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

January 6, 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 a large number of 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 rule that maps the temporary equilibrium correspondence into a set of admissible forecasts and that provides a trade-off between the accuracy of the forecast and its computational complexity.

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, Gaussian Process Regression.

JEL Classification: C63, C68 , D50, D52.

*We thank Michael Reiter, Victor Rios-Rull, Tom Sargent, Tony Smith 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", and the SFI. Simon Scheidegger gratefully acknowledges support from the Cowles Foundation at Yale University.

1 Introduction

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 (2012)). Unfortunately, for stochastic general equilibrium models with a large number of 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 (see, e.g., Brumm et al. (2017b), and references therein). In this paper, we develop an alternative to rational expectations equilibria and consider temporary equilibria with forecasting functions that depend on 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, and we show that by restricting the complexity of agents' forecasts one can numerically approximate them for 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 by performing a non-parametric regression using a finite set of points on the temporary equilibrium correspondence. Our crucial assumption is that the agent's forecasts are stable over time in the sense that using a larger data set or a more complicated class of functions for the regressions does not improve the long-run average quality of the forecasts sufficiently to justify the associated increase in computational costs.

In the temporary equilibrium, the agents might make significant mistakes, and the concept does not require identical expectations and forecasts across agents. However, in our construction of self-justified equilibria, we ensure that the only reason that prevents agents' forecasts from being arbitrarily accurate is the computational cost associated with more accurate forecasts.

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 (e.g., along the lines of Brumm et al. (2017)) is conceptually straightforward.

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 and that agents' portfolio-choices in the current period induce a probability distribution over this grid. This assumption can be viewed as a technical approximation to a continuous model, but one can also think of bounded rationality justifications. For example,

one might want to assume that at the beginning of a period, an agent cannot measure his financial wealth with arbitrary precision and makes small errors in rounding.

To develop a tractable version of the model, 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 RKHS (reproducing kernel Hilbert space) 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)). Using this combination allows us to construct a method that determines an economically intuitive projection for a fixed dimension of the subspace. For dynamic economic problems, this method was first introduced by Scheidegger and Bilionis (2017).

This method 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. We first consider the simplest case where an agent only uses his own asset-holding (together with the exogenous shock) to forecast future utilities (i.e., the asset holdings across all agents are projected into own asset holdings). This turns out to work very well for some of the agents in our model economy. However, for other agents, this simple method leads to large forecasting errors. We then exploit active subspace methods (see Constantine et al. (2014)) to show that adding one additional explanatory variable, that consists of a weighted mean of asset holdings across agents, reduces forecasting errors for these agents substantially. This observation will allow us to use the methods developed here to tackle models with hundreds to thousands of agents.

There is a large and diverse body of work exploring deviations from rational expectation (see, e.g., Sargent (1993), Kurz (1994), Woodford (2013), Gabaix (2014), Adam et al. (2016)). 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. 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.[1]

As Sargent (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 embodies 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.

There is also a large body of literature on the numerical approximation of rational expectations equilibria in models with heterogeneous agents (see Maliar and Maliar (2014) for a comprehensive overview). In the description of our numerical method, we will point out some relevant papers in this literature.

Applied dynamic general equilibrium modeling has been criticized for its failure to take into account the considerable heterogeneity in tastes and technologies across agents. Farmer and Foley (2009) make this point forcefully and strongly advocate the use of so-called agent-based models to understand macroeconomic dynamics. As they point out, in agent-based models, the agents can be as diverse as needed, but behavioral rules are often arbitrary. Up to now—especially in the presence of aggregate and idiosyncratic shocks—it seemed too complicated to incorporate substantial heterogeneity into large-scale dynamic general equilibrium models because existing solution methods are not able to handle this amount of heterogeneity (see Brumm et al. (2017b), Scheidegger et al. (2018)). Using the concept of self-justified equilibria, one can incorporate large-scale heterogeneity into general equilibrium models, potentially improve their usefulness for applied work and bridge the gap between agent-based modeling and applied general equilibrium.

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 consider a concrete example of the concept which has the attractive features that it is tractable and directly leads to a numerical method to compute forecasts. In Section 5 we describe our computational strategy. In Section 6 we give a simple example to illustrate both the concept of self-justified equilibria and our computational method.

2 A general dynamic Markovian economy

We consider a Bewley-style overlapping generations model (see Bewley (1984)) with incomplete financial markets and a continuum of agents. Time is 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

[1] The methods developed in Krusell and Smith (1997) and Evans and Phillips (2014) can also be interpreted to arise from this motivation, and there are some important similarities to our work.

probability $\pi(z'|z)$. A history of shocks up to some date t is denoted by $z^t = (z_0, z_1, \ldots, z_t)$ 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, $y_1, ..., y_A$, 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 $y^a = (y_1, ..., y_a)$ 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 (distinguishing themselves by age and history of shocks), 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.[2] Aggregate (labor) endowments are $e(z) = \sum_{y^a \in \mathbf{I}} \eta(y^a) e_{y^a}(z)$. Each agent who is born at some node z^t has a time-separable expected utility function

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

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, $j \in \mathbf{J} = \{1, \ldots, J\}$ 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 $f_j : \mathbb{R}_+^J \times \mathbf{Z} \to \mathbb{R}_+$. If asset j is a Lucas tree (i.e., an asset in positive net supply), then $f_j(q, z) = q_j + div_j(z)$ for some dividends $div_j : \mathbf{Z} \to \mathbb{R}_+$. 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 $div(z_t) = \bar{\theta} \cdot f(q(z^t), z_t) - \bar{\theta} \cdot q(z^t)$.

[2] As opposed to the standard formulation 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 $\mathbf{Y}$.

At each z^t and 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

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

The agent y^a faces trading constraints $\theta_{y^a} \in \mathbf{\Theta}_{y^a} \subset \mathbb{R}^J$, where $\mathbf{\Theta}_{y^A} = \{0\}$ for all $y^A \in \mathbf{Y}^A$. To simplify notation we write $\vec{\theta} = (\theta_{y^a})_{y^a \in \mathbf{I}}$, $\vec{\theta}^- = (\theta^-_{y^a})_{y^a \in \mathbf{I}}$ and $\vec{x} = (x_{y^a})_{y^a \in \mathbf{I}}$.

It is useful to define the set of possible portfolio holdings with market-clearing built-in as

$$\mathbf{\Theta} = \{\vec{\theta} : \sum_{y^a \in \mathbf{I}} \eta(y^a)\theta_{y^a} = \bar{\theta}, \quad \theta_{y^a} \in \mathbf{\Theta}_{y^a} \text{ for all } y^{a-1} \in \mathbf{I}\}.$$

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

$$\mathbf{\Theta}^- = \{\vec{\theta}^- : \theta^-_{y^1} = 0, \sum_{y^{a-1} \in \mathbf{I}} \eta(y^{a-1})\theta^-_{y^a} = \bar{\theta} \text{ and } \theta^-_{y^a} \in \mathbf{\Theta}_{y^{a-1}} \text{ for all } y^a\}.$$

We define the state space to be $\mathbf{S} = \mathbf{Z} \times \mathbf{\Theta}^-$ with Borel σ-algebra $\mathcal{S}$. The law of motion of the exogenous shock, π, and current choices $\vec{\theta}$ determine a probability distribution over next period's state - we write $\mathbb{Q}(\cdot|z, \vec{\theta})$. We will make assumptions on this probability distribution below which turn out to simplify the analysis and allow us to prove existence.

2.1 Self justified equilibria

In a competitive environment, agents choose asset-holdings in the current period to maximize 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 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 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

$$M_{y^a} : \mathbf{S} \times \mathbb{R}^I_+ \times \mathbf{\Theta} \times \mathbb{R}^J_+ \to \mathbb{R}^J_+,$$

that predicts marginal utilities of assets in the next period on the basis of the current state, current prices, and current consumptions and portfolio-holdings across agents. In our formulation, the agent forecasts marginal utilities from asset holdings. It might seem more standard to assume that the agent forecasts prices and then solves his life-cycle optimization problem on the basis of forecasted prices. 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.[3]

We denote by $\vec{M} = (M_{y^a})_{y^a \in \mathbf{I}}$ the forecasting functions across all agents. Throughout this paper, we assume that $M_{y^A}(\cdot) = 0$ for all $y^A \in \mathbf{Y}^A$, 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 $\theta^-_{y^a}$, we can write an agent y^a's maximization problem as

$$\begin{aligned} \max_{x \in \mathbb{R}_+, \theta \in \Theta_{y^a}} & \; u_{y^a}(x) + M_{y^a}(s, \vec{x}, \vec{\theta}, q) \cdot \theta \;\; \text{s.t.} \\ & x + \theta \cdot q - e_{y^a}(z) - \theta^-_{y^a} \cdot f(q, z) \leq 0. \end{aligned} \tag{1}$$

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

Given forecasting functions across agents, $\vec{M}$, we define the temporary equilibrium correspondence

$$\mathbf{N}_{\vec{M}} : \mathbf{S} \rightrightarrows \mathbb{R}^I_+ \times \mathbf{\Theta} \times \mathbb{R}^J$$

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.,

$$\begin{aligned} \mathbf{N}_{\vec{M}}(s) \;=\; \{ & (\vec{x}, \vec{\theta}, q) \in \mathbb{R}^I_+ \times \mathbf{\Theta} \times \mathbb{R}^J_+ : \\ & (x_{y^a}, \theta_{y^a}) \in \arg \max_{x \in \mathbb{R}_+, \theta \in \Theta_{y^a}} u_{y^a}(x) + M_{y^a}(s, \vec{x}, \vec{\theta}, q) \cdot \theta \text{ s.t.} \\ & \quad x + \theta \cdot q - e_{y^a}(z) - \theta^-_{y^a} \cdot f(q, z) \leq 0 \quad \text{for all } y^a \in \mathbf{I} \}. \end{aligned} \tag{2}$$

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

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

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

[3] It 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, but this would still be substantially more complicated than merely forecasting a number.

the correspondence. In that framework, "sunspots"would play an important role. In this paper, 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' ability to forecast is constrained by two factors. First, they can only use a finite amount of information to form the forecasts. Secondly, they cannot evaluate (or store) arbitrarily complicated functions, but instead, approximate the equilibrium forecasts by "simple" functions.

To formalize this, we denote by $\mathbf{F}^N$ the set of all Borel-measurable functions from $\mathbf{S}$ to $\mathbb{R}^I_+ \times \mathbf{\Theta} \times \mathbb{R}^J_+$, and by $\mathbf{F}^M$ the set of Borel-measurable functions from $\mathbf{S} \times \mathbb{R}^I_+ \times \mathbf{\Theta} \times \mathbb{R}^J_+$ to $\mathbb{R}^J_+$. We assume that there is a collection of possible "regressions" $(R^{n,d}_{y^a})_{n,d\in\mathbb{N}}$, where each $R^{n,d}_{y^a} : \mathbf{F}^N \to \mathbf{F}^M$ uses n points on the equilibrium function, $N \in \mathbf{F}^N$ to determine a forecasting-function, $M_{y^a} \in \mathbf{F}^M$ whose computational complexity is measured by d. It is often useful to denote the range of $R^{n,d}_{y^a}(\cdot)$ by $\mathbf{M}^{n,d}_{y^a}$ and call this the set of admissible forecasting functions.

