Under consideration for publication in Knowledge and Information
Systems

A Computational Model of Labor Market Participation with Health Shocks and Bounded Rationality†

Alessandro Moro[1] and Paolo Pellizzari[2]

[1]Bank of Italy, Statistical Data Collection and Processing Directorate, External Statistics Division, Rome, Italy, alessandro.moro2@bancaditalia.it;
[2]Ca' Foscari University, Department of Economics, Venice, Italy, paolop@unive.it

Abstract. This paper presents a computational agent-based model of labor market participation, in which a population of agents, affected by adverse health shocks that impact the costs associated with working efforts, decides whether to leave the labor market and retire. This decision is simply taken by looking at the working behaviors of the other agents, comparing the respective levels of well-being and imitating the more advantageous decision of others. The analysis reveals that such mechanism of social learning and imitation suffices to replicate the existing empirical evidence regarding the decline in labor market participation of older people. As a consequence, the paper demonstrates that it is not necessary to assume perfect and unrealistic rationality at the individual level to reproduce a rational behavior in the aggregate.

Keywords: Labor market participation; Health shocks; Bounded rationality; Agent-based modeling

1. Introduction

This paper is an extended and enriched version of the contribution presented by Moro and Pellizzari (2016) at the Conference on Practical Applications of Agents and Multi-agent Systems (PAAMS), Seville, 2016. In particular, this work outlines an agent-based model of labor market participation, in which a population of agents is naturally affected over time by adverse shocks that reduce

† The views expressed in the paper are those of the authors and do not involve the responsibility of the Bank of Italy.

Received 27 Sep 2016
Revised 22 Apr 2017
Accepted 17 Jun 2017

health and magnify the costs associated with working efforts. Agents decide whether or not to leave the labor market and retire (or access an otherwise available source of support, like a social security benefit or pension). This decision is made by looking at the working behaviors of other randomly met agents and comparing the respective levels of well-being, rather than maximizing the utility function in a standard and fully rational way.

The analysis reveals that such mechanism of learning-by-meeting and imitation, even in the presence of otherwise entirely naive agents, can replicate some of the existing empirical evidence regarding the decline in labor market participation of older people. As a consequence, the paper shows that it is not necessary to assume full rationality at the individual level to reproduce a rational behavior in the aggregate.

The relevance of this claim can be understood by considering that the mainstream literature on labor supply and retirement choices is dominated by structural models, in which fully rational agents take decisions by maximizing their life-time utility. Examples of this kind of modeling approach, that also incorporates health shocks, are Alonso-Ortiz (2014), Bound et al (2010), Burkhauser et al (2004), French (2005), Gilleskie (1998), Gustman and Steinmeier (2002), Gustman and Steinmeier (2015), Haan and Prowse (2014), Heyma (2004), Laun and Wallenius (2015), Laun and Wallenius (2016), Pelgrin and St-Amour (2016), Rust and Phelan (1997). To give some flavor, the agents in the paper by Rust and Phelan optimize their behavior working backward the optimal employment decision (i.e. how much to work) at any time t between the age of 57 and 102 and choosing whether to apply for Social Security. The situation faced by workers is described by 7 state variables, and optimality is determined in principle in each of over 14,000 possible future occurrences maximizing the value of expected discounted utility.[1]

There is overwhelming evidence in experimental economics and psychology that shows how framing effects and contextual factors limit human rationality (Rabin, 1998) and how poorly humans behave in problems that involve backward induction (Camerer, 1997). Whether the previously described rational optimization can be performed by realistic agents is a no-brainer: consumer decision making is affected by information overload as soon as more than 10 options and 15 information items are provided (Lurie, 2004), (Malhotra, 1982), and it is well known that humans do not have good performances when dealing with (non trivial) probabilistic setups or fail to grasp even the meaning of simple ratios (Reyna and Brainerd, 2007).

An alternative view, that incorporates the main findings of the empirical literature and much better accommodates with the psychological limitations of real agents, makes use of agent-based models where most often boundedly rational agents act using heuristics and cognitively sound adaptive rules of thumb. For instance, Lettau (1997) presents a model of portfolio decisions with agents that learn through a genetic algorithm how to behave based on random innovations

[1] This is an Herculean task requiring, among other things, to estimate transition probabilities from one state to all possible future states based on beliefs which, in turn, depend on 26 parameters to be determined. Other questionable assumptions include "rational expectations", a notion according to which all individual subjective probabilities equate the "objectively estimable population probability" and that requires the capability to anticipate regulatory changes. The paper acknowledges several times that only some justification for such assumptions can be provided.

and imitation. In particular, it is shown that investors in mutual funds exhibit the same pattern of investments of the adaptive agents described in the paper and that the model closely matches the mutual fund data.

Recently, Goudet et al (2015) have developed an agent-based model of the French labor market, in which boundedly rational workers can be in different states (inactive, unemployed, employed, . . .) and decide to move from one state to another by comparing the utility level of the current state with the expected utility in the new one.

More germane to our topic, Axtell and Epstein (1999) describe a model where a relatively small number of workers are rational, whereas a small proportion of them retire at random, and the vast majority just imitates the behaviors of the other agents in their social network. Again, high levels of optimal behavior can result in the aggregate despite low levels of individual rationality. Our paper supports a similar conclusion in a framework that also incorporates health shocks and proposes a different mechanism of social learning.

It is worth stressing that the present paper proposes a mechanism of imitation of behaviors. This is the main difference with respect to the dispositional contagion framework proposed by Epstein (2014), in which each agent is not influenced by the actions of his peers, but by their dispositions. In our study the former approach is preferred because behaviors are more easily observed than dispositions.

