Review of Economic Dynamics 29 (2018) 172–194

Contents lists available at ScienceDirect

Review of Economic Dynamics

www.elsevier.com/locate/red

# The job ladder and its implications for earnings risk ☆

Joachim Hubmer

*Department of Economics, Yale University, 28 Hillhouse Avenue, New Haven, CT 06511, USA*

ARTICLE INFO

*Article history:*
Received 13 September 2016
Received in revised form 10 January 2018
Available online xxxx

*JEL classification:*
E24
J24
J31
J64

*Keywords:*
Earnings dynamics
Life-cycle earnings risk
Non-Gaussian shocks
On-the-job search

ABSTRACT

This paper analyzes the ability of a job ladder framework to explain recent evidence on life-cycle earnings dynamics. Using administrative data, Guvenen et al. (2015) document several new facts about the distribution of earnings growth, most notably large negative skewness and high excess kurtosis, rejecting the frequently used log-normal framework. I show that these new facts can be well explained by a standard structural representation of a frictional labor market, a life-cycle version of the job ladder model, in combination with a simple human capital process. Furthermore, I identify endogenous search effort, risk aversion and wealth accumulation, and skill loss in unemployment as key model features that interact with the labor market friction to jointly reconcile the evidence.



## 1. Introduction

Incorporating heterogeneous agents has been one of the key advances in macroeconomic equilibrium modeling in the past decades. Idiosyncratic earnings risk, which is quantitatively significant and to a large extent uninsurable, is the major ingredient in these models. Capturing the salient features of earnings risk is therefore essential. While previously earnings dynamics have been mainly described by parsimonious log-normal earnings processes estimated from survey data, recent evidence based on administrative data reveals several new facts that contradict the assumption of log-normality.[1] In particular, Guvenen et al. (2015) utilize a rich dataset from the U.S. Social Security Administration to document that earnings growth exhibits large negative skewness and very high kurtosis.[2] Furthermore, they find sizable variation in these higher order moments across different population groups defined by age and the level of earnings. Negative skewness means that negative changes are typically larger than positive changes. Excess kurtosis implies that most individuals experience rather small (or no) earnings changes, while a few are hit by very large shocks. Both features are in stark contrast to the log-normal framework, which implies symmetric changes and rather mild risk.

☆ I am grateful to Giuseppe Moscarini and Tony Smith for invaluable guidance and support. I thank Fatih Guvenen, Christopher Huckfeldt, Per Krusell, Aleh Tsyvinski, and seminar participants at the 2016 SED Meeting in Toulouse as well as at Yale for helpful comments and suggestions. I thank in particular the editor, Gianluca Violante, and an anonymous referee for their generous input, which greatly improved the paper.

*E-mail address:* joachim.hubmer@yale.edu.

[1] Administrative data has several advantages over survey data, including the absence of measurement error and often much larger sample sizes, thus allowing for estimating higher order moments and in general applying nonparametric methods.

[2] Earlier work on earnings dynamics going beyond the log-normal framework includes Geweke and Keane (2000) and Bonhomme and Robin (2010).

https://doi.org/10.1016/j.red.2018.01.003
1094-2025/

This paper aims to show that a standard structural representation of a frictional labor market, the job ladder model, can explain these new facts. Moreover, it sheds light on which model features have to be added to the basic job ladder model outlined by Burdett (1978) for this purpose. From a statistical point of view, these new facts can also be modeled by error component models, in the tradition of a rich literature going back to MaCurdy (1982) and Abowd and Card (1989), among many others. It is, however, from an economist's point of view more satisfactory to explore to what extent a behavioral model can account for the new evidence. The job ladder framework is a natural starting point as its defining element is expected to generate negative skewness and excess kurtosis: In the absence of a separation shock, workers only switch to more productive firms, moving up the ladder relatively slowly. A layoff implies falling down the ladder all the way, hence on-the-job search implies negative skewness. Moreover, both job-to-job transitions as well as movements between employment and unemployment are potentially large and infrequent changes, thus generating excess kurtosis.

