

Self-justified equilibria: Existence and computation*

Felix Kubler

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

Simon Scheidegger

DBF, University of Zurich
simon.scheidegger@uzh.ch

June 18, 2018

Abstract

In this paper, we introduce the concept of “self-justified equilibrium” as a tractable alternative to rational expectations equilibrium in stochastic general equilibrium models with a large number of heterogeneous agents. A self-justified equilibrium is a temporary equilibrium where 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. The crucial assumption is that forecasting functions lie within a given class of simple functions and that they minimize long-run average forecasting errors among all functions in the class. 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 agents project current endogenous variables into a lower dimensional subspace and where the dimension of this subspace can be viewed as optimally trading off the accuracy of the forecast and its complexity. Using Gaussian process regression coupled with active subspaces as in Scheidegger and Bilonis (2017), we can solve models with hundreds to thousands of heterogeneous agents.

*We thank Victor Rios-Rull, Tom Sargent as well as seminar participants at Toulouse and at Louvain-La-Neuve for useful conversations and comments, and gratefully acknowledge financial support from PASC and the SFI.

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. In this paper, we develop an alternative to rational expectations equilibria and consider temporary equilibria with forecasting functions that are optimal within a given class, but that might lead to incorrect 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 the approach is as follows. In a temporary equilibrium, agents use current endogenous variables and the shock to forecast future marginal utilities for assets; prices for commodities and assets in the current period ensure that markets clear. The forecasts are assumed to be functions that lie in a pre-specified class (a simple example are semi-algebraic functions of fixed description complexity) - the agent chooses a function to minimize a loss function of average realizations of marginal utilities along the equilibrium path and the forecasts. In the temporary equilibrium, these expectations might be far from correct and agents might make significant mistakes. However, their forecasts are optimal given the agents’ constraints. The concept does not require identical expectations or identical forecasts across agents. Different types of agents can have different expectations and different forecasting functions.

To prove the existence of 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. In any case, while the assumption is necessary for the technical argument it does not affect the computed solutions since all computations are necessarily using finite precision arithmetics.

In our application, we consider a specific form for the forecasting function in that we assume that each agent projects the current endogenous variables into a relatively low dimensional subspace and approximates forecasts over this subspace globally.¹ Following Scheidegger and Bilonis (2017), we

¹We use the term “global solution” for a solution that is computed using equilibrium conditions at many points in the state space of a dynamic model—in contrast to a “local solution”, which rests on a local approximation around a

achieve this by combining Gaussian process regression (see, e.g., Rasmussen and Williams (2005))—a tool from supervised machine learning that can be used to approximate functions with prominent local features—with the exploitation of so-called active subspaces (see, e.g., Constantine et al. (2014)). Using this combination allows us to construct a method that determines an economically intuitive linear projection for a fixed dimension of the subspace. This combination 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.

We demonstrate that our computational method can be applied to large-scale heterogeneous agents models by applying it to an economy with 120 agents and segmented financial markets. We first consider the simplest case where an agent only uses his own asset-holding (together with the 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 in standard calibrations of the model. However, once we assume sufficient heterogeneity in tastes across generations, this simple method leads to large forecasting errors. We then use active subspace methods from 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 to almost zero. It is subject to further research to explore models where the dimension of the active subspace is larger.

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.

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 agents’ expectations can be computed by the modeler.

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 macro-economy dynamics. An agent-based model is a computerized simulation of a

steady state of the model.

number of decision-makers and institutions, which interact through prescribed rules. The agents can be as diverse as needed but in these agent-based models, behavioral rules are often arbitrary. Up to now, 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. 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 rest 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 special case which has the attractive features that it is tractable and that forecasts can be viewed as a trade-off between complexity and accuracy. 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 probability $\pi(z'|z)$. A history of shocks up to some date t is denoted by $z^t = (z_0, z_1, \dots, z_t)$ and called a date event. Whenever convenient, we use t instead of z^t .

At each date event, a continuum of ex-ante identical agents enter the economy, live for A periods, and differ ex-post by the realization of their idiosyncratic shocks. Each agent faces idiosyncratic shocks, y_1, \dots, y_A , that have support in a finite set \mathbf{Y}^A . We denote by $\eta_{y^a}(y_{a+1})$ the (conditional) probability of idiosyncratic shock y_{a+1} for an agent with shock history y^a , $\eta_0(y_1)$ to denote the probability of idiosyncratic shock y_1 at the beginning of life, and, $\eta(y^a)$ to denote the probability of a history of idiosyncratic shocks. We assume that the idiosyncratic shocks are independent of the aggregate shock, that they are identically distributed across agents within each type and age and, as in the construction in Proposition 2 in Feldman and Gilles (1985), that they “cancel out” in the aggregate, that is, the joint distribution of idiosyncratic shocks within a type ensures that at each history of aggregate shocks, z^t , for any $y^a \in \mathbf{Y}^a$ the fraction of agents with history $y^a = (y_1, \dots, y_a)$ is $\eta(y^a)$. This allows the 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 measurable functions of the current aggregate shock.² Each agent who can be identified by his date-event of birth, z^t , has a time-separable expected utility function

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

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

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

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

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

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

We define the state space to be $\mathbf{S} = \mathbf{Z} \times \Theta^-$ 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 but which are not standard.

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

²As 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} .

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 shock. While in rational expectations these expectations are always correct, the definition of a self-justified equilibrium simply requires them to be optimal within a restricted class of forecasting functions, given an agent's approximation to the invariant distribution. We, therefore, allow these forecasts to be imperfect and heterogeneous across agents.

