Wealth distribution over the business cycle
A mean-field game with common noise

Thomas Bourany*

Master in Mathematics & Applications, UPMC – Sorbonne Université, Paris

M.A. in Social Sciences – Economics, The University of Chicago

Link for the complete version

Master thesis UPMC-Sorbonne / Paris-Diderot

September 2018

Abstract

This paper revisits the standard "heterogeneous agents" model with aggregate shocks by applying mathematical methods from Mean Field Games (MFG). If the Aiyagari-Bewley-Huggett framework is a typical example of MFG, as formulated Achdou, Han, Lasry, Lions, and Moll (2017), it is denoted "MFG with common noise" in presence of aggregate risk. Solving such class of models usually suffers from the "curse of dimensionality" arising when behaviors and the distribution of agents interact, and we develop novel methods to address this issue. By relying on a discretization procedure for aggregate shocks – using tree structures and optimal quantization with a finite number of shocks – we keep the full dimensionality of the MFG system. This contrasts with standard methods that consider a finite set of moments (bounded-rationality à la Krusell-Smith) or use Projection and Perturbation (à la Reiter). We apply our method to the standard Krusell-Smith framework, while providing examples of two extensions (with endogenous labor supply or the one-asset HANK model, cf. longer version). Our method is relevant to analyze the transmission of large shocks on the wealth distribution. Uninsurable risk causes wealthier agents to accumulate precautionary savings – increasing investment and capital – while poorer Hand-to-Mouth agents expand consumption in response to aggregate shocks. Both channels matter quantitatively, implying amplification effects of aggregate fluctuations.

* *thomas.bourany@uchicago.edu*, *thomasbourany.github.io*

I would like to thank my supervisor Yves Achdou for his guidance all throughout this work, and to Pierre Cardaliaguet and Xavier Ragot for introducing me to Mean Field Games and Heterogeneous agents models. I also thank the PhD students of LJLL and LPSM at Paris-Diderot and UPMC-Sorbonne and at Sciences Po Paris: Thibaut Montes for his introduction to optimal quantization and Julien Pascal and Ziad Kobeissi for interesting discussions about MFG.
All errors are mine.

1 Introduction

One recent development of macroeconomics has been to integrate agent heterogeneity to study the income and wealth distribution of households. Since Bewley (1986) and Aiyagari (1994), this "heterogenous agents" literature has provided relevant answers on topics ranging from the causes and dynamics of inequality to the distributional consequences of monetary or fiscal policies[1]. One aspect of this literature investigates the implications of aggregate shocks for income and wealth distribution. The present paper follows this tradition by providing new methods to solve such type of model using tools from the Mean Field Games literature.

The first treatment of this topic goes back to Krusell and Smith (1998): when studying incomplete market economies with both uninsurable idiosyncratic risk and aggregate risk, the model features "approximate aggregation" in the behavior of agents: the macroeconomic aggregates are almost perfectly described using only the mean of the wealth distribution, and controls of agent are almost independent of their states. Therefore, one is naturally inclined to wonder whether agent heterogeneity substantially changes the results at stake in "representative agent" models. This paper investigate these questions.

In this project, we explore the implication of aggregate uncertainty – or "common noise" – as in Krusell-Smith. As most heterogeneous agents model in economics, this framework is an interesting example of Mean Field Games (MFG)[2]. A Mean-Field Game can be described a game with a large number of "small" symmetric players: the interaction between them is only reflected by the interaction between each agent and the "distribution" – i.e. the measure m_t – of the other agents:

$$\sup_{\alpha} \mathbb{E}\left[\int_0^\infty e^{-\rho t} f(t, X_t, m_t, \alpha_t) dt\right]$$
$$dX_t = b(t, X_t, m_t, \alpha_t)dt + \sigma(t, X_t, m_t)dB_t + \sigma^0(t, X_t, m_t)dW_t^0$$

where B_t and W_t^0 refers respectively to the idiosyncratic noise and the common noise.
This stochastic control problem boils down to a system of coupled partial differential equations:

- First, a Hamilton-Jacobi-Bellman (HJB) equation – where the agents make their choices $\alpha^\star$ taking the measure m_t of agents as given
- Second, a Fokker-Plank (FP, or Kolmogorov Forward) equation – where the measure m_t evolves, given optimal choices and the controlled process $X^{\alpha^\star}$.

The standard Aiyagari-Bewley model has recently been reformulated as an example of such Mean Field Game (MFG) by Achdou, Han, Lasry, Lions, and Moll (2017). The economic problem is the following: households solve a typical consumption-saving problem and are subject

[1]One could mention for instance Gabaix, Lasry, Lions, and Moll (2016) for dynamics of inequality, Kaplan and Violante (2014) for consumption and redistributive effects of fiscal stimulus, Heathcote, Storesletten, and Violante (2014) for the influence of incomplete insurance on risk-sharing, consumption and labor supply, Benhabib, Bisin, and Zhu (2011) on how both labor and capital idiosyncratic risk can generated Pareto wealth distribution and Bhandari, Evans, Golosov, and Sargent (2018) on optimal monetary and fiscal policy in incomplete markets

[2]A summary on how to obtain the standard MFG system is recalled in the long version of this paper. Developed simultaneously by Lasry and Lions (2007) and Huang, Malhamé, and Caines (2006), interested readers can also find a comprehensive treatment in the lecture notes Cardaliaguet (2018).

to idiosyncratic risk on their income. In an incomplete market and unable to hedge this risk, the agents use the only asset at their disposal – saving in capital – to self-insure. This capital stock is also used as an input factor by firms and 'priced' competitively at an interest rate. However, there is a limited amount households can borrow, and this 'credit constraint' forces the agents to (over)- self-insure: that what we call "precautionary saving".

Our approach is to simulate the model by solving numerically this system of two equations, using a finite-difference scheme. The two main particularities of the model are (i) a credit constraint that acts as a "state-constraint", imposing a specific treatment of the PDE on the boundaries and (ii) the mean-field interaction through the interest rate acting as a coupling between the two PDEs that requires the search for an equilibrium fixed point.

When introducing aggregate uncertainty, this model is known as a Mean Field Game with "common noise", i.e. common source of random shocks that applies to the MFG system as the whole[3]. The treatment of such common noise is as old as the Krusell and Smith (1998) article but impose challenging difficulties in mathematics.

When common noise interacts with both the behavior and the distribution of agents, solving the model suffers from what is sometimes called "curse of dimensionality" by computational economists. As explained in Carmona and Delarue (2018): "*In order to account for the dependence of the equilibria upon the realization of the common noise, it is necessary to enlarge the space in which the fixed point has to be sought.*" The control problem becomes infinite-dimensional and solving this model has appeared to be notoriously difficult.

The present master thesis offers a new numerical method to deal with such issue and to simulate the MFG system with common noise. The main idea is the discretize the aggregate shocks. Following the approach of Y. Achdou, we approximate the common noise – that can be a two-state jump process or a Brownian motion – as a finite number of shocks using a tree structure. Between each shock – i.e. on each branch of the tree – the MFG is a standard deterministic system of two evolution PDEs, the HJB and the FP coupled by the interest rate. The main challenge is to link the behavior of the MFG in function of the anticipation of future aggregate shocks and in function of the evolution of past state variables (i.e. saving). A way to include this is to (i) change the terminal condition for the HJB: the final value function is set to be the expectation of future value functions over the different paths of common noise and (ii) define the initial condition of the FP with the past wealth distribution before the realization of the common noise. With these two methods, we can compute the evolution of the MFG for different trajectories of aggregate productivity – i.e. the common noise in our setting.

This new method has several comparative advantages: it is the first method to keep track of both the full heterogeneity of agents and their anticipation of future aggregate shocks. However, the treatment of problem of MFG with common noise has a long history discussed in two separate literatures.

[3]Here we will use indifferently the term "common noise" and "aggregate risk/shocks/uncertainty"

First, the literature of computational economics have searched to reduce to dimensionality of the problem: Originally, Krusell and Smith (1998) approximated the wealth distribution using finite number of moments and solving the control problem using this finite/bounded-rational view of the economy. Later, Reiter (2010) used projection methods to approximate the measure of agents with histograms and implemented linearization and perturbation methods to simulate the economy for (necessarily small) aggregate shocks. With advances by Algan, Allais, and Den Haan (2008) and Winberry (2016) who use parametric families to approximate the measure, the combinaison of these methods have been relatively successful since the recent article by Ahn, Kaplan, Moll, Winberry, and Wolf (2018). One need to note that truncation methods have also been developed and offer parsimonious intuitions, and can be simulated using Dynare, see Challe, Matheron, Ragot, and Rubio-Ramirez (2017) or the recent Ragot (2018).

In the mathematical literature, the goal has rather been to reformulate the problem in infinite-dimension. A first approach using the *master equation* was introduced by P. L. Lions in his lectures at the *Collège de France*. This amounts to define a PDE on the state-space *times* the space of probability measures on the state and relying on differential calculus on the Wasserstein space. As such, it allows to deal with problem without or with common noise by treating the master equation of first or second order. This method, developed extensively in Cardaliaguet, Delarue, Lasry, and Lions (2017) proves to be very powerful.

Another approach by Carmona, Delarue, and Lacker (2016) is to discretize the common noise in order to treat the problem with a finite space. The idea, building around the notion of weak MFG solution can be summarized as follow: First of all, without common noise, the standard fixed-point procedure used in MFG consists in (i) fixing a measure m_t, (ii) solving the optimal control problem $\alpha_t^\star$, (iii) finding the corresponding probability law of the controlled process $X_t^{\alpha^\star}$, (iv) iterating to find a fixed point $m_t = \mathcal{L}(X_t)$. With the introduction of the common noise, the measure considered is now a measure *conditional* on the common noise $\mathcal{L}(X_t|B_t^0)$. However, conditioning introduces a major technical difficulty in the procedure above, and the discretization is a way to bypass such complication. In section 3.1, we describe briefly these approaches with a tentative to adapt these results to the Krusell-Smith model.

This project and the numerical scheme we developed relies on discretizing the common noise to approximate the solution of the stochastic MFG system, as rationalized by this last article. This method has the first and main advantage to preserve the non-linearity of the system. It is the first technique able to describe precisely the evolution of policy functions and distribution of agents along the different trajectories of the shock. This bypasses the traditional use of linearization techniques as in perturbation methods. Moreover, the second strength of our procedure is to be unambiguous on the anticipations of agents. With our discretization of the common noise, we compute effectively the different trajectories of aggregate risk and measure how agents expect and react to these future shocks. Even with a rough discretization, we have a precise view on how agents hedge against aggregate risk. This can be proved relevant in future research in economics and mathematical finance.

However, one of the main challenge of this approach is when the discretization procedure is made finer: the number of trajectories can grow exponentially – and so does the computation time. A goal for further research will be to develop novel techniques to be able to reach a greater accuracy of the MFG with common noise. In the long version of this paper, we offer novel ideas to tackle this problem by combining recent techniques of the mathematical literature.