The rest of the paper is organized as follows: the next section illustrates the main stylized facts that the model is able to reproduce, it describes the model, its calibration as well as the results of the simulations; finally, the last section concludes stressing the reasons why the present model can be considered a valid alternative to the mainstream approaches.

2. Analysis

This section of the paper is divided in five parts: the first one describes a few traits of the empirical evidence regarding the decline in the labor market participation of elderly workers affected by adverse health shocks; the second subsection describes in detail the proposed agent-based computational model; the third part shows the results of the model simulation and, in particular, how the model is able to capture some relevant stylized facts; then, it is proved that the main conclusions of the model are robust to different values of the parameters and different specifications of the model; finally, the last subsection presents some welfare implications of our analysis.

2.1. Stylized Facts

Data from the first wave of the English Longitudinal Study of Ageing (ELSA) are chosen in order to illustrate the empirical evidence regarding the labor market participation of elderly people and the main stylized facts that the model is able to replicate. This dataset contains information on both economic and labor supply variables, as well as on self-assessed health status.

The ELSA sample is representative of people aged 50 and over, living in private households in England. The interviews of the first wave began in March 2002 and spanned 12 months, completing in March 2003.

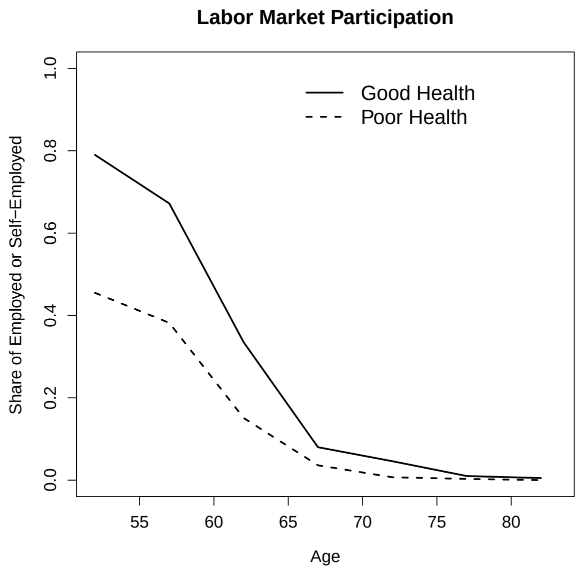

Fig. 1. Labor market participation by age and health. The solid line represents the participation of people with a good health, while the dashed line shows the participation of people affected by poor health conditions. Source: first wave of the English Longitudinal Study of Ageing (ELSA).

Figure 1 shows how labor market participation, measured as the proportion of people employed or self-employed, declines with the age of individuals. Specifically, each curve is associated with a different level of health and displays, in correspondence to different age classes, the percentage of people in the ELSA sample that are inside the labor market.[2]

Although labor market participation decreases with age in each health group, it is evident that people affected by poor health conditions have much lower participation rates than the individuals with a good health status, and this is true in every age group. In fact, people affected by adverse health shocks who are eligible to disability benefits or retirement have a strong incentive to leave the labor market. Issues related to the effects of various benefits are still debated by policymakers as they may artificially alter the participation to the labor market or offer awkward incentives to workers.

2.2. The Model

In the model two types of agent, workers and pensioners, interact in a bidimensional torus space. At the beginning of the simulations ($t = 0$), the proportion of pensioners is equal to p. Workers receive a labor income y^l, which is greater than

[2] The data used to draw the curves in Figure 1 are taken from the Table 4A.3 in the Annex 4.1 - Tables on Work and Retirement of the first wave of the ELSA survey. http://www.elsa-project.ac.uk/reportWave1

the pension y^p, but their working effort e^l is bigger than the one of pensioners e^p.[3]

Each agent i starts with an age $a_{i,0}$, randomly drawn from a uniform distribution on the (50, 89) interval, and receives a health $h_{i,0}$, drawn from a uniform distribution on the $\left(0, \frac{50}{a_{i,0}}\right)$ interval.

Agents derive well-being $U^s_{i,t}$ from their working decisions that assumes the following expression:

$$U^s_{i,t} = y^s_{i,t} - \frac{e^s_{i,t}}{h_{i,t}} \tag{1}$$

where $s \in \{l, p\}$ denotes the state, labor or pension, of the individual. Thus, the well-being is an increasing function of income, whereas it decreases when the working effort increases. Moreover, the negative effect of exerting effort is greater when an agent suffers from a bad health, i.e. when $h_{i,t}$ is low.

This definition of well-being is consistent with the framework proposed by the OECD (2013). In particular, this institution measures well-being in terms of outcomes achieved in two broad domains: material living conditions (income and wealth, jobs and earnings, ...) and quality of life (physical and mental health, work-life balance, ...).

In each time step, an agent is randomly selected (asynchronous activation) and he moves to a random site of the bidimensional space, where he can meet another agent. The movement rule is very simple: the agent turns right at an angle that is randomly picked in $\{0, 1, 2, \dots, 359\}$ degrees and moves forward f patches (steps). If the moving agent "lands" on a patch hosting another person and the two are similar in terms of age and health status, they exchange information about their working choices, represented by the vector (y, e), and their levels of well-being U. The similarity between the agents i and j is established if the following inequality holds:

$$w \frac{\mid a_{i,t} - a_{j,t} \mid}{100} + (1 - w) \mid h_{i,t} - h_{j,t} \mid < d, \tag{2}$$

where $0 \leq w \leq 1$ measures the relative importance of age with respect to health, and d represents a similarity threshold. A low (high) value of w means that agents are more likely to imitate the behaviors of those with similar health conditions (age). Additionally, a higher value of the threshold d implies that agents are more willing to copy the working choices of dissimilar players. The difference of ages is divided by 100 to obtain a measure, whose magnitude can be compared to that of the difference of health levels.

