Wage Inequality and the Business Cycle*

Gadi Barlevy
Department of Economics
Northwestern University
Evanston, IL 60208
USA

Daniel Tsiddon
Department of Economics
Tel Aviv University
Tel Aviv, 69978
Israel

March 23, 2000

Abstract

Previous work on wage inequality over the cycle has treated long-run trends in inequality as an orthogonal process that is unrelated to cyclical fluctuations. This intuition implies cyclical inequality can be measured by regressing a detrended measure of inequality on cyclical indicators. We develop a model in which cyclical and trend inequality are related. In our model, recessions increase inequality when inequality is increasing, but decrease it when inequality is decreasing. This cautions against using simple detrending in measuring the effects of business cycles on wage inequality.

*The authors would like to acknowledge Kiminori Matsuyama for useful discussion.

Introduction

Economists studying business cycles are primarily interested in questions that relate to welfare: for example, how do aggregate fluctuations affect the well-being of individuals, and can society be made better off by pursuing a stabilization policy to offset them?[1] An equally interesting but much less studied question is whether business cycles have a differential impact on different members of society, and what consequences if any this has on the degree of inequality within the greater population. For example, do business cycles promote inequality by hurting those who are already least well off? or do they alleviate inequality by bringing down the wages of those who are the most well off in line with lower-paid workers? With the increasing availability of datasets that provide information on the distribution of earnings at business cycle frequencies, economists have begun to tackle some of these questions. Various studies now exist that use individual earnings data on wages and income from the past forty years to address precisely such questions, including work by Blinder and Esaki (1978), Blank and Blinder (1986), Burtless (1990), and Cutler and Katz (1991). All of these papers seem to reach a similar conclusion: recessions promote greater inequality since they have the most adverse impact on those who are already least well off.[2]

In interpreting these findings, it is important to realize that the recent time period for which data on the distribution of earnings is readily available is also marked by dramatic secular changes in the earnings distribution. In particular, wage and income inequality in the United States have dramatically increased during the past 40 years. Consequently, any attempt to measure the evolution of inequality over the business cycle must take a stance on how to account for secular trends when identifying the effects of cyclical fluctuations. The empirical literature described above has typically assumed that the process which drives long-run changes in wage inequality is orthogonal to cyclical phenomena. Under this logic, we simply need to include a time trend to capture the effects of long-run developments in conventional measures of inequality such as

[1] See, for example, the analysis in Lucas (1987) and the subsequent literature his discussion generated. A review of this literature is contained in Barlevy (2000).

[2] These studies deal with income inequality. Wage dynamics over the cycle are examined in Swanson (1998) and Rubinstein and Tsiddon (2000), although they do not relate their findings to conventional measures of inequality. Rather, they examine whether wage cyclicality systematically differs across particular sub-groups. Both find that wages are most cyclically sensitive among those who are least well off.

the Gini index or wage ratios at different percentiles, and we can identify the effects of business cycles on inequality as deviations of these indices from trend. This is precisely the approach taken in the papers cited above. However, there is little if any discussion in these papers that attempts to justify the assumption that long-run trends and cyclical fluctuations in inequality are unrelated. This naturally raises questions about the robustness of such exercises. For example, do the findings of these papers hinge on this identifying assumption? And, if long-run and short-run developments in earnings inequality are in fact related, are the conclusions of these papers influenced by the fact that they focus on a period of increasing long-run inequality?

This paper tries to address these questions, at least with respect to earnings inequality among employed workers. Our starting point is a model that exhibits secular long-run changes in wage inequality. We then introduce cyclical fluctuations into this model in order to study the interaction between trend and cycle in contributing to changes in the distribution of wages among employed workers. More specifically, we model changes in long-run wage inequality as being driven by ability-biased technical change, which has been used by previous authors to account for the secular increase in wage inequality over the past 40 years. That is, we assume individuals differ in their ability to learn how to use new technologies, where learning is assumed to take time away from production. Those who are quicker to learn become more productive sooner, leaving their less able counterparts behind. This results in increasing inequality over time. When we introduce business cycles into this environment of increasing inequality, intertemporal substitution implies that in less productive times, individuals will devote more of their time to learning and less to production. Those who are more able to learn will drift apart more rapidly from their less able counterparts during downturns. Hence, recessions contribute to more rapid wage inequality than booms, just as we observe in the data.

It would be misleading, however, to infer from this observation that recessions generally exacerbate inequality, since our model really implies that recessions amplify the underlying trend of increasing inequality. To illustrate this point, we modify the model to allow for decreasing inequality over time. Once we do this, downturns have the opposite effect: rather than exacerbating inequality, they accelerate the pace of wage equalization. This is because we generate decreasing inequality in our model by assuming there is a limit on how productive individuals can be under the new technology. In that case, the initial increase in inequality that occurs when a new technology is introduced will eventually be reversed as those who are quickest to learn master the new technology and their wages stop growing. At this point, less able workers begin to catch up to their more able colleagues. Thus, inequality initially rises and then declines, just as predicted by the celebrated Kuznets (1955) curve — although the process we describe occurs with each new innovation rather than over the general course of economic development. Along the downward sloping part of the Kuznets curve, the same forces of intertemporal substitution will cause recessions to accelerate the rate at which the less able catch-up with the most able: recessions still encourage less able workers to invest in learning more skills, but have little or no effect on the more able workers who are sufficiently proficient. Hence, our model cautions against looking for a uniform effect of business cycles on the distribution of wages, as the current literature has set out to do. Instead, the effect of recessions will be intimately connected with underlying long-run trends: recessions should cause the wage distribution to fan out if inequality is increasing, but should also cause the wage distribution to become more compressed if inequality is decreasing.

To assess the plausibility of this assertion that business cycles do not have a uniform impact on the wage distribution, we compare the effects of the Great Depression on inequality to those of more recent macroeconomic downturns. The Depression provides a natural setting to test our hypothesis. It came on the heals of what was arguably a period of ability-biased technological change, namely the shift towards use of electricity in large scale production. David (1991) argues that the phasing in of such technologies was largely concentrated between 1910-1920. At this time, income and wage inequality were generally increasing, which is consistent with our basic assumption that new technologies are biased towards certain segments of the population. This inequality climaxed in the late 1920s, at which point the tide of increasing inequality began to recede. The Depression occurred just as this convergence process was getting underway. If the forces we describe in our model capture an important channel through which business cycles affect the distribution of wages, we would expect to see some evidence of declining inequality during the Great Depression, or at least some evidence of closing wage gaps. We review existing evidence for this period, and argue that while the evidence is far from conclusive, there is some support for the notion that the Great Depression was associated with an acceleration in the pace of declining inequality. At the very least, the Depression was not associated with a dramatic increase in inequality one would have expected given the severity of the Great Depression and the seemingly robust conclusion from recent empirical work that recessions tend to increase inequality.

The paper is organized as follows. Section 1 develops a model of increasing wage inequality and illustrates the interaction between long-run inequality and cyclical fluctuations. Section 2 modifies the model to allow for decreasing wage inequality, and illustrates that recessions will act to decrease inequality in such circumstances. Section 3 relates this prediction to evidence on inequality during the Great Depression. Section 4 concludes.

1. The Case of Increasing Wage Inequality

We begin our discussion with the more straightforward case where long-run wage inequality is increasing. As noted in the Introduction, our explanation for increasing inequality is based on ability-biased technical change. That is, we assume that at some initial date, a new technology is introduced which requires workers to learn a set of skills in order to use it more effectively to produce output. Individuals vary in their ability to learn new skills, and given their innate ability must decide how much time to devote to working and how much time to devote to mastering additional skills. More able workers who have an innate advantage in adapting to new technologies will acquire skills at a more rapid rate, leaving their slower colleagues behind and creating an increasing gap in wages (as well as earnings) over time. Previous work has already raised the possibility that rising inequality over the past 40 years can be attributed to the arrival of new technologies, coupled with the tendency of some individuals to take better advantage of these technologies than others. For example, Galor and Tsiddon (1997) and Tsiddon and Rubinstein (1998) develop models where individuals with high innate productivity switch to new technologies that are more complementary with innate talent. This sorting magnifies differences in innate productivity to make for a more skewed wage distribution than before the new technology was introduced.[3] Our model is more closely related to Caselli (1999). He assumes that when a new technology arrives, individuals must pay a fixed cost before they can use it to produce. The cost of learning differs across workers, so only a fraction of the population shifts to the new technology. The difference between his approach and ours is that we model the learning process as ongoing rather than a one-time prerequisite for production. Formally, we model learning as building up a base of knowledge on how to use a new technology,

[3] Roy (1951) originally observed that sorting across occupations would make for a more skewed wage distribution than the underlying distribution of innate productivity. The contribution of this more recent literature is in relating this observation to technological change.

in much the same way as previous authors model the acquisition of human capital, specifically Ben-Porath (1967). The model we present in this section is essentially a variant of his which allows for stochastic fluctuations in aggregate productivity.[4]

Consider a world in which labor is the only factor of production. At date 0, a new technology is introduced for producing output from labor. The productivity of a given worker depends on how familiar he is with the various facets of this new technology. We assume that at date 0, all individuals begin with equal knowledge of how to use this technology. Let $s_{it} \in [0, \infty)$ denote a measure of how skilled individual i is on this new technology by date t. As a concrete example, consider the introduction of computers. In that case, s_{it} can reflect the competency of a person with computer technology, e.g. the number of different applications he can execute or the number of programming languages he is fluent in. Alternatively, s_{it} could reflect an ordering of different production methods within a general purpose technology, so that individuals move on to more advanced techniques over time. This latter interpretation shows that our model can capture the intuition in previous models where inequality results from individuals sorting across technologies as opposed to distinguishing themselves within the same technology. Our assumption that all individuals begin with an equal footing amounts to the assumption that $s_{i0} = s_0$ for all i. For reasons that will become clear below, we require $s_0 > 0$.