To prove how this method can be relevant for MFG and heterogeneous agents models (H.A.) in economics, we first describe our algorithm in section 4 and apply this method to the main Krusell-Smith framework in section 5. When solving the model, we study the importance of aggregate shocks on the shape of the wealth distribution and how precautionary saving can change the reaction of households and the dynamics of aggregate variables.

In a brief description below and in the long version of this paper, we provide extensions when labor supply in endogenous or prices are rigid. These H.A. models – analogous of the RBC and the NK models – display even greater amplification channels of productivity shocks.

2 Model: the Aiyagari-Bewley-Huggett framework as a MFG

In one of its simplest formulation, the model shows the economy, composed of a continuum of households. They face idiosyncratic and uninsurable income shocks, and are subject to a credit (borrowing) constraint. They thus solve the following stochastic control problem:

$$\sup_{c_t} \mathbb{E}_0 \int_0^\infty e^{-\rho t} u(c_t) dt$$

$$\text{subject to :} \qquad da_t = (z_t w_t + r_t a_t - c_t)dt \qquad \text{(Budget constraint)}$$

$$\text{and} \qquad a_t \geq \underline{a} \qquad \text{(Credit constraint)}$$

where a is wealth – the state variable we'll focus on – c_t consumption, ρ rate of time preference, $u(\cdot)$ a utility function, supposed increasing and concave ($u' > 0, u'' < 0$). For the simulation, the utility will be a Constant Relative Risk Aversion (CRRA) utility, i.e. $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$.

The income is composed of a wage w_t and productivity factor z_t. The wage w_t, and interest rate r_t, will adjust in general equilibrium, considering the firm's side (cf. next subsection). The productivity z_t is subject to shocks that are (i) idiosyncratic and (ii) uninsurable (agents can not hedge/cover against this risk). z intuitively represents the state of the agent, for instance, employment $z_{high} = z_2$ and unemployment $z_{low} = z_1 = \frac{1}{2} z_2$. In this presentation, it is modeled simply as a jump processes with two states $\{z_1, z_2\}$ (with intensities λ_1, λ_2, the higher the intensity, the higher the proba to jump). However, one can generalize it any stochastic process: e.g. diffusion $dz_t = b(z_t)dt + \sigma(z_t)dB_t$, Poisson or Levy processes. The only numerical constraint is that it has a bounded domain[4].

One of the specificities of this problem is the credit constraint, which is a *state* constraints: $a_t \geq \underline{a}$. State constraint is a complicated problem for control theory. In our situation, intuitively,

[4]In case of diffusion, it will be reflected after a limit

the optimal strategy of the agent might be (and will be) to move on the constraint ($\partial\Omega$) and stay there (poverty trap). Mathematically, it is not possible to find a PDE and a boundary condition on $\partial\Omega$ even in the sense of distribution.

According to Soner (1986), the mapping from the state to the set of admissible controls $a \mapsto \mathcal{A}_a$(and $c \in \mathcal{A}_a$) "will have a complicated structure", and its regularity may not be insured in general. In our situation, it will be the case and we will implicitly impose a constraint on the derivative of the value function at the boundary.

This will result on both (i) a Dirac mass on the boundary and (ii) an explosion near the boundary. Economically, there is a lower bound on the value of this borrowing limit: $\underline{a} \geq -z_1 w_t / r_t$. The latter term represents the natural borrowing limit – the credit that an agent could repay if it would fall permanently in the low income state (and repaying its debt at rate r_t.

The supply side – a neoclassical firm

The supply side is driven by a representative firm, produce goods following a production function of aggregate productivity, capital and labor $F(A, K, L)$, supposed to be concave in the two last variables. Given δ the depreciation of capital and A productivity level, the firm is producing in a perfectly competitive environment, and thus maximize its profit:

$$\Pi(K_t, L_t) = \sup_{K,L} \; F(A_t, K_t, L_t) - (r_t + \delta)K_t - w_t L_t$$

The price of demanded input factors – i.e. r_t interest rate, w_t wage – will be determined by the First-Order-Condition of this optimization problem:

$$r_t = \partial_K F(A_t, K_t, L_t) - \delta \qquad\qquad w_t = \partial_L F(A_t, K_t, L_t)$$

In this article, we suppose that (i) production function is simply given by the Cobb-Douglas production function: $F(A_t, K, L) = A_t \, K_t^{\alpha} \, L_t^{1-\alpha}$, and (ii) the effective labor supply is fixed and equal to the average productivity of households: $z_{av} = \frac{z_1\lambda_2 + z_2\lambda_1}{\lambda_1+\lambda_2}$.

The optimality relation described above reduce to :

$$\begin{aligned}
K(r) &:= \left(\frac{\alpha A}{r+\delta}\right)^{\frac{1}{1-\alpha}} z_{av} && \text{Capital demand} \\
w &= (1-\alpha)\, A\, K^{\alpha} z_{av}^{-\alpha} && \text{Wage} \\
r &= \alpha\, A\, K^{\alpha-1} z_{av}^{1-\alpha} - \delta && \text{Interest rate}
\end{aligned}$$

Wage and interest rate are indeed constant when productivity A and all parameters are constant. Let us look at the Mean Field Game formulation of the problem.

A mean field game formulation

This stochastic control problem boils down to a system of two partial differential equation: (i) a Hamilton-Jacobi-Bellman (HJB) equation and (ii) a Fokker-Plank (or Kolmogorov Forward – FP). When parameters are constant, the stationary equilibrium system is described by the equations:

$$\rho v_j(a) = \max_c u(c) + \partial_a v_j(a)(z_j w + ra - c) + \lambda_j(v_{-j}(a) - v_j(a)) \qquad \text{[HJB]}$$

$$0 = \frac{d}{da}[s_j(a)\, g_j(a)] + \lambda_j g_j(a) - \lambda_{-j} g_{-j}(a) \qquad \text{[FP]}$$

$$S(r) := \int_a^\infty a\, g_1(a) da + \int_a^\infty a\, g_2(a) da = K(r) \qquad \text{[Market clearing]}$$

where j indicate the z-state of the agent and $-j$ the opposite state.

However, when the system is subject to shocks – under either (i) perfect foresight, i.e. deterministic transitions, (ii) MIT shocks i.e. unanticipated (zero probability) shocks, the system dynamics is described by the two following PDEs:

$$\rho\, v_j(t,a) = \partial_t v_j(t,a) + \max_c u(c) + \partial_a v_j(t,a)\, s_j(t,a) + \lambda_j(v_{-j}(t,a) - v_j(t,a)) \qquad \text{[HJB]}$$

$$0 = \partial_t g_j(t,a) + \frac{d}{da}[s_j(t,a)\, g_j(t,a)] + \lambda_j g_j(t,a) - \lambda_{-j} g_{-j}(t,a) \qquad \text{[FP]}$$

$$S(t,r) := \int_{\underline{a}}^\infty a\, g_1(t,a) da + \int_{\underline{a}}^\infty a\, g_2(t,a) da = K(t,r) \qquad \text{[Market clearing]}$$

$$s_j(t,a) = z_j w_t + r_t\, a - c_j(t,a) \qquad c_j(t,a) = (u')^{-1}(\partial_a v_j(t,a)) \qquad \text{[FOC]}$$

$$v_j(T,\cdot) = v^\infty \qquad g_j(0,\cdot) = g^0 \qquad \partial_a v_j(\underline{a}) \geq u'(z^j w + r\underline{a}) \qquad \text{[Boundary conditions]}$$

In these two cases, the third equations describe the market clearing, i.e. the Walrasian adjustment of prices (i.e. interest rate) to equalize supply of saving $S(r)$ and demand for capital $K(r)$. The last equation represents the optimality condition of the control variable c. The HJB actually features an optimization problem, and, since the objective $u(c) - pc$ is concave in c, the optimum is reached for: $u'(c^\star) = p$ and thus $c^\star = (u')^{-1}(p)$.

Here the state-constraint does not show up in the HJB (as could be the case in discrete time) but it appears in the boundary condition. It comes from the optimality of the maximization in the HJB – holding everywhere, and in particular at $\underline{a}$ – and the FOC is given by $u'(c^j(\underline{a})) = \partial_a v^j(\underline{a})$, and the state-constraint affecting the control: $s^j(a) = z_t^j w + r\, a - c^j(a) \geq 0$. Since u is concave, its derivative is decreasing, yielding the boundary condition:
$u'(z^j w + r\underline{a})) \leq u'(c^j(\underline{a})) = \partial_a v^j(\underline{a})$.

3 Introducing aggregate shocks

We now consider the same framework, when introducing the aggregate shock, or common noise – to use to mathematics idiom. Here, the level of productivity A_t (TFP: total factor productivity) is subject to random fluctuations, and this affects firms in their production process, and thus prices of the factors (interest rate and wages) and in turn the behavior of households.

The form of the stochastic process for aggregate productivity will matter a lot for the anticipations of agents (who are perfectly forward looking and thus draw expectations). The simplest way would be to consider a simple Brownian motion dB_t starting at a initial level $\bar{A}_0$ and stopped at a deterministic stopping time T (assuming to be the stationary equilibrium). This will be the situation considered in the following mathematical analysis and the first simulation.

Moreover, to study the transmission mechanisms after a transitory shock analogous to an $AR(1)$ process. The closest process in continuous time would be a "Generalized Ornstein Uhlenbeck" or "Jump-Drift Process": $dX_t = -\theta(X_t - \mu)dt + \varepsilon dN_t$ where dN_t is a jump process ($dN = 1$ with intensity λ) and $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. This latest is the one we consider for the plot of the Impulse Response Functions (IRF) in the results in section 5.

3.1 General framework

To provide the general framework with common noise, we can consider the case where the evolution of idiosyncratic state z_t and common state A_t are diffusion processes. We denote m_t the measure of the states (a_t, z_t, A_t) and $\langle m_t, a\rangle$ the first moment of the measure w.r.t the states a_t.

$$\sup_{c_t} \mathbb{E}_0 \int_0^\infty e^{-\rho t} u(c_t) dt \tag{1}$$

$$\text{s.t.} \quad da_t = \Big(z_t \underbrace{\left(A_t\, \kappa_w\, \langle m_t, a\rangle^\alpha\right)}_{\equiv w_t} + \underbrace{\left(A_t\, \kappa_r \langle m_t, a\rangle^{\alpha-1} - \delta\right)}_{\equiv r_t} a_t - c_t\Big) dt \qquad \& \qquad a_t \geq \underline{a} \tag{wealth}$$

$$dz_t = b^{idio}(z_t)dt + \sqrt{2\sigma^{idio}} dB_t \tag{labor productivity}$$

$$dA_t = b^{com}(A_t)dt + \sqrt{2\sigma^{com}} dW_t^0 \tag{aggreg. productivity}$$

where we normalized the labor supply $z_{av} = \langle m_t, z\rangle$ to a constant and defined $\kappa_w = (1-\alpha) z_{av}^{-\alpha}$ and $\kappa_r = \alpha z_{av}^{1-\alpha}$ to alleviate a bit the notations. Consumption c_t is the control of the agent.

