

Self-justified equilibria: Existence and Computation*

Felix Kubler

DBF, University of Zurich
and Swiss Finance Institute
fkubler@gmail.com

Simon Scheidegger

Department of Finance, HEC Lausanne
simon.scheidegger@unil.ch

October 22, 2019

Abstract

In this paper, we introduce the concept of “self-justified equilibria” as a tractable alternative to rational expectations equilibria in stochastic general equilibrium models with heterogeneous agents. A self-justified equilibrium is a temporary equilibrium where, in each period, agents trade in assets and commodities to maximize the sum of current utility and expected future utilities that are forecasted on the basis of current endogenous variables and the current exogenous shock. Agents’ characteristics include a loss function that prescribes how the agent trades off the accuracy and the computational complexity of possible forecasts. We provide sufficient conditions for the existence of self-justified equilibria, and we develop a computational method to approximate them numerically. For this, we focus on a convenient special case where we use Gaussian process regression coupled to active subspaces to model agents’ forecasts. We demonstrate that this framework allows us to solve stochastic overlapping generations models with hundreds of heterogeneous agents and very accurate forecasts.

Keywords: Dynamic General Equilibrium, Rational Expectations, Active Subspaces.

JEL Classification: C63, C68, D50, D52.

*We thank Filipe Martins-da-Rocha, Michael Reiter, Victor Rios-Rull, Tom Sargent, Tony Smith, Thomas Winberry, as well as seminar participants at various universities and conferences for useful conversations and comments. Moreover, this work was generously supported by grants from the Swiss National Supercomputing Centre (CSCS) under project ID s885, the Swiss Platform for Advanced Scientific Computing (PASC) under project ID “Computing equilibria in heterogeneous agent macro models on contemporary HPC platforms”. Simon Scheidegger gratefully acknowledges support from the MIT Sloan School of Management.

1 Introduction

In this paper, we develop an alternative to rational expectations equilibria and consider temporary equilibria with forecasting functions that approximate the temporary equilibrium correspondence, but that might lead to imprecise forecasts at any given time. We derive simple sufficient conditions that ensure the existence of these “self-justified” equilibria, we develop an algorithm to approximate self-justified equilibria numerically, and we give an example that shows that by restricting the complexity of agents’ forecasts one can often solve models with very many agents.

The basic idea of our proposed approach is as follows: In a temporary equilibrium, agents use current endogenous variables and the exogenous shock to forecast future marginal utilities for assets; prices for commodities and assets in the current period ensure that markets clear. Forecasting functions are assumed to lie in a pre-specified class, and an agent chooses a function that minimizes the agents’ loss function, which depends on the mean-squared long-run errors of the forecasts and on the computational complexity of the forecasting function.

The assumption of rational expectations and the use of recursive methods to analyze dynamic economic models has revolutionized financial economics, macroeconomics, and public finance (see, e.g., Ljungqvist and Sargent (2000)). Discrete-time, infinite-horizon, general equilibrium models with heterogeneous agents and incomplete financial markets now play an important role in macroeconomics and public finance. Unfortunately, for these stochastic general equilibrium models with heterogeneous agents, rational expectations equilibria are generally not tractable, computational methods to approximate these equilibria numerically are often ad-hoc, and a rigorous error analysis seems impossible. In particular, it is generally not possible to make statements about how close an approximate equilibrium is to an exact equilibrium (see Kubler (2011)). Therefore computational methods typically focus on computing rational expectation ϵ -equilibria, i.e., allocations and prices that clear markets and satisfy agents optimality conditions (i.e., Euler equations) up to some $\epsilon > 0$. Errors in Euler equations provide a good method to analyze solutions to dynamic optimization problems (see Santos (2000) or Judd (1992)). However, they become meaning-less in models with heterogeneous agents: Agents’ incorrect choices need to be coordinated to ensure the definition of an approximate equilibrium. Agents’ mistakes cannot be systematic or random, but they are determined by the requirement that at any time, agents rational expectations from the previous period need to turn out to be correct this period and that markets clear. In an approximation that is theoretically sound, one would hope that it is explicit what determines agents choices and that the fact that they make mistakes is part of the model.

As Sargent et al. (1993) points out, “when implemented numerically ... rational expectations models impute more knowledge to the agent within the model ... than is possessed by an econometrician”, and a sensible approach to relax rational expectations is “expelling rational agents from our model environment and replacing them with ‘artificially intelligent’ agents who behave like

econometricians.”

This quote motivates the idea underlying self-justified equilibria—to construct a tractable model of the macro-economy that takes into account substantial heterogeneity across agents, one needs to assume that the modeler can compute agents’ expectations. As argued above, for the case of rational expectations this is, in general, only possible if one forces agents to make the “right mistakes” that ensure that previous expectations turn out to be correct. In the temporary equilibrium we consider in this paper, the agents might make significant mistakes in their forecasts, but these stem from the definition of equilibrium. They make no mistakes in choices, given their forecasts, and the only reason that prevents agents’ forecasts from being arbitrarily accurate is the computational cost associated with more accurate forecasts.

In order to ensure a mathematically simple concept and the existence of a stationary equilibrium, we need to assume that our agents are “hyper-intelligent” in the sense that they solve optimization problems and integrals exactly.

We introduce the concept of self-justified equilibria in the context of an infinite horizon pure exchange economy with overlapping generations, a single perishable commodity, and aggregate uncertainty. This allows us to investigate the properties of a self-justified equilibrium with as little notation as possible. An extension to production economies with several commodities and to economies with infinitely lived agents is conceptually straightforward (e.g., along the lines of Brumm et al. (2017)).

To prove the existence of a self-justified equilibrium, we make the simplifying assumption that accounting is finite. That is to say, we assume that beginning-of-period portfolios across agents lie on some finite (arbitrarily fine) grid, that agents’ portfolio-choices in the current period induce a probability distribution over next period’s state, and that the support of this distribution is a subset of the grid of possible beginning-of-period portfolios across agents. This assumption can be viewed as a technical approximation to a continuous model, but one can also think of bounded rationality justifications. Our preferred interpretation is that at the beginning of each period, there are (small) transfers across agents that depend only on the aggregate state of the economy. In this framework, our definition of a self-justified equilibrium applies directly.

We argue that it is necessary to allow for the possibility that the set of admissible forecasting functions is non-convex. In order to obtain existence, we assume that the set decomposes into a finite union of convex sets and that agents make discrete choices. When one introduces discrete choices into a standard dynamic stochastic model with heterogeneous agents, individual best responses are not longer convex-valued. Following Starr (1969), who establishes the existence of Arrow Debreu equilibria in economies with non-convex preferences, the assumption of a continuum of agents is crucial to ensure convexity of the best response correspondences.

To develop an algorithm to compute self-justified equilibria, we consider a specific form for

the forecasting functions and the associated non-parametric regression. We assume that each agent projects the current endogenous variables into a relatively low dimensional subspace and approximates forecasts over this subspace by regularized least squares with a reproducing kernel Hilbert space (RKHS) regularization. Computationally this amounts to combining Gaussian process regression (see, e.g., Rasmussen and Williams (2005)) with the exploitation of active subspaces (see, e.g., Constantine et al. (2014)). For dynamic economic problems, this method was first introduced by Scheidegger and Bilonis (2019). It directly gives rise to a simple algorithm that trades off complexity and simplicity of the forecasting function and allows us to approximate self-justified equilibria numerically. Moreover, the error analysis becomes simple since we can reverse-engineer a cost-function of computational complexity, which rationalizes the computed approximation as a self-justified equilibrium. We demonstrate that our computational method can be applied to large-scale heterogeneous agents models by solving for self-justified equilibria in an overlapping generations economy with segmented financial markets. We assume that agents live for 60 periods and that there are three types of agents per generation, resulting in 180 agents altogether. The three types distinguish themselves by preferences, endowments, and trading restrictions.

There is a large and diverse body of work exploring deviations from rational expectation (see, e.g., Sargent et al. (1993), Kurz (1994), Woodford (2013), Gabaix (2014), Adam et al. (2016), Molavi (2019), Geng (2018)). Much of this work is motivated by insights from behavioral economics about agents' behavior or by the search for simple economic mechanisms that enrich the observable implications of standard models. Much of the work also focuses on single-agent models where many of the technical difficulties discussed in this paper disappear. The motivation of this paper is rather different in that we want to develop a simple alternative to rational expectations that allows researchers to rigorously analyze stochastic dynamic models with a very large number of heterogeneous agents.

The methods developed in Krusell and Smith (1998) and Evans and Phillips (2014) can also be interpreted to arise from this motivation, and there are some important similarities to our work. In fact, the equilibrium concept in Krusell and Smith (1998) can be interpreted as that of a self-justified equilibrium (although they do not call it that and do not define it formally). To defend their computational strategy, Krusell and Smith (1998) write “the calculated object satisfies all the standard equilibrium conditions except the agents ability to make perfect forecasts... The accuracy is so high that we find it very hard to argue on the basis of the irrationality' of the agents in our model that our approximate equilibrium is a less satisfactory economic model than an exact equilibrium”. Unfortunately, there is no formal way to relate the computed equilibrium to a rational expectations equilibrium. It should be noted, however, that our forecasting errors can be viewed as errors in Euler equations. While every self-justified equilibrium is a rational expectations ϵ -equilibrium, most rational expectation ϵ -equilibria are not self-justified equilibria. In this sense,

our concept can be viewed as a refinement of approximate equilibria that naturally leads to an efficient way to compute them numerically and to straightforward error analysis.

The remainder of the paper is organized as follows. In Section 2, the general economy is introduced, and a self-justified equilibrium is defined. In Section 3, we prove existence. In Section 4, we explain some basic ideas on how to numerically model a self-justified equilibrium and describe our algorithm. In Section 5, we give a numerical example to illustrate the concept. Section 6 concludes.

2 A general dynamic Markovian economy

We consider a Bewley-style overlapping generations model (see Bewley (1992)) with incomplete financial markets and a continuum of agents. Time is discrete and indexed by $t \in \mathbb{N}_0$. Exogenous shocks z_t realize in a finite set $\mathbf{Z} = \{1, \dots, Z\}$, and follow a first-order Markov process with transition probability $\pi(z'|z)$. A history of shocks up to some date t is denoted by $z^t = (z_0, z_1, \dots, z_t)$ and called a date event.