The amount of output a given worker can produce is proportional to his skill at using the new technology and to the amount of time he actually spends working. Formally, an individual has one unit of time each instant which he can allocate between working and learning additional skills.[5] Let $n_{it} \in [0, 1]$ denote the fraction of time individual i spends in acquiring additional skills at a given point in time, so that $1 - n_{it}$ denotes the amount of time he spends working.

[4] Dellas and Sakellaris (1997) also develop a model of human capital accumulation with fluctuations in aggregate productivity. Their model differs from ours in that they assume concave utility over consumption and a linear accumulation technology, while we assume a linear utility and a concave accumulation technology. Moreover, they focus on the predictions of the model for schooling decisions, whereas we focus on the implications for wage growth and inequality.

[5] This specification assumes that learning occurs off the job. However, all we really need for our results is that there is a net tradeoff between engaging in production and acquiring skill, so it is possible for there to be some learning about the new technology through active work. Under our alternative interpretation, this assumption implies upgrading to a new technology requires time away from production.

We assume that the output an individual can produce is given by

$$y_{it} = z_t s_{it} (1 - n_{it}) \tag{1.1}$$

where z_t denotes the productivity of the underlying technology. Since the output of a worker only depends on his own labor and skill, we are implicitly assuming that production exhibits constant returns to scale, i.e. the aggregate accumulation of skills across the population at large will not affect the incentives of any given individual to accumulate skills. This assumption allows us to study the decision of each individual in isolation, which means we can apply the analysis of Ben-Porath (1967) for a single individual in order to study an aggregate model with a population of individuals.[6] To allow for aggregate fluctuations, we assume z_t follows a Markov process. That is, z_t can take on two values, $Z_0 < Z_1$, and aggregate productivity switches between these two values at a constant rate μ per unit of time. Although it is debatable whether productivity shocks are the underlying source of business cycle fluctuations, these shocks influence the decisions of agents in our model only insofar as they generate procyclical movements in wages. This procyclical pattern is consistent with empirical evidence on wage dynamics such as Solon, Barsky, and Parker (1994), and so assuming productivity shocks to be the source of fluctuations seems reasonable.

Next, we specify the process of learning how to produce under the new technology. We adopt Ben-Porath's functional form by assuming that

$$\dot{s}_{it} = a_i (s_{it} n_{it})^b \tag{1.2}$$

where $a_i > 0$ measures the speed at which individual i learns new skills, and $b \in (0, 1)$ is a constant that is the same across all individuals and captures diminishing returns in learning new skills. This specification assumes that the rate at which an individual acquires new skills depends on how many skills he already knows, as well as on the amount of time they spend learning. We choose this learning technology for two reasons. First, it is one that many readers are likely to be familiar with. Second, it is one of the few functional forms for which we can get closed form solutions. A key assumption in what follows is that while all individuals start

[6] Allowing for decreasing returns in the aggregate amount of skill could generate interesting spillovers; for example, as able individuals accumulate skill, they would decrease the returns to skill for the rest of the population. This would presumably amplify inequality. However, solving such a dynamical system would be complicated, and is beyond the scope of this paper.

out at the same skill level, they differ in a_i. That is, some individuals are inherently quick to learn how to use the new technology: given the same level of skills and the same amount of time as others, they can become more proficient in using the new technology. Clearly, unless individuals start out with a positive amount of skills on the new technology, no learning can ever be initiated. This warrants our earlier restriction that $s_0 > 0$.

Finally, assuming individuals have linear preferences over consumption and discount the future at rate ρ and omitting the individual subscript, we can express the problem of each individual in this economy as

$$\max_{n_t} E_0 \left[\int_0^{\infty} z_t s_t \left(1 - n_t\right) e^{-\rho t} dt \right] \tag{1.3}$$

subject to

1. $0 \leq n_t \leq 1$
2. $\dot{s} = a\left(sn\right)^b$
3. z_t follows a Markov process
4. Initial conditions z_0, s_0

Note that we have specified the maximization problem as though agents consume their entire income at each instant. This is because under our assumption of risk-neutrality, individuals will be indifferent about the timing of consumption whenever the interest rate is equal to their discount rate. Since all individuals share a common discount rate, the only equilibrium is one where the interest rate on bonds is the same as the discount rate. Hence, in equilibrium, no individual will wish to borrow or lend even if credit markets were available, and our description of the individual maximization problem (1.3) accords with the proper notion of an equilibrium.

We can solve the above maximization problem using standard techniques in stochastic dynamic programming. Let $V_j\left(s\right)$ denote the expected discounted utility of the individual when he acts optimally, given that his skill level this instant is equal to s and that the level of productivity at this instant is equal to Z_j. This pair of functions for $j \in \{0, 1\}$ must satisfy the system of asset equations

$$\rho V_j\left(s\right) = \max_n \left\{ Z_j s\left(1-n\right) + V_j'\left(s\right) a\left(sn\right)^b + \mu\left[V_{-j}\left(s\right) - V_j\left(s\right)\right] \right\} \tag{1.4}$$

where $V_{-j}(s)$ is the value function evaluated at the same skill level at the complementary level of productivity. Taking the first order condition with respect to n, we have

$$sn = \left[\frac{abV_j'(s)}{Z_j}\right]^{\frac{1}{1-b}} \tag{1.5}$$

Substituting this first order condition into (1.4) yields a system of two differential equations, one for each j:

$$\rho V_j(s) = Z_j s - \left[\frac{abV_j'(s)}{Z_j^b}\right]^{\frac{1}{1-b}} + \frac{1}{b}\left[\frac{abV_j'(s)}{Z_j^b}\right]^{\frac{1}{1-b}} + \mu\left[V_{-j}(s) - V_j(s)\right]$$

To solve this system of equations, we use the method of undetermined coefficients. We guess that the solution of this system is a linear function of s, i.e.

$$V_j(s) = \alpha_j + \beta_j s$$

It can be shown that this solution is in fact unique.[7] Using the solution for β_j and substituting it into the first order condition (1.5) yields the following expression for the growth rate of skills in terms of the primitives for each individual:

$$\frac{ds}{dt} = a^{\frac{1}{1-b}}\left[\frac{b}{\rho}\frac{(\rho+\mu)Z_j + \mu Z_{-j}}{(\rho+\mu)Z_j + \mu Z_j}\right]^{\frac{b}{1-b}} \tag{1.6}$$

A few remarks about (1.6) are in order. First, it implies that within a given regime of aggregate productivity, skills are accumulated at a constant linear rate, i.e. $\frac{ds}{dt}$ is equal to a constant. Thus, the individual acquire a constant amount of new skills each instant, where the number of additional skills depends on the level of aggregate productivity. We should note that since accumulation occurs at a linear rate, a worker contributes to his existing skill base at an ever diminishing rate in percentage terms. This has some consequences for the extent of inequality we could ever observe in this economy, but does not appear to be important for any of our qualitative results. Second, the rate at which an individual accumulates new skills rises with

[7] To prove uniqueness, we can use the phase diagram diagram in Figure A1. This diagram shows that the value of learning is positive and finite only if it is constant (i.e. $x(s)$ must always stay at the steady state). This implies that $V(s) - \frac{Zs}{\rho}$ is a constant, so that $V(s)$ is uniquely linear.

the ability parameter a. Workers who are better able to learn choose to absorb skills at a faster rate than those who are less able, which would lead to a growing divergence in skill levels across workers of different ability. Finally, since $Z_1 > Z_0$, (1.6) implies individuals learn new skills at a faster rate when aggregate productivity is low. This result reflects optimal intertemporal substitution on the part of the agents. Intuitively, individuals take advantage of the fact that they are relatively less productive during recessions to accumulate skills and better prepare themselves for the future, when they expect they will be more productive on average. The claim that investment in skills should be countercyclical has been advanced elsewhere in the literature, although with little or no emphasis on the implications of this result on inequality.[8]

To relate this model to wage inequality, define the wage of an individual as his income each instant divided by the amount of time he spends working. Thus, the wage of individual i is given by

$$w_{it} = z_t s_{it} \tag{1.7}$$

This wage is proportional to the individual's skill level, where the constant of proportionality is the level of aggregate productivity and is thus the same for all individuals. Hence, the wage ratio of individuals at different percentiles will be the same as the ratio of their skill levels. Since individuals with a better ability choose to accumulate skills more rapidly, the distribution of skills fans out over time. This is illustrated graphically in Figure 1, which traces the time path of the skill level of various individuals with different abilities to learn a_i. Since all individuals start out at the same skill level and the change in skills $\frac{ds}{dt}$ is strictly increasing in a regardless of aggregate conditions, the ranking of skills across individuals is preserved over time. Hence, we can think of each path in Figure 1 as the skill level of a particular percentile within the skill distribution (and consequently the wage distribution). As can be seen in the figure, the skill level at higher percentiles drifts apart from the skill level at lower percentiles. Thus, conventional measures of wage inequality such as interpercentile wage ratios (e.g. the ratio of

[8] See, for example, Hall (1991), Cooper and Haltiwanger (1993), Dellas and Sakellaris (1997), Aghion and Saint Paul (1998), and Canton and Uhlig (1999). In much earlier work, Becker (1975) discusses the role of intertemporal substitution for human capital accumulation, although he does not relate his results to aggregate fluctuations. Instead, he argues low wages early in life provide incentive to accumulate human capital early in life. Still, the analogy to low wages during recessions is straightforward. Becker does examine the implications of this channel for inequality over the life cycle. He argues low wages at a young age generate greater inequality among workers when old since they have sorted themselves according to their human capital. This mirrors some of our discussion below for why recessions foster inequality when inequality is increasing. However, the effect of recessions when inequality is decreasing does not have a parallel in Becker's discussion of the life cycle.

