequences of deviations $\{x_t, y_t\}_{t=0}^{\infty}$ such that $\{x_t + x^*, y_t + y*\}_{t=0}^{\infty} \in \mathcal{H}_1 = \mathcal{H}_x \times \mathcal{H}_y$ from time 0 to ∞ endowed with the norm $\|\{x_t, y_t\}_{t=0}^{\infty}\|_{\ell^1(\mathcal{H}_1)} = \sum_{t=0}^{\infty} \|(x_t - x^*, y_t - y^*)\|_{\mathcal{H}_1}$ having finite norm. Define also $\ell^1_{y_0}(\mathcal{H}_1)$ as the subspace of $\ell^1(\mathcal{H}_1)$ consisting of sequences $\{y_0, \{x_t, y_t\}_{t=1}^{\infty}\}$, that is, excluding x_0, endowed with the relative topology. These are Banach spaces since each element of a sequence is a member of a Banach space and the space of norm summable sequences is complete. We define an equilibrium as a sequence $\{x_t, y_t\}_{t=0}^{\infty} \in \ell^1(\mathcal{H}_1)$ satisfying $F(x_t + x^*, y_t + y^*, x_{t+1} + x^*, y_{t+1} + y^*, 0) = 0 \ \forall t \geq 0$ and endow the space of norm-summable sequences in $\mathcal{H}_2$ with the $\ell^1(\mathcal{H}_2)$ norm $\|\{z_t\}_{t=0}^{\infty}\|_{\ell^1(\mathcal{H}_2)} = \sum_{t=0}^{\infty} \|z_t\|_{\mathcal{H}_2}$, making it also a Banach space, and the operator $F^{\infty}(\{x_t, y_t\}_{t=0}^{\infty}) = \{F(x_t + x^*, y_t + y^*, x_{t+1} + x^*, y_{t+1} + y^*, 0)\}_{t=0}^{\infty}$ a map from $\ell^1(\mathcal{H}_1) \to \ell^1(\mathcal{H}_2)$ (where boundedness follows if $F(.,.,.,.,0)$ is bounded from $\mathcal{H}_1 \times \mathcal{H}_1$ endowed with norm $\|x, y\|_{\mathcal{H}_1} + \|x', y'\|_{\mathcal{H}_1}$ is bounded). By assumption $\{x_t, y_t\}_{t=0}^{\infty} = \{0, 0\}_{t=0}^{\infty}$ satisfies $F^{\infty}(\{0, 0\}_{t=0}^{\infty}) = 0$. F^{∞} is continuous on an $\ell^1(\mathcal{H}_1)$ neighborhood of $\{0, 0\}_{t=0}^{\infty}$ since for any pair of sequences $(\{x_t, y_t\}_{t=0}^{\infty})_i$, $(\{x_t, y_t\}_{t=0}^{\infty})_j$, we have $\|F^{\infty}((\{x_t, y_t\}_{t=0}^{\infty})_i) - F^{\infty}((\{x_t, y_t\}_{t=0}^{\infty})_j)\|_{\ell^1(\mathcal{H}_2)} = \sum_{t=0}^{\infty} \|F(x_t^i + x^*, y_t^i + y^*, x_{t+1}^i + x^*, y_{t+1}^i + y^*, 0) - F(x_t^j + x^*, y_t^j + y^*, x_{t+1}^j + x^*, y_{t+1}^j + y^*, 0)\|_{\mathcal{H}_2}$ which is bounded by the $\ell^1(\mathcal{H}_1)$ norm of the difference in arguments times the modulus of uniform continuity of $F(x_t + x^*, y_t + y^*, x_{t+1} + x^*, y_{t+1} + y^*, 0)$, which is finite in a bounded neighborhood of $0, 0, 0, 0$ due to continuity of the Fréchet derivatives of F. Split the arguments of F^{∞} into $(x_0, s) \in \mathcal{H}_x \times \ell^1_{y_0}(\mathcal{H}_1)$ and define $\frac{\partial F^{\infty}}{\partial x_0}$ and $\frac{\partial F^{\infty}}{\partial s}$ as the partial Fréchet derivatives of F^{∞} with respect to x_0 and s, respectively. These exist and are continuous at $\{x^*, y^*\}_{t=0}^{\infty}$. This is easily seen for $\frac{\partial F^{\infty}}{\partial x_0} = \{F_x(x^*, y^*, x^*, y^*), 0, 0, 0, ...\}$ since F_x is continuous by assumption. For s, this is a little less straightforward: its rows are given by

$$[\frac{\partial F^{\infty}(\{x_t, y_t\}_{t=0}^{\infty})}{\partial s}]_0 = [F_y(x_0 + x^*, y_0 + y^*, x_1 + x^*, y_1 + y^*)$$

$$F_{x'}(x_0 + x^*, y_0 + y^*, x_1 + x^*, y_1 + y^*) \quad F_{y'}(x_0 + x^*, y_0 + y^*, x_1 + x^*, y_1 + y^*) \quad 0 \quad 0 \quad \dots \ ]$$

$$[\frac{\partial F^\infty(\{x_t,y_t\}_{t=0}^\infty)}{\partial s}]_1 =$$

$$[\ 0 \quad F_x(x_1+x^*,y_1+y^*,x_2+x^*,y_2+y^*) \quad F_y(x_1+x^*,y_1+y^*,x_2+x^*,y_2+y^*)$$

$$F_{x'}(x_1+x^*,y_1+y^*,x_2+x^*,y_2+y^*) \quad F_{y'}(x_1+x^*,y_1,x_2+x^*,y_2) \quad \ldots\ ]$$

$$[\frac{\partial F^\infty(\{x_t,y_t\}_{t=0}^\infty)}{\partial s}]_2 = [\ 0 \quad 0 \quad 0 \quad F_x(x_2+x^*,y_2+y^*,x_3+x^*,y_3+y^*)$$

$$F_y(x_2+x^*,y_2y^*,x_3+x^*,y_3+y^*) \quad F_{x'}(x_3+x^*,y_3+y^*,x_4+x^*,y_4+y^*)$$

$$F_{y'}(x_3+x^*,y_3,x_4+x^*,y_4) \quad \ldots]$$

etc.

It is a linear operator mapping sequences $s = \{y_0, x_1, y_1, x_2, y_2, \ldots\} \in \ell^1_{y_0}(\mathcal{H}_1)$ to $\ell^\infty(\mathcal{H}_2)$, as each block is bounded and each row consists of a finite set of blocks. Continuity of this operator (with respect to operator norm on the space $\mathcal{L}(\ell^1_{y_0}(\mathcal{H}_1) \to \ell^1(\mathcal{H}_2))$ is given by bounding

$$\sup_{\|s\|_{\ell^1_{y_0}(\mathcal{H}_1)}=1} \|[\frac{\partial F^\infty(\{x_t,y_t\}_{t=0}^\infty)_i}{\partial s} - \frac{\partial F^\infty(\{x_t,y_t\}_{t=0}^\infty)_j}{\partial s}]s\|_{\ell^1(\mathcal{H}_2)}$$

which, if $(\{x_t,y_t\}_{t=0}^\infty)_i$ and $(\{x_t,y_t\}_{t=0}^\infty)_j$ are in the neighborhood over which the Fréchet derivatives are uniformly continuous, is less than the sum of the moduli of uniform continuity of F_x, F_y, $F_{x'}$, and $F_{y'}$ (which is by assumption finite) times the $\ell^1(\mathcal{H}_1)$ distance between the sequences. This inequality holds since for any row of $\frac{\partial F^\infty}{\partial s}$, only the set of elements of the sequence corresponding to times t and $t+1$ enters into the value of the operator.

Finally, to apply the implicit function theorem to F^∞ to solve for s as a function

of x_0, we must show that $\frac{\partial F^\infty(\{0,0\}_{t=0}^\infty)}{\partial s}$ is invertible. To do this, we show that it is surjective and injective. We prove that it is surjective constructively, by constructing, for any $\{a_i\}_{i=0}^\infty \in \ell^\infty(\mathcal{H}_2)$ an element $s \in \ell^1_{y_0}(\mathcal{H}_1)$ such that $\frac{\partial F^\infty(\{0,0\}_{t=0}^\infty)}{\partial s} s = \{a_i\}_{i=0}^\infty$. We may construct s recursively, using g_X and h_X as defined in the theorem to solve for the values of y, x', and y' consistent with values of x and y at previous times. To find the starting value, we solve $a_0 = F_y y_0 + F_{x'} x_1 + F_{y'} y_1$ by finding y_0 and x_1 consistent with a_0 under the assumption that y_1 may be constructed by the law of motion given by g_X, permitting all subsequent time periods to be solved for recursively. That is, we find y_0, x_1 solving $a_0 = F_y y_0 + F_{x'} x_1 + F_{y'} g_X x_1 = M(y_0, x_1)$. By invertibility of M, we may define $(y_0, x_1) = M^{-1} a_0$ and choose $y_1 = g_X x_1$ so that the first row of $\frac{\partial F^\infty(\{0,0\}_{t=0}^\infty)}{\partial s}[\{y_0, \{x_t, y_t\}_{t=1}^\infty\}]$ equals a_0 by construction. To ensure that the next row holds, with x_1 and $y_1 = g_X x_1$ given, we must solve $a_1 = \begin{bmatrix} F_x & F_y \end{bmatrix} \begin{bmatrix} x_1 \\ g_X x_1 \end{bmatrix} + \begin{bmatrix} F_{x'} & F_{y'} \end{bmatrix} \begin{bmatrix} x_2 \\ y_2 \end{bmatrix}$. This equation has 'fundamental solution' for $a_2 = 0$ given by the recursive update rule $x_2 = h_X x_1$ and $y_2 = g_X h_X x_1$. To find a general solution, we may add to this values solving $a_1 = \begin{bmatrix} F_{x'} & F_{y'} \end{bmatrix} \begin{bmatrix} \tilde{x}_2 \\ g_X \tilde{x}_2 \end{bmatrix}$. Applying the Schur decomposition, this equals $a_1 = Q^* \begin{bmatrix} S_{11} & S_{12} \\ 0 & S_{22} \end{bmatrix} \begin{bmatrix} (U_{11} + U_{12} g_X)\tilde{x}_2 \\ (U_{21} + U_{22} g_X)\tilde{x}_2 \end{bmatrix} = Q_1^* S_{11}(U_{11} + U_{12} g_X)\tilde{x}_2$ using that $U_{21} + U_{22} g_X = 0$ by construction. Q_1^* is invertible by unitarity of Q, S_{11} is invertible by Γ-regularity of the derivative pair, and $(U_{11} + U_{12} g_X)$ is invertible by Lemma 1 in the main text, so $\tilde{x}_2 = (U_{11} + U_{12} g_X)^{-1} S_{11}^{-1} Q_1 a_1$ and $x_2 = h_X x_1 + \tilde{x}_2$ and $y_2 = g_X h_X x_1 + g_X \tilde{x}_2$. We may then iterate these forward to find a fundamental solution to the next row, and then add $\tilde{x}_3 = (U_{11} + U_{12} g_X)^{-1} S_{11}^{-1} Q_1 a_2$ and $g_X \tilde{x}_3$ to the fundamental solution to find x_3 and y_3. This process can be continued indefinitely, resulting in the solution $s(\{a_t\}_{t=0}^\infty) = \{y_0, \{x_t, y_t\}_{t=1}^\infty\}$ given by $(y_0, x_1) = M^{-1} a_0$, $x_k = h_X^{k-1} x_1 + \sum_{j=1}^{k-1} h_X^{k-1-j} (U_{11} + U_{12} g_X)^{-1} S_{11}^{-1} Q_1 a_j$ for all $k > 1$ and $y_k = g_X x_k$ for

all $k > 0$. We may demonstrate that $s(\{a_t\}_{t=0}^{\infty}) \in \ell^1_{y_0}(\mathcal{H}_1)$ since

$$\sum_{k=1}^{\infty} \|x_k\|_{\mathcal{H}_X} \leq \|(M^{-1}a_0)\| \sum_{k=0}^{\infty} \|h_X^k\| + \|(U_{11}+U_{12}g_X)^{-1}S_{11}^{-1}Q_1\|(\sum_{j=0}^{\infty} \|h_X^j\|) \sum_{k=1}^{\infty} \|a_k\| < \infty$$

since $\|h_X^k\|$ must, since its eigenvalues are bounded in modulus by 1, eventually be of norm less than one, and so $\sum_{j=1}^{\infty} \|h_X^j\|$ is a geometric series with a finite bound. Similarly, $\sum_{k=0}^{\infty} \|y_k\|_{\mathcal{H}_X} \leq \|y_0\| + \|g_X\| \sum_{k=1}^{\infty} \|x_k\|_{\mathcal{H}_X} < \infty$. So, we have shown that $\frac{\partial F^{\infty}(\{x^*,y^*\}_{t=0}^{\infty})}{\partial s}$ is surjective map in $\mathcal{L}(\ell^1_{y_0}(\mathcal{H}_1) \to \ell^1(\mathcal{H}_2))$.

To show invertibility, by the bounded inverse theorem it is sufficient to show that this map is unique. We prove this by contradiction. First, note that completeness of U_{22} and the restriction of the domain of $\frac{\partial F^{\infty}(\{0,0\}_{t=0}^{\infty})}{\partial s}$ to $\ell^1_{y_0}(\mathcal{H}_1)$ rule out other 'recursive' solutions taking the same form as above but for different generalized Schur decompositions. First, completeness ensures a unique solution to $U_{21} + U_{22}g_X = 0$. Further, for any generalized Schur decomposition generated by Cauchy curve Γ generating a Riesz projector which projects onto a subspace other than that generated by Γ equal to the complex unit circle, one of two issues may occur. One possibility is that there exists a spectral subspace corresponding to an element of the spectrum outside of the unit circle, in which case T_{11} has an element of its spectrum with modulus greater than one, and so the sequence $\{x_t\}_{t=0}^{\infty}$ defined by x_0 given, $x_{t+1} = h_X x_t$ does not converge to 0 for each value of x_0, by Gohberg *et al.* (1990, Theorem IV.3.1) and so a fortiori is not in $\ell^1_{y_0}(\mathcal{H}_1)$. The other possibility is that no elements of the spectrum from outside the unit circle are brought inside Γ, but it shrinks so that some spectral subspace corresponding to element of the spectrum with modulus less than one, call it $\mathfrak{M}$, is brought outside of Γ, in which case, U_{22} which was an invertible linear operator from $\mathcal{H}_y$ to $\mathcal{H}_{12}$ must now have range space $\mathcal{H}_{12} \oplus \mathfrak{M}$ and the same domain, so it cannot be invertible. The case where an element of the spectrum has modulus exactly one is ruled out by assumption. So, only one solution of the posited

recursive form exists. Next we must show that no solution other than the recursive one exists.

To show $\frac{\partial F^\infty}{ds}$ is injective, simply show that $\frac{\partial F^\infty}{ds}s = 0$ implies $s = \{y_0, \{x_t, y_t\}_{t=1}^\infty\} = 0$. Suppose not. Then, there is some t such that either $x_t \neq 0$ or $y_t \neq 0$. Consider the first such time t, and suppose that $t \geq 1$. Then $F_{x'}x_t + F_{y'}y_t = 0$ imposes

$$\begin{bmatrix} S_{11}(U_{11}x_t + U_{12}y_t) + S_{12}(U_{21}x_t + U_{22}y_t) \\ S_{22}(U_{21}x_t + U_{22}y_t) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

. If $y_t = g_X x_t$, then the second row cancels and $x_t = 0$. Instead, it must be that $(U_{21}x_t + U_{22}y_t) \in Ker(S_{22})$, where this null space is potentially nontrivial. If S_{22} is complete, $y_t = g_X x_t$, the second row cancels, and $x_t = y_t = 0$. If it is not, there may be some non-zero $z_t \in Null(S_{22})$ such that $U_{21}x_t + U_{22}y_t = z_t$, and so $y_t = U_{22}^{-1}U_{21}x_t + U_{22}^{-1}z_t = g_X x_t + U_{22}^{-1}z_t$, and, plugging this into the first row, obtain $x_t = (U_{11} + U_{12}g_X)^{-1}[U_{11}U_{22}^{-1} + S_{11}^{-1}S_{12}]z_t$, which we write as $x_t = m_Z z_t$ and so the value of x_t, y_t must be of this form. We may then consider what this implies for x_{t+1}, y_{t+1}. We then have $F_x x_t + F_y y_t + F_{x'}x_{t+1} + F_{y'}y_{t+1} = 0$ imposes

$$\begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix} \begin{bmatrix} ((U_{11} + U_{12}g_X)m_Z + U_{12}U_{22}^{-1})z_t \\ z_t \end{bmatrix} = \begin{bmatrix} S_{11} & S_{12} \\ 0 & S_{22} \end{bmatrix} \begin{bmatrix} U_{11}x_{t+1} + U_{12}y_{t+1} \\ U_{21}x_{t+1} + U_{22}y_{t+1} \end{bmatrix}.$$

Values of $(U_{21}x_{t+1} + U_{22}y_{t+1})$ solving the second equation $T_{22}z_t = S_{22}(U_{21}x_{t+1} + U_{22}y_{t+1})$ exist so long as the range space of S_{22} contains $T_{22}z_t$, in which case multiple solutions exist, given by a minimum norm solution $S_{22}^*(S_{22}S_{22}^*)^{-1}T_{22}z_t$ plus some element of the null space of S_{22}, which we may call ϵ_{t+1}, and so for some $\epsilon_{t+1} \in Ker(S_{22})$ $z_{t+1} = S_{22}^*(S_{22}S_{22}^*)^{-1}T_{22}z_t + \epsilon_{t+1}$ is a solution, and so we have $y_{t+1} = g_X x_{t+1} + U_{22}^{-1}z_{t+1}$. We may then solve the first row for x_{t+1} to obtain $x_{t+1} = m_Z z_{t+1} + (h_X m_Z + (U_{11} + U_{12}g_X)^{-1}S_{11}^{-1}T_{11}U_{12}U_{22}^{-1} + T_{12})z_t$. Repeating this pro-

cess indefinitely, $y_{t+k} = g_X x_{t+k} + U_{22}^{-1} z_{t+k}$ where $z_{t+k} = S_{22}^*(S_{22}S_{22}^*)^{-1}T_{22}z_{t+k-1} + \epsilon_{t+k}$ for some $\epsilon_{t+k} \in Ker(S_{22})$, and x_{t+k} is given by $m_Z z_{t+k}$ plus terms in z_{t+i} for i from 0 to $k-1$. We therefore have a sequence which at all t has leading term defined by z_{t+k}, which satisfies $T_{22}^{-1}S_{22}z_{t+k} = z_{t+k-1}$ for all $k \geq 1$. By construction of the generalized Schur decomposition, (T_{22}, S_{22}) has spectrum strictly bounded outside the unit circle, and so the series $T_{22}^{-1}S_{22}z_{t+k} = z_{t+k-1}$ declines exponentially fast toward zero when run backwards in time, and so since $z_t \neq 0$, this means that the series must grow exponentially in norm over time. As a result, this conjectured solution is not in $\ell^1_{y_0}(\mathcal{H}_1)$. It now suffices to rule out the case where the first nonzero element is y_0. If this is the case, the first row requires $y_1 = g_X x_1 + U_{22}^{-1} z_1$ for $z_1 = S_{22}^*(S_{22}S_{22}^*)^{-1}T_{22}y_0 + \epsilon_1$, $\epsilon_1 \in \text{Ker}(S_{22})$ and $x_1 = (U_{11} + U_{12}g_X)^{-1}[S_{11}^{-1}(T_{11}U_{12} + T_{12}U_{22})T_{22}^{-1}S_{22} - S_{11}^{-1}S_{12} - U_{12}U_{22}^{-1}]z_1$, and so we have the same form as before, and so again this results in an explosive series. So, the kernel of $\frac{\partial F^\infty}{ds}$ on $\ell^1_{y_0}(\mathcal{H}_1)$ consists only of 0, and so $\frac{\partial F^\infty}{ds}$ has a unique linear inverse, which is bounded so long as F_X^{-1} is.