Once inequality (2) is satisfied (and only in this case), the two agents compare their respective well-being. Then, if $U^s_{i,t} < U^s_{j,t}$, agent i will imitate the working choices $\left(y^s_{j,t}, e^s_{j,t}\right)$ of agent j, and vice versa. Social learning is therefore implemented with this kind of imitation mechanism involving no optimization or information processing on the part of the agents, who resemble walkers that randomly encounter other hikers and have a chance to discuss different experiences related to working satisfaction and opportunity to retire.

[3] As mentioned previously, pensioners and pension income can also be interpreted as people who decide to leave the labor market and receive a subsidy and the corresponding disability benefit, respectively.

It is worth emphasizing that the movement in the bidimensional space can also be interpreted in metaphoric terms as the intensity of the social relationships of each agent: according to this view, a low value of the movement parameter f means that agents are shy and more likely to often meet the same close peers and, accordingly, their social network is very limited; conversely, more mobile and lively agents have higher chances of stumbling upon different peers, their acquaintances' network is more widespread and can, on average, obtain more information about retirement decisions.

Quite naturally, age evolves deterministically over time ($a_{i,t+1} = a_{i,t} + 1$), while health is affected by random shocks over the life cycle:

$$h_{i,t+1} = gh_{i,t} + (1 - g)x_{i,t+1} \tag{3}$$

where $0 \leq g \leq 1$ is a parameter capturing the persistence of health status, and $x_{i,t+1}$ is a random variable distributed with a uniform distribution on the $\left(0, \frac{50}{a_{i,t+1}}\right)$ interval: thus, the future level of health is a weighted average of the current health status and the random shock. This expression implies that an older agent is more likely to experience bad health shocks and, consequently, a decreasing level of health.

Finally, it is assumed that agents die at age 90 and they are replaced by new agents with an age equal to 50 and a health drawn from a uniform distribution on the $(0, 1)$ interval. Moreover, these new agents are pensioners or workers with a probability equal to p and $1-p$, respectively. A deterministic death is preferred because it allows to observe enough agents with bad health even in the oldest age classes, enabling the estimation of the average proportion of workers for such age ranges. Conversely, if the probability of death depends on the age of agents, we may have a selection problem, because only agents with a good health status would have survived.

The model is implemented in NetLogo (Wilensky, 1999). The user can choose the values for the model parameters, agents are randomly located in the bidimensional torus space, and each agent receives a random age, health and working condition. The code is available at http://tinyurl.com/paams16health and a screenshot of the Netlogo's graphical user interface is provided in Figure 2.

In every time period the simulation evolves through the following steps:

1. an agent is randomly selected;
2. the chosen agent, following the movement rule, moves to a random site, where he may meet another agent;
3. if there is an encounter and the two agents are reasonably similar in terms of age and health conditions, the one with the lowest level of well-being will imitate the working choices of the agent with the highest well-being.

These simple rules are replicated for all agents. Then, once all agents have performed their tasks, the age and health status of each player is updated. Agents with an age equal to 90 die, they are replaced by new ones endowed with a random health, time advances by one unit and the procedure starts again until the user quits the simulation or some specified time is reached. It is possible to monitor the evolution of the system by looking at some time series graphs that show the dynamics of the main model variables, as well as in the bidimensional space, in which workers and pensioners are depicted with different colors.

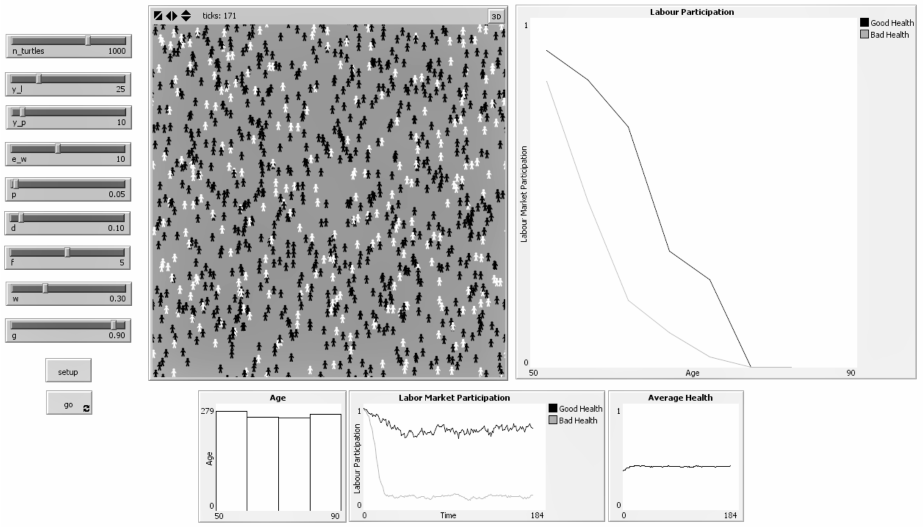

Fig. 2. Screenshot of Netlogo's graphical user interface. The larger monitor on the left shows the bidimensional torus space and the different agents: workers (white) and pensioners (black). The bigger monitor on the right displays the simulated labor market participation curves of the agents with a good health status (black) and those affected by bad health conditions (gray). The other pictures present some histograms and time series graphs that allow to monitor the evolution of the system. It is possible to change the values of the model parameters using the sliders on the left.