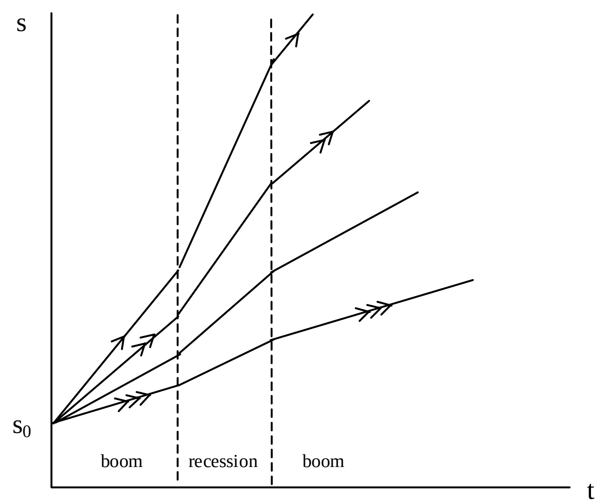

Figure 1: Skill Distribution Over Time

Note: Each path represents the evolution of skills for an individual at a given percentile in the ability (and consequently skill) distribution.

the wage of the 90th percentile to the wage of the 10th percentile) will increase over time as the skill distribution becomes more fanned out. Moreover, the distribution of skills fans out more rapidly during recessions as all individuals shift to acquiring more skills, and so conventional measures of inequality will tend to increase at a faster rate during downturns. We can formalize these observations with the following Proposition, whose proof is contained in an Appendix.

Proposition 1: Consider two individuals i and j where $a_i > a_j$.

1. The wage ratio $\frac{w_{it}}{w_{jt}} = 1$ at $t = 0$. For any realization of z_τ defined for $\tau \in [0, \infty)$, the ratio $\frac{w_{it}}{w_{jt}}$ is increasing and continuous in t, and converges asymptotically to a finite upper bound,

$$\lim_{t \to \infty} \frac{w_{it}}{w_{jt}} = \left(\frac{a_i}{a_j} \right)^{\frac{1}{1-b}} > 1$$

2. For any realization of z_τ where $\tau \in [0, t)$, the change in $\frac{w_{it}}{w_{jt}}$ is decreasing in z_t, i.e.

$$\frac{d}{dt} \left(\frac{w_{it}}{w_{jt}} \right) \bigg|_{z_t = Z_0} > \frac{d}{dt} \left(\frac{w_{it}}{w_{jt}} \right) \bigg|_{z_t = Z_1} > 0$$

Our model therefore describes one particular channel that causes recessions to exacerbate wage inequality. However, this effect depends on the fact that inequality is increasing over the long run: downturns provide greater incentive for agents to sort themselves into different skill levels, but sorting into skill levels is the source of inequality to begin with. In other words, recessions amplify the underlying increase in inequality rather than create it. Breaking down inequality into orthogonal trend and cycle components is entirely artificial in our framework. Of course, there may very well be other channels through which aggregate fluctuations affect inequality over the cycle and which are orthogonal to long-run trends in the distribution of wages, and these may very well have motivated previous researchers to treat long-run trend and cycle as distinct. But this does not deny the mechanism we have described here.

Finally, we should point that our discussion so far has only focused on wage inequality. We can similarly investigate the evolution of income inequality in our model, where income is defined by $y_{it} = z_t s_{it} (1 - n_{it})$. Here, two observations are in order. First, our model abstracts from unemployment, which presumably plays an important role in changes in the distribution of income over the cycle. Thus, any predictions our model makes for income inequality are only appropriate as statements about inequality among those who remain employed during recessions. That said, if we examine interpercentile income ratios in our model, we can establish that the second part of Proposition 1 will continue to hold. That is, income inequality among employed workers rises more rapidly during recessions than during booms. However, in contrast with the first part of the Proposition, income ratios are not continuous, since work hours n_{it} are not continuous over time. Instead, a negative productivity shock will cause income ratios to jump upwards on impact. The reason for this is that when a negative productivity shock hits, workers with high ability will reduce the time they spend working by a greater proportion than less able workers. Hence, on impact, the income of more able workers will fall relative to less able workers when the economy enters a downturn, reducing inequality. But from that point on, interpercentile income ratios will increase at a more rapid rate as the earnings of more able workers grow apart from those of less able workers. As long as such discontinuous jumps are small and observations are sufficiently far apart so that jumps are not easily discernible, income inequality will appear to follow the same pattern as wage inequality.

2. The Case of Decreasing Wage Inequality

In the previous section, we argued that recessions cause the wage distribution to fan out more rapidly. However, it is important to recognize that this does not indicate recessions inherently contribute to greater inequality. Rather, the above model implies recessions amplify an underlying trend towards increasing inequality. This section underscores this point by illustrating that recessions will be associated with faster wage equalization when wage inequality exhibits a declining trend. To allow for a decreasing trend in wage inequality, we modify the model above to include diminishing returns to skill at the individual level. With this assumption, more able individuals eventually slow down in acquiring skills, which allows their less able counterparts to catch up and close the wage gap. The same channel of intertemporal substitution will now cause recessions to hasten the decline in inequality: downturns provide greater incentive to concentrate on skill accumulation for less able workers than for more able workers, and so recessions accelerate the rate at which lower paid workers catch up with their higher paid colleagues.

We should again point out that our observation that ability-biased technical change generates only temporary increases in inequality has been noted by previous authors. Galor and Tsiddon (1997) argue that inequality will eventually decline as minor innovations on new technologies are introduced that allow access to new technologies to less innately productive workers. However, they model the arrival of innovations as exogenous, and so their framework is inadequate for studying how this process of closing wage gaps would be affected by aggregate fluctuations. Caselli (1999) also presents a scenario in which income and wage gaps might disappear in the long run. In his model, ongoing capital accumulation increases the inequality in wages between new and old technologies, drawing increasingly more workers into the new technology. For certain parameter restrictions, this process culminates with all workers in the population switching to the new technology, at which point there is a discrete jump to complete equality. In contrast with his model, our approach allows for a gradual convergence towards equality that depends on the endogenous decisions of workers which can be influenced by aggregate fluctuations. Thus, our model is more suitable for studying the effects of business cycles on inequality when long-run inequality is on the decline. It is precisely the application to aggregate fluctuations that makes our model intriguing: if our argument is correct, it would follow that the main reason economists have found such clear evidence that recessions promote inequality over the past 40 years is because this was a period of rising inequality; had they investigated the same questions under a different lamppost, they would likely have been led to very different conclusions.[9]

To introduce diminishing returns to skill at the individual level, we impose an upper bound on the skill level, i.e. $s_{it} \leq \overline{s}$. Thus, at some point, accumulating additional skills no longer contributes to productivity, an extreme form of diminishing returns. Intuitively, our assumption implies there is only so much to learn about a new technology before one becomes fully proficient in it. Alternatively, if we interpret s_{it} as an ordering of production processes within a general purpose technology, our assumption would reflect the fact that there is some upper bound on how productive any particular production process can be for a given general purpose

[9] Greenwood and Yorukoglu (1997) offer a related explanation for why the trend towards increasing inequality at the onset of a new technology is eventually reversed. In their model, new technologies can only be improved by skilled workers. Workers can become skilled by paying a fixed cost which varies across the population. The initial increase in inequality is due to increased demand for high skill workers. Over time, the technologies are sufficiently productive, demand for skilled workers declines, and inequality will be reversed. Their story also has the feature that the transition into a new technology favors workers who have an easier time becoming skilled, although introducing aggregate fluctuations into their framework appears to be a more complicated undertaking.

technology. Adding even such a minor modification makes the mathematical analysis substantially more complicated. We begin our discussion of this scenario by abstracting from aggregate fluctuations; that is, we consider the case where aggregate productivity is constant. This case turns out to involve a differential equation that can still be solved in closed form. We then introduce business cycles into this framework, which involves a two-dimensional dynamical system that has no closed form solution, at least as far as we know. However, we can still characterize its main properties, which allows us to discuss the evolution of earnings inequality over the cycle when long-run trend inequality is declining.