At each date event, a continuum of ex-ante identical agents enter the economy, live for A periods, and differ ex-post by the realization of their idiosyncratic shocks. Each agent faces idiosyncratic shocks, y_1, \dots, y_A , that have support in a finite set \mathbf{Y}^A . We denote by $\eta_{y^a}(y_{a+1})$ the (conditional) probability of idiosyncratic shock y_{a+1} for an agent with shock history y^a , $\eta_0(y_1)$ to denote the probability of idiosyncratic shock y_1 at the beginning of life, and, $\eta(y^a)$ to denote the probability of a history of idiosyncratic shocks. We assume that the idiosyncratic shocks are independent of the aggregate shock, that they are identically distributed across agents with the same history of shocks and, as in the construction in Proposition 2 in Feldman and Gilles (1985), that they “cancel out” in the aggregate, that is, the joint distribution of idiosyncratic shocks within a type ensures that at each history of aggregate shocks, z^t , for any $y^a \in \mathbf{Y}^a$ the fraction of agents with history $y^a = (y_1, \dots, y_a)$ is $\eta(y^a)$. This allows us to focus on equilibria for which prices and aggregate quantities only depend on the history of aggregate shocks, z^t . We denote the set of all date events at time t by \mathbf{Z}^t and, taking z_0 as fixed, we write $z^t \in \mathbf{Z}^t$ for any $t \in \mathbb{N}_0$ (including $t = 0$). At each z^t , there are finitely many different agents actively trading, that distinguish themselves by age and history of shocks, and who are collected in a set $\mathbf{I} = \cup_{a=1}^A \mathbf{Y}^a$. A specific agent at a given node z^t is denoted by $y^a \in \mathbf{I}$.

At each date event, there is a single perishable commodity, the individual endowments are denoted by $e_{y^a}(z^t) \in \mathbb{R}_+$ and assumed to be time-invariant and functions of the current aggregate shock.¹ Aggregate (labor) endowments are given by $e(z) = \sum_{y^a \in \mathbf{I}} \eta(y^a) e_{y^a}(z)$. Each agent who is

¹As opposed to the standard formulation (Bewley (1992)), where an agent’s fundamentals are functions of his current idiosyncratic shock, y , we assume that they are functions of the history of all shocks. Clearly, these formulations are equivalent if one allows for a sufficiently rich set \mathbf{Y} .

born at some node z^t has a time-separable expected utility function

$$U_{z^t}((x_{t+a})_{a=0}^{A-1}) = \sum_{a=1}^A \sum_{z^{t+a-1} \succeq z^t} \sum_{y^a} \eta(y^a) \pi(z^{t+a-1}|z^t) u_{y^a}(x_{y^a}(z^{t+a-1})), \quad (1)$$

where $x_{y^a}(z^{t+a-1}) \in \mathbb{R}_+$ denotes the agent y^a 's (stochastic) consumption at date $t+a-1$.

There are J assets, $j \in \mathbf{J} = \{1, \dots, J\}$ traded at each date event. Assets can be infinitely lived Lucas trees in positive net supply or one-period financial assets in zero net supply. The net supply of an asset j is denoted by $\bar{\theta}_j \geq 0$. Assets are traded at prices q , and their (non-negative) payoffs depend on the aggregate shock and possibly on the current prices of the assets $f_j : \mathbb{R}_+^J \times \mathbf{Z} \rightarrow \mathbb{R}_+$. If asset j is a Lucas tree (i.e., an asset in positive net supply), then $f_j(q, z) = q_j + \text{div}_j(z)$ for some dividends $\text{div}_j : \mathbf{Z} \rightarrow \mathbb{R}_+$. Asset j could also be a collateralized loan whose payoff depends on the value of the underlying collateral, or an option, or simply a risk-free asset. The aggregate dividends of the trees are defined as

$$\text{div}(z_t) = \bar{\theta} \cdot f(q(z^t), z_t) - \bar{\theta} \cdot q(z^t). \quad (2)$$

At each z^t an agent y^a enters the period with a portfolio $\theta_{y^a}^-(z^t)$ and chooses a new portfolio $\theta_{y^a}(z^t)$ and consumes

$$x_{y^a}(z^t) = e_{y^a}(z_t) + \theta_{y^a}^-(z^t) \cdot f(q(z^t), z_t) - \theta_{y^a}(z^t) \cdot q(z^t). \quad (3)$$

The agent y^a faces trading constraints $\theta_{y^a} \in \Theta_{y^a} \subset \mathbb{R}^J$, where $\Theta_{y^A} = \{0\}$ for all $y^A \in \mathbf{Y}^A$.

2.1 Self justified equilibria

In a competitive environment, agents choose asset-holdings in the current period to maximize their expected lifetime utility, and current prices ensure that markets clear. To understand how today's asset choices affect future utilities, the agent needs to form some expectations about future prices and to compute his optimal life-cycle asset-holdings under these prices. As already mentioned, it turns out to be useful to model the forecasting of prices and the recursive solution of the agents' problem in one step, and assume that the agent makes a current decision given expectations over the next period's marginal utility of asset holdings. These expectations are based on current endogenous variables and the exogenous shock. While in a rational expectations equilibrium, these expectations are always correct, we allow them here to be imprecise and heterogeneous across agents.

In a temporary equilibrium, the expectations of each agent, $y^a \in \mathbf{I}$, are characterized by a function M_{y^a} that predicts marginal utilities of assets in the next period based on the current state, current prices, and current consumption and portfolio-holdings across agents.

For what follows, it will turn out to be useful to allow for otherwise identical agents to use different forecasting functions. We assume that there is a finite number, K , of different forecasts used by the agents over their life-cycle, and that the measure of agents that use forecast k is ν_K .

We write $M_{y^a}^k(z, \cdot) = M^k(y^a, z, \cdot)$ to denote the forecasting function of agent (y^a, k) given the shock z .

To simplify the notation, we write $\vec{\theta} = (\theta_{y^a, k})_{y^a \in \mathbf{I}, k=1, \dots, K}$, $\vec{\theta}^- = (\theta_{y^a, k}^-)_{y^a \in \mathbf{I}, k=1, \dots, K}$, and $\vec{x} = (x_{y^a, k})_{y^a \in \mathbf{I}, k=1, \dots, K}$. Next, it is useful to define the set of possible portfolio holdings with market-clearing built-in as

$$\Theta = \{\vec{\theta} : \sum_{y^a \in \mathbf{I}, k=1 \dots K} \eta(y^a) \nu_k \theta_{y^a, k} = \bar{\theta}, \quad \theta_{y^a, k} \in \Theta_{y^a} \text{ for all } y^a \in \mathbf{I}, k = 1 \dots K\}. \quad (4)$$

Similarly, let the set of all beginning-of-period portfolio holdings be

$$\Theta^- = \{\vec{\theta}^- : \theta_{y^1}^- = 0, \quad \sum_{y^{a-1} \in \mathbf{I}, k=1 \dots K} \nu_k \eta(y^{a-1}) \theta_{y^a, k}^- = \bar{\theta} \text{ and } \theta_{y^a, k}^- \in \Theta_{y^{a-1}} \text{ for all } y^a k = 1 \dots K\}. \quad (5)$$

In the following we assume that M_{y^a} only depends on current asset holdings across agents and the current exogenous shock, i.e.,

$$M_{y^a}^k : \mathbf{Z} \times \Theta \rightarrow \mathbb{R}_+^J. \quad (6)$$

In our formulation, the agent forecasts the marginal utilities from asset holdings. It might seem more natural to assume that the agent forecasts prices and then solves his life-cycle optimization problem based on forecasted prices (as, for example in Krusell and Smith (1998)). However, this turns out to be much more complicated because he has to forecast prices over his entire life-cycle, and not just one-period ahead. Moreover, we illustrate in a simple example below that forecasting prices might be more complicated than forecasting marginal utilities from asset-holdings. Finally, one could argue that the agent might forecast his value function in the next period to solve the maximization problem. This turns out to be too complicated since he has to forecast an entire function.²

We denote by $\vec{M} = (M_{y^a}^k)_{y^a \in \mathbf{I}, k=1, \dots, K}$ the forecasting functions across all agents. Throughout this paper, we assume that $M_{y^A}^k(\cdot) = 0$ for all $y^A \in \mathbf{Y}^A$, $k = 1, \dots, K$, forecasts of agents of age A are irrelevant.

Assuming concavity of utility, the first order conditions are necessary and sufficient for agents' optimality and, given prices q and beginning-of-period asset-holdings $\theta_{y^a, k}^-$, we can write an agent (y^a, k) 's maximization problem as

$$\begin{aligned} \max_{x \in \mathbb{R}_+, \theta \in \Theta_{y^a}} \quad & u_{y^a}(x) + M_{y^a}^k(z, \vec{\theta}) \cdot \theta \quad \text{s.t.} \\ & x + \theta \cdot q - e_{y^a}(z) - \theta_{y^a, k}^- \cdot f(q, z) \leq 0. \end{aligned} \quad (7)$$

The agent takes as given the current portfolio- and consumption choices across all agents, $\vec{\theta}$, \vec{x} , and current prices q . For now, the functions $M_{y^a}^k(\cdot)$ are given—we endogenize this for our definition of the self-justified equilibrium below.

²It is true that one could approximate the value function by a finitely parameterized family of functions, and the agent forecasts the finite-dimensional vector of parameters. However, this would still be substantially more complicated than merely forecasting a number.

We define the state space to be $\mathbf{S} = \mathbf{Z} \times \Theta^-$, with Borel σ -algebra \mathcal{S} . It is useful to notice that the law of motion of the exogenous shock, π , and current choices $\vec{\theta}$ determine a probability distribution over next period's state, and to write $\mathbb{Q}(\cdot|z, \vec{\theta})$ to denote this probability distribution.