The integers n and d are chosen to trade off the complexity of the forecast, which we measure by a cost function, $c_{y^a}(n,d)$, and the long run accuracy, which we measure by

$$\int_{s\in\mathbf{S}} \left(\frac{M_{y^a}(s,N(s)) - m_{y^a}(z,N_\theta(s))}{m_{y^a}(z,N_\theta(s))} \right)^2 d\mathbb{Q}^*(s), \tag{3}$$

where $\mathbb{Q}^*(s)$ is an invariant measure over states and, for each y^a, m_{y^a} denotes the actual realized marginal utility of assets in the subsequent period—that is,

$$m_{y^a}(z,\vec{\theta}) = \int_{s'\in\mathbf{S}} f(N_q(s'),z') \sum_{y_{a+1}\in\mathbf{Y}} \eta_{y^a}(y_{a+1}) u'_{y^{a+1}}(N_{x_{y^{a+1}}}(s')) d\mathbb{Q}(s'|z,\vec{\theta}). \tag{4}$$

We assume that for each agent y^a, the available non-parametric regressions, $(R^{n,d}_{y^a})_{n,d\in\mathbb{N}}$, have an universal approximating property in that whenever $m_{y^a}(z,N_\theta(z,\vec{\theta}^-))$ is continuous in $\vec{\theta}^-$, for each $\epsilon > 0$, there is a $\hat{n}, \hat{d}$ such that

$$\sup_{s\in\mathbf{S}} \| \frac{\hat{M}_{y^a}(s,N(s)) - m_{y^a}(z,N_\theta(s))}{m_{y^a}(z,N_\theta(s))} \| < \epsilon,$$

where $\hat{M}_{y^a} = R^{\hat{d},\hat{n}}_{y^a}(N)$. 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.

DEFINITION 1 *A self-justified equilibrium consists of 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, $a < A$, *there is some* n_{y^a}, d_{y^a} *such that*

$$M_{y^a} = R_{y^a}^{n_{y^a},d_{y^a}}(N),$$

and there is no $(\bar{n}, \bar{d}) > (n_{y^a}, d_{y^a})$ *with*

$$c_{y^a}(\bar{n}, \bar{d}) + \int_{s\in\mathbf{S}} \left(\frac{\bar{M}_{y^a}(s, N(s)) - m_{y^a}(z, N_\theta(s))}{m_{y^a}(z, N_\theta(s))} \right)^2 d\mathbb{Q}^*(s) <$$
$$c_{y^a}(n_{y^a}, d_{y^a}) + \int_{s\in\mathbf{S}} \left(\frac{M_{y^a}(s, N(s)) - m_{y^a}(z, N_\theta(s))}{m_{y^a}(z, N_\theta(s))} \right)^2 d\mathbb{Q}^*(s),$$

where $\bar{M}_{y^a} = R_{y^a}^{\bar{n},\bar{d}}(N)$.

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 a non-parametric regression and that the forecasts trade off accuracy and computational costs in that more accurate forecasts impose prohibitively higher costs. Note that we do not require that (n, d) minimize the sum of the forecasting error and computational cost, as we only compare the current forecast to more complex forecasts. This is crucial for the existence of an equilibrium, but also makes economic sense if we assume that the agents regard the costs associated with producing the current forecasting function as sunk.

Similarly to the concept of "self-confirming" equilibrium (see e.g. Fudenberg and Levine (1993) or Cho and Sargent (2009)), 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

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

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 optimziation problem (1) and agents forecast future prices perfectly.

Under the assumptions stated in the next section, as the set of admissible regressions, $R_{y^a}^{n,d}(\cdot)$, becomes sufficiently rich, a self-justified equilibrium converges to a rational expectations equilibrium if costs are zero. The main contribution of this paper is to explore what happens if the agent is unable to approximate m_{y^a} perfectly. In this case, self-justified equilibria can be arbitrarily far from a rational expectations equilibrium.

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 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*

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

Individual endowments are positive, i.e.,

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

2. *The set $\mathbf{\Theta}$ is compact, and for each $y^a \in \mathbf{I}$, the set $\mathbf{\Theta}_{y^a}$ is a closed convex cone containing $\mathbb{R}^J_+$.*

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

4. *For all $y^a \in \mathbf{I}$ and all $\theta^-_{y^a} \in \mathbf{\Theta}_{y^a}$*

$$\theta^-_{y^a} \cdot f(q, z) \geq 0 \text{ for all } q \in \mathbb{R}^J_+, z \in \mathbf{Z}.$$

Assumption 1.4 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.

The crucial and non-standard assumption of the 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 $\widehat{\mathbf{S}} \subset \mathbf{S}$ 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 $\vec{\theta} \in \mathbf{\Theta}$.*

2. The measure $\mathbb{Q}(\cdot|z,\vec{\theta})$ is continuous in $\vec{\theta}$ for all $z \in \mathbf{Z}$, $\vec{\theta} \in \mathbf{\Theta}$.

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, $\mathbf{\Delta}^{ZG-1}$. Assumption 2.2 then simply states that this vector changes continuously in $\vec{\theta}$.

From a practical point of view, the 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 $\mathbf{\Theta}^-$. The assumption will allow us to obtain simple existence results below, however, it is certainly not a standard assumption in this strand of literature, and it is not compatible with full rationality of individuals.

Assuming finite accounting has several economic justifications. One interpretation is that actual portfolios lie in $\mathbf{\Theta}^-$, 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 beginning-of-period portfolios always lie on a finite grid as a technological constraint. This viewpoint seems natural when one thinks of the grid to be extremely fine. For this interpretation, let $\widehat{\mathbf{\Theta}}^- \subset \mathbf{\Theta}^-$ be a finite set, and assume that given $\vec{\theta}(z^t)$, we have

$$\vec{\theta}^-(z^{t+1}) \in \arg\min_{\vec{\theta}^- \in \widehat{\mathbf{\Theta}}^-} \|\bar{\theta} + \epsilon_{t+1} - \vec{\theta}^-\|_2,$$

with $\bar{\theta}_{y^a} = \theta_{y^{a-1}}$ for all $a = 2, \dots A$, $y^a \in \mathbf{Y}^a$, and $\bar{\theta}_{y^1} = 0$ for all $y^1 \in \mathbf{Y}$. In this formulation, ϵ_t should be interpreted as a (small) rounding error, and it is assumed that the support of $\epsilon(\cdot)$ is centered around zero, convex, and sufficiently small. We assume that ϵ_t is i.i.d. and that it only affects the current rounding error. In this formulation, it is easy to verify that Assumption 2.2 holds whenever ϵ_t has a continuous density function. Of course, the formulation of the agent's problem in (1) now potentially (depending on the set of admissible forecasting functions, $\mathbf{M}_{y^a}$) builds in another layer of bounded rationality, since the correct dynamic programming problem of an agent is no longer a standard convex program.

Assumption 2 guarantees that each $m^a_y(z,\vec{\theta})$ as defined in (4) is continuous in $\vec{\theta}$ for any selection of the equilibrium correspondence, $N(\cdot)$. Since we assume $\widehat{\mathbf{S}}$ to be finite and to contain GZ elements, for fixed $\vec{M} \in \mathbf{M}$, a selection of the temporary equilibrium correspondence can be viewed as a vector $N \in \left(\mathbb{R}^I_+ \times \mathbf{\Theta} \times \mathbb{R}^J_+\right)^{GZ}$. We write $M^{n,d}_{y^a}(\cdot;N) = R^{n,d}_{y^a}(N)$, and it is useful to note that for each z, the function $M^{n,d}_{y^a}(z,\cdot)$ is defined on a subset of the Euclidean space. To make this more explicit, we sometimes write the forecasting function as $M^{n,d}_{y^a}(z,\vec{\theta}^-,\vec{x},\vec{\theta},q,\nu)$, where $\nu \in \left(\mathbb{R}^I_+ \times \mathbf{\Theta} \times \mathbb{R}^J_+\right)^{GZ}$.

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

Assumption 3

1. *For all* $n, d \in \mathbb{N}$, *for all* $\nu \in \left(\mathbb{R}^I_+ \times \mathbf{\Theta} \times \mathbb{R}^J_+\right)^{GZ}$, *all* $s \in \widehat{\mathbf{S}}$ *and all* $\vec{\theta} \in \mathbf{\Theta}$, $\vec{x} \in \mathbb{R}^I_{++}, q \in \mathbb{R}^J_{++}$ *the function* $M^{n,d}_{y^a}(s, \vec{x}, \vec{\theta}, q; \nu)$ *is jointly continuous in* $\vec{x}$, $\vec{\theta}$, q, ν.

2. *For each agent* $y^a \in \mathbf{I}$, *and all* $n, d \in \mathbb{N}$ *all functions in in the range of* $R^{n,d}_{y^a}$ *are uniformly bounded above, i.e., there is some* $\bar{m}$ *such that*

$$M_{y^a,j}(z, \vec{\theta}^-, q, \vec{\theta}, \vec{x}) < \bar{m} \text{ for all } z \in \mathbf{Z}, \vec{\theta}^-, q, \vec{\theta}, \vec{x}, j \in \mathbf{J} \text{ and all } M \in \mathbf{M}^{d,n}_{y^a}.$$

3. *For each* y^a, $c_{y^a}(n, d)$ *is strictly increasing in* (n, d), *and for any* d, $(c_{y^a}(n, d) - c_{y^a}(n-1, d)) \to \infty$ *as* $n \to \infty$, *and for any* n, $(c_{y^a}(n, d) - c_{y^a}(n, d-1)) \to \infty$ *as* $d \to \infty$.

Assumption 3.1 is relatively standard and very likely to be satisfied in applied settings. Assumption 3.2 might appear to be rather strong. However, with enough structure on the sets $\mathbf{M}_{y^a}$, and with a more concrete description of the economy, one can typically find these bounds in an overlapping-generations setting. Clearly, with strictly positive endowments and borrowing constraints, all functions in $\mathbf{M}_{y^{A-1}}$ are bounded. A backward induction argument can then be used to justify Assumption 3.2. It is clear that in a framework with infinitely lived agents, this becomes much more difficult.

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. We fix (n_{y^a}, d_{y^a}) for all agents y^a, decompose the economy into sub-economies for each $s \in \mathbf{S}$ and construct a map from a compact and convex set of all agents' choices, prices, probabilities, μ, and forecasts, M_s, into itself. Using Kakutani's theorem (see Border (1989)), we can show that this map has a fixed point, and we finish the proof by demonstrating that for sufficiently large n and d, this is a self-justified equilibrium.