Introducing an upper bound and assuming a constant level of productivity Z changes (1.3) into the following maximization problem:

$$\max_{n_t} \int_0^{\infty} Z s_t \left(1 - n_t\right) e^{-\rho t} dt \tag{2.1}$$

subject to

1. $0 \leq n_t \leq 1$
2. $\dot{s} = \begin{cases} a\left(sn\right)^b & \text{if } s < \overline{s} \\ 0 & \text{if } s = \overline{s} \end{cases}$
3. Initial condition s_0

When $s_t = \overline{s}$, there is no point to learning, and so individuals will set $n_t = 0$. Thus, an individual will earn a constant flow $Z\overline{s}$ per unit of time once he has reached $\overline{s}$, which discounted at a rate ρ establishes that

$$V\left(\overline{s}\right) = \frac{Z\overline{s}}{\rho} \tag{2.2}$$

For $s < \overline{s}$, we again use techniques of continuous time dynamic programming to solve for each worker's maximization problem.. For $s < \overline{s}$, the asset equation for $V(s)$ is given by

$$\rho V\left(s\right) = \max_n \left\{ Zs\left(1-n\right) + V'\left(s\right) a\left(sn\right)^b \right\}$$

The first-order condition is given by

$$sn = \left[\frac{abV'\left(s\right)}{Z}\right]^{\frac{1}{1-b}}$$

which when substituted back into the asset equation yields the differential equation

$$\rho V(s) = Zs + \frac{1-b}{b}\left[\frac{ab}{Z^b}\right]^{\frac{1}{1-b}} \left(V'(s)\right)^{\frac{1}{1-b}} \tag{2.3}$$

with (2.2) as its boundary condition.

Equation (2.3) is obviously related to the non-linear differential equation in the previous section. However, by imposing a boundary condition at $s = \overline{s}$, we rule out a linear functional form for $V(s)$. Still, this differential equation can be readily analyzed by using a transformation of $V(s)$. Define a new variable $x(s)$ as

$$x(s) = V(s) - \frac{Zs}{\rho}$$

This variable $x(s)$ has a ready economic interpretation: it is the value of the option to learn additional skills when the current skill level is s. This is because if the agent was prevented from learning any additional skills, he would allocate all of his time endowment to working, giving him a discounted utility of $\frac{Zs}{\rho}$. The difference between this and $V(s)$ reflects how much the worker would be willing to pay for the option of learning additional skills beyond his current level. After some algebraic manipulation, we can rewrite (2.3) as a differential equation in terms of $x(s)$:

$$\frac{dx(s)}{ds} = \left(\frac{\rho x(s)}{k}\right)^{1-b} - \frac{Z}{\rho} \tag{2.4}$$

where $k = \frac{1-b}{b}\left[\frac{ab}{Z^b}\right]^{\frac{1}{1-b}}$. The phase diagram for $x(s)$ is illustrated in Figure 2, which plots $x(s)$ against $\frac{dx(s)}{ds}$. This system has a unique steady state at

$$x^* = \frac{k}{\rho}\left(\frac{Z}{\rho}\right)^{\frac{1}{1-b}} \tag{2.5}$$

A simple calculation confirms that x^* above is equal to the value of learning when s is unbounded. Since the value of learning will be weakly lower when s is bounded above than when it is unbounded, we know that when skill is bounded above, it must be the case that $x(s_0) \leq x^*$. Moreover, it is easy to show by contradiction that $x(s_0) \neq x^*$. In terms of Figure 2, this implies $x(s_0)$ will lie to the left of this steady state x^*. In this region, $\frac{dx(s)}{ds} < 0$. To ascertain determine how $x(s_t)$ changes over time, we need to determine the laws of motion for s. As

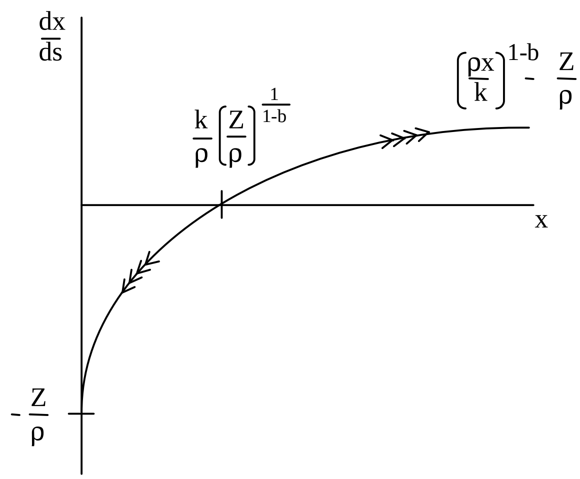

Figure 2: Phase Diagram for the Option Value of Learning (Bounded Skill)

illustrated in the Appendix, we can express the dynamics of s directly in terms of the value of learning $x(s)$:

$$\frac{ds}{dt} = a^{\frac{1}{1-b}} \left[\frac{b}{Z}\right]^{\frac{b}{1-b}} \left(\frac{\rho x(s)}{k}\right)^{b}$$

The above equation has an appealing interpretation, namely that individuals accumulate skills more rapidly when the option to acquire additional skills is particularly valuable, i.e. when $x(s)$ is large.[10] Since the value of this option is positive for $s < \overline{s}$, the individual will choose to accumulate more skills if he has not yet mastered the new technology, i.e. $\frac{ds}{dt} > 0$. Thus, the individual will accumulate more skills and move closer towards $\overline{s}$ over time. However, since $\frac{dx(s)}{ds} < 0$, the option to learn becomes less valuable over time, and the individual accumulates additional skills at a progressively slower rate. Intuitively, this is because our assumptions on the accumulation technology imply there accumulating skills facilitates subsequent accumulation later on. As the individual edges closer to $\overline{s}$ and there are only a few additional skills left to learn, this effect becomes increasingly less important, giving the worker less incentive to accumulate skills over time. The time profile of the skill level of an individual will look like one of the paths illustrated in panel (a) of Figure 3, which traces the evolution of skill for individuals with different ability parameters a, just as in Figure 1 above. The figure is drawn so that the ranking of individuals by skill remains constant over time, i.e. workers who are more able to learn will always rank higher than workers who are less able to learn, and where individuals reach the upper bound $\overline{s}$ in finite time. We confirm that this is in fact the case in the following Proposition.

Proposition 2: For an individual with arbitrary learning parameter $a > 0$, we have

1. The time path of the skill level s_t is concave in t, and $\lim_{t\to\infty} s_t = \overline{s}$.

2. There exists a $T < \infty$ such that $s_{it} = \overline{s}$ for all $t \geq T$, i.e. the upper bound on s is attained in finite time.

3. For any two individuals i and j where $a_i > a_j$, we have $s_{it} \geq s_{jt}$ for all t, with strict inequality for $s_0 < s_{it} < \overline{s}$.

[10] While this result is intuitive, it can be misleading. When we allow for business cycles, we can show that for s sufficiently close to $\overline{s}$, the individual will accumulate more skills during recessions even when the option to learn is more valuable during booms. We return to this point below.

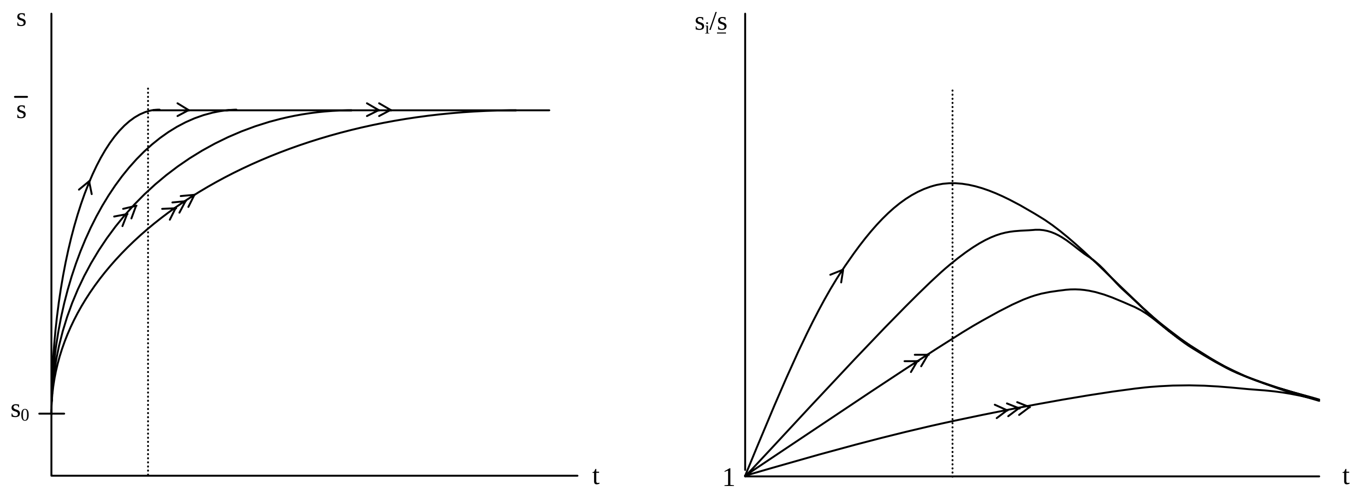

(a) (b)

Figure 3: Skill Levels and Ratios with Bounded Skills

Note: Each path in panel (a) represents the evolution of skills for an individual at a given percentile in the ability (and consequently skill) distribution. Each path in panel (b) represents the ratio of skill between a given percentile in the ability distribution and the skill level of the least able individual in the population.