Given forecasting functions across agents, $\vec{M} = (M_{y^a}^k)$, we define the temporary equilibrium correspondence

$$\mathbf{N}_{\vec{M}} : \mathbf{S} \rightrightarrows \mathbb{R}_+^{IK} \times \Theta \times \mathbb{R}^J \quad (8)$$

as a map from the current state to current prices and choices that clear markets and that are optimal for the agents, given their forecasting functions, i.e.,

$$\begin{aligned} \mathbf{N}_{\vec{M}}(s) = \{ & (\vec{x}, \vec{\theta}, q) \in \mathbb{R}_+^{IK} \times \Theta \times \mathbb{R}^J : \\ & (x_{y^a, k}, \theta_{y^a, k}) \in \arg \max_{x \in \mathbb{R}_+, \theta \in \Theta_{y^a}} u_{y^a}(x) + M_{y^a}^k(z, \vec{\theta}) \cdot \theta \text{ s.t.} \\ & x + \theta \cdot q - e_{y^a}(z) - \theta_{y^a, k}^- \cdot f(q, z) \leq 0 \text{ for all } y^a \in \mathbf{I} \}. \end{aligned} \quad (9)$$

Assuming that for a given \vec{M} , the set $\mathbf{N}_{\vec{M}}(s)$ is non-empty for all $s \in \mathbf{S}$, and that there exists a single-valued (Borel-measurable) selection $N_{\vec{M}}(s)$, we write

$$N_{\vec{M}}(s) = N(s) = (N_{\vec{x}}(s), N_{\vec{\theta}}(s), N_q(s)). \quad (10)$$

The function $N(s)$ depends on \vec{M} , but to simplify notation, we often drop the subscript.

In what follows, we assume that all agents base their forecasting functions on the selection, $N(\cdot)$. In principle, one could imagine equilibria where different agents use different selections of the correspondence. In such a framework, ‘‘sunspots’’ would play an important role. In this paper, however, we focus on the ‘‘spot-less’’ case, where the possible multiplicity of temporary equilibria plays no role.

The crucial innovation of this paper is to allow for heterogeneous and possibly imprecise forecasts across agents while still allowing for the possibility that they are rational. For this, we assume that the agents cannot evaluate (or store) arbitrarily complicated functions, but instead, approximate the equilibrium forecasts by ‘‘simple’’ functions, for which we will give examples in the next section.

To formalize these ideas, we assume that the agents forecasting functions are in a given set of (continuous) functions \mathbf{M} , and that each agent y^a takes as given a distribution over $\vec{\theta}$ given z , $\mathbb{P}(\vec{\theta}|z)$, and chooses the forecast $M_{y^a}^k$ that minimizes a loss function that depends on the complexity of the forecast, i.e.

$$M_{y^a}^k(z, \cdot) \in \arg \min_{M \in \mathbf{M}} \mathcal{L}_{y^a} \left(\int_{\vec{\theta} \in \Theta} \left\| M(\vec{\theta}) - m_{y^a}(z, \vec{\theta}) \right\|_p^p d\mathbb{P}(\vec{\theta}|z), M \right), \quad (11)$$

where $\|x\|_p = (\sum_i |x_i|^p)^{\frac{1}{p}}$ denotes the p -norm, for some $p \geq 1$, and where the marginal utility from assets purchased at s for agent y^a, k -as read off from the temporary equilibrium correspondence—is given by

$$m_{y^a, k}(z, \vec{\theta}) = \int_{s'} f(N_q(s'), z') \sum_{y_{a+1}} \eta_{y^a}(y_{a+1}) u'_{y_{a+1}}(N_{x_{y_{a+1}}}(s')) d\mathbb{Q}(s'|z, \vec{\theta}). \quad (12)$$

We assume that the functions in \mathbf{M} have an universal approximating property in that whenever $m_{y^a}(z, \vec{\theta})$ is continuous in $\vec{\theta}$, for each $\epsilon > 0$, there is a $M \in \mathbf{M}$ such that

$$\sup_{\vec{\theta} \in \Theta} \|M(\vec{\theta}) - m_{y^a}(z, \vec{\theta})\| < \epsilon. \quad (13)$$

In principle, an agent can make arbitrarily accurate forecasts if the equilibrium map is continuous. However, in a self-justified equilibrium, computational costs prevent this.

We have the following definition for our concept of self-justified equilibrium.

DEFINITION 1 *A self-justified equilibrium consists of an integer K of different forecasts used across otherwise identical agents, measures $\nu_k > 0$ of agents that use these forecasts, as well as forecasts \vec{M} , a selection $N(\cdot)$ of the temporary equilibrium correspondence, $\mathbf{N}_{\vec{M}}(\cdot)$, and measure \mathbb{Q}^* on $(\mathbf{S}, \mathcal{S})$, such that*

1. \mathbb{Q}^* is invariant given the law of motion induced by $N(\cdot)$ and by $\mathbb{Q}(\cdot, \cdot)$. That is to say, for all $\mathbf{B} \in \mathcal{S}$

$$\mathbb{Q}^*(\mathbf{B}) = \int_{s \in \mathbf{S}} \mathbb{Q}(\mathbf{B}|z, N_{\vec{\theta}}(s)) d\mathbb{Q}^*(s).$$

2. For each y^a , $a < A$, each $k = 1, \dots, K$,

$$M_{y^a}^k(z, \cdot) \in \arg \min_{M \in \mathbf{M}} \mathcal{L}_{y^a} \left(\int_{s \in \mathbf{S}} \|M(N_{\vec{\theta}}(s)) - m_{y^a, k}(z, N_{\vec{\theta}}(s))\|_p^p d\mathbb{Q}^*(s|z), M \right).$$

Part 1 of Definition 1 is defining an invariant measure that is needed to compute the long-run forecasting error. Part 2 ensures that each agent's forecasting function is obtained by minimizing a loss function that trades off the accuracy of the forecast and its computational complexity. In what follows, we assume for simplicity that $p = 2$, in principle it could be heterogeneous across agents and take any positive integer value.

Similarly to the concept of “self-confirming” equilibrium (see e.g. Fudenberg and Levine (1993) or Cho and Sargent (2008)), a self-justified equilibrium can be interpreted as a stationary point of a learning process which itself is not modeled in the theory. The crucial difference is that in a self-justified equilibrium, an agent's forecasts can be incorrect in every step.

Both rational expectations equilibria and self-justified equilibria are special cases of a temporary equilibrium in this model. For the special case where

$$m_{y^a}(z, N_{\vec{\theta}}(s)) = M_{y^a}(z, N_{\vec{\theta}}(s)) \text{ for all } s \in \mathbf{S}, \quad (14)$$

we obtain a standard rational expectations equilibrium if we assume concave utility. In this case, the first-order conditions that describe agents' optimal choices are also necessary and sufficient conditions for the optimization problem stated in (7), and agents forecast future prices perfectly.

Under the assumptions stated in Section 4, when the set of admissible functions, \mathbf{M} is sufficiently rich, and losses from computationally expensive functions are low, a self-justified equilibrium

converges to a rational expectations equilibrium. The main contribution of this paper is to explore what happens if the agent is unable to approximate $m_{y^a,k}$ perfectly.

3 Existence

To prove the existence of simple equilibria in heterogeneous agents models with incomplete markets, one needs to impose strong assumptions on fundamentals. Brumm et al. (2017) present one possible set of strong assumptions and argue that without strong assumptions, simple equilibria might fail to exist (Kubler and Polemarchakis (2004) provide simple counterexamples). We show now that under the (strong) assumption of finite accounting, proving existence is relatively straightforward.

3.1 Assumptions

We first make a number of fairly standard assumptions on fundamentals that are used to prove the existence of a temporary equilibrium for given forecasting functions.

ASSUMPTION 1

1. For each $y^a \in \mathbf{I}$, the Bernoulli-utility function $u_{y^a}(\cdot)$ is continuously differentiable, strictly increasing, strictly concave, and satisfies an Inada condition

$$u'_{y^a}(x) \rightarrow \infty \text{ as } x \rightarrow 0.$$

Individual endowments are positive, i.e.,

$$e_{y^a}(z) > 0 \text{ for all } z \in \mathbf{Z}.$$

2. For any $\epsilon > 0$ for each $y^a \in \mathbf{I}$, $z \in \mathbf{Z}$, and any (finite) $\theta_{y^a,k}^- \in \mathbb{R}^J$ the set

$$\{(x, \theta) \in \mathbb{R}_+ \times \Theta_{y^a}(x - e_{y^a}(z)) + \theta \cdot \frac{1}{p_s} q_s - \theta_{y^a,k}^- \cdot f(\frac{1}{p_s} q_s, z) \leq 0\}$$

is a compact convex set containing the point $(e_{y^a}(z), 0)$, whenever $p, q_j \geq \epsilon$ for all j and $p + \sum_{j=1}^J q_j = 1$.

3. The payoff functions, $f : \mathbb{R}_+^J \times \mathbf{Z} \rightarrow \mathbb{R}^J$, are non-negative valued and continuous. Moreover, for any $i = 1, \dots, J$ and $j = 1, \dots, J$, the payoff $f_j(q, z)$ only depends on q_i if $\bar{\theta}_i > 0$.

Assumption 1.1 is a standard assumption in general equilibrium analysis (see, e.g., Kubler and Polemarchakis (2004)). Assumption 1.2 is the crucial assumption that restricts agents' trades. Part of it is motivated by collateral and default. These constraints ensure that agents cannot borrow against future endowments. In our formulation, this is true independently of prices, and could be justified if we allow for default (see Kubler and Schmedders (2003) for a detailed motivation), or if

agents face appropriate borrowing constraints. Although it is a strong, reduced-form assumption, one can verify that it is satisfied in many variations of the model. Assumption 1.3 subsumes models with Lucas trees and one-period assets.