First, we need to find a suitable, convex and compact domain for the map. Assumption 1.3 implies that there exist numbers $l, r \in \mathbb{R}$ such that whenever $\vec{\theta} \in \mathbf{\Theta}$,

$$l \leq \theta_{y^a,j} \leq r \text{ for all } y^a \in \mathbf{I}, j \in \mathbf{J}.$$

Let the set of admissible asset holdings be $\mathbf{T} = [l, r]^J$, and let the set of admissible consumptions be

$$\mathbf{X} = \left[0, \max_{z \in \mathbf{Z}, y^a \in \mathbf{I}} \frac{e(z) + div(z)}{\eta(y^a)}\right].$$

We construct a upper-hemi-continuous, non-empty and convex-valued correspondence, $\mathbf{\Phi}$, 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,

$$\mathbf{\Phi} : (\mathbf{X}^I \times \mathbf{T}^I \times \Delta^J)^{GZ} \times \mathbf{\Delta}^{GZ} \rightrightarrows (\mathbf{X}^I \times \mathbf{T}^I \times \Delta^J)^{GZ} \times \mathbf{\Delta}^{GZ}.$$

For this construction, for all $y^a \in \mathbf{I}$ and all $s \in \widehat{\mathbf{S}}$, let

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

where

$$\widetilde{M}_{y^a} = R_{y^a}^{n_{y^a},d_{y^a}}(\nu), \quad \nu = (x_s, \theta_s, \frac{1}{p_s}q_s)_{s\in\widehat{\mathbf{S}}} \quad . \tag{5}$$

Define the price-players best response as

$$\mathbf{\Phi}_{0,s}(\vec{\theta}_s, \vec{x}_s) = \arg \max_{(p,q)\in\mathbf{\Delta}^J} p(\sum_{y^a\in\mathbf{I}} \eta(y^a)(x_{y^a,s} - e_{y^a}(z) - div(z))) + q \cdot (\sum_{y^a\in\mathbf{I}} \eta(y^a)(\theta_{y^a,s} - \bar{\theta})),$$

and let

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

Assumptions 1 - 3 guarantee that the mapping

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

is non-empty, convex valued, and upper hemi-continuous. By Kakutani's fixed point theorem, there exists a fixed point. Assumption 1 guarantees that the additional constraints imposed by forcing choices to lie in $\mathbf{T} \times \mathbf{X}$ are not binding, and hence the forecasting functions defined by (5) at the fixed point, together with $\mathbb{Q}^* = \mu$ and the equilibrium values constitute a candidate self-justified equilibrium for the given (n_{y^a}, d_{y^a}). Assumption 3.3 implies directly that since forecasting errors are always finite this must be an actual self-justified equilibrium if (n_{y^a}, d_{y^a}) are sufficiently large. □

The discretization of the state-space enables us to prove a very simple result. Without this, strong assumptions would be needed to ensure the existence of a recursive rational expectations equilibrium (see Brumm et al. (2017)), and the existence of a self-justified equilibrium thus would remain an open problem.

4 (Artificially) intelligent agents

Our definition of self-justified equilibria is very general, and it puts very little structure on agents' forecasts. We now want to examine economies were individuals have access to very good forecasting

technologies. We use concepts from the machine-learning literature (see, e.g., Scheidegger and Bilionis (2017)) to specify good regressions for the artificially intelligent agents in our model.

For this, we first need to specify a flexible functional form for forecasting functions and define our regression method $R(\cdot)$.

4.1 Least squares regularization

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 $k : \mathbf{X} \times \mathbf{X} \to \mathbb{R}$ a (positive definite) *kernel* if for any finite sequence $(x_j)_{j=1,...,n}$ 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 $\overline{\mathbf{X}} \subset \mathbf{X}$, any continuous function $f : \overline{\mathbf{X}} \to \mathbb{R}$ and any $\epsilon > 0$ there are finitely many $(x_i, c_i) \in \overline{\mathbf{X}} \times \mathbb{R}$ such that

$$\sup_{x \in \overline{\mathbf{X}}} |\sum_i c_i k(x, x_i) - f(x)| < \epsilon.$$

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.

$$k_{\text{SE}}(x, x') = \sigma^2 \exp\left\{ -\frac{1}{2} \sum_{i=1}^{k} \frac{(x_i - x_i')^2}{\ell_i^2} \right\}. \tag{6}$$

In this formulation the $\sigma^2, \ell_1, \ldots, \ell_k \in \mathbb{R}_+$ 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. Rasmussen and Williams (2005, Chapter 6)) endowed with an inner product $\langle ., . \rangle_{\mathbf{H}}$ which for $f = \sum_{i=1}^{s} \alpha_i k(., x_i)$ and $g = \sum_{j=1}^{r} \beta_j k(., t_j)$ satisfies

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

Given a data set $\{\left(x^{(i)}, y^{(i)}\right) | i = 1, ..., n\}$ consisting of n vectors $x^{(i)} \in \mathbb{R}^d$ and corresponding, potentially noisy, observations,

$$y^{(i)} = f(x^{(i)}) + \epsilon_i, \tag{7}$$

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., Evgeniou et al. (2000), and there references therein) solves the following problem:

$$\min_{f \in \mathbf{H}_k} \frac{1}{n} \sum_{i=1}^{n} (y^{(i)} - f(x^{(i)}))^2 + \lambda \|f\|_k^2, \tag{8}$$

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 (8) can be written as

$$\hat{f}(x) = \sum_{i=1}^{n} \alpha_i k(x, x_i), \tag{9}$$

where α solves

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

This construction gives us the first ingredient of the regression function $R^{n,d}$—that is, the agent takes n points along an equilibrium process and solves the regression problem (8). The optimal value of λ depends on the properties of the noise term. The determinants of this parameter are described below.

Unfortunately, this procedure is hardly feasible if the domain of the forecasting function is very high-dimensional. Therefore we have to use a method to reduce this dimension which will result in the determination of the second parameter, d.

4.2 Overcoming the curse of dimensionality

To make the concept of self-justified equilibrium tractable in this setting, it is essential to find a simple domain for agents' forecasts. So far, we allowed forecasts to depend on all current endogenous variables which is clearly too general to be useful in applications. In particular, the agents will face a *curse of dimensionality* (Bellman (1961)) when trying to approximate and evaluate functions on very high-dimensional domains.

The structure of the equilibrium suggests that it suffices to base forecasts only on the current shock and on (new) portfolio-choices across agents. As we will argue in the examples below, this often yields excellent results and is well suited for computational purposes. For the remainder of the paper, we assume that agents' forecasts do not depend on the current endogenous state, on prices, or consumption choices, and we write

$$M_{y^a} : \mathbf{Z} \times \mathbf{\Theta} \to \mathbb{R}^J_+.$$

There are obvious alternatives which are not considered in this paper. For example, following Brumm and Kubler (2014), forecasts could depend on current consumption-choices across agents. This would make the dimension of the domain independent of the number of assets and therefore be useful in models with many financial securities. From an economic point of view, it might make more sense to assume that households base their forecasts on current prices, and possibly on lagged shocks since these are easily observable. However, it is clear that current portfolio-choices determine the (natural) endogenous state in the next period, and it is, therefore, a good choice from the perspective of the computational modeler.

In many applications, the set of current asset holdings, $\mathbf{\Theta}$, will be very high dimensional (it is a subset of $\mathbb{R}^{IJ}$). Both as a matter of realism and for tractability, it seems advantageous to

assume that the agents only take a low dimensional subspace of the actual state-space and use this for their forecasts. In our tractable version of the model, we assume that agents project $\vec{\theta}$ into a lower dimensional subspace and use the latter for the forecasts. That is to say, M_{y^a} is actually not defined on $\mathbf{\Theta}$, but instead on a subset of $\mathbb{R}^d_+$, with d typically being much smaller than IJ.

4.2.1 What are good projections

Given a $d \times IJ$ projection matrix $W_{y^a,z,j}$ for a given agent y^a, shock z and asset j, we define the set of admissible forecasting-functions, i.e. the range of $R^{n,d}(\cdot)$, to be

$$\mathbf{M}^{d,n}_{y^a,z,j} \subset \{f : \mathbf{\Theta}^W_{y^a,z} \to \mathbb{R}\},$$

where

$$\mathbf{\Theta}^W_{y^a,z} = \{\phi \in \mathbb{R}^d : \phi = W^T_{y^a,z,j}\theta, \theta \in \mathbf{\Theta}_{y^a}\}.$$

For each shock $\bar{z} \in \mathbf{Z}$ and each asset $j \in \mathbf{J}$ the agent's regression $R^{n,d}_{y^a}$ then maps into $\mathbf{M}^{d,n}_{y^a,z,j}$. In choosing $W_{y^a,z,j}$, two extremes are conceivable. First, one could view the projection matrices, $W_{y^a,z,j}$, $y^a \in \mathbf{I}$, $z \in \mathbf{Z}$, $j \in \mathbf{J}$ as fundamentals–agents have certain technologies that allow them to observe projections of the state into lower dimensional subspaces (for example, they observe parts of the wealth distribution). Second, one could take d as given and require that the matrices $W_{y^a,z,j}$ are optimal in the sense that they minimize the mean squared error in Equation (3). This would fit our definition of a self-justified equilibrium, but unfortunately, it turns out to amount to a non-convex optimization problem that is generally not tractable.[4] The problem is so complicated that it is not consistent with the whole idea of boundedly rational agents. Moreover, we wish to develop a theory of optimal projections that are independent of the sets $\mathbf{M}^{n,d}_{y^a}$ and only depend on the function that needs to be approximated. This allows us to disentangle the methods used to approximate a d-dimensional function from the method used to find an optimal projection of the IJ-dimensional vector $\vec{\theta}$ into $\mathbb{R}^d$.

In the following, we take an approach that lies between the two extremes, and we believe that it has an elegant micro-foundation and turns out to be very tractable. In that approach, agents choose a matrix $W_{y^a,z,j}$ to minimize the unexplained part of the variations in $\hat{m}_{y^a,z,j}$ as measured by the mean squared derivative of $\hat{m}$ with respect to the orthogonal complement of the variables used for forecasting.

To this end, we assume that each agent y^a uses his own portfolio as the primary factor that influences next period's marginal utilities. This is a natural assumption, and if asset prices would only depend on the current and lagged shock, this would yield an optimal solution. However, in our model, asset prices vary with the distribution of assets in the economy. We therefore write θ_{-y^a} to denote the portfolio of all other agents in the economy besides agent y^a, and we write

[4] In Appendix A, we describe some of the difficulties that arise if one requires the matrix to be chosen optimally.

$\vec{\theta} = (\theta_{y^a}, \theta_{-y^a})$. θ_{-y^a} influences the agent's marginal utility for assets because it influences all future prices. We assume that the agent assesses the variability of his future marginal utility that is caused by this variability of future prices by the mean squared gradient and chooses a projection to ensure that the unexplained part of fluctuations is minimized.

