Wealth inequality under “keeping up with the joneses” preferences

Published version:

DOI:10.1016/j.econlet.2016.06.035

(Article begins on next page)

Wealth inequality under "keeping up with the Joneses" preferences *

Claudio Campanale

Universidad de Alicante

CeRP (Collegio Carlo Alberto)

June 27, 2016

Abstract

In the present paper I introduce "keeping up with the Joneses preferences" in an otherwise standard heterogeneous agent economy. The model simulations show that this kind of preferences can generate a substantial increase in wealth inequality compared to an equal model with standard expected utility.

Keywords: "Keeping up with the Joneses", consumption cascades, heterogeneous agents, wealth inequality

JEL codes: E21, D21

*Claudio Campanale, Departamento de Fundamentos del Análisis Económico, Universidad de Alicante, Campus San Vicente del Raspeig, 03690, Alicante, Spain. Phone: +34 965903614 ext. 3262. E-mail: claudio@ua.es.

I wish to thank Asier Mariscal, Francesco Serti and Francesco Turino for comments. Financial support by the Spanish Ministerio de Econiomía y Competitividad (project ECO2015-70540-P), the Generalitat Valenciana (project PROMETEO/2013/037) and Collegio Carlo Alberto is gratefully acknowledged. All errors and inconsistencies are my own.

1 Introduction

Standard heterogeneous agent models assume that utility depends on the level of personal consumption. Despite this there is a long tradition in economic thinking dating back to Veblen (1899) and Duesenberry (1949) that has recognized that interpersonal comparisons can play an important role in motivating economic actions including consumption and saving behavior. In the present paper I apply this classical insight about the role of interpersonal comparisons to a standard modern quantitative model of saving, that is, the precautionary saving model.

More precisely I propose a version of the classical model of Aiyagari (1994) where I assume that an agent's utility depends not only on her personal consumption but also on an external reference consumption level defined as the average consumption of agents in the earnings group immediately above its own. I find that the model can increase wealth inequality by a substantial amount. The magnitude of the increase depends both on the strength of the external habit motive and on the persistence of the earnings process with a high persistence needed for this magnitude to be large.

A number of recent papers have found empirical support for the old idea that interpersonal comparison of economic outcomes matter. As an example Luttmer (2005) studies the impact of average local income on self-reported well-being and finds a negative relationship as strong as the positive relationship between the latter variable and own income. His results also suggest that the channel is through utility functions that depend on relative consumption in addition to absolute consumption. In a slightly different vein, Bertrand and Morse (2013) and Frank et al. (2014) test empirically the relationship between other's income, as expressed by the top income shares, and different economic choices like consumption of middle-income households or bankruptcy rates at the state or county level. In all cases they find a positive relationship. The cited literature provides an empirical motivation for the present study. The distinctive contribution of this paper consists of extending the study of the role of inter-personal comparisons in consumption, that have so far been studied in empirical or theoretical settings, to a quantitative model and to explore their role in shaping wealth inequality.

2 The Model

I study the steady-state properties of a neoclassical economy with no aggregate uncertainty. The economy is populated by a measure one of infinitely lived households. Households are endowed with a unit of time that they supply inelastically to the labor market. Each period they receive a shock to their efficiency units of labor that I denote with e. I assume that e belongs to a finite set $E = \{e_1, ..., e_n\}$ and that it follows a first-order Markov process that can described with a transition probability matrix π. Households evaluate the utility of a flow of consumption c by using the function:

$$U(c, C) = \frac{c^{1-\alpha} C^{\gamma\alpha}}{1-\alpha} \tag{1}$$

where α is the coefficient of relative risk aversion with respect to own consumption and γ defines the impact on the household's utility of the average consumption level of a reference group. When $\gamma > 0$ any addition to the consumption of the reference group raises the marginal utility of own individual consumption: consumption becomes more valuable since it helps "keeping up with the Joneses". This utility function has been used in Galí (1994) and is also a special case of the function used by Abel (1990). In those articles the reference consumption group was the set of all households in the economy. In the present context with no aggregate uncertainty the average consumption is constant and would wash out of the agents' first order conditions, making it irrelevant. It is thus assumed that for agents endowed with an earnings shock inferior to the top one, the reference group is the set of all agents with the next higher earnings shock. Agents with the highest earnings shocks on the other hand do not have an external reference group. Formally, for agents with earnings shock $i \in \{1, 2, ..., n\}$