In a temporary equilibrium each agent, $y^a \in \mathbf{I}$, is characterized by a function

$$M_{y^a} : \mathbf{S} \times \mathbb{R}_+^I \times \Theta \times \mathbb{R}_+^J \rightarrow \mathbb{R}_+^J,$$

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.

We denote by $\vec{M} = (M_{y^a})_{y^a \in \mathbf{I}}$ the forecasting functions across all agents. Throughout 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

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

The agent takes as given current average 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

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

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.,

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

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 selection $N(s)$, we write

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

It should be kept in mind that the function $N(s)$ also depends on \vec{M} . However, to simplify notation, we do not make this explicit.

The crucial innovation of this paper is to allow for heterogeneous and possibly incorrect forecasts across agents while still maintaining the assumption that agents are rational. For this, we assume that the agents deviate from rational expectations with respect to one crucial aspect: They cannot evaluate (or store) arbitrarily complicated functions, but instead approximate the equilibrium forecasts by “simple” functions. These functions could be relatively simple because they aggregate $\vec{\theta}$ into a lower dimensional vector (cf. Section 4 below), or because they belong to some convenient class of functions - a simple example would be semi-algebraic functions of fixed description complexity. For the definition of a self-justified equilibrium we therefore assume that agents are characterized by sets of admissible forecasting functions, \mathbf{M}_{y^a} , $y^a \in \mathbf{I}$, and we write $\mathbf{M} = \times_{y^a \in \mathbf{I}} \mathbf{M}_{y^a}$.

To make optimal current choices, agents need to know the marginal utility of their asset holdings in the next period. This is an equilibrium object since it depends on all future prices over the agent’s life-cycle. Given a selection $N(s)$ of the equilibrium correspondence, it is given by

$$m_{y^a}(z, \vec{\theta}) = \int_{s' \in \mathbf{S}} f(N_q(s'), z') \sum_{y_{a+1} \in \mathbf{Y}} \eta_{y^a}(y_{a+1}) u'_{y_{a+1}}(N_{x_{y_{a+1}}}(s')) d\mathbb{Q}(s'|z, \vec{\theta}) \quad (3)$$

Each agents y^a ’s forecast, M_{y^a} , is chosen from a (exogenously given) set of functions \mathbf{M}_{y^a} to minimize the average of the squared difference between the forecasted marginal utility and realized marginal utility, m_{y^a} , along an invariant distribution.

We then have the following definition.

DEFINITION 1 *A self-justified equilibrium consists of forecasts $\vec{M} \in \mathbf{M}$, 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 , $a < A$, M_{y^a} provides the best average approximation under this measure, i.e.

$$M_{y^a} \in \arg \min_{M \in \mathbf{M}_{y^a}} \int_{s \in \mathbf{S}} (M(s, N(s)) - m(z, N_\theta(s)))^2 d\mathbb{Q}^*(s)$$

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 the outcome 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, as long as they are the best forecasts the agent can choose.

For the special case where

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

we obtain a standard rational expectations equilibrium. The main contribution of this paper is to explore what happens if the agent is unable to approximate m_{y^a} perfectly.

3 Existence

To prove the existence of simple equilibria in heterogeneous agents models with incomplete markets, one needs to impose strong assumptions on fundamentals. Brumm et al. (2017) argue that without strong assumptions, simple equilibria might fail to exist (Kubler and Polemarchakis (2004) provide simple counterexamples).

3.1 Assumptions

We first make a number of fairly standard assumptions on fundamentals:

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 conditions

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

individual endowments are positive, i.e.,

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

2. The set Θ is compact, and for each $y^a \in \mathbf{I}$, the set Θ_{y^a} is a closed convex cone containing \mathbb{R}_+^J .
3. The payoff functions, $f : \mathbb{R}_+^J \times \mathbf{Z} \rightarrow \mathbb{R}^J$, are non-negative valued and continuous. Moreover, for any $i, j = 1, \dots, J$ 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 $\theta_{y^a}^- \in \Theta_{y^a}$

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

Assumptions 1.1-1.3 are standard (see, e.g., Kubler and Schmedders (2003)). 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 again 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 $\hat{\mathbf{S}} \subset \mathbf{S}$ such that the support of the transition function $\mathbb{Q}(\cdot|z, \vec{\theta})$ is a subset of $\hat{\mathbf{S}}$ for all $z \in \mathbf{Z}$ and all $\vec{\theta} \in \Theta$.

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

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

From a practical point of view, 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 not clear which of our results hold true in the limit as the grid becomes dense in Θ^- . The assumption will allow us to obtain simple existence results below, but it comes at the cost of some opaqueness.

Assuming finite accounting has several economic justifications. One interpretation is that actual portfolios lie in Θ^- but that agent cannot measure portfolios arbitrarily finely and make their decisions based on rounded values, exhibiting some degree of bounded rationality. Our preferred interpretation is that agents take the fact that 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 $\hat{\Theta}^- \subset \Theta^-$ be a finite set, and assume that given $\vec{\theta}(z^t)$, we have

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

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.

Since we assumed $\widehat{\mathbf{S}}$ 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 $N \in (\mathbb{R}_+^I \times \Theta \times \mathbb{R}_+^J)^{GZ}$. We make the following reduced-form assumption on forecasting and loss functions:

ASSUMPTION 3

1. For all $\mu \in \Delta^{ZG-1}$ and all $N \in (\mathbb{R}_+^I \times \Theta \times \mathbb{R}_+^J)^{GZ}$, the following

$$\widetilde{M}_{y^a}(N, \mu) = \arg \min_{M \in \mathbf{M}_{y^a}} \sum_{s \in \mathbf{S}} \mu(s) (M(s, N(s)) - m(z, N_\theta(s)))^2,$$

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

is well defined (i.e., the arg min exists and is unique). $\widetilde{M}_{y^a}(N, \mu)$ is a function of $(s, N(s))$ and is assumed to be jointly continuous in $(N, \mu, \vec{\theta}_-, q, \vec{\theta}, \vec{x})$ for all $z \in \mathbf{Z}$.