The crucial and non-standard assumption of this paper is that accounting is finite, i.e., that beginning of period portfolios lie in a finite set (or at least that agents perceive them to lie in a finite set). This simplifies the analysis dramatically, and we will argue below that it has few practical disadvantages. Formally, we make the following assumptions:

ASSUMPTION 2

1. *There is a finite set $\widehat{\mathbf{S}} = \mathbf{Z} \times \widehat{\Theta}^- \subset \mathbf{S}$ such that the support of the transition function $\mathbb{Q}(\cdot|z, \vec{\theta})$ is a subset of $\widehat{\mathbf{S}}$ for all $z \in \mathbf{Z}$ and all $\vec{\theta} \in \Theta$.*
2. *The measure $\mathbb{Q}(\cdot|z, \vec{\theta})$ is upper hemi-continuous and convex valued in $\vec{\theta}$ for all $z \in \mathbf{Z}$, $\vec{\theta} \in \Theta$.*

Assuming that $\widehat{\mathbf{S}}$ contains ZG elements, we then can take $\mathbb{Q}(\cdot|z, \vec{\theta})$ to be a vector in the $ZG - 1$ dimensional unit simplex, Δ^{ZG-1} . Assumption 2.2 then simply states that this vector changes continuously in $\vec{\theta}$.

From a practical point of view, this assumption seems innocuous. Because of finite precision arithmetic in scientific computations, almost any numerical method will lead to $\vec{\theta}^-$ lying on a (possibly very fine) grid. Assumption 2.2 then states that there is some randomness in the rounding error. However, from a technical point, the assumption turns out to be crucial. It is subject to further research to see which of our results hold in the limit as the grid becomes dense in Θ^- . The assumption will allow us to obtain simple existence results below. However, it is certainly not a standard assumption in this strand of literature. Note, however, that in the technically very related literature on Markov-perfect equilibria in stochastic games this is the classical assumption, already made by Shapley (1953), that can only be relaxed in special cases.

Assuming finite accounting has several economic justifications. One interpretation is that actual portfolios lie in Θ^- , but that agent cannot measure portfolios arbitrarily finely and make their decisions based on rounded values, exhibiting some degree of bounded rationality. Our preferred interpretation is that agents take the fact that at the beginning of each period, there are random (small) transfers of assets between individuals. The transfers only depend on aggregate variables, i.e., average asset-holdings of all agents within one type y^a . An individual maximizes utility, taking as given these random transfers. The transfers are designed to ensure that the resulting after-transfer asset holdings like in $\widehat{\Theta}^-$.

Assumption 2 guarantees that $\widehat{\mathbf{S}}$ is finite and we can take it to contain ZG elements. For fixed $\vec{M} \in \mathbf{M}$, a selection of the temporary equilibrium correspondence can then be viewed as a vector $N \in (\mathbb{R}_+^I \times \Theta \times \mathbb{R}_+^J)^{ZG}$.

We make the following reduced-form assumption on forecasting-functions:

ASSUMPTION 3

1. There is a finite W and a set $\widehat{\mathbf{M}} \subset \mathbf{M}$ with

$$\widehat{\mathbf{M}} = \cup_{i=1}^W \mathbf{M}_i$$

that has the following properties

(a) For all $N \in (\mathbb{R}_+^I \times \Theta \times \mathbb{R}_+^J)^{ZG}$, all distributions over s , $\mathbb{P}(s)$ and all y^a and all z ,

$$\arg \min_{M \in \widehat{\mathbf{M}}} \mathcal{L}_{y^a} \left(\int_{s \in \mathbf{S}} M(N_{\vec{\theta}}(s)) - m_{y^a, k}(z, N_{\vec{\theta}}(s)) d\mathbb{P}(s|z), M \right) \subset \widehat{\mathbf{M}}$$

(b) For all y^a , for all $N \in (\mathbb{R}_+^I \times \Theta \times \mathbb{R}_+^J)^{ZG}$, all $s \in \widehat{\mathbf{S}}$ and all $\vec{\theta} \in \Theta$, the function $M_{y^a}^k(z, \vec{\theta}; N)$ is jointly continuous in $\vec{\theta}$, and N whenever

$$M_{y^a}^k(z, \cdot; N) \in \arg \min_{M \in \mathbf{M}^k} \mathcal{L}_{y^a} \left(\int_{s \in \mathbf{S}} M(N_{\vec{\theta}}(s)) - m_{y^a, k}(z, N_{\vec{\theta}}(s)) d\mathbb{Q}^*(s|z), M \right)$$

2. All functions in \mathbf{M} are uniformly bounded above.

Assumption 3.1 is the second key ingredient of our existence result. While the examples above clearly demonstrate that it is too strong to assume that \mathbf{M} is a convex and well-behaved set, we assume that it can be written as the finite union of well-behaved sets, taking into account that it is never optimal to choose certain $M \in \mathbf{M}$ because the loss due to their complexity cannot possibly outweigh the accuracy gain. The assumption allows us to assume that finitely many and otherwise identical, agents make different choices and, with the right assignment of measures, can take the best responses to be convex-valued. Assumption 3.2 is a weak assumption on the set of admissible forecasts. In an overlapping generations economy this is without loss of generality.

3.2 The main theoretical result

With these assumptions, the existence of a self-justified equilibrium reduces to the existence of a finite-dimensional fixed point. The main result of this section thus reads as follows:

THEOREM 1 *Under Assumptions 1-3, there exists a self-justified equilibrium.*

Proof. Assumption 3.1 implies that we can define a finite $\tilde{W} \in \mathbb{N}$, $\tilde{W} \leq W^{IZ}$, to be the maximal number of possible forecasts an agent can use over his life-cycle. The total maximal number of ex-post different agents in the economy is then given by $I\tilde{W}$.

We take prices at each $s \in \widehat{\mathbf{S}}$ to lie in the trimmed simplex $(p, q) \in \Delta_\epsilon^J = \{(p, q) \in \mathbb{R}_+^{J+1}, p + \sum_{j=1}^J q_j = 1, p \geq \epsilon, q_j \geq \epsilon, j = 1, \dots, J\}$. For any $\epsilon > 0$, Assumption 1.2 together with non-negative consumption implies that we can take a compact set \mathbf{T}^- to denote the possible beginning of period portfolio holdings of all $I\tilde{W}$ agents. Note that this is generally a much larger set than Θ^- , which

depends on measures ν_k and has market clearing built-in. Let $\widehat{\mathbf{T}}^-$ denote a discretization and—in a slight abuse of notation—let $\widehat{\mathbf{S}} = \mathbf{Z} \times \widehat{\mathbf{T}}^-$ denote the set of states, extended to all states that can be reached without market-clearing.

We decompose the economy into sub-economies for each $s \in \widehat{\mathbf{S}}$ and construct a map from a compact and convex set of all agents' choices, prices, probabilities, μ , and forecasts, M_s , into itself. We show that this map is upper-hemi-continuous (uhc) and convex valued, and using Kakutani's theorem (see Border (1985)), we can show that this map has a fixed point. As ϵ becomes sufficiently small, one can prove market-clearing. We can then construct an equilibrium by choosing the correct weights of different agents, ν_k , that correspond to the agents that are active in markets.

First, we need to find a suitable, convex, and compact domain for the map. Assumption 1.2 implies that the set

$$\mathbf{B} = \cup_{(p,q) \in \Delta^\epsilon} \left\{ (x, \theta) \in \mathbb{R}_+ \times \mathbb{R}^J : (x - e_{y^a}(z)) + \theta \cdot \frac{1}{p_s} q_s - \theta_{y^a, k}^- \cdot f\left(\frac{1}{p_s} q_s, z\right) \leq 0 \right\} \quad (15)$$

is compact and convex.

We construct a uhc, non-empty and convex-valued correspondence, Φ , mapping choices and prices at each element in $\widehat{\mathbf{S}}$ as well as a probability measure over $\widehat{\mathbf{S}}$, to itself, which has a fixed point,

$$\Phi : \mathbf{B}^{ZGI\tilde{W}} \times \Delta^{ZGJ} \times \Delta^{ZG} \rightrightarrows \mathbf{B}^{ZGI\tilde{W}} \times \Delta^{ZGJ} \times \Delta^{ZG-1}. \quad (16)$$

For this construction, for all $y^a \in \mathbf{I}$, all $k = 1, \dots, \tilde{W}$ and all $s \in \widehat{\mathbf{S}}$, let

$$\begin{aligned} \Phi_{y^a, k, s}((x_t, p_t, q_t)_{t \in \widehat{\mathbf{S}}}) &= \arg \max_{(x, \theta) \in \mathbf{B} \cap \Theta_{y^a}} u_{y^a}(x) + \widetilde{M}_{y^a}^k(z, \vec{\theta}_s) \cdot \theta \\ &\text{s.t.} \\ &(x - e_{y^a}(z)) + \theta \cdot \frac{1}{p_s} q_s - \theta_{y^a, k}^- \cdot f\left(\frac{1}{p_s} q_s, z\right) \leq 0, \end{aligned} \quad (17)$$

where

$$\widetilde{M}_{y^a}^k(z, \cdot) = \arg \min_{M \in \mathbf{M}^k(y^a, z)} \mathcal{L}_{y^a} \left(\int_{s \in \mathbf{S}} M(N_{\vec{\theta}}(s)) - m_{y^a, k}(z, N_{\vec{\theta}}(s)) d\mathbb{Q}^*(s|z), M \right), \quad (18)$$

and where $k(y^a, z)$ denotes agent k 's choices of forecasting sets at y^a, z , for each $k = 1, \dots, \tilde{W}$.

We say an agent k is *active* if his forecasts are always optimal, i.e. for all y^a, z

$$M_{y^a}^k(z) \in \arg \min_{M \in \mathbf{M}} \mathcal{L}_{y^a} \left(\int_{s \in \mathbf{S}} M(N_{\vec{\theta}}(s)) - m_{y^a, k}(z, N_{\vec{\theta}}(s)) d\mathbb{Q}^*(s|z), M \right). \quad (19)$$

Define a best response correspondence for all active agents of history y^a at s as

$$\Phi_{y^a, s}((x_t, p_t, q_t)_{t \in \widehat{\mathbf{S}}}) = \cup_{k=1, \dots, \tilde{W}, k \text{ active}} \Phi_{y^a, k, s}. \quad (20)$$

Our definition of a self-justified equilibrium obviously requires that agents that are not active have measure zero and their choices does not influence market clearing.

Let

$$(\tilde{\Phi}_{y^a, s}(\cdot))_{y^a \in \mathbf{I}, s \in \tilde{\mathbf{S}}} = \text{convex hull} \left((\Phi)_{y^a, s}(\cdot)_{y^a \in \mathbf{I}, s \in \tilde{\mathbf{S}}} \right). \quad (21)$$

This correspondence is convex valued and uhc.