To formalize this idea, let $D = IJ - J$ be the dimension of other agents' asset holdings and write forecasts as $M_{y^a,j}(z, \theta_{y^a}, W_{y^a,z,j}\theta_{-y^a})$. Without loss of generality, we assume that $W_{y^a,z,j}$ is an element of the d-dimensional Stiefel-manifold in $\mathbb{R}^D$, i.e.,

$$W_{y^a,z,j} \in \mathbf{V}_d(\mathbb{R}^D) = \left\{ A \in \mathbb{R}^{D \times d} : A^T A = I_{d \times d} \right\},$$

where $I_{d\times d}$ is the $d \times d$ identity matrix. Given a candidate $d \times D$ matrix $V_1 \in \mathbf{V}_d(\mathbb{R}^D)$, there is a $V_2 \in \mathbf{V}_{D-d}(\mathbb{R}^D)$ such that

$$[V_1, V_2] \begin{bmatrix} V_1^T \\ V_2^T \end{bmatrix} = I_{D \times D},$$

and we can write

$$m_{y^a,j}(z, \vec{\theta}) = m_{y^a,j}\left(z, \left(\theta_{y^a}, [V_1 V_2] \begin{bmatrix} V_1^T \\ V_2^T \end{bmatrix} \theta_{-y^a}\right)\right) = m_{y^a,j}\left(z, \theta_{y^a}, V_1 V_1^T \theta_{-y^a} + V_2 V_2^T \theta_{-y^a}\right).$$

Defining $\phi_1 = V_1^T \theta_{-y^a}$ and $\phi_2 = V_2^T \theta_{-y^a}$, we obtain a function

$$\widehat{m}_{y^a,j}(z, \theta_{y^a}, \phi_1, \phi_2) = m_{y^a,j}\left(z, \theta_{y^a}, V_1\phi_1 + V_2\phi_2\right).$$

Strengthening Assumption 2.2, assume for now that $\widehat{m}_{y^a}$ is continuously differentiable in θ_{-y^a} ($\mathbb{Q}^*$-a.e.). Given our justification for finite accounting, this simply amounts to assuming that the transition probability $\mathbb{Q}(\cdot|z, \vec{\theta})$ is continuously differentiable in θ and therefore does not seem substantially stronger than the original assumption. Nevertheless we will relax the assumption below.

We assume that the agent approximates the function $\widehat{m}_{y^a,j}$ using only (θ_{y^a}, ϕ_1). For simplicity assume for now that $\mathbf{M}_{y^a,j}$ consists of all (Borel-measurable) functions. For his case, we obtain

$$M_{y^a,j}(z, \theta_{y^a}, \phi_1) = \int_{\phi_2} \hat{m}_{y^a,j}(z, \phi_1, \phi_2) d\mathbb{Q}^*(\phi_2|z, \theta_{y^a}, \phi_1),$$

where $\widehat{\mathbb{Q}}^*(z, (\theta_{y^a}, \phi_1, \phi_2))$ denotes the invariant distribution over

$$(z, (\theta_{y^a}\phi_1, \phi_2)) = (z, N_{\theta_{y^a}}(s), V_1 N_{\theta_{-y^a}}(s), V_2 N_{\theta_{-y^a}}(s)),$$

which is induced by $\mathbb{Q}^*$, and $\widehat{\mathbb{Q}}^*(\phi_2|\theta_{y^a}, \phi_1, z)$ denotes the invariant distribution of ϕ_2 conditional on z, θ_{y^a}, and ϕ_1.

This approximation is justified if the impact of ϕ_2 on the function $\widehat{m}_{y^a,j}$ is relatively small. How do agents decide that the effect of ϕ_2 on next period's marginal utility is small? We assume in

this paper that they use the squared derivative with respect to ϕ_2, averaged along the stationary distribution, to measure the variability[5] with respect to ϕ_2, and define

$$\xi_{y^a,z,j}(V_1,V_2)=\int_{(\theta_{y^a},\phi_1,\phi_2)}\left(\nabla_{\phi_2}\widehat{m}_{y^a,j}\left(z,\theta_{y^a},\phi_1,\phi_2\right)\right)^T\left(\nabla_{\phi_2}\widehat{m}_{y^a,j}\left(z,\theta_{y^a},\phi_1,\phi_2\right)\right)d\widehat{\mathbb{Q}}^*(\theta_{y^a},\phi_1,\phi_2|z),$$

where for $x\in\mathbb{R}^D$,

$$\nabla_x f(x,y)=\begin{pmatrix}\frac{\partial f}{\partial x_1}\\ \vdots\\ \frac{\partial f}{\partial x_D}\end{pmatrix},$$

and the partial derivatives are taken to be one-sided derivatives at the boundary of the domain.

We assume that $W_{y^a,z,j}=V_1$ solves

$$\min_{(V_1V_2)\in\mathbf{V}_{IJ-J}(\mathbb{R}^{IJ-J})}\xi_{y^a,z,j}(V_1,V_2). \tag{10}$$

It turns out that there is a straightforward characterization of an optimal solution to this minimization problem—in stark contrast to the case where the projection is chosen to minimize the mean squared average forecasting error. In computational sciences, this is used in the "classical" active subspace approach (see Constantine et al. (2014)) which we discuss in detail in Section 5.

Note that

$$\nabla_{\phi_2}\widehat{m}_{y^a,j}(z,\theta_{y^a},\phi_1,\phi_2)=\nabla_{\phi_2}m_{y^a,j}\left(z,\theta_{y^a},V_1\phi_1+V_2\phi_2\right)=V_2^T\nabla_{\theta_{-y^a}}m_{y^a,j}(z,\theta_{y^a},\theta_{-y^a}),$$

and that

$$\begin{aligned}&\int_{(\theta_{y^a},\phi_1,\phi_2)}\left(\nabla_{\phi_2}\widehat{m}_{y^a,j}\left(z,\theta_{y^a},\phi_1,\phi_2\right)\right)^T\left(\nabla_{\phi_2}\widehat{m}_{y^a,j}\left(z,\theta_{y^a},\phi_1,\phi_2\right)\right)d\widehat{\mathbb{Q}}^*(\theta_{y^a},\phi_1,\phi_2|z)=\\&\int_{(\theta_{y^a},\phi_1,\phi_2)}tr\left(\left(\nabla_{\phi_2}\widehat{m}_{y^a,j}\left(z,\theta_{y^a},\phi_1,\phi_2\right)\right)^T\left(\nabla_{\phi_2}\widehat{m}_{y^a,j}\left(z,\theta_{y^a},\phi_1,\phi_2\right)\right)\right)d\widehat{\mathbb{Q}}^*(\theta_{y^a},\phi_1,\phi_2|z)=\\&\int_{(\theta_{y^a},\phi_1,\phi_2)}tr\left(\left(\nabla_{\phi_2}\widehat{m}_{y^a,j}\left(z,\theta_{y^a},\phi_1,\phi_2\right)\right)\left(\nabla_{\phi_2}\widehat{m}_{y^a,j}\left(z,\theta_{y^a},\phi_1,\phi_2\right)\right)^T\right)d\widehat{\mathbb{Q}}^*(\theta_{y^a},\phi_1,\phi_2|z).\end{aligned}$$

Therefore, solving (10) amounts to solving

$$\min_{V_2\in\mathbf{V}_{D-d}(\mathbb{R}^D)}tr\left(V_2^TC_{y^a,z,j}V_2\right),$$

where

$$C_{y^a,z,j}=\int_{(\theta_{y^a},\theta_{-y^a})}\left(\nabla_{\theta_{-y^a}}m_{y^a,j}\left(z,\theta_{y^a},\theta_{-y^a}\right)\right)\left(\nabla_{\theta_{-y^a}}m_{y^a,j}\left(z,\theta_{y^a},\theta_{-y^a}\right)\right)^Td\widehat{\mathbb{Q}}^*(\vec{\theta}|z). \tag{11}$$

Denoting $\lambda_1\geq\lambda_2\geq\ldots\geq\lambda_D$ as the eigenvalues of $C_{y^a,z,j}$, it follows from the Courant-Fischer Theorem (see, e.g., Horn and Johnson (1985), Theorem 4.2.11) that since $C_{y^a,z,j}$ is a symmetric

[5]Sobol and Kucherenko (2009) discuss several different approaches to estimate the influence of individual factors and groups of factors and show that many of them can be effectively bounded by the average squared gradient of the function.

matrix the minimum is given by $\sum_{i=d+1}^{D} \lambda_i$ and one (not unique) minimizer is given by the matrix of associated eigenvectors.

This suggests the following construction of projection-matrices: Since $C_{y^a,z,j}$ is symmetric positive definite, it admits the form

$$C_{y^a,z,j} = V \Lambda V^T, \tag{12}$$

where $\Lambda = \text{diag}(\lambda_1, \cdots, \lambda_D)$ is a diagonal matrix containing the eigenvalues of C in decreasing order, $\lambda_1 \geq \cdots \geq \lambda_D \geq 0$, and $V \in \mathbf{V}_D(\mathbb{R}^D)$ is an orthonormal matrix whose columns correspond to the eigenvectors of C, separating the d largest eigenvalues from the rest,

$$\Lambda = \begin{bmatrix} \Lambda_1 & 0 \\ 0 & \Lambda_2 \end{bmatrix}, \quad V = \begin{bmatrix} V_1 & V_2 \end{bmatrix}, \tag{13}$$

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

$$W_{y^a,z,j} = V_1. \tag{14}$$

Intuitively, $W_{y^a,z,j}$ rotates the input space in such a manner that the directions associated with the largest eigenvalues correspond to directions of maximal function variability (see Constantine (2015)).

Above's discussion gives directly rise to the following proposition (which is Lemma 2.2 in Constantine et al. (2014)), which makes the active subspace method very attractive for our model.

PROPOSITION 1 *The mean squared gradients of $\hat{m}$ with respect to ϕ_1 and ϕ_2 satisfy*

$$\int_{(\theta_{y^a},\phi_1,\phi_2)} \left(\nabla_{\phi_1} \widehat{m}_{y^a,j}\left(z, \theta_{y^a}, \phi_1, \phi_2\right)\right)^T \left(\nabla_{\phi_1} \widehat{m}_{y^a,j}\left(z, \theta_{y^a}, \phi_1, \phi_2\right)\right) d\widehat{\mathbb{Q}}^*(\theta_{y^a}, \phi_1, \phi_2|z) \leq \lambda_1 + \ldots + \lambda_d$$

and

$$\int_{(\theta_{y^a},\phi_1,\phi_2)} \left(\nabla_{\phi_2} \widehat{m}_{y^a,j}\left(z, \theta_{y^a}, \phi_1, \phi_2\right)\right)^T \left(\nabla_{\phi_2} \widehat{m}_{y^a,j}\left(z, \theta_{y^a}, \phi_1, \phi_2\right)\right) d\widehat{\mathbb{Q}}^*(\theta_{y^a}, \phi_1, \phi_2|z) \geq \lambda_{d+1} + \ldots + \lambda_{IJ-J}.$$

Both inequalities hold with equality if (V_1, V_2) are chosen according to (12) and (13).

As mentioned, the matrix consisting of the eigenvectors associated with the d largest eigenvalues is not the unique solution to (10). Nevertheless, we will assume in the following that each projection matrix $W_{y^a,z,j}$, $y^a \in \mathbf{I}$, $z \in \mathbf{Z}$, is determined by (12) and (13).

4.3 Self-justified equilibrium with (artificially) intelligent agents

The regression of an agent, $R^{n,d}$ now consists of two steps. First, the agent needs to estimate the matrices $C_{y^a,z,j}$ in Equation (11) using a finite number of observations. Then, the agent undertakes a regularized least square regression using a subset of these finite number of observations.