Since the last part of Proposition 2 implies skill levels can be ordered by ability, we can treat each path in Figure 3(a) as the path of skill for a particular percentile within the skill distribution. Moreover, since the wages of all workers are equally proportional to skill at all times, Figure 3(a) captures up to a scaling factor the paths of wages across the population over time. As is clear from the figure, allowing for an upper bound generates a pattern of wage inequality that is consistent with Kuznets' (1955) discussion of the relationship between growth and the distribution of earnings: as productivity grows over time, inequality first rises and then falls. The reason this occurs in our framework is that when the new technology first arrives, individuals learn at different rates how to use it, causing the distribution of skills across individuals to fan out. But eventually, workers who are more able to learn slow down as they move closer to the maximum skill level $\overline{s}$. Productivity continues to grow, albeit at a slower rate than before, but wage inequality declines as less able workers begin to catch up to their peers. At this point, the wage distribution becomes more compressed. An alternative way of illustrating this reduction in inequality is by looking at skill ratios rather than skill levels, as illustrated in panel (b) of Figure 3. This figure shows the ratio of the skill level of a given individual relative to the skill level of the individual with the lowest ability parameter $\underline{a}$. For each percentile, the ratio begins at 1, rises up to some maximum level, and then declines gradually back to 1 as even the slowest individual in the economy reaches $\overline{s}$.[11] Note that our model predicts the distribution of wages becomes compressed from the top of the distribution down. In particular, relative wages begin to decline first at the top of the distribution as individuals there slow down in their accumulation of skills. At the same time, the relative wages of workers in the middle of the distribution continues to rise, increasing the gap between the middle and the bottom of the distribution while decreasing the gap between the middle and the top of the distribution.

To study the effect of business cycles when the wage distribution is becoming compressed, we introduce fluctuations in aggregate productivity into this model. While our claim that recessions should accelerate the process of wage convergence may appear to be a straightforward extension of our previous arguments, establishing it formally is not trivial. The remainder of this section traces our steps in establishing this claim, along with some observations about the nature of

[11] The fact that all individuals eventually reach the same skill level is clearly unrealistic. This result is due to our assumption that all individuals have a common ceiling $\overline{s}$. Our analysis would continue to go through if each individual had his own ceiling $\overline{s}_i$, provided that these ceilings still generate some eventual compression in the distribution of wages across individuals (i.e. individuals with a better capacity to learn must reach their ceilings first to allow less able individuals to catch up, at least in part).

our result that will be useful in interpreting empirical evidence later on. We begin with the claim that just as in the case of unbounded skill accumulation, workers will acquire skills at a more rapid rate during recessions than during booms.

Proposition 3: Consider an individual with a skill level $s < \overline{s}$. For any realization z_τ defined for $\tau \in [0, t)$, the individual will accumulate skills more rapidly if aggregate productivity at date t is low, i.e.

$$\left.\frac{ds_t}{dt}\right|_{z_t = Z_0} > \left.\frac{ds_t}{dt}\right|_{z_t = Z_1}$$

The above Proposition states that regardless of history of aggregate productivity, as long as a worker is still accumulating skills (i.e. his skill level is below the bound $\overline{s}$), he will learn at a faster pace during a recession than during a boom. The intuition, once again, is due to intertemporal substitution: it is optimal to concentrate skill acquisition in less productive times. However, the extension of this argument to the case where skills are bounded above is not obvious. In particular, when a worker is already near the upper bound, the option to learn becomes more valuable during booms, i.e. $x_1(s) > x_0(s)$ for s sufficiently close to $\overline{s}$. This is because when the agent has only a few more skills left to learn, the value of acquiring additional skills derives primarily from the additional output the worker can produce rather than from the fact that a higher skill base facilitates the acquisition of further skills. But a worker is more productive during a boom than during a recession. This could in principle change the incentives of the worker to 'reach for the finish line' more quickly when aggregate productivity is high. Proposition 3 establishes that this is not the case under our assumptions.

The next step is to apply Proposition 3 to argue that recessions serve to reduce inequality between less able individuals who are still far behind and more able individuals who are already close to the upper bound $\overline{s}$ or have already reached it. However, it is difficult to compare learning speeds across workers of different ability in the absence of closed form expressions for the learning rate. Our approach is to first establish that even with aggregate productivity shocks, individuals reach $\overline{s}$ in finite time. After an individual has already exhausted all possible learning, it is clear that his wage relative to the wages of his less able colleagues will decline: his wages stop growing altogether, while those of his less able colleagues continue to grow.

Likewise, it is obvious that the decline in the relative wage of more able workers will be more rapid during recessions. A simple continuity argument establishes that both patterns continue to hold when the more able individual is sufficiently close to $\overline{s}$ but has not reached it yet. This is captured in the following Proposition:

Proposition 4: Consider two individuals i and j where $a_i > a_j$. Then $s_{it} \geq s_{jt}$ for all t, with strict inequality if $s_0 < s_{it} < \overline{s}$. Moreover, there exists an $\varepsilon > 0$ such that

1. If $s_{it} - \overline{s} = \varepsilon$, then $\dfrac{w_{it}}{w_{jt}}$ decreases with t for any continuation realization z_s defined over $s \in [t, \infty)$.

2. Wage inequality declines more rapidly during recessions, i.e. for any realization z_τ defined over $\tau \in [0, t)$ which is compatible with $s_{it} \leq \overline{s} - \varepsilon$, we have

$$\frac{d}{dt}\left(\frac{w_{it}}{w_{jt}}\right)\bigg|_{z_t = Z_0} < \frac{d}{dt}\left(\frac{w_{it}}{w_{jt}}\right)\bigg|_{z_t = Z_1} < 0$$

Note that Proposition 4 only characterizes the effect of recessions on inequality between any two individuals when the more able individual is close to the bound $\overline{s}$. Thus, a recession could very well accelerate wage convergence between the very top percentiles and the remainder of the distribution, even as it accelerates wage dispersion within the bottom of the distribution, where individuals are far away from the upper bound and wage inequality is still growing. In looking for evidence on how cyclical fluctuations affect the distribution of wages, we must keep in mind that the effects may be dissimilar at different parts of the earnings distribution.

To summarize, we have presented a model in which the dynamics of individual wages depend on the endogenous decisions of workers on the rate at which to adapt to a new technology. We showed that our model could be made compatible with either increasing or decreasing wage inequality across a population of heterogeneous workers. When we allow for fluctuations in aggregate productivity in this model, we find that recessions tend to amplify whatever trend is already occurring within the population: if inequality is increasing because of the choices made by different agents, recessions create incentives for individuals to behave in a more extreme fashion that generates greater inequality. This prediction should caution against trying to infer general properties about the way in which business cycles affect inequality. Studying the effects of business cycles on the distribution of earnings requires taking into account broader trends that affect the wage distribution and which may interact with cyclical fluctuations.

3. Wage Inequality and the Great Depression

Obviously, there is little we can do to confirm our conjecture that the effect of recessions on wage inequality hinges on long-run trends in wage inequality from data taken over the last 40 years, since this period was associated with increasing inequality in the distribution of wages. However, the history of wage inequality for the United States includes various episodes in which wage inequality declined over extended periods. For example, economic historians have identified the middle of the 20th century as a period of declining wage and income inequality; Williamson and Lindert (1980) refer to this period as the Income Revolution, while Goldin and Margo (1992) refer to it as the Great Compression. It so happens that this period includes the Great Depression of 1929. This allows us to examine the effects of a particular cyclical episode at a time of declining wage inequality, and ask whether the fact that the Depression occurred in an era of decreasing earnings inequality caused it to hasten the decrease in wage inequality as our model predicts.

The middle part of the 20th century appears to be a particularly relevant time period to test the predictions of our model. First, as noted by David (1991), the beginning of the century marked the introduction of a new general purpose technology, namely electricity. This technology was gradually phased in between 1910-1920. At this time, income and wage inequality increased, just as the model predicts should occur at the onset of a new technology. However, various measures of earnings inequality peaked around the late 1920s and early 1930s, at which point these measures declined dramatically. This compression in the distribution of earnings continued through World War II and beyond it, leading historians to cast doubt on either government policies or World War II as the primary source of shrinking inequality. Instead, they have emphasized the role of technological progress among less advanced sectors and workers; this, for example, is the theme advanced by Williamson and Lindert (1980) in their analytical study of this period. This era therefore appears to capture the forces we describe in our model. In what follows, we interpret empirical evidence on earnings during this period in light of the model we developed above. First, we examine whether and how the distribution of earnings changed over the course of the Great Depression to establish that the Depression did in fact occur in a period of declining wage inequality. Second, we turn to the role of aggregate fluctuations by examining whether the Depression was associated with a more rapid decline in inequality. Finally, we examine whether the Depression was a unique period of accelerated compression in wages or whether business cycles generally had a different impact on earnings inequality prior to World War II.

To test the predictions of the model outlined in the previous sections, we would need data that allows us to trace the ratio of wages for particular percentiles in the wage distribution. Such comprehensive data on the entire distribution of wages is unfortunately not available for the relevant time period; but economic historians have carefully constructed wage series for particular occupation groups that may serve as a reasonable substitute. Williamson and Lindert (1980) compile series on the wage rates of self-employed physicians, associate professors, public school teachers, skilled workers in building trades, and skilled workers in manufacturing, all expressed relative to the wage rates of a reference group of unskilled workers. Goldin and Margo (1992) compile additional wage series for the relative pay of skilled and unskilled workers in the railroad industry, as well as for clerical workers in New York relative to unskilled workers. Since all of the wage series together cover a diverse set of occupations, we can try to use the wages of occupations associated with different pay levels as proxies for the paths of different wage percentiles. In doing so, however, we need to be mindful that the wage rates for particular occupations may be affected by circumstances specific to those occupations but which are not representative of changes in the earnings of a particular wage percentile.