For each $s \in \mathbf{S}$, define the price-players best response as

$$\Phi_{0, s}(\vec{\theta}_s, \vec{x}_s) = \arg \max_{(p, q) \in \Delta_c^J} p \left(\sum_{y^a \in \mathbf{I}} \eta(y^a)(x_{y^a, s} - e_{y^a}(z) - \text{div}(z)) \right) + q \cdot \left(\sum_{y^a \in \mathbf{I}} \eta(y^a)(\theta_{y^a, s} - \bar{\theta}) \right), \quad (22)$$

and let the invariant measure, $\mu(s)$ be determined by

$$\Phi_\mu((\vec{\theta}_s)_{s \in \mathbf{S}}, \mu) = \left(\sum_{s' \in \mathbf{S}} \mu(s') \mathbb{Q}(s|z', \vec{\theta}_{s'}) \right)_{s \in \mathbf{S}}. \quad (23)$$

Assumptions 1 - 3 guarantee that the mapping

$$\Phi = \times_{s \in \mathbf{S}, y^a \in \mathbf{I}} \Phi_{y^a, s} \times \times_{s \in \mathbf{S}} \Phi_{0, s} \times \Phi_\mu \quad (24)$$

is non-empty, convex valued, and uhc. By Kakutani's fixed point theorem, there exists a fixed point with prices $(\bar{p}_s, \bar{q}_s)_{s \in \hat{\mathbf{S}}}$.

As $\epsilon \rightarrow 0$, Assumption 1 guarantees that there will be a strictly positive ϵ such that

$$(\bar{p}_s, \bar{q}_s)_{s \in \hat{\mathbf{S}}} \in \arg \max_{(p, q) \in \Delta_0^J} p \left(\sum_{y^a \in \mathbf{I}} \eta(y^a)(x_{y^a, s} - e_{y^a}(z) - \text{div}(z)) \right) + q \cdot \left(\sum_{y^a \in \mathbf{I}} \eta(y^a)(\theta_{y^a, s} - \bar{\theta}) \right). \quad (25)$$

By a standard argument, markets clear. Since for each y^a , $(x_{y^a, s})_{s \in \hat{\mathbf{S}}}$ lies in the convex hull of best responses of all agents with forecasts M^k , $k = 1, \dots, \tilde{W}$ that are active in the sense that their forecasts yield the same minimal value of the loss function. We can write each

$$((x_{y^a, s}, \theta_{y^a, s})_{s \in \mathbf{S}}) = \sum_{i=1}^{\tilde{K}} \nu_i (x_{y^a, i, s})_{s \in \mathbf{S}}, \quad (26)$$

for some $\tilde{K} \leq \tilde{W} \in \mathbb{N}$ and some $\nu_i \geq 0$, $\sum_i \nu_i = 1$. From this, we can construct the appropriate weights for all active agents. This finishes our proof of existence. \square

The discretization of the state-space enables us to prove a relatively strong result. Without this, strong assumptions would be needed to ensure the existence of a recursive rational expectations equilibrium (Brumm et al. (2017)), and the existence of a self-justified equilibrium thus would remain an open problem. With the assumption, the proof of existence reduces to essentially a finite-dimensional problem. The only technical difficulty lies in the non-convexity in agents' choices in the forecast. The key to overcoming the obstacle of non-existence is to allow otherwise identical agents to make different choices in forecasting functions.

Note that the assumption of a discrete set of states forces us to reduce the non-convexity of \mathbf{M} to a finite set of discrete choices.

Note also that the assumption and our result directly show the existence of recursive rational expectations equilibria in this framework.

4 Approximate SJE

Just like the concept of rational expectations equilibrium, our concept is an idealization that cannot be computed exactly. In particular, it is clear that the invariant distribution in Definition 1.1 can only be approximated by a finite number of simulated steps and that the integral in definition 1.2 can only be approximated by Monte-Carlo integration. Clearly, in a computed SJE, the condition 1.2 will hold with some error that can be estimated by simulating a new, larger sample, and then comparing solutions. However, this error can be made arbitrarily small by simulating for sufficiently many periods. It is therefore comparable to errors in agents' optimization and market clearing, which can be pushed close to machine precision.

Our definition of self-justified equilibria is very general, and it puts very little structure on agents' forecasts. In particular, so far we made no assumptions on how the loss-functions $\mathcal{L}_{y^a}(\cdot)$ depends on the complexity of the function M . In general, we want to examine economies where individuals have access to very good forecasting technologies. We then reverse-engineer the loss function to rationalize the good forecasts as optimal forecasts. Alternatively (in particular in cases where good forecasts might not be easily available) we could measure the complexity of a function by the number of floating point operations needed to evaluate it at a point.

Next, we employ concepts from numerical analysis to specify agents' forecasts. Abstractly, the problem of finding suitable forecasting functions reduces to how to approximate a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ by a suitable weighted sum of basis functions, given observations, $(x_i, y_i = f(x_i))_{i=1}^L$.

4.1 Polynomial approximation

The classical Weierstrass Approximation Theorem states that any continuous function on a closed and bounded interval can be uniformly approximated on that interval by polynomials to any degree of accuracy. This theorem is a special case of the Stone–Weierstrass Theorem (see, e.g., Judd (1998)), which implies an analogous result for multidimensional real-valued functions. The set of polynomials, therefore, has the universal approximation property that we require above.

In light of the classical results, polynomial interpolation is a natural starting point for an overview of global approximation methods. Polynomials offer the advantage that they can be evaluated, differentiated, and integrated easily and in finitely many steps using only the basic arithmetic operations of addition, subtraction, and multiplication. This allows us to simply measure the computational complexity of an approximation by the number of additions and multiplications required to obtain the value of the forecasting function at a point.

The polynomial approximation would solve

$$\min_{c \in \mathbb{R}^N} \sum_{l=1}^L \left(\sum_{j=1}^N (c_j \phi_j(x_i) - y_i) \right)^2, \quad (27)$$

where $N = \binom{n+d}{d}$ and $(\phi_j)_{j=1}^N$ denotes a basis of polynomials of total degree d on \mathbb{R}^n .

4.2 Radial basis functions

Unfortunately, polynomial regression is often ill-conditioned, and this is not the computationally most convenient method to approximate our forecasting functions. To describe a convenient family of forecasting-functions mapping some set $\mathbf{X} \subset \mathbb{R}^k$ to the real numbers, we call a function $k : \mathbf{X} \times \mathbf{X} \rightarrow \mathbb{R}$ a (positive definite) *kernel* if for any finite sequence $(x_j)_{j=1,\dots,n}$ the $n \times n$ matrix $K_x = (k(x_i, x_j))_{i,j}$ is positive semi-definite. We assume that the kernel is “universal” in that it has the following universal approximating property. Given any compact $\bar{\mathbf{X}} \subset \mathbf{X}$, any continuous function $f : \bar{\mathbf{X}} \rightarrow \mathbb{R}$ and any $\epsilon > 0$ there are finitely many $(x_i, c_i) \in \bar{\mathbf{X}} \times \mathbb{R}$ such that

$$\sup_{x \in \bar{\mathbf{X}}} \left| \sum_i c_i k(x, x_i) - f(x) \right| < \epsilon. \quad (28)$$

To fix ideas, it is useful to give a concrete example, namely the so-called *square exponential* (SE) kernel, which we use in our computations below.

$$k_{\text{SE}}(x, x') = \exp \left\{ -\frac{1}{2} \sum_{i=1}^k \frac{(x_i - x'_i)^2}{\ell_i^2} \right\}. \quad (29)$$

In this formulation the $\ell_1, \dots, \ell_k \in \mathbb{R}_+$ are so-called hyper-parameters and can be chosen depending on the specific features of the data. As Micchelli et al. (2006) show, this is a universal kernel.

Given any kernel, k , we consider the (unique) associated reproducing kernel Hilbert space \mathbf{H}_k (see, e.g., Williams and Rasmussen (2006), Chapter 6) endowed with an inner product $\langle \cdot, \cdot \rangle_{\mathbf{H}}$, which for $f = \sum_{i=1}^s \alpha_i k(\cdot, x_i)$ and $g = \sum_{j=1}^r \beta_j k(\cdot, t_j)$ satisfies

$$\langle f, g \rangle_{\mathbf{H}} = \sum_i \sum_j \alpha_i \beta_j k(x_i, t_j). \quad (30)$$

Given a data set $\{(x^{(i)}, y^{(i)}) \mid i = 1, \dots, L\}$ consisting of L vectors $x^{(i)} \in \mathbb{R}^n$ and corresponding observations, $y^{(i)} = f(x^{(i)})$, agents want to construct a function \hat{f} that trades off smoothness and approximation in an optimal way.

Given a reproducing kernel Hilbert space, \mathbf{H}_k , with a positive definite kernel $k(x, y)$, classical regularization theory (see, e.g., Williams and Rasmussen (2006), and the references therein) solves the following problem:

$$\min_{f \in \mathbf{H}_k} \frac{1}{n} \sum_{i=1}^n (y^{(i)} - f(x^{(i)}))^2 + \lambda \|f\|_k^2, \quad (31)$$

where $\|f\|_k = \langle x, x \rangle_{\mathbf{H}_k}$ is the norm defined by $k(\cdot)$. It can be shown that the solution to (31) can be written as

$$\hat{f}(x) = \sum_{i=1}^n \alpha_i k(x, x_i), \quad (32)$$

where α solves

$$(K + \lambda I)\alpha = y, \quad (K)_{ij} = k(x_i, x_j), \quad y = (y^{(1)}, \dots, y^{(n)})^T. \quad (33)$$

Using (32) to obtain agents' forecasting functions, M , turns out to be computationally convenient because this expression is identical to the posterior mean of a Gaussian process (GPs; see Rasmussen and Williams (2005)). Formulating the problem in a Gaussian framework has the crucial advantage that it naturally leads to systematic ways for choosing the hyper-parameters of the kernel, $k(\cdot)$, as well as the regularization parameter λ in (31) via maximum likelihood estimation. Moreover, the standard deviation of the GP can be used as an indication of goodness of fit and can give some indication on whether a higher value of n can lead to much higher accuracy.

