Explosive earnings dynamics: Whoever has will be given more

Sarah Meyer and Mark Trede
Universität Münster

May 19, 2016

Abstract

This paper suggests a model of explosive earnings dynamics where positive deviations tend to increase the growth rate even further. This "Matthew effect" can explain a number of empirical regularities. First, we show that the explosive model might resemble a model with heterogeneous earnings profiles in terms of its covariance structure. Second, we derive the optimal consumption and savings behaviour under explosiveness and compare it to other models. Third, we present a panel test against explosiveness and apply it to German and U.S. earnings data. We find that the null hypothesis of no explosiveness can be rejected. However, the proportion of explosive profiles is small.

JEL codes: J31, D91, C33
Keywords: labour income, idiosyncratic risk, explosive stochastic processes

1 Introduction

Income and earnings dynamics receive a lot of attention not only in the economics literature but also in public debate. There is a widespread concern internationally about the rise in earnings inequality since the 1980s. In real terms, workers at the bottom of the distribution have lost in many countries during the last decades, while we witnessed a large increase in the share of earnings accruing to the top decile and, even more pronounced, the top percentile, especially in English-speaking countries (Atkinson, Piketty and Saez 2011). Earnings risk is an important part of the overall economic risk facing individuals and households. Although earnings risk is to a large extent idiosyncratic, it is commonly considered as uninsurable due to the obvious moral hazard problem.

There are two main approaches to modelling earnings dynamics, and both have a long tradition. On the one hand, earnings are regarded as following a deterministic path over the life-cycle which is disturbed by random shocks with weak or moderate persistence. Since different individuals are assumed to have different expected earnings profiles, this class of models is often termed "heterogeneous income profiles" (Guvenen 2009). Evidence in favour of the heterogeneous model is found e.g. in Lillard and Weiss (1979), Baker (1997), Haider (2001), Guvenen (2009), Browning, Ejrnæs and Alvarez (2010).

On the other hand, many studies consider earnings shocks as having a unit root, e.g. MaCurdy (1982), Abowd and Card (1989), Baker and Solon (2003), Meghir and Pistaferri (2004), Hryshko (2012). Random changes are not only transitory (although transitory shocks may of course also be present) but have a permanent effect for the rest of the working life. Heterogeneity of deterministic growth rates is often ruled out in these models (Baker and Solon (2003) and Hryshko (2012) embed both heterogeneity and unit roots in their models). Due to the restricted heterogeneity (Guvenen 2009) labels these models as "restricted income profiles". It is surprisingly difficult to distinguish clearly between the heterogeneous and the unit root model on a statistical basis (Ejrnæs and Browning 2014). Despite the fact that earnings dynamics has been investigated in a huge number of papers, no consensus has been reached yet as to the preferable model specification.

This paper suggests a new simple model: explosive earnings dynamics. Positive deviations tend to boost the growth rate such that the deviation from a common trend will get even larger. In the same way, negative shocks will let the growth rate drop. In sociology, such self-reinforcing inequalities where the rich get richer and the poor get poorer are sometimes called "Matthew effect" (after Matthew 13:12: "Whoever has will be given more, and they will have an abundance. Whoever does not have, even what they have will be taken from them"). Explosiveness is ruled out in previous studies on earnings dynamics even though the extreme divergence of some earnings profiles has been noted before (Guvenen, Karahan, Ozkan and Song 2015). From a formal point of view, we relax the restriction that the persistence parameter must not be larger than unity but keep the restriction that there is no deterministic growth rate heterogeneity.

We investigate the properties of the explosive model and show that it combines many features of the heterogeneous and the unit root models. While long run earnings predictions are subject to large and ever increasing uncertainty as in the unit root case, heterogeneity in expected earnings profiles emerges after some initial shocks as in the heterogeneous case. The explosive model can mimic the variances and covariances of the heterogeneous model closely, both in earnings levels and differences. Hence, observations that seem to corroborate the heterogeneous model may in fact stem from an explosive process.

An important implication of the stochastic properties of the earnings process is the optimal consumption and savings behaviour. It is well known (Zeldes 1989) that no closed form solution of this stochastic dynamic optimization problem is available. Instead we solve it numerically in a simplified setting that allows to compare the three models. If earnings follow an explosive process, the optimal consumption and savings behaviour is similar to the heterogeneous model in certain situations (in particular, if the individual happens to experience a series of positive shocks early in their working life), but more resembles the unit root case if the starting periods are less favourable. On the micro level, some individuals will behave as if they are liquidity constrained, while others follow the classic pattern of building up wealth when young in order to smooth consumption over their life cycle. How an individual behaves does not depend on preferences (which we assume to be identical for all) but on the realizations of their earnings process.

We suggest to use a panel unit root test against the alternative of explosiveness to distinguish between the earnings models. Statistical tests against explosiveness exist for univariate time series (Phillips, Wu and Yu 2011) but have not been applied to earnings or other panel data. Gustavsson and Österholm (2014) investigate the issue of unit roots in a panel data set but estimate the persistence parameter separately for each individual. Following a suggestion by Hanck (2013) we construct a panel unit root test based on Simes' (1986) intersection test, and apply the test procedure to earnings data from the cross-national equivalent files of the German Socio-Economic Panel (GSOEP) and the U.S. Panel Study of Income Dynamics (PSID) data sets. We find that the null hypothesis of stationarity or unit roots can be rejected in both countries. Statistically significant explosiveness can be found in the data at least for a fraction of the population.

A well specified statistical models of earnings dynamics is of interest for economists and policy makers for several reasons. An appropriate description of consumption behaviour has to be based on a correctly specified earnings model. How individuals respond to variations in their earnings depends to a large degree on the persistence of shocks and whether there are deterministic components. The more uncertain future earnings, the higher the incentive to build a wealth buffer stock. Evidence about the correct specification of earnings models can also help validate economic theories that endogenize earnings such as Mincer's (1958) classical human capital theory. From a policy point of view, it is important to understand the nature of earnings profiles. If shocks are permanent – and not insurable – policy makers should aim to introduce or facilitate risk sharing mechanisms. On the other hand, if the evolution of earnings is mainly deterministic an obvious policy response is to improve education in order to shift workers who would end up on low-performing trajectories onto higher profiles.

The paper is organized as follows. In section 2 we introduce the explosive model and compare its statistical properties to the case of heterogeneous earnings profiles with transitory shocks and to the unit root model with both transitory and permanent shocks. Section 3 is concerned with the optimal savings and consumption behaviour of individuals facing either of the three models. In section 4 we suggest a panel test against explosiveness. The empirical applications are reported in section 5. We apply the panel test to GSOEP and PSID earnings data. Section 6 concludes.

2 Models of earnings dynamics

Let $y^i_{h,t}$ denote person i's residual of a regression of log earnings in period t on a cubic polynomial in potential experience h^i and dummy variables for the level of education (with levels: less than high school/high school/more than high school). This is the usual way of eliminating common effects that influence all individuals in period t (Guvenen 2009). Since the regressions are run separately for each period, they also capture other time-specific variations in the labour market such as increasing returns to education (Autor 2014). In the following, $y^i_{h,t}$ will be referred to simply as "earnings".

Earnings are modelled as

$$y^i_{h,t} = \alpha^i + \beta^i h + z^i_{h,t} + \varepsilon^i_{h,t} \tag{1}$$

where α^i and β^i are individual specific random effects with zero expectations, variances σ^2_α and σ^2_β and covariance $\sigma_{\alpha,\beta}$. The random shock $\varepsilon^i_{h,t}$ is assumed to be homoscedastic and independent of α^i and β^i. It represents the short-term, transitory earnings shocks. Shocks that are longer lasting are modelled by

$$z^i_{h,t} = \rho^i z^i_{h-1,t-1} + \eta^i_{h,t} \tag{2}$$

where $\eta^i_{h,t}$ is a (homoscedastic) random shock. The initial value is set to $z^i_{0,t} = 0$. In contrast to the literature we do not impose any restrictions on the AR coefficient ρ^i (apart from regarding negative values as implausible). In particular, we allow ρ^i to be larger than unity.

In the literature on earnings dynamics, two special cases are commonly considered. First, the random slope effects are set to constant zero (where $\sigma_\beta^2 = 0$) such that expected earnings are initially the same for every individual. Earnings differences are entirely due to random and persistent shocks affecting the individuals during their working life. In this setting, the process $z_{h,t}^i$ is modelled as a unit root process, i.e. $\rho^i = 1$.