Combining these results, the implicit function theorem on Banach spaces applies, and there exists a neighborhood of $x_0 = 0$ around which there exists for every x_0 in this neighborhood a sequence $s(x_0) = \{y_0(x_0), \{x_t(x_0), y_t(x_0)\}_{t=1}^\infty\}$ such that $x_0,\ s(x_0)$ satisfy the equilibrium conditions for all t. Further, this function is continuous and differentiable, with inverse given by $-(\frac{\partial F^\infty}{ds})^{-1}\frac{\partial F^\infty}{dx_0}$. To characterize the result of applying $(\frac{\partial F^\infty}{ds})^{-1}$ to $\frac{\partial F^\infty}{dx_0} = \{F_x, 0, 0, 0, \ldots\}$, we may note that $M^{-1}(-F_x)$ is given by the values of (y_0, x_1) solving $F_x + F_y y_0 + F_{x'}x_1 + F_{y'}g_X x_1 = 0$, which, by generalized Schur decomposability and the invertibility of U_{22}, is solved uniquely by the operators $y_0 = g_X$, $x_1 = h_X$. Applying this and the formula for subsequent values of x_t and y_t, we obtain the recursive form $\frac{\partial x_k(x_0)}{\partial x_0} = h_X^k$ for all $k \geq 1$ and $\frac{\partial y_k(x_0)}{\partial x_0} = g_X h_X^k$ for all $k \geq 0$, as claimed.

The above result showed that a local equilibrium exists for all starting values local to the steady state, and its derivatives follow the given recursive forms. It remains

to show that the equilibrium itself is recursive. To see this, note that the implicit function theorem implies that the function $s(x_0)$ is unique. Further, we have that the function given by the restriction of F^∞ to $t \geq 1$, $F^\infty_{t\geq 1}(x_1, \{y_1, \{x_t, y_t\}^\infty_{t=2}\}) = \{F(x_t + x^*, y_t + y^*, x_{t+1} + x^*, y_{t+1} + y^*, 0)\}^\infty_{t=1}$, has the property that the pair solving x_0, $s(x_0)$ to $F^\infty(x_0, s) = 0$ has restriction to $t \geq 1$ $x_1(x_0)$, $\{y_1(x_0), \{x_t(x_0), y_t(x_0)\}^\infty_{t=2}\}$ which is a solution to $F^\infty_{t\geq 1}(x_1, \{y_1, \{x_t, y_t\}^\infty_{t=2}\}) = 0$. But since $F^\infty_{t\geq 1}$ is identical to $F^\infty = \{F(x_t + x^*, y_t + y^*, x_{t+1} + x^*, y_{t+1} + y^*, 0)\}^\infty_{t=0}$, by the implicit function theorem applied to $F^\infty_{t\geq 1}$, for any x_1 in a neighborhood of x^*, there is a unique $\{\tilde{y}_1(x_1), \{\tilde{x}_t(x_1), \tilde{y}_t(x_1)\}^\infty_{t=2}\}$ such that $F^\infty_{t\geq 1}(x_1, \{y_1(x_1), \{x_t(x_1), y_t(x_1)\}^\infty_{t=2}\}) = 0$, and this is unique and equal to $s(x_1)$. As a result, for $x_1 = x_1(x_0)$, the function $\tilde{x}_2(x_1)$ for $F^\infty_{t\geq 1} = 0$ must equal both $x_2(x_0)$ and $x_1(x_1(x_0))$. Since this may be repeated infinitely often, we have that, for x_0 in a neighborhood of 0, the solution to $F^\infty = 0$ satisfies $x_t(x_0) = x_1(x_1(...(x_1(x_0))))$, i.e., the function x_1 applied t times. So, we are justified in defining the function $h(x)$ (on values of $x \in \mathcal{H}_x$ instead of deviations) as $x_1(x - x^*) + x^*$, which may be applied recursively to find $x_t(x_0)$. Similarly, by uniqueness and recursion $y_t(x_0) = y_0(x_t(x_0))$ for all t, and so we may define $g(x)$ as $y_1(x - x^*) + y^*$. So, for any x_0 in $\mathcal{N}$, there exists a recursion $x_{t+1} = h(x_t)$, $y_t = g(x_t)$ setting $F(x_t, y_t, x_{t+1}, y_{t+1}) = 0$ $\forall t \geq 0$, with the claimed properties.

The claim that this recursion converges back to steady state follows from the fact that $\{x_0, s(x_0)\} \in \ell^1(\mathcal{H}_1)$ by construction. □

The following minor result, possibly not new, ensures that the paths induced by products of sequences of linear operators which converge, such as when the chain rule is applied to find derivatives with respect to initial conditions of an object constructed recursively along a convergent sequence, exhibit the stability properties of their limiting operator. It is used to ensure that the invertibility conditions required for the implicit function theorem hold even along sequences starting away from the steady state.

Lemma .12. *Let $\{A_k\}_{k=1}^{\infty}$ be a sequence of operators in $\mathcal{L}(\mathcal{H} \to \mathcal{H})$ uniformly bounded by $C < \infty$ and converging in operator norm to a fixed limit A such that $\Sigma(A)$ is strictly inside the complex unit circle Γ. Then $\| \prod_{k=1}^{n} A_k \| \to 0$ as $n \to \infty$, and furthermore this convergence is at an exponential rate.*

Remark. Under the conditions of the lemma, $\|A^n\| \to 0$ exponentially by Gohberg *et al.* (1990, Thm. IV.3.1), and for k sufficiently large, $\|A_k^n\| \to 0$ also, so this should be interpreted as saying that the stability of recursively constructed sequences is unaffected by changes along the sequence which are eventually negligible. Note that this statement is trivial also for sequences of normal operators, for which the operator norm is given by the spectral radius and so $\| \prod_{k=1}^{n} A_k \| \leq \prod_{k=1}^{n} \|A_k\| \searrow 0$.

Proof. Convergence is given by a 'blocking'-type argument. While the elements of $\prod_{k=1}^{n} A_k$ converge to A, this convergence does not imply that the product of the element converges to the product of the limit. Instead, the sequence may be separated into blocks $i = 1 \dots m$ of finite length J_i, which, since a convergent sequence is Cauchy, each block contains a set of elements of diameter going to 0. Since asymptotically the spectrum of A_k is strictly inside the unit circle, for large enough i and for long enough J_i, the j_i-fold product of any element within a block has norm bounded away from 1, and for small enough diameter this is still true for the product over the block itself. Since the norm of the product is bounded by the product of the norms of the blocks, each less than 1, it decays exponentially as more blocks are added.

Denote the j^{th} element in block i as A_{i_j}. For any set of blocks, obtain $\| \prod_{k=1}^{n} A_k \| \leq \prod_{i=1}^{m} \| \prod_{j=1}^{J_i} A_{i_j} \|$. By convergence of A_k in operator norm, their spectra also converge, and so for any small enough $\epsilon > 0$, $\exists K_1$ such that$\forall k > K_1$, $\sup |\Sigma(A_k)| < \sup |\Sigma(A_k)| + \epsilon < 1$. By Gohberg *et al.* (1990, Thm. IV.3.1) and uniform boundedness of the A_k, this implies that for any $\delta > 0$, for $k > K_1$, there exists some $J(\delta)$ such that $\|A_k^{J(\delta)}\| < \delta$ uniformly over k. For any block i, $\| \prod_{j=1}^{J_i} A_{i_j} - A_{i_{J_i}}^{J_i} \| \leq C^{J_i - 1} \sum_{i=1}^{J_i - 1} \|A_{i_j} - A_{i_{J_i}}\|$. Since A_k is a Cauchy sequence, for any $\epsilon_2 > 0$ there exists

K_2 such that for all $k_1, k_2 > K_2$ $\|A_{k_1} - A_{k_2}\| < \epsilon_2$. For some $\rho < 1$ we may choose δ, ϵ_2 such that for $k > \max\{K_1(\epsilon), K_2(\delta)\}$, $\| \prod_{j=1}^{J(\delta)} A_{k+j}\| \leq C^{J(\delta)-1}(J(\delta)-1)\epsilon_2 + \delta = \rho$. Let $m = \lfloor (n - \max\{K_1(\epsilon), K_2(\delta)\})/J(\delta) \rfloor + 1$, $J_i = J(\delta)$ for all $i \geq 2$ and $J_1 = n - (m-1)J(\delta)$, i.e., the first $\max\{K_1(\epsilon), K_2(\delta)\}$ elements, plus the remainder of n after adding the largest feasible integer number of blocks of size $J(\delta)$ after that. Then by uniform boundedness and the fact that the remainder contains no more than $J(\delta)$ blocks, $\| \prod_{j=1}^{J_1} A_{1j}\|$ is bounded by a constant C_2, and so for $n > \max\{K_1(\epsilon), K_2(\delta)\})$, $\| \prod_{k=1}^{n} A_k\| \leq C_2 \rho^{m-1}$ which decays to 0 exponentially in m, and so in n as well. $\square$

Theorem .5

Proof. Consider the map $\mathcal{M}(g, h, \sigma)(x) = \mathbb{E}F(x, g(x), h(x) + \sigma\eta z', g(h(x) + \sigma\eta z'), \sigma)$ mapping $\mathcal{B}_1$ to $\mathcal{W}^{2,\infty}(\mathcal{N} \to \mathcal{H}_2)$, the space of operators from $\mathcal{H}_x$ to $\mathcal{H}_2$ bounded on $\mathcal{N} \subset \mathcal{H}_x$ with bounded first and second Fréchet derivatives on the same region, where $\mathcal{B}_1$ is the Cartesian product of $\mathcal{W}^{2,\infty}(\mathcal{N} \to \mathcal{H}_y)$, the space of operators from $\mathcal{H}_x$ to $\mathcal{H}_y$ bounded and with bounded first and second derivatives on $\mathcal{N} \subset \mathcal{H}_x$, $\mathcal{W}^{1,\infty}(\mathcal{N} \to \mathcal{H}_x)$, the Banach space of operators from $\mathcal{H}_x$ to $\mathcal{H}_x$ bounded and with bounded derivatives on $\mathcal{N} \subset \mathcal{H}_x$, and $[-\epsilon, \epsilon] \subset \mathbb{R}$. We hope to solve $\mathcal{M}(g, h, \sigma)(x) = 0$ implicitly for g and h as a function of σ (where 0 is the operator mapping all $x \in \mathcal{H}_x$ to $0 \in \mathcal{H}_2$). To do this, we must show $\frac{\partial}{\partial(g,h)}\mathcal{M}(g^*, h^*, 0)(x)$ is invertible, where g^*, h^* are the deterministic operators found in the previous theorem. To do this, it is sufficient to show that for every $\psi(x) \in \mathcal{W}^{2,\infty}(\mathcal{N} \to \mathcal{H}_2)$, there exists (unique) $\phi(x) = (\phi_g(x), \phi_h(x)) \in \mathcal{W}^{2,\infty}(\mathcal{N} \to \mathcal{H}_y) \times \mathcal{W}^{2,\infty}(\mathcal{N} \to \mathcal{H}_x)$ such that $\frac{\partial}{\partial(g,h)}\mathcal{M}(g^*, h^*, 0)(x)[(\phi_g(x), \phi_h(x))] = \psi(x)$. Applying the definition of $\mathcal{M}$, noting that at $\sigma = 0$ F is deterministic and so the expectation disappears, and rearranging,

obtain

$$F_y(x, g^*(x), h^*(x), g^*(h^*(x)))[\phi_g(x)] + F_{x'}(x, g^*(x), h^*(x), g^*(h^*(x)))[\phi_h(x)] + \\ = \psi(x) - F_{y'}(x, g^*(x), h^*(x), g^*(h^*(x)))\frac{\partial}{\partial x} g^*(h^*(x))[\phi_h(h^*(x))] \\ -F_{y'}(x, g^*(x), h^*(x), g^*(h^*(x)))[\phi_g(h^*(x))] \quad (12)$$

To simplify notation, we write this as $G_1(x)\phi(x) = \psi(x) - G_2(x)\phi(h^*(x))$. At $x = x^*$, this reduces to $F_y[\phi_g(x^*)]+F_{x'}[\phi_h(x^*)] = \psi(x^*)-F_{y'}[\phi_g(x^*)]-F_{y'}g_X[\phi_h(x^*)]$, or $G_1(x^*)\phi(x^*) = \psi(x^*) - G_2(x^*)\phi(x^*)$. By assumption, $G_1(x^*) = [F_y, F_{x'}]$ is invertible, and so this equals $\phi(x^*) + G_1^{-1}(x^*)G_2(x^*)\phi(x^*) = G_1^{-1}(x^*)\psi(x^*)$. This has a unique solution if and only if $-1 \in \rho(G_1^{-1}(x^*)G_2(x^*))$ so $I + G_1^{-1}(x^*)G_2(x^*)$ is invertible, which holds because we assume that $\Sigma(G_1^{-1}(x^*)G_2(x^*))$ is inside the complex unit circle. Note that this is stronger than necessary for a unique solution at x^*. However, away from x^*, the system becomes no longer time reversible, and components of the spectrum outside the unit circle correspond to iterating x backward in time along $h^*(x)$. As $h^*(x)$ generally has unbounded or even nonexistent inverse, these do not generate bounded solutions for $\phi(x)$ away from x^*.

We next seek the value of $\phi(x)$ away from x^* by using continuity and solving forward. Since $G_1^{-1}(x^*)$ is assumed bounded and $G_2(x^*)$ is bounded since $F_{y'}$ and g_X are, the resolvent set of $G_1^{-1}(x^*)G_2(x^*)$ is open. By continuity of the derivatives of F and g^* with respect to x we therefore have that there is a neighborhood of x^* on which $G_1^{-1}(x)G_2(x)$ is bounded and, by continuity of the spectrum of bounded operators, has spectrum inside the unit circle. If so desired, we may restrict this neighborhood so that the spectrum is bounded away from the unit circle. Since $h^*(x)$ is stable and continuous in a neighborhood of x^*, there exists a neighborhood $\mathcal{U}_0$ contained in the above neighborhood and $\mathcal{N}$ such that $h^*(x) \in \mathcal{U}_0 \; \forall x \in \mathcal{U}_0$. As a result, for any $x \in \mathcal{U}_0$, $\phi(x) = G_1^{-1}(x)\psi(x)-G_1^{-1}(x)G_2(x)\phi(h^*(x))$, and we may iterate this forward to obtain

$\phi(h^*(x)) = G_1^{-1}(h^*(x))\psi(h^*(x)) - G_1^{-1}(h^*(x))G_2(h^*(x))\phi(h^*(h^*(x)))$ and continue to iterate to find an expression for $\phi(x)$ in terms of an infinite series. Formally we have

$$\phi(x) = \sum_{k=0}^{\infty} \gamma_k \psi(x_k)$$

where x_k and γ_k are defined recursively by $x_0 = x$, $x_{k+1} = h^*(x_k)$, $\gamma_1 = G_1^{-1}(x_0)$ $\gamma_k = -\gamma_{k-1}G_2(x_{k-1})G_1^{-1}(x_k)$. Since $\psi \in \mathcal{W}^{1,\infty}(\mathcal{N} \to \mathcal{H}_2)$, $\psi(x_k)$ is uniformly bounded by a constant on $\mathcal{U}$. By stability $\|x_k - x^*\| \to 0$, and by continuity of $G_1^{-1}(x)G_2(x)$, $G_1^{-1}(x_k)G_2(x_k) \to G_1^{-1}(x^*)G_2(x^*)$ with spectrum bounded away from the unit circle, the operator norm of γ_k must eventually decline exponentially by .12. As a result, the series converges in $\mathcal{H}_2$-norm. Further, since this boundedness is true for all x, and the uniform boundedness over x implies that the bound in.12 can likewise be made uniform over x, $\sup_{x\in\mathcal{U}} \|\phi(x)\|_{\mathcal{H}_2} < \infty$, and thus $\phi(.)$ is bounded on $\mathcal{U}_0$.

To show that $\phi(.) \in \mathcal{W}^{2,\infty}(\mathcal{U} \to \mathcal{H}_2)$, we must first show that its first derivatives are bounded on some set. To see this, first note that $\psi(x) \in \mathcal{W}^{2,\infty}(\mathcal{N} \to \mathcal{H}_2)$ and so has uniformly bounded derivatives on $\mathcal{U}_0$. Next, note that, by assumption, the derivatives of F with respect to x, y, x' and y' themselves have uniformly bounded derivatives, and, by the previous theorem extended to apply the implicit function theorem to an operator which is twice continuously differentiable, h^* and g^* are themselves twice continuously differentiable and so have uniformly bounded derivatives on a neighborhood of x^*. As a result, $G_1^{-1}(x)$ and $G_2(x)$ have derivatives which are uniformly bounded in a neighborhood of x^*. By the stability of h^* on $\mathcal{U}_0$ and the uniform boundedness of its derivatives, writing $\frac{d}{dx}x_k(x) = \prod_{i=1}^{k} \frac{d}{dx}h^*(h^{*k-i}(x))$ by the chain rule, and noting that $h^{*k-i}(x) \to x^*$ and so by continuity $\frac{d}{dx}h^*(h^{*k-i}(x)) \to h_X$, which has spectrum inside the unit circle by assumption, Lemma .12 applies and $\|\frac{d}{dx}x_k(x)\| \to 0$ exponentially. Denote the set over which all of these properties hold as $\mathcal{U}$. Applying the product rule, obtain $\frac{d}{dx}\phi(x) = \sum_{k=0}^{\infty}((\frac{d}{dx}\gamma_k)\psi(x_k) + \gamma_k\frac{d}{dx}\psi(x_k))$. The sum over the second term converges uniformly in x by the convergence of γ_k and the fact that

$\frac{d}{dx}\psi(x)$ is uniformly bounded. By .12, products of j instances of $G_1^{-1}(x)G_2(x)$ evaluated at different points are eventually bounded by λ^j for a constant $\lambda < 1$. By the decay of the derivative of x_k to 0, and the uniform boundedness of G_1^{-1}, G_2 and their derivatives, for k larger than $\bar{k}$, $\|\frac{d}{dx}[G_1^{-1}(x_k)G_2(x_k)]\| \leq \lambda$ also. Applying the product rule to the recursive formulation of γ_k, the first term of the derivative is bounded by a constant times $\sup_{x\in\mathcal{U}}\|\psi(x)\|_{\mathcal{H}_2}$ times $\sum_{k=0}^{\infty}(\sum_{i=0}^{\bar{k}}\|\frac{d}{dx}[G_1^{-1}(x)G_2(x)]\|^i\lambda^{i-\bar{k}} + \sum_{i=\bar{k}+1}^{k}\lambda^i)$. The first part is bounded by a constant times a convergent geometric series, the second is bounded by a constant times $\sum_{k=0}^{\infty} k\lambda^i$ which is also a convergent series. As a result, the series converges uniformly over $x \in \mathcal{U}$, and so $\frac{d}{dx}\phi(.)$ is bounded.