2.3. Calibration and Results

The model is simulated many times and under different conditions. It is assumed that the working effort of pensioners e^p is equal to 0 and the benefits pensioners receive y^p are kept constant and equal to 10: seen from the perspective of the workers, this means that the option of retirement is independent from health conditions and takes a normalized value of 10. The parameter g is also constant and equal to 0.9, implying that the health status of agents resembles a persistent and slowly decaying AR(1) process, which appears a rather sensible assumption.

A discrete grid is created for the other model parameters, and for each point of the grid, i.e. for each possible combination of parameter values, the model is simulated 100 times.[4] In every simulation the proportion of workers in the age range [50, 54], [55, 59], [60, 64], [65, 69], [70, 74], [75, 79], [80, 89] is recorded,[5] distinguishing between agents with a good health, defined as those agents with $h_{i,t} \geq 0.5$, and the ones with a bad health, i.e. those with $h_{i,t} < 0.5$. All simulations are halted after 200 time steps and the proportion of workers ($lmp^{pred}_{a,h,s}$) in the age class a and health group h is evaluated at the end of the simulation s. Then, for each combination of parameter values, the average proportion

[4] Simulations are performed using the BehaviorSpace tool in NetLogo, which allows to specify the grid of parameter values and the number of simulations for each point in the grid.

[5] These seven classes of age correspond to the ones presented in the Table 4A.3 in the Annex 4.1 - Tables on Work and Retirement of the first wave of the ELSA survey.

Table 1. Calibrated values of the model parameters.

Parameter	Description	Value
p	Initial share of pensioners	0.15
y^l	Workers' income	30
y^p	Pensioners' income	10
e^l	Workers' effort	15
e^p	Pensioners' effort	0
f	Movement	5
w	Age weight	0.7
d	Similarity threshold	0.1
g	Health persistence	0.9

is obtained averaging the $S = 100$ simulations, for each age range and health status:

$$lmp^{pred}_{a,h} = \frac{1}{S}\sum_s lmp^{pred}_{a,h,s} \tag{4}$$

The values of the parameters are selected in order to minimize the absolute distance between the observed labor market participation rates in the ELSA sample $lmp^{obs}_{a,h}$, in the different age and health groups, and those predicted by the model $lmp^{pred}_{a,h}$, divided by the product between the total number of age A and health H classes. Formally:

$$D = \frac{1}{AH}\sum_{a,h} \mid lmp^{obs}_{a,h} - lmp^{pred}_{a,h} \mid \tag{5}$$

The smallest value for D is 0.036, attained at the values listed in Table 1. Hence, on average, the absolute deviation between observed and simulated labor market participation rates is about 3.6%.

Figure 3 compares the average labor market participation rates predicted by the model ($lmp^{pred}_{a,h}$) for the different health and age groups, when the values contained in Table 1 are assigned to the model parameters, with the labor market participation rates observed in the ELSA sample. In this picture, a 90% confidence interval is also added by calculating the 5th and 95th percentiles of the simulated labor market participation rates (5th and 95th sample percentiles of $lmp^{pred}_{a,h,s}$) in each health and age class. Clearly, the model accurately fits the decreasing trend of the labor market participation curves with respect to age in both health groups. The observed labor market participation rates almost always lie inside the confidence bands and the average simulated rates are very close to those observed in the real sample.

Real and simulated data are listed in Table 2, together with the absolute differences that build the error measure D. Values in the last column of the table and the visual inspection of the pictures in Figure 3 show that most of modeling error is due to a unique entry (0.31, boldfaced), the participation rate of workers aged [50, 54] with bad health. This mismatch is due to the fact that the model underestimates the difference of working behaviors between the two health groups in this initial age class.

The participation of agents affected by adverse health shocks decreases more rapidly than that of agents with a good health. This aspect is again consistent

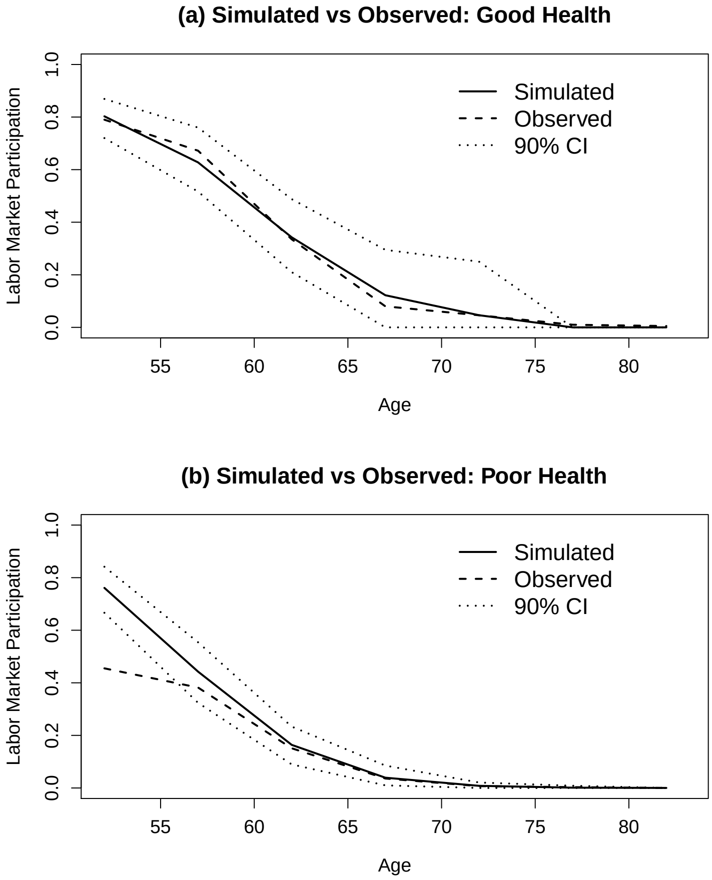

Fig. 3. Comparison between real and simulated labor market participation curves for the two health groups: (a) group of the agents with a good health ($h_{i,t} \geq 0.5$); (b) group of agents affected by bad health conditions ($h_{i,t} < 0.5$). The solid line represents the simulated labor market participation rates, while the dashed line shows the ones observed in the ELSA sample. The dotted lines are the 90% confidence bands.