The second special case allows for heterogeneity in earnings profiles ($\sigma_\beta^2 > 0$) and restricts the persistence parameter to $\rho^i \leq 1$, often strictly so. There is no consensus yet in the literature as to which model describes real-world earnings processes better.

The explosive case $\rho^i > 1$ has not been considered in the literature on earnings dynamics even though it exhibits a number of attractive properties. We proceed to show that the common restriction $\rho^i \leq 1$ is in fact unnecessary. To streamline notation, the indices t for time and i for individuals will be dropped in the following.

The variances of the persistent component z_h are

$$Var(z_h) = \sigma_\eta^2 + \rho^2\sigma_\eta^2 + \ldots + \rho^{2(h-1)}\sigma_\eta^2 + \rho^{2h}Var(z_0). \tag{3}$$

For $Var(z_0) = 0$ and $\rho \neq 1$, (3) can be simplified to

$$Var(z_h) = \sigma_\eta^2 \cdot \frac{1-\rho^{2h}}{1-\rho^2}. \tag{4}$$

For $\rho = 1$ the variance (4) is to be interpreted as the limit for $\rho \to 1$ which is simply $t\sigma_\eta^2$. The covariance between z_h and z_{h+n} for $n \geq 1$ is

$$Cov(z_h, z_{h+n}) = \rho^n Var(z_h).$$

Hence, the variance-covariance structure of earnings is

$$\begin{aligned} Var(y_h) &= \sigma_\alpha^2 + t^2\sigma_\beta^2 + 2h\sigma_{\alpha,\beta} + \sigma_\eta^2\frac{1-\rho^{2h}}{1-\rho^2} + \sigma_\varepsilon^2 & (5)\\ Cov(y_h, y_{h+n}) &= \sigma_\alpha^2 + h(h+n)\sigma_\beta^2 + (2h+n)\sigma_{\alpha,\beta} + \rho^n\sigma_\eta^2\frac{1-\rho^{2h}}{1-\rho^2}. & (6) \end{aligned}$$

In terms of growth rates the covariance between Δy_h and Δy_{h+n} for $n \leq 2$ is

$$Cov(\Delta y_h, \Delta y_{h+n}) = \sigma_\beta^2 + \sigma_\eta^2\left[\rho^{n-1}\frac{\rho-1}{\rho+1}(1+\rho^{2h-1})\right]. \tag{7}$$

The results (5), (6) and (7) do not require any restriction on the persistence parameter ρ. If $|\rho| < 1$ and h is sufficiently large, then the final factor in round brackets is close to 1, an approximation that is often used in the heterogeneous model (MaCurdy 1982, Guvenen 2009).

Of course, if the transitory component of earnings is modelled in a more sophisticated way or if measurement error is added, then (5) and (6) have to be modified accordingly.

Figure 1 shows the variance of earnings for $\rho = 0.8, 1, 1.03$ as a function of experience h if there is no heterogeneity, $\sigma_\beta^2 = \sigma_\alpha^2 = 0$. The other parameters are $\sigma_\eta^2 = 0.1^2$ and $\sigma_\varepsilon^2 = 0.25^2$. If there is only weak persistence ($\rho = 0.8$) the cross-sectional variance of earnings converges quickly to a constant. In the case of a unit root the variance increases linearly in experience. If the persistence parameter is greater than 1, the variance is a convex function of experience. The dotted red line shows the cross-sectional variance for heterogeneity in earnings profiles while the persistence parameter is set to 0. The heterogeneity parameters $\sigma_\alpha^2, \sigma_\beta^2$ and $\sigma_{\alpha,\beta}$ have been chosen such that the fit is close. Concerning covariances, it is of course also possible to find a parametrization of the heterogeneous model that results in covariances (6) resembling those of the explosive model.

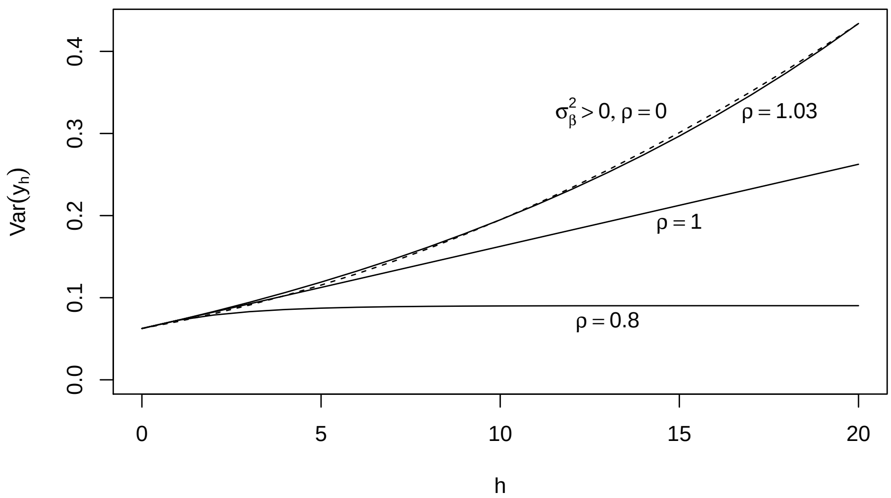

Figure 1: Variance of earnings as a function of experience

The graph indicates that explosive earnings profiles and heterogeneous earnings profiles are alike, at least in some respects. Individuals who are moving along an earnings profile above (below) the average tend to grow faster (slower). The main difference between the heterogeneous and the explosive model is the role played by risk. In the explosive model it is a matter of good luck to be shifted onto a fast growing profile at the start of your career. A stroke of bad luck may still push you back below average, unless you are already on a comfortably high profile. In the heterogeneous model, chance just once determines the growth rate β which then remains constant during your life-time.

The different role of risk in the heterogeneous and explosive model is also apparent in conditional earnings distributions. Let the information set I_h include the variables α, β and z_h. For $n \geq 1$ (and any ρ) the conditional distribution of earnings is normal with expectation

$$\begin{aligned} E(y_{h+n}|I_h) &= \alpha + \beta(h+n) + E(z_{h+n}|I_h) \\ &= \alpha + \beta(h+n) + \rho^n z_h \end{aligned}$$

and variance

$$\begin{aligned} Var(y_{h+n}|I_h) &= Var(\alpha + \beta h + z_{h+n} + \varepsilon_{h+n}|I_h) \\ &= Var(z_{h+n}|I_h) + \sigma_\varepsilon^2 \\ &= \sigma_\eta^2 \cdot \frac{\rho^{2n}-1}{\rho^2-1} + \sigma_\varepsilon^2. \end{aligned} \tag{8}$$

The most important difference between the explosive and the heterogeneous model is evident in (8). Under heterogeneity, where $\rho < 1$ and $\sigma_\beta^2 > 0$, the long-run risk converges to σ_ε^2 and, hence, does not increase without bounds when the forecasting horizon is extended. In the explosive regime, where $\rho > 1$ and $\sigma_\beta^2 = 0$, the risk is ever increasing in a convex fashion in the length of the forecasting horizon n. Risk is also increasing in the unit root case ($\rho = 1$), but only linearly.

To illustrate the large effect of long-term risk, Figure 2 depicts the probability to fall under the (arbitrary) poverty line -0.4 as a function of the initial value z_t. The dotted-dashed line indicates the poverty probability in the explosive regime without heterogeneity. If you start near the poverty line there is a high risk to be poor in 5 years, whereas poverty is very unlikely for individuals starting at the higher end. The dashed line shows that the poverty risk in the unit root case is close to the explosive case. The solid line depicts the povery risk under heterogeneity for an individual with a positive slope coefficient of $\beta = 0.02$ (and intercept $\alpha = 0$) when the persistence parameter is $\rho = 0.8$. Since the rise in earnings is deterministic, the long-run risk of poverty is low and does not much depend on the initial position. Of course, the picture would be different for an individual with a negative slope who would end up in poverty after a sufficiently long time almost with certainty.