To show boundedness of second derivatives, essentially similar procedures can be followed. By the product rule, $\frac{d^2}{dx^2}\phi(x) = \sum_{k=0}^{\infty}((\frac{d^2}{dx^2}\gamma_k)\psi(x_k) + \gamma_k\frac{d^2}{dx^2}\psi(x_k) + 2(\frac{d}{dx}\gamma_k)\frac{d}{dx}\psi(x_k))$. Since $\frac{d}{dx}\psi(x)$ is uniformly bounded, the summation of the last term is bounded on an appropriate neighborhood by the exact procedures used to show $\sum_{k=0}^{\infty}(\frac{d}{dx}\gamma_k)\psi(x_k)$ is bounded, and the second is bounded by the assumption that $\psi(x) \in \mathcal{W}^{2,\infty}(\mathcal{N} \to \mathcal{H}_2)$ and so has uniformly bounded second derivatives. To show the first part, we must control the second derivatives of the recursive construction of γ_k. First note that for k large enough, $\|\frac{d^2}{dx^2}x_k(x)\|_{\mathcal{L}(\mathcal{U}\to\mathcal{L}(\mathcal{U}\to\mathcal{H}_x))} \to 0$ exponentially. To see this, note

$$\begin{aligned}\frac{d^2}{dx^2}x_k(x)[a][b] &= \frac{d}{dx}[\prod_{i=1}^{k}\frac{d}{dx}h^*(h^{*k-i}(x))[a]][b] \\ &= \sum_{j=1}^{k}\prod_{i=1}^{j-1}\frac{d}{dx}h^*(h^{*k-i}(x)) \cdot \frac{d^2}{dx^2}h^*(h^{*k-j}(x))[\prod_{i=j+1}^{k}\frac{d}{dx}h^*(h^{*k-i}(x))[a]][b]\end{aligned}$$

By the uniform boundedness of $\frac{d^2}{dx^2}h^*()$ in a neighborhood of x^* which applies by the implicit function theorem used to construct it extended to apply to a three times continuously differentiable operator and by the convergence of iterated first derivatives by the construction of the blocks in .12, this is bounded in operator norm by k

times an exponentially decaying quantity in k, and so itself is exponentially decaying. Similarly, by three times continuous differentiability of F, the second derivatives of G_1^{-1} and G_2 are also uniformly bounded on a neighborhood of x^*, and so using the exponential convergence of $\frac{d}{dx}x_k(x)$ and $\frac{d^2}{dx^2}x_k(x)$, the product rule and the chain rule, $\|\frac{d^2}{dx^2}[G_1^{-1}(x_k)G_2(x_k)]\| \to 0$ exponentially also. So, by the product rule again, $\frac{d^2}{dx^2}\gamma_k$ is the sum of k exponentially decaying components and so also declines exponentially in k in operator norm. Uniform boundedness of $\|\psi_k(.)\|$ and the continuity of the second derivative of γ_k then imply the convergence of the geometric sum $\sum_{k=0}^{\infty}((\frac{d^2}{dx^2}\gamma_k)\psi(x_k)$ uniformly over x in a neighborhood of x^*. As a result, $\frac{d^2}{dx^2}\phi(x)$ is bounded and so $\phi(.) \in \mathcal{W}^{2,\infty}(\mathcal{U} \to \mathcal{H}_2)$.

So, restricting all operators on $\mathcal{N}$ to $\mathcal{U}$, we see that $\frac{\partial}{\partial(g,h)}\mathcal{M}(g^*, h^*, 0)(x)$ has a bounded inverse on $\mathcal{W}^{2,\infty}(\mathcal{U} \to \mathcal{H}_2)$.

Continuity of $\mathcal{M}(g, h, \sigma)(x)$ and continuity of $\frac{\partial}{\partial(g,h)}\mathcal{M}(g, h, \sigma)(x)$ with respect to g, h, and σ in a neighborhood of g^*, h^*, 0 are guaranteed by the bounded support condition on z, continuous differentiability of $\mathbb{E}F$ with respect to its arguments and by the twice continuous differentiability of g, which holds at g^* as a corollary of the implicit function theorem used to construct it, extended to three times continuously differentiable F and locally in a neighborhood of g^* since we consider only operators in $\mathcal{W}^{2,\infty}(\mathcal{U} \to \mathcal{H}_y)$. To see the importance of the bounded support condition, note that σ enters $g(h(x)+\sigma\eta z')$ and so to ensure that $x' \in \mathcal{U}$ for all $x \in \mathcal{U}$, it is sufficient, since $h(x) \in \mathcal{U}$ and $\mathcal{U}$ is open, there exists a radius ϵ_s such that $\|\sigma\eta z'\| < \epsilon_s$, which is true if $\|z'\| < \infty$ for σ sufficiently small. As $g(x)$ and $g_x(x)$, into which x' enters in $\frac{\partial}{\partial(g,h)}\mathcal{M}(g, h, \sigma)(x)$, are guaranteed to be bounded and continuous only over a set $\mathcal{U}$, allowing z' to take unbounded support would result in the possibility of unbounded changes for small changes in σ if no further conditions were imposed on g and g_x and so could violate continuity. Combining the above conditions, the implicit function theorem in Banach space implies that there exists a neighborhood $(-\epsilon, \epsilon)$ of σ around

0 in which there exist continuous, differentiable functions $g(.,\sigma)$, $h(.\sigma)$ from $(-\epsilon,\epsilon) \rightarrow$ $\mathcal{W}^{2,\infty}(\mathcal{U} \rightarrow \mathcal{H}_y) \times \mathcal{W}^{2,\infty}(\mathcal{U} \rightarrow \mathcal{H}_x)$ satisfying $\mathcal{M}(g(.\sigma), h(.,\sigma), \sigma)(x) = 0$. $\square$

Appendix C: Proofs of Propositions in Main Text

Chapter 1 Proofs

Section 1.2 Proof

Proof. of (1.1). Suppose $h(x,z) := h(x,\sigma) + \sigma z$ is a measurable function from $(\mathcal{B}_x \times \mathcal{B}_z, \Sigma_x \otimes \Sigma_z)$, the product space of $\mathcal{B}_x \times \mathcal{B}_z$ equipped with a product sigma field, to $(\mathcal{B}_x, \Sigma_x)$. We want conditions on the space, the function, and the sigma fields such that it induces a measurable stochastic process on the product space of $\mathcal{B}_x$. We may assume z is drawn independently of x according to measure μ^z on $(\mathcal{B}_z, \Sigma_z)$, and may ask for the initial distribution of x to be given by μ^x. For each x, we can define the pushforward measure on $(\mathcal{B}_x, \Sigma_x)$ by $\mu_x^{x'}(f(x')) := \mu^z(f(h(x,\sigma) + \sigma z)$ for any $f \in \mathcal{M}^+(\mathcal{B}_x, \Sigma_x, \bar{\mathbb{R}}, \mathscr{B}(\bar{\mathbb{R}}))$ nonnegative measurable functions from x to the real line equipped with the Borel sigma field. If the family $(\mu_x^{x'})_x$ of measures satisfies $x \to \mu_x^{x'}(A)$ is a measurable function from $(\mathcal{B}_x, \Sigma_x) \to (\bar{\mathbb{R}}, \mathscr{B}(\bar{\mathbb{R}}))$ for any $A \in \Sigma_x$, then this is a probability kernel and by, e.g, the Ionescu Tulcea extension theorem, the family induces a measurable stochastic process for x_t on the countable product space $\otimes_{t=1}^{\infty}(\mathcal{B}_x, \Sigma_x)$.

To show measurability of the family of measures $(\mu_x^{x'})_x$, consider a λ-class argument. The measure μ^z maps the class of measurable rectangles $\{x \in A^1, z \in A^2\}$

for $A^1 \in \Sigma_x, A^2 \in \Sigma_z$ to (nonnegative multiples of) indicators of sets Σ_x, which are therefore measurable. The class of measurable rectangles generates the product sigma field $\Sigma_x \otimes \Sigma_z$ and is stable under pairwise intersections. The class of bounded nonnegative functions $f(x,z) \in \mathcal{M}^+(\mathcal{B}_x \times \mathcal{B}_z, \Sigma_x \otimes \Sigma_z, \bar{\mathbb{R}}, \mathscr{B}(\bar{\mathbb{R}}))$ such that $\mu^z f(x,z)$ is $(\mathcal{B}_x, \Sigma_x)$ measurable can be shown to form a λ-cone (Pollard, 2002, 2.11 Def. <43>) and so by these facts (Pollard, 2002, 2.11 Lemma <44>), μ^z maps $\mathcal{M}^+(\mathcal{B}_x \times \mathcal{B}_z, \Sigma_x \otimes \Sigma_z, \bar{\mathbb{R}}, \mathscr{B}(\bar{\mathbb{R}}))$ to $\mathcal{M}^+(\mathcal{B}_x, \Sigma_x, \bar{\mathbb{R}}, \mathscr{B}(\bar{\mathbb{R}}))$. In particular, let $h(x,z)$ be $\mathcal{B}_x \times \mathcal{B}_z, \Sigma_x \otimes \Sigma_z \to \mathcal{B}_x, \Sigma_x$ measurable, then $\mu^z(f(h(x,z))$ is $(\mathcal{B}_x, \Sigma_x)$ measurable for any $f \in \mathcal{M}^+(\mathcal{B}_x, \Sigma_x, \bar{\mathbb{R}}, \mathscr{B}(\bar{\mathbb{R}}))$ and in particular, $x \to \mu_x^{x'}(A)$ is a measurable function from $(\mathcal{B}_x, \Sigma_x) \to (\bar{\mathbb{R}}, \mathscr{B}(\bar{\mathbb{R}}))$ for any $A \in \Sigma_x$. As a result, $(\mu_x^{x'})_x$ is a probability kernel.

To construct a measurable stochastic process, consider the i.i.d. sequence $\{z_t\}_{t=0}^{\infty}$ such that z_t each have identical marginal measure μ_t^z and, beginning with initial measure μ^x, construct the sequence of probability kernels on $\otimes_{t=1}^{\infty}(\mathcal{B}_x, \Sigma_x)$ by iterating the identical kernels defined by $\mu_{xt}^{x'}(f(x')) := \mu_t^z(f(h(x,z))$. This generates a sequence $x_0 \sim \mu_0^x$, $x_t = h(x_{t-1}, z_t)$. By the Ionescu Tulcea extension theorem, the sequence of kernels induces a measurable stochastic process on the countable product space $\otimes_{t=1}^{\infty}(\mathcal{B}_x, \Sigma_x)$ with finite-dimensional distributions generated by the iterated probability kernels. Note that the only assumptions made on $(\mathcal{B}_x, \Sigma_x)$, $(\mathcal{B}_z, \Sigma_z)$ and $h(x,z)$ are that $h(x,z)$ is jointly measurable from the product sigma field over x and z to the sigma field over x. In particular, because the probability kernel was constructed explicitly, no topological assumptions needed to be made on the spaces or sigma fields, as are usually required to invoke the Kolmogorov extension theorem. This permits, among other constructions, the use of nonseparable function spaces or non-Borel sigma fields, which may alleviate some difficulties when working in infinite-dimensional space.

By measurability of $g(x,\sigma)$ and F, the measurability of the probability ker-

nels defining the conditional distribution of the random variables $y_t = g(x_t, \sigma)$ and $F(x_t, g(x_t, \sigma), h(x_t, \sigma) + \sigma z_{t+1}, g(h(x_t, \sigma) + \sigma z_{t+1}, \sigma)$ given x and from there the corresponding stochastic processes can be established in an analogous fashion, ensuring that (x_t, y_t) is product measurable and $\mathbb{E}F(x, g(x, \sigma), h(x, \sigma) + \sigma\eta z', g(h(x, \sigma) + \sigma\eta z', \sigma), \sigma)$ coincides with the conditional expectation of $F(x_t, g(x_t, \sigma), h(x_t, \sigma) + \sigma\eta z_{t+1}, g(h(x_t, \sigma) + \sigma\eta z_{t+1}, \sigma), \sigma)$ at time t given $x_t = x$, as claimed. □

Section 1.4 Proofs

Proof. Of Theorem (1.1). The proof proceeds in two steps: first, showing that the generalized Schur decomposition is continuous with respect to the approximation, and then showing the policy operators are continuous in the generalized Schur decomposition.

First, note that

$$\|(\tilde{B}^K, \tilde{A}^K) + (I - \pi^K)(B_I, A_I)(I - \pi^K) - (B, A)\|_{\mathcal{B}} \leq$$
$$\|(\tilde{B}^K, \tilde{A}^K) - (B^K, A^K)\|_{\mathcal{B}} + \|(B^K, A^K) + (I - \pi^K)(B_I, A_I)(I - \pi^K) - (B, A)\|_{\mathcal{B}} =$$
$$\|(\tilde{B}^K, \tilde{A}^K) - (B^K, A^K)\|_{\mathcal{B}} +$$
$$\|(I-\pi^K)(B_I, A_I)(I-\pi^K) - (I-\pi^K)(B, A)(I-\pi^K) - (I-\pi^K)(B, A)\pi^K - \pi^K(B, A)(I-\pi^K)\|_{\mathcal{B}} =$$
$$\|(\tilde{B}^K, \tilde{A}^K) - (B^K, A^K)\|_{\mathcal{B}} +$$
$$\| - (I - \pi^K)(B_C, A_C)(I - \pi^K) - (I - \pi^K)(B_C, A_C)\pi^K - \pi^K(B_C, A_C)(I - \pi^K)\|_{\mathcal{B}}$$
$$= \|(\tilde{B}^K, \tilde{A}^K) - (B^K, A^K)\|_{\mathcal{B}} + \|\pi^K(B_C, A_C)\pi^K - (B_C, A_C)\|_{\mathcal{B}} \leq \zeta_K + \eta_K$$

where the third expression follows from the decomposition of $(B, A) = (B_I, A_I) + (B_{C,}, A_C)$, and the construction of π^K so that $(I-\pi^K)(B_I, A_I)\pi^K = 0$ and $\pi^K(B_I, A_I)(I-\pi^K) = 0$.

The consistency of the approximation of (B, A) implies consistency of the components of the Schur decomposition by (.2) and (.3) and the bound on dif(B, A). Note that the generalized Schur decomposition of $(\tilde{B}^K, \tilde{A}^K)$ separately is equivalent to (one ordering of) the generalized Schur decomposition of their sum. More precisely,

$$\begin{bmatrix} \tilde{Q}_1^{*K} & 0 & ,\tilde{Q}_2^{*K} & 0 \\ 0 & Q_1^{*I} & 0 & Q_2^{*I} \end{bmatrix} \begin{bmatrix} \tilde{T}_{11}^K & 0 & \tilde{T}_{12}^K & 0 \\ 0 & T_{11}^I & 0 & T_{12}^I \\ 0 & 0 & \tilde{T}_{22}^K & 0 \\ 0 & 0 & 0 & T_{22}^I \end{bmatrix}, \begin{bmatrix} \tilde{S}_{11}^K & 0 & \tilde{S}_{12}^K & 0 \\ 0 & S_{11}^I & 0 & S_{12}^I \\ 0 & 0 & \tilde{S}_{22}^K & 0 \\ 0 & 0 & 0 & S_{22}^I \end{bmatrix} * \begin{bmatrix} \tilde{U}_{11}^K & 0 & \tilde{U}_{12}^K & 0 \\ 0 & U_{11}^I & 0 & U_{12}^I \\ \tilde{U}_{21}^K & 0 & \tilde{U}_{22}^K & 0 \\ 0 & U_{21}^I & 0 & U_{22}^I \end{bmatrix}$$

where an I superscript indicates a component corresponding to the Schur decomposition on Ker π^K of (B^I, A^I), is a generalized Schur decomposition of $(\tilde{B}^K, \tilde{A}^K) + (I - \pi^K)(B_I, A_I)(I - \pi^K)$ corresponding to curve Γ. Note that by operator norm convergence, for sufficiently large K, $\sigma_{\min}(\lambda\tilde{B}^K - \tilde{A}^K + (I - \pi^K)(\lambda B_I -, A_I)(I - \pi^K)) \geq \sigma_{\min}(\lambda B - A) - 2(\zeta_K + \eta_K) > 0$ uniformly in $\lambda \in \Gamma$ by Weyl's inequality and the compactness of Γ and so $(\tilde{B}^K, \tilde{A}^K) + (I - \pi^K)(B_I, A_I)(I - \pi^K)$ is Γ-regular and so the generalized Schur decomposition described exists.

To bound $\|g_K - g_X\|$, note

$$g_K = -(\tilde{U}_{22}^K)^{-1}\tilde{U}_{21} - (U_{22}^I)^{-1}U_{21}^I = -\begin{pmatrix} \tilde{U}_{22}^K & 0 \\ 0 & U_{22}^I \end{pmatrix}^{-1} \begin{pmatrix} \tilde{U}_{21}^K & 0 \\ 0 & U_{21}^I \end{pmatrix} = -(\tilde{U}_{22})^{-1}\tilde{U}_{21}.$$

By (.2),

$$\begin{aligned}\|\tilde{U}_1 - U_1\| &\leq \|U_1\|\|I - (I + P^*P)^{-\frac{1}{2}}\| + \|P\| \left\|U_2(I + P^*P)^{-\frac{1}{2}}\right\| \\ &\leq C\|P\| + o(\|P\|) \leq C\frac{2(\zeta_K + \eta_K)}{\delta}\end{aligned}$$

for some constant $C < 2 + \epsilon$ for any ϵ, for K sufficiently large, where $\delta > 0$ by the assumption that dif(B,A)>0. As a result, by invertibility of U_{22}, Weyl's inequality, and the triangle inequality, $\left\|-\tilde{U}_{22}^{-1}\tilde{U}_{21} + U_{22}^{-1}U_{21}\right\| \leq C\frac{2(\zeta_K+\eta_K)}{\delta}$ for some constant C for K large enough.

Similarly, we have

$$\begin{aligned}h_K = &\left(\begin{pmatrix} \tilde{U}_{11}^K & 0 \\ 0 & U_{11}^I \end{pmatrix} - \begin{pmatrix} \tilde{U}_{12}^K & 0 \\ 0 & U_{12}^I \end{pmatrix}\begin{pmatrix} \tilde{U}_{22}^K & 0 \\ 0 & U_{22}^I \end{pmatrix}^{-1}\begin{pmatrix} \tilde{U}_{21}^K & 0 \\ 0 & U_{21}^I \end{pmatrix}\right)^{-1} * \\ &\begin{pmatrix} \tilde{S}_{11}^K & 0 \\ 0 & S_{11}^I \end{pmatrix}^{-1}\begin{pmatrix} \tilde{T}_{11}^K & 0 \\ 0 & T_{11}^I \end{pmatrix} * \\ &\left(\begin{pmatrix} \tilde{U}_{11}^K & 0 \\ 0 & U_{11}^I \end{pmatrix} - \begin{pmatrix} \tilde{U}_{12}^K & 0 \\ 0 & U_{12}^I \end{pmatrix}\begin{pmatrix} \tilde{U}_{22}^K & 0 \\ 0 & U_{22}^I \end{pmatrix}^{-1}\begin{pmatrix} \tilde{U}_{21}^K & 0 \\ 0 & U_{21}^I \end{pmatrix}\right) \\ &= (\tilde{U}_{11} + \tilde{U}_{12}g_K)^{-1}(\tilde{S}_{11})^{-1}\tilde{T}_{11}(\tilde{U}_{11} + \tilde{U}_{12}g_K)\end{aligned}$$

Applying the triangle inequality, (.2), (.3), and convergence of g_K, this implies that for some constant C, for K large enough, $\|h_K - h_x\|_{op} \leq C\frac{\zeta_K+\eta_K}{\delta}$, as claimed. □

A demonstration that for appropriately smooth functions wavelet representations provide the necessary error control to ensure consistency follows from some standard estimates regarding wavelet coefficients.