Figure 4 illustrates the wage ratios of the various occupations described above. Some of the relative wage series are not directly comparable since they are expressed relative to different reference groups. But the reference group is always comprised of unskilled workers in particular occupations, which tend to be earn roughly similar wages. The wages of physicians rank at the top, followed by associate professors, clerks, teachers, and then skilled workers in building trades and manufacturing. Since the ranking of these occupations is relatively stable over time, it seems reasonable to assume that different occupations are drawn from different percentiles of the wage distribution, and hopefully reflect general trends at these percentiles. Since unskilled

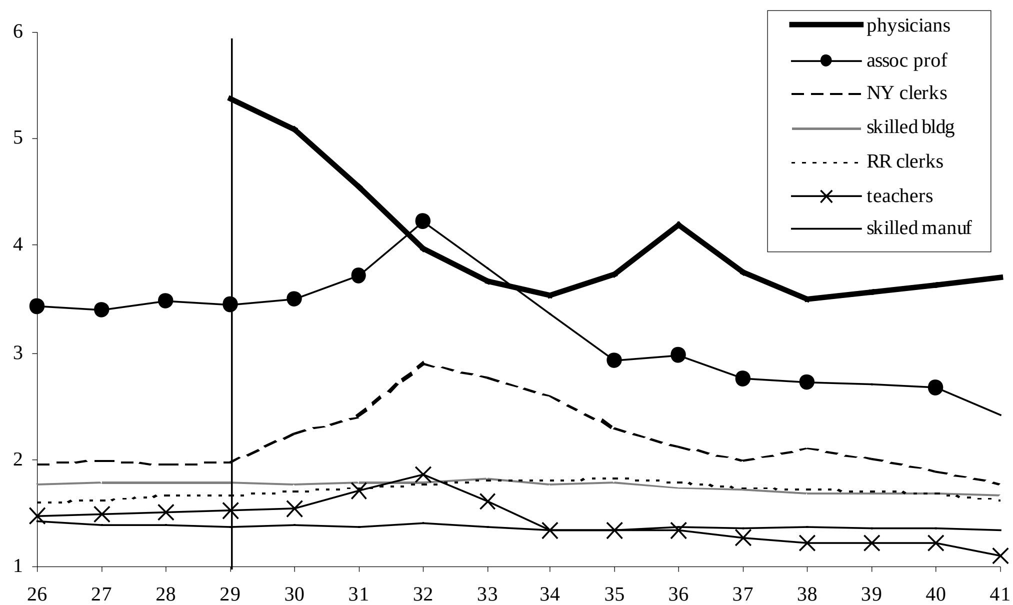

Figure 4: Relative Wages during the Great Depression

Notes: The wage series for public school teachers, associate professors, self-employed physicians, and skilled workers in building trades are series (8), (10), (11), and (13) respectively in Willimason and Lindert (1980). The first three series compare the annual income of teachers, professors, and physicians with 2000 hours of work at the hourly wage rate of unskilled workers basd primarily on data reported by the National Industrial Conference Board (NICB). The series for skilled workers in building are hourly wages of skilled workers \relative to the hourly wage rates of journeymen, helpers, and laborers in the building trades for various cities. The series on skilled wokrers in manufacturing, clerks in the railroad industry, and clerks in New York State are columns (1), (4.2), and (6) respectively from Table VII of Goldin and Margo (1992). The wages of skilled manufacturing and clerks in New York are expressed relative to the unkilled NICB wages, although the first series is based on hourly wage rates while the latter is based on weekly wage rates. The series on clerks in the railroad industry denote hourly wages of clerks relative to machinists and laborers.

labor at the bottom of the earnings distribution, Figure 4 should generate analogous patterns to those illustrated in Figure 3(b). Indeed, the paths of relative wages from the onset of the Great Depression bear close resemblance to the right half of Figure 3(b): the wages of the highest paid workers, physicians, decline relative to the wages of other workers, while the wages of occupations in the middle of the distribution, including professors, clerks, and teachers drift apart from those of unskilled workers but grow closer to those of the highest paid workers. That is, the wage distribution compiled from occupational wages suggests a significant compression at the top of the distribution, just as our model predicts. However, at least between 1929 and 1932, relative wages decline only for one occupation, physicians. This could reflect the fact that relative wages decline only at the top of the distribution first, but it may equally reflect a fall in relative demand for physician services during the Depression that had nothing to do with patterns in the distribution of earnings. This conjecture is contradicted by evidence on the health care industry during this time period. In her chronicling of the history of hospitals, Stevens (1989) observes:

> Overall, despite the increased burden of free care, hospitals fared less badly than many other sectors of the economy. To some extent, the Depression provided a shake-up and consolidation of the hospital system, emphasizing the advantage of concentrating the fixed costs of radiology, pathology, and other technical services in larger institutions; the average general hospital grew from 84 beds in 1929 to 104 in 1940. And although patient payments to nongovernmental hospitals and sanitoria (of all types) declined by 17 percent between 1929 and 1933, this was less than half the decline in consumer expenditures as a whole. Patient payments to governmental hospitals of all types actually increased by 21 percent between 1929 and 1933, as local general hospitals and state and local psychiatric hospitals sought free-paying patients. None of these figures, moreover, includes adjustments for declining prices. Even under Depression conditions, then, there was a promising market for new forms of consumer financing, including prepayment of hospital insurance. (p148-9)

Stevens further argues that demand for health care, particularly psychiatric care, may very well have increased during recessions. The option to work in hospitals should have similarly prevented a decline in the relative wages of self-employed physicians. This leaves wage gains among workers in the bottom part of the earnings distribution as a viable explanation for the patterns in Figure 4. The relative wages of other occupations such as professors, clerks, and teachers also decline beginning before the trough of the Depression in 1933, and continue to decline throughout the decade even though productivity, employment, and output all recover very slowly during this time.[12] Even if much of the compression in the middle of the wage distribution occurred near the end of the Depression and in the subsequent recovery phase, it is striking that individuals at the bottom of the distribution were able to gain so much relative to workers further up in the wage distribution despite the continuing dire economic conditions throughout the 1930s. This is quite a departure from the experience of low wage workers during recessions in the latter part of the century.

To check the robustness of our inferences from wage data, we next turn to evidence on income inequality. Recall that income inequality and wage inequality should roughly evolve in parallel, although there may be some discrepancies because of changes in hours worked on the impact of aggregate productivity shocks. Still, we would expect that trends in wage inequality should be mirrored in income inequality when looking over a period of several years. For this, we would need data on labor income for a representative sample of employed workers. Once again, the data historians have been able to collect has some important shortcomings for our purposes. The primary data sources on income for this period are tax returns, which are analyzed extensively by Kuznets (1953). While this source provides a representative sample of income earners, albeit a distorted one that may be affected by differential patterns in tax avoidance, a major problem with income tax data is that it includes all sources of income, including non-labor income which is orthogonal to the forces we describe in our model. This is particularly troublesome in light of the stock market crash of 1929, whose effects were probably more pronounced for the top income groups. Fortunately, Kuznets defines a measure of "economic income" that excludes changes in asset prices. This measure continues to include non-labor income that is not tied to asset values. Still, if the patterns of pay ratios suggested in Figure 3 are representative of changes in the wage distribution, we might expect to see them reflected in income measures that include labor income. In particular, if the distribution of wages becomes compressed at the top, the share of income of the highest income group should have also declined. Figure 5 illustrates the share of income that belongs to the top 1% and the top 2-5% of the population, both of which are taken from Kuznets (1953). These figures confirm the pattern from wage

[12] Cole and Ohanian (1999) discuss the exceptional nature of the slow recovery from the Great Depression.

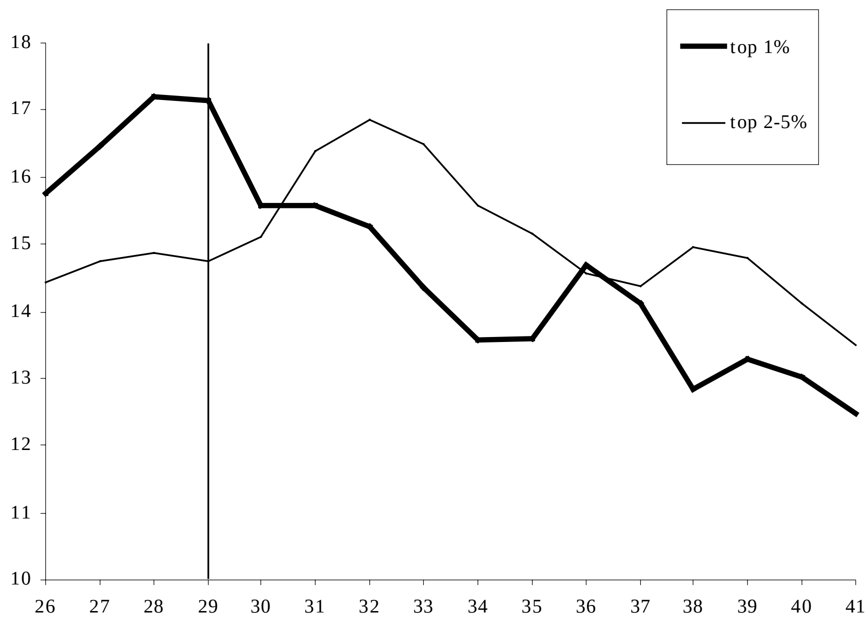

Figure 5: Income Shares of Top Taxpayers during the Great Depression

Notes: Each figure denotes the percentage of total reported income in each year that belonged to the top x% of all taxpayers. The data are taken from Table 122 in Kuznets (1953).

data: the incomes share of the very top of the distribution began to decline at the onset of the Depression, consistent with evidence on the relative wages from occupational groups. The income share of workers further down the earnings distribution declined only from 1932, which matches the patterns in wages for occupations in the middle of the earnings distribution, and continued to decline at a significant pace during the sluggish recovery in the remainder of the decade.