The different roles played by earnings risk have an impact on the optimal behaviour of individuals. In the absence of insurance opportunities for idiosyncratic earnings risk, the larger the risk the more individuals have to save in order to build up a wealth buffer that can protect them against strokes of bad luck. In the next section we derive the optimal dynamic savings behaviour

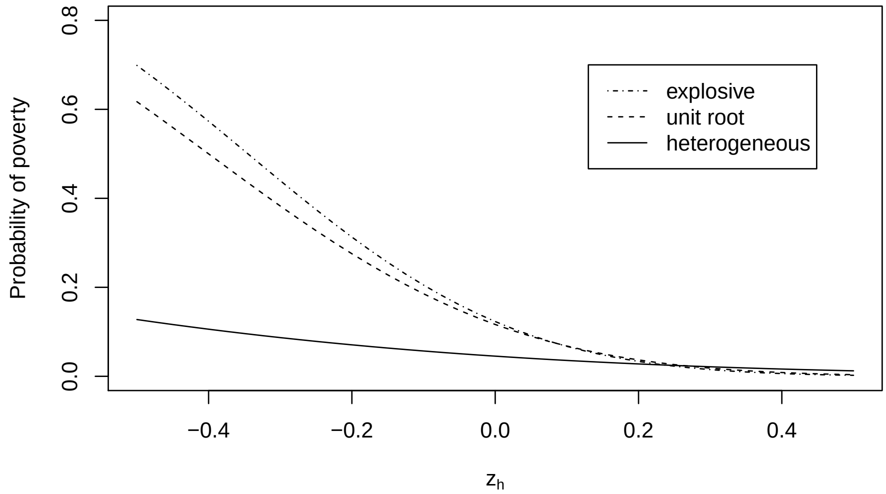

Figure 2: Poverty risk in 5 years conditional on z_h. The plot shows the probability to fall below the poverty line -0.4 given that current earnings is z_h, the solid line is for the heterogeneous regime, the dashed line for the unit root case, and the dotted-dashed line for the explosive regime.

in a simplified setting and investigate the micro and macro-economic impacts of different model assumptions.

3 Optimal savings behaviour

Since future earnings are a stochastic process, the individuals face a stochastic dynamic optimisation problem. Their saving and consumption behaviour depends on the structure of the earnings process. Without the capability to insure against idiosyncratic earnings shocks, individuals have an incentive to save and build up a wealth buffer stock as a means of self-insurance. Theoretical foundations of buffer stock savings models are provided by Carroll (2011). Gourinchas and Parker (2002) model the consumers' optimization problem under a random walk (with drift but no heterogeneity) assumption for permanent earnings. They show that young individuals mainly save out of the precautionary motive. The main motive switches at an age of around 40, and older households mostly save for retirement and bequests. Guvenen and Smith (2014) enrich the model by a more complex retirement structure, partial insurance and heterogeneous earnings profiles with learning about the initially uncertain individual slope parameter.

We derive and compare the optimal savings behaviour in both the model with heterogeneity and short-term persistence ($\rho < 1$) and the explosive model ($\rho > 1$). The unit root case is considered as the intermediate case ($\rho = 1$). In order to delineate the main effects we abstract the earnings dynamics from other features that are relevant for the savings behaviour in the real world.

Let $U(c) = c^{1-\gamma}/(1-\gamma)$ denote the instantaneous utility function with coefficient of relative risk aversion γ. The Bellman equation for an individual facing uncertain future earnings is

$$V_t(W_t, z_t) = \max_{C_t} \{U(C_t) + E_t(V_{t+1}(W_{t+1}, z_{t+1}))\} \tag{9}$$

where the state variable W_t is the individual's wealth, and the state variable z_t is the earnings shock with dynamics given in (2). Consumption C_t is a control variable, and $E_t(\cdot)$ is the expectation conditional on the information in period t. The transition equation for wealth is

$$W_{t+1} = W_t + e^{y_t} - C_t. \tag{10}$$

Earnings dynamics are modelled as

$$\begin{aligned} y_{t+1} &= \beta(t+1) + z_{t+1} \\ z_{t+1} &= \rho z_t + \eta_{t+1}. \end{aligned}$$

Note that there is no random effect intercept α. Neither are there measurement errors nor shocks with short-term effects (the variance of ε_t in (1) is set to 0). Further, we assume that the expectation of the future (log) earnings path is normalized to $E(y_{t+n}) = 0$ for all $n > 0$.

Explosiveness is an unrealistic assumption in infinite horizon models. Hence, we assume that the working life of the individual is finite and ends in period T. We model the impact of the retirement period after the working life very crudely by assuming that the last period's value function is

$$V_T(W_T, Y_T) = R \cdot U((W_T + Y_T)/R), \tag{11}$$

i.e. the individuals distribute their final cash-at-hands equally over their remaining R years of life without leaving any bequests and without any remaining debt. The length of the retirement period is assumed to be deterministic and known.

Since we focus on the self-insurance motive of savings, we dropped both a subjective discount factor (which would appear in front of the conditional expectation in the Bellman equation) and an interest rate (which would appear in the budget constraint (10)). A discount factor, representing impatience, would tend to shift consumption towards the early periods in life. A positive interest rate would give an incentive to postpone consumption in order to earn capital

	Earnings dynamics model		
Parameter	Heterogeneity	Unit root	Explosive
ρ	0.80	1	1.02
σ_β	0.02	0	0
σ_η	0.1	0.129	0.083
γ	2	2	2
T	40	40	40
R	15	15	15

Table 1: Parametrization of the three models of earnings dynamics

income in addition to labour income. Both aspects distract from the main focus of our analysis and are therefore cut out. Individuals with CRRA instantaneous utility and an earnings process that might results with positive probability in earnings arbitrarily close to zero for the rest of their working life, will decide not to get into debt (Zeldes 1989, Gourinchas and Parker 2002). Hence, it is not necessary to explicitly model any liquidity constraints.

The optimal consumption behaviour cannot be derived in closed form but has to be computed numerically.[1] Table 1 shows the parametrizations of the three models (heterogeneity, unit root, explosive). In order to isolate the effects of model assumptions on consumption behaviour the parameters have to be chosen such that the earnings paths are comparable in terms of their riskiness. The parameters are chosen such that the unconditional variance $Var(y_T)$ is about the same in all models. Of course, the unconditional expected earnings paths are constantly zero in all models. The conditional variances, given either β or some realized earnings at the start of the working life, are different. So are the unconditional variances for $t < T$.

The numerical solution of the dynamic optimization problem yields a policy function $C_t^*(W_t, z_t)$ returning the optimal level of consumption in period t given current wealth W_t and earnings y_t (or, equivalently, z_t). In the heterogeneous model, the value of β is assumed to be known to the individual from the outset, see Guvenen (2007) for a heterogeneous model where individuals do not know their β but learn about it by observing their own earnings path. The policy function $C_t^*(\cdot, \cdot)$ and the earnings dynamics jointly determine the evolution of wealth and consumption both for an individual and the total cohort (or population). The individual earnings, consumption and wealth profiles are of course subject to random shocks (caused by η_t). For a large

[1] See the appendix for details about the numerical solution.

cohort, average earnings, consumption and wealth are close to their expected paths since the random shocks are independent from each other and thus average out.

Figure 3 depicts the average wealth paths for a cohort of size $N = 5000$ for the three models. The solid line shows the evolution for the heterogeneous model, the dashed line for the unit root model, and the dash-dotted line for the explosive model. Under all regimes, the wealth path is strictly increasing over the entire working life and reaches its highest level when the retirement period starts. The average level of wealth is highest in the unit root model, and lowest in the explosive model.