2. For each agent $y^a \in \mathbf{I}$, all functions in \mathbf{M}_{y^a} are uniformly bounded above, i.e., there is some \bar{m} such that

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

Assumption 3.1 is relatively standard and very likely to be satisfied in applied settings. Assumption 3.2 is a bit more problematic. 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 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, we can show that this map has a fixed point, and we finish the proof by demonstrating that 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 l, r such that whenever $\vec{\theta} \in \Theta$,

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

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

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

We construct an upper-hemi-continuous, non-empty and convex-valued correspondence, Φ , mapping choices and prices at each element in $\hat{\mathbf{S}}$ as well as a probability measure over $\hat{\mathbf{S}}$, $(\mathbf{X}^I \times \mathbf{T}^I \times \Delta^J)^{GZ} \times \Delta^{GZ}$ to itself, which has a fixed point. For all $y^a \in \mathbf{I}$ and all $s \in \hat{\mathbf{S}}$, let

$$\begin{aligned} \Phi_{y^a, s}((x_s, p_s, q_s)_{s \in \hat{\mathbf{S}}}) &= \arg \max_{x \in \mathbf{X}, \theta \in \Theta_{y^a} \cap \mathbf{T}} u_{y^a}(x) + \widetilde{M}_{y^a}(z, \vec{\theta}_s^-, \vec{q}_s, \vec{\theta}_s, \vec{x}_s) \cdot \theta \text{ s.t.} \\ &\quad (x_{y^a} - e_{y^a}(z)) + \theta_{y^a} \cdot \frac{1}{p_s} q_s - \theta_{y^a}^- \cdot f\left(\frac{1}{p_s} q_s, z\right) \leq 0 \end{aligned}$$

where

$$\widetilde{M}_{y^a} = \arg \min_{M \in \mathbf{M}_{y^a}} \sum_{s \in \mathbf{S}} \mu(s) \left(M(z, \vec{\theta}_s^-, \vec{q}_s, \vec{\theta}_s, \vec{x}_s) - m(s) \right)^2, \quad (4)$$

with

$$m(s) = \sum_{s' \in \mathbf{S}} \mathbb{Q}(s' | z, \vec{\theta}_s) f\left(\frac{1}{p_{s'}} q_{s'}, z'\right) \sum_{y^{a+1} \in \mathbf{Y}} \eta_{y^a}(y^{a+1}) u'_{y^{a+1}}(x_{y^{a+1}}(s')).$$

Define the price-players best response as

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

and let

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

Assumptions 1 - 3 guarantee that the mapping $\Phi = \times_{s \in \mathbf{S}} ((\times_{y^a \in \mathbf{I}} \Phi_{y^a, s}) \times \Phi_{0, s}) \times \Phi_\mu$,

$$\Phi : (\mathbf{X} \times \mathbf{T} \times \Delta^J)^{GZ} \times \Delta^{GZ-1} \rightrightarrows (\mathbf{X} \times \mathbf{T} \times \Delta^J)^{GZ} \times \Delta^{GZ-1}$$

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 (4) at the

fixed point, together with $\mathbb{Q}^* = \mu$ and the equilibrium values constitute a self-justified equilibrium.

□

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 A tractable version of the model

To make the concept of self-justified equilibrium tractable, it is essential to find a simple domain for agents' forecasts. The structure of the equilibrium suggests that this might consist of 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. However, note that in principle it is also possible to use other variables for computations. However, this is beyond the scope of the present work.

For the rest of the paper we assume that agents forecasts do not depend on the current endogenous state, on prices, or on consumption choices and we write

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

In many applications, the set of current asset holdings Θ will be very high dimensional. Both as a matter of realism and for tractability, it seems advantageous to assume that the agents only take a low dimensional part of the actual state-space and use this for their forecasts. In our tractable version of the model, we assume that agents take a linear projection of $\vec{\theta}$ into a lower dimensional subspace and use the latter for the forecasts. That is to say, M_{y^a} is actually not defined on Θ , but instead on a subset of \mathbb{R}_+^d , with d typically being much smaller than IJ . The agents use so-called ridge functions (see Pinkus (2015)) to approximate future marginal utilities.

4.1 Discovering the relevant dimensions of the state space

Given a $d \times IJ$ projection matrix $W_{y^a, z}$ for a given agent y^a and shock z , we define

$$\mathbf{M}_{y^a, z}(d) = \{f : \Theta_{y^a, z}^W \rightarrow \mathbb{R}^J\},$$

where

$$\Theta_{y^a, z}^W = \{\phi \in \mathbb{R}^d : \phi = W_{y^a, z}^T \theta, \theta \in \Theta_{y^a}\}.$$

For each $\bar{z} \in \mathbf{Z}$, the agent's forecasting function solves

$$\min_{M \in \mathbf{M}_{y^a, \bar{z}}(d)} \int_{\vec{\theta} \in \Theta} \left(M(\bar{z}, W_{y^a, \bar{z}}^T N_{\vec{\theta}}(\bar{z}, \vec{\theta})) - m(\bar{z}, N_{\vec{\theta}}(\bar{z}, \vec{\theta})) \right)^2 d\mathbb{Q}^*(\vec{\theta} | \bar{z}). \quad (5)$$

At this point, we impose no restrictions on the set \mathbf{M}_{y^a} (as we will explain below, this implies that the solution to (5) is simply the conditional expectation) but focus on the question of a sensible choice of d and the matrices $W_{y^a,z}$.