4.3 Ridge approximation and active subspaces

Unfortunately, neither polynomial nor radial basis function approximation is feasible if the domain of the forecasting function is very high-dimensional. Therefore we have to use a method to reduce this dimension. We assume that agents project the very high-dimensional argument of the forecasting function into a lower-dimensional space.

To approximate a very high dimensional function $f : \mathbb{R}^D \rightarrow \mathbb{R}$, assume that it can be reasonably well approximated with the following form:

$$f(x) \approx h(W^T x), \quad (34)$$

where the matrix $W \in \mathbb{R}^{D \times d}$ projects the high-dimensional input space, \mathbb{R}^D , to the low-dimensional *active subspace*, \mathbb{R}^d , $d \ll D$, and $h : \mathbb{R}^d \rightarrow \mathbb{R}$ is a d -dimensional function which we will denote as the *link* function. Note that the representation of (34) is not unique. All matrices W whose columns span the same subspace of \mathbb{R}^D yield identical approximations. Thus, without loss of generality, we restrict our attention to matrices in the Stiefel manifold, $W \in \mathbf{V}_d(\mathbb{R}^D)$.

We want to find an optimal W , and an optimal link function, h , in the sense minimizing the L^2 norm of the difference of $f(x)$ and $h(W^T x)$ with respect to a probability density over \mathbb{R}^D , $\rho(x)$. We can write a matrix $V \in W \in \mathbf{V}_d(\mathbb{R}^D)$ as (V_1, V_2) , i.e. or given $V_1 \in \mathbf{V}_d(\mathbb{R}^D)$, we can define $V_2 = I_{D \times D} - V_1 V_1^T$ and write $x = V_1 y + V_2 z$ for $y = V_1^T x$, $z = V_2^T x$. We can define $\tilde{\rho}(y, z) = \rho(V_1 y + V_2 z)$ and marginal and conditional densities by the standard equations. The conditional expectation is then defined as

$$\mathbb{E}(f(x)|y) = \int f(V_1 y + V_2 z) \tilde{\rho}(z|y) dz. \quad (35)$$

The optimal V_1 solves the following optimization problem:

$$\min_{V_1 \in \mathbf{V}_d(\mathbb{R}^D)} \int_x (f(x) - \mathbb{E}(f(x)|V_1^T x))^2 \rho(x) dx. \quad (36)$$

It is conceptually easy to couple this idea with polynomial least square approximation. Hokanson and Constantine (2018) note that the polynomial least-squares problem can be solved explicitly,

reducing the search for an optimal projection, V_1 , to an optimization problem over the Stiefel manifold.

Unfortunately, it turns out that the method there cannot be employed for large-scale problems. In particular, the optimization-problem (36) is a complicated, non-convex optimization problem, and even the search for a stationary point turns out to be very costly in high dimensions. Constantine et al. (2017) propose to use active subspace methods to obtain an approximation for a stationary point, and we use their method to obtain good projections in our computational strategy.

4.3.1 Active subspaces

Constantine et al. (2014) give a simple method to choose the projection matrix, W , which we briefly review. Let $\rho(x)$ be the probability density function of the relevant invariant distribution. Define a matrix

$$C := \int (\nabla f(x))(\nabla f(x))^T \rho(x) dx, \quad (37)$$

where

$$\nabla f(\cdot) = \left(\frac{\partial f(\cdot)}{\partial x_1}, \dots, \frac{\partial f(\cdot)}{\partial x_D} \right). \quad (38)$$

Since C is symmetric positive definite, it admits the form

$$C = V\Lambda V^T, \quad (39)$$

where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_D)$ is a diagonal matrix containing the eigenvalues of C in decreasing order, $\lambda_1 \geq \dots \geq \lambda_D \geq 0$, and $V \in \mathbb{R}^{D \times D}$ is an orthonormal matrix whose columns correspond to the eigenvectors of C . The classical active subspace approach in Constantine et al. (2014) suggests separating the d largest eigenvalues from the rest,

$$\Lambda = \begin{bmatrix} \Lambda_1 & 0 \\ 0 & \Lambda_2 \end{bmatrix}, \quad V = \begin{bmatrix} V_1 & V_2 \end{bmatrix}, \quad (40)$$

(here $\Lambda_1 = \text{diag}(\lambda_1, \dots, \lambda_d)$, $V_1 = [v_{11} \dots v_{1d}]$, and Λ_2, V_2 are defined analogously), and setting the projection matrix to

$$W = V_1. \quad (41)$$

We can then write $y = V_1^T x$ and $z = V_2^T x$ and

$$f(x) = f(VV^T x) = f(V_1 V_1^T x + V_2 V_2^T x) = f(V_1 y + V_2 z). \quad (42)$$

Active subspace methods are attractive in practice because it turns out that for many multivariate functions in engineering models and in the natural sciences, one observes sharp drops in the spectrum of C at relatively small values of d (see Constantine (2015) and the references therein).

Constantine et al. (2014) prove the following theoretical result, which makes concrete how well active subspace methods lead to a good approximation. Let $\tilde{\rho}(y, z) = \rho(V_1 y + V_2 z)$ and define the conditional expectation of the function value, given y as

$$G(y) = \int_z f(V_1 y + V_2 z) \tilde{\rho}(z|y) dz. \quad (43)$$

Then, we have

$$\int_x (f(x) - G(V_1^T x))^2 \rho(x) dx \leq C(\lambda_{d+1} + \dots + \lambda_D), \quad (44)$$

where C is the Poincaré constant that depends on the probability density ρ .

Unfortunately, as Parente et al. (2019) point out, the Poincaré constant depends on an underlying probability distribution that weights sensitivities of the investigated function and what is crucial for the error bounds is a conditional distribution, conditioned on a so-called active variable, that naturally arises in the context. They propose a framework that allows for upper bounds on this constant.

4.4 An algorithm

To numerically approximate a self-justified equilibrium, we focus on the case where agents use optimal projections to form their forecasts. For each agent y^a and each $z \in \mathbf{Z}$, we, therefore need to find forecasting functions

$$M_{y^a}(z, \vec{\theta}) = \hat{M}(z, W_{y^a}^T \vec{\theta}), \quad (45)$$

The main computational issues are then how to solve for the projection matrices W_{y^a} and for the forecasting functions \hat{M}_{y^a} . To this end, following Scheidegger and Billionis (2019), we couple GPs to active subspaces.

For the first step, it is natural to assume that the integral (37) is approximated via Monte Carlo, that is, assuming that the observed inputs are drawn from a simulated path of the economy, and to assume that one approximates the gradients via finite differences—that is,

$$\hat{C}_{y^a, z, j} = \frac{1}{N} \sum_{i=1}^N g^{(i)} \left(g^{(i)} \right)^T, \quad (46)$$

where

$$g^{(i)} = \left(\frac{m_{y^a, j}(z, \theta_{y^a}^i, \theta_{-y^a}^i) - m_{y^a, j}(z, \theta_{y^a}^i, \theta_{y^a}^i + h, \theta_{-(y^a, \tilde{y}^a)}^i)}{h} \right)_{\tilde{y}^a \neq y^a}. \quad (47)$$

Given d and $d \times IJ$ projections $W_{y^a, j, z}$, the agent uses a regularized least squares method to find a good fit for $x^{(i)} = \left(\theta_{y^a}(z^{t(i)}), W_{y^a, z, j} \vec{\theta}_{-y^a}(z^{t(i)}) \right)$ and $y^{(i)} = m_{y^a, j}(z, \vec{\theta}(z^{t(i)}))$, $i = 1, \dots, n$, where $(z^{t(i)})$ are nodes with the current shock $z_{t(i)} = z$. Due to our projection, there is now a noise-component which determines the parameter λ in (31). In our computational examples below, we determine this by maximum-likelihood.

In our setup, the computation of self-justified equilibria is straightforward and reduces to GP regression and the repeated solution of non-linear systems of equations. In particular, we employ an iterative simulation scheme to solve for the optimal forecasting functions. In many respects, our method is very close to standard stimulation based projection-methods pioneered by – Den Haan and Marcet (1990) (see also Judd et al. (2011)). The basic steps of the algorithm are the following:

1. Fix a stopping criterium, η , the size of the training sample, an upper bound on iteration $\overline{\text{iter}}$, as well as the number of samples used for estimating C_N —that is, N .

The initial guess for each agent’s forecasting:

Fix an initial size of the training sample, n . Assume that agents only use own asset holdings to forecast, i.e., $d = J$ and each $IJ \times d$ projection matrix $W_{y^a, z}$ project on agent y^a ’s asset holdings. Next, construct the GP whose posterior means approximate

$$M_{y^a, z'}^0 : \mathbf{Z} \times \mathbb{R}^d \rightarrow \mathbb{R}_+. \quad (48)$$

Then, choose an approximation accuracy ξ and choose an initial condition $z_0, \vec{\theta}(z^{-1})$.

2. Iteration step:

Simulate a temporary equilibrium path for given forecasts \vec{M}^0 .

For $i = 1, n$

(a) Solve numerically for a temporary equilibrium, set $\vec{x}_i, \vec{\theta}_i, q_i$ to the equilibrium values and set $z_i = z$.

(b) Using pseudo random numbers, draw a new z' and set $\theta_{y^a}^- = \theta_{y^a-1}$ for all agents y^a .

3. For each y^a , regress the equilibrium values of $f(q_i, z_i)u'(x_{y^a+1, i})$ on $W_{y^a, z_{i-1}} \vec{\theta}_{i-1}$ and z_{i-1} to obtain a new GP whose posterior mean gives a new forecasting function $M_{y^a}^1$.

4. If

$$\|M^1 - M^0\| < \eta, \quad (49)$$

then set $M^* = M^1$. Elseif number of iteration steps below $\overline{\text{iter}}$ set $M^0 = M^1$ and repeat time iteration step 2. Else increase n and repeat iteration step 2.