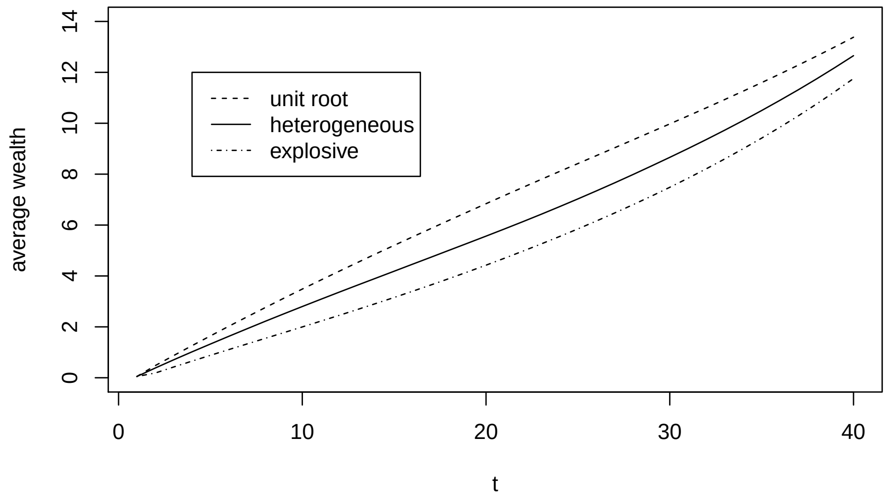

Figure 3: Average wealth paths for the heterogeneous model (solid line), the unit root model (dashed), and the explosive model (dot-dashed)

The aggregate wealth paths do not reveal the substantial differences between individuals that are due to different shock histories or, in the case of the heterogeneous model, to different slopes in expected earnings. Figure 4 shows individual earnings, consumption and wealth paths for each model. For comparability, earnings are not shown in logarithms here. The left panels show "unfortunate" individuals with negative slope or a series of negative shocks early in their working life. The panels on the right show the corresponding profiles for "lucky" individuals.

The plots show that there is considerable variation in behaviour. Individuals with a negative slope β (top left) know that their life-time earnings are low, they start saving immediately to keep their level of consumption more or less constant. When they enter retirement, they have a wealth stock that allows them to keep up their consumption over the retirement period. Individuals

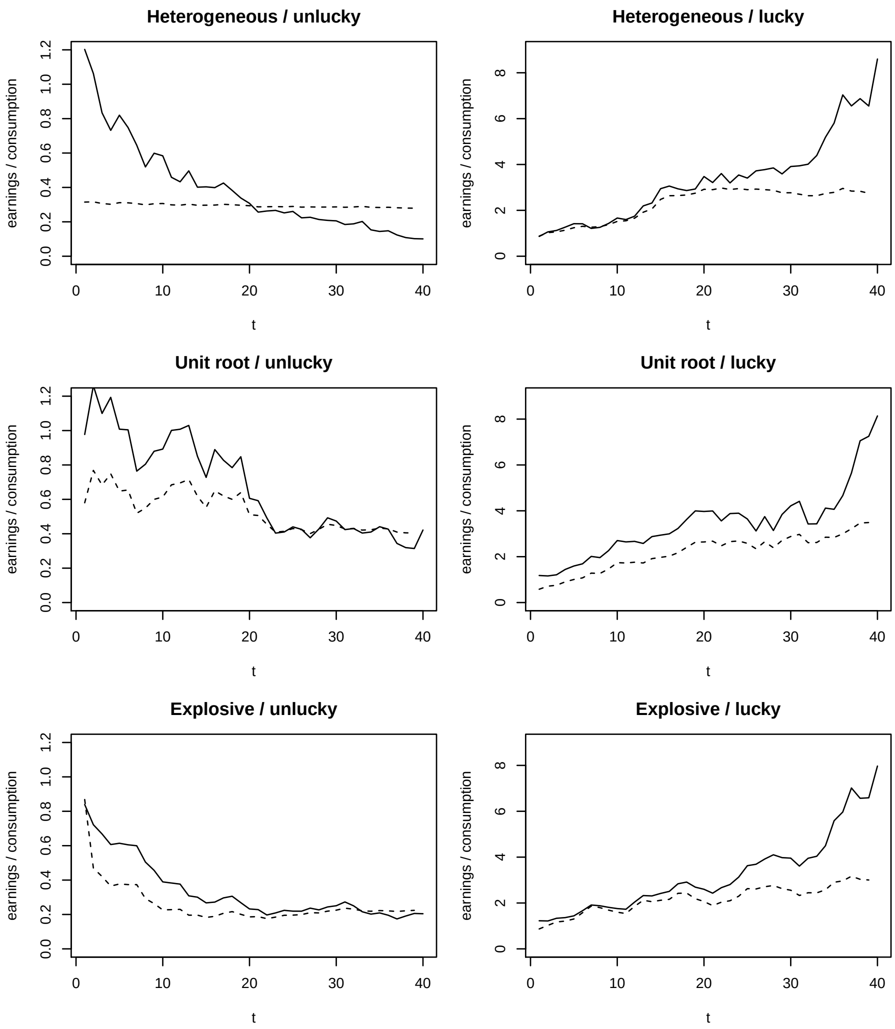

Figure 4: Individual earnings, consumption and wealth paths. The left panels show earnings (solid lines) and consumption (dashed) for "unlucky" individuals, the right panels show "lucky" individuals; the heterogeneous model is shown in the top row, the unit root model in the middle, and the explosive model in the bottom row.

with a positive slope (top right) behave as if they are liquidity constrained. Their consumption is almost identical to earnings in the first half of their working life. Only then do they start to save for retirement.

In the unit root case, the individuals face a large amount of earnings uncertainty. As a consequence, individuals are not able to keep their consumption constant. They will react relatively strongly to positive or negative shocks since each shock as a large impact on life-time earnings. Both lucky and unlucky individuals (middle row, right and left) have an incentive to save early in their working life to build up a wealth buffer.

In the explosive case, individuals behave similarly to the unit root case when they face negative shocks. In contrast, a series of positive shocks will shift them onto a higher growth rate path and hence decreases uncertainty, at least in the sense that future poverty becomes much less likely. Therefore, as in the heterogeneous model, lucky individuals tend to behave almost as if they are liquidity constrained.

Of course, the paths shown in figure 4 are just randomly drawn examples of lucky and unlucky individuals. In order to compare how the distribution of consumption and wealth differs between the three models we simulate a large number ($N = 5000$) of earnings paths over the entire life cycle for each model (heterogeneous, unit root, and explosive) and compute the implied N individual consumption and wealth paths.

In a stationary equilibrium, each cohort is populated by a fraction of $1/T$ of the total population. Hence, pooling all $(N \times T)$ simulated consumption values gives the time-invariant cross-sectional distribution of consumption, and similarly for wealth. Figure 5 depicts the Lorenz curves of consumption and wealth under each model. The consumption Lorenz curve of the heterogeneous model has least inequality, the unit root model implies the highest level of inequality. The Lorenz curve of the explosive model is in between but much closer to the heterogeneous model than to the unit root case.

As to wealth, there is no clear ordering since the Lorenz curves intersect. At the bottom part of the wealth distribution, inequality is highest under the heterogeneous model and almost identical under the other two models. For the richer part of the population, the heterogeneous model implies the least inequality whereas, again, the unit root model implies the highest level of inequality.

The results of this section suggest that the explosive model is not just a still riskier version of the unit root model. On the contrary, a persistence parameter $\rho > 1$ implies micro and macro effects that are similar to the heterogeneous model. We conclude that an empirical investigation of earnings dynamics should take into account the possibility that $\rho \leq 1$. In the next section we suggest to use panel unit root tests against explosiveness as a relatively robust tool to determine the suitable earnings dynamics specification.

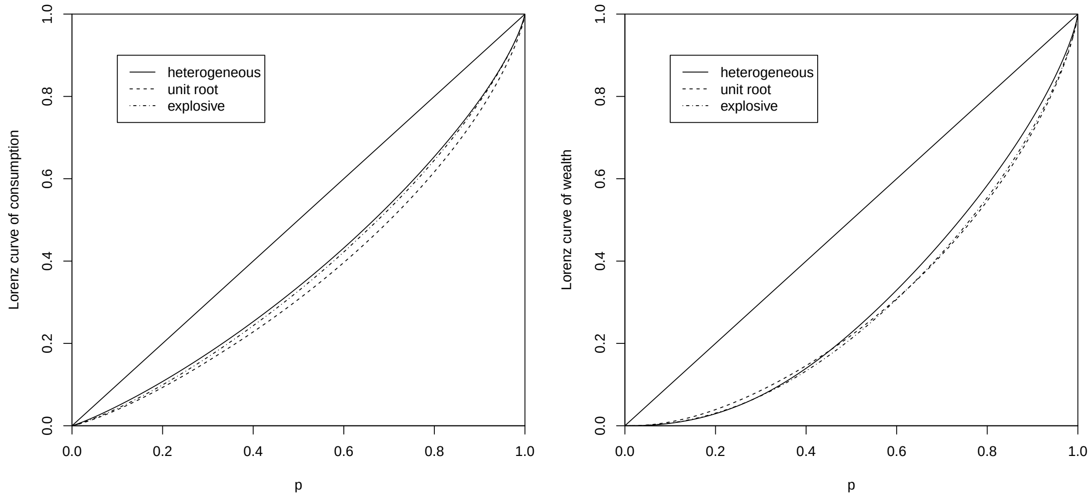

Figure 5: Cross-sectional Lorenz curves of consumption (left) and wealth (right)