Without loss of generality we assume that $W_{y^a,z}$ is an element of the d -dimensional Stiefel-manifold, i.e.,

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

where $I_{d \times d}$ is the $d \times d$ identity matrix. In choosing $W_{y^a,z}$, two extremes are conceivable. First, one could view the projection matrices, $W_{y^a,z}$, $y^a \in \mathbf{I}$, $z \in \mathbf{Z}$, as fundamentals—agents have certain technologies that allow them to observe projections of the state into lower dimensional subspaces (for example, they observe the mean wealth distribution as well as conditional means). Second, one could take d as given and require that the matrices $W_{y^a,z}$ are optimal in the sense that they minimize some mean squared error. In the following, we take the approach that lies between the two extremes, and we believe that it has an elegant micro-foundation. In that approach, agents are “satisfied” with a given projection matrix $W_{y^a,z}$ if there are not apparent improvements possible. In the Appendix, we describe some of the difficulties that arise if one requires the matrix to be chosen optimally. While the problem is in principle well-posed, its solution is so complicated that it is not consistent with the whole idea of boundedly rational agents.

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 $\vec{\theta} = (\theta_{y^a}, \theta_{-y^a})$. Clearly, 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 future prices by the mean squared gradient, and chooses an “active subspace” (see Section 5.1 below) to ensure that the unexplained part of fluctuations is at most a ϵ -fraction of total fluctuations.

Formally, given a candidate $n \times (IJ - J)$ projection matrix $V_1 \in \mathbf{V}_n(\mathbb{R}^{IJ-J})$, there is a $V_2 \in \mathbf{V}_{IJ-J-n}(\mathbb{R}^{IJ-J})$ such that

$$[V_1, V_2] \begin{bmatrix} V_1^T \\ V_2^T \end{bmatrix} = I_{(IJ-J) \times (IJ-J)},$$

and we can write

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

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

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

Strengthening Assumption 2.2, we assume that \widehat{m}_{y^a} is continuously differentiable in θ_{-y^a} . 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.

We assume that the agent approximates the function \widehat{m}_{y^a} using only (θ_{y^a}, ϕ_1) , i.e.,

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

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} 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 with respect to ϕ_2 . 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.

We assume that the agent is satisfied with a $(IJ - J) \times n$ matrix V_1 if it explains a fraction $1 - \epsilon$ of the total variation of future marginal utilities, i.e.,

$$\frac{\int_{(\theta_{y^a}, \phi_1, \phi_2)} (\nabla_{\phi_2} \widehat{m}_{y^a}(z, \theta_{y^a}, \phi_1, \phi_2))^T (\nabla_{\phi_2} \widehat{m}_{y^a}(z, \theta_{y^a}, \phi_1, \phi_2)) d\widehat{\mathbb{Q}}^*(z, \theta_{y^a}, \phi_1, \phi_2)}{\int_{(\theta_{y^a}, \phi_1, \phi_2)} (\nabla_{\phi_1} \widehat{m}_{y^a}(z, \theta_{y^a}, \phi_1, \phi_2))^T (\nabla_{\phi_1} \widehat{m}_{y^a}(z, \theta_{y^a}, \phi_1, \phi_2)) d\widehat{\mathbb{Q}}^*(z, \theta_{y^a}, \phi_1, \phi_2)} < \epsilon, \quad (6)$$

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

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

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

In some of the numerical examples below, the agent will be satisfied with $V_1 = 0$, i.e., only take his own asset-position to forecast future marginal utilities. In this case, we have that ϕ_2 is of full dimension and

$$\frac{\int_{(\theta_{y^a}, \phi_2)} (\nabla_{\phi_2} \widehat{m}_{y^a}(z, \theta_{y^a}, \phi_2))^T (\nabla_{\phi_2} \widehat{m}_{y^a}(z, \theta_{y^a}, \phi_2)) d\widehat{\mathbb{Q}}^*(z, \theta_{y^a}, \phi_2)}{\int_{(\theta_{y^a}, \phi_2)} (\nabla_{\theta_{y^a}} \widehat{m}_{y^a}(z, \theta_{y^a}, \phi_2))^T (\nabla_{\theta_{y^a}} \widehat{m}_{y^a}(z, \theta_{y^a}, \phi_2)) d\widehat{\mathbb{Q}}^*(z, \theta_{y^a}, \phi_2)} < \epsilon.$$

4.2 Self-justified equilibrium in a tractable economy

With this, an economy is described by assets, trading constraints, preferences and endowments, but also $\epsilon_{y^a}(d)$, $d = J, \dots, IJ$ for all (active) agents $y^a \in \mathbf{I}$. We allow ϵ to depend on the dimension d to

incorporate the possibility that an agent prefers to explain little with a low dimensional projection than explaining a lot using a very high-dimensional function. We also allow ϵ to depend on the agent to incorporate heterogeneity in forecasts into the model.