We can rewrite our framework, using the standard notations from stochastic control/MFG[5] Therefore, the dynamics of the state is the following:

$$\begin{aligned} dX_t &= b(X_t, m_t, \alpha_t^\star)dt + \widetilde{\sigma}^{idio} dB_t + \widetilde{\sigma}^{com} dW_t^0 \\ dX_t &= -D_p \mathcal{H}\big(t, X_t, m_t, D_x v_t(X_t)\big)dt + \widetilde{\sigma}^{idio} dB_t + \widetilde{\sigma}^{com} dW_t^0 \end{aligned} \tag{2}$$

where v_t is the value function of the control problem and $\mathcal{H}$ the hamiltonian given below and B_t and W_t^0 respectively the idiosyncratic and common noise – both are 1-dimensional Wiener processes (Brownian motion), due to the value of $\widetilde{\sigma}^{idio}$ and $\widetilde{\sigma}^{com}$. This formalism is closer to the usual notation of the MFG literature.

By posing the Hamiltonian $\mathcal{H}(t, x, m, p) = \max_\alpha \big(u(\alpha) + b(x, m, \alpha) \cdot p\big)$ we see that we are in a "strongly coupled" Mean Field Game problem: the coupling between agents states and the measure of other players depends on their strategies, through the prices (interest rate and wages). For results on strongly coupled MFG, see Bertucci et al. (2018).

[5] With $X_t \equiv (a_t, z_t, A_t)$ the state is in $\mathbb{R}^3$, the optimal control $\alpha^\star(t, x, m, p) \equiv c_t^\star$ and we use the vector formulation:

$$b(X_t, m_t, \alpha_t^\star) \equiv \begin{pmatrix} z_t\, A_t\, \kappa_w\, \langle m_t, a\rangle^\alpha + (A_t\, \kappa_r \langle m_t, a\rangle^{\alpha-1} - \delta)\, a_t - c_t^\star \\ b^{idio}(z_t) \\ b^{com}(A_t) \end{pmatrix}, \qquad \begin{aligned} \widetilde{\sigma}^{idio} &= (0, \sqrt{2\sigma^{idio}}, 0)^T \\ \widetilde{\sigma}^{com} &= (0, 0, \sqrt{2\sigma^{com}})^T \end{aligned}$$

Finite vs. infinite-horizon. Economics models are often formalized in infinite horizon, discounting the future at a rate ρ. Considering first a stationary equilibrium – a "steady-state" – usual methods analyze transmission effects of a one-time shock – an "unexpected and transitory deviation from steady state" – in a given number of periods (in discrete time). In the following, we adopt an intermediary setting: for a finite time period $[0,T]$, the system is subject to aggregate shocks, and thereafter, when $t > T$, the MFG system returns to the case without aggregate uncertainty – the aggregate uncertainty is "switched off". As a result, the main difference is to consider multiple trajectories of aggregate noise: there are thus many different "steady-states". The advantage is also to consider the MFG with common noise in finite-time horizon, a setting more familiar for applied mathematicians.

Therefore, the MFG without common noise, when $\widetilde{\sigma}^{com} = \underline{0}$, is characterized by the usual system of PDE:

$$\begin{cases} -\partial_t v(t,x) + \rho\, v(t,x) = \mathcal{H}\big(t,x,m,D_x v\big) + \sigma^{idio}\partial^2_{x_2} v(t,x) \qquad \text{on } [0,T]\times\mathbb{X} \\ \partial_t m(t,x) - \operatorname{div}\Big[D_p\mathcal{H}\big(t,x,m,D_x v\big)\, m(t,x)\Big] - \sigma^{idio}\partial^2_{x_2} m(t,x) = 0 \qquad \text{on } [0,T]\times\mathbb{X} \\ v(T,\cdot) = v^\infty \qquad m(0,\cdot) = m^0 \qquad \text{on } \mathbb{X} \end{cases} \tag{3}$$

where $\mathbb{X} := [\underline{a},\infty)\times\mathbb{R}\times[0,\infty)$, and v^∞ is the value function for the stationary control problem. The Hamiltonian is given above by $\mathcal{H}(t,x,m,p) = \max_\alpha \big(u(\alpha) + b(x,m,\alpha)\cdot p\big)$. Note that we adopt this notation since, in our special setting, $\sigma^{idio}\partial^2_{x_2} v(t,x) = \frac{1}{2}\operatorname{tr}\big(\widetilde{\sigma}\widetilde{\sigma}'^{idio} D^2_x v(t,x)\big)$ and $\sigma^{idio}\partial^2_{x_2} m(t,x) = \frac{1}{2} D^2_x\big(\widetilde{\sigma}\widetilde{\sigma}'^{idio} m(t,x)\big)$. We emphasize the fact that the MFG equilibria is deterministic is this context. Let us figure out what would be the setting when a transitory – finite-horizon – aggregate shock occurs.

3.2 MFG with common noise – Different mathematical approaches

When $\widetilde{\sigma}^{com} > \underline{0}$ the effect of the common noise is to randomize the MFG system. Therefore, the measure m_t becomes a random flow of measures, and can be now considered as a flow of conditional marginal measures of $(X_t)_{t\geq 0}$ given the realization of the common noise W^0_t, i.e. formally $m_t = \mathcal{L}(X_t|W^0)$. Many different approaches exist in the mathematical literature to handle such framework.

Given the forward dynamics of the state eq. (2), solving the Mean Field Game usually consists in four steps:

1. Fix an arbitrary measure m_t, that should now be a stochastic process (over $\mathcal{P}(\mathbb{X})$) adapted to the filtration $\mathcal{F}^0$ generated by the common noise W^0.
2. Solve the standard stochastic control problem eq. (1) with random coefficient, subject to the dynamics eq. (2)
3. When an optimal control exists in feedback form $\alpha^\star(t,x,m)$, plug it into the Fokker Planck to obtain the evolution of m'_t.
4. Find a fixed point where this solution m' is precisely the m we started from. It reduces to a search for a flow of random measure such that $m_t = \mathcal{L}(X_t\,|\mathcal{F}^0)$

In presence of common noise, the difficulty is twofold. First, the fixed point problem is performed in an infinite dimensional space $[\mathcal{C}([0,T],\mathcal{P}(\mathbb{X}))]^{\Omega}$ where Ω denotes the underlying probability space carrying the common noise. This space being too large, the use of compactness arguments may fail, preventing the proof of existence by mean of Schauder's Theorem – the standard method to solve the case without common noise. Second, the control problem should be solved in the space of stochastic processes, i.e. adapted to $\mathcal{F}^0$. Let us see now how to deal with such issue.

First, as introduced above, this MFG system with common noise could be view as a system of *Stochastic Partial Differential Equation*. Following the approach of Cardaliaguet, Delarue, Lasry, and Lions (2017) and Carmona and Delarue (2014), we can the model as a system of Forward-Backward SDE or Stochastic PDE :

$$\begin{cases} dv_t &= \Big[\rho\, v_t - \mathcal{H}\big(t,x,m,D_x v_t\big) - \sigma^{idio}\partial^2_{x_2} v_t - \sigma^{com}\partial^2_{x_3} v_t - \sqrt{2\sigma^{com}}\partial_{x_3} Z^0_t\Big]\,dt + \, Z^0_t dW^0_t \qquad \text{on } \mathbb{X} \quad \& \; t\in[0,T] \\ dm_t &= \Big[\operatorname{div}\big(D_p\mathcal{H}\big(t,x,m,D_x v\big)\,m_t\big) + \sigma^{idio}\partial^2_{x_2} m_t + \sigma^{com}\partial^2_{x_3} m_t\Big]dt - \sqrt{2\sigma^{com}}\,\partial_{x_3} m_t\, dW^0_t \qquad \text{on } \mathbb{X} \quad \& \; t\in[0,T] \\ v_T &= v^{\infty} \qquad\qquad m_0 = m^0 \qquad\qquad \text{on } \mathbb{X} \end{cases} \tag{4}$$

where (v_t, Z^0_t) is solution of the first Backward Stochastic Differential equation (BSDE) and Z^0 the (1-dim) random field that allows the solution to be adapted to the filtration $\mathcal{F}^0$ generated by the common noise W^0. Said differently, it could be understood as the "effect of common uncertainty" on value and behavior of agents.

Second, another solution covered in the recent literature of Mean Field Game consists in rewriting the problem in infinite-dimension using the *Master equation*, as described in Cardaliaguet, Delarue, Lasry, and Lions (2017). Set over the state-space $\mathbb{X}$ times the space of probability measure $\mathcal{P}(\mathbb{X})$, this equation describe the Nash equilibrium in a single equation of the value $U(t,x,m)$ and constitutes the appropriate setting to understand MFG with common noise. The MFG system of SPDE can be viewed as a characteristic of the master equation and in our framework, the latter becomes:

$$\begin{cases} -\partial_t U + \rho\, U - \mathcal{H}(\cdot, D_x U) - \sigma^{idio}\partial^2_{x_2} U - \sigma^{com}\partial^2_{x_3} U - \displaystyle\int_{\mathbb{X}} \sigma^{idio}\partial_{y_2}\partial_{m_2} U(\cdot;y) + \sigma^{com}\partial_{y_3}\partial_{m_3} U(\cdot;y)\, m(dy) \\ \qquad + \displaystyle\int_{\mathbb{X}} \partial_{m_1} U(\cdot;y)\,\partial_{p_1}\mathcal{H}\big(\cdot,y,\cdot,D_x U(\cdot,y,\cdot)\big) m(dy) - 2\sigma^{com}\int_{\mathbb{X}} \partial_{x_3} D_{m_3} U(\cdot,x,\cdot;y) m(dy) \\ \qquad -\sigma^{com}\displaystyle\int_{\mathbb{X}\times\mathbb{X}} \partial^2_{m_3} U(\cdot,x,\cdot;y,y')\, m_t^{\otimes 2}(dy,dy') = 0 \qquad \text{in } [0,T]\times\mathbb{X}\times\mathcal{P}(\mathbb{X}) \\ U(t,x,m) = v^{\infty}(x,m) \qquad\qquad \text{in } \mathbb{X}\times\mathcal{P}(\mathbb{X}) \end{cases} \tag{5}$$

where U and $\mathcal{H}$ are taken in (t,x,m) and $U(\cdot,x,\cdot;y) = U(t,x,m;y)$ where there could be ambiguities. Note that to transform the MFG system of SPDE into the master equation, we identified the random field Z^0_t with :

$$Z^0_t(x) = \int_{\mathbb{X}} D_m U(t,x,m;y) m_t(dy)$$

A complete description of this master equation is dedicated in the longer version of this paper. Note that it is analogous to the infinite-dimensional HJB expressed in equation (43) of the Appendix A.1. of Ahn, Kaplan, Moll, Winberry, and Wolf (2018). It includes both local (the five first) terms and non-local (the five last) terms. The latter are specific to MFG and include how agents expect the evolution of the distribution of the different states (wealth x_1, idiosyncratic x_2 or aggregate x_3 productivity), with aggregate uncertainty.