5. Compute on a test-sample with $n^t \gg n$ an equilibrium sequence of length n^t and the realized forecasting error for all agents. If the average error is below some threshold, exit. Else
6. Compute C_{y^a} as defined in Equation 46 and its eigenvalues, λ . At sharp drops of the spectrum, form an active subspace, and check if the improvement in accuracy is large given the old sample of points. If no, exit. Else, include the relevant eigenvectors of C_N into the projection matrix, W_{y^a} , make a new initial guess for GPs and go to time iteration step 2.

Using active subspaces as a dimension-reduction technique turns out to fit well our economic model and produces excellent results in our examples below. Reiter (2010) considers an alternative approach that is better suited for models with 100,000 agents which differ only in their asset holdings, but it does not fit well into our framework where we target models with 100 - 1000 heterogeneous agents. Building on Reiter (2010), Ahn et al. (2018) solve heterogeneous agent macro models in continuous time by applying dimension reduction techniques to reduce the linear system of PDEs that characterizes their equilibrium. While their setting substantially differs from the one we are targeting here, a comparison of these approaches would be an interesting subject for further research.

5 A numerical example

We assume that all agents live for $A = 60$ periods, that aggregate shocks take two values, $z = 1, 2$, and that an idiosyncratic shock only occurs in the first period of an agent's life. We assume that this initial idiosyncratic shock can take three values $y = 1, 2, 3$ and that $\eta_0(y) = \frac{1}{3}$, $y = 1, \dots, 3$. The initial shock can be interpreted as the type of the agent. The types of agents distinguish themselves by trading constraints, endowment risk over the life-cycle, and preferences. An agent is then characterized by (y, a) , where $y = 1, 2, 3$ denotes the initial shock, and $a = 1, \dots, 60$ denotes an agent's age. Taken together, there are $3 \cdot 60 = 180$ agents trading in commodity- and asset markets in each period.

Type 1 agents ($y = 1$) can trade in a single Lucas-tree and a full set of Arrow securities (or options on the tree). In our framework, it is useful to assume that the Arrow-securities pay in the Lucas-tree (as in Gottardi and Kubler (2015)). Type 2 and 3 agents ($y = 2, 3$) can only trade in the Lucas tree. All agents face borrowing constraints, which (in this simple model) are equivalent to short-sale constraints. The model is a simplified OLG-version of Chien et al. (2011).

We assume that agents have CRRA utility functions with

$$u_{y,a}(c) = \beta^a \frac{c^{1-\gamma_y}}{1-\gamma_y}. \quad (50)$$

We choose $\beta = 0.98$, $\gamma_1 = 0.5$, and $\gamma_y = 1.5$ for $y = 2, 3$. Individual endowments are

$$e_{1,a}(z) = 0.4 + a/500, \text{ for } a < 50, \quad e_{1,a}(z) = 0.3 \text{ for } a \geq 50, z = 1, 2,$$

$$e_{2,a}(1) = \frac{e_{1,a}}{1.2}, \quad e_{2,a}(2) = 1.2e_{1,a} \text{ for } a = 1, \dots, A,$$

$$e_{3,a}(1) = 1.2e_{1,a}, \quad e_{3,a}(2) = \frac{e_{1,a}}{1.2} \text{ for } a = 1, \dots, A.$$

The dividends of the single tree are given by $div(z) = 3$ for $z = 1, 2$, and we take its supply to be $\bar{\theta} = 7$ – since the number of agents who hold the tree is fairly large this turns out to be numerically more stable than the standard value $\bar{\theta} = 1$.

We assume that the Markov transition matrix for the aggregate shock is

$$\pi = \begin{pmatrix} 0.95 & 0.05 \\ 0.05 & 0.95 \end{pmatrix}. \quad (51)$$

For concreteness, it is useful to define the temporary equilibrium system of inequalities as the system of all agents' KKT-conditions together with the market clearing conditions, i.e.,

$$\begin{aligned} -u'_{1,a}(e_{1,a}(z) + \theta_{(1,a-1),z}^- (\sum_{z' \in \mathbf{Z}} q_{z'} + \text{div}(z)) - q \cdot \theta_{1,a}) + \beta M_{1,a}(z, z', W_{1,a} \vec{\theta}) + \kappa_{1,a}, & \quad \text{for all } a, z' \quad (52) \\ & \kappa_{1,a} \cdot \theta_{1,a} \\ -u'_{y,a}(e_{y,a}(z) + \theta_{y,a-1}^- (\sum_{z' \in \mathbf{Z}} q_{z'} + \text{div}(z)) - \sum_{z' \in \mathbf{Z}} q_{z'} \theta_{y,a}) + \beta M_{y,a}(z, W_{y,a} \vec{\theta}) + \kappa_{y,a} & \quad \text{for all } a; y = 2, 3 \\ & \kappa_{y,a} \theta_{y,a}, \quad y = 2, 3, a = 1, \dots, A \\ & \sum_a (\theta_{(1,a),z} + \theta_{2,a} + \theta_{3,a}) - \bar{\theta}, \quad \text{for all } z \in \mathbf{Z}. \end{aligned}$$

We can combine $\kappa_{y,a}$ and $\theta_{y,a}$ into one variable and obtain a system with $(A-1)Z+2(A-1)+Z = 238$ equations and unknowns.

5.1 A simple self-justified equilibrium

As mentioned above, we start by assuming that agents only use their own asset holdings to forecast future marginal utilities. It is natural to assume that agent 1 (who can trade in two assets) assumes that his holdings in asset z (that pays if shock z realizes) only affects marginal utility in shock z for each $z = 1, 2$. Agents 2 and 3 base their forecasts on their Lucas-tree holdings. For all three agents, forecasts are, therefore, a function of the current shock and a single continuous variable. This is the simplest candidate self-justified equilibrium in our framework, and the question is how high do costs of a more accurate approximation have to be to rationalize this as an equilibrium.

With this specification, forecasts are, somewhat surprisingly, very good for most agents. This is somewhat reminiscent of the results by Krusell and Smith (1998), where very simple forecasts also turn out to be very accurate in the calibrated model.

Figure 1 depicts the forecasts for the marginal utility of the Lucas tree of a 59-year-old agent of type 2 plotted against the average realized marginal utility of the tree for the exogenous shock being $z = 1$. As can be seen in this figure, there is almost a perfect overlap between forecasted values and realized values. The mean-squared forecasting errors are below 5×10^{-4} for all agents of types 2 and 3 and all ages. The forecasting functions are obtained by a GP regression using approximately 200 points³. Despite the fact that the asset holdings of all other agents will affect prices and hence the marginal utility of the tree, these seem to play almost no role for accurate forecasts.

³In our simulation approach, the actual number of points varies in each iteration.




Figure 1: This figure shows the forecasts (red crosses) for the marginal utility of asset 1 of a 57-year-old agent of type 2 plotted against the average realized marginal utility (black bullets) of the tree for the exogenous shock being $z = 1$. Clearly, the own asset-choice gives an excellent forecasts.




Figure 2: This figure depicts the variation in prices (given the current shock is 1 and the previous shock also was 1) which is not explained by shocks.

Asset prices certainly do vary as the wealth-distribution varies along the equilibrium path. Figure 2 shows the variation of the tree price, given the current shock is 1, and the previous shock also was 1. This must be caused by changes in the wealth distribution over time. Why does it not affect the forecasts of type 2 and 3 agents? Note that the variation in prices is relatively small, and while this variation does affect forecasts, the effects are quantitatively tiny. The reason for this is that the marginal utility of agents of types 2 and 3 (which needs to be forecasted) is given by $\frac{q+div}{c^{1.5}}$ which turns out to vary much less than the price q . A relative increase of the price by a factor of $1 + \epsilon$ for some small $\epsilon > 0$ will lead to a much smaller increase of consumption (for younger agents because they save more, for the old agents because they have labor income) and therefore to a variation in marginal utility that is significantly smaller than $\sqrt{1 + \epsilon}$. This can be seen easily for the agents of age 60, where $c = e_{60,a}(z_t) + \theta^-(q + div)$ and magnitude of $\theta^-(q + div)$ is about the same as of $e_{y,60}$.

For agents of type 1, however, the situation is different. Figure 3 depicts the forecasts of a 59-year-old agent of type 1 plotted against the (average) realized marginal utilities of the 60-year-old agent. There are variables in addition to the own asset holdings that have significant effects on the marginal utilities⁴




Figure 3: Forecasts of a 59-year-old agent of type 1 (red crosses) plotted against the average realized marginal utilities (black bullets) of the 60-year-old agent. It becomes obvious from this figure that the own asset-choice is insufficient for a good forecast.

The average (squared) forecasting error is around 4×10^{-3} for agents of ages 58 and 59 and

⁴One should note the scale; the variation in own asset holdings is rather small, the overall variation of marginal utilities is also relatively small.

type 1 and, therefore, about an order of magnitude larger than for agents of types 2 and 3.

In particular, the variation in marginal utilities for the 59-year-olds of type 1 is relatively large compared to type 2 agents because of the utility function: the marginal utility is given by $\frac{q+div}{c^{0.5}}$ and, as for the agents of types 2 and 3, a relative increase in the price by $1 + \epsilon$ will lead to a much smaller increase in consumption. However, this means that marginal utility will vary by much more than $\sqrt{1 + \epsilon}$. The variation in prices, therefore, causes significant variations in the marginal utilities of the old agents of type 1. This is what is depicted in Figure 3. A similar effect comes into play for agents of ages 55-58, but it becomes quantitatively small for younger agents. In particular, it is important to note that for younger agents, this problem is much less severe—that is, the average (squared) forecasting errors of agents under the age of 55 are below 6×10^{-4} .

5.2 Finding the active subspaces

Suppose that the costs of moving from a one-dimensional to a higher dimensional domain of forecasting functions are relatively low. In particular, let us assume that agents whose average forecasting errors are above 10^{-3} search for a higher-dimensional active subspace.

It turns out that for this specification, there exists a two-dimensional active subspace for agents of type 1 and ages 55-60. In addition to an agent’s own asset holding, a single one-dimensional variable is needed to obtain accurate forecasts. The additional variable turns out to be a weighted sum of asset holdings across all agents, (roughly) weighted by the agents’ marginal propensity to consume. Employing a higher-dimensional ($d > 2$) space to forecast future marginal utilities turns out to improve the accuracy of the forecasts by very little.