$$U_i(c, C) = \frac{c^{1-\alpha} C^{\gamma\alpha}}{1 - \alpha} \tag{2}$$

where $C = 1$ if $i = n$ and $C = \overline{C}(e_{i+1})$ if $i < n$ where $C = \overline{C}(e_{i+1})$ is the average consumption for agents with earnings shock $i + 1$.

The choice to use earnings one step above as the characteristic that defines reference groups seems natural given that the consumption externality plausibly arises as a consequence of exposure to other groups' consumption. In light of this it is reasonable to think that earnings poor households are not likely to be much exposed to the top earners life-styles since they are likely not to live in the same neighborhood or share the same workplace but for the same reasons they are more likely to know the consumption possibilities of other households that earn a bit more than them. Also the choice of earnings rather than wealth or income is convenient because it leads to a straightforward extension of well-known methods to solve the model.[1] Having utility depend on consumption of the earnings group immediately above implies that consumption at the top of the earnings distribution will indirectly affect consumption at all points of that distribution down to the bottom. This idea was termed "consumption cascades"by Frank et al. (2014).

There are no state contingent markets to insure household specific earnings risk. In order to save, the household has access to a single asset that pays interest at a rate r. I denote the amount of assets held by the household by a and I assume that $a \in A \equiv [\underline{a}, \infty)$.

[1]Defining the consumption reference by wealth groups would be problematic from a numerical point of view since it would lead to discontinuities in the value function around the thresholds defining the groups. The same applies to income which depends on wealth through the earned interest component.

Given the preferences and asset structure specified above we directly write the household's optimization problem in dynamic programming form. The state variables of this problem are the shock to its endowment of efficiency units of labor and its assets at the beginning of the period, that is, the pair $\{a, e\}$. The problem thus reads:

$$V(a,e) = \max_{c,a'} \left\{ u(c,C) + \beta \mathbf{E} V(a',e') \right\} \tag{3}$$

subject to the resource constraint

$$a' = ew + (1+r)a - c \tag{4}$$

a forecasting rule for the reference group's consumption

$$C = \mathbf{C}(e) \tag{5}$$

and the no-borrowing constraint

$$a \geq 0 \tag{6}$$

In the resource constraint (4) w is the rental rate for each efficiency unit of labor and r is the rental rate on capital. In the value function equation (3) β is the standard subjective discount factor and $\mathbf{E}$ is the expectation operator.

Aggregate output is produced by a representative firm operating under perfect competition via a standard neoclassical, constant return to scale production function $Y = F(K, L)$, where Y is aggregate output, K is the total amount of capital and L is the total amount of labor used in production. The output can be indifferently used for investment and consumption. Capital depreciates at a constant rate $\delta \in [0, 1]$.

The equilibrium for this economy can be defined in the usual way and is thus omitted for the sake of brevity.[2]

[2]The only slight variation is that we need to insure consistency between the forecasting rule for the consumption of reference groups $\mathbf{C}(e)$ with the actual average consumption of those groups $\overline{\mathbf{C}}(e)$.