with the empirical evidence contained in the ELSA survey. In fact, given the functional form chosen for the well-being function, a worker with a bad health who meets a similar and already retired agent is more likely to change his status than a worker with a good health. Conversely, a retired agent with a good health may decide to enter again in the labor market when meeting a worker more easily than a pensioner with a bad health. This explains the differences in participation between agents with good and bad health conditions and the reason why labor market participation declines in both groups as health deteriorates due to the aging of the population.

One of the main contributions of this paper is indeed the demonstration that no hyper-rational agent, like the ones included in structural models, is needed in the model to replicate the existing evidence as it is sufficient to assume that simple agents randomly and repeatedly meet and mimic each other's behav-

Table 2. Real and simulated labor market participation rates.

Health	Age	Real	Simulated	Abs. diff.
Good	[50,54]	0.79	0.80	0.01
	[55,59]	0.67	0.63	0.04
	[60,64]	0.34	0.34	0.01
	[65,69]	0.08	0.12	0.04
	[70,74]	0.05	0.05	0.00
	[75,79]	0.01	0.00	0.01
	[80,89]	0.01	0.00	0.01
Poor	[50,54]	0.46	0.76	**0.31**
	[55,59]	0.38	0.44	0.06
	[60,64]	0.15	0.16	0.01
	[65,69]	0.04	0.04	0.00
	[70,74]	0.01	0.01	0.00
	[75,79]	0.00	0.00	0.00
	[80,89]	0.00	0.00	0.00

ior. Moreover, the model appears to fit the empirical evidence better than (or, at least, as well as) other contributions in the literature and, remarkably, with fewer parameters. For instance, this conclusion emerges from the comparison of our results with Figures 2 and 3 in French (2005): in fact, the structural model presented in that paper, calibrated on the US Panel Study of Income Dynamics data, visibly underestimates the labor market participation rate of healthy individuals in the age range $[60, 70]$ and quite significantly overestimates that of unhealthy workers in the age class $[45, 65]$. The paper by Laun and Wallenius (2016) uses a different data set, i.e. the Luxembourg Income Study dataset, and, even if too high participation is predicted in some cases, the match of empirical data and the results of the model are good and similar to what we have obtained. However, even when the simulated moments closely fit the empirical ones, such as in Rust and Phelan (1997), it can be argued that structural models need an abundance of parameters to obtain acceptable outcomes.

2.4. Sensitivity Analysis

In order to illustrate the effects of changes in the values of the model parameters and prove the robustness of the obtained results, Figure 4 compares the baseline calibration, reported in the first panel of the figure, with other scenarios characterized by some deviations from the values contained in Table 1.

Panel 4b shows the effect of a higher income, i.e. y^l equals 35: it is evident that the labor market participation curves decline more slowly, especially for workers with a good health, because higher salaries make the outside option of retirement less attractive in relative terms. In fact, in an even more transparent way, the same pattern is observed in Panel 4c, in which working effort e^l is reduced to 10 (keeping income fixed at 30). Therefore, a policy implication that can be derived according to the model is that the reduction of the working effort represents a more powerful incentive than the increase of wages in order to prolong the labor market participation, especially for the workers with a good health status.

A lower value of the age weight w means that agents are more likely to imitate the behaviors of those who are similar in terms of health conditions than the ones with the same age. This leads to an increase in the divergence of the