For the agents $y = 1, a > 54$ we compute the matrix C_N (cf. (46)) by employing Monte-Carlo draws and finite differences, and we find that one single eigenvalue (in addition to the ones associated with own asset holdings) dominates all others. In Figure 4, we plot the 18 largest eigenvalues on a \log_{10} -scale (for the agent (1,59) whose realized marginal utilities are plotted in Figure 3 above). The figure confirms that all other eigenvalues are negligibly small compared to the one that corresponds to the weighted sum of asset holdings across agents – the jump from the largest to next largest eigenvalue is in the order of 10,000. This suggests that there is a two dimensional active subspace. We, therefore, re-compute a self-justified equilibrium with agents of type 1 and ages 55 to 60 using a two-dimensional active subspace. The optimal projection matrices $W_{y^a,j}$ will obviously change since the equilibrium prices and allocations change with better forecasts. We start using the active subspace resulting from the computation of C_N in the old equilibrium and recompute the matrix C_N twice as we iterate towards the new equilibrium.

In the new equilibrium, the one-dimensional subspace continues to work very well for all agents of types 2 and 3—the error for those types is almost the same as above. In addition, the average forecasting errors of type 1 agents are now uniformly below 5×10^{-4} . Figure 5 depicts the analog of

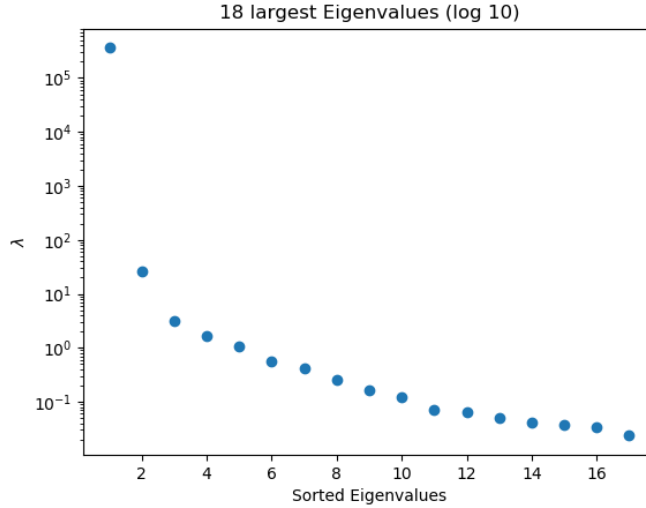


Figure 4: This figure depicts the largest eigenvalues for agent (1,59) and shock 2.

Figure 3 for the case of a two-dimensional active subspace. As can be seen, also for the 59-year-old agents of type 1, the forecasts are now almost exact. Forecasts of agents of types 2 and 3 look almost the same as the ones depicted in Figure 1.

It turns out that the variation in prices is well explained by weighting all agent’s asset holdings by their marginal propensities to consume. Asset prices are high if the young agents are relatively wealthy, and asset prices are low if the old agents are relatively wealthy. Moreover, we find that the projection matrix W obtained through the eigenvector associated with the largest eigenvalue of the matrix C_N captures this mechanism almost perfectly. It remains to be the case that all other eigenvalues of C_N are several orders of magnitude smaller than the largest eigenvalue, confirming that we have found the active subspace.

In the new equilibrium, the one-dimensional subspace continues to work very well for all agents of types 2 and 3—the error for those types is almost the same as above. In addition, the average forecasting errors of type 1 agents are now uniformly below 5×10^{-4} . Figure 5 depicts the analog of Figure 3 for the case of a two-dimensional active subspace. As can be seen, also for the 59-year-old agents of type 1, the forecasts are now almost exact. Forecasts of agents of types 2 and 3 look almost the same as the ones depicted in Figure 1.

It turns out that the variation in prices is well explained by weighting all agent’s asset holdings by their marginal propensities to consume. Asset prices are high if the young agents are relatively wealthy, and asset prices are low if the old agents are relatively wealthy.

Note that in the computed equilibrium, the forecasting errors are so small that one might be tempted to view it as an approximation to a rational expectations equilibrium. As explained in the introduction, our computational strategy actually produces excellent ϵ rational expectations equilibria.




Figure 5: Forecasts of a 59-year-old agent of type 1 (red crosses) plotted against the average realized marginal utilities (black bullets) of the 60-year-old agent as a function of the two-dimensional active subspace. The forecasts based on a two-dimensional active subspace are now extremely accurate.

6 Conclusion

This paper makes two contributions. First, we define the concept of self-justified equilibria as a natural generalization of rational expectations equilibrium, and we provide sufficient conditions for their existence. Second, we develop a numerical algorithm to approximate self-justified equilibrium numerically.

An example shows that our approximations to self-justified equilibria satisfy all of the conditions that, in the literature, typical approximations to rational expectations equilibria satisfy. Hence our computational method can also be viewed as a competitive method to approximate rational expectations ϵ -equilibria in models with many agents.

References

- Adam, K., A. Marcet, and J. P. Nicolini (2016). Stock market volatility and learning. *The Journal of Finance* 71(1), 33–82.
- Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf (2018). When inequality matters for macro and macro matters for inequality. *NBER Macroeconomics Annual* 32(1), 1–75.

- Bewley, T. (1992). Fiscal and monetary policy in a general equilibrium model. In *Equilibrium and Dynamics*, pp. 55–122. Springer.
- Border, K. C. (1985). *Fixed point theorems with applications to economics and game theory*. Cambridge university press.
- Brumm, J., D. Kryczka, and F. Kubler (2017). Recursive equilibria in dynamic economies with stochastic production. *Econometrica* 85(5), 1467–1499.
- Chien, Y., H. Cole, and H. Lustig (2011). A multiplier approach to understanding the macro implications of household finance. *The Review of Economic Studies* 78(1), 199–234.
- Cho, I.-K. and T. J. Sargent (2008). Self-confirming equilibrium. *The new Palgrave dictionary of economics*.
- Constantine, P. G. (2015). *Active subspaces: Emerging ideas for dimension reduction in parameter studies*, Volume 2. SIAM.
- Constantine, P. G., E. Dow, and Q. Wang (2014). Active subspace methods in theory and practice: applications to kriging surfaces. *SIAM Journal on Scientific Computing* 36, A1500–A1524.
- Constantine, P. G., A. Eftekhari, J. Hokanson, and R. A. Ward (2017). A near-stationary subspace for ridge approximation. *Computer Methods in Applied Mechanics and Engineering* 326, 402–421.
- Den Haan, W. J. and A. Marcet (1990). Solving the stochastic growth model by parameterizing expectations. *Journal of Business & Economic Statistics* 8(1), 31–34.
- Evans, R. W. and K. L. Phillips (2014). Olg life cycle model transition paths: Alternate model forecast method. *Computational Economics* 43(1), 105–131.
- Feldman, M. and C. Gilles (1985). An expository note on individual risk without aggregate uncertainty. *Journal of Economic Theory* 35(1), 26–32.
- Fudenberg, D. and D. K. Levine (1993). Self-confirming equilibrium. *Econometrica* 61, 523–46.
- Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics* 129(4), 1661–1710.
- Geng, R. (2018). Recursive equilibria in dynamic economies with bounded rationality.
- Gottardi, P. and F. Kubler (2015). Dynamic competitive economies with complete markets and collateral constraints. *The Review of Economic Studies* 82(3), 1119–1153.
- Hokanson, J. M. and P. G. Constantine (2018). Data-driven polynomial ridge approximation using variable projection. *SIAM Journal on Scientific Computing* 40(3), A1566–A1589.

- Judd, K. L. (1992). Projection methods for solving aggregate growth models. *Journal of Economic Theory* 58(2), 410–452.
- Judd, K. L. (1998). *Numerical methods in economics*. The MIT press.
- Judd, K. L., L. Maliar, and S. Maliar (2011). Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models. *Quantitative Economics* 2(2), 173–210.
- Krusell, P. and A. A. Smith, Jr (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy* 106(5), 867–896.
- Kubler, F. (2011). Verifying competitive equilibria in dynamic economies. *The Review of Economic Studies* 78(4), 1379–1399.
- Kubler, F. and H. Polemarchakis (2004). Stationary markov equilibria for overlapping generations. *Economic Theory* 24(3), 623–643.
- Kubler, F. and K. Schmedders (2003). Stationary equilibria in asset-pricing models with incomplete markets and collateral. *Econometrica* 71(6), 1767–1793.
- Kurz, M. (1994). On the structure and diversity of rational beliefs. *Economic theory* 4(6), 877–900.
- Ljungqvist, L. and T. Sargent (2000). *Recursive macroeconomic theory*. Mit Press.
- Micchelli, C. A., Y. Xu, and H. Zhang (2006). Universal kernels. *Journal of Machine Learning Research* 7(Dec), 2651–2667.
- Molavi, P. (2019). Macroeconomics with learning and misspecification: A general theory and applications. *Unpublished manuscript, MIT*.
- Parente, M. T., J. Wallin, and B. Wohlmuth (2019). Generalized bounds for active subspaces. *arXiv preprint arXiv:1910.01399*.
- Reiter, M. (2010). Solving the incomplete markets model with aggregate uncertainty by backward induction. *Journal of Economic Dynamics and Control* 34(1), 28–35.
- Santos, M. S. (2000). Accuracy of numerical solutions using the euler equation residuals. *Econometrica* 68(6), 1377–1402.
- Sargent, T. J. et al. (1993). Bounded rationality in macroeconomics: The arne ryde memorial lectures. *OUP Catalogue*.
- Scheidegger, S. and I. Billionis (2019). Machine learning for high-dimensional dynamic stochastic economies. *Journal of Computational Science* 33, 68–82.

- Shapley, L. S. (1953). Stochastic games. *Proceedings of the national academy of sciences* 39(10), 1095–1100.
- Starr, R. M. (1969). Quasi-equilibria in markets with non-convex preferences. *Econometrica: journal of the Econometric Society*, 25–38.
- Williams, C. K. and C. E. Rasmussen (2006). *Gaussian processes for machine learning*, Volume 2. MIT press Cambridge, MA.
- Woodford, M. (2013). Macroeconomic analysis without the rational expectations hypothesis. *Annu. Rev. Econ.* 5(1), 303–346.