For the first step, it is natural to assume that the integral (11) 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,

$$\hat{C}_{y^a,z,j} = \frac{1}{N}\sum_{i=1}^{N} g^{(i)}\left(g^{(i)}\right)^T, \tag{15}$$

where

$$g^{(i)} = \left(\frac{m_{y^a,j}\left(z,\theta^i_{y^a},\theta^i_{-y^a}\right) - m_{y^a,j}\left(z,\theta^i_{y^a},\theta^i_{\widetilde{y^a}}+h,\theta^i_{-(y^a,\widetilde{y^a})}\right)}{h}\right)_{\widetilde{y^a}\neq y^a}$$

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 $y^{(i)} = m_{y^a,j}(z,\vec{\theta}(z^{t(i)}))$, $i=1,\ldots,n$, 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 (8). In our computational examples below, we determine this by maximum-likelihood.

To prove existence for fixed d, n, a version of the proof of Theorem 1 can be used. An important point to note is that the eigenvectors of C will change continuously as elements of C change continuously (keeping the matrix symmetric and definite)—see Horn and Johnson (1985). Unfortunately, the eigenvector associated with the largest eigenvalue will not change continuously at points where eigenvalues coincide. To prove existence, we, therefore, need to define the projections to stay constant outside of a region where the largest eigenvalue switches—a case that in practice never occurs.

5 Computation

To numerically approximate a self-justified equilibrium in a model where agents use optimal projections to form their forecasts, the main computational issues are (i) how to solve for the projection matrices W_{y^a} and for the forecasting functions M_{y^a}, and (ii) how to determine the parameters n and d of each agents' regression.

Clearly, the specification of the cost-function $c_{y^a}(\cdot)$ crucially determines how difficult it is to compute a self-justified equilibrium. For very high costs, computations become (almost) trivial, for very low costs they become impossible. The average error in Equation (3) can only become arbitrarily low in the limit as agents evaluate their forecasts using all other agents' asset holdings and as n becomes very large.

As mentioned above, our computational strategy is to solve for the parameters n and d that ensure a relatively low forecasting error and that have the property that a small increase in d does not change the quality of the forecast significantly. It turns out that Proposition 1 above often gives a simple rule for the determination of d, which is merely the dimension of the active subspace

(Constantine et al. (2014)). Furthermore, our representation of forecasting functions, M, can also be obtained as the posterior mean of a Gaussian process (see Rasmussen and Williams (2005)). The advantages of that formulation are 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 (8). Moreover, as it will become clear below, the standard deviation of the Gaussian process can be used as an indication of goodness of fit. This can give some indication on whether a higher value of n can lead to much higher accuracy.

To explain these two methods in some detail, we briefly discuss active subspace methods in Section 5.1, and Gaussian process regression in Section 5.2. Illustrative examples of their joint workings are provided in Appendix B. In Section 5.3 we then discuss our solution algorithm which combines active subspaces and GP-regression with a fairly standard, simulation-based technique to solve for the relevant points on the temporary equilibrium correspondence.

5.1 Active subspaces

As explained in detail above, for given d we compute optimal projections from the eigenvectors of $C_{y^a,z,j}$ (cf. (11)). Just like our economic agents, we cannot evaluate the matrices $C_{y^a,z,j}$ exactly. Instead, the usual procedure is to approximate the integral in (11) via Monte Carlo methods, and the gradients by finite differences.[6]

Computing the eigenvalues of $C_{y^a,z,j}$ often gives a simple way to determine d. If we observe sharp drops in the magnitude of the eigenvalue at the d-th largest eigenvalue, then this is a good candidate for a dimension that trades off accuracy and complexity. Active subspace methods are attractive in practice because it turns out that for many multivariate functions that occur for example in engineering models and the natural sciences, one observes such sharp drops in the spectrum of C at relatively small values of d (see, e.g., Constantine (2015), and the references therein). The active subspace is then a subspace of the input space which suffices for a good approximation of the underlying function.

In our iterative computational strategy described below, we start with $W_{y^a,z} = 0$, i.e., the agents only use their own asset holdings to forecasts future marginal utilities. If the resulting forecasting errors are large, we determine a sharp drop in the spectrum of C and increase d accordingly. In our computational examples, the first eigenvalue turns out to be several orders of magnitude large than all others. Accordingly, we project $\vec{\theta}_{-y^a}$ into a one-dimensional space. This results in a two-dimensional active subspace consisting of $\theta_{y^a}, W\theta_{-y^a}$.

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

[6] Alternatively, one may use the Bayesian information criterion to discover the active subspace. For the latter, see Tripathy et al. (2016).

approach which 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. (2017) solve heterogeneous agent macro models in continuous time by applying dimension reduction techniques to reduce the linear system of PDEs that characterizes their equilibrium—a setting that substantially differs from the one we are targeting here.

5.2 Gaussian process regression

We provide a very brief introduction to Gaussian process regression based on Rasmussen and Williams (2005) (see also Scheidegger and Bilionis (2017) for a more detailed introduction). Gaussian process (GP) regression is a method from Bayesian statistics (see, e.g., Barry (1986)) and now often used in supervised machine learning (see, e.g., Rasmussen and Williams (2005)). There are many examples in the literature where the combination of GPs and active subspaces proves very fruitful (see, e.g., Tripathy et al. (2016), or Scheidegger and Bilionis (2017)).

A GP is a collection of random variables, any finite number of which have a joint Gaussian distribution. We say that $f(\cdot)$ is a GP with *mean function* $m(\cdot)$ and *covariance function* $k(\cdot,\cdot)$, and write

$$f(\cdot) \sim \mathrm{GP}\left(m(\cdot), k(\cdot,\cdot)\right) \tag{16}$$

As mentioned above, we use a squared exponential kernel as defined in Equation (6). The numbers $\ell_i > 0$ and $\sigma > 0$ in the formula for the kernel denote the characteristic length-scale of the i-th input, and the signal strength. The "hyper-parameters" of the covariance function are typically estimated by maximum-likelihood (see Scheidegger and Bilionis (2017)). In our implementation, we use a self-customized version of the software package Limbo (see Cully et al. (2018)), which provides several options for this step.

The specification of the mean function $m(\cdot)$ is similar to the specification of a prior in Bayesian statistics. In our numerical examples below, we set $m(\cdot) = 0$. Note that this does not imply that the posterior mean (which we use as our approximating function) is zero. Rasmussen and Williams (2005, Chapter 2.7) discuss several ways to model a mean function.

Let us define the matrix of so-called "training inputs" as

$$X = \left\{x^{(1)}, \ldots, x^{(n)}\right\}. \tag{17}$$

Given X, we have a Gaussian prior on the corresponding response outputs,

$$\vec{f} = \left\{f\left(x^{(1)}\right), \ldots, f\left(x^{(n)}\right)\right\}.$$

In particular,

$$\vec{f}|X \sim \mathcal{N}\left(m, K\right), \tag{18}$$

where $m := m(X) \in \mathbb{R}^n$ being the mean function evaluated at all points in X, and $K \in \mathbb{R}^{n \times n}$ is the covariance matrix with

$$K_{ij} = k(x^{(i)}, x^{(j)}), \tag{19}$$

and $k(x^{(i)}, x^{(j)})$ given by (6).

In order to derive an explicit expression for the likelihood, we assume that the noise-terms ϵ_i in (7) are i.i.d. normal with mean zero and variance s^2. This assumption is not going to be satisfied in our application. However, it turns out that the method works well even if the noise is not i.i.d. normal. Using the independence of the observations, we obtain

$$y|\vec{f}, s \sim \mathcal{N}\left(y\middle|\vec{f}, s^2 I_n\right). \tag{20}$$

The *likelihood*-function of the observations is then given by

$$y|X, s \sim \mathcal{N}\left(y|m, K + s^2 I_n\right). \tag{21}$$

Bayes' rule combines the prior GP (see (16)) with the likelihood (see (21)) and yields the *posterior* GP

$$f(\cdot)|X, y, s \sim \mathcal{GP}\left(f(\cdot)\middle|\tilde{m}(\cdot), \tilde{k}(\cdot,\cdot)\right), \tag{22}$$

where the *posterior* mean and covariance functions are given by

$$\tilde{m}(x) = m(x) + K(x, X)\left(K + s^2 I_n\right)^{-1}(y - m) \tag{23}$$

and

$$\begin{aligned} \tilde{k}(x, x') &:= \tilde{k}(x, x';, s) \\ &= k(x, x') - K(x, X)\left(K + s^2 I_n\right)^{-1} K(X, x), \end{aligned} \tag{24}$$

respectively.

To carry out interpolation tasks when iterating on policies, one has to work with the predictive (marginal) distribution of the function value $f(x^*)$ for a single test input x^*. That is, given our posterior for the GP $f(\cdot)$, we can derive the marginal distribution of $f(\cdot)$ at any point. It reads

$$f(x^*)|X, y, s \sim \mathcal{N}\left(\tilde{m}(x^*), \tilde{\sigma}(x^*)\right), \tag{25}$$

where $\tilde{m}(x^*) = \tilde{m}(x^*)$ is the *predictive mean* given by (23), and $\tilde{\sigma}^2(x^*) := \tilde{k}(x^*, x^*; s)$ is the *predictive variance*.

Throughout our computations, we use the predictive mean as the value of the unknown function. Hence, we derive the same formula as in (9). As mentioned above, the advantage of this procedure is that we can use the maximum likelihood to estimate the hyper-parameters and s^2 from our training data. It can be shown (see Rasmussen and Williams (2005, Chapter 5.2)) that in the Gaussian setting, the marginal likelihood has the following form:

$$\log \mathcal{L}(y|X, \xi) = \frac{1}{2} y^T (K_f + s^2 I) y - \frac{1}{2} \log |(K_f + s^2 I)| - \frac{n}{2} \log 2\pi.$$

The latter can me maximized with respect to ξ.

Standard GPs are not able to deal with very high input dimensions because they rely on the Euclidean distance to define input-space correlations. Since the Euclidean distance becomes uninformative as the dimensionality of the input space increases, the number of observations required to learn the function grows exponentially. To this end, following Scheidegger and Bilionis (2017), we couple GPs to active subspaces, which is consistent with our economic modeling.

5.3 The basic computational strategy

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

   $$M^0_{y^a,z'} : \mathbf{Z} \times \mathbb{R}^d \to \mathbb{R}_+.$$

   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 $M^1_{y^a}$.

4. If

   $$\|M^1 - M^0\| < \eta$$

   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_{N,y^a} as defined in Equation 15 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.

The computation of the temporary equilibrium is done using a simple Newton-method, the derivatives needed for the computation of C_N (cf. (15)) are approximated using one-sided finite differences, and a self-customized version of Limbo (Cully et al. (2018)) is employed for the GP regressions.

6 An example

To illustrate the concept of self-justified equilibria and our general computational strategy, it is useful to focus on a specific simple example. Concretely, 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 $\eta_0(y) = \frac{1}{3}$, $y = 1, \ldots, 3$. 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 $a = 1, \ldots, 60$ 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

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

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

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

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

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

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

$$\pi = \begin{pmatrix} 0.95 & 0.05 \\ 0.05 & 0.95 \end{pmatrix}.$$

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.,

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

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. This system has to be solved at every simulation step 2(a) in our algorithm (see. Sec. 5.3).