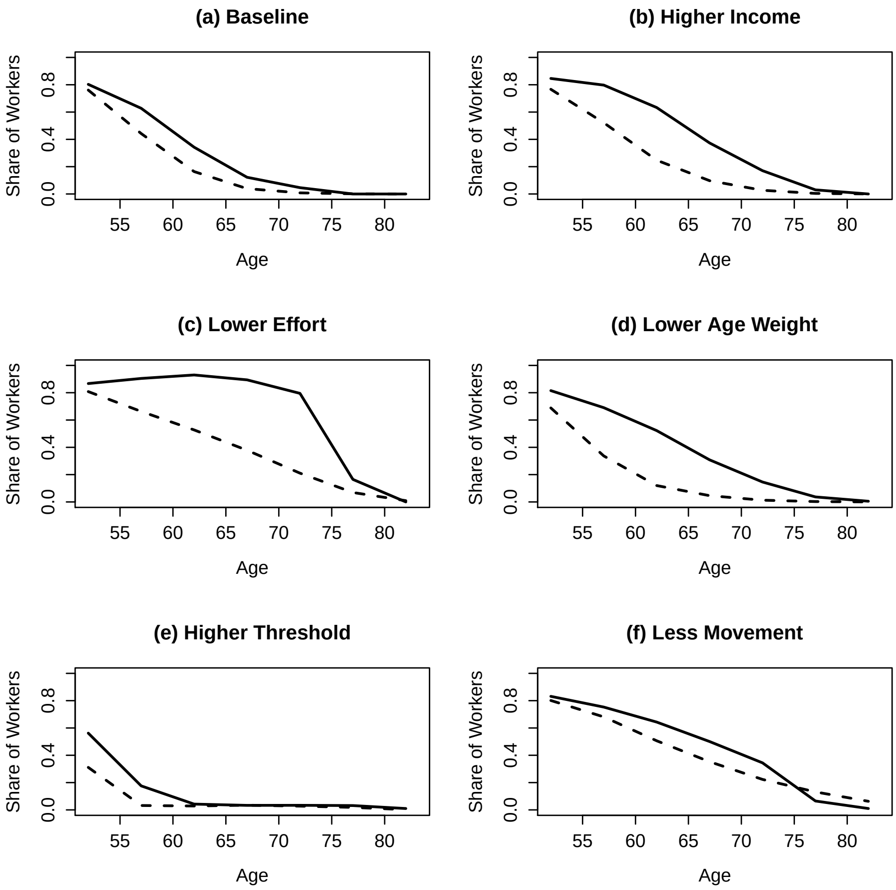

Fig. 4. Simulated labor market participation by age and health in different scenarios: (a) baseline calibration: values of Table 1 assigned to the model parameters; (b) higher labor income: $y^l = 35$; (c) lower working efforts: $e^l = 10$; (d) lower age weight: $w = 0.3$; (e) higher threshold: $d = 0.3$; (f) less movement: $f = 1$. In each panel, the solid line represents the participation of agents with a good health ($h_{i,t} \geq 0.5$), while the dashed line shows the participation of agents affected by poor health conditions ($h_{i,t} < 0.5$).

labor market participation curves with respect to the baseline scenario, as shown in Panel 4d. A higher value of the similarity threshold is associated with steeper participation curves, as can be seen in Panel 4e, because workers tend to imitate the behaviors of more dissimilar agents that may have already decided to retire.

The last panel of the figure illustrates the effect of the mobility parameter f, showing that less mobile (or less socially connected) agents enjoy more infrequent meetings and, consequently, flatter labor market participation curves. Remarkably, nearly all of the graphs contained in Figure 4 always preserve the same pattern, i.e. participation rates that decline when age increases, and in which workers affected by adverse health shocks are more likely to leave the labor market.

In the previous subsections it has been assumed that agents outside the labor market can decide to become active again when meeting a worker with a higher level of well-being. However, the real world may allow less flexibility and it was pointed out that there are legal systems in which the choice of leaving the labor market is irreversible: a typical example is the retirement decision in some countries.

In order to explicitly allow for these restrictions, as a further robustness check, the model is modified imposing that agents that decide to leave the labor market cannot change their status again. Figure 5 shows that the model is again able to fit the labor market participation curves with the same calibrated values used in the almost indistinguishable Figure 3. Understandably, the simulated rates are slightly lower because there are, on average, more pensioners given the assumption that these agents cannot have another chance to change status, and the value of the error criterion D raises by 0.2% only.

The results of the model are therefore robust to different values of the parameters and alternative specifications of the model.

2.5. Welfare Considerations

An important issue is whether the system described above reaches an equilibrium, and how this equilibrium can be described with respect to the two groups of agents identified by their health status.

Figure 6 presents, separately for the good health and poor health group, the time series of the average income, effort and well-being in a typical simulation of the calibrated model. It is possible to observe that after about twenty time steps the model reaches a long term stable state in which the agents experiencing a good health status tend to remain in the labor market earning a higher income and exerting a bigger effort than the people affected by poor health conditions (see Panel a and b of Figure 6). These considerations are completely in line with the previously described results and depend on the considerable cost in terms of well-being that the working efforts represent for the agents with a low health.

Panel 6c shows the time series of the average well-being for the two health groups. An interesting feature is represented by the evident shocks in the time series of the poor health group compared to the smooth dynamics of the agents with good health. This behavior is generated by the newborn workers endowed with very bad health conditions who do not meet a pensioner when moving around. In this case, the welfare (1) of these agents may even reach negative values that strongly affect the average well-being of the group, given the sensitiveness of the mean to such outlying and sporadic cases. However, the two time