Proof. of Theorem (1.2). First we demonstrate bounds on η_K, the error induced by truncating to a K term wavelet series, using results on wavelet coefficients and oper-

ator norm bounds from Johnstone (2013), then bounds on ζ_K, the error induced by calculating the inner products with the wavelet basis by quadrature using quadrature error estimates from Beylkin *et al.* (1991).

First, denoting the blocks of (B_C, A_C) as $K_{r,ij}$, $\max\{\left\|B_C^K - B_C\right\|_{op}, \left\|A_C^K - A_C\right\|_{op}\} \leq J\max_{r,i,j} \left\|K_{r,ij}^K - K_{r,ij}\right\|_{op}$ by definition of operator norm. Because an orthonormal basis is used, $\pi^K(B_I, A_I)\pi^K$ is simply expressed in terms of identity matrices on this space, and so can be evaluated exactly.

The projection of $K_{r,ij}$ onto the space of the first $K_i \times K_j$ wavelet coefficients can be expressed using the inner product with the tensor product over the first $K_i \times K_j$ orthonormal basis functions $\{\phi_k\}_{k=1}^{K_i}$ and $\{\phi_k\}_{k=1}^{K_j}$ as

$$\pi^{K_i} K_{r,ij} \pi^{K_j}[f(y)] = \sum_{k=1}^{K_j} \sum_{l=1}^{K_i} \langle K_{r,ij}(x,y), \phi_k(x)\phi_l(y)\rangle \langle \phi_k(y), [f(y)]\rangle \phi_l(x)$$

$= \int \hat{K}_{r,ij}(x,y) f(y) dy$ where $\hat{K}_{r,ij}(x,y) = \sum_{k=1}^{K_j} \sum_{l=1}^{K_i} \langle K_{r,ij}(x,y), \phi_k(x)\phi_l(y)\rangle \phi_k(y)\phi_l(x)$ is the $K_i \times K_j$ term projection of the kernel of the integral operator onto the wavelet basis. Since $K_{r,ij}(x,y) \in \Lambda^{\alpha_{r,ij}}([0,1]^{d_i} \times [0,1]^{d_j})$ and ϕ_k are a standard wavelet basis, we can use norm bounds to control the error in this projection. Sup norm bounds available in Chen & Christensen (2015), show that under the $\alpha_{r,ij}-$Hölder assumption,

$$\left\|\hat{K}_{r,ij}(x,y) - K_{r,ij}(x,y)\right\|_{L^\infty([0,1]^{d_i} \times [0,1]^{d_j})} = O((K_i K_j)^{-\alpha_{r,ij}/(d_i+d_j)})$$

when wavelets satisfying (1.4.3) are used. In particular, adapting the proof of their Lemma 2.4, letting

$$\ell_{K_i K_j} = \sup_{f \in L^\infty([0,1]^{d_i} \times [0,1]^{d_j})} \left\| \sum_{k=1}^{K_j} \sum_{l=1}^{K_i} \langle f(x,y), \phi_k(x)\phi_l(y)\rangle \phi_k(y)\phi_l(x) \right\|_{L^\infty} / \left\|f(x,y)\right\|_{L^\infty}$$

be the Lebesgue constant for the tensor product wavelet basis

$$\left\|\hat{K}_{r,ij}(x,y) - K_{r,ij}(x,y)\right\|_{L^\infty([0,1]^{d_i}\times[0,1]^{d_j})} \;\leq\; (1 \,+\, \ell_{K_iK_j})O((K_iK_j)^{-\alpha_{r,ij}/(d_i+d_j)}),$$

and by their Theorem 5.1 applied in the case of uniform density, $\ell_{K_iK_j}$ is bounded uniformly in K_i and K_j.

By compactness of the domain, we have

$$\begin{aligned}\int \left|\hat{K}_{r,ij}(x,y) - K_{r,ij}(x,y)\right| dx &\leq C \left\|\hat{K}_{r,ij}(x,y) - K_{r,ij}(x,y)\right\|_{L^\infty([0,1]^{d_i}\times[0,1]^{d_j})} \\ \int \left|\hat{K}_{r,ij}(x,y) - K_{r,ij}(x,y)\right| dy &\leq C \left\|\hat{K}_{r,ij}(x,y) - K_{r,ij}(x,y)\right\|_{L^\infty([0,1]^{d_i}\times[0,1]^{d_j})}\end{aligned}$$

almost surely, so by Young's inequality (Johnstone, 2013, Theorem C.26)

$$\begin{aligned}\sup_{\|f\|=1} \left\|\int (\hat{K}_{r,ij}(x,y) - K_{r,ij}(x,y))f(y)dy\right\| &\leq C \left\|\hat{K}_{r,ij}(x,y) - K_{r,ij}(x,y)\right\|_{L^\infty([0,1]^{d_i}\times[0,1]^{d_j})} \\ &\leq O((K_iK_j)^{-\alpha_{r,ij}/(d_i+d_j)})\end{aligned}$$

As this holds for each r, i, j, we have

$$\eta_K = \max\{\left\|B_C^K - B_C\right\|_{op}, \left\|A_C^K - A_C\right\|_{op}\} \leq O(J\max_{r,i,j}(K_iK_j)^{-\alpha_{r,ij}/(d_i+d_j)})$$

as claimed, by bounding the operator norm by the Frobenius norm of the $J \times J$ matrix with i, j element equal to the operator norm of the i, j block.

To use this result to bound the number of basis functions needed to obtain a total operator norm error of order ϵ, letting $\bar{\alpha} = \min_{r,ij} \frac{2\alpha_{r,ij}}{d_i+d_j}$, by setting $\{K_j\}_{j=1}^J$ all equal and proportional to $(\frac{J}{\epsilon})^{\frac{1}{\bar{\alpha}}}$, obtain $J\max_{r,i,j}(K_iK_j)^{-\alpha_{r,ij}/(d_i+d_j)} = O(\epsilon)$. This results in a basis of size $K = \sum_{j=1}^J K_j$ proportional to $J(\frac{J}{\epsilon})^{\frac{1}{\bar{\alpha}}}$ as claimed.

Next, bound ζ_K, the error induced by approximating each integral operator in (B_C, A_C) by a matrix with entries given by the discrete wavelet transform of $K_{r,ij}(x_s, y_t)$.

For convenience, define the level of the d_j-tensor product of multiresolution analyses of Im π^{K_j} in each dimension as $\{n_{jp}\}_{p=1}^{d_j}$, and let the total number of basis functions in the tensor product basis satisfy $K_j = \Pi_{p=1}^{d_j} 2^{n_{jp}}$.[10] The discrete wavelet transform in one dimension is a unitary mapping on the space spanned by the scaling functions $\phi_{n_j,s} := 2^{-n_j/2}\phi(2^{-n_j}x - s + 1)$ at multiresolution level n_j from vectors whose entries are inner products with these scaling functions to vectors whose entries are inner products with the orthonormal wavelet basis spanning the same space, and in multiple dimensions it maps the tensor product of scaling functions representation to the tensor product of wavelets representation. As the operator norm is unitarily invariant, it therefore suffices to bound the operator norm error in terms of the error in the representation defined in terms of scaling function coefficients. By the compact support, vanishing moment condition, and Hölder exponent bound, Beylkin *et al.* (1991) show by a Taylor expansion argument that if a scaling function with the property $\int \phi(x+\tau)x^m dx = 0$ for all integers $m \leq \alpha + 1$, for some integer τ, is used, then any $f(x) \in \Lambda^\alpha[0,1]$ satisfies $2^{-n/2} f(2^{-n}(k-1+\tau)) = \int f(x)\phi_{n,k}(x)dx + O(2^{-n(\alpha+\frac{1}{2})})$ uniformly in k, and for multivariate functions $f(x^1,\ldots,x^d) \in \Lambda^\alpha[0,1]^d$, a straightforward extension shows

$$(2^{-n_1/2}\ldots 2^{-n_d/2}) f(2^{-n}(k_1-1+\tau),\ldots,2^{-n}(k_d-1+\tau)) =$$
$$\int\ldots\int f(x_1,\ldots,x_d)\phi_{n,k^1}(x_1)\ldots\phi_{n,k^d}(x_d)dx_1\ldots dx_d + O(\prod_{p=1}^{d} 2^{-n_p/2}\sum_{p=1}^{d} 2^{-\alpha n_p})$$

[10] One can avoid restricting to powers of 2 by using a larger number of functions at the finest level, at the cost of more cumbersome notation. The order of all asymptotic results remains the same.

Applying this to $\frac{1}{\sqrt{K_i K_j}} K_{r,ij}(x_s, y_t)$ we see that its entries satisfy

$$\left| \frac{1}{\sqrt{K_i K_j}} K_{r,ij}(x_s, y_t) - \left\langle K_{r,ij}(x,y), \Pi_{p=1}^{d_i} \phi_{n_{ip},s+\tau}(x_p) \Pi_{p=1}^{d_j} \phi_{n_{jp},t+\tau}(y_p) \right\rangle \right|$$
$$= O((\prod_{p=1}^{d_i} 2^{-n_{ip}/2} \prod_{p=1}^{d_j} 2^{-n_{jp}/2})(\sum_{p=1}^{d_i} 2^{-n_{ip}\alpha_{r,ij}} + \sum_{p=1}^{d_j} 2^{-n_{jp}\alpha_{r,ij}}))$$

uniformly in s, t. To control the operator norm error induced by this approximation to the matrix of scaling function coefficients, we again use Young's inequality, combined with the fact that the scaling functions $\phi_{n_j,s}$ are rescaled translations of a single bounded and compactly supported function over a regular grid, to bound the operator norm error in the quadrature approximation of the finite projection of $K_{r,ij}(x,y)$ onto a finite tensor product wavelet basis. In particular, denoting $\theta_{ijst} := \left\langle K_{r,ij}(x,y), \Pi_{p=1}^{d_i} \phi_{n_{ip},s+\tau}(x_p) \Pi_{p=1}^{d_j} \phi_{n_{jp},t+\tau}(y_p) \right\rangle$ and $\hat{\theta}_{ijst} := \frac{1}{\sqrt{K_i K_j}} K_{r,ij}(x_s, y_t)$ the L^∞ norm error induced by quadrature in the $K_i \times K_j$ term representation of the kernel is equal to

$$\sup_{x,y \in [0,1]^{d_i} \times [0,1]^{d_j}} \left| \sum_{s=1}^{K_i} \sum_{t=1}^{K_j} (\theta_{ijst} - \hat{\theta}_{ijst}) \Pi_{p=1}^{d_i} \phi_{n_{ip},s+\tau}(x_p) \Pi_{p=1}^{d_j} \phi_{n_{jp},t+\tau}(y_p) \right|$$

As noted in Chen & Christensen (2015, Section 6), by the assumption that the one-dimensional scaling function ϕ has support within a compact interval, with length no greater than $3N+1$ for a fixed integer N (depending order of the wavelet used), at most $3N+1$ scaling functions at any fixed level n_j may overlap on any set of positive Lebesgue measure, and so over the $d_i + d_j$-dimensional tensor product space, no point

x, y is covered by more than $(3N+1)^{d_i+d_j}$ scaling functions.[11] As a result

$$\begin{aligned}
&\sup_{x,y\in[0,1]^{d_i}\times[0,1]^{d_j}} \left| \sum_{s=1}^{K_i}\sum_{t=1}^{K_j} (\hat{\theta}_{ijst} - \theta_{ijst}) \prod_{p=1}^{d_i} \phi_{n_{ip},s+\tau}(x_p) \prod_{p=1}^{d_j} \phi_{n_{jp},t+\tau}(y_p) \right| \\
&\quad \leq (3N+1)^{d_i+d_j} \max_{s,t} \left| \hat{\theta}_{ijst} - \theta_{ijst} \right| \sup_{x,y} \left| \prod_{p=1}^{d_i} \phi_{n_{ip},s+\tau}(x_p) \prod_{p=1}^{d_j} \phi_{n_{jp},t+\tau}(y_p) \right| \\
&\quad \leq (3N+1)^{d_i+d_j} O(\prod_{p=1}^{d_i} 2^{-n_{ip}/2} \prod_{p=1}^{d_j} 2^{-n_{jp}/2}) (\sum_{p=1}^{d_i} 2^{-n_{ip}\alpha_{r,ij}} + \sum_{p=1}^{d_j} 2^{-n_{jp}\alpha_{r,ij}})) * \\
&\qquad \prod_{p=1}^{d_i} 2^{n_{ip}/2} \prod_{p=1}^{d_j} 2^{n_{jp}/2} \sup_x |\phi(x)| \\
&\qquad = O((3N+1)^{d_i+d_j} (\sum_{p=1}^{d_i} 2^{-n_{ip}\alpha_{r,ij}} + \sum_{p=1}^{d_j} 2^{-n_{jp}\alpha_{r,ij}}))
\end{aligned}$$

by boundedness and the definition of $\phi_{n_{jp},s}$. When the number of basis functions used in each dimension is identical for all dimensions $p = 1 \ldots d_i$ and $1 \ldots d_j$, this term is bounded by

$$O((3N+1)^{d_i+d_j}(d_i+d_j)(K_iK_j)^{-\alpha_{r,ij}/(d_i+d_j)})$$

This is the same order as the projection result, except for a multiplicative constant depending on dimension. Let $\bar{d} = \max_j 2d_j$. Then, if the number of basis functions is set so that $\{K_j\}_{j=1}^J$ all equal and proportional to $(\frac{(3N+1)^{\bar{d}}\bar{d}J}{\epsilon})^{\frac{1}{\bar{\alpha}}}$, the above bound along with Young's inequality gives an operator norm error bound bound for each block no greater than $O(\frac{\epsilon}{J})$. With each of J^2 blocks bounded by no more than this quantity, obtain the bound

$$\zeta_K \leq O(\epsilon)$$

[11] The vanishing moments property characterizing Coiflets also requires that the length of the filter defining the scaling function be longer by a factor of 1.5 than the filter for the corresponding standard Daubechies wavelet. This results in an larger constant in front of the quadrature error and the running time of the discrete wavelet transform, but does not affect the rate of convergence.

exactly as claimed. □

Chapter 2 Proofs

Section 2.2 Results

We apply the implicit function theorem to calculate $\frac{d\omega}{d\lambda}$. Taking derivatives of equations (2.3),(2.4), (2.5), and (2.6) evaluated at the steady state, obtain

$$
\begin{aligned}
\frac{dY}{d\lambda} &= \mu\bar{w}(x)[.] \\
\frac{dY}{dw} &= \mu\bar{\lambda}(x)[.] \\
\frac{dT}{d\lambda} &= [\tilde{c}]^{\frac{1}{1-\sigma}} \frac{1}{1-\sigma}\left[\int_G \bar{\lambda}(z)\bar{w}(z)^{1-\sigma}e^{\tilde{\tau}(1-\sigma)\tau(x,z)}dz\right]^{\frac{\sigma}{1-\sigma}} \int_G [.]\bar{w}(z)^{1-\sigma}e^{\tilde{\tau}(1-\sigma)\tau(x,z)}dz \\
\frac{dT}{dw} &= [\tilde{c}]^{\frac{1}{1-\sigma}} \frac{1}{1-\sigma}\left[\int_G \bar{\lambda}(z)\bar{w}(z)^{1-\sigma}e^{\tilde{\tau}(1-\sigma)\tau(x,z)}dz\right]^{\frac{\sigma}{1-\sigma}} (1-\sigma)\int_G [.]\bar{\lambda}(z)\bar{w}(z)^{-\sigma}e^{\tilde{\tau}(1-\sigma)\tau(x,z)}dz \\
\frac{dw}{dY} &= [\tilde{c}]^{\frac{1}{\sigma}} \frac{1}{\sigma}\left[\int_G \bar{Y}(z)\bar{T}(z)^{\sigma-1}e^{-\tilde{\tau}(\sigma-1)\tau(x,z)}dz\right]^{\frac{1-\sigma}{\sigma}} \int_G [.]\bar{T}(z)^{\sigma-1}e^{-\tilde{\tau}(\sigma-1)\tau(x,z)}dz \\
\frac{dw}{dT} &= [\tilde{c}]^{\frac{1}{\sigma}} \frac{1}{\sigma}\left[\int_G \bar{Y}(z)\bar{T}(z)^{\sigma-1}e^{-\tilde{\tau}(\sigma-1)\tau(x,z)}dz\right]^{\frac{1-\sigma}{\sigma}} (\sigma-1)\int_G [.]\bar{Y}(z)\bar{T}(z)^{\sigma-2}e^{-\tilde{\tau}(\sigma-1)\tau(x,z)}dz \\
\frac{d\omega}{dw} &= \bar{T}(x)^{-\mu}[.] \\
\frac{d\omega}{dT} &= -\mu\bar{T}(x)^{-\mu-1}\bar{w}(z)[.]
\end{aligned}
$$

By the chain rule, we can express the derivative of the real wage with respect to the population distribution as

$$\frac{d\omega}{d\lambda} = \frac{d\omega}{dw}\frac{dw}{d\lambda} + \frac{d\omega}{dT}\left(\frac{dT}{dw}\frac{dw}{d\lambda} + \frac{dT}{d\lambda}\right) \tag{13}$$

where by the implicit function theorem in Banach space and the chain rule repeatedly

applied,

$$\frac{dw}{d\lambda} = (I - \frac{dw}{dY}\frac{dY}{dw} - \frac{dw}{dT}\frac{dT}{dw})^{-1}(\frac{dw}{dY}\frac{dY}{d\lambda} + \frac{dw}{dT}\frac{dT}{d\lambda}).$$

Section 2.3 Proofs

Proposition. *Derivation of* $\frac{\hat{d\omega}}{d\lambda}_{\phi}$*:* $\frac{\hat{d\omega}}{d\lambda}_{\phi} = (1 - \mu H(\phi))\frac{\mu H(\phi) - H(\phi)^2}{\sigma - \mu H(\phi) - (\sigma-1)H(\phi)^2} + \frac{\mu}{\sigma-1}H(\phi)$*, with* $H(\phi) := \frac{(\sigma-1)^2}{(\sigma-1)^2 + \tau^{-2}\phi^2}$

Proof. $\frac{d\omega}{d\lambda}$ is shown by Equation (13) to equal, in the uniform steady state case, $\frac{dw}{d\lambda} - \mu(\frac{dT}{dw}\frac{dw}{d\lambda} + \frac{dT}{d\lambda})$ which is a composition of convolution operators and their inverses and so can also be expressed as multiplication by the Fourier transform of some function. To construct the Fourier transform of the function, simplify the integrals in equations (2.3),(2.4), (2.5), and (2.6) and denote

$$H(\phi) = \frac{(\sigma-1)^2}{(\sigma-1)^2 + \tau^{-2}\phi^2}$$

the Fourier transform of the Laplace distribution in the convolution operator

$$\frac{\tau(1-\sigma)}{2}\int_G [.]e^{\tau(1-\sigma)|x-z|}dz.$$

This yields the formulas $\frac{\hat{dw}}{dT} = \frac{\sigma-1}{\sigma}H$, $\frac{\hat{dw}}{dY} = \frac{1}{\sigma}H$, $\frac{\hat{dT}}{d\lambda} = \frac{1}{1-\sigma}H$, $\frac{dT}{dw} = H$. Substituting into the expressions for partial derivatives, obtain $\frac{\hat{dw}}{d\lambda} = \frac{-\frac{\mu}{\sigma}H + \frac{1}{\sigma}H^2}{1 - \frac{\mu}{\sigma}H - \frac{\sigma-1}{\sigma}H^2}$ and

$$\frac{\hat{d\omega}}{d\lambda}_{\phi} = (1 - \mu H(\phi))\frac{\mu H(\phi) - H(\phi)^2}{\sigma - \mu H(\phi) - (\sigma-1)H(\phi)^2} + \frac{\mu}{\sigma-1}H(\phi). \tag{14}$$

This is almost the same as Krugman (1996)'s equation (A.44) for this term, but differs slightly due to what appears to be an algebra error in the text. □

Proof. of (2.1). The proof applies the machinery and notation of Stewart (1973). While rates of convergence are obtained, no attempt is made to ensure that these are

optimal. First, note that $\gamma_\phi = \|(B_\phi, A_\phi) - (B_I^i, A_I^i)\|_F \to 0$ by assumption, and so all submatrices also converge at least as rapidly in Frobenius norm. Next note that (B_I^i, A_I^i) has generalized Schur decomposition

$$\left(Q^{*\infty} \begin{bmatrix} S_{11}^\infty & S_{22}^\infty \\ 0 & S_{22}^\infty \end{bmatrix} \begin{bmatrix} U_{11}^\infty & U_{12}^\infty \\ U_{21}^\infty & U_{22}^\infty \end{bmatrix}, Q^{*\infty} \begin{bmatrix} T_{11}^\infty & T_{22}^\infty \\ 0 & T_{22}^\infty \end{bmatrix} \begin{bmatrix} U_{11}^\infty & U_{12}^\infty \\ U_{21}^\infty & U_{22}^\infty \end{bmatrix} \right)$$

where $(S^\infty, T^\infty) = \left(\begin{bmatrix} 1 & 0 & 0 \\ 0 & \sqrt{2} & 0 \\ 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1/\sqrt{2} \\ 0 & 0 & 1/\sqrt{2} \end{bmatrix} \right)$ and $U^\infty = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$.