3 Parameter Calibration

The model is calibrated taking a year as the length of the period. Technology is defined by a standard Cobb-Douglas production function:

$$Y = K^{\theta} L^{1-\theta} \tag{7}$$

and the capital share θ is set at 0.35. The depreciation rate of capital δ is fixed at 0.06. The Markov chain for the efficiency units of labor is obtained by discretizing an AR(1) process in logarithms:

$$\ln e_t = \rho \ln e_{t-1} + \varepsilon_t \tag{8}$$

where ε is a normal i.i.d. random variable, independently distributed across agents, with mean 0 and variance σ_{ε}^2. For the autocorrelation coefficient I use two values. As a baseline I set $\rho = 0.95$ which is the value estimated by Storesletten et al. (2004). I also explore the quantitative properties of the model under a lower persistence scenario where ρ is set to the value of 0.9. The value of σ_{ε} is set implicitly by fixing the coefficient of variation of earnings. This takes the value of 0.2 based on Aiyagari (1994). With respect to preferences, β is set at 0.96 and the coefficient of relative risk aversion is 1.5, close to the estimate by Attanasio et al. (1999). The remaining parameter to be set is γ which determines the strength of the impact of the reference group's average consumption on the household's marginal utility. Since it was not possible to find estimates in the literature I resort to performing several experiments with different values of γ and in particular I consider values of 0.35, 0.65 and 0.8. Given the functional form of the period utility index, the special case $\gamma = 0$ corresponds to the standard utility case and will be presented for comparison.

4 Results

I explore the behavior of the wealth distribution in the model when the strength of the externality in consumption and the persistence of the earnings process vary and how the effects of γ and ρ interact.

Table 1 and table 2 describe the results. The tables also report the corresponding figures for the data, taken from Campanale (2007) and based on the 1998 issue of the Survey of Consumer Finances. Table 1 reports the value of the Gini index. In the first line of the bottom sub-panel we can see that in the higher earnings persistence case the Gini index of wealth increases from 0.572 in the standard utility model to 0.773 in the expenditure cascades model, corresponding to a 35 percent increase. The increase occurs along the whole range of values of γ and is such that even with an underlying earnings distribution that is more equal than the one in the data the model can generate the same Gini index of wealth found in the data.[3] The consumption externality thus turns out to be a powerful mechanism to increase wealth inequality. The second line of the panel reports the effect of γ on wealth inequality when earnings are less persistent. The Gini index increases, thus confirming the previous result qualitatively. At the quantitative level the increase is more modest: from 0.494 in the standard utility case to 0.550 in the consumption cascades model, an increase of 11.3 percent. The comparison of the results in the two panels suggests that while the concern for relative consumption may be a powerful factor to increase wealth inequality, a high persistence of earnings plays a crucial role in magnifying this factor.[4]

[3]According to Campanale (2007) the Gini index for earnings is 0.63 in the data. In the model it is 0.2 by construction.

[4]The intuition for this result is that in the present model the earnings shock affects not only the budget constraint but also the utility of current consumption, adding an extra degree of uncertainty to the standard Aiyagari (1994) model. This increased uncertainty leads to higher wealth inequality. In turn this extra uncertainty, hence the increase in wealth inequality, becomes larger when the consumption distribution

Table 1: **Gini index of wealth: data and selected values of ρ and γ**

Data				
	0.78			
Model	$\gamma = 0$	$\gamma = 0.35$	$\gamma = 0.65$	$\gamma = 0.8$
$\rho = 0.95$	0.572	0.608	0.700	0.773
$\rho = 0.9$	0.494	0.502	0.518	0.550

Table 2: **The distribution of wealth for selected values of ρ and γ: percentage share of wealth by percentiles**

Percentiles	bottom 40	top 20	top 10	top 5	top 1
Data	1.35	79.5	66.1	53.5	29.5
$\rho = 0.95$					
$\gamma = 0$	4.1	57.1	35.6	21.0	5.5
$\gamma = 0.35$	3.5	61.6	40.0	24.6	7.0
$\gamma = 0.65$	1.4	72.3	49.9	32.3	10.4
$\gamma = 0.80$	0.2	80.9	58.9	40.2	14.8
$\rho = 0.9$					
$\gamma = 0$	8.3	50.6	30.9	18.0	4.7
$\gamma = 0.35$	8.0	51.5	31.6	18.5	4.9
$\gamma = 0.65$	7.4	53.1	32.9	19.5	5.3
$\gamma = 0.8$	6.0	56.0	35.7	21.7	6.3