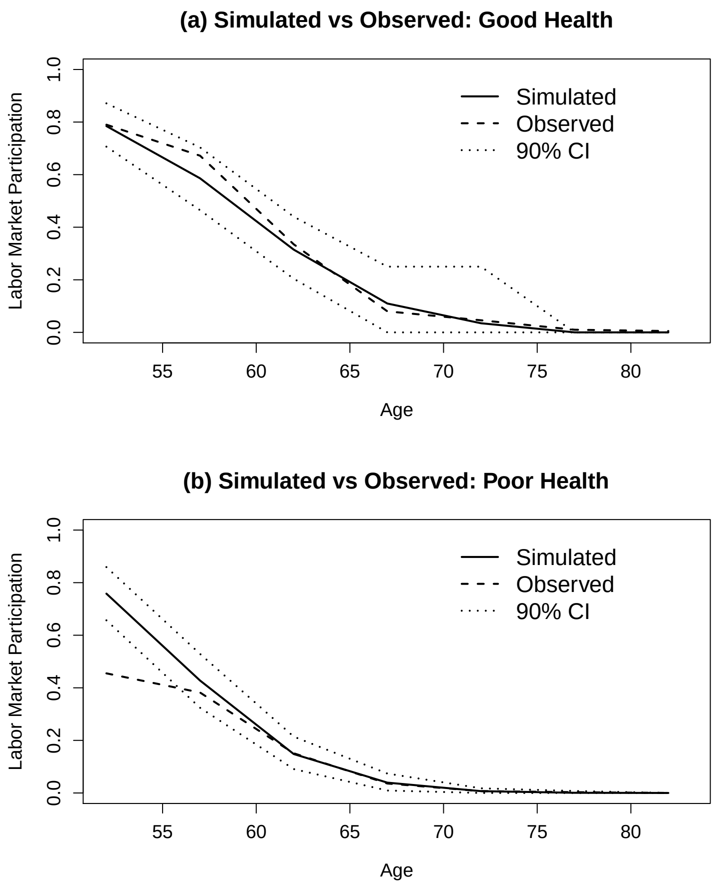

Fig. 5. Comparison between real and simulated labor market participation curves for the two health groups when agents that decide to leave the labor market cannot change their status again: (a) group of the agents with a good health ($h_{i,t} \geq 0.5$); (b) group of agents affected by bad health conditions ($h_{i,t} < 0.5$). The solid line represents the simulated labor market participation rates, while the dashed line shows the ones observed in the ELSA sample. The dotted lines are the 90% confidence bands.

series appear to asymptotically converge to levels of welfare that are close to each other. In fact, in our calibrated model, the well-being experienced by the agents affected by poor health conditions outside the labor market is comparable to that of workers with a good health status who stay in the labor market and the majority of the members of the two groups avoids changing its working conditions. Hence, despite the bounded rationality of these naive agents, they are able to reach a self-sorting equilibrium in a completely decentralized way.[6]

It is worth stressing that the calibrated value of the labor income y^l is such

[6] In this setup the term "equilibrium" differs from the concept of Nash equilibrium because our agents, keeping fixed the choices of their peers, are always willing to change their working status as long as they meet another similar agent with a higher well-being.

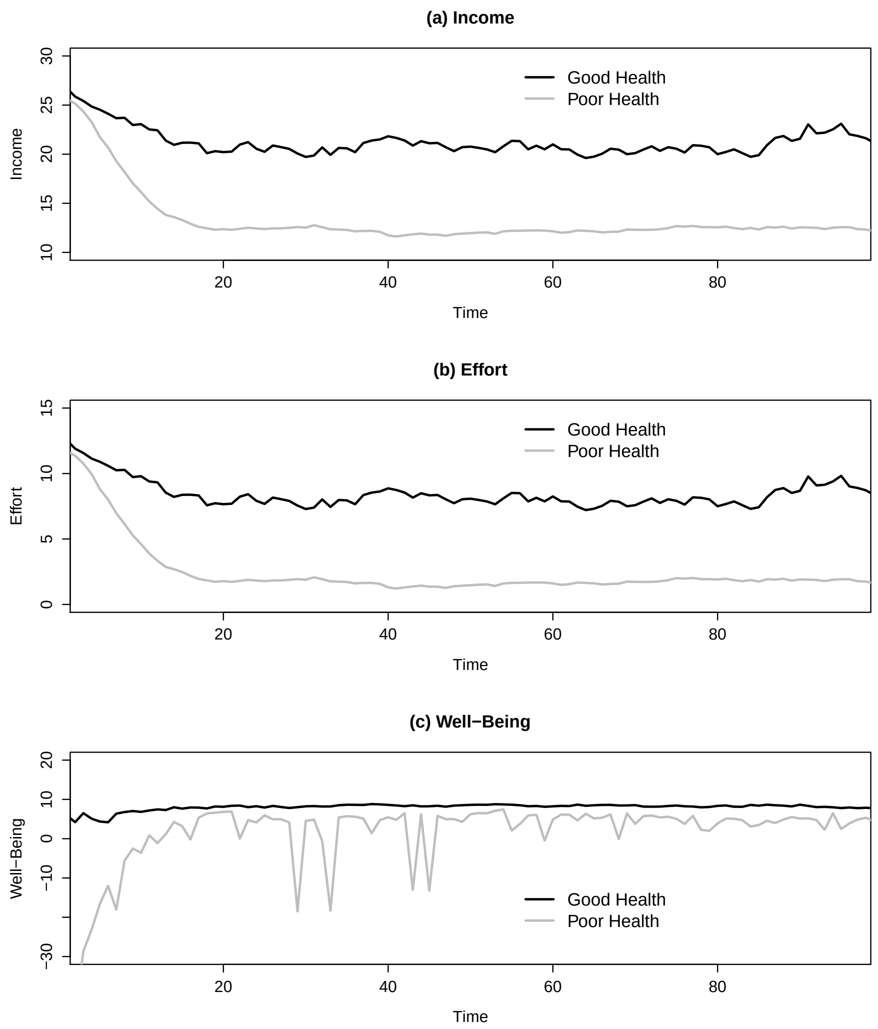

Fig. 6. Time series of three variables of the model for the two health groups: (a) average income; (b) average effort; (c) average well-being. In each panel the black line refers to the good health group ($h_{i,t} \geq 0.5$), while the grey line is associated to the poor health group ($h_{i,t} < 0.5$).

that workers obtain, on average, the same level of well-being of the pensioners.[7] This suggests that, at least according to this simple model of the English welfare state, retirement or other social benefits offer a valuable support to the people outside the labor market. Recalling that the calibration of the model is based on the fitting of the labor market participation curves without any reference to

[7] On the contrary, the pension income y^p is assumed fixed to an arbitrary value in the model simulations.

welfare considerations, the previously mentioned result is unexpected and, for this reason, very powerful.

3. Conclusion

The unsophisticated model described in this paper is able to replicate some relevant stylized facts contained in the ELSA survey, concerning the labor market participation of elderly people.

The core of the model is represented by naive agents moving in a bidimensional torus space, exchanging information on their work and health conditions, and making labor market choices according to this information: although extremely simple, the proposed mechanism of imitation and social learning is able to reproduce an aggregate behavior that can be considered rational and consistent with the empirical evidence, despite the bounded rationality of the individual agents.

