The historical evolution of the wealth distribution: A quantitative-theoretic investigation

Authors: Joachim Hubmer, Per Krusell, and Anthony A. Smith, Jr.

Extended abstract

We derive qualitative and quantitative predictions of a microfounded model of wealth inequality and look at how the model's predictions compare with actual outcomes in the United States over the postwar period. The model's microeconomic core emphasizes that a household faces important earnings risks (deriving from the risk of unemployment and from uncertain wages) and cannot fully insure against these risks. To capture these features, we base our work on the broadly used Bewley-Huggett-Aiyagari (BHA) setup. Our model also features stochastic movements in saving rates that go beyond the mere precautionary-savings motive inherent in the BHA setup: we posit some randomness in discount factors, as in Krusell and Smith (1998). We demonstrate that such randomness gives rise to a Pareto-shaped right tail of the wealth distribution, a feature that appears to approximate the data well and that has been discussed much recently and derived in a variety of highly stylized models. Piketty and Zucman (2014) develop one such model and we argue that it can be viewed as a reduced form for the model we present here, and indeed also for the one in Krusell and Smith (1998).

We calibrate our discount-rate process so as to match key statistics of the average wealth distribution, including the Pareto coefficient characterizing its right tail, and we examine whether there is a tight relationship between this coefficient and other measures of wealth concentration, such as the percentage of society's wealth held by the 0.1% richest. We also examine to what extent calibrations of the earnings distribution that yield large wealth inequality but do not feature discount-factor (or return) shocks—such as that in Castañeda, Días-Giménez, and Ríos-Rull (2003)—can deliver a Pareto shape for the right tail of the wealth distribution: as far as we can tell, at least standard versions of such calibrations cannot.

A key comparative-statics exercise we perform (quantitatively) involves the relation between earnings inequality and wealth. We show, in particular, that at least many forms of increases in the earnings dispersion lead to *decreases* in long-run wealth inequality, and we examine the key mechanisms behind this finding. However, a comparative-dynamics exercise also demonstrates that in the shorter run, wealth dispersion may *increase*, and hence the net impact on any given horizon is hard to predict without fully specifying the model quantitatively. Other comparative-statics exercises here concern changes in population growth and technology growth.

For our quantitative work, we feed in measures of the changes in earnings inequality (net of taxes and transfers) experienced during this period, as measured by Heathcote, Storesletten, and Violante (2010) using a standard process for earnings consisting of a sum of a persistent and a temporary component and we take into account changes in taxes as well. Here, we consider a range of assumptions regarding the extent to which the changes

in the net-earnings distribution were foreseen. The goal here is to investigate whether these observed changes (i.e., changes in the distribution of gross earnings and changes in taxes) suffice in accounting for the evolution of wealth inequality over this period in the United States. [As of February 15, 2015, the work is ongoing, still without definitive results.]