4 Panel tests for explosiveness

The application of unit root tests to panel data has attracted much attention in recent years. The literature distinguishes two "generations" of panel unit root tests. First generation tests are based on the assumption that individual time series are cross-sectionally independent. For instance, Maddala and Wu (1999) applied a method of aggregating individual tests, which was originally suggested by Fisher (1925). Doing so, they tested the joint null hypothesis that all individual processes have a unit root against the alternative that at least one process is stationary. The null hypothesis is rejected at level α if the test statistic, which combines the p-values of N ADF tests, is larger than a given critical value. The critical value was shown to be the $(1-\alpha)$-quantile of the χ^2 distribution with $2N$ degrees of freedom. However, this aggregation method is restricted to test statistics (and hence p-values) which are cross-sectionally independent. Further first generation panel unit root tests were suggested by Levin, Lin and James Chu (2002) and Im, Pesaran and Shin (2003).

For economic applications it is rather inappropriate to assume cross section units to be independent as they are often contemporaneously correlated for a variety of reasons such as common factors or spatial spillover effects. Numerous panel unit root tests have been developed that allow for different forms of cross section dependence, such as Chang (2002), Phillips and Sul (2003), Bai and Ng (2004), Breitung and Das (2005) and Moon and Perron (2004).

Chang (2002) proposes a test based on a nonlinear instrumental variable estimation of the common augmented Dickey-Fuller regression. With nonlinear transformations of lagged levels used as instruments, she shows that individual ADF statistics are asymptotically independent. The test statistic is defined as a standardized sum of the individual IV t-ratios. However, her test was shown (see Im and Pesaran (2004)) to be valid only if the number of cross section units N is fixed as $T \to \infty$.

Phillips and Sul (2003), Moon and Perron (2004) and Bai and Ng (2004) approach the problem in a similar fashion. They make use of a residual factor model which takes into account cross section dependence. Phillips and Sul (2003) suggest an orthogonalization procedure which may asymptotically remove the common factors. Standard panel tests can then be applied to the transformed series. In line with this, Moon and Perron (2004) suggest to first de-factor the panel data, which is accomplished by means of a principal components estimation of the factor loadings. They show that their proposed test has good asymptotic power properties if there are no deterministic trends. While both previous approaches only allow the idiosyncratic components in the factor model to have unit roots, Bai and Ng (2004) additionally allow for the possibility of unit roots and cointegration in common factors. This method, however, requires large panels with $N/T \to 0$. Again the common factors are estimated by principal components. Both the idiosyncratic components and common factors are then separately tested for unit roots. The approaches based on factor models are particularly attractive if the number of cross section units (N) is large compared to the number of time periods (T).

Breitung and Das (2005) propose a robust version of the OLS Dickey-Fuller t-statistic which still performs well if the number of time periods is less than the number of cross-sectional units. As an alternative, they suggest a GLS approach obtained from an OLS estimation of a transformed model. The GLS approach, however, is only feasible if $T > N$. Pesaran (2007) introduces the cross section augmented Dickey-Fuller (CADF) test. It augments the standard ADF test with the cross-section averages of both lagged levels and first differences of the individual series. First generation unit root tests can then be applied to the results of the individual CADF tests (e.g. Maddala and Wu (1999)).

All panel unit root tests outlined above are tests against stationarity; none of them tests against explosiveness. In the literature, tests against explosiveness are only applied for the detection of financial bubbles, and they are restricted to univariate time series data. Among them are Bhargava (1986) and Phillips et al. (2011). Combining the p-value approach of Simes (1986) and the p-values of univariate right-tailed unit root tests, we obtain a procedure to test for explosiveness in panel data sets.

Our approach follows suggestions by Hanck (2013) who proposes a panel unit root test based on Simes (1986) classical intersection test. Hanck sets up the global null hypothesis H_0 that all individual null hypotheses $H_{i,0},\ i = 1, ..., N$ are true. The method is easy to implement since it only requires the p-values of N time series unit root tests. They are ordered ascendingly and compared to increasing critical values. The panel null hypothesis is rejected if at least one p-value is smaller than the corresponding critical value. Hanck's method further enables us to identify those units of the panel, for which $H_{i,0}$ is rejected. Moreover, it accounts for the multiple testing nature since it controls the Familywise Error Rate (FWER), i.e. the probability of falsely rejecting at least one individual null hypothesis, at level α.

Our test procedure is based on the first-order panel autoregression implied by (1) and (2),

$$y^i_{h,t} - \alpha^i = \rho^i(y^i_{h-1,t-1} - \alpha^i) + u^i_{h,t} \qquad i = 1, \dots, N;\ t = 1, \dots, T \tag{12}$$

where T denotes the number of time series observations on each of the N individuals, and $u^i_{h,t}$ is a stochastic process with $E(u^i_{h,t}) = 0$ and $Var(u^i_{h,t}) = \sigma^2_u$. The earnings model implies

$$u^i_{h,t} = \eta^i_{h,t} + \varepsilon^i_{h,t} - \rho^i \varepsilon^i_{h-1,t-1}$$

and, hence, $u^i_{h,t}$ has an MA component. Before we continue with the panel case, we consider the univariate unit root test against explosiveness, i.e. the null hypothesis $H_0 : \rho^i \leq 1$ against $H_1 : \rho^i > 1$. The augmented Dickey-Fuller test for individual i is based on the regression

$$\Delta y^i_{h,t} = \mu^i + \beta^i h + \phi^i y^i_{h-1,t-1} + \sum_{k=1}^{K_i} \Delta y^i_{h-k,t-k} + v^i_{h,t}. \tag{13}$$

where $\phi^i = \rho^i - 1$ and $\mu^i = \alpha^i(1 - \rho^i)$. The additional lagged differences are included to capture autocorrelations in the error term $u^i_{h,t}$ such that $v^i_{h,t}$ is uncorrelated. Since time series with an MA component can be approximated by AR processes with a large number of lags, the number of additional lagged differences K_i must not be set too low. Gustavsson and Österholm (2014) suggest to use the information criterion BIC to determine the number of lags. In the appendix, we show the BIC may fail to choose the correct number of lags in relatively short time series. We set $K_i = 3$ which is large enough to pick up the autocorrelation of the error term, but still small enough to ensure a sufficiently large number of observations.

The regression (13) does not only include a constant μ^i but also a trend. The trend is necessary to allow for diverging earnings profiles under the null hypothesis in case of stationarity, i.e. if $\rho^i < 1$. In the absence of a trend, $\rho^i < 1$ would imply a stationary time series around a constant level.

As in ADF tests against stationarity, the test statistic in the test against explosiveness is the t-statistic belonging to the regression coefficient ϕ^i in (13). Its distribution under the null hypothesis is non-standard but tabulated. Since we need precise individual p-values even if they are extremely small, we derived the distribution of the test statistic under the null hypothesis by Monte-Carlo simulations with 1 million replications.

The global null hypothesis states that *all* N time series are either unit root processes or stationary,

$$H_0 : \rho^1 \leq 1, \rho^2 \leq 1, \ldots, \rho^N \leq 1.$$

The global null hypothesis is the intersection over N individual hypotheses, $H_0 = \bigcap_{i=1,\ldots,N} H_{i,0}$ with $H_{i,0} : \rho^i \leq 1$. The alternative hypothesis states that at least one process is explosive. Simes (1986) provides a modification of Bonferroni's procedure for testing multiple hypotheses. The method is less conservative than the classical Bonferroni procedure which lacks power if the test statistics are correlated. The modified version overcomes this problem and controls the FWER even if the test statistics are not independent. If they happen to be independent, the type I error probability is equal to α.