Applying standard formulas for policy functions, obtain $\hat{g}_\infty = -U_{22}^{\infty -} U_{21}^\infty = (0, 0)$ and

$$\hat{h}_\infty = (I_2 + \hat{g}_\infty^* \hat{g}_\infty)^{-1} (\begin{pmatrix} I_2 \\ \hat{g}_\infty \end{pmatrix}^* U_1^{\infty *} S_{11}^{\infty -1} T_{11}^\infty U_1^\infty \begin{pmatrix} I_2 \\ \hat{g}_\infty \end{pmatrix}) = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$$

As generalized eigenvalues corresponding to the stable subspace are equal to 0 and the generalized eigenvalue corresponding to the unstable subspace is ∞, the measure of subspace separation defined in Stewart (1973), which ensures that Schur subspaces are numerically stable, is given by $\delta = \text{dif}(S_{11}^\infty, T_{11}^\infty, S_{22}^\infty, T_{22}^\infty) > 0$. As a result, by Stewart (1973), Theorem 5.7 and 5.3, $\left\| \sin \Theta(U_1^{\infty *}, U_1^{\phi *}) \right\|_F \leq 2 \frac{\gamma_\phi}{\delta - 2\gamma_\phi}$ for γ_ϕ small enough, and similarly for U_2^∞, where $\Theta(U_1^{\infty *}, U_1^{\phi *})$ is the matrix of principal angles between the span of $U_1^{\infty *}$ and $U_1^{\phi *}$. While this does not imply that $\left\| U_1^{\infty *} - U_1^{\phi *} \right\|_F \to 0$, as the span does not uniquely define the basis, it does imply, because $U_2^{\phi *}$ and $U_2^{\infty *}$ have a one-dimensional span and norm 1, that

$$\begin{aligned} \left\| U_2^{\phi *} - U_2^{\infty *} \right\|_F^2 &= 2 - 2 \left| \cos \Theta(U_2^{\phi *}, U_2^{\infty *}) \right| \\ &= 2 - 2\sqrt{1 - \sin^2 \Theta(U_2^{\phi *}, U_2^{\infty *})} \\ &\leq 2 - 2\sqrt{1 - (\frac{2\gamma_\phi}{\delta - 2\gamma_\phi})^2} = O(\gamma_\phi) \to 0. \end{aligned}$$

Since $U_{22}^{\infty} = 1$ is invertible, the policy function $\hat{g}_\phi = -U_{22}^{\phi -} U_{21}^{\phi}$ therefore satisfies the bound $\|\hat{g}_\phi - \hat{g}_\infty\|_F^2 \leq O(\gamma_\phi) \to 0$, as claimed.

Further, it is possible to show that for each ϕ, there exists a unitary (2×2) transformation R_ϕ of $U_1^{\infty *}$ such that $\left\| U_1^{\infty *} R_\phi - U_1^{\phi} \right\|_F \to 0$. Applying the definition of principal angles, for each ϕ there exist unitary matrices $R_\phi^1 = [R_{\phi 1}^1, R_{\phi 2}^1]$ and $R_\phi^2 = [R_{\phi 1}^2, R_{\phi 2}^2]$ such that

$$[\cos \Theta(U_1^{\infty *}, U_1^{\phi *})]_{11} = \left\langle U_1^{\infty *} R_{\phi 1}^1, U_1^{\phi *} R_{\phi 1}^2 \right\rangle$$

and

$$[\cos \Theta(U_1^{\infty *}, U_1^{\phi *})]_{22} = \left\langle U_1^{\infty *} R_{\phi 2}^1, U_1^{\phi *} R_{\phi 2}^2 \right\rangle,$$

so

$$\begin{aligned}
\left\| U_1^{\infty *} R_\phi - U_1^{\phi} \right\|_F^2 &:= \left\| U_1^{\infty *} R_\phi^1 R_\phi^{2*} - U_1^{\phi} \right\|_F^2 \\
&= \left\| U_1^{\infty *} R_\phi^1 - U_1^{\phi} R_\phi^2 \right\|_F^2 \\
&= 2(1 - [\cos \Theta(U_1^{\infty *}, U_1^{\phi *})]_{11} + 1 - [\cos \Theta(U_1^{\infty *}, U_1^{\phi *})]_{22}) \\
&\leq 4 - 4\sqrt{1 - (\frac{2\gamma_\phi}{\delta - 2\gamma_\phi})^2} = O(\gamma_\phi) \to 0.
\end{aligned}$$

Equivalent results show that for a different unitary transform R_ϕ^Q, $\left\| R_\phi^Q Q_1^{\infty} - U_1^{\phi} \right\|_F^2 = O(\gamma_\phi)$. Combining these results and applying the triangle inequality,

$$\left\| S_{11}^{\phi} - S_{11}^{\infty(\phi)} \right\|_F := \left\| S_{11}^{\phi} - R_\phi^Q Q_1^{\infty} A_\infty U_1^{\infty *} R_\phi \right\|_F \leq O(\gamma_\phi^{\frac{1}{2}})$$

and

$$\left\| T_{11}^{\phi} - T_{11}^{\infty(\phi)} \right\|_F := \left\| T_{11}^{\phi} - R_\phi^Q Q_1^{\infty} B_\infty U_1^{\infty *} R_\phi \right\|_F \leq O(\gamma_\phi^{\frac{1}{2}})$$

also, gives convergence of the generalized Schur components of the finite order ma-

trices along a triangular array to unitary transformations $(S_{11}^{\infty(\phi)}, T_{11}^{\infty(\phi)})$ of the generalized Schur components of the limit pencil. Noting that unitary transformations leave singular values unaffected and that S_{11}^{∞} is invertible, S_{11}^{ϕ} is also asymptotically invertible, so by Weyl's inequality

$$\left\| S_{11}^{\phi-1} - S_{11}^{\infty(\phi)-1} \right\|_F \leq \left\| S_{11}^{\phi-1} \right\|_{op} \left\| S_{11}^{\infty(\phi)-1} \right\|_{op} \left\| S_{11}^{\phi} - S_{11}^{\infty(\phi)} \right\|_F \leq O(\gamma_\phi^{\frac{1}{2}}).$$

Using the unitarity of R_ϕ and applying the triangle inequality, one can see that

$$\left\| U_1^{\phi *} S_{11}^{\phi-1} T_{11}^{\phi} U_1^{\phi} - U_1^{\infty *} S_{11}^{\infty-1} T_{11}^{\infty} U_1^{\infty} \right\|_F = \left\| U_1^{\phi *} S_{11}^{\phi-1} T_{11}^{\phi} U_1^{\phi} - U_1^{\infty *} R_\phi S_{11}^{\infty(\phi)-1} T_{11}^{\infty(\phi)} R_\phi^* U_1^{\infty} \right\|_F \leq O(\gamma_\phi^{\frac{1}{2}}),$$

and so the fact that Schur vectors do not converge does not affect the convergence of the policy function, which is invariant to unitary transformations of these vectors. Finally, defining

$$\hat{h}_\phi = (I_2 + \hat{g}_\phi^* \hat{g}_\phi)^{-1} \left(\begin{pmatrix} I_2 \\ \hat{g}_\phi \end{pmatrix}^* U_1^{\phi *} S_{11}^{\phi-1} T_{11}^{\phi} U_1^{\phi} \begin{pmatrix} I_2 \\ \hat{g}_\phi \end{pmatrix} \right)$$

the above results and the triangle inequality imply that $\left\| \hat{h}_\phi - \hat{h}_\infty \right\|_F \leq O(\gamma_\phi^{\frac{1}{2}})$.

To show compactness, it suffices to show that the singular values converge to 0. As $g[.]$ and $h[.]$ are block-diagonal, it suffices to show that the operator norm of each block converges to 0. As the operator norm is bounded by the Frobenius norm, each block has operator norm at most $O(\gamma_\phi^{\frac{1}{2}}) \to 0$ and so compactness holds.

(ii) To show that an $h[.]$ is Hilbert Schmidt, $\text{Tr}(h^* h) < \infty$, it suffices to show that the sum of squared singular values converges. As the sum of squared singular values for each block is equal to the square of its Frobenius norm, which is $O(\gamma_\phi)$ for large $|\phi|$, convergence holds so long as $\sum_{\phi=n}^{\infty} \gamma_\phi < \infty$ for some finite n. Superlinear

convergence $\gamma_\phi = O(|\phi|^{-(1+\epsilon)})$ for some $\epsilon > 0$ is sufficient for this sum to be finite. $\square$

Chapter 3 Proofs

Proof. of (3.1) Unitarity of U provides the following facts: since $U^* = U^{-1}$, we have $U^*U = I$. Decomposing U into U_{11}, U_{12}, U_{21}, and U_{22} obtain

$$\begin{bmatrix} U_{11}^*U_{11} + U_{21}^*U_{21} & U_{11}^*U_{12} + U_{21}^*U_{22} \\ U_{12}^*U_{11} + U_{22}^*U_{21} & U_{12}^*U_{12} + U_{22}^*U_{22} \end{bmatrix} = I = \begin{bmatrix} I_x & 0 \\ 0 & I_y \end{bmatrix}$$

Where $I_x = \varphi^{X*}\varphi^X$ and $I_y = \varphi^{Y*}\varphi^Y$ are the identity operators on $\mathcal{H}_x$ and $\mathcal{H}_y$ respectively. To see this more formally, consider $U_{11}^*U_{12} + U_{21}^*U_{22}$. It can be written as

$$\begin{aligned} \varphi^{X*}U_1^*U_1\varphi^Y + \varphi^{X*}U_2^*U_2\varphi^Y &= \varphi^{X*}(U_1^*U_1 + U_2^*U_2)\varphi^Y \\ &= \varphi^{X*}\varphi^Y = 0 \end{aligned}$$

Equivalent calculations describe the other identities.

Using these identities we can express

$$\begin{aligned} (U_{11} + U_{12}g_X)^*(U_{11} + U_{12}g_X) &= U_{11}^*U_{11} + U_{11}^*U_{12}g_X + g_X^*U_{12}^*U_{11} + g_X^*U_{12}^*U_{12}g_X \\ &= I_x - U_{21}^*U_{21} - U_{21}^*U_{22}g_X + g_X^*U_{22}^*U_{21} + g_X^*(I_y - U_{22}^*U_{22})g_X \\ &= I_x - U_{21}^*U_{21} + U_{21}^*U_{22}U_{22}^*(U_{22}U_{22}^*)^{-1}U_{21} \\ &\quad + U_{21}^*(U_{22}U_{22}^*)^{-1*}U_{22}U_{22}^*U_{21} + g_X^*I_yg_X - g_X^*U_{22}^*U_{22}g_X \\ &= I_x - U_{21}^*U_{21} + U_{21}^*U_{21} + U_{21}^*U_{21} + g_X^*I_yg_X - U_{21}^*U_{21} \\ &= I_x + g_X^*I_yg_X \\ &= I_x + {g_X}^*g_X \end{aligned}$$

As a result, post-multiplying by $(U_{11}+U_{12}g_X)^{-1}$ and inverting $(I_x+g_x{}^*g_x)$, obtain

$$(U_{11}+U_{12}g_X)^{-1} = (I_x + g_X{}^*g_X)^{-1}(U_{11}+U_{12}g_X)^*$$

□

Lemma 3.2

Proof. By the Courant-Fisher min-max characterization of eigenvalues of self-adjoint matrices,

$$\begin{aligned}
\lambda_{\min}(PAP) &= \min\{\max\{\langle PAPx, x\rangle \,|\, \|x\| = 1,\ x \in U\} | U \subseteq \mathrm{Sp}(P),\ \dim U = 1\} \\
&= \min\{\max\{\langle APx, Px\rangle \,|\, \|x\| = 1,\ x \in U\} | U \subseteq \mathrm{Sp}(P),\ \dim U = 1\} \\
&= \min\{\max\{\langle Ax, x\rangle \,|\, \|x\| = 1,\ x \in U\} | U \subseteq \mathrm{Sp}(P),\ \dim U = 1\} \\
&\geq \inf\{\max\{\langle Ax, x\rangle \,|\, \|x\| = 1,\ x \in U\} | U \subseteq \mathcal{H},\ \dim U = 1\} \\
&= \lambda_{\min}(A) \\
&= \lambda^{-1}_{\max}(A^{-1}) \\
&= \frac{1}{\|A^{-1}\|}
\end{aligned}$$

where the fact that orthonormal projections are self-adjoint was used in line 2, line 5 follows from Helmberg (1969) §30 corollary 8.1, line 6 follows from the subsequent exercise, and the final line follows from the same corollary and the assumption of positivity, so $\|A^{-1}\| = \max\{|\lambda|,\ \lambda \in \Sigma(A^{-1})\} = \max\{\lambda,\ \lambda \in \Sigma(A^{-1})\}$. □

Lemma 3.3

Proof. $\pi^{k_1} \overset{s.o.t.}{\to} I_{\mathcal{H}_2}$ and $\pi^{k_2} \overset{s.o.t.}{\to} I_{\mathcal{H}_1}$ by monotonicity and Helmberg (1969) §30 Theorem 5. As a result, for any sequence $k_1,\ k_2 \to \infty$, $\tilde{A} \overset{s.o.t.}{\to} A$ and $\tilde{B} \overset{s.o.t.}{\to} B$ since

$\forall x \in \mathcal{H}_1$

$$\|(\pi^{k_1} A \pi^{k_2} - A)x\| \leq \|\pi^{k_1}\| \, \|A\| \, \|(\pi^{k_2} - I_{\mathcal{H}_1})x\| + \|(\pi^{k_1} - I_{\mathcal{H}_2})Ax\| \to 0$$

since $\|\pi^{k_1}\| \leq 1$ and A is a bounded operator so $\|A\| < \infty$. Similarly for B and so also A^* and B^*. More generally, strong operator topology convergence is continuous with respect to application of uniformly bounded sequences of linear operators.

Next, we find a uniform bound on $\|((\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B}))^{-1}\|$. For this, first we note that by the definition of Γ-regularity, $\forall \zeta \in \Gamma$, $\zeta A - B$ is a bounded operator with bounded inverse, and by compactness of Γ, and the fact that the resolvent set is open (Gohberg *et al.* , 1990, IV.1) and so the inverse is continuous over a curve inside of it, $\max_{\zeta \in \Gamma} \|(\zeta A - B)^{-1}\| < \infty$. $(\zeta A^* - B^*)(\zeta A - B)$ is therefore a bounded, invertible, self-adjoint positive operator with bounded inverse, uniformly over $\zeta \in \Gamma$. Applying 3.2, obtain that for any k_2,

$$\max_{\zeta \in \Gamma} \|(\pi^{k_2}(\zeta A^* - B^*)(\zeta A - B)\pi^{k_2})^{-1}\| \leq \max_{\zeta \in \Gamma} \|((\zeta A^* - B^*)(\zeta A - B))^{-1}\| < \infty.$$

To extend this to $((\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B}))^{-1}$, note that for a fixed k_2, $\pi^{k_2}(\zeta A^* - B^*)(\zeta A - B)\pi^{k_2}$ is a finite-dimensional and so compact operator which may be represented by a $k_2 \times k_2$ matrix. Taking the k_2-dimensional set of right singular vectors of this matrix on the space spanned by π^{k_2}, say $\{s_i\}_{i=1}^{k_2}$, $(\zeta A - B)\pi^{k_2}$ may be represented by the operator $\sum_{i=1}^{k_2} \langle s_i, (.) \rangle (\zeta A - B)s_i$. By $\pi^{k_1} \overset{s.o.t.}{\to} I_{\mathcal{H}_2}$, for fixed k_2, as $k_1 \to \infty$ $\max_{i \in 1 \ldots k_2} \|(\pi^{k_1} - I_{\mathcal{H}_2})(\zeta A - B)s_i\| \to 0$ and so

$$\|(\pi^{k_1} - I_{\mathcal{H}_2})(\zeta A - B)\pi^{k_2}\| \leq k_2 \max_{i \in 1 \ldots k_2} \|(\pi^{k_1} - I_{\mathcal{H}_2})(\zeta A - B)s_i\| \to 0.$$

The same logic applies to the adjoint, and so for fixed k_2,

$$\|\pi^{k_2}(\zeta A^* - B^*)\pi^{k_1}(\zeta A - B)\pi^{k_2} - \pi^{k_2}(\zeta A^* - B^*)(\zeta A - B)\pi^{k_2}\| \to 0. \tag{15}$$

This applies uniformly over ζ by compactness. Thus, all singular values also converge, and so the singular values of the inverse converge as well. Thus, by taking a sequence of k_1 increasing in k_2, we may choose $k_1(k_2)$ large enough such that for any chosen $\delta > 0$,

$$\max_{\zeta\in\Gamma}\|((\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B}))^{-1}\| \leq \max_{\zeta\in\Gamma}\|((\zeta A^* - B^*)(\zeta A - B))^{-1}\| + \delta < \infty$$

uniformly in k_2.