The evidence on wage ratios and income shares above are consistent with our prediction that the earnings distribution should eventually become compressed, starting from the top of the distribution down. However, our primary concern is not whether the compression of wages occurred during the Great Depression, but whether there was a *more rapid* compression of wages during the Great Depression than otherwise. McLean (1991) makes a similar distinction between the effects of the Great Depression proper from the observation that inequality declined over the longer period between 1929 and 1950. He concludes that the Great Depression contributed to earnings inequality rather than ameliorating it. However, in reaching this conclusion, McLean implicitly assumes cyclical fluctuations have the same effect on relative wages at different parts of the wage distribution. Our model instead suggests that recessions should amplify compression at the top of the distribution even as they facilitate the fanning out of the bottom of the wage distribution. The relevant question, then, is whether the Depression was associated with a more rapid decline in inequality at the top of the distribution. To address this question, consider first Table 1, which is taken directly from Kuznets (1953). This table reports the changes in the share of total income that belongs to the upper income groups. Our model predicts that when this share is declining, it should decline more rapidly during recessions than during booms. During the Great Depression, i.e. between 1929 and 1932, the share of the top 1% declined at a rate of 0.53% per year. In the subsequent recovery between 1932 and 1937, it actually increased by 0.02% per year. In the next contraction between 1937 and 1938, the share of income of the top 1% fell by 1.46% per year, while in the subsequent recovery it fell by only 0.48% per year. Turning to the top 5% of taxpayers, the Depression was associated with only a small decline in their share of total income, much smaller than the decline in inequality in the subsequent recovery. However, the subsequent contraction in 1937-1938 was associated with a sharp decline in the income share of this group that exceeded the comparable decline in the subsequent expansion: 1.13% per year compared with 1.02% per year. This conforms with our claim that recessions should accelerate the pace of compression in earnings among workers at the top of the distribution first, and only then at lower parts of the distribution.

Table 1: Change per Year in Shares of Upper Income Groups during Business Cycles

year	cycle	top 1%	top 5%
1918-19	contraction	0.27	0.44
1919-20	expansion	-0.50	-0.84
1920-21	contraction	1.16	3.40
1921-23	expansion	-0.61	-1.29
1923-24	contraction	0.63	1.40
1924-26	expansion	0.51	0.48
1926-27	contraction	0.46	0.72
1927-29	expansion	0.05	0.06
1929-32	contraction	**-0.53**	-0.03
1932-37	expansion	0.02	-.038
1937-38	contraction	**-1.46**	**-1.13**
1938-44	expansion	-0.48	-1.02

Source: Kuznets (1953), Table 15

A related analysis using the same income tax data is presented in McLean (1991). He uses the methodology of Blinder and Esaki (1978) and Blank and Blinder (1986) to assess the effects of cyclical fluctuations on the distribution of earnings. Essentially, this involves regressing the income share of various quintiles on a time trend and unemployment, where the latter captures the effects of macroeconomic fluctuations. McLean repeats this exercise for the share of the top income groups from Kuznets' data. This makes for a particularly interesting comparison, since the approach is identical to more recent empirical work that has argued recessions exacerbate income inequality. In contrast with findings for more recent periods, McLean finds the opposite pattern (although statistically insignificant) for the top income group during the Great Depression, i.e. their share of income declines when unemployment rises. This reversal is precisely what our model predicts.

The above evidence appears to suggest that the Great Depression did not have the same impact on the distribution of wages as recessions in the latter part of the century. At the very least, this bolsters our claim that business cycles do not have a consistent or uniform effect on the earnings distribution, and cautions against drawing too much inference from evidence during the past 40 years. However, we have specifically tried to argue that a reason business cycles do not have a uniform effect on wage inequality is because of patterns in long-run wage inequality. This raises the question of whether the difference between the Great Depression and more recent cyclical fluctuations attributable to differences in long-run trends in the wage distribution. An alternative hypothesis is that differences in changes in the distribution of earnings during the Great Depression reflect more general differences between modern business cycles and those that occurred under very different institutional and policy environments that existed before World War II. Table 1 provides some evidence that bears on this hypothesis. The two contractions in 1929-1933 and 1937-1938 occurred at a time when the share of income among the top 1% of taxpayers was declining. Both contributed to a more rapid decline in the share of the top 1%. But contractions earlier in the century occurred when the income share of the top 1% was still increasing, and contributed to a more rapid increase in the share of the top 1%. Thus, recessions that occurred prior to World War II but during a time of increasing income inequality had a similar impact on the distribution of wages as modern business cycles in the last 40 years have had. Further evidence that inequality intensified during recessions prior to the Great Depression comes from evidence on wage data. Williamson and Lindert (1980) argue that during the 'uneven plateau' between the Civil War to the Great Depression, the wages of various occupation groups relative to those of unskilled labor tended to be countercyclical. They write:

> Like our measures of income dispersion, the pay ratios show a generally countercyclical pattern. The pay ratios tend to drop in booms and to rise in recessions. This tendency is much more pronounced when the boom or contraction comes rapidly than when it takes a few years to gather momentum. This tendency is also a little more pronounced for teachers, whose pay contracts are longer in term. (p82)

In sum, historical evidence provides some support for our contention that business cycles do not have a uniform effect on earnings inequality: the Great Depression accelerated the closing of income gaps between the top of the earnings distribution and the remainder of the distribution, and inequality further down the earnings distribution accelerated in the subsequent contraction later on that decade. More generally, the distribution of earnings became much more compressed during the 1930s despite dismal economic conditions and a very weak recovery between 1932 and 1937. The evidence is also consistent with our claim that what distinguished the Great Depression from other recessions was the fact that long run inequality was declining: contractions that occur both before and after the Depression when inequality was on the rise contribute to greater inequality, unlike the Depression which occurred at a time of decreasing inequality.

4. Conclusion

An important aspect of our paper is the usual admonition against ignoring the lessons of history: whereas recent work has argued that recessions contribute to increasing inequality, economic historians have already come to appreciate that the Great Depression was a period of declining rather than increasing inequality, as well as the fact that the Depression may have even contributed to more rapid compression at the very top of the earnings distribution. However, aside from introducing historical evidence into the discussion on how business cycles affect the earnings distribution, our paper provides a consistent theoretical framework that can reconcile modern evidence on the effects of business cycles with the glaring absence of increasing earnings inequality during the severe Great Depression. Specifically, our explanation argues that recessions merely amplifying underlying trends in earnings inequality rather than inherently contributing to greater earnings dispersion. This suggests that using simple deviations from trend to identify the effects of business cycles may be very misleading over long horizons which include periods of both increasing and decreasing inequality.

We close with two observations. First, our model only focuses on a particular mechanism, according to which long run trends and cyclical fluctuations in the wage distribution are related. There are other mechanisms through which business cycles could affect the earnings distribution which are entirely unrelated to long-run trends. For example, Beaudry and DiNardo (1991) construct a model in which booms reduce earnings inequality because they induce contract renegotiation among lower paid workers which allows them to extract wage increases from their employers. Barlevy (1999) argues that recessions generate a first-order stochastic shift of the distribution of wages to the left by slowing down the rate at which workers move into jobs where they are more productive. Even some of the work on the welfare effects of business cycles is premised on the claim that recessions increase income dispersion, which implies individual consumption streams are more volatile and therefore more unpleasant than is reflected by per-capita consumption. Our analysis does not deny any these effects. But even if these changes are orthogonal to long-run trends, our results higlight the fact that the earnings distribution does not evolve in a vacuum. Identifying any effects of business cycles on the distribution of wages, regardless of the particular channel they represent, requires careful treatment of whatever governs underlying changes in the wage distribution.

Second, our analysis focuses on inequality among employed workers. Thus, we ignore variation in income due to changes in unemployment. This is probably an important component in determining earnings inequality, since unemployment is highly cyclically sensitive, and at least in recent times has tended to be concentrated among lower paid workers. By omitting unemployment, we are probably neglecting an important component that determines income inequality. But we view our approach as providing a foundation of a model where these additional features should be incorporated, especially since it seems to be quite useful in explaining variation in inequality among employed workers at different phases of U.S. history.

Appendix

Proof of Proposition 1: not yet written up...