6.1 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 in Krusell and Smith (1998) and Storesletten et al. (2007) where very simple forecasts also turn out to be very accurate in the calibrated model. It also fits with the results in Chien et al. (2011) who show that asset prices can often be forecasted accurately by a small number of lagged shocks.

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 perfect overlap between forecasted values

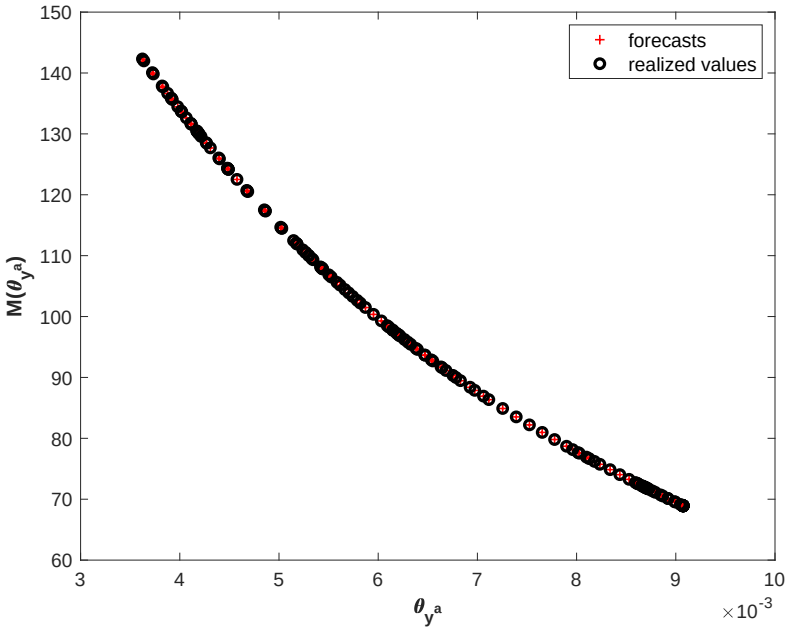

Figure 1: This figure shows the forecasts (red crosses) for the marginal utility of asset 1 of a 57-year-old agent of type 2 plotted against the average realized marginal utility (black bullets) of the tree for the exogenous shock being $z = 1$. Clearly, the own asset-choice gives an excellent forecasts.

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 points[7]. 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.

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

[7]In our simulation approach, the actual number of points varies in each iteration.

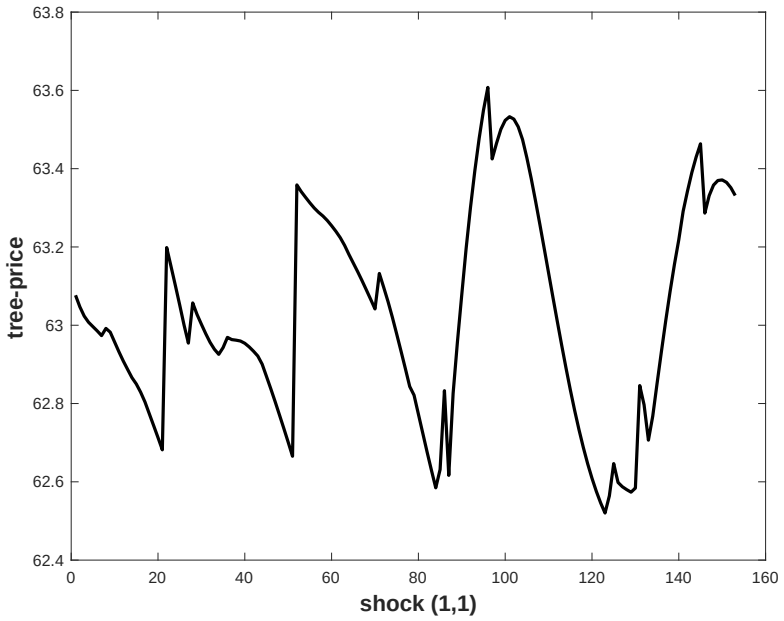

Figure 2: This figure depicts the variation in prices (given the current shock is 1 and the previous shock also was 1) which is not explained by shocks.

agent. There are variables in addition to the own asset holdings that have significant effects on the marginal utilities[8]

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

6.2 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.

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

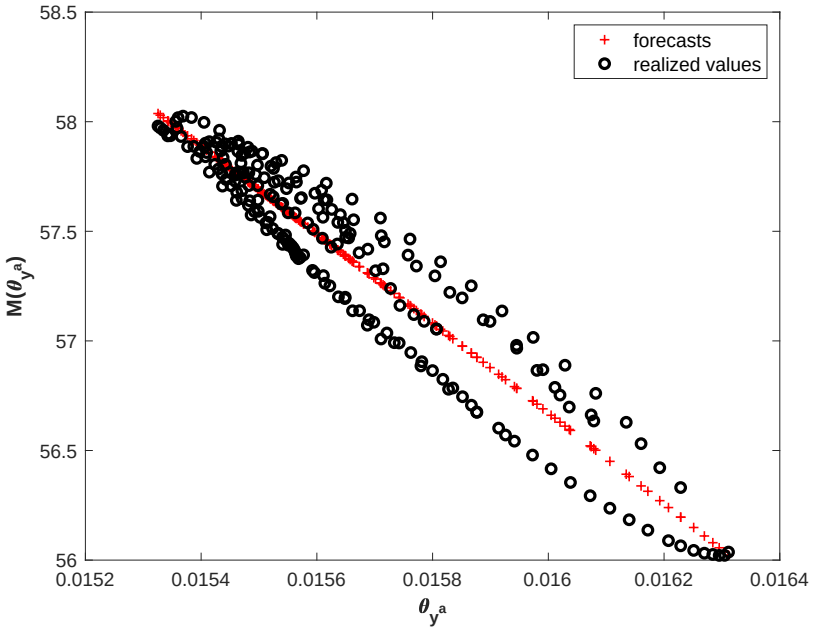

Figure 3: Forecasts of a 59-year-old agent of type 1 (red crosses) plotted against the average realized marginal utilities (black bullets) of the 60-year-old agent. It becomes obvious from this figure that the own asset-choice is insufficient for a good forecast.

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. (15)) 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 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

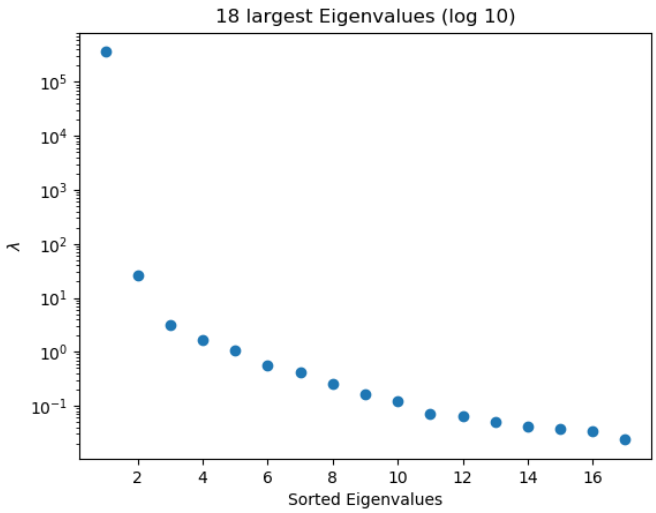

Figure 4: This figure depicts the largest eigenvalues for agent (1,59) and shock 2.

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. 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.

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. 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. This is why it is important to define the requirements of a self-justified equilibrium formally. Our analysis shows that increasing the dimension of the domain of forecasting functions further will lead to a tiny improvement in the quality of the forecasts. Even tiny costs suffice to rationalize the computed solution as a self-justified equilibrium.

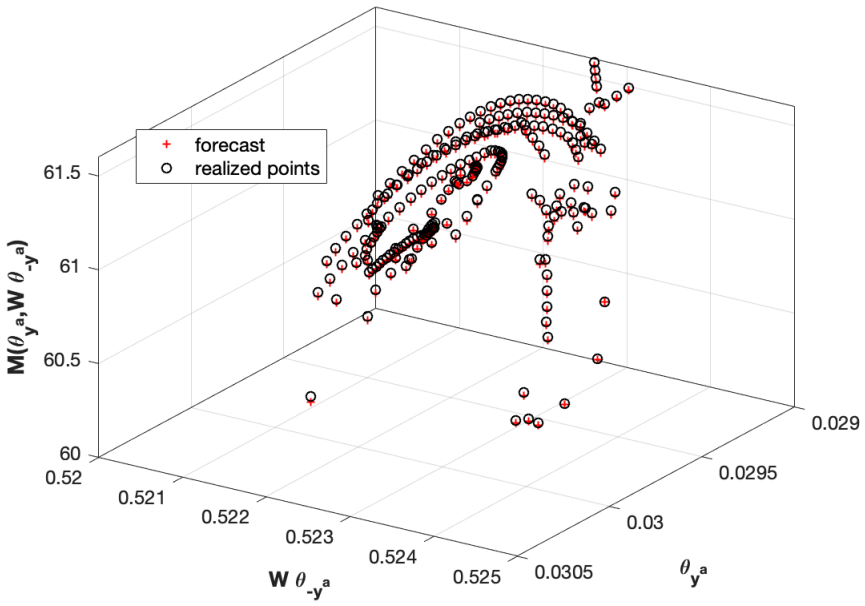

Figure 5: Forecasts of a 59-year-old agent of type 1 (red crosses) plotted against the average realized marginal utilities (black bullets) of the 60-year-old agent as a function of the two-dimensional active subspace. The forecasts based on a two-dimensional active subspace are now extremely accurate.

In this example, it was sufficient to consider forecasting functions which map from a two-dimensional domain. Increasing the number of assets will naturally increase the dimension of the domain of forecasting functions but it is important to note that in practical applications, standard GPs easily scale up to about 10 dimensions. The computation of the matrices C_N turns out to be relatively costly, but this does not have to be done in every iteration. Thus, the example in this section implies that GPs in conjunction with active subspaces are powerful enough to deliver very accurate forecasts in much more complicated models, where active subspaces might be substantially larger than one or two dimensions.

7 Conclusion

This paper makes three 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 argue that active subspace methods provide a natural way to formalize bounded rationality in very high dimensional models. Third, we provide an implementation to approximate self-justified equilibria numerically. In a relatively small model with 180 agents, we show that the method can potentially be used for large-scale applications.

We allow for the possibility of idiosyncratic shocks and a continuum of agents. However, in

our current implementation, when solving for the temporary equilibrium, we compute the optimal demand for each agent in the economy. If there is a continuum of agents (that differ ex-post by the realization of an idiosyncratic shock), one needs to aggregate groups of agents with similar wealth levels into one type of agent to make this step feasible. This adds another layer of approximation to our method but is very simple in practice.

In future research, we plan to consider production economies as well as economies with several consumption goods. While this is conceptually straightforward, it is not clear if the dimension of the active subspace is as low as in the example in Section 6.