A self-justified equilibrium with satisficing projections then consists of $(IJ - J) \times d_{y^a}$ matrices $W_{y^a, z}$ for each agent, y^a , and each shock, z , such that for each agent and each shock, inequality (12) holds with $V_1 = W_{y^a, z}$ and $\epsilon = \epsilon_{y^a}(d)$ as well as a selection $N(\cdot)$ of the temporary equilibrium correspondence, $\mathbf{N}_{\vec{M}}(\cdot)$, and a measure \mathbb{Q}^* on $(\mathbf{S}, \mathcal{S})$, such that

- \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 $\mathbf{B} \in \mathcal{S}$

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

- For each y^a $a < A$ and each $z \in \mathbf{Z}$, $M_{y^a, z}(z, W_{y^a, z} \vec{\theta})$ is equal to the conditional expectation of \hat{m} , given ϕ_2 , i.e.,

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

where ϕ_1, ϕ_2 , and \hat{m} are as defined above, and $\hat{\mathbb{Q}}^*(\phi_2|z, \theta_{y^a}, \phi_1)$ denotes the invariant distribution of ϕ_2 conditional on z, θ_{y^a} and ϕ_1 .

We make use of the well-known fact that the conditional expectation solves the least-squares problem (5).

In the next part of the paper, we will describe computational methods to solve for this self-justified equilibrium efficiently. For this, it is important to first note that we impose “too much” rationality on the agent to be able to solve his problem exactly. The fact that forecasts minimize the least-squared error under the (a priori unknown) invariant distribution makes it impossible to compute the forecast exactly. Instead, we will have to resort to Monte-Carlo simulations and approximate the invariant distribution by finitely many draws. At the same time, we are apparently also unable to compute the conditional expectation exactly. Hence, we will need to approximate this using a numerical method. To this end, we choose Gaussian process regressions. The latter has been proven to be very useful in other contexts (see Scheidegger and Bilonis (2017)).

5 Computation

To numerically approximate a self-justified equilibrium in a model where agents use satisficing projections to form their forecasts the main computational issues are (i) how to find satisficing projections, (ii) how to approximate the (low) dimensional forecasting functions well, and (iii) how to solve for them.

5.1 Finding $W_{y^a, z}$

In order to compute satisficing projections, we use so-called active subspace methods developed by Constantine et al. (2014) (see also Scheidegger and Bilonis (2017)).

We first review the basics of the approach: to approximate a very high dimensional function $f : \mathbb{R}^D \rightarrow \mathbb{R}$, we assume that it can be reasonably well approximated by the following form:

$$f(x) \approx h(W^T x), \quad (7)$$

where the matrix $W \in \mathbb{R}^{D \times d}$ projects the high-dimensional input space, \mathbb{R}^D , into a low-dimensional *active subspace*, \mathbb{R}^d , $d \ll D$. $h : \mathbb{R}^d \rightarrow \mathbb{R}$ is a d -dimensional function that is commonly termed *link* function. Note that the representation of Eq. 7 is not unique. All matrices W whose columns span the same subspace of \mathbb{R}^D yield identical approximations. Thus, without loss of generality, we restrict our attention to matrices in the Stiefel manifold, $W \in \mathbf{V}_d(\mathbb{R}^D)$.

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

$$C := \int (\nabla f(x)) (\nabla f(x))^T \rho(x) dx, \quad (8)$$

where

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

Since C is symmetric positive definite, it admits the form

$$C = \Lambda V^T, \quad (9)$$

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

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

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

$$W = V_1. \quad (10)$$

Intuitively, W rotates the input space in such a manner that the directions associated with the largest eigenvalues correspond to directions of maximal function variability (Constantine (2015)).

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

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

It is impossible to evaluate Eq. 8 exactly. Instead, the usual practice is to approximate the integral in Eq. 8 via Monte Carlo, that is, assuming that the observed inputs are drawn from $\rho(x)$, one approximates C using the observed gradients by

$$C_N = \frac{1}{N} \sum_{i=1}^N g^{(i)} \left(g^{(i)} \right)^T. \quad (11)$$

In practice, the eigenvalues and eigenvectors of C_N are found using the singular value decomposition of C_N . Clearly in our framework, the gradient, G^i cannot be evaluated analytically (in fact they are not guaranteed to exist), so we generally approximate (11) by finite differences.³

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 sharp drops in the spectrum of C at relatively small values of d (see Constantine (2015) and the references therein).

Constantine et al. (2014) prove the following theoretical result which makes the active subspace method very attractive for our model.

LEMMA 1 *The mean squared gradients of f with respect to y and z satisfy*

$$\int (\nabla_y f)^T (\nabla_y f) \rho(x) dx = \lambda_1 + \dots + \lambda_n$$

and

$$\int (\nabla_z f)^T (\nabla_z f) \rho(x) dx = \lambda_{n+1} + \dots + \lambda_d.$$

This Lemma now allows us to construct the desired projection matrices $W_{y^a, z}$ for a self-justified equilibrium in the tractable economy. Lemma 1 states that in order to find the desired projection matrices for a given agent y^a and a given shock z , we simply have to find a d such that

$$\frac{\lambda_{d+1} + \dots + \lambda_{IJ}}{\lambda_1 + \dots + \lambda_d} < \epsilon_{y^a}(d), \quad (12)$$

where the λ_i are the eigenvalues of C_n as defined in (11), where g^i is the finite difference gradient of $m(z_i, N_\theta(s_i))$ with respect to θ_{-y^a} , and where $(s_i)_{i=1}^N$ denotes a simulated path of equilibrium realizations of the state.

In our iterative computational strategy described below, we start with a simple guess for $W_{y^a, z}$ and update along the iterations.

To make the algorithm operable, we first need to understand how to conveniently approximate functions on arbitrary domains. For this, we use so-called Gaussian process (GP) regression, which is a method from supervised machine learning (see, e.g, Rasmussen and Williams (2005)) There are many examples in the literature where the combination of GP-regression and active subspaces proves very fruitful (see, e.g., Tripathy et al. (2016), or Scheidegger and Bilonis (2017)).