Let $p_{(1)} \leq \ldots \leq p_{(N)}$ denote the ordered p-values of the tests belonging to the individual hypotheses $H_{i,0}$. Simes' test rejects the global H_0 at level α if

$$p_{(j)} \leq j \cdot \frac{\alpha}{N} \text{ for some } j = 1, \ldots, N,$$

i.e., one compares p-values, sorted from most to least significant, to gradually increasing points $j \cdot \alpha/N$. The global null hypothesis is rejected if there exists at least one p-value which is sufficiently small.

The main advantage of this p-value combination approach is that we only require p-values from univariate test statistics, even if these are not independent. At the same time, as found by e.g. Maddala and Wu (1999), p-value combination tests are typically competitive in terms of power and size to computationally much more demanding procedures. We employ standard right-tailed Augmented Dickey-Fuller (ADF) tests to each time series separately. The p-values are computed as the probabilities of obtaining larger test statistics than those actually resulting from the ADF tests under the null hypothesis.

In contrast to previously suggested panel unit root tests, this approach allows to specify the fraction and the identity of the rejected units by means of a procedure suggested by Hommel

(1988). He criticizes that Simes' test does not allow statements about individual hypotheses since the FWER is not controlled in this case.[2] To overcome this problem, Hommel introduced an improved multiple test procedure which allows to make statements about individual hypotheses. He applies Simes' test to each intersection hypothesis of a closed test procedure. His procedure controls the FWER at level α, provided Simes' test has level α. The test decisions for each individual hypothesis is performed according to the following algorithm:

1. Compute $j = \max\{(i = 1, \ldots, N) : p_{(N-i+k)} > k\alpha/i \text{ for all } k = 1, \ldots, i\}$.
2. If $p_{(N)} \leq \alpha$, reject all $H_{i,0}$. Else, reject all $H_{i,0}$ with $p_i \leq \alpha/j$.

Like other p-value combination approaches that are based on transformed sums of p-values, Simes' procedure is consistent as $T \to \infty$ for any $N < \infty$ (e.g. Hanck (2013)). To prove this, let I_A be the index set of hypotheses for which the alternatives hold true. Since $|\rho_{(i)}| > 1$ for at least one i, the index set is not empty ($I_A \neq \emptyset$). When the univariate p-values are inferred from consistent test statistics, there exists some p_j with $j \in I_A$, such that $p_j \equiv p_{(1)} \to_p 0$. Consequently, Simes' procedure is consistent with $\lim_{T\to\infty} \Pr(p_{(1)} < \alpha/n) = 1$. Many other panel unit root tests (e.g. Im et al. (2003), Pesaran (2007)) further require $n_1/n \to \theta_1 > 0$ (as $n \to \infty$) for test consistency. That is, the fraction of rejected individual null hypotheses must not converge against 0. However, this is not necessary for tests which are based on the combination of p-values like Simes' test, since the global null is already rejected if one p-value is sufficiently small.

5 Empirical applications

For our empirical analysis we use earnings data from the cross-national equivalent files of the German SOEP (GSOEP) and the U.S. Panel Study of Income Dynamics (PSID). More precisely, we employ the annual 1984-2012 waves of the GSOEP and the 1970-1997 waves of the PSID, which hence comprise a maximum of 29 and 28 years of observations. For both data sets, we restrict our analysis to individuals aged between 20 and 64 who worked at least 520 hours per year. The maximum amount of hours worked is restricted to 5110, resulting in an average number of hours worked of 1852 (that is approximately 35 hours per week). Moreover, we only consider individuals having an hourly wage rate of larger than, or equal to, 3 Euro/h. Finally,

[2] Simes proposed to reject $H_{1,0}, \ldots, H_{k,0}$ where $k = \max\{j : p_{(j)} \leq j\alpha/N\}$. This procedure was, however, shown not to control the FWER at level α.

we only take into account individuals with at least 15 observations which, however, need not be consecutive. Following these restrictions, we obtain a data set of size $N = 4270$ for the GSOEP and of size $N = 4472$ in case of the PSID. Table 2 gives a brief descriptive overview over both data sets.

	GSOEP	PSID
number of observations	4270	4472
ø years with earnings obs.	20.0	19.8
ø age in first wave	29.3	24.9
ø age in last wave	53.8	50.4

Table 2: Data description

To eliminate possible time effects including inflation and a potential rise in the skill premium, we run a cubic regression of log income on experience h and a dummy variable for the level of education (less than high school, completed high school, or more than high school). This is the usual way to eliminate common effects affecting all individuals equally (Guvenen 2009). The regressions are carried out separately for each wave in each of the data sets. Doing so, experience is approximated in the usual way by age minus school years minus 6. The residuals resulting from this regression constitute the earnings $y^i_{h,t}$ which are modelled in equation (1). To test if explosiveness is evident in earnings, we make use of the panel unit root test suggested in the previous section. We start with the GSOEP data set. The first two columns in table 3 outline the smallest ordered p-values of the univariate right-tailed unit root tests (left column) and compares them to Simes' cutoff-values (right column).[3] The results show that the global null hypothesis, that all time series are unit root or stationary processes, is rejected at the 5% level. Figure 6 (upper panel) illustrates this result: The first 49 p-values for the GSOEP, and the first 29 p-values for the PSID, are below Simes' cutoff line, which is represented by the dashed line.

If we turn around the null and alternative hypotheses, the null hypothesis that all individuals have non-stationary (i.e. unit root or explosive) processes is also rejected at the 5% significance level. Simes' cutoff yields 85 rejections (i.e. significantly stationary processes) for the PSID and 123 rejections for the GSOEP.

Apparently, while explosiveness is evident in a statistically significant manner in both data sets, it does not play a quantitatively important role. We conclude that explosiveness cannot be

[3]The number of individuals reduces to $N = 4061$ due to time series where the NAs are positioned such that no unit root test could be run.

	GSOEP		PSID	
	p-values	Simes' cutoff	p-values	Simes' cutoff
$p_{(1)}$	0.000000	0.000012	0.000000	0.000012
$p_{(2)}$	0.000000	0.000025	0.000000	0.000024
$p_{(3)}$	0.000000	0.000037	0.000003	0.000035
$p_{(4)}$	0.000000	0.000049	0.000013	0.000047
⋮	⋮	⋮	⋮	⋮
$p_{(12)}$	0.000004	0.000148	0.000071	0.000141
$p_{(13)}$	0.000006	0.000160	0.000077	0.000153
$p_{(14)}$	0.000014	0.000172	0.000087	0.000165
⋮	⋮	⋮	⋮	⋮
$p_{(28)}$	0.000188	0.000345	0.000293	0.000329
$p_{(29)}$	0.000194	0.000357	0.000321	0.000341
$p_{(30)}$	0.000201	0.000369	0.000389	0.000353
$p_{(31)}$	0.000226	0.000382	0.000449	0.000365
⋮	⋮	⋮	⋮	⋮
$p_{(48)}$	0.000590	0.000591	0.001081	0.000565
$p_{(49)}$	0.000597	0.000603	0.001081	0.000576
$p_{(50)}$	0.000723	0.000616	0.001084	0.000588
$p_{(51)}$	0.000724	0.000628	0.001087	0.000600
$p_{(52)}$	0.000794	0.000640	0.001099	0.000612
N_{new}	4061		4251	