Third, a last approach of the mathematical literature to handle common noise is to study *weak equilibria of the MFG*. The strategy developed in Carmona, Delarue, and Lacker (2016) relies on pure probabilistic arguments and a procedure to discretize the common noise. Focusing on the case where the common noise has a finite number of outcomes, one can use the "usual" procedure for MFG existence, this time conditioning on a finitely-supported approximation of the aggregate noise. The existence when refining the discretization is thus only proved as a *weak equilibrium*, exactly as weak solutions for stochastic differential equations, where the measure m_t of the controlled process may not be adapted to the aggregate shocks. However, under the assumption of pathwise uniqueness, the authors can prove an analogous version of the Yamada-Watanabe theorem for SDE: any weakly existing equilibria that satisfy pathwise uniqueness – i.e. indistinguishable trajectories – is in fact a strong solution of the MFG.

This approach could be seen as a rationale for the numerical scheme we develop in the next section.

4 The solution method: a discretization procedure

The main idea of our method is to approximate the stochastic process of common noise. In this setting, the only state affected by this aggregate shock is the productivity level A_t. We use a tree structure to approximate this process in time with a finite number M of simple shocks, and in space with a finite number K of possible states. With this method, one can approximate any process using a finite number K^M of trajectories. When we consider a simple Brownian Motion dW^0 or other diffusion processes, we can use different tree structures: (i) binomial, trinomial or "K-nomial" trees and (ii) optimal quantization trees.

Building trees

More precisely, in our case, the process for productivity A_t starts at $\bar{A}_0$. We first discretize the time: $[0,T]$ is divided $M+1$ periods $[t_m, t_{m+1}]$ with $m \in \{1,\dots,M\}$, uniformly every $\Delta T = \frac{T}{M+1}$. At each date t_m, the process A_t switches between K deterministic trajectories $(A^k_t)_{t\in[t_m,t_{m+1}]}$ with $k \in \{1,\dots,K\}$. The probability of transition from one state A_{t_m} to one of these K trajectories is $\pi_{k_{m+1}|k_m} := \mathbb{P}(A_t = A^k_t\,,\ \forall t \in [t_m, t_{m+1}]\,\big|A_{t_m})$

This way, we can "build" a *tree* of different trajectories of common noise, and the discretized process is still stochastic, with the probability of transition given by the coefficient of the hypermatrix $\pi_{k_{m+1}|k_m}$ – whose coefficients should be carefully chosen as we will see. When taking $\Delta T \to 0$, you can approximate any process.

Solving the MFG system – grafting branches

A way to 'solve' the MFG with common noise is to compute the evolution of the MFG system – HJB and FP equations – on *each branch*, i.e. on each deterministic trajectories $(A^k_t)_{t\in[t_m,t_{m+1}]}$ For a given trajectory, the value function at time t_m depends on the state A_{t_m} and writes:

$$
\begin{aligned}
v(t_m, x, A_{t_m}) &:= \sup_{c_t} \mathbb{E}\Big[\int_{t_m}^{T} e^{-\rho t} u(c_t)dt \;\Big|\; A_{t_m}\Big] \\
&= \sup_{c_t} \mathbb{E}\Big[\int_{t_m}^{t_{m+1}} e^{-\rho t} u(c_t)dt + e^{-\rho t_{m+1}} v(t_{m+1}, X_{t_{m+1}}, A_{t_{m+1}}) \;\Big|\; A_{t_m}\Big]
\end{aligned}
$$

by applying standard dynamic programming arguments. Therefore, by denoting t_m^- the time before revelation of the shock, and t_m^+ when the future trajectory is revealed, the HJB equation with random coefficients $A^{k_m}_t$ becomes:

$$
\begin{cases}
-\partial_t v(t,x,A^{k_m}_t) + \rho\, v(t,x,A^{k_m}_t) = \mathcal{H}\big(t,x,m,D_x v\big) + \sigma^{idio}\partial^2_{x_2} v(t,x,A^{k_m}_t) & \text{on } [t_m,t_{m+1}]\times\mathbb{X} \\
v(t^-_{m+1}, x, A^{k_m}_t) = \mathbb{E}\Big[v(t^+_{m+1}, x, A_{t_{m+1}}) \;\Big|\; A^{k_m}_{t_m}\Big] & \text{on } \mathbb{X}
\end{cases}
$$

Therefore, the dynamics of A_t matter a lot on how the agents consider this terminal condition, which can thus be written, when the common noise is discretized, as follow:

$$
v(t^-_{m+1}, x, A^{k_m}) = \sum_{k=1}^{K} \pi_{k|k_m}\, v(t^+_{m+1}, x, A^k_{t_{m+1}})
$$

In particular, since agents are forward looking (and rational!), they form expectations over the different set of future branches, and that will be accounted in the value function. The value function "jumps" between t^-_m and t^+_m to correct the (past) expectations after the shock of information. Also, the presence of this conditional expectation $\mathbb{E}[\cdot \mid A_{t_m}]$ allows the value function to be adapted to the filtration generated by the common noise A_t.

Note that, at the last period $[t_M, T]$, there is no more uncertainty and the terminal condition becomes:

$$
v(T, \cdot, A^{k_M}_T) = v^{\infty, k_M}
$$

where v^{∞,k_M} is the stationary value function of one of the K^M terminal equilibria. To compute these value function, the algorithm will therefore the HJBs backward, starting from the terminal condition and correcting the value function at each node by taking the conditional expectation w.r.t. the common noise.

Consider now the associated Fokker Planck equation, on each interval $[t_m, t_{m+1}]$, when the controlled drift is given by $D_p\mathcal{H}$. Again, the measure depends on the random coefficient A^{k_m} through the controls:

$$
\begin{cases}
\partial_t m(t,x,A^{k_m}) - \operatorname{div}\Big[D_p\mathcal{H}\big(t,x,m,D_x v(\cdot,A^{k_m})\big)\, m(t,x,A^{k_m})\Big] - \sigma^{idio}\partial^2_{x_2} m(t,x,A^{k_m}) = 0 & \text{on } [t_m,t_{m+1}]\times\mathbb{X} \\
m(t^+_m, \cdot, A^{k_m}_{t^+_m}) = m(t^-_m, \cdot, A^{k_{m-1}}_{t^-_m}) & \text{on } \mathbb{X}
\end{cases}
$$

where $m(t_1^+, \cdot) = m^0$ for the first period starting at $t_1 := 0$. However, in contrast of the value function, the measure is backward-looking and naturally accounts for the past drift values. It is therefore continuous in time t. To compute the measure, the F.P. equations will be solved forward, starting from the initial condition for each trajectory of the common noise.

4.1 Binomial/trinomial tree structure

When we consider a simple Brownian Motion dW^0 or other diffusion processes[6], an obvious approximation would be to consider a simple random walk, a process that rises or drops with probability $1/2$. Moreover, one could also think about the "normal"-random-walks:

$$S_n = \sum_k^n \varepsilon_k \qquad \text{with} \;\; \varepsilon_k \sim \mathcal{N}(0,1) \qquad \text{and} \;\; W_t^{(n)} := \frac{S_{\lfloor nt \rfloor}}{\sqrt{n}}$$

The Donsker's theorem provides a strong justification of this approximation: the process $W_t^{(n)}$ as a random variable on the Skorokhod space $\mathcal{D}([0,T])$ will converge in law toward a standard Brownian Motion, when time increment goes to zero, i.e. when $n \to \infty$.

Parametrization

How to choose the trajectories $(A_t^{k_m})_{t \in [t_m, t_{m+1}]}$ and the transition probability $\pi_{k_{m+1}|k_m}$? Here, we take advantage of the fact that the increment of the random walk follows a normal distribution. Hence, one should refer to the "Gaussian case" in the literature on optimal quantization, cf. Pagès and Printems (2003) and Pagès (2017).

Concretely, thanks to the database on the website http://quantize.maths-fi.com, the procedure is relatively straighforward. To approximate the distribution of the aggregate shock, one should choose a Voronoï quantization grid $\Gamma = \{a_1, \dots, a_K\}$ in order to minimize the L^2-mean quantization error $e(A_t, \Gamma) = \big|\big| \min_{\{a_t\} \in \Gamma} |A_t - a_t| \big|\big|_{L^2}$ between the random variable A_t and the nearest neighbor[7] a_t in Γ. The problem consists thus in the minimization over all grids Γ of size K:

$$e_{K,L^2}(A_t) := \inf \left\{ e_{L^2}(A_t, \Gamma) \;\middle|\; \Gamma \subset \mathbb{R}, |\Gamma| \leq K \right\}$$

Now, considering this grid Γ, with the Voronoï partition $\big(C_k(\Gamma)\big)_{1 \leq k \leq K}$ such that

$$C_k(\Gamma) \subset \left\{ z \in \mathbb{R} : |z - a_k| \leq \min_{1 \leq j \leq K} |z - a_j| \right\}$$

we compute the probability weight p_k by measuring the probability of the random variable falling into this particular Voronoi cell

$$p_k := \mathbb{P}\big(A_t \in C_k(\Gamma)\big)$$

Moreover, (a) the Brownian motion has stationary increment, one could apply the same method

[6]If one would approximate diffusions, e.g. Geometric Brownian motion or Ornstein-Uhlenbeck processes with bi/trinomial trees, one could also refer to the simulation by T. Montes, cf. http://simulations.lpsm.paris/trinomial_trees/

[7]Therefore, this approach looks similar to clustering methods in machine leaning

for every node and every level $1 \leq m \leq M$ of the tree and (b) the increments being independent, the probability of transition is simply $\pi_{k_{m+1}|k_m} = p_{k_{m+1}}$.