As a result, we have that $\forall x \in \mathcal{H}_1$

$$\begin{aligned}
&\|((\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B}))^{-1}x - ((\zeta A^* - B^*)(\zeta A - B))^{-1}x\| \leq \\
&\qquad \left\|((\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B}))^{-1}\right\| * \\
&\|(\zeta\tilde{A}^*-\tilde{B}^*)(\zeta\tilde{A}-\tilde{B})((\zeta\tilde{A}^*-\tilde{B}^*)(\zeta\tilde{A}-\tilde{B}))^{-1}x-((\zeta\tilde{A}^*-\tilde{B}^*)(\zeta\tilde{A}-\tilde{B}))((\zeta A^*-B^*)(\zeta A-B))^{-1}x\| \\
&+ \|(I - ((\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B}))^{-1}(\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B}))((\zeta A^* - B^*)(\zeta A - B))^{-1}x\| \leq \\
&\qquad (\max_{\zeta\in\Gamma}\|((\zeta A^* - B^*)(\zeta A - B))^{-1}\| + \delta)* \\
&\qquad (\|(\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B})((\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B}))^{-1}x - x\|+ \\
&\qquad \|x - ((\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B}))((\zeta A^* - B^*)(\zeta A - B))^{-1}x\|)+ \\
&+ \|(I - ((\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B}))^{-1}(\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B}))((\zeta A^* - B^*)(\zeta A - B))^{-1}x\|
\end{aligned}$$

The term on the penultimate line goes to zero since $(\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B}) \overset{s.o.t.}{\to} (\zeta A^* - B^*)(\zeta A - B)$. The final and third from last lines are given by projections of a fixed vector onto the orthogonal complements of the image and the domain, respectively, of $(\zeta\tilde{A}^* - \tilde{B}^*)(\zeta\tilde{A} - \tilde{B})$, which by the self-adjointness of this operator are

identical. We may then note that by 15, for any k_2, there exists k_1 such that these projections are within any ϵ of the projections onto the orthogonal complement of the image and domain (again identical) of $\pi^{k_2}(\zeta A^* - B^*)(\zeta A - B)\pi^{k_2}$. As this is a projection onto a monotonically decreasing sequence of subspaces whose intersection converges to the null set as k_2 grows, Helmberg (1969) §30 Theorem 5 applies, and so both these terms also converge to 0. Again this result holds uniformly in $\zeta \in \Gamma$. As a result, passing through the product with $(\zeta\tilde{A}^* - \tilde{B}^*)\tilde{A}$, $\forall x \in \mathcal{H}_1$

$$\max_{\zeta\in\Gamma}\|((\zeta\tilde{A}^*-\tilde{B}^*)(\zeta\tilde{A}-\tilde{B}))^{-1}(\zeta\tilde{A}^*-\tilde{B}^*)\tilde{A}x-((\zeta A^*-B^*)(\zeta A-B))^{-1}(\zeta A^*-B^*)Ax\| \to 0.$$

This uniform convergence then implies by dominated convergence that

$$\tilde{P}_1^{k_1(k_2),k_2} \overset{s.o.t.}{\to} \frac{1}{2\pi\iota}\int_\Gamma((\zeta A^* - B^*)(\zeta A - B))^{-1}(\zeta A^* - B^*)Ad\zeta.$$

Then, by the invertibility of $(\zeta A - B)$,

$$((\zeta A^* - B^*)(\zeta A - B))^{-1} = (\zeta A - B)^{-1}(\zeta A^* - B^*)^{-1}$$

and so

$$\frac{1}{2\pi\iota}\int_\Gamma((\zeta A^* - B^*)(\zeta A - B))^{-1}(\zeta A^* - B^*)Ad\zeta = \frac{1}{2\pi\iota}\int_\Gamma(\zeta A - B)^{-1}Ad\zeta = P_1$$

and so the lemma holds. □

Lemma 3.4

Proof. By strong operator topology convergence of $\tilde{P}_1^{k_1(k_2),k_2}$ to P_1, we have that for any fixed Z^{k_3} and any $\delta > 0$, there exists, by 3.3, $k_2(k_3)$ and $k_1(k_2(k_3))$, such that $\|\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)}Z^{k_3} - P_1Z^{k_3}\| < \delta$. $P_1Z^{k_3}$ is a fixed quasimatrix of rank no greater than k_3, and so, unless it is identically 0, it has a minimal non-zero singular value,

call it $\sigma_{\min}^{k_3}$. We may therefore set $\varepsilon_{k_3} = \frac{1}{2}\sigma_{\min}^{k_3}$ if $P_1 Z^{k_3}$ is not identically 0 and some small constant ε if it is. By Weyl's inequality for singular values (see, e.g., Tao (2011, Ch. 1.3)), we have by above that there exist $k_2(k_3)$ and $k_1(k_2(k_3))$ such that

$$|\sigma_i(\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)} Z^{k_3}) - \sigma_i(P_1 Z^{k_3})| \leq \|\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)} Z^{k_3} - P_1 Z^{k_3}\| < \varepsilon_{k_3} \ \forall i,$$

where σ_i is the i^{th} singular value in order of decreasing absolute value. As a result, if $\sigma_i(P_1 Z^{k_3}) = 0$, the numerical singular value will be below the threshold, and if $\sigma_i(P_1 Z^{k_3}) \neq 0$ (and so $\varepsilon_{k_3} \neq 0$), the numerical singular value will be greater than $\sigma_{\min}^{k_3} - \varepsilon_{k_3} = \varepsilon_{k_3}$ and so above the threshold. As a result, for sufficiently large k_2, regularizing selects exactly the singular vectors corresponding to nonzero singular values. Next, we assess the bias induced in the singular vectors themselves by the numerical error in $\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)} Z^{k_3}$ using a variant of the Davis-Kahan theorem for eigenvectors of self-adjoint operators. In particular, we may use the construction of Bosq (2000) providing error bounds on the norm convergence of approximate eigenvectors by noting that the left singular vectors of $\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)} Z^{k_3}$ are the eigenvectors of $(\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)} Z^{k_3})(\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)} Z^{k_3})^*$, and so the i^{th} column of $U_1^{*k_3,\varepsilon_{k_3}}$, denoted $U_{1i}^{*k_3,\varepsilon_{k_3}}$ converges (up to a sign change) to the i^{th} singular vector of $P_1 Z^{k_3}$, suggestively denoted $U_{1i}^{*(k_3)}$. Formally,

$$\|\text{sign}(\langle U_{1i}^{*k_3,\varepsilon_{k_3}}, U_{1i}^{*(k_3)}\rangle) U_{1i}^{*k_3,\varepsilon_{k_3}} - U_{1i}^{*(k_3)}\| \leq$$
$$\frac{2\sqrt{2}}{\min\{\sigma_i - \sigma_{i+1}.\sigma_i - \sigma_{i-1}\}} \|(\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)} Z^{k_3})(\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)} Z^{k_3})^* - (P_1 Z^{k_3})(P_1 Z^{k_3})^*\|$$

following from Bosq (2000, Lemma 4.3), where σ_i is the i^{th} singular value of $P_1 Z^{k_3}$. This formula applies in the case that the gap in the singular values is non-zero. If some eigenvalues have multiplicity greater than one, Bosq (2000, Lemma 4.4) shows that the sample singular vectors converge also to vectors spanning the spectral subspace corresponding to that eigenvalue, where the measure of the spectral gap used is then

given by the gap between the corresponding eigenvalue and that of the closest distinct eigenvalue. Note that for fixed k_3, the operators of interest are compact. As a result, for any fixed k_3, there exists a minimal spectral gap $C_{k_3} = \min_{i<k_3}\{\sigma_i - \sigma_{i+1}\}$, if the sequence of singular values is arranged so that repeated singular values occur only once in the sequence. So, by the operator norm convergence of $\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)} Z^{k_3}$ to $P_1 Z^{k_3}$, $\forall \delta > 0$ there exists a $k_2(k_3)$ such that

$$\|(\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)} Z^{k_3})(\tilde{P}_1^{k_1(k_2(k_3)),k_2(k_3)} Z^{k_3})^* - (P_1 Z^{k_3})(P_1 Z^{k_3})^*\| < \frac{C_{k_3}}{2\sqrt{2}}\delta$$

and so $\max_{i=1\ldots k_3} \|\text{sign}(\langle U_{1i}^{*k_3,\varepsilon_{k_3}}, U_{1i}^{*(k_3)}\rangle) U_{1i}^{*k_3,\varepsilon_{k_3}} - U_{1i}^{*(k_3)}\| < \delta$.

We may define for each k_3 the operator $U_1^{(k_3)}$ by $U_1^{*(k_3)} = [U_1^{*(k_3),a}, U_1^{*(k_3),b}]$ where $U_1^{*(k_3),a}$ is given by the quasimatrix with left singular vectors of $P_1 Z^{k_3}$ corresponding to non-zero singular values as columns, without loss of generality assumed to take the sign as $U_{1i}^{*k_3,\varepsilon_{k_3}}$, so we may ignore the indeterminacy of the sign in the convergence result, and $U_1^{*(k_3),b}$ is constructed by taking a complete orthonormal basis of $Im(P_1)/Im(U_1^{(k_3),a})$ as columns, so $U_1^{(k_3)}$ is an analysis operator on $Im(P_1)$ representing the coefficients of an element of this space with respect to a complete orthonormal basis on $Im(P_1)$, and mapping elements of $\mathcal{H}_1/Im(P_1)$ to 0.

So, since all the columns converge and the correct (corresponding to non-zero singular value) columns are chosen asymptotically, we obtain along the sequence $k_1(k_2(k_3))$, $k_2(k_3)$, and ε_{k_3}, $\|U_1^{k_3,\varepsilon_{k_3}} - U_1^{(k_3),a}\| \to 0$. To show $\forall x \in \mathcal{H}_1$, $\|(U_1^{(k_3),a} - U_1^{(k_3)})x\| \to 0$, note that

$$\|(U_1^{(k_3),a} - U_1^{(k_3)})x\| = \|U_1^{(k_3),b} x\| = \|U_1^{*(k_3),b} U_1^{(k_3),b} x\|$$

because $U_1^{*(k_3),b}$ as a matrix of orthonormal columns is an isometry on $Im(U_1^{(k_3),b})$.

By the definition of $U_1^{(k_3),b}$, $U_1^{*(k_3),b}U_1^{(k_3),b}$ is an orthogonal projection onto

$$Im(P_1)/Im(U_1^{(k_3),a}) = Im(P_1)/Im(P_1 Z^{k_3})$$

which is, by construction of Z^{k_3}, a monotonically decreasing sequence of subsets converging to the empty set, and so goes to 0 by Helmberg (1969) §30 Theorem 5.

Combining the above, obtain $\forall x \in \mathcal{H}_1$, $\|(U_1^{k_3,\varepsilon_{k_3}} - U_1^{(k_3)})x\| \leq \|U_1^{k_3,\varepsilon_{k_3}} - U_1^{(k_3),a}\|\|x\| + \|(U_1^{(k_3),a} - U_1^{(k_3)})x\| \to 0$, as desired. □

Lemma 3.5

Proof. It suffices to show that $I - U_1^{k_3,\varepsilon_{k_3}*}U_1^{k_3,\varepsilon_{k_3}}$ converges in the strong operator topology to an orthogonal projection onto $\mathcal{H}_1/Im(P_1)$, as the subsequent steps of the proof are identical to those of the previous one with $\tilde{P}_1^{k_1(k_2),k_2}$ replaced by $I - U_1^{k_3,\varepsilon_{k_3}*}U_1^{k_3,\varepsilon_{k_3}}$. For this, note simply that for any k_3, $I - U_1^{(k_3)*}U_1^{(k_3)}$ is identical and an orthogonal projection onto $\mathcal{H}_1/Im(P_1)$, by the construction of $U_1^{(k_3)}$. Strong operator topology convergence of $U_1^{k_3,\varepsilon_{k_3}}$ to $U_1^{(k_3)}$ was established by the previous lemma, and by construction, since $U_1^{k_3,\varepsilon_{k_3}*}$ is given by a set of orthonormal column vectors, it satisfies $\|U_1^{k_3,\varepsilon_{k_3}*}\| \leq 1$ uniformly in k_3. Further, for all $x \in Im(U_1^{(k_3)})$, a space isometrically isomorphic to and so w.lo.g. represented by ℓ_2,

$$\begin{aligned}\|(U_1^{k_3,\varepsilon_{k_3}*} - U_1^{(k_3)*})x\|_{\ell_2} \leq & \|U_1^{k_3,\varepsilon_{k_3}*} - U_1^{(k_3),a*}\|\|x\|_{\ell_2} + \|(U_1^{(k_3),a*} - U_1^{(k_3)*})x\|_{\ell_2} \\ = & \|U_1^{k_3,\varepsilon_{k_3}} - U_1^{(k_3),a}\|\|x\|_{\ell_2} + \|(U_1^{(k_3),a*} - U_1^{(k_3)*})x\|_{\ell_2}\end{aligned}$$

since the operator norm is invariant under adjoints. The first term goes to zero by the argument in the previous lemma. By the isometry property of $U_1^{(k_3),b}$ on $Im(U_1^{(k_3)*})$, the second term is identical in norm to $\|U_1^{(k_3),b}U_1^{(k_3),b*}x\|_{\ell_2}$ which is a sequence of projections onto spaces monotonically decreasing to zero and so converges to 0 by Helmberg (1969) §30 Theorem 5. Combining these results, $I - U_1^{k_3,\varepsilon_{k_3}*}U_1^{k_3,\varepsilon_{k_3}} \overset{s.o.t.}{\to}$

$I - U_1^{(k_3)*}U_1^{(k_3)}$ as desired. □

Lemma 3.6

Proof. First, show that $\mathrm{Im}(\pi_2) = \mathrm{Im}(AP_1)$. By Gohberg *et al.* (1990, Thm IV.1.1), $\pi_2 A = AP_1$, so if we can show $\mathrm{Im}(\pi_2) = \mathrm{Im}(\pi_2 A)$, this holds. Clearly, $\mathrm{Im}(\pi_2 A) \subseteq \mathrm{Im}(\pi_2)$. Consider any $z \in \mathrm{Im}(\pi_2)$. By construction, A_1 maps $\mathrm{Im}(P_1) \to \mathrm{Im}(\pi_2)$ and is invertible. So since $z \in \mathrm{Im}(\pi_2)$, $A_1^{-1}z \in \mathrm{Im}P_1$ and so $AA_1^{-1}z = A_1A_1^{-1}z = z$ and so $\pi_2 AA_1^{-1}z = \pi_2 z$. Then because π_2 is a projection, $\pi_2 z = z$. So, for any $z \in \mathrm{Im}(\pi_2)$, there exists an element, $A_1^{-1}z$ such that $\pi_2 AA_1^{-1}z = z$, and so $\mathrm{Im}(\pi_2) \subseteq \mathrm{Im}(\pi_2 A)$. So $\mathrm{Im}(\pi_2) = \mathrm{Im}(\pi_2 A) = \mathrm{Im}(AP_1)$, and so we can use the span of AP_1 to construct Q_1.

By 3.4, for fixed k_3, as $k_2,\ k_1(k_2) \to \infty$, $\|U_1^{k_3,\varepsilon_{k_3}*} - U_1^{(k_3),a*}\| \to 0$. As this is a finite-dimensional matrix, we have by the boundedness and pointwise consistency of $\tilde{A}$ and $\tilde{P}_1^{k_1(k_2),k_2}$ that

$$\|\tilde{A}\tilde{P}_1^{k_1(k_2),k_2}U_1^{k_3,\varepsilon_{k_3}*} - AP_1U_1^{(k_3),a*}\| \to 0$$

as $k_2,\ k_1(k_2) \to \infty$. By construction, the columns of $U_1^{(k_3),a*}$ are orthonormal elements of $\mathrm{Im}(P_1)$, and AP_1 has bounded inverse on $\mathrm{Im}(P_1)$. As a result, uniformly over all k_3, the minimum singular value of $AP_1U_1^{(k_3),a*}$ is bounded away from 0. So, by the continuity of the QR decomposition with respect to the operator norm at points with minimal singular value bounded away from 0 (Golub & van Loan, 1996, Ch. 5), for fixed k_3, $\|Q_1^{k_3*} - Q_1^{(k_3),a*}\| \to 0$ where $Q_1^{(k_3),a*}R^{(k_3)} = qr(AP_1U_1^{(k_3),a*})$. $Q_1^{(k_3),a*}$ is a finite rank operator whose columns are orthonormal elements of $\mathrm{Im}(AP_1) = \mathrm{Im}(\pi_2)$. As with $U_1^{(k_3)}$, we may define $Q_1^{(k_3)*}$ by $Q_1^{(k_3)*} = [Q_1^{(k_3),a*}, Q_1^{(k_3),b*}]$ where $Q_1^{(k_3),b*}$ is constructed by taking a complete orthonormal basis of $Im(\pi_2)/Im(Q_1^{(k_3),a*})$ as columns, so $Q_1^{(k_3)}$ is an analysis operator on $Im(\pi_2)$ representing the coefficients of an element of this space with respect to a complete orthonormal basis on $Im(\pi_2)$, and mapping elements of $\mathcal{H}_2/Im(\pi_2)$ to 0. To show $\forall x \in \mathcal{H}_2$, $\|(Q_1^{(k_3),a} - Q_1^{(k_3)})x\| \to 0$,

note that

$$\|(Q_1^{(k_3),a} - Q_1^{(k_3)})x\| = \|Q_1^{(k_3),b}x\| = \|Q_1^{(k_3),b*}Q_1^{(k_3),b}x\|$$

because $Q_1^{(k_3),b*}$ as a matrix of orthonormal columns is an isometry on $Im(Q_1^{(k_3),b})$. By the definition of $Q_1^{(k_3),b*}$, $Q_1^{(k_3),b*}Q_1^{(k_3),b}$ is an orthogonal projection onto

$$Im(\pi_2)/Im(Q_1^{(k_3),a*}) = Im(\pi_2)/\mathrm{Im}(AP_1U_1^{(k_3),a*})$$

which is, by construction of $U_1^{(k_3),a}$, which spans a sequence of subspaces monotonically converging to $\mathrm{Im}(P_1)$, a monotonically decreasing sequence of subsets converging to the empty set, and so goes to 0 by Helmberg (1969) §30 Theorem 5.