Table 3: First sorted p-values and Simes' cutoff

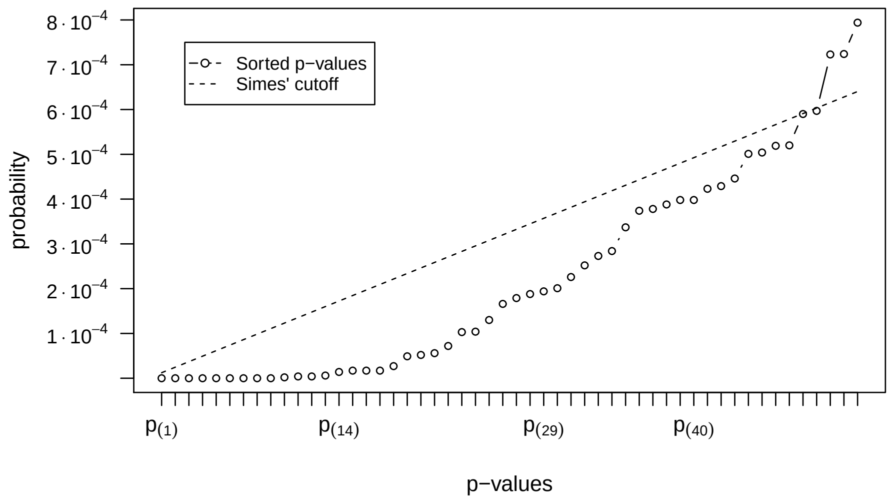

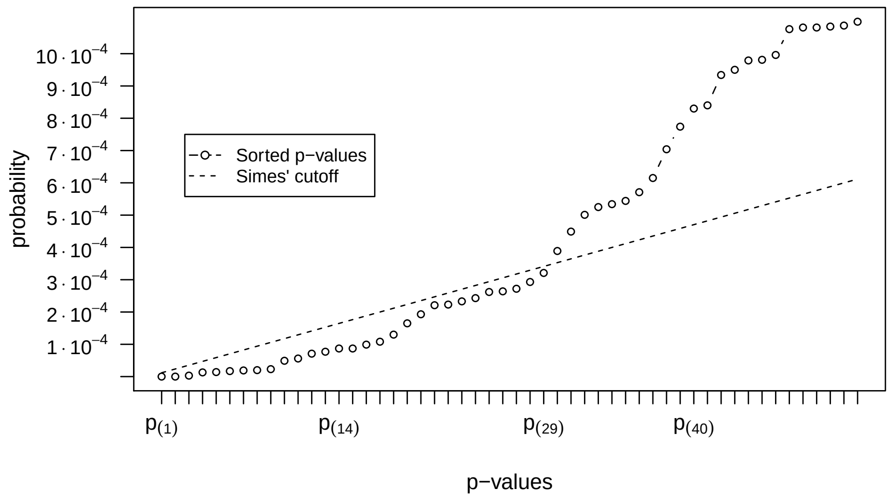

Figure 6: First sorted p-values compared to Simes' cutoff line, upper panel: GSOEP, lower panel: PSID

ruled out for a limited number of individuals. Hence, heterogeneity should not only be taken into account with respect to the earnings profiles but also with respect to the persistence properties. This finding is in line with other studies that stress the importance of heterogeneity.

6 Conclusion

In many countries earnings inequality has risen sharply in the last decades. Workers at the bottom of the distribution lost, while the top earnings increased ever more. The idea that the rich get richer and the poor get poorer, is formalized by model with explosive earnings dynamics. Individuals with earnings above a certain threshold experience higher growth rates than individuals below that line. This case is ruled out by the two standard models of earnings dynamics, the unit root model and the model with heterogeneous profiles.

We show that, in some respects, the explosive model may resemble models with heterogeneous earnings profiles. In particular, the two models can have very similar variance-covariance structures of earnings. If the true model is explosive the variances and covariances might therefore look as if they belong to a model with profile heterogeneity, and vice versa. Taking into account consumption or savings data does not necessarily solve the identification problem. We derive the optimal consumption behaviour for the explosive model and compare it to the standard models. It turns out that individuals facing explosive earnings dynamics behave in a similar way as individuals facing heterogeneous profiles. Hence, identification based on a structural model is difficult.

A straightforward method to distinguish between explosive and non-explosive (i.e. stationary or unit root) dynamics is a panel unit root test against explosiveness. We suggest a procedure based on Simes' (1986) intersection test. The global null hypothesis states that all individuals have either stationary earnings processes or unit root processes. The test procedure is illustrated for U.S. and German earnings panel data ranging over almost 30 years. We find that the null hypothesis can be rejected at usual significance levels, but the number of significantly explosive earnings profiles is small.

Our findings support other studies that stress the importance of heterogeneity for models of earnings dynamics. Heterogeneity should not only be implemented in terms of the expected earnings profile, but, perhaps even more important, also in regard to the persistence properties such that both mean-reverting, unit root, and explosive processes are present.

References

Abowd, J. and Card, D. (1989). On the covariance structure of earnings and hours changes, *Econometrica* **57**(2): 411–445.

Atkinson, A. B., Piketty, T. and Saez, E. (2011). Top incomes in the long run of history, *Journal of Economic Literature* **49**: 3–71.

Autor, D. H. (2014). Skills, education, and the rise of earnings inequality among the "other 99 percent", *Science* **344**(6186): 843–851.

Bai, J. and Ng, S. (2004). A PANIC attack on unit roots and cointegration, *Econometrica* **72**(4): 1127–1177.

Baker, M. (1997). Growth-rate heterogeneity and the covariance structure of life-cycle earnings, *Journal of Labor Economics* **15**(2): 338–375.

Baker, M. and Solon, G. (2003). Earnings dynamics and inequality among Canadian men, 1976-1992: Evidence from longitudinal income tax records, *Journal of Labor Economics* **21**(2): 289–321.

Bhargava, A. (1986). On the theory of testing for unit roots in observed time series, *Review of Economic Studies* **53**(3): 369–384.

Breitung, J. and Das, S. (2005). Panel unit root tests under cross-sectional dependence, *Statistica Neerlandica* **59**(4): 414–433.

Browning, M., Ejrnæs, M. and Alvarez, J. (2010). Modelling income processes with lots of heterogeneity, *Review of Economic Studies* **77**: 1353–1381.

Carroll, C. D. (2011). Theoretical foundations of buffer stock saving, CFS Working Paper, No. 2011/15.

Chang, Y. (2002). Nonlinear IV unit root tests in panels with cross-sectional dependency, *Journal of Econometrics* **110**(2): 261–292.

Ejrnæs, M. and Browning, M. (2014). The persistent-transitory representation for earnings processes, *Quantitative Economics* **5**: 555–581.

Fisher, R. A. (1925). *Statistical methods for research workers*, Genesis Publishing Pvt Ltd.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life-cycle, *Econometrica* **70**(1): 47–89.

Gustavsson, M. and Österholm, P. (2014). Does the labor-income process contain a unit root? Evidence from individual-specific time series, *Journal of Economic Dynamics and Control* **47**: 152–167.

Guvenen, F. (2007). Learning your earning: Are labor income shocks really very persistent?, *American Economic Review* **97**(3): 687–712.

Guvenen, F. (2009). An empirical investigation of labor income processes, *Review of Economic Dynamics* **12**(1): 58–79.

Guvenen, F., Karahan, F., Ozkan, S. and Song, J. (2015). What do data on millions of U.S. workers reveal about life-cycle earnings risk?, NBER Working Paper 20913.

Guvenen, F. and Smith, A. A. (2014). Inferring labor income risk and partial insurance from economic choices, *Econometrica* **82**(6): 2085–2129.

Haider, S. J. (2001). Earnings instability and earnings inequality in the United States, 1967–1991, *Journal of Labor Economics* **19**(4): 799–836.

Hanck, C. (2013). An intersection test for panel unit roots, *Econometric Reviews* **32**(2): 183–203.

Hommel, G. (1988). A stagewise rejective multiple test procedure based on a modified Bonferroni test, *Biometrika* **75**(2): 383–386.

Hryshko, D. (2012). Labor income profiles are not heterogeneous: Evidence from income growth rates, *Quantitative Economics* **3**: 177–209.

Im, K. S. and Pesaran, M. H. (2004). On the panel unit root tests using nonlinear instrumental variables, Working Paper.

Im, K. S., Pesaran, M. H. and Shin, Y. (2003). Testing for unit roots in heterogeneous panels, *Journal of econometrics* **115**(1): 53–74.

Levin, A., Lin, C.-F. and James Chu, C.-S. (2002). Unit root tests in panel data: asymptotic and finite-sample properties, *Journal of econometrics* **108**(1): 1–24.

Lillard, L. A. and Weiss, Y. (1979). Components of variation in panel earnings data: American scientists, 1960-70, *Econometrica* **47**(2): 437–454.

MaCurdy, T. E. (1982). The use of time series processes to model the error structure of earnings in a longitudinal data analysis, *Journal of Econometrics* **18**(1): 83–114.

Maddala, G. S. and Wu, S. (1999). A comparative study of unit root tests with panel data and a new simple test, *Oxford Bulletin of Economics and statistics* **61**(S1): 631–652.

Meghir, C. and Pistaferri, L. (2004). Income variance dynamics and heterogeneity, *Econometrica* **72**(1): 1–32.