In the following, we will use several concrete examples:

(i) A simple tree, where $K = 3$ branches per nodes and $M = 3$ "waves" of shocks,

(ii) A "long" tree, where $K = 2$ and $M = 9$, where the time discretization is finer

(iii) a "large tree", where $K = 7$ and $M = 3$ the (space) quantization grid is finer

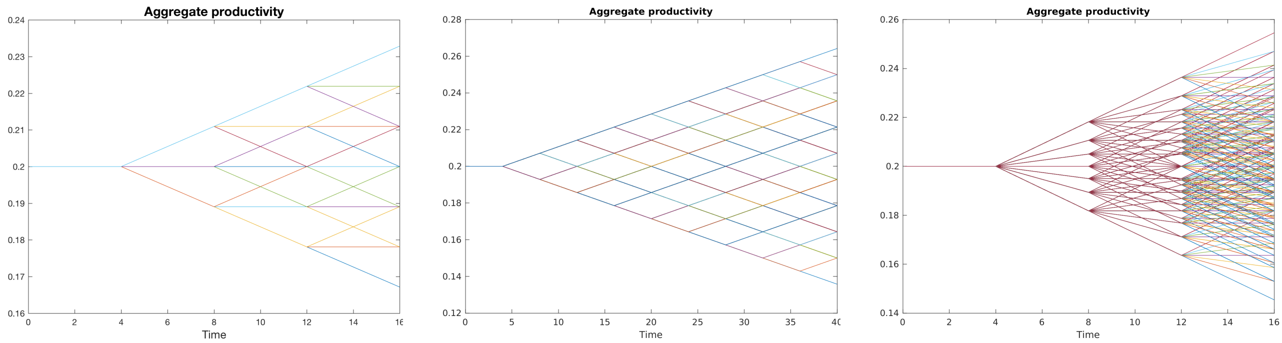

4.2 Quantization tree

Previously, we used optimal quantization for every increment of the tree, which were gaussian. We could also consider quantization trees for the *whole* process. The main idea is approximate the process using a skeleton of the distribution, supported by a tree whose branches are not identical anymore. The description of this approach – used a lot in Math-Finance – is provided in the longer version of this paper[8].

4.3 Algorithm – solving the MFG on each branch of the tree

The objective is to find equilibrium of the MFG i.e. the value function v_j and the measure $g_j \equiv m_t(\cdot, z_j)$ $(j = 1, 2)$ solving the two PDEs and the interest rate r clearing the market.

A complete description can be found in the long version of the paper. We summarize here the general method to find the equilibrium fixed point, iterating on ℓ:

1. Guess interest rate r^ℓ, compute capital demand $K(r^\ell)$ & wages $w(K)$
2. Solve the HJB using finite differences (semi-implicit method): obtain the controlled drift $s_j^\ell(a)$ and then the value function v_j^ℓ, by solving a system of sort:

$$-\frac{\mathbf{v}^{n+1} - \mathbf{v}^n}{\tau} + \rho\, \mathbf{v}^{n+1} = \mathbf{u}(\mathbf{v}^n) + \mathbf{A}(\mathbf{v}^n; r)\mathbf{v}^{n+1}$$

   - In the stationary equilibrium, solve for v_∞. In the transition case, compute the path of v^n, starting from the *terminal* condition $v^N = v_\infty$ and iterating *backward.*

3. Using $\mathbf{A}^T$, solve the FP equation, via the finite diff. system:
$\frac{\mathbf{g}^{n+1} - \mathbf{g}^n}{\tau} + \mathbf{A}(\mathbf{v}^n; r)^T \mathbf{g}^{n+1} = 0$, and obtain g_j

   - In the stationary equilibrium, solve for g_∞. In the transition case, compute the path of g^n, starting from the *initial* condition $g^0 = g_0$ and iterating *forward.*

[8]One could refer to Bally, Pagès, and Printems (2005) for a more general approach

4. Compute the capital supply $S(\mathbf{g}, r) = \int_{\underline{a}}^{\infty} a\, g_1(a) da + \int_{\underline{a}}^{\infty} a\, g_2(a) da$
5. If $S(r) > K(r)$, decrease $r^{\ell+1}$, updating (using bisection method, and conversely, and come back to step 2.
6. Stop if $S(r) \approx K(r)$

Duality: Solving the FP using the HJB

In this problem, the FP is the adjoint equation of the HJB equation. As a result, the operator matrix in the Fokker Plank equation is the transpose (A^T) of the HJB operator matrix A – *taking optimal consumption as given.* Therefore, the most important would be to solve the HJB equation numerically and then to solve FP by "simply" transposing the matrix A.

Common noise

The algorithm is rationalized and explained above: it consists in choosing an appropriate tree structure to approximate the common noise A_t. On this tree representing K^M trajectories – M waves of shocks and K branches for each shock – one would solve the time-varying MFG equilibria, using the right boundary conditions:

- A terminal condition for the HJB: $\mathbf{v}^{k_N} = \sum_k \pi_{k_{N+1}|k_N} \mathbf{v}^{k_{N+1}}$
- An initial condition for the FP: $g^0 = g_0$ and solve the FP on the full trajectory on $[0, T]$

The equilibrium on one branch of the tree, with a deterministic trajectory of A_t, is exactly the one described above.

4.4 Relative advantages and computational challenge

As observed above, the MFG equilibrium strongly depends on past values of aggregate shocks: two equilibria with the same value A_{t_m} differ if the past values $A_{t_{m-1}}$ were different. This past value indeed affects the evolution of the measure – which is backward looking and thus path-dependent. This is a concrete sign that MFG equilibria are not Markovian in $x \in \mathbb{X}$. However, it will Markovian when enlarging the space to include the measure, i.e. $\mathbb{X} \times \mathcal{P}(\mathbb{X})$. This observation justify the use of the Master equation to describe MFG equilibria.

The main drawback of this method is the computational difficulty of simulating the stochastic system of forward-backward PDEs. Indeed, the tree structure complexity increases exponentially due to the size K^M of the discretization grid. This has two effects: (i) it increases the number of trajectories K^M one has to simulate, the number of both HJB and FP equations to solve, and thus the time of computation, and (ii) it increases the size of the data, since the value function and the measure are array in four dimensions: $[0, T] \times \mathbb{X}$ and Ω the space supporting the underling common noise. This memory therefore increases exponentially when the discretization is made finer.

These two difficulties can be tackled by different arrangements for example by parallelizing the algorithm – as the tree structure would allow to compute the trajectories on different cores.

5 Results – Precautionary saving along the business cycle

In this section, we will cover the result of our simulation for the standard Krusell-Smith framework as well as a brief note on the extensions with endogenous labor supply (à la RBC) or nominal rigidities (à la New Keynesian, i.e. HANK)
We provide results of the model with aggregate uncertainty in TFP A_t:

- with a Brownian common noise $A_t = \bar{A}_0 + B_t$.
- with a common noise following a Generalized Ornstein-Ulhenbeck, or "Jump-Drift" process : $d\widetilde{A_t} = -\theta(\widetilde{A_t} - \widetilde{A_0})dt + \varepsilon dN_t$ where dN_t is a jump process ($dN = 1$ with intensity λ) and $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. We consider such process to provide economic intuitions, as it allows to plot the Impulse Response Functions (IRF) after a one-time temporary deviation from steady-state.
- we then compare the previous results with standard IRF of the representative agent model (i.e. single player in the game), simulated using Dynare

5.1 Discretized Brownian common noise

We presents the simulation of the MFG system with common noise of the "simple tree" (i) in the following, to provide intuitions and economic results. The trees (ii)-(v) described in section 4 are displayed in the long version of this paper.

The evolution of aggregate variables
How agent heterogeneity interact with aggregate dynamics in presence of aggregate shocks? In the following graph we plot the evolution of the main variables – capital supply, consumption and saving, and prices, interest rate and wages – along each branch of the tree of productivity.

Note that the colors of the trajectories are linked between the different graphs: for instance the sky blue line corresponds to the "best scenario" of productivity on all the plots.

As a result of the continuity of the wealth distribution g in time, the capital stock adjust continuously, unlike the consumption and saving that feature jumps at each nodes of the tree. In particular, saving is subject to two opposite effects:
(i) at the revelation of the common noise, there is a shift in expectations of households, and they reduce their precautionary saving during the boom, decreasing demand for asset/capital and raising interest rate.
(ii) since the productivity rises over time, the capital is more productive and the demand for capital from firms slowly increases as well, raising interest rate and wages. This mechanically increases disposable income, saving and investment.

The first motive is instantaneous but typical of heterogenous agents model while the second is revealed with the change in productivity and prices and is similar to standard Brock-Mirman dynamics. The shift from the first to the second will change the trend of capital evolution, from a drop to a surge during boom periods.

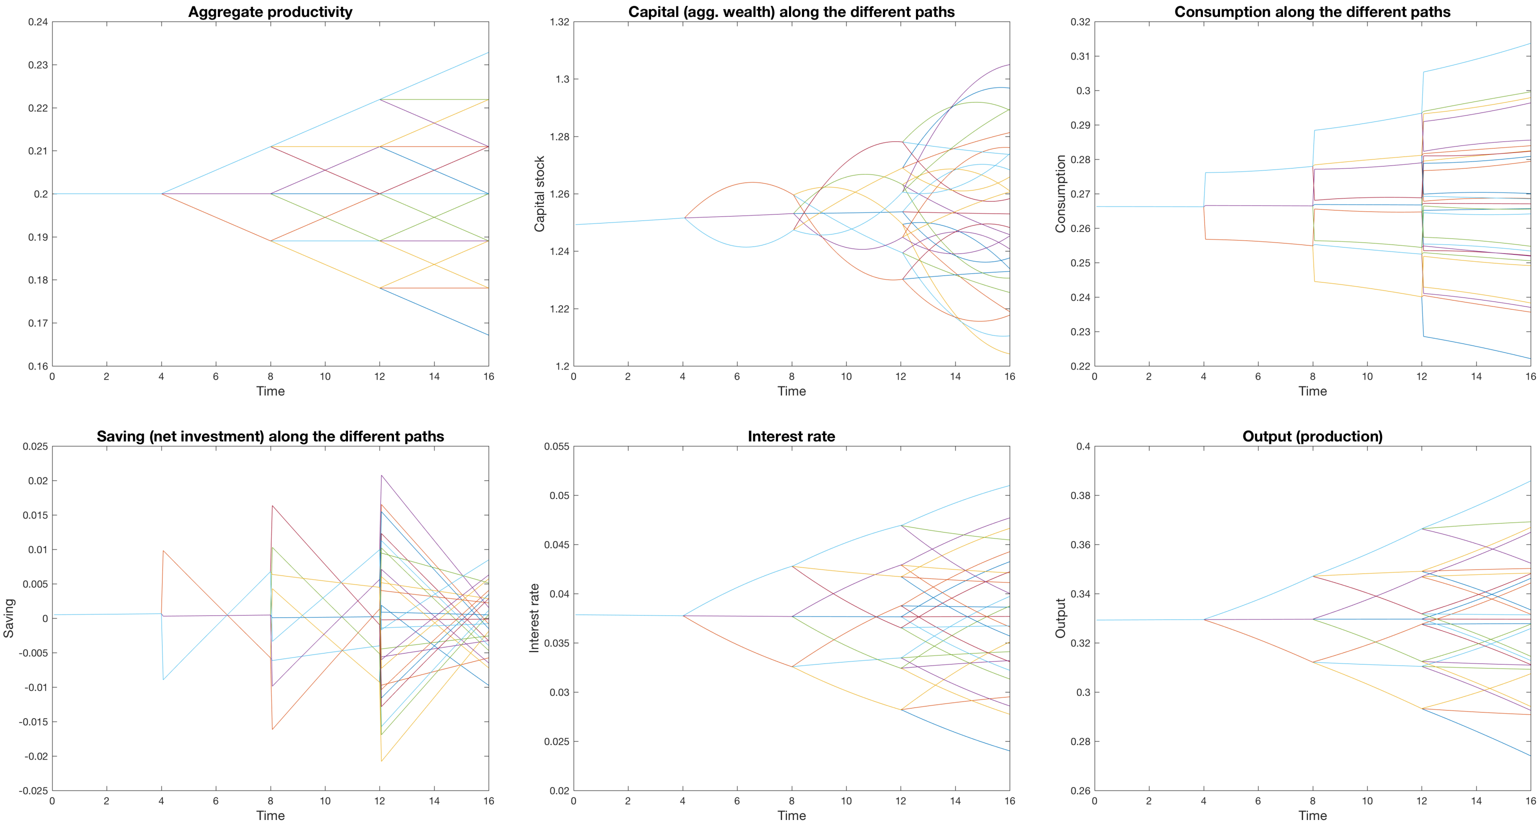

Figure 1: Evolution of aggregate variables with a Brownian TFP in the Krusell-Smith model

The value function and wealth distribution across time

As explained in section 3.1, both the value function and the measure are now random variables adapted to the filtration generated by the common noise $\sigma(B_t^0)$. Since these objects are infinite dimensional and "functions" of the common noise (i.e. adapted to the "alea" generated by this noise), the task to display the evolution might not be obvious.