Combining the above, obtain $\forall x \in \mathcal{H}_2$,

$$\|(Q_1^{k_3} - Q_1^{(k_3)})x\| \leq \|Q_1^{k_3} - Q_1^{(k_3),a}\|\|x\| + \|(Q_1^{(k_3),a} - Q_1^{(k_3)})x\| \to 0,$$

as desired. By the construction of $Q_1^{(k_3)*}$ from orthonormal vectors and the invariance of operator norm under adjoints, $\forall x \in \mathrm{Im}\ Q^{(k_3)}$,

$$\|(Q_1^{k_3*} - Q_1^{(k_3)*})x\| \leq \|Q_1^{k_3*} - Q_1^{(k_3),a*}\|\|x\| + \|(Q_1^{(k_3),a*} - Q_1^{(k_3)})x\| \to 0$$

as well. □

Lemma 3.7

Proof. Part (i):

By 3.6 and the proof of 3.4

$$\begin{aligned}\|Q_1^{k_3}\tilde{A}U_1^{k_5,\varepsilon_{k_5}*} - Q_1^{(k_3)}AU_1^{(k_5),a*}\| \;\leq\; & \|Q_1^{k_3}\tilde{A}\|\|U_1^{k_5,\varepsilon_{k_5}*} - U_1^{(k_5),a*}\| + \\ & \|Q_1^{k_3}\|\|(\tilde{A} - A)U_1^{(k_5),a*}\| + \\ & \|(Q_1^{k_3} - Q_1^{(k_3)})AU_1^{(k_5),a*}\|\end{aligned}$$

goes to 0 as $k_3 \to \infty$ for fixed k_5, since $U_1^{(k_5),a*}$ is finite dimensional. By 3.2, and the invertibility of S_{11},

$$\|((Q_1^{(k_3)} A U_1^{(k_5),a*})^* Q_1^{(k_3)} A U_1^{(k_5),a*})^{-1}\|^{-1} \geq \|(S_{11}^* S_{11})^{-1}\|^{-1} > 0,$$

and so for any $\delta > 0$, there exists $k_3(k_5)$ such that

$$\|((Q_1^{k_3} \tilde{A} U_1^{k_5,\varepsilon_{k_5}*})^* Q_1^{k_3} \tilde{A} U_1^{k_5,\varepsilon_{k_5}*})^{-1}\| \leq \|(S_{11}^{(k_5)*} S_{11}^{(k_5)})^{-1}\|^{-1} + \delta.$$

So, by this and 3.4, $(S_{11}^{k_5*} S_{11}^{k_5})^{-1} S_{11}^{k_5*}$ converges in strong operator topology to $S_{11}^{(k_5)-1}$ and, similarly,

$$U_1^{(k_5)*} (S_{11}^{k_5*} S_{11}^{k_5})^{-1} S_{11}^{k_5*} Q_1^{k_3} \overset{s.o.t.}{\to} U_1^{(k_5)*} S_{11}^{(k_5)-1} Q_1^{(k_3)}.$$

Strong operator topology convergence of $T_{11}^{k_3}$ follows from boundedness and convergence of its components, again by 3.6 and the proof of 3.4, similarly by the convergence of the adjoints, $Q_1^{k_3*} T_{11}^{k_3} U_1^{k_3}$ also converges in strong operator topology to $Q_1^{(k_3)*} T_{11}^{(k_3)} U_1^{(k_3)}$. Combining these results,

$$\begin{aligned}(U_1^{k_5*} (S_{11}^{k_5*} S_{11}^{k_5})^{-1} S_{11}^{k_5*} Q_1^{k_3} (Q_1^{k_3*} T_{11}^{k_3} U_1^{k_3}) &\overset{s.o.t.}{\to} U_1^{(k_5)*} S_{11}^{(k_5)-1} Q_1^{(k_3)} Q_1^{(k_3)*} T_{11}^{(k_3)} U_1^{(k_3)} \\ &= U_1^{(k_5)*} S_{11}^{(k_5)-1} T_{11}^{(k_3)} U_1^{(k_3)}\end{aligned}$$

which holds because $Q_1^{(k_3)} Q_1^{(k_3)*}$ acts as the identity on Im $Q_1^{(k_3)} \supseteq$ Im $T_{11}^{(k_3)}$ by definition of $T_{11}^{(k_3)}$.

Part (ii):

Assume now also that Ω_1 and Ω_2 are compact operators. Consider arbitrary $x \in$ Im(P_1). By Gohberg *et al.* (1990, Theorem IV.1.1), $(Bx, Ax) = (B_1 x, A_1 x)$, and by the generalized Schur decomposition, $(Bx, Ax) = (Q_1^* T_{11} U_1 x, Q_1^* S_{11} U_1 x)$. Applying

A_1^{-1} to both, obtain

$$(\Omega_1 x, x) = (A_1^{-1} Q_1^* T_{11} U_1 x, A_1^{-1} Q_1^* S_{11} U_1 x).$$

As a result, $\forall x \in \mathrm{Im}(P_1)$ $A_1^{-1} Q_1^* T_{11} U_1 x = \Omega_1 x$, and since for all $x \in \mathcal{H}_1/\mathrm{Im}(P_1)$, $U_1 x = 0$ by construction and $U_1^* U_1$ projects onto $\mathrm{Im}(P_1)$, $A_1^{-1} Q_1^* T_{11} U_1 = \Omega_1 U_1^* U_1$. Rearranging, obtain $T_{11} = Q_1 A_1 \Omega_1 U_1^*$. Since Q_1, A_1, and U_1^* are bounded and Ω_1 is compact by assumption, T_{11} is a compact operator. Similarly, $S_{11} = Q_1 A_1 U_1^*$ and so is bounded and invertible since A_1 is, so $S_{11}^{-1} T_{11}$ is also a compact operator.

While strong operator topology convergence of $T_{11}^{k_3}$ follows from boundedness and convergence of its components, compactness allows $Q_1^{k_3 *} T_{11}^{k_3} U_1^{k_3}$ to converge in operator norm. Pre- and post-multiplication by $Q_1^{k_3 *}$ and $U_1^{k_3}$ respectively generate a quantity that depends only on the domain of U_1 and Q_1 and not the range, as $U_1^* U_1$ and $Q_1^* Q_1$ are orthogonal projections in $\mathcal{H}_1$, as are their finite-dimensional analogues. For fixed k_3 and noting that the domain of $U_1^{(k_3),a*}$ is isometrically isomorphic to k_3 dimensional Euclidean space and so we may write a basis in this space as $\{e_i\}_{i=1}^{k_3}$, we obtain

$$\begin{aligned} \|Q_1^{k_3} \tilde{B} U_1^{k_3 *} - Q_1^{(k_3),a} B U_1^{(k_3),a*}\| &\leq \\ \sup_{(\sum_{i=1}^{k_3} z_i^2)^{\frac{1}{2}}=1} \|\sum_{i=1}^{k_3} (Q_1^{k_3} \tilde{B} U_1^{k_3 *} - Q_1^{(k_3),a} B U_1^{(k_3),a*}) z_i e_i\| &\leq \\ \sum_{i=1}^{k_3} \|(Q_1^{k_3} \tilde{B} U_1^{k_3 *} - Q_1^{(k_3),a} B U_1^{(k_3),a*}) e_i\| &\to 0 \end{aligned}$$

where the last line follows from the triangle inequality, the fact that $|z_i| \leq 1$ $\forall i$, and boundedness and pointwise convergence. Defining $T_{11}^{(k_3),a} = Q_1^{(k_3),a} B U_1^{(k_3),a*}$, this may be expressed as $\|T_{11}^{k_3} - T_{11}^{(k_3),a}\| \to 0$ for fixed k_3. Pre- and post- multiplying to get an operator acting on $\mathcal{H}_1$, obtain

$$\begin{aligned}\|Q_1^{k_3*}T_{11}^{k_3}U_1^{k_3} - Q_1^{(k_3),a*}T_{11}^{(k_3),a}U_1^{(k_3),a}\| \quad &\leq \\ \|Q_1^{k_3*}\|\|U_1^{k_3}\|\|T_{11}^{k_3} - T_{11}^{(k_3),a}\| + \|U_1^{k_3}\|\|(Q_1^{k_3*} - Q_1^{(k_3),a*})T_{11}^{(k_3),a}\| & \\ +\|Q_1^{(k_3),a*}T_{11}^{(k_3),a}(U_1^{k_3} - U_1^{(k_3),a})\| \quad &\to 0\end{aligned}$$

as $k_1(k_2)$ and $k_2 \to \infty$ for fixed k_3. The first term goes to 0 since each element is bounded and $\|T_{11}^{k_3} - T_{11}^{(k_3),a}\| \to 0$. The second term goes to 0 by again bounding the supremum by a sum over a finite set times the maximum coefficient and applying the triangle inequality, then using 3.6 to ensure pointwise convergence. The final term may be seen to be bounded by $\|Q_1^{(k_3),a*}T_{11}^{(k_3),a}\|\|U_1^{k_3} - U_1^{(k_3),a}\|$, where the first term is bounded and the second was shown to converge to 0 for fixed k_3 in the proof of 3.4. We then have that

$$\begin{aligned}Q_1^{(k_3),a*}T_{11}^{(k_3),a}U_1^{(k_3),a} = Q_1^{(k_3),a*}Q_1^{(k_3),a}BU_1^{(k_3),a*}U_1^{(k_3),a} = & \\ Q_1^{(k_3),a*}Q_1^{(k_3),a}Q_1^{(k_3)*}Q_1^{(k_3)}BU_1^{(k_3)*}U_1^{(k_3)}U_1^{(k_3),a*}U_1^{(k_3),a} = & \\ Q_1^{(k_3),a*}Q_1^{(k_3),a}Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}U_1^{(k_3),a*}U_1^{(k_3),a} &\end{aligned}$$

is a compression of the compact operator $Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}$ projecting domain and range onto a sequence of subspaces monotonically increasing to the full space and so $\|Q_1^{(k_3),a*}T_{11}^{(k_3),a}U_1^{(k_3),a} - Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}\| \to 0$ as $k_3 \to \infty$. Combining these two results, $\|Q_1^{k_3*}T_{11}^{k_3}U_1^{k_3} - Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}\| \to 0$ as $k_3 \to \infty$.

In the case where $\sigma_{k_6}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}) > \sigma_{k_6+1}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)})$, convergence of $\|\text{thresh}_{k_6}(Q_1^{k_3*}T_{11}^{k_3}U_1^{k_3}) - \text{thresh}_{k_6}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)})\|$ to 0 as $k_3 \to \infty$ then follows by the consistency of the singular value decomposition with respect to the operator norm, and in particular, Weyl's inequality for singular values and the SinΘ theorem

for singular vectors as in the proof of 3.4, implying

$$\|\text{thresh}_{k_6}(Q_1^{k_3*}T_{11}^{k_3}U_1^{k_3}) - \text{thresh}_{k_6}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)})\| \leq$$
$$C/(\sigma_{k_6}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}) - \sigma_{k_6+1}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}))\|Q_1^{k_3*}T_{11}^{k_3}U_1^{k_3} - Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}\|$$

where C is a universal constant. Note also that by compactness of $Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}$, the spectrum of $(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)})^*Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}$ is discrete, decays to 0, and all eigenvalues have finite multiplicity. As a result, there exists an infinite subsequence $\{k_{6,i}\}_{i=1}^{\infty}$ of singular values, with indices each separated by a finite integer, along which $\sigma_{k_{6,i}}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}) > \sigma_{k_{6,i}+1}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)})$ holds, and this subsequence also converges to 0, so we may take this as our increasing sequence k_6. We then have that, because

$$\|\text{thresh}_{k_6}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}) - Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}\| \to 0$$

as $k_6 \to \infty$ since this is an increasing sequence of compressions converging to a compact operator,

$$\|\text{thresh}_{k_6,i}(Q_1^{k_3*}T_{11}^{k_3}U_1^{k_3}) - Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}\| \to 0$$

as $i \to \infty$.

Note that for fixed k_6, $\text{thresh}_{k_6}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)})$ is a finite dimensional object which, despite the notation, does not depend on the value of k_3. This is important because k_3 must grow with k_6 fixed for convergence to hold, and so a finite dimensional object like $\text{thresh}_{k_6}(T_{11}^{(k_3)})$ which does depend on k_3, when a pointwise convergent operator is applied to it, is not guaranteed to converge in norm. Combining the

above, obtain

$$\|U_1^{(k_5)*}(S_{11}^{k_5*}S_{11}^{k_5})^{-1}S_{11}^{k_5*}Q_1^{k_3}\text{thresh}_{k_6}(Q_1^{k_3*}T_{11}^{k_3}U_1^{k_3})- \\ U_1^{(k_5)*}S_{11}^{(k_5)-1}Q_1^{(k_3)}\text{thresh}_{k_6}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)})\| \to 0$$

for fixed k_6. Since

$$\|U_1^{(k_5)*}S_{11}^{(k_5)-1}Q_1^{(k_3)}\text{thresh}_{k_6}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}) - U_1^{(k_5)*}S_{11}^{(k_5)-1}T_{11}^{(k_3)}U_1^{(k_3)}\| \leq \\ \|U_1^{(k_5)*}S_{11}^{(k_5)-1}Q_1^{(k_3)}\|\|\text{thresh}_{k_6}(Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}) - Q_1^{(k_3)*}T_{11}^{(k_3)}U_1^{(k_3)}\|$$

and the first component is bounded while the second goes to 0 as $k_6 \to \infty$, the final part of the lemma holds. □

Theorem 3.1

Proof. Part (i):

Consistency of $g_X^{k_4}$ follows from application of 3.2 and the proof of 3.5. In particular, for fixed k_4 as k_3 increases, $\|U_2^{k_4,\varepsilon_{k_4}} - U_2^{(k_4),a}\| \to 0$, and so by the continuity of singular values and the boundedness of φ^Y, for any k_4 and any $\delta > 0$, there exists K_3 such that $\forall k_3 > K_3$

$$\left|\sigma_{\min}(U_2^{k_4,\varepsilon_{k_4}}\varphi^Y\varphi^{Y*}U_2^{k_4,\varepsilon_{k_4}*}) - \sigma_{\min}(U_2^{(k_4),a}\varphi^Y\varphi^{Y*}U_2^{(k_4),a*})\right| < \delta.$$

Then by the assumption that $U_2\varphi^Y\varphi^{Y*}U_2^*$ has bounded inverse and 3.2, since

$$U_2^{(k_4),a}\varphi^Y\varphi^{Y*}U_2^{(k_4),a*}$$

is a compression of $U_2\varphi^Y\varphi^{Y*}U_2^*$, it satisfies

$$\sigma_{\min}(U_2^{(k_4),a}\varphi^Y\varphi^{Y*}U_2^{(k_4),a*}) \geq \|(U_2\varphi^Y\varphi^{Y*}U_2^*)^{-1}\|^{-1}$$

for all k_4. Thus, there exists a sequence of $k_3(k_4)$ such that

$$\sigma_{\min}(U_2^{k_4,\varepsilon_{k_4}}\varphi^Y\varphi^{Y*}U_2^{k_4,\varepsilon_{k_4}*}) > \|(U_2\varphi^Y\varphi^{Y*}U_2^*)^{-1}\|^{-1} - \delta > 0$$

for all k_4 and so $\|(U_2^{k_4,\varepsilon_{k_4}}\varphi^Y\varphi^{Y*}U_2^{k_4,\varepsilon_{k_4}*})^{-1}\|$ is uniformly bounded. $U_2^{k_4,\varepsilon_{k_4}}$ converges in strong operator topology to $U_2^{(k_4)}$ by 3.5, and, furthermore, $U_2^{k_4,\varepsilon_{k_4}*}$ converges pointwise to $U_2^{(k_4)*}$ by steps identical to those used to show $U_1^{k_3,\varepsilon_{k_3}*} \overset{s.o.t.}{\to} U_1^{(k_3)*}$ in the proof of 3.5. Since each component in the composition is uniformly bounded in operator norm and converges in strong operator topology, $x \in \mathcal{H}_X$, there is a sequence along which

$$\|(g_X^{k_4} + \varphi^{Y*}U_2^{(k_4)*}(U_2^{(k_4)}\varphi^Y\varphi^{Y*}U_2^{(k_4)*})^{-1}U_2^{(k_4)}\varphi^X)x\| \to 0.$$

Note that for any $x \in \mathcal{H}_X$, $-\varphi^{Y*}U_2^{(k_4)*}(U_2^{(k_4)}\varphi^Y\varphi^{Y*}U_2^{(k_4)*})^{-1}U_2^{(k_4)}\varphi^X x$ is identical for any k_4 regardless of the basis used to construct $U_2^{(k_4)}$ and so there is no loss in dropping the superscript in the notation for the limit object g_X.

Convergence of $h_X^{k_4,k_5}$ and $h_X^{k_6}$ depend on the convergence of $g_X^{k_4}$ but also of its adjoint. By the invariance of operator norm with respect to adjoints and the pointwise convergence of both $U_2^{k_4,\varepsilon_{k_4}}$ and $U_2^{k_4,\varepsilon_{k_4}*}$, $g_X^{k_4*} \overset{s.o.t.}{\to} g_X^*$ by the boundedness and pointwise convergence of its components. Strong operator topology convergence of $\varphi^{X*}\varphi^X + g_X^{k_4*}g_X^{k_4}$ to $\varphi^{X*}\varphi^X + g_X^*g_X$ then follows as well. To see that this convergence is stable, note that $\varphi^{X*}\varphi^X + g_X^{k_4*}g_X^{k_4}$ is a sum of quadratic forms and so is a self-adjoint positive operator. Weyl's inequality then gives that

$$\sigma_{\min}(\varphi^{X*}\varphi^X + g_X^{k_4*}g_X^{k_4}) \geq \sigma_{\min}(\varphi^{X*}\varphi^X) + \sigma_{\min}(g_X^{k_4*}g_X^{k_4}) \geq 1 + 0 = 1,$$

so $\|(\varphi^{X*}\varphi^X + g_X^{k_4*}g_X^{k_4})^{-1}\| \leq 1$ for all k_4.

Strong operator topology convergence of $h_X^{k_4,k_5}$ follows from the triangle inequality

and boundedness and pointwise convergence of

$$(\varphi^{X*}\varphi^{X} + g_X^{k_4*} g_X^{k_4})^{-1}$$

$$(g_X^{k_4*}\varphi^{Y*} + \varphi^{X*})$$

$$U_1^{k_5*}(S_{11}^{k_5*} S_{11}^{k_5})^{-1} S_{11}^{k_5*} Q_1^{k_3} (Q_1^{k_3*} T_{11}^{k_3} U_1^{k_3})$$

$$\text{and } (\varphi^{X} + \varphi^{Y} g_X^{k_4})$$

by the above and 3.7 part (i). Note that as k_4 and k_5 appear in sequentially applied components, k_4 need not be taken as a function of k_5 or vice versa, so long as k_3 (and subsequently k_2 and k_1) are chosen as at least the maximum required value for pointwise convergence over each component.

Part (ii):

Operator norm convergence for $h_X^{k_6}$ follows by applying pointwise convergence of $g_X^{k_4}$ and

$$U_1^{k_5*}(S_{11}^{k_5*} S_{11}^{k_5})^{-1} S_{11}^{k_5*} Q_1^{k_3}$$

on the finite dimensional operator $\text{thresh}_{k_6}(Q_1^{(k_3)*} T_{11}^{(k_3)} U_1^{(k_3)})$ to achieve norm convergence for any finite dimension k_6, and then letting k_6 grow. Note that this requires k_4 to be large relative to k_6, but does not impose that k_4 be large relative to k_5, as pointwise consistent operators depending on k_4 and k_5 are not constructed sequentially. In summary, the sequences k_j which ensure norm convergence of $h_X^{k_6}$ satisfy $k_1 \gg k_2 \gg k_3 \gg k_4 \gg k_6$ and $k_1 \gg k_2 \gg k_3 \gg k_5 \gg k_6$, where $a \gg b$ indicates that for each fixed b in the sequence, there is some a needed as a function of b which may be arbitrarily large. Informally, this means that a is taken to be large relative to b.