Appendix A: Optimal ridge approximation and active subspaces

In our economic model, agents do not search for the optimal projection but are satisfied with finding an active subspace that reduces most of the "noise" from the forecasts. It turns out that the problem of finding an optimal projection is a difficult non-convex problem, but that the active subspace methods our agents use often provide reasonable approximations to an optimal projection.

Constantine et al. (2014) have 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

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

Theorem 3.1 in Constantine et al. (2014) now states that

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

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

Unfortunately, in this framework, Poincaré bounds are known to be far away from tight upper bounds (the exception being the standard normal distribution). Therefore, Theorem 3.1 in Constantine et al. (2014) does not tell us much about how far we are from an optimal projection.

The situation is slightly different if ρ is standard normal. In this case, the Poincaré constant is known to be 1, and it is easy to see that it can be obtained in a worst-case scenario. As Zahm et al. (2018) point out, this can be extended to non-standard normal densities. Assuming that the normal density has covariance matrix Σ, they show that If one takes as projection matrix

$$P = (\sum_i v_i v_i^T)\Sigma^{-1},$$

where (λ_i, v_i) solves

$$Cv_i = \lambda_i \Sigma^{-1} v_i,$$

one can obtain to following upper bound:

$$\int_x (f(x) - G(P^T x))^2 \rho(x)dx \leq (\lambda_{d+1} + \ldots + \lambda_D).$$

While our ergodic distributions are unlikely to be normal, the result is useful, since mixture of normal distributions typically can describe the distributions in our model.

An optimal projection can easily be defined, but hardly ever computed in higher dimensions. Suppose that for a given function $f : \mathbb{R}^d \to \mathbb{R}$ and a given $n \ll d$, one wants to find a a $n \times d$ matrix $V_1 \in \mathbf{V}_n(\mathbb{R}^d)$ that allows for an "optimal "approximation of $f(\cdot)$ by a function $g : \mathbb{R}^n \to \mathbb{R}$, setting

$$f(x) \simeq g(V_1 x).$$

We want to define optimality as minimizing the L^2 norm with respect to a probability density over $\mathbb{R}^d$, $\rho(x)$. For given V_1, we can define $V_2 = I - V_1 V_1^T$ and write $x = V_1^T y + V_2^T z$ for $y = V_1^T x$, $z = V_2^T y$. We can define $\widetilde{\rho}(y, z) = \rho(V_1 x + V_2 y)$ and marginal and conditional densities by the standard equations. The conditional expectation is

$$\mathbb{E}\left(f(x)|y\right) = \int f(V_1 y + V_2 z)\widetilde{\rho}(z|y)dz.$$

The optimal V_1 solves the following optimization problem:

$$\min_{V_1 \in \mathbf{V}_n(\mathbb{R}^d)} \int_x (f(x) - \mathbb{E}(f(x)|V_1^T x))^2 \rho(x) dx. \tag{26}$$

Unfortunately, this is a very complicated, non-convex optimization problem, and even the search for a stationary point turns out to be very costly in high dimensions (see e.g. Cohen et al. (2012)). Constantine et al. (2017) propose to use active subspace methods to obtain an approximation for a stationary point. Since the problem is non-convex, there is, unfortunately, no guarantee that the stationary point is, in fact, a minimum. However, Constantine et al. (2017) also provide various examples to illustrate that one can sometimes expect to obtain a good approximation from active subspaces.

Appendix B: Analytic examples

In this Appendix, we provide three instructive, synthetic examples to demonstrate the joint workings of Gaussian process regression and active subspaces.

As a first test case, we choose a function $f : \Omega \to \mathbb{R}$ with $\Omega = [-1, 1]^2$, observe it at a finite number of points, discover the active subspace, and then fit the GP hyper-parameters by maximizing the likelihood. Subsequently, we use the active subspace posterior mean as a surrogate of $f(\mathbf{x})$. Then, we randomly generate $N = 1,000$ test points uniformly drawn from Ω, denoted by $\mathbf{X} = \{\mathbf{x}^{(i)} : i = 1, \ldots, N\}$, and finally compute the average error, which we define as

$$e = \sqrt{\sum_{i=1}^{N} \left(f\left(\mathbf{x}^{(i)}\right) - \tilde{m}\left(\mathbf{x}^{(i)}\right)\right)^2 / f\left(\mathbf{x}^{(i)}\right)^2}. \tag{27}$$

Following [17], we choose a 2-dimensional function, namely,

$$f(x, y) = \exp(0.3x + 0.7y). \tag{28}$$

The analytical function given by (28) is depicted in the left panel of Figure 6. The arrows in the right panel of Figure 6 indicate that $[0.3, 0.7]$ is the direction in which this function varies the most, whereas in its orthogonal direction $[-0.7, 0.3]$, $f(\cdot)$ it is constant. The projection matrix $\mathbf{W}$ of the

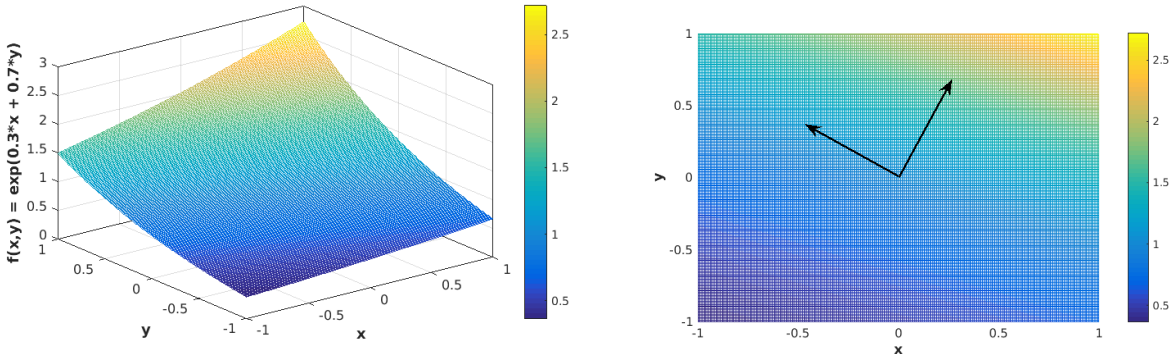

Figure 6: The left panel shows $f(x, y) = \exp(0.3x+0.7y)$. The right panel displays arrows indicating that the test function given by (28) varies the most in the direction [0.3, 0.7] and is constant in the orthogonal direction.

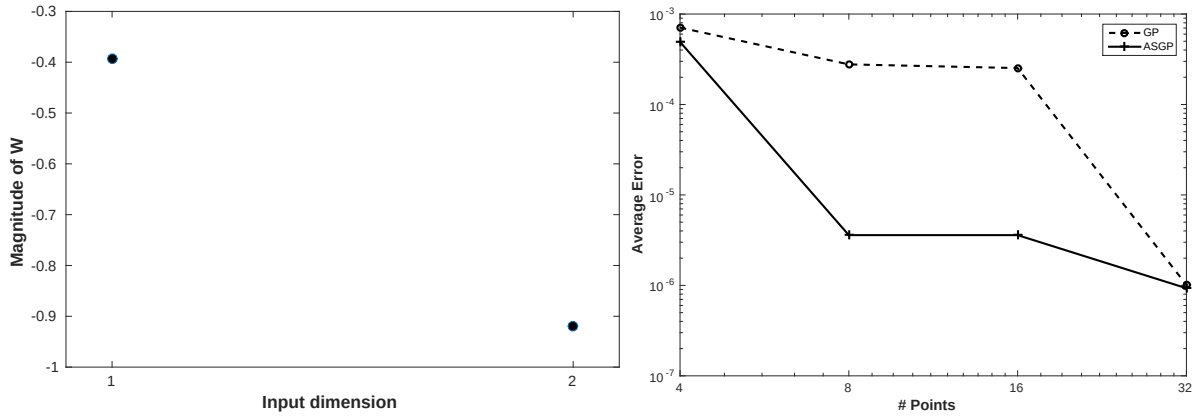

Figure 7: The left panel shows the projection matrix $\mathbf{W}$ of the 1-dimensional active subspace. The right panel displays a comparison of the interpolation error for 2-dimensional GPs (labelled as "GP" in the legend) and an 1-dimensional active subspace (labelled as "ASGP" in the legend) of varying resolution, respectively.

underlying 1-dimensional active subspace is shown in the left panel of Figure 7. The right panel of Figure 7 illustrates how the convergence of the interpolator constructed jointly by active subspaces and GPs compares to the performance of pure GPs with an increase in the number of sampling points. We see that the GPs in conjunction with active subspaces yield very competitive results with a considerably reduced computational burden. As a second example, we construct with active subspaces and GPs an interpolator of $f : [-1, 1]^{10} \to \mathbb{R}$

$$f(x_1, ..., x_{10}) = \exp(0.01x_1 + 0.7x_2 + 0.02x_3 + 0.03x_4 + 0.04x_5 + 0.05x_6 + 0.06x_7 + 0.08x_8 + 0.09x_9 + 0.1x_{10}). \quad (29)$$

Note that we deliberately put substantially more weight on the second dimension x_2. The left panel of Figure 8 shows the sorted eigenvalues of the matrix $\mathbf{C}_N$. The gap after the first eigenvalue reveals

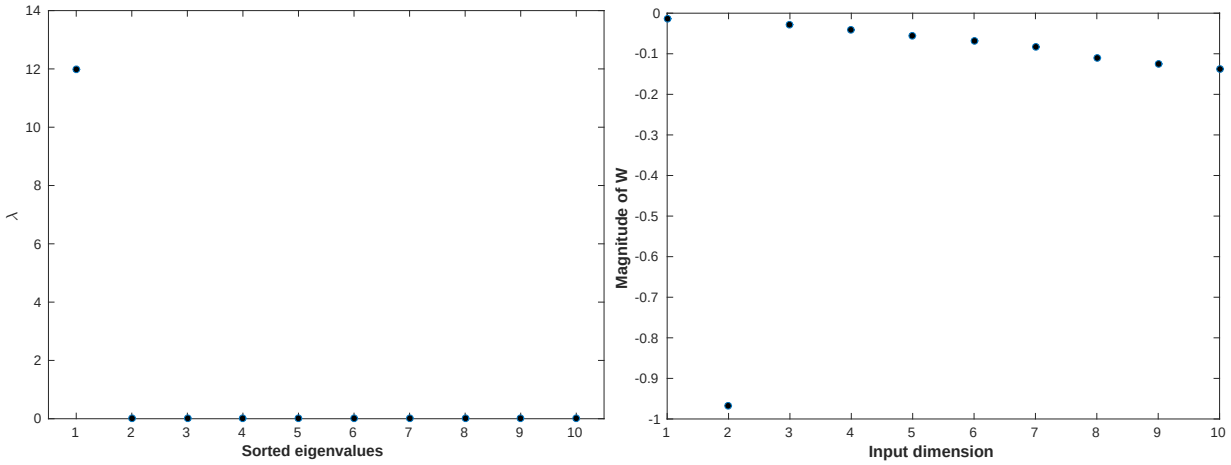

Figure 8: The left panel shows the sorted eigenvalues of $\mathbf{C}_N$ (see (15)), whereas the right panel displays the components of the projection matrix $\mathbf{W}$ of a 1-dimensional AS.

that $f(\cdot)$ has a 1-dimensional AS. In the right panel of Figure 8, we display again the projection matrix $\mathbf{W}$ of the 1-dimensional AS. It can be seen that, as one would expect from (29), the most dominant dimension is the second one. Third, we construct with active subspaces and GPs an interpolator of $f : [-1, 1]^{10} \to \mathbb{R}$

$$f(x_1, ..., x_{10}) = x_2 \cdot x_3 \cdot \exp(0.01x_1 + 0.7x_2 + 0.02x_3 + 0.03x_4 + 0.04x_5 + 0.05x_6 + 0.06x_7 + 0.08x_8 + 0.09x_9 + 0.1x_{10}). \quad (30)$$

The left panel of Figure 9 displays the sorted eigenvalues of the matrix $\mathbf{C}_N$. One can see that the gap in the spectrum only occurs after the third eigenvalue, indicating that the active subspace is a 3-dimensional space. The right panel of Figure 9 illustrates how the convergence (see (27)) of the interpolator with active subspace dimensions $d = \{1, 2, 3\}$ performs with an increasing number of training inputs. We can see that only the interpolator with an active subspace dimension 3 quickly converges to a satisfactory level of accuracy.