With our discretization procedure, we choose to use the tree (ii) as described in section 4, with $K = 2$ and $M = 9$ to make the result of the evolution of v_t and m_t clearer. Note that the simulation for this tree (ii) is displayed in the long version of this paper.

The probability distribution of the value or measure would be $\mathbb{E}_t\big(v_t|\sigma(B_t^0)\big) \approx h^v(B_t^0)$ and $\mathbb{E}_t\big(v_t|\sigma(B_t^0)\big) \approx h^m(B_t^0)$, where $h^v : \Omega_t \mapsto \mathcal{C}(\mathbb{X})$ and $h^m : \Omega_t \mapsto \mathcal{P}(\mathbb{X})$, where Ω_t is heuristically the probability space supporting the common noise at time t.

These random variables have value in $\mathcal{C}(\mathbb{X})$ or $\mathcal{P}(\mathbb{X})$, where $\mathbb{X}$ is in the 3 dimensions. Instead of plotting one dimension at a time, we choose to display two trajectories ω_t :

- The "best case" scenario, where aggregate productivity only increases at each tree node
- The "worst case" scenario, where the economy is only subject to TFP contraction

First of all, let us show the value function and the measure at the initial point of the game, i.e. at $t = 0$ and at the end of the game – when the randomness is switched off as explained in our remark at the beginning of section 3.1 – i.e. at $t = T$ in these two scenarios for productivity.

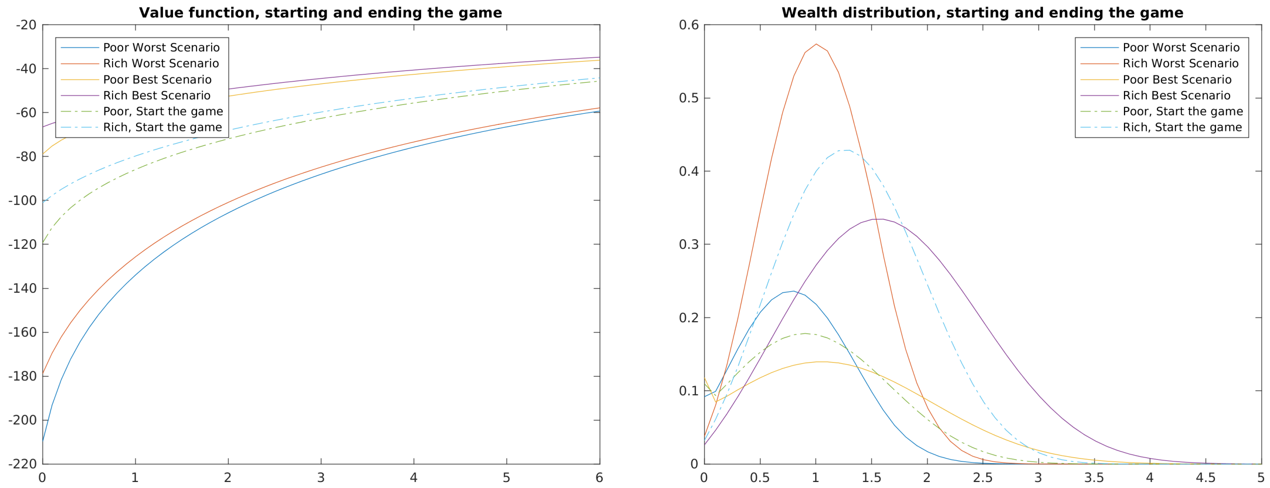

Figure 2: Value functions and distribution of both type of agents at the time boundaries

We see the aggregate productivity distorts both the measure of the agents and shift upward or downward the value function for both high and low income agents. The mechanism is intuitive: the value increases in boom and decreases in contraction, and the measure shifts to the right – people become richer on average – in boom, and shift to the left – concentrate more on low wealth – when economy is depressed. We now describe more precisely the underlying mechanisms for these two trajectories.

As described before, for each 'wave' of shock $m \in \{1, \ldots, M\}$, there is a change of information about productivity A_t between t_m^- and t_m^+. Therefore, the value function is discontinuous at this point t_m to account for the revelation of the shock: this jump in the value function is linked to the jump in consumption and saving since the agents will evaluate differently the need of precautionary saving – depending of the probability of future shocks, thus changing the conditional expectation.

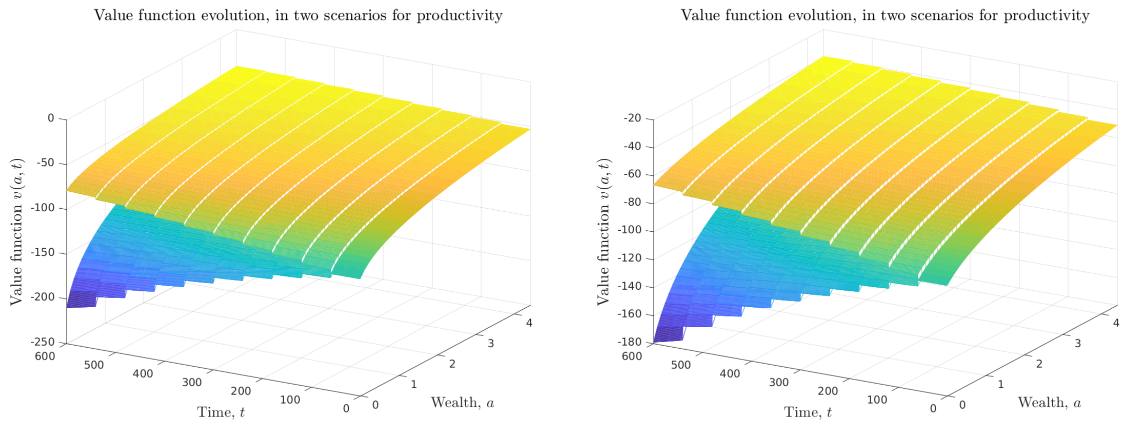

Figure 3: Value functions both type of agents in two trajectories of agg. productivity over time

For the two scenarios of productivity, the previous two graphs display the evolution of value functions for poor agents (low income $j = 1$, LHS) and rich agents $j = 2$ (RHS). Note that the scale may change the shape of the graphs: despite having different state z_j the poor and rich agents have almost similar value function, as displayed in the fig. 2

On the contrary, the wealth distribution is continuous over time. The control and the controlled drift, $b(t, x, m, \alpha^{\star}) = -D_p\mathcal{H}(t, x, m, \nabla_x v)$, will also jump – as function of the value function – but the measure have a smooth evolution as only the infinitesimal variation in the Fokker Planck distorts the shape of the distribution.

In the next four graphs, the wealth distribution is represented as a function of time for the two scenario: best case (LHS) and worst-case scenario (RHS), for the poor agent (first line), or the rich (second line). Note that the wealth distribution shift leftward in situation of recession (decrease in productivity), since the income in mechanically lower: both wages and interest rate are decreasing due to the structural change in productivity A_t.

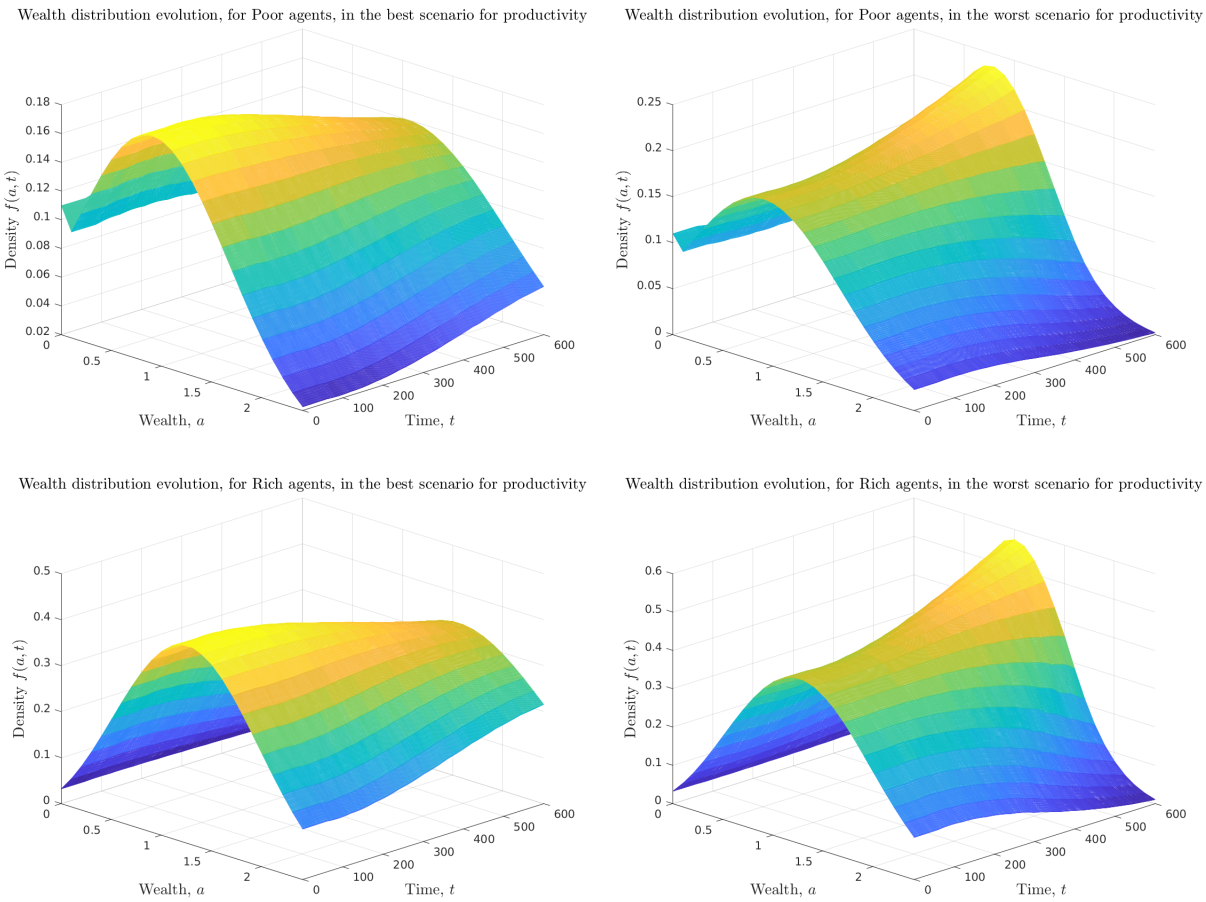

Figure 4: Wealth distribution of both type of agents, in two trajectories of agg. TFP over time

However, we claim that – due to precautionary saving – this leftward shift is attenuated by the willingness of both rich and poor agents to accumulate relatively more assets. Indeed, they self-insure against potential risk of further recessions or risk of falling (or staying) into the low income state. This mechanism changing the control has a smoothing effect on capital evolution – the first moment of the distribution – and the evolution of consumption and saving has general equilibrium effects on wages and interest rate, as we explain next.