Mincer, J. (1958). Investment in human capital and personal income distribution, *Journal of Political Economy* **66**(4): 281–302.

Moon, H. R. and Perron, B. (2004). Testing for a unit root in panels with dynamic factors, *Journal of econometrics* **122**(1): 81–126.

Pesaran, M. H. (2007). A simple panel unit root test in the presence of cross-section dependence, *Journal of Applied Econometrics* **22**(2): 265–312.

Phillips, P. C. and Sul, D. (2003). Dynamic panel estimation and homogeneity testing under cross section dependence, *Econometrics Journal* **6**(1): 217–259.

Phillips, P. C., Wu, Y. and Yu, J. (2011). Explosive behavior in the 1990s NASDAQ: When did exuberance escalate asset values?, *International economic review* **52**(1): 201–226.

Simes, R. J. (1986). An improved Bonferroni procedure for multiple tests of significance, *Biometrika* **73**(3): 751–754.

Zeldes, S. (1989). Optimal consumption with stochastic income: Deviations from certainty equivalence, *Quarterly Journal of Economics* **104**(2): 275–298.

A Numerical solution of dynamic optimization

There are two state variables: the wealth level W_t and the deviation z_t. All parameters (β, γ, ρ, σ_η) are assumed to be known to the individual. Hence, it is also possible to define earnings Y_t or log earnings y_t rather than z_t as the second state variable. The stochastic dynamic optimization problem (9) with constraint (10) is solved by backward recursion. In the last period T the value function is given by (11) which is completely specified for given R and γ.

As noted by Zeldes (1989) and others, it is possible to merge the state variables W_t and Y_t into their ratio W_t/Y_t and thus reduce the number of states from 2 to 1 which makes the numerical computations substantially easier. Unfortunately, this method hinges on the unit root property, it does not work once explosiveness is taken into account.

The value function is numerically represented by a bivariate grid over W_t with $n_W = 42$ equidistant points ranging over the interval $[0.001, 15]$, and Y_t with $n_Y/2 = 20$ equidistant points over the interval $[0.0001, 0.5]$ followed by another 20 equidistant points over $[1, 40]$,

$$\tilde{V}_T(\tilde{W}_i, \tilde{Y}_j), \qquad i = 1, \dots, n_W, \quad j = 1, \dots, n_Y$$

where $\tilde{W}_i$ and $\tilde{Y}_j$ are the grid points.

For periods $t = T-1, T-2, \dots, 1$ the value function and the optimal consumption level are found by a golden section search over consumption C_t. For a given value of C_t the maximand $U(C_t) + E_t V_{t+1}(W_{t+1}, z_{t+1})$ is calculated as follows. Since the future level of wealth W_{t+1} solely depends on the nonstochastic variables W_t, C_t and Y_t the conditional expectation is only relevant for future earnings. We compute the conditional expectation of next periods value function, given $\tilde{W}_i$, $\tilde{Y}_j$ and C_t, by Gaussian Hermite quadrature (with 26 points, so that even very extreme movements by almost 9 standard deviations are taken into account). Let $\tilde{\eta}_1, \dots, \tilde{\eta}_{26}$ denote the grid points and $\omega_1, \dots, \omega_{26}$ their weights. For $i = 1, \dots, n_W$, $j = 1, \dots, n_Y$ and $k = 1, \dots, 26$, next periods earnings levels are

$$\tilde{Y}_{t+1,i,j,k} = \exp(\rho \log(\tilde{Y}_j) + \tilde{\eta}_k + (1-\rho)t\beta + \beta)$$

and next periods wealth levels are

$$\tilde{W}_{t+1,i,j} = \tilde{W}_i + \tilde{Y}_j - C_t.$$

Since next periods earnings and wealth levels are, in general, not located on the grid points we need to interpolate and extrapolate. As the number of periods is rather large ($T = 40$), even small approximation errors in late periods will accumulate and the value function approximation for early periods can become very poor. We transform the value function by

$$\mathcal{V}_{t+1}(\tilde{W},\tilde{Y}) = (-V_{t+1}(\tilde{W},\tilde{Y}))^{-\rho}$$

in order to reduce the amount of nonlinearity and approximate the transformed value function by bilinear interpolation and extrapolation. After transforming back we obtain

$$\tilde{V}_{t+1,i,j,k} = V_{t+1}(\tilde{W}_{t+1,i,j},\tilde{Y}_{t+1,i,j,k})$$

and the conditional expectation is computed as $\sum_{k=1}^{26}\omega_k\tilde{V}_{t+1,i,j,k}$ for all i and j. It is now straightforward to optimized $U(C_t)+E_tV_{t+1}(W_{t+1},z_{t+1})$ over C_t for all i and j. The maximum is the value of the value function $V_t(\tilde{W}_i,\tilde{Y}_j)$. The optimal consumption level $C_t^*(\tilde{W}_i,\tilde{Y}_j)$ is also recorded.

B Lag order determination for the ADF test

The panel test against explosiveness is based on N univariate ADF tests. The number of lags K_i to be included in regression (13) can be determined by the Bayesian (or Schwarz) information criterion (Gustavsson and Österholm 2014). However, if there are measurement errors, the BIC does not, in general, choose a sufficiently large lag order as can be demonstrated by means of a simple Monte-Carlo simulation.

Figure 7 (solid line) shows the distribution function of the ADF test statistic (with constant and trend) under a unit root when there is no measurement error. The time series is generated as $y_t = z_t + t$ with

$$z_t = z_{t-1} + \eta_t, \qquad t = 1,\dots,50$$

where $\eta_t \sim N(0,1)$ is white noise and $z_0 = 0$. The number of lags to be included is chosen by the BIC (of course, the correct number is 0). The number of Monte-Carlo replications is $R = 10000$.

The dashed line in figure 7 depicts the distribution function of the ADF test statistics if white noise measurement error is added to the time series,

$$y_t = z_t + t + \varepsilon_t$$

where $\varepsilon_t \sim N(0,1)$. Measurement error leads to an MA component in the first differences of y_t. The number of lags has again been selected by the BIC. Apparently, the distribution is shifted to the left and does not equal the null distribution (solid line). We conclude that the BIC does

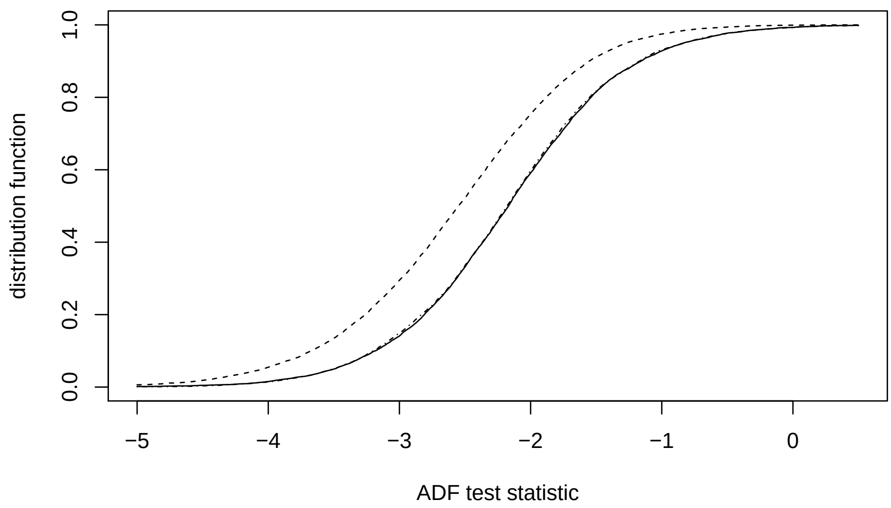

Figure 7: Distribution function of the ADF test statistic; solid line: distribution under a unit root without measurement error and lag order determined by BIC; dashed line: distribution with measurement error and lag order determined by BIC; dotted-dashed line: distribution with measurement error and lag order set to 3.

not succeed in selecting the correct lag order. The same result holds for the Akaike information criterion AIC (not shown).

If we set the number of lags to $K_i = 3$ the resulting distribution of the test statistic is shown by the dotted-dashed line in figure 7. Apparently, it virtually equals the true null distribution (solid line). In our empirical application we therefore set the lag order to $K_i = 3$ for all individuals.