For fixed k_6, obtain

$$\|h_X^{k_6} - h_X\| \leq$$
$$(\|(\varphi^{X*}\varphi^X + g_X^{k_4*}g_X^{k_4})^{-1}(g_X^{k_4*}\varphi^{Y*} + \varphi^{X*})U_1^{k_5*}(S_{11}^{k_5*}S_{11}^{k_5})^{-1}S_{11}^{k_5*}Q_1^{k_3}\| + \|\varphi^X + \varphi^Y g_X^{k_4}\|)*$$
$$\|\text{thresh}_{k_6}(Q_1^{k_3*}T_{11}^{k_3}U_1^{k_3}) - \text{thresh}_{k_6}(Q_1^*T_{11}U_1)\|+$$
$$\|(\varphi^{X*}\varphi^X + g_X^{k_4*}g_X^{k_4})^{-1}(g_X^{k_4*}\varphi^{Y*} + \varphi^{X*})U_1^{k_5*}(S_{11}^{k_5*}S_{11}^{k_5})^{-1}S_{11}^{k_5*}Q_1^{k_3}\|*$$
$$\|\text{thresh}_{k_6}(Q_1^*T_{11}U_1)((\varphi^X + \varphi^Y g_X^{k_4}) - (\varphi^X + \varphi^Y g_X))\|+$$
$$\|\varphi^X + \varphi^Y g_X\| * \|[(\varphi^{X*}\varphi^X + g_X^{k_4*}g_X^{k_4})^{-1}(g_X^{k_4*}\varphi^{Y*} + \varphi^{X*})U_1^{k_5*}(S_{11}^{k_5*}S_{11}^{k_5})^{-1}S_{11}^{k_5*}Q_1^{k_3})-$$
$$(\varphi^{X*}\varphi^X + g_X^*g_X)^{-1}(g_X^*\varphi^{Y*} + \varphi^{X*})U_1^*S_{11}^{-1}Q_1]\text{thresh}_{k_6}(Q_1^*T_{11}U_1)\|+$$
$$\|(\varphi^{X*}\varphi^X + g_X^*g_X)^{-1}(g_X^*\varphi^{Y*} + \varphi^{X*})\| * \|\varphi^X + \varphi^Y g_X\|*$$
$$\|U_1^*S_{11}^{-1}Q_1\text{thresh}_{k_6}(Q_1^*T_{11}U_1) - U_1^{(k_5)*}S_{11}^{(k_5)-1}T_{11}^{(k_3)}U_1^{(k_3)}\|$$

Uniform boundedness of $g_X^{k_4}$, $(\varphi^{X*}\varphi^X + g_X^{k_4*}g_X^{k_4})^{-1}$, and $U_1^{k_5*}(S_{11}^{k_5*}S_{11}^{k_5})^{-1}S_{11}^{k_5*}Q_1^{k_3*}$ and 3.7 ensure that the first term in the sum goes to 0. For the second term, note that $\text{thresh}_{k_6}(Q_1^*T_{11}U_1) = \sum_{i=1}^{k_6} d_i \langle v_i, (.)\rangle u_i$ for some singular values d_i and singular vectors u_i and v_i, and so

$$\|\text{thresh}_{k_6}(Q_1^*T_{11}U_1)((\varphi^X + \varphi^Y g_X^{k_4}) - (\varphi^X + \varphi^Y g_X))\| =$$
$$\|\sum_{i=1}^{k_6} d_i \left\langle v_i, ((\varphi^X + \varphi^Y g_X^{k_4}) - (\varphi^X + \varphi^Y g_X))(.)\right\rangle u_i\| =$$
$$\|\sum_{i=1}^{k_6} d_i \left\langle ((g_X^{k_4*}\varphi^{Y*} - \varphi^{X*}) - (g_X^*\varphi^{Y*} + \varphi^{X*}))v_i, (.)\right\rangle u_i\| \leq$$
$$\sup_{i=1\ldots k_6} \|((g_X^{k_4*}\varphi^{Y*} - \varphi^{X*}) - (g_X^*\varphi^{Y*} + \varphi^{X*}))v_i\| \sum_{i=1}^{k_6} |d_i|$$

For fixed k_6, $\sup_{i=1\ldots k_6} \|((g_X^{k_4*}\varphi^{Y*} - \varphi^{X*}) - (g_X^*\varphi^{Y*} + \varphi^{X*}))v_i\| \to 0$ by strong operator topology convergence of $g_X^{k_4*}$, and $\sum_{i=1}^{k_6} |d_i|$ is bounded, so this term goes to 0 also. Note that the rate at which k_4 must grow relative to k_6 depends on, among other

issues, the (partial) sum of the singular values of $Q_1 T_{11} U_1$. By compactness, $|d_i| \to 0$ as i grows, so $\sum_{i=1}^{k_6} |d_i| = O(k_6)$ at worst. In the case where T_{11} is a trace-class operator, which is not uncommon in practice, this term is instead $O(1)$ and the rate is determined by the growth of $\sup_{i=1\ldots k_6} \|((g_X^{k_4 *}\varphi^{Y*} - \varphi^{X*}) - (g_X^* \varphi^{Y*} + \varphi^{X*}))v_i\|$. For the third term, we again have that by the finite dimensionality of $\mathrm{thresh}_{k_6}(Q_1 T_{11} U_1)$ and the strong operator topology convergence of $U_1^{k_5 *}(S_{11}^{k_5 *} S_{11}^{k_5})^{-1} S_{11}^{k_5 *} Q_1^{k_3 *}$ due to 3.7 and of $(\varphi^{X*}\varphi^X + g_X^{k_4 *} g_X^{k_4})^{-1}$ and $g_X^{k_4 *}\varphi^{Y*} + \varphi^{X*}$, we may find for any fixed k_6 a k_5 and a k_4 such that this term is smaller than any given $\epsilon > 0$. The final summand goes to 0 by 3.7. $\square$

Bibliography

Acemoglu, Daron, & Jensen, Martin Kaae. 2012 (June). *Robust Comparative Statics in Large Dynamic Economies*. Working Paper 18178. National Bureau of Economic Research.

Achdou, Yves, Lasry, Jean-Michel, Lions, Pierre-Lous, & Moll, Benjamin. 2013. *Heterogeneous Agent Models in Continuous Time*. Tech. rept. Princeton University.

Adcock, Ben, Hansen, Anders, Roman, B., & Teschke, G. 2013. Generalized sampling: stable reconstructions, inverse problems and compressed sensing over the continuum. *arXiv*.

Adcock, Ben, Hansen, Anders, Roman, Bogdan, & Teschke, Gerd. 2014a. Generalized sampling: Stable reconstructions, inverse problems and compressed sensing over the continuum. *Advances in Imaging and Electron Physics*, **182**(viii), 187–279.

Adcock, Ben, Hansen, Anders, & Poon, Clarice. 2014b. On optimal wavelet reconstructions from Fourier samples: linearity and universality of the stable sampling rate. *Applied and Computational Harmonic Analysis*, **36**(3), 387–415.

Aiyagari, S. Rao. 1994. Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, **109**(3), 659–684.

Akerlof, George A., & Kranton, Rachel E. 2000. Economics and identity. *Quarterly Journal of Economics*, 715–753.

Allen, Treb, & Arkolakis, Costas. 2014. Trade and the Topography of the Spatial Economy. *Quarterly Journal of Economics*.

Allen, Treb, & Arkolakis, Costas. 2015. *Lecture 8 : Trade and Migration*.

Andrews, Donald W. K. 2011. Examples of L2-Complete and Boundedly Complete Distributions. *Cowles Foundation Working Paper*, May.

Artuç, Erhan, Chaudhuri, Shubham, & McLaren, John. 2010. Trade Shocks and Labor Adjustment: A Structural Empirical Approach. *The American Economic Review*, **100**(3), 1008–1045.

Benatia, David, Carrasco, Marine, & Florens, Jean-Pierre. 2015. *Functional linear regression with functional response*.

Bergin, J., & Bernhardt, D. 1995. Anonymous sequential games: Existence and characterization of equilibria. *Economic Theory*, **5**, 461–489.

Berlinet, Alain, & Thomas-Agnan, Christine. 2004. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Kluwer.

Bewley, Truman. 1986. Stationary monetary equilibrium with a continuum of independently fluctuating consumers. *In: Contributions to Mathematical Economics in Honor of Gerard Debreu*.

Beylkin, G., Coifman, R., & Rokhlin, V. 1991. *Fast Wavelet Transforms and Numerical Algorithms I*.

Bhatia, R., & Rosenthal, P. 1997. How and Why to Solve the Operator Equation AX-XB = Y. *Bulletin of the London Mathematical Society*, **29**(1), 1–21.

Blanchard, Olivier Jean, & Kahn, Charles M. 1980. The Solution of Linear Difference Models under Rational Expectations. *Econometrica*, **48**(5), pp. 1305–1311.

Bosq, Denis. 2000. *Linear Processes in Function Spaces: Theory and Applications*. Springer.

Boyd, John P. 2000. *Chebyshev and Fourier Spectral Methods*. 2nd edn. Mineola, NY: Dover Publications.

Caballero, Ricardo J., & Engel, Eduardo M. R. a. 1999. Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach. *Econometrica*, **67**(4), 783–826.

Caliendo, Lorenzo, Dvorkin, Maximiliano, & Parro, Fernando. 2015. *Trade and Labor Market Dynamics*.

Canay, Ivan A, Santos, Andres, & Shaikh, Azeem M. 2013. On the testability of identification in some nonparametric models with endogeneity. *Econometrica*, **81**(6), 2535–2559.

Carl, Bernd, & Stephani, Imtraud. 1990. *Entropy, compactness, and the approximation of operators*. Cambridge Tracts in Mathematics, no. 98. Cambridge University Press.

Chang, Yoosoon, Kim, Chang Sik, & Park, Joon Y. 2014. *Nonstationarity in Time Series of State Densities*. Tech. rept. Indiana University.

Chatelin, Françoise. 2011. *Spectral Approximation of Linear Operators*. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics.

Chen, Xiaohong, & Christensen, Timothy M. 2015. Optimal uniform convergence rates and asymptotic normality for series estimators under weak dependence and weak conditions I. *Journal of Econometrics*, **188**(2), 447–465.

Christensen, Timothy. 2014. Nonparametric Identification of Positive Eigenfunctions. *Econometric Theory*.

Chung, Hess. 2007. *Perturbative Solutions for Heterogeneous Agent DSGE Models*. Indiana University Department of Economics.

Çivril, A. 2014. Column Subset Selection Problem is UG-hard. *Journal of Computer and System Sciences*, 849–859.

Cochrane, John. 2011. Determinacy and Identification with Taylor Rules. *Journal of Political Economy*, **119**(3), 565–615.

Coifman, Ronald R., & Maggioni, Mauro. 2006. Diffusion Wavelets. *Applied and Computational Harmonic Analysis*, **21**(1), 53–94.

Conway, John B. 1978. *Functions of One Complex Variable*. Graduate Texts in Mathematics, no. 11. Springer.

Craven, B.D., & Nashed, M.Z. 1982. Generalized Implicit Function Theorems When the Derivative Has No Bounded Inverse. *Nonlinear Analysis, Theory, Methods, and Applications*, **6**(4), 375–387.

Den Haan, Wouter J. 2010. Comparison of solutions to the incomplete markets model with aggregate uncertainty. *Journal of Economic Dynamics and Control*, **34**(1), 4–27.

Diamond, Rebecca. 2016. The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000. *American Economic Review*, **106**(3), 479–524.

Dornbusch, Rudiger. 1976. Expectations and Exchange Rate Dynamics. *Journal of Political Economy*, 1161–1176.

Ferraty, Frédéric, & Romain, Yves. 2011. *The Oxford Handbook of Functional Data Analysis*.

Flett, T.M. 1980. *Differential analysis : differentiation, differential equations, and differential inequalities*. Cambridge: Cambridge University Press.

Glaeser, Edward L., & Gottlieb, Joshua D. 2009. The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States. *Journal of Economic Literature*, **47**(4), 983–1028.

Gohberg, Israel, Goldberg, Seymour, & Kaashoek, Marinus A. 1990. *Classes of Linear Operators Vol. I*. Operator Theory: Advances and Applications, vol. 49. Birkhauser Verlag.

Gohberg, Israel, Goldberg, Seymour, & Kaashoek, Marinus A. 1993. *Classes of Linear Operators Vol. II*. Operator Theory: Advances and Applications, vol. 63. Birkhauser Verlag.

Golub, Gene H., & van Loan, Charles F. 1996. *Matrix Computations*. Third edn. Johns Hopkins University Press.

Golub, Gene H., & Van Loan, Charles F. 1996. *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press.

Gordon, Grey. 2011. *Computing Dynamic Heterogeneous-Agent Economies: Tracking the Distribution*.

Guillas, Serge. 2001. Rates of convergence of autocorrelation estimates for autoregressive Hilbertian processes. *Statistics and Probability Letters*, 281–291.

Guvenen, Fatih. 2011 (November). *Macroeconomics With Heterogeneity: A Practical Guide*. Working Paper 17622. National Bureau of Economic Research.

Guvenen, Fatih, Ozkan, Serdar, & Song, Jae. 2012. *The Nature of Countercyclical Income Risk*. Tech. rept. Working Paper No. 18035. National Bureau of Economic Research.

Hairer, Martin. 2009. An Introduction to Stochastic PDEs.

Heathcote, Jonathan, Storesletten, Kjetil, & Violante, Giovanni L. 2009. Quantitative Macroeconomics with Heterogeneous Households. *Annual Review of Economics*.

Helmberg, Gilbert. 1969. *Introduction to Spectral Theory in Hilbert Space*. Dover edn. New York: Wiley.

Hong, Y.P., & Pan, C.-T. 1992. Rank Revealing QR Factorizations and the Singular Value Decomposition. *Mathematics of Computation*, **58**(197), 213–232.

Horowitz, Joel L. 2013. Ill-Posed Inverse Problems in Economics. *Annual Review of Economics*, **6**(1), 140128174544008.

Horváth, Lajos, & Kokoszka, Piotr. 2012. *Inference for Functional Data with Applications*. Springer Series in Statistics, no. 200. Springer.

Huggett, Mark. 1993. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, **17**(5–6), 953 – 969.

Jin, He-Hui, & Judd, Kenneth. 2002. *Perturbation methods for general dynamic stochastic models*. Tech. rept. Stanford University.

Johnstone, Iain M. 2013. *Gaussian estimation: Sequence and wavelet models*. Unpublished manuscript.

Kato, Tosio. 1976. *Perturbation Theory for Linear Operators*. 2 edn. Springer.

Kesavan, S. 2004. *Nonlinear Functional Analysis: A First Course*. New Delhi: Hindustan Book Agency.

Klein, Paul. 2000. Using the generalized Schur form to solve a multivariate linear rational expectations model. *Journal of Economic Dynamics and Control*, **24**(10), 1405 – 1423.

Krieger, Stefan. 2001. *Bankruptcy Costs, Financial Constraints, and the Business Cycle*. Tech. rept. Yale University.

Krugman, Paul. 1979. A Model of Balance-of-Payments Crises. *Journal of Money, Credit, and Banking*, 311–325.

Krugman, Paul. 1991. Increasing Returns and Economic Geography. *Journal of Political Economy*, **99**(3), 483–499.

Krugman, Paul. 1996. *The Self-Organizing Economy*. Cambridge, MA: Blackwell Publishers.

Krusell, Per, & Smith, Anthony A. 1998. Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy*, **106**(5), pp. 867–896.

Krusell, Per, & Smith, Anthony A. 2006. Quantitative Macroeconomic Models with Heterogeneous Agents. *In:* Blundell, Richard, Newey, Whitney K., & Persson, Torsten (eds), *Advances in Economics and Econometrics: Theory and Applications*, vol. I. Cambridge University Press.

Kurbatova, I. V. 2009. Banach algebras associated with linear operator pencils. *Mathematical Notes*, **86**(3-4), 361–367.

Ljungqvist, Lars, & Sargent, Thomas. 2004. *Recursive Macroeconomic Theory*. 2nd edn. MIT Press. Chap. 17 Incomplete Markets Models.

Maddison, Chris J, Tarlow, Daniel, & Minka, Tom. 2014. A* Sampling. *Advances in Neural Information Processing Systems*, 3086–3094.

Mallat, Stephane. 2008. *A wavelet tour of signal processing: the sparse way*. Third edn. Academic press.

Marshall, Alfred. 1890. Principles of political economy. *Maxmillan, New York*.

Mertens, Thomas, & Judd, Kenneth. 2012. *Equilibrium Existence and Approximation for Incomplete Market Models with Substantial Heterogeneity*. Tech. rept. New York University.

Miao, Jianjun. 2006. Competitive equilibria of economies with a continuum of consumers and aggregate shocks. *Journal of Economic Theory*, **128**(1), 274 – 298.

Mohri, Mehryar, Rostamizadeh, Afshin, & Talwalkar, Ameet. 2012. *Foundations of machine learning*. MIT Press.

Morris, Jeffrey S. 2014. *Functional Regression*. Tech. rept. arXiv:1406.4068v1. The University of Texas, MD Anderson Cancer Center.

Park, Joon Y, & Qian, Junhui. 2012. Functional regression of continuous state distributions. *Journal of Econometrics*, **167**(2), 397–412.

Piyapromdee, Suphanit. 2015 (jun). *The Impact of Immigration on Wages, Internal Migration and Welfare*.

Pollard, David. 2002. *A User's Guide to Measure Theoretic Probability*. Cambridge University Press.

Reiter, Michael. 2009. Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, **33**(3), 649 – 665.

Roback, Jennifer. 1982. Wages, Rents, and the Quality of Life. *Journal of Political Economy*, **90**(6), 1257–1278.

Rotemberg, Julio J. 1982. Monopolistic Price Adjustment and Aggregate Output. *Review of Economic Studies*, **49**(4), 517–531.

Santos, Andres. 2012. Inference in Nonparametric Instrumental Variables With Partial Identification. *Econometrica*, **80**(1), 213–275.

Schmitt-Grohe, Stephanie, & Uribe, Martin. 2004. Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control*, **28**(4), 755 – 775.

Sethna, James P. 2006. *Statistical Mechanics: Entropy, Order Parameters, and Complexity*. Oxford UP.

Skorohod, A.V. 1984. *Random Linear Operators*. D. Reidel Pub. Co.

Stewart, G W. 1973. Error and Perturbation Bounds for Subspaces Associated with Certain Eigenvalue Problems. *SIAM Review*, **15**(4), 727–764.

Stewart, G.W., & Sun, Ji-guang. 1990. *Matrix Perturbation Theory*. Boston Academic Press.

Stokey, Nancy L. 2008. *The Economics of Inaction: Stochastic Control Models with Fixed Costs*. Princeton University Press.

Sweldens, Wim, & Piessens, Robert. 1994. Quadrature Formulae and Asymptotic Error Expansions for Wavelet Approximations of Smooth Functions. *SIAM Journal on Numerical Analysis*, **31**(4), 1240–1264.

Tao, Terence. 2011. *Topics in random matrix theory*. Chap. 1.3 Eigenvalues and sums of Hermitian matrices.

Townsend, Alex, & Trefethen, Lloyd N. 2014. Continuous analogues of matrix factorizations. *SIAM Review*.

Veracierto, Marcelo. 2014. Adverse Selection, Risk Sharing and Business Cycles. *Federal Reserve Bank of Chicago Working Paper*.

Wei, Dong. 1998. *Coiflet Type Wavelets: Theory, Design, and Application*. Ph.D. thesis, University of Texas at Austin.

Winberry, Thomas. 2014. *Lumpy Investment, Business Cycles, and Stimulus Policy.*