³Alternatively, one may use the Bayesian information criterion to discover the active subspace. For the latter, see Tripathy et al. (2016).

5.2 Gaussian process regression

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

$$y^{(i)} = f(x^{(i)}) + \epsilon_i, \quad (13)$$

we 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} 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:

$$\min_{f \in \mathbf{H}} \frac{1}{n} \sum_{i=1}^n (y_i - f(x^{(i)}))^2 + \lambda \|f\|_K^2, \quad (14)$$

where $\|\cdot\|_K$ is the norm defined by $K(\cdot)$. It can be shown that the solution to Eq. 14 can be written as

$$\hat{f}(x) = \sum_{i=1}^n \alpha_i K(x, x_i), \quad (15)$$

where α solves

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

As Rasmussen and Williams (2005) point out, the representation of f can also be obtained as the posterior mean of a Gaussian process. The advantages of that formulation are that it naturally leads to systematic ways for choosing $K(\cdot)$ and λ and that the standard deviation of the Gaussian process can be used as an indication of goodness of fit. We provide a very brief introduction to Gaussian process regression based on Rasmussen and Williams (2005) (see also Scheidegger and Bilonis (2017) for a more detailed introduction).

A Gaussian process 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

$$f(\cdot) \sim \text{GP}(m(\cdot), k(\cdot, \cdot)) \quad (16)$$

The covariance function can be chosen, but must be positive semi-definite and symmetric. Throughout our work, we either use the so-called *square exponential* (SE)

$$k_{\text{SE}}(x, x') = \sigma^2 \exp \left\{ -\frac{1}{2} \sum_{i=1}^r \frac{(x_i - x'_i)^2}{\ell_i^2} \right\}, \quad (17)$$

or the Matern-3/2 covariance kernel:

$$k_{\text{mat}}(x, x') = \sigma^2 \left(1 + \sqrt{3} \sum_{i=1}^l \frac{(x_i - x'_i)^2}{\ell_i^2} \right) \exp \left(-\sqrt{3} \sum_{i=1}^l \frac{(x_i - x'_i)^2}{\ell_i^2} \right), \quad (18)$$

where $\ell_i > 0$ and $\sigma > 0$ in both kernels denotes 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 Bilonis (2017)). In our implementation we use 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 implementation below we take $m(\cdot) = 0$. Note that this does not imply 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

$$X = \{x^{(1)}, \dots, x^{(n)}\}. \quad (19)$$

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

$$\vec{f} = \{f(x^{(1)}), \dots, f(x^{(n)})\}.$$

In particular,

$$\vec{f}|X \sim \mathcal{N}(m, K), \quad (20)$$

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

$$K_{ij} = k(x^{(i)}, x^{(j)}), \quad (21)$$

and $k(x^{(i)}, x^{(j)})$ given by Eqs. 17 or 18.

In order to derive an explicit expression for the likelihood, we assume that the noise-terms ϵ_i in Equation (13) are i.i.d. normal with mean zero and variance s^2 . Clearly, 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

$$y|\vec{f}, s \sim \mathcal{N}\left(y|\vec{f}, s^2 I_n\right). \quad (22)$$

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

$$y|X, s \sim \mathcal{N}\left(y|m, K + s^2 I_n\right). \quad (23)$$

Bayes' rule combines the prior GP (see Eq. 16) with the likelihood (see Eq. 23) and yields the *posterior* GP

$$f(\cdot)|X, y, s \sim \mathcal{GP}\left(f(\cdot)|\tilde{m}(\cdot), \tilde{k}(\cdot, \cdot)\right), \quad (24)$$

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

$$\tilde{m}(x) = m(x) + K(x, X) (K + s^2 I_n)^{-1} (y - m) \quad (25)$$

and

$$\begin{aligned} \tilde{k}(x, x') &:= \tilde{k}(x, x'; , s) \\ &= k(x, x') - K(x, X) (K + s^2 I_n)^{-1} K(X, x), \end{aligned} \quad (26)$$

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. We obtain,

$$f(x^*)|X, y, s \sim \mathcal{N}(\tilde{m}(x^*), \tilde{\sigma}(x^*)), \quad (27)$$

where $\tilde{m}(x^*) = \tilde{m}(x^*)$ is the *predictive mean* given by Eq. 25, 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 Equation (15). The advantage of this procedure is that we can use maximum likelihood to estimate the hyper-parameters and s^2 from our training data. In principle, it would be useful also to make use of the variance-covariance term that indicates how accurate the forecast is at that point. Incorporating this into our economic model is subject to further research.

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 enormously. To this end, following Scheidegger and Bilonis (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 Gaussian regression and the repeated solution of non-linear systems of equations. In particular, we employ an iterative scheme to solve for the optimal forecasting functions.

The basic details of the algorithm are then as follows:

1. Initial guess for each agent's forecasting:

Initially, we 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 Gaussian processes whose posterior means approximate

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

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 Gaussian process whose posterior mean gives a new forecasting function $M_{y^a}^1$

4. If

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

then set $M^* = M^1$. Else set $M^0 = M^1$ and repeat time iteration step 2.

5. Compute C_N as defined in Equation 11 and its eigenvalues, λ . If all agents' satisficing criteria (12) are satisfied, terminate. Else include one more eigenvector of C_N into the projective matrix W_{y^a} , make a new initial guess for Gaussian processes 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 are approximated using one-sided finite differences.