Note that the effect is qualitatively reversed when the productivity increases – in the best case scenario – but the quantitative result is not symmetrical, since the MFG with common noise is a highly non linear system.

To provide explicit argument for our claim – that precautionary saving have smoothing effect on capital – we know display the evolution of the control – $c_t^\star$ and the controlled drift $s_t = b(t, a, j, \bar{m}, c_t^\star)$ for the low income agents for the worst-case scenario.

Note that the mechanism is similar for rich agents and reversed in the case of economic growth (increase in A_t). To display the clear mechanism, we come back on our "simple" tree (i), where there is only $M = 3$ shocks, and the dynamics are more obvious.

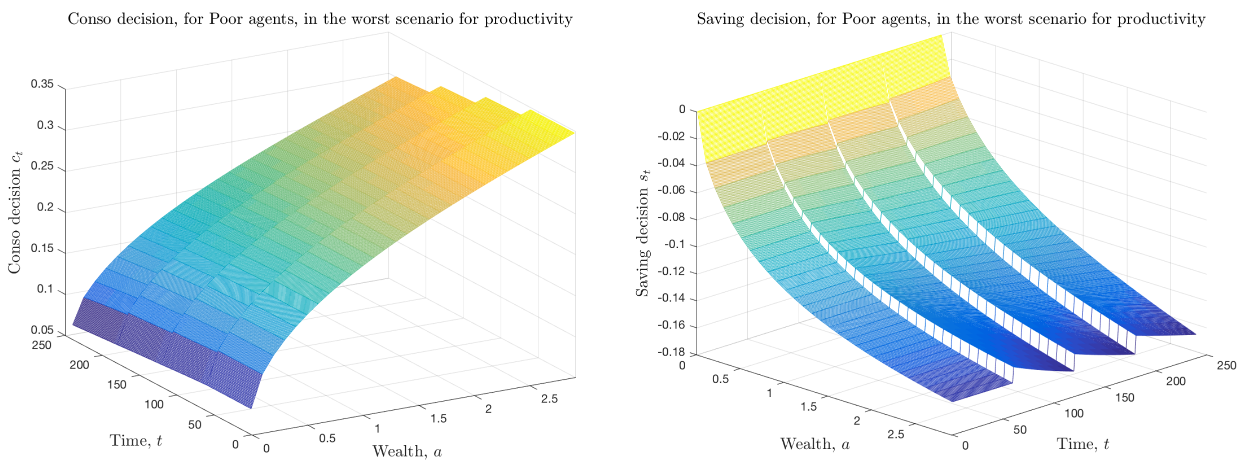

Figure 5: Control – c_t – and saving – s_t – over time for low income agents in the worst case scenario for productivity

As explained above in the description of the aggregate variables, the revelation of the common noise change abruptly the behavior of household: in recession, the saving jumps up and the consumption jumps down. This effect is strong, even though the aggregate productivity have not changed yet – due to the continuity of the Brownian motion $A_{t_m^-} = A_{t_m^+}$.

However, the important underlying reason is the change in conditional expectation of future shocks: there is higher probability of falling is a future depression when the TFP is already low – thanks to the martingale property of the Brownian motion.

The effect of such jump has general equilibrium effect: as displayed in fig. 1, the decrease in capital is delayed – with capital increasing at first before decreasing subsequently – and the interest rate drops more that what would be strictly implied by the change in productivity – almost 15% drop when the change in TFP in only 5%.

The transmission channels and the quantitative effects should be analyzed with Impulse Response Function (IRF) and that is the task we cover in the next section.

5.2 Jump-drift process and Impulse Responses

In this section, we consider the aggregate shock, affecting firm's productivity, that follows a "Jump Drift" Process. This stochastic process, resembling closely to an Ornstein Uhlenbeck process, is not a diffusion, but feature jumps and a mean-reverting behavior. Such dynamics allows to plot the Impulse Response Functions (IRF) of the system. More precisely, the aggregate productivity and the deviation from steady-state follows:

$$A_t = \bar{A}_0 \, e^{\widetilde{A}_t} \qquad\qquad d\widetilde{A}_t = -\theta \widetilde{A}_t dt + \; \varepsilon \, dN_t$$
$$dN = 1 \quad \text{with intensity} \;\; \lambda \qquad\qquad \text{and} \;\; \varepsilon \sim \mathcal{N}(0, \sigma^2)$$

For this experiment, we plot one trajectories – among all the branches of the tree – that represents the dynamics of the economy after a one-time temporary deviation from steady-state, i.e. a jump $dN_t = 1$ with $\varepsilon = \sigma$. On the first plot (showing $\widetilde{A}_t$) we display the upper and lower bound of our approximation tree for this process. After the shock, the system is not perturbed and comes back to the steady-state value of the aggregate productivity.

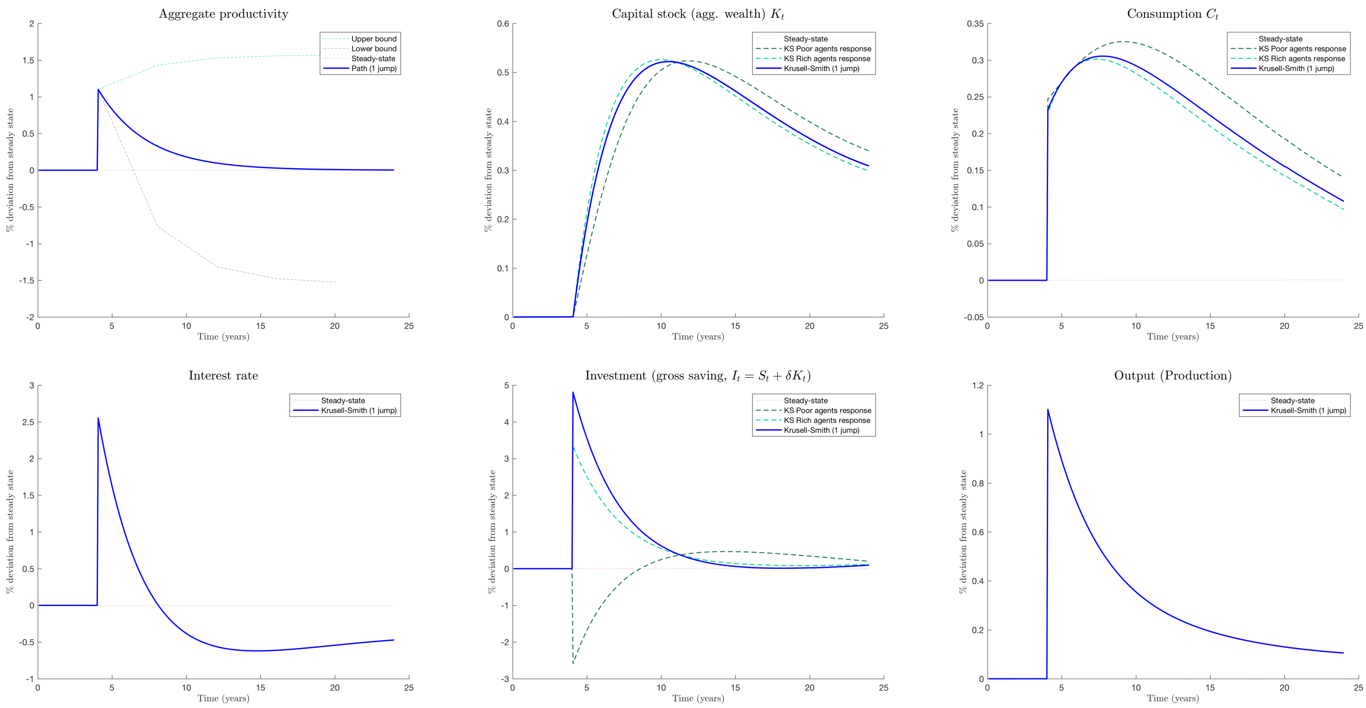

Figure 6: Impulse Responses to a one-time TFP shock in the Krusell-Smith model

An interesting feature, is the fact that poor agents, who are more likely to be credit-constraint and hand-to-mouth, are more willing to reduce their precautionary saving when the positive shock – while it is the reverse for rich agents – and increases relatively more their consumption than the rich (high-income) agents.

Moreover, the increase in capital decreases its marginal productivity over time, causing the interest rate to drop below its pre-shock level. This feature is standard in the Brock Mirman model. We now want to now if the two models can provide quantitative differences.

5.3 Comparison with representative agent model

In this section, we compare the previous model to the standard Brock Mirman model. We simulate the model using DYNARE (cf. the CEPREMAP working paper Adjemian et al. (2011)), a software solving rational expectations models using perturbation methods.