A more detailed picture of wealth inequality is given in table 2. The table is organized in three sub-panels. In the top sub-panel I report the share of wealth held by different percentiles of the distribution in the data. The middle sub-panel reports the same figures for the model with high earnings persistence and the bottom sub-panel does the same for the model with low earnings persistence. Focusing on the second sub-panel we can see that the share of wealth held by the bottom 40 percent of the distribution falls from 4.1 percent in the standard model to 1.4 percent in the model with γ equal to 0.65 and to only 0.2 percent in the model with γ equal to 0.8. Moving to the other columns of the panel we see that the share of wealth of the top 20 percent increases from 57.1 percent to 80.9 percent, the share of the top 5 percent almost doubles from 21 percent to 40.2 percent and the share of the top 1 percent almost triples from 5.5 percent to 14.8 percent when we move from the model with standard utility to the model with the highest value of the external habit parameter. The figure for the top 1 percent is still only half its empirical counterpart, however as it was mentioned before the model earnings concentration falls quite short of the data.

In the case with ρ equal to 0.9 the introduction of external habit can still reduce the share of wealth of the bottom 40 percent and increase the share held by the top groups but the effects are quantitatively smaller: as examples, the third column shows that the share of the top 20 percent increases from 50.6 to 56 percent and the sixth column that the share of the top 1 percent increases from 4.7 percent to 6.3 percent when we move from γ equal to 0 to γ equal to 0.8. These results thus confirm that introducing external habits may have powerful effects in increasing wealth inequality when coupled with substantial earnings persistence.[5]

becomes more dispersed which is what happens in this class of models when earnings persistence increases.

[5]I performed similar simulations with different values of risk-aversion and of σ_ε. The results were similar and are thus not reported for the sake of brevity.

5 Conclusions

In the present paper I have introduced “keeping up with the Joneses ”preferences in a Bewley type of model that is a workhorse of modern macroeconomics. It has been shown that the model can generate a sizeable increase in wealth inequality provided that the strength of the external habit is relatively high and is combined with substantial persistence in the earnings process. This result suggests that the introduction of external habits may have important effects in the predictions of standard heterogeneous agent models. This warrants further research in the area. A natural extension would be to recast the model in a life-cycle framework with bequests and a more elaborate and realistic earnings process and study the effect of these preferences in this richer setting.

References

[1] Abel A.B., 1990. Asset prices under habit formation and catching up with the Joneses. American Economic Review. 80 (2). 38-42.

[2] Aiyagari, S.R., 1994. Uninsured idiosyncratic risk and aggregate savings. Quarterly Journal of Economics. 109. 659-684.

[3] Attanasio O.P., Banks J., Meghir C., Weber W. 1999. Humps and bumps in lifetime consumption. Journal of Business and Economics Statistics. 17 (January). 22-35.

[4] Bertrand M., Morse A. 2013. Trickle-down consumption. NBER working paper, 18883.

[5] Campanale C. 2007. Increasing returns to savings and wealth inequality. Review of Economic Dynamics. 10. 646-675.

[6] Duesenberry J.S. 1949. Income, savings and the theory of consumer behavior. Cambridge, MA. Harvard University Press.

[7] Frank R.H., Seth Levine A., Dijk, O. 2014. Expenditure cascades. Review of behavioral economics. 1 (1-2). 55-73.

[8] Gali J. 1994. Keeping up with the Joneses: consumption externalities, portfolio choice and asset prices. Journal of Money, Credit and Banking. 26 (1). 1-8.

[9] Luttmer E.F.P. 2005. Neighbors as negatives: relative earnings and well-being. Quarterly Journal of Economics. 120 (3). 963-1002.

[10] Storesletten K., Telmer C., Yaron A. 2007. Asset pricing with idiosyncratic risk and overlapping generations. Review of Economic Dynamics. 10 (4). 519-548.

[11] Veblen T. 1899. The theory of the leisure class. New York, NY. Macmillan.