6 A simple example

In order to illustrate the concept of self-justified equilibria our general computational strategy, it is useful to focus on a specific simple example. In the simplest example, we assume that agents live for A periods and that there are two types of agents per generation and no idiosyncratic shocks. An agent is then characterized by (y, a) , where $y = 1, 2$ denotes the initial shock. The agents distinguish themselves by trading constraints and preferences. Type 1 agents can trade in a single Lucas-tree and in Arrow securities. In our framework, it is useful to assume that the Arrow-securities pay in the Lucas-tree (as in Gottardi and Kubler (2015) or Chien and Lustig (2011)). Type 2 agents can only trade in the Lucas tree. Both agents face borrowing constraints.

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.,

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

We can combine $\kappa_{i,a}$ and $\theta_{i,a}$ into one variable and obtain a system with $(A - 1)Z + (A - 1) + Z$ equations and unknowns. This system has to be solved at every simulation step 2 (a) in our algorithm and is the most time-consuming part of the computation.

6.1 A simple self-justified equilibrium with accurate forecasts

For the simplest example, assume that $A = 60$, $Z = 2$. All agents have CRRA utility functions with $u_{y,a}(c) = \beta_y^a \frac{c^{1-\gamma_y}}{1-\gamma_y}$. We take $\beta_y = 0.96$ for $y = 1, 2$, and $\gamma_1 = 2$, $\gamma_2 = 0.5$. Individual endowments are

$$e_{y,a}(1) = 0.4 + a/500, e_{y,a}(2) = 0.9 * (0.4 + a/500) \text{ for } a < 50,$$

$$e_a(1) = e_a(2) = 0.3 \text{ for } a \geq 50.$$

Moreover, we also assume that $d(z) = 2$ for both $z = 1, 2$, and that $\pi(1) = \pi(2) = \frac{1}{2}$.

We start off 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 1 (that pays if shock 1 realizes) only affects marginal utility in shock 1 and asset 2 only affect marginal utility in shock 2.

In the computed self-justified equilibrium, forecasting errors, as measured by the maximal relative deviation between forecasted marginal utilities and realized marginal utilities, are tiny. In particular, they are the smallest for young agents (around 0.001) and the largest for old agents (around 0.01). Average errors across agents are about 0.001. Moreover, forecasts are almost linear. In Figure 1a, we show the forecast of agent of age 5 and type 1 for the next period as a function of his asset holding in asset 1.

This result is, of course, consistent with many examples in the literature, where one finds pseudo aggregation (most notably Krusell and Smith (1996)) and Chien and Lustig (2011), but also Storesletten et al. (2007)). The main reason why the simple forecasts are well in this example is that there is almost no variation in asset holdings and that asset prices are mostly a function of the current and past exogenous shock. In Figure 1b, we show the asset prices for the 2 exogenous shock and confirm that there is indeed very little variation.

-FIGURE 1 ABOUT HERE-

Clearly the linearity of forecasts is an artifact of this particular example. Within our simple model we need to introduce more heterogeneity in tastes in order to obtain larger price-volatility. As it turns out, the result that linear forecasts are quite accurate holds true for a wide variety of parameter specifications. To go beyond this simple model, we, therefore construct an example where forecasts that do not take into account the wealth distribution across agents do not do a very good job.

6.2 Moving away from the simple example

One particular case where the simplicity of forecasts breaks down can be obtained by assuming that agents across generations have different subjective beliefs over the aggregate shocks. While this does not completely fit our model and does not fit the idea that the agents know invariant distributions, it gives us a modeling testbed to compare different algorithms. In particular, it can be incorporated into our model if we assume that the Bernoulli utility also depends on a history of shocks. In this case, all formal results in this paper go through (but the notation becomes more cumbersome).

In the concrete case, we modify the simple example above by assuming that agents of type 1 and ages 55-59 have incorrect probabilities in that

$$\pi^a(1) = 0.8, \pi^a(2) = 0.8 \text{ for } a = 50, \dots, 59$$

All other agents have the correct beliefs. While this does not exactly fit our model description where we assumed that all agents have identical beliefs, it is easy to modify the model for this specification.

With this specification, forecasts are systematically misspecified—not only because they are linear, but mainly because future marginal utilities for asset holding do not only depend on own choices. Figure 2 depicts the same forecasting function as Figure 1, but for this specification with heterogeneous beliefs. We can see that linear functions do not do a good job. Moreover, it seems that other variables have to be added to make forecasts accurate.

[Figure 2 about here]

Surprisingly, the active subspace is two-dimensional. In addition to an agent’s own asset holding, as 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, weighted (roughly) by the agents marginal propensity to consume. Employing a higher dimensional space to forecast future marginal utilities turns out to add very little. We compute the matrix $C_N(11)$ by Monte-Carlo draws and finite differences and find that one single Eigenvalue (in addition to the ones associated with own asset holdings) dominates all others. Increasing the dimension of the projective space from 2 to 3 hence has thus has almost effect on the forecasting power

[Figure 3 about here]

In Figure 3, we plot all Eigenvalues on a log-scale. The figure confirms that all other Eigenvalues are negligibly small compared to the one that corresponds to the agent’s own asset holdings and the weighted sum of asset holdings across agents.

Adding the additional variable then turns out to reduce forecasting errors to almost zero, comparable to the case in the Section above.

In principle GP-regressions scale up to 7-10 dimensions. The simple example in this section illustrates that this is likely to be enough to obtain very accurate forecasts even in much more

complicated models.