It has already been emphasized that the dynamics underlying the model can be interpreted in figurative terms: for instance, the movement in the bidimensional space is a measure of the intensity of the social relationships of each agent. Following this idea, agents that explore less the surrounding space are more likely to meet too often the same peers and fail to be faced with alternative and possibly better decisions regarding retirement. Conversely, more mobile agents have higher chances of seeing in their more widespread network diverse examples to imitate and improving their own well-being.

The model captures the most prominent stylized facts regarding the decline in labor market participation of aged workers, and well fits the participation curves with relatively few parameters. This parsimony is in sharp contrast with the mainstream structural models, which usually require numerous parameters to replicate the empirical evidence (John von Neumann famously stated that "With four parameters I can fit an elephant. And with five I can make him wiggle his trunk.")

The analysis of the paper proves that the main results are robust to changes in the parameter values and to an alternative model specification that bans the transit from retirement to work: the labor market participation curves predicted by the model always exhibit the same patterns of decline when age increases, and in which workers with bad health conditions are more likely to leave the labor market. This situation results in a self-sorting separation equilibrium: in fact, agents affected by bad health decide to become pensioners, receiving a lower income but exerting zero effort and, in this way, they end up enjoying a higher level of well-being, comparable to that of workers with a good health status.

Moreover, it is reasonable to think that more sophisticated forms of interaction between agents, such as the dispositional contagion mechanism assumed in the Epstein's notion of Agent_Zero (Epstein, 2014), would likely deliver an outcome similar to our results.

In conclusion, we believe that the model presented in this paper is a behaviorally plausible alternative to the standard modeling approaches that unrealistically assume the full rationality of individuals, despite an overwhelming psychological and experimental literature showing the limits of human reasoning.

References

Alonso-Ortiz J (2014) Social Security and Retirement Across the OECD. Journal of Economic Dynamics and Control 47: 300–316

Axtell RL, Epstein JM (1999) Coordination in Transient Social Networks: An Agent-Based Computational Model of the Timing of Retirement. In Aaron HJ (ed). Behavioral Dimensions of Retirement Economics, Brookings Institution Press, Washington DC, pp 161–183

Bound J, Stinebrickner TR, Waidmann TA (2010) Health, Economic Resources and the Work Decisions of Older Men. Journal of Econometrics 156: 106–129

Burkhauser RV, Butler JS, Gumus G (2004) Dynamic Programming Model Estimates of Social Security Disability Insurance Application Timing. Journal of Applied Econometrics 19: 671–685

Camerer CF (1997) Progress in Behavioral Game Theory. Journal of Economic Perspectives 11: 167–188

Epstein JM (2014) Agent_Zero: Toward Neurocognitive Foundations for Generative Social Science, Princeton University Press, Princeton

French E (2005) The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour. Review of Economic Studies 72: 395–427

Gilleskie DB (1998) A Dynamic Stochastic Model of Medical Care Use and Work Absence. Econometrica 66: 1–45

Goudet O, Kant JD, Ballot G (2015) Forbidding Fixed Duration Contracts: Unfolding the Opposing Consequences with a Multi-Agent Model of the French Labor Market. In Amblard F, et al (eds). Advances in Artificial Economics, Springer International Publishing, pp 151–167

Gustman AL, Steinmeier TL (2002) Retirement and the Stock Market Bubble. NBER Working Paper No. 9404, National Bureau of Economic Research, Cambridge

Gustman AL, Steinmeier TL (2015) Effects of Social Security Policies on Benefit Claiming, Retirement and Saving. Journal of Public Economics 129: 51-62

Haan P, Prowse V (2014) Longevity, Life-Cycle Behavior and Pension Reform. Journal of Econometrics 178: 582–601

Heyma A (2004) A Structural Dynamic Analysis of Retirement Behaviour in the Netherlands. Journal of Applied Econometrics 19: 739–759

Laun T, Wallenius J (2015) A Life Cycle Model of Health and Retirement: The Case of Swedish Pension Reform. Journal of Public Economics 127: 127–136

Laun T, Wallenius J (2016) Social Insurance and Retirement: A Cross-Country Perspective. Review of Economic Dynamics 22: 72–92

Lettau M (1997) Explaining the Facts with Adaptive Agents: the Case of Mutual Fund Flows. Journal of Economic Dynamics and Control 21: 1117–1147

Lurie NH (2004) Decision Making in Information-Rich Environments: The Role of Information Structure. Journal of Consumer Research 30: 473–486

Malhotra NK (1982) Information Load and Consumer Decision Making. Journal of Consumer Research 8: 419–430

Moro A, Pellizzari P (2016) An Agent-Based Model of Labor Market Participation with Health Shocks. In Demazeau Y, et al (eds.). Advances in Practical Applications of Scalable Multi-agent Systems, The PAAMS Collection, Springer International Publishing, pp 157-168

OECD (2013) How's Life? 2013: Measuring Well-Being, OECD Publishing, Paris

Pelgrin F, St-Amour P (2016) Life Cycle Responses to Health Insurance Status. Journal of Health Economics 49: 76–96

Rabin M (1998) Psychology and Economics. Journal of Economic Literature 36: 11–46

Reyna VF, Brainerd CJ (2007) The Importance of Mathematics in Health and Human Judgment: Numeracy, Risk Communication, and Medical Decision Making. Learning and Individual Differences 17: 147–159

Rust J, Phelan C (1997) How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets. Econometrica 65: 781–831

Wilensky U (1999) NetLogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston http://ccl.northwestern.edu/netlogo/