References

[1] Adam, K, A. Marcet and J.P. Nicolini, (2016), "Stock Market Volatility and Learning", *Journal of Finance*, 71, 33–82.

[2] Ahn, S, G. Kaplan, B. Moll, T. Winberry and C. Wolf, (2018), "When Inequality Matters for Macro and Macro Matters for Inequality", NBER Macroeconomic Annual,32, 1–75

[3] Barry, D., (1986), "Nonparametric Bayesian Regression". *Annals of Statistics*, 14, 934–953.

[4] Bellman, R., (1961), *Adaptive Control Processes: A Guided Tour*, Princeton University Press.

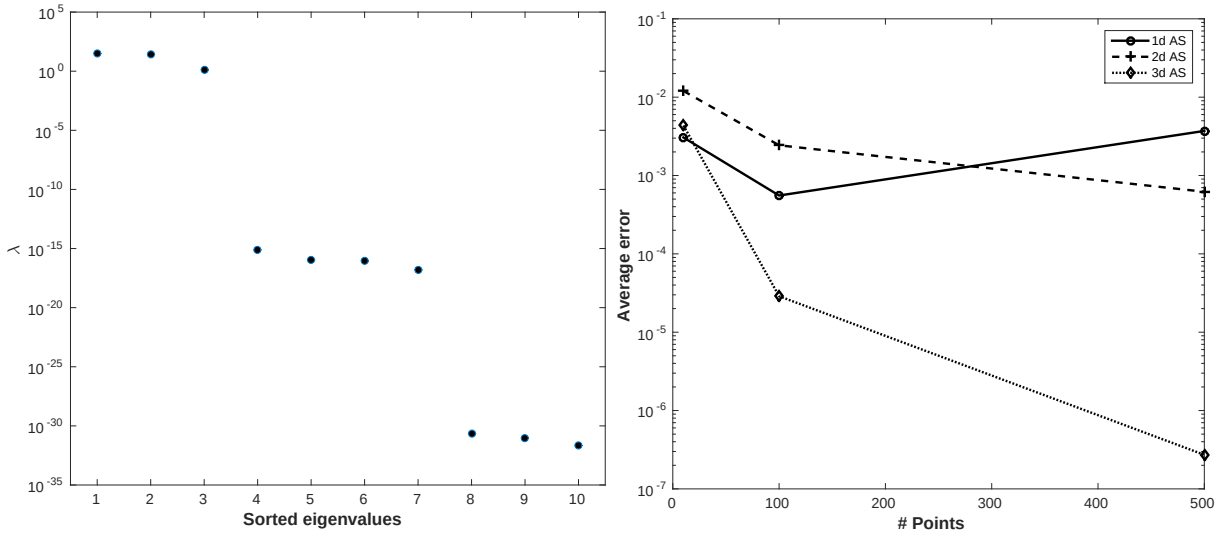

Figure 9: The left panel shows sorted eigenvalues of $\mathbf{C}_N$ (see (15)), whereas in the right panel we compare the interpolation error for $1, 2$, and 3-dimensional active subspaces (labelled by "AS" in the legend) as a function of increased sample size.

[5] Bengio, Y., O. Delalleau and N.Le Roux, (2005), "The curse of highly variable functions for local kernel machines", *Neural Information Processing Systems.*

[6] Bewley, T. (1984), "Fiscal and Monetary Policy in a General Equilibrium Model", Cowles Foundation Discussion Paper 690.

[7] Blei, D. and M Jordan,(2005), "Variational inference for Dirichlet process mixtures" *Bayesian Analysis* 1, 121–144.

[8] Border, K. C., (1989), *Fixed Point Theorems with Applications to Economics and Game Theory.* Cambridge University Press.

[9] Brumm, J. and F. Kubler (2014), "Applying Negishi's method to stochastic models with overlapping generations", working paper, University of Zurich.

[10] Brumm, J. , D. Kryczka and F. Kubler, (2017), "Recursive Equilibria in Dynamic Economies with Stochastic Production", (2017), *Econometrica*, 85, 1467–1499.

[11] Brumm, J. , F. Kubler and S. Scheidegger, (2017b), "Computing equilibria in dynamic stochastic macro-models with heterogeneous agents", *Advances in Economics and Econometrics, Eleventh World Congress*, 185–230.

[12] Chien, Y. and H. Lustig (2011), "The Market Price of Aggregate Risk and the Wealth Distribution," *The Review of Financial Studies* 23, 1596–1650.

[13] Chien, Y, H. Cole and H. Lustig, (2011), "A Multiplier Approach to Understanding the Macro Implications of Household Finance", *Review of Economic Studies* 78, 199–234.

[14] Cho, I.-K. and T. Sargent (2009), "Self-confirming equilibria"in *The New Palgrave Dictionary of Economics, 2nd edition*

[15] Citanna, A. and P. Siconolfi (2012), "Recursive equilibria in stochastic OLG economies: Incomplete markets," *Journal of Mathematical Economics*, 48, 322–337.

[16] Cohen, A., I. Daubechies, R. DeVore, G. Kerkyacharian and D. Picard, (2012), "Capturing ridge functions in high dimensions from point queries", *Constructive Approximation*, 35, 225 –243.

[17] Constantine, P., (2015), *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies* SIAM, Philadelphia.

[18] Constantine, P. 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.

[19] Constantine, P., A. Eftekhari, J. Hokanson and R. Ward, (2017), "A near-stationary subspace for ridge approximation", *Computer Methods in Applied Mechanics and Engineering*, 326, 402–421.

[20] Cully, A., K. Chatzilygeroudis, F. Allocati and J-B Mouret, (2018), "Limbo: A Flexible High-performance Library for Gaussian Processes modeling and Data-Efficient Optimization", preprint.

[21] Den Haan, W. J. and A. Marcet, (1990), "Solving the Stochastic Growth Model by Parameterizing Expectations," *Journal of Business and Economic Statistics*, 8, 31–34.

[22] Evans, R. and K.L. Phillips, (2014), "OLG Life Cycle Model Transition Paths: Alternate Model Forecast Method", *Computational Economics*, 43, 1095–1110.

[23] Evgeniou, T., M. Pontil and T. Poggio, (2010), "Regularization Networks and Support Vector Machines", *Advances in Computational Mathematics*, 13 : 1.

[24] Farmer, D.F and D. Foley, (2009), "The economy needs agent-based modeling", Nature Vol 460, 485–486

[25] Fudenberg, D. and D. Levine (1993), "Self-confirming equilibria", *Econometrica*, 61, 523–545.

[26] Gabaix, X, (2014), "A sparsity based model of bounded rationality", *The Quarterly Journal of Economics*, 129, 1661–1710.

[27] Gottardi, P. and F. Kubler (2015) "Dynamic Competitive Economies with Complete Markets and Collateral Constraints", *Review of Economic Studies*, 82, 1119–1153.

[28] Horn, R.A. and C.R. Johnson, (1985), *Matrix Analysis*, Cambridge University Press.

[29] Judd, K.L., L. Maliar and S. Maliar, (2011) "Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models", *Quantitative Economics* 2, 173–210.

[30] Krueger, D. and F. Kubler, "Computing equilibrium in OLG models with stochastic production", (2004), *Journal of Economic Dynamics and Control*, 28, 1411-1436.

[31] Krusell, P. and A. Smith (1998), "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896.

[32] Kubler, F. and K. Schmedders (2003), "Stationary Equilibria in Asset-Pricing Models with Incomplete Markets and Collateral," *Econometrica*, 71, 1767–1795.

[33] Kubler, F. and H.M. Polemarchakis (2004), "Stationary Markov Equilibria for Overlapping Generations," *Economic Theory*, 24, 623–643.

[34] Kurz, (1994), "On Rational Belief Equilibria, " *Economic Theory*, 4, 859–876.

[35] Ljungqvist, L. and T.J. Sargent, (2012), *Recursive Macroeconomic Theory.* MIT Press.

[36] Maliar, L. and S. Maliar (2014), "Numerical Methods for Large Scale Dynamic Economic Models", *Handbook of Computational Economics,* Volume 3, 325 – 477.

[37] Micchelli, C., Y. Xu and H. Zhang, (2006), "Universal Kernels", *Journal of Machine Learning Research*, 7, 2651–2667.

[38] Rasmussen, C. and C. Williams, Christopher, (2005), *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, MIT Press.

[39] Reiter, M., (2010), "Approximate and Almost-Exact Aggregation in Dynamic Stochastic Heterogeneous-Agent Models", IHS Economics Series 258.

[40] Sargent, T.J., (1993), *Bounded Rationality in Macroeconomics*, Oxford University Press.

[41] Scheidegger, S. and I. Bilionis (2017). "Machine learning for high-dimensional dynamic stochastic economies", working paper, University of Lausanne. Available at SSRN: http://dx.doi.org/10.2139/ssrn.2927400.

[42] Scheidegger, J. D. Mikushin, F. Kubler and O. Schenk, "Rethinking large-scale economic modeling for efficiency: optimizations for GPU and Xeon Phi clusters", (2018), *IPDPS 2018.*

[43] Sobol, I.M. and S. Kucherenko, (2009), "Derivative based local sensitivity measures and their link with global sensitivity indices". *Mathematics and Computers in Simulation*, 79, 3009–3017.

[44] Storesletten, K., C. I. Telmer, and A. Yaron (2007) "Asset pricing with idiosyncratic risk and overlapping generations." *Review of Economic Dynamics*, 10, 519 548.

[45] Tripathy, R., and I. Bilionis, and M. Gonzalez (2016) "Gaussian processes with built-in dimensionality reduction: Applications to high-dimensional uncertainty propagation ". *Journal of Computational Physics*, 321, 191–223.

[46] Woodford, M., (2013), "Macroeconomic Analysis without the Rational Expectations Hypothesis", *Annual Review of Economics*, 5, 303–346.

[47] Zahm, O., P. Constantine, C. Prieur and Y. Marzouk (2018), "Gradient-based dimension reduction of multivariate vector-valued functions", preprint.