7 Conclusion

This paper makes three contributions. First, we define the concept of self-justified equilibrium as a natural generalization of rational expectations equilibrium, and we provide sufficient conditions for 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 120 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 optimal demand for each agent in the economy. If a continuum of agents, one needs to aggregate 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.

Future research includes production economies as well as economies with several consumption goods.

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

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

where C is the Poincaré constant that depends on the pdf 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 worse 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

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

where (λ_i, v_i) solves

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

one can obtain to following upper bound:

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

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. It is subject to further research to explore this in more detail. In any case, even this is not the optimal projection.

An optimal projection can easily be defined, but hardly ever computed in higher dimensions. Suppose that for a given function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and a given $n \ll d$, one wants to find 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 $g : \mathbb{R}^n \rightarrow \mathbb{R}$, 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 $\tilde{\rho}(y, z) = \rho(V_1 x + V_2 y)$ and marginal and conditional densities by the standard equations. The conditional expectation is

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

The optimal V_1 solves the following optimization problem:

$$\min_{V_1 \in \mathbf{V}_n(\mathbb{R}^d)} \int_x (f(x) - \mathbb{E}(f(x)|V_1^T x))^2 \rho(x) dx. \quad (28)$$

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.

References

- [1] Adam, K, A. Marcat and J.P. Nicolini, (2016), “Stock Market Volatility and Learning”, *Journal of Finance*, 71, 33–82.
- [2] Bengio, Y., O. Delalleau and N.Le Roux, (2005), “The curse of highly variable functions for local kernel machines”, *Neural Information Processing Systems*.
- [3] Bewley, T. (1984), “Fiscal and Monetary Policy in a General Equilibrium Model”, Cowles Foundation Discussion Paper 690.
- [4] Blei, D. and M Jordan,(2005), “Variational inference for Dirichlet process mixtures” *Bayesian Analysis* 1, 121–144.
- [5] Brumm, J. , D. Kryczka and F. Kubler, “Recursive Equilibria in Dynamic Economies with Stochastic Production”, (2017), *Econometrica*, 85, 1467–1499.
- [6] Chien, Y. and H. Lustig (2011), “The Market Price of Aggregate Risk and the Wealth Distribution,” *The Review of Financial Studies* 23, 1596-1650.
- [7] Cho, I.-K. and T. Sargent (2009), “Self-confirming equilibria”in *The New Palgrave Dictionary of Economics, 2nd edition*
- [8] Citanna, A. and P. Siconolfi (2012), “Recursive equilibria in stochastic OLG economies: Incomplete markets,” *Journal of Mathematical Economics*, 48, 322–337.
- [9] 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.
- [10] Constantine, P., (2015), *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies* SIAM, Philadelphia.
- [11] 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.
- [12] 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.
- [13] 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.

- [14] Evgeniou, T., M. Pontil and T. Poggio, (2010), “Regularization Networks and Support Vector Machines”, *Advances in Computational Mathematics*, 13 : 1.
- [15] Farmer, D.F and D. Foley, (2009), “The economy needs agent-based modeling”, *Nature* Vol 460, 485–486
- [16] Fudenberg, D. and D. Levine (1993), “Self-confirming equilibria”, *Econometrica*, 61, 523–545.
- [17] Gabaix, X, (2014), “A sparsity based model of bounded rationality”, *The Quarterly Journal of Economics*, 129, 1661–1710.
- [18] Gottardi, P. and F. Kubler (2015) “Dynamic Competitive Economies with Complete Markets and Collateral Constraints”, *Review of Economic Studies*, 82, 1119–1153.
- [19] Krueger, D. and F. Kubler, “Computing equilibrium in OLG models with stochastic production”, (2004), *Journal of Economic Dynamics and Control*, 28, 1411-1436.
- [20] Krusell, P. and A. Smith (1998), “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106, 867–896.
- [21] Kubler, F. and K. Schmedders (2003), “Stationary Equilibria in Asset-Pricing Models with Incomplete Markets and Collateral,” *Econometrica*, 71, 1767–1795.
- [22] Kubler, F. and H.M. Polemarchakis (2004), “Stationary Markov Equilibria for Overlapping Generations,” *Economic Theory*, 24, 623–643.
- [23] Kurz, (1994), “On Rational Belief Equilibria,” *Economic Theory*, 4, 859–876.
- [24] Ljungqvist, L. and T.J. Sargent, (2012), *Recursive Macroeconomic Theory*. MIT Press.
- [25] Pinkus, A. *Ridge Functions*, Cambridge University Press, Cambridge, 2015.
- [26] Rasmussen, C. and C. Williams, Christopher, (2005), *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, MIT Press.
- [27] Sargent, T.J., (2002), *Bounded Rationality in Macroeconomics*
- [28] Scheidegger, S., Bilonis, I. (2017). “Machine learning for high-dimensional dynamic stochastic economies”, working paper, University of Zurich.
- [29] 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.
- [30] Storesletten, K., C. I. Telmer, and A. Yaron (2007) “Asset pricing with idiosyncratic risk and overlapping generations.” *Review of Economic Dynamics*, 10, 519–548.

- [31] Tripathy, R., and I. Bilonis, and M. Gonzalez (2016) “Gaussian processes with built-in dimensionality reduction: Applications to high-dimensional uncertainty propagation”. *Journal of Computational Physics*, 321, 191–223.
- [32] Woodford, M., (2013), “Macroeconomic Analysis without the Rational Expectations Hypothesis”, *Annual Review of Economics*, 5, 303–346.
- [33] Zahm, O., P. Constantine, C. Prieur and Y. Marzouk (2018), “Gradient-based dimension reduction of multivariate vector-valued functions”, preprint.