In the proposed life-cycle model, heterogeneous and risk averse workers search for job opportunities at heterogeneous firms while unemployed as well as on-the-job. To smooth consumption and partially insure against labor market shocks, they have access to a risk-free asset. Workers differ by age, wealth, and human capital, modeled as a Gaussian random walk. Firms are characterized by their productivity level. Crucially, I do not calibrate the model to match the evidence on earnings dynamics. Instead, I target aggregate labor market transition rates and the earnings distribution in levels, plus the average long-term earnings loss associated with displacement from a job. In turn, I evaluate the performance of the model against the earnings dynamics moments documented in Guvenen et al. (2015). Overall, the model is successful in explaining the magnitude of negative skewness and excess kurtosis. Moreover, it captures that negative skewness and kurtosis are increasing in the level of earnings and over the life-cycle.

I identify three key model features that are essential to fit the data by way of shutting each one of them down separately, re-estimating the respective alternative model, and comparing the resulting earnings dynamics to the benchmark model. The first one is endogenous search. In the benchmark model, search is costly and agents optimally choose search intensity. It turns out that if contact rates are exogenous (and uniform across agents), then the model generates much less negative skewness as long-term unemployment (and thus very large negative shocks) becomes extremely rare. Furthermore, the model is no longer able to replicate that negative skewness is increasing in age. Primarily, this pattern holds in the benchmark model because older agents optimally search less. Thus, they re-climb the ladder more slowly after a separation shock.

The second key feature is risk aversion and wealth accumulation. If workers are risk neutral, I find that both negative skewness as well as kurtosis decline substantially. While the risk neutral model is explaining the data well for the bottom 40% of the earnings distribution, it fails to match the increase in negative skewness and kurtosis for higher earners. In the benchmark version, there is a wealth effect: higher earners tend to be richer, thus they optimally search at lower intensity, which prolongs non-employment spells and thus generates larger negative earnings changes.

The third important ingredient concerns the dynamics of human capital. The benchmark model assumes that the growth rate of human capital depends on the employment status of a worker. In particular, human capital depreciates in unemployment in the spirit of Ljungqvist and Sargent (1998). The size of this decline is disciplined by targeting the average earnings loss, in present value terms, associated with displacement from a job as estimated by Davis and von Wachter (2011). If the growth rate of human capital is required to be the same for employed and unemployed workers, the model can no longer generate large enough earnings losses, while simultaneously matching aggregate labor market transition rates. Consequently, negative earnings changes are smaller and skewness is less negative. Similarly, as the largest changes are negative, kurtosis decreases.

There are some aspects of the Guvenen et al. (2015) data that the benchmark model is missing. First, the magnitude of earnings fluctuations, as measured by the standard deviation of earnings changes, is too low. At least to some extent, this is to be expected as the benchmark model excludes some sources of transitory fluctuations such as multiple jobholding, varying hours worked, and bonus payments. Structurally modeling all those would be daunting.[3] Second, the model is creating too much negative skewness and excess kurtosis for the highest earners, those in the top 10% and especially even farther out in the right tail. This failure of the model is related to two features of the benchmark model that are otherwise crucial in explaining earnings dynamics for the bottom 90% of earners: risk aversion and wealth accumulation, as well as endogenous search effort. Very high earners tend to be very rich, and thus in the model they have very little incentive to search for a new job once hit by a separation shock, creating very large negative earnings changes. There is, perhaps, an analogy to the literature on explaining wealth inequality that struggles with explaining wealth concentration at the very top. In that literature, the standard behavioral model has trouble explaining why the very rich are not consuming more of their wealth. Here, the model has trouble explaining why the very rich are not searching less, and working less. One could imagine that a taste for work is at play for top athletes and artists who enjoy performing, and

[3] One can, of course, enrich the dynamics of human capital by adding a transitory component that statistically accounts for these missing elements. This avenue is pursued in an extended model presented in Appendix A. Relative to the benchmark, this model generates more earnings fluctuations. As the added exogenous transitory shock is symmetric, skewness is increasing (i.e., less negative), thus providing a better fit to the data (the baseline model is overshooting). The kurtosis of earnings growth is not affected materially.