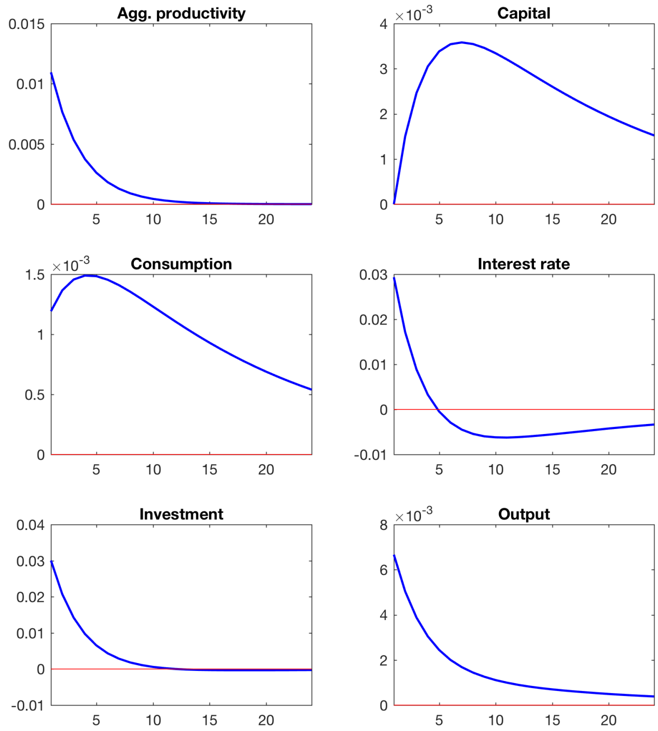

Figure 7: Impulse to a TFP shock in the Brock Mirman model

If the quantitative difference may not be obvious, note that the capital stock increase by 0.5 % in the Krusell Smith model compared to 0.3 % in the Brock Mirman. This is due to the larger increase in investment – caused by precautionary saving of the rich agents – that is almost twice as large, 5% compared to 3 % and that imply a stronger amplification effect of the productivity shock on output. Indeed the production reacts one to one to the 1.1% productivity shock, while in the Brock Mirman the effect is dampened by consumption smoothing (c.f. Euler Equation). Due to the presence of Hand-to-Mouth consumers, the reaction of consumption is also much stronger in the Krusell-Smith model, with 0.25% compared to only 0.125 % in the Brock-Mirman model.

5.4 Extensions

In two extensions, we consider a similar framework with :

(i) Endogenous labor supply

(ii) Price-rigidities: One-Asset HANK model

We use these two models to draw direct analogies with the most standard frameworks of the macroeconomic literature: the RBC model and the NK model. By comparing the representative agents models and their heterogeneous agents counterpart, we show that labor supply also response to precautionary saving motives and that the interaction with nominal rigidities could amplify aggregate fluctuation. Simulating the models using the same methods, we provide here a short summary of these results, knowing that the complete description of the models and their outcome can be found in the longer version.

In presence of endogenous labor supply, when the shocks are revealed, the change in expectation affects the work effort of households in different ways. The agents with low income/low wealth and close to the borrowing constraint react by supplying more labor. This implies a slight amplification effect on the aggregate labor supply and output, that rises higher (twice higher and 20 % respectively for labor and output) and last longer that the representative agents (RBC) counterpart.

In the New Keynesian version of this model, firm's Rotemberg pricing and Taylor-rule monetary policy also affect the response to productivity shocks. The rigidity in prices when associated with the precautionary saving can have a destabilizing effect on aggregate output. This amplified transmission effect is known in HANK model, as explained in Challe, Matheron, Ragot, and Rubio-Ramirez (2017) where the model is simulated with reduced heterogeneity. Compared to the representative agent NK model, the predictions for labor supply, output "gap" and inflation are reversed: positive TFP shock induces increased work efforts and consumptions by poor households, generating amplification, inflation and monetary tightening (instead of easing). The response of monetary policy with Taylor rule is also far from being optimal, as discussed thoroughly in Bhandari, Evans, Golosov, and Sargent (2018)

6 Conclusion

In this project, we used novel methods to simulate the Mean Field Game system with common noise. This system is composed of two stochastic PDEs: a Hamilton-Jacobi-Bellman describing the evolution of the value function and a Fokker-Planck describing the evolution of the distribution, and both depends on the realization of the common noise. After a discretization of this process using different tree structures, these two PDEs were solved using a Finite-Difference Scheme. This method can be applied to a large class of models, e.g. the Krusell-Smith model and extensions.

Our main result is to show that agents heterogeneity matters quantitatively for aggregate fluctuations through different channels. First, wealthier agents are hoarding larger quantity of precautionary savings to self-insure against both idiosyncratic and aggregate uncertainty, implying a surge in saving-investment after productivity shocks. Second, poorer agents are more reactive to aggregate fluctuations – in income, wage and interest rate – and exhibit typical hand-to-mouth behavior in consumption. This contrasts substantially with Krusell-Smith result on "approximate-aggregation": in our simulation, it does not appear that controls of agents – consumption and saving – exhibit any particular linearity in their states – income and wealth.

The strength of our approach is to keep the full heterogeneity of agents in both policy function and distribution, and to be clear on anticipations of agents on future paths of shocks.

More specifically, pure anticipation effects of uncertainty and general equilibrium effects could be easily disentangled using our method. Our conclusion matters quantitatively for macroeconomic dynamics as its creates amplification mechanisms of productivity shocks: in comparison to the representative agent framework where permanent income hypothesis prevails and propagation mechanisms are usually weak, our model feature one-to-one response of output to TFP shocks.

References

Achdou, Y., F. Camilli, and I. Capuzzo-Dolcetta (2013). Mean field games: convergence of a finite difference method. *SIAM Journal on Numerical Analysis 51*(5), 2585–2612.

Achdou, Y. and I. Capuzzo-Dolcetta (2010). Mean field games: Numerical methods. *SIAM Journal on Numerical Analysis 48*(3), 1136–1162.

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2017, November). Income and wealth distribution in macroeconomics: A continuous-time approach. *R & R, Review of Economic Studies* (NBER 23732).

Adjemian, S., H. Bastani, M. Juillard, F. Mihoubi, G. Perendia, M. Ratto, and S. Villemot (2011). Dynare: Reference manual, version 4.

Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf (2018). When inequality matters for macro and macro matters for inequality. *NBER Macroeconomics annual 32*(1), 1–75.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics 109*(3), 659–684.

Algan, Y., O. Allais, and W. J. Den Haan (2008). Solving heterogeneous-agent models with parameterized cross-sectional distributions. *Journal of Economic Dynamics and Control 32*(3), 875–908.

Bally, V., G. Pagès, and J. Printems (2005). A quantization tree method for pricing and hedging multidimensional american options. *Mathematical Finance 15*(1), 119–168.

Barles, G. and P. E. Souganidis (1991). Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic analysis 4*(3), 271–283.

Benhabib, J., A. Bisin, and S. Zhu (2011). The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica 79*(1), 123–157.

Benhabib, J., A. Bisin, and S. Zhu (2015). The wealth distribution in bewley economies with capital income risk. *Journal of Economic Theory 159*, 489–515.

Bertucci, C., J. M. Lasry, and P. L. Lions (2018). Some remarks on mean field games.

Bewley, T. (1986). Stationary monetary equilibrium with a continuum of independently fluctuating consumers. *Contributions to mathematical economics in honor of Gérard Debreu 79.*

Bhandari, A., D. Evans, M. Golosov, and T. J. Sargent (2018). Inequality, business cycles, and monetary-fiscal policy. Technical report, National Bureau of Economic Research.

Capuzzo-Dolcetta, I. and P.-L. Lions (1990). Hamilton-jacobi equations with state constraints. *Transactions of the American Mathematical Society 318*(2), 643–683.

Cardaliaguet, P. (2013/2018). Notes on mean field games. *Lecture notes from P.L. Lions' lectures at College de France and P. Cardaliaguet at Paris Dauphine*.

Cardaliaguet, P., F. Delarue, J.-M. Lasry, and P.-L. Lions (2017). The master equation and the convergence problem in mean field games. *arXiv preprint arXiv:1509.02505*.

Carmona, R. and F. Delarue (2014). The master equation for large population equilibriums. In *Stochastic Analysis and Applications 2014*, pp. 77–128. Springer.

Carmona, R. and F. Delarue (2018). *Probabilistic Theory of Mean Field Games with Applications I-II.* Springer.

Carmona, R., F. Delarue, and A. Lachapelle (2013). Control of mckean–vlasov dynamics versus mean field games. *Mathematics and Financial Economics 7*(2), 131–166.

Carmona, R., F. Delarue, and D. Lacker (2016). Mean field games with common noise. *The Annals of Probability 44*(6), 3740–3803.

Challe, E., J. Matheron, X. Ragot, and J. F. Rubio-Ramirez (2017). Precautionary saving and aggregate demand. *Quantitative Economics 8*(2), 435–478.

Corlay, S. (2011). *Some aspects of optimal quantization and applications to finance.* Ph. D. thesis, Université Pierre et Marie Curie-Paris VI.

Fleming, W. H. and H. M. Soner (2006). *Controlled Markov processes and viscosity solutions*, Volume 25. Springer Science & Business Media.

Gabaix, X., J.-M. Lasry, P.-L. Lions, and B. Moll (2016). The dynamics of inequality. *Econometrica 84*(6), 2071–2111.

Galí, J. (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications.* Princeton University Press.

Heathcote, J., K. Storesletten, and G. L. Violante (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review 104*(7), 2075–2126.

Huang, M., R. P. Malhamé, and P. E. Caines (2006). Large population stochastic dynamic games: closed-loop mckean-vlasov systems and the nash certainty equivalence principle. *Communications in Information & Systems 6*(3), 221–252.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to hank. *American Economic Review 108*(3), 697–743.

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Krusell, P. and A. A. Smith, Jr (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy 106*(5), 867–896.

Lasry, J.-M. and P.-L. Lions (2007). Mean field games. *Japanese journal of mathematics 2*(1), 229–260.

Luschgy, H. and G. Pagès (2002). Functional quantization of gaussian processes. *Journal of Functional Analysis 196*(2), 486–531.

Pagès, G. (2017). Introduction to numerical probability for finance. *UPMC, Polycopié du 'Master 2 Probabilités et Finance'.*

Pagès, G. and J. Printems (2003). Optimal quadratic quantization for numerics: the gaussian case. *Monte Carlo Methods and Applications mcma 9*(2), 135–165.

Pham, H. and X. Wei (2017). Dynamic programming for optimal control of stochastic mckean–vlasov dynamics. *SIAM Journal on Control and Optimization 55*(2), 1069–1101.

Ragot, X. (2018). Heterogeneous agents in the macroeconomy: reduced-heterogeneity representations. In *Handbook of Computational Economics*, Volume 4, pp. 215–253. Elsevier.

Reiter, M. (2010). Solving the incomplete markets model with aggregate uncertainty by backward induction. *Journal of Economic Dynamics and Control 34*(1), 28–35.

Soner, H. M. (1986). Optimal control with state-space constraint i. *SIAM Journal on Control and Optimization 24*(3), 552–561.

Souganidis, P. E. (1985). Approximation schemes for viscosity solutions of hamilton-jacobi equations. *Journal of differential equations 59*(1), 1–43.

Winberry, T. (2016). A toolbox for solving and estimating heterogeneous agent macro models. *Forthcoming Quantitative Economics.*

Yong, J. and X. Y. Zhou (1999). *Stochastic controls: Hamiltonian systems and HJB equations*, Volume 43. Springer Science & Business Media.

Young, E. R. (2010). Solving the incomplete markets model with aggregate uncertainty using the krusell–smith algorithm and non-stochastic simulations. *Journal of Economic Dynamics and Control 34*(1), 36–41.