Proof of Proposition 2: Although one can obtain all of these results using general characterization results directly on (2.4), it is easier to establish our claims by transforming (2.4) into a dynamical system where x is directly a function of t (as opposed to s_t). The latter can be solved in closed form, which yields the various statements of the Proposition almost directly. We first require the laws of motion of s, using the first order condition in the text:

$$\begin{aligned}
\dot{s} &= a\left(sn\right)^{b} \\
&= a\left(\frac{abV'\left(s\right)}{Z}\right)^{\frac{b}{1-b}} \\
&= a\left(\frac{ab}{Z}\left[x'\left(s\right)+\frac{Z}{\rho}\right]\right)^{\frac{b}{1-b}} \\
&= a\left[\frac{ab}{Z}\left(\frac{\rho x\left(s\right)}{k}\right)^{1-b}\right]^{\frac{b}{1-b}} \\
&= a\left[\frac{\rho b x\left(s\right)}{\left(1-b\right)Z}\right]^{b}
\end{aligned}$$

We then substitute this into (2.4) to generate an equation for $\frac{dx}{dt}$:

$$\begin{aligned}
\frac{dx}{dt} &= \frac{dx}{ds}\cdot\frac{ds}{dt} \\
&= \left(x\left(\frac{\rho x}{k}\right)^{1-b}-\frac{Z}{\rho}\right)a\left[\frac{\rho b x}{\left(1-b\right)Z}\right]^{b} \\
&= k^{b-1}\left[\frac{b}{\left(1-b\right)Z}\right]^{b}a\rho x-a\left(\frac{Z}{\rho}\right)^{1-b}\left(\frac{b}{1-b}\right)^{b}x^{b} \\
&= \frac{\rho}{1-b}x-\left[a\left(\frac{Z}{\rho}\right)^{1-b}\left(\frac{b}{1-b}\right)^{b}\right]x^{b}
\end{aligned}$$

This differential equation is known as a Bernoulli equation, which can be solved in closed form using a simple transformation of variables. In particular, define $y = x^{1-b}$. Then we have

$$\frac{dy}{dt} = \rho y - \left[a \left(\frac{(1-b)\, Z}{\rho} \right)^{1-b} b^b \right]$$

which has the solution

$$y = \left[a \left(\frac{(1-b)\, Z}{\rho} \right)^{1-b} \left(\frac{b}{\rho} \right)^b \right] + Ce^{\rho t} \tag{4.1}$$

for some constant C. From the phase diagram, we know that $y \downarrow 0$, which establishes part of the first statement in the proposition, but also implies $C < 0$. Since

$$\dot{s} = a \left[\frac{\rho b}{(1-b)\, Z} \right]^b y^{\frac{b}{1-b}}$$

and $\dot{y} < 0$, then $\ddot{s} = \frac{ab}{1-b} \left[\frac{\rho b}{(1-b)\, Z} \right]^b y^{-\frac{1}{1-b}} \dot{y} < 0$, establishing s_t is concave. The fact that $C < 0$ insures there exists a finite T for which $y_T = 0$. Finally, we can apply the fact that two non-identical concave graphs intersect in at most two points and the observation that $s_{i0} = s_{j0} = s_0$ and that there exists a T large enough such that $s_{jt} = s_{it} = \overline{s}$ for all $t > T$ to conclude that either $s_{it} \geq s_{jt}$ or $s_{it} \leq s_{jt}$ for all t, with strict inequality for $s_0 < s_{it} < \overline{s}$. Since the option to learn is weakly more valuable for those who are more able, it follows that $y_{i0} \geq y_{j0}$. This implies $\dot{s}_i > \dot{s}_j$ at $t = 0$, which establishes the claim. ■

Proof of Proposition 3: For an individual with skill level $s < \overline{s}$, the asset equations are just as before:

$$\rho V_j(s) = \max_n \left\{ Z_j s (1-n) + V_j'(s)\, a\, (sn)^b + \mu \left[V_{-j}(s) - V_j(s) \right] \right\}$$

with a first order condition

$$sn = \left[\frac{ab V_j'(s)}{Z_j} \right]^{\frac{1}{1-b}}$$

which can be substituted in to yield a system of differential equations

$$\rho V_j(s) = Z_j s + \left(\frac{1-b}{b} \right) \left[\frac{ab}{Z_j^b} \right]^{\frac{1}{1-b}} \left(V_j'(s) \right)^{\frac{1}{1-b}} + \mu \left[V_{-j}(s) - V_j(s) \right]$$

Define $x_j(s)$ as the value of the option to learn additional skills:

$$x_j(s) = V_j(s) - \frac{(\rho+\mu) Z_j + \mu Z_{-j}}{\rho(\rho+2\mu)} s$$

We can rewrite the asset equation in terms of $x_j(s)$:

$$x_j'(s) = \left[\frac{(\rho+\mu) x_j(s) - \mu x_{-j}(s)}{k_j}\right]^{1-b} - \frac{(\rho+\mu) Z_j + \mu Z_{-j}}{\rho(\rho+2\mu)}$$

where $k_j = \frac{1-b}{b}\left[\frac{ab}{Z_j^b}\right]^{\frac{1}{1-b}}$. The phase diagram for (x_0, x_1) is illustrated in Figure A1, which is defined as long as $(\rho+\mu) x_j(s) - \mu x_{-j}(s) > 0$, a condition which holds at the optimum and which insures $[(\rho+\mu) x_j(s) - \mu x_{-j}(s)]^{1-b}$ is well-defined.

Using the first-order condition above, the rate at which the individual acquires skills given a particular realization of productivity $z_t = Z_j$ is given by

$$\left.\frac{ds_t}{dt}\right|_{z_t=Z_j} = a(sn)^b = a\left(\frac{b}{1-b}\right)^b \left[\frac{\rho x_j(s) + \mu[x_j(s) - x_{-j}(s)]}{Z_j}\right]^b$$

From this equation, it follows that the agent acquires skills more rapidly in a recession, i.e. when $z_t = Z_0$, if and only if

$$\frac{\rho x_0(s) + \mu[x_0(s) - x_1(s)]}{Z_0} \geq \frac{\rho x_1(s) + \mu[x_1(s) - x_0(s)]}{Z_1}$$

or

$$x_1 \leq \frac{(\rho+\mu) Z_1 + \mu Z_0}{\mu Z_1 + (\rho+\mu) Z_0} x_0 \tag{4.2}$$

which describes a region in (x_1, x_0) space whose boundary is a line with slope greater than 1 that cuts through the origin. The boundary is illustrated as a dotted line in the phase diagram of Figure A1, and the region where agents acquire skills more rapidly during recessions is to the right of this half-line.

Next, since $\left.\frac{ds_t}{dt}\right|_{z_t=Z_j} \geq 0$ and $x_j'(s) \leq 0$ (i.e. the option value of learning must weakly decrease with s), the phase diagram implies that when the agent behaves optimally, $(x_0, x_1) \to (0, 0)$,

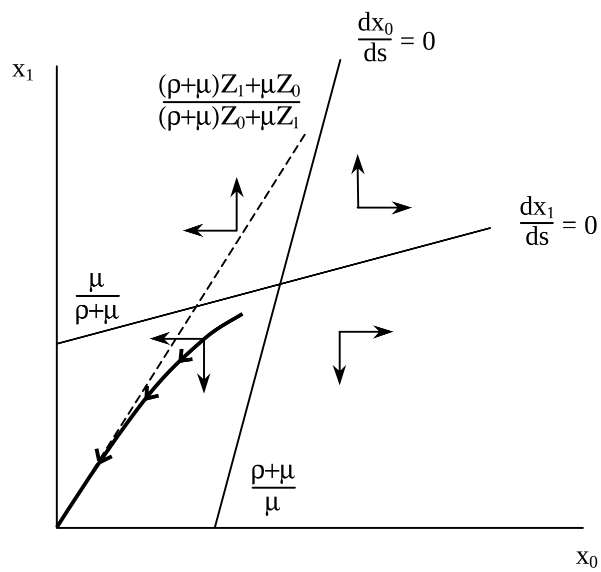

Figure A1: Dynamics with Aggregate Shocks and Bounded Skill which implies $s \to \overline{s}$ for all realizations of z_t. At this point, we can apply L'Hopital's rule to evaluate

$$\lim_{s \to \overline{s}} \frac{\rho x_0(s) + \mu[x_0(s) - x_1(s)]}{\rho x_1(s) - \mu[x_1(s) - x_0(s)]} = \frac{\rho x_0'(s) + \mu[x_0'(s) - x_1'(s)]}{\rho x_1'(s) - \mu[x_1'(s) - x_0'(s)]} = \frac{Z_0}{Z_1}$$

Thus, as $(x_0(s), x_1(s)) \to (0,0)$, the speed of learning in both levels of productivity converges to the same rate. Graphically, the path of $(x_0(s), x_1(s))$ that is consistent with optimal behavior is tangent to the line defined by (4.2), as illustrated by the thick line in Figure A1. We now establish that along the optimal path $(x_0(s), x_1(s))$ for $s < \overline{s}$, (4.2) is satisfied, i.e. the thick line in Figure A1 does in fact lie to the left of the dashed line as illustrated. To see this, pick a point (x_0, x_1) for which (4.2) holds, and evaluate the derivatives $x_0'(s)$ and $x_1'(s)$ that are consistent with optimality, i.e. which satisfy the asset equations:

$$\begin{aligned}
x_0' &= \left[\frac{(\rho+\mu)x_0 - \mu x_1}{k_0}\right]^{1-b} - \frac{(\rho+\mu)Z_0 + \mu Z_1}{\rho(\rho+2\mu)} \\
x_1' &= \left[\frac{(\rho+\mu)x_1 - \mu x_0}{k_1}\right]^{1-b} - \frac{(\rho+\mu)Z_1 + \mu Z_0}{\rho(\rho+2\mu)} \\
&= \left[\frac{Z_1}{Z_0}\frac{k_0}{k_1}\frac{(\rho+\mu)x_0 - \mu x_1}{k_0}\right]^{1-b} - \frac{(\rho+\mu)Z_1 + \mu Z_0}{\rho(\rho+2\mu)} \\
&= \frac{Z_1}{Z_0}\left[\frac{(\rho+\mu)x_0 - \mu x_1}{k_0}\right]^{1-b} - \frac{(\rho+\mu)Z_1 + \mu Z_0}{\rho(\rho+2\mu)}
\end{aligned}$$

and for which $x_j'(s) < 0$. If this path satisfies the first order condition, then

$$\begin{aligned}
\frac{dx_1}{dx_0} &= \frac{dx_1/ds}{dx_0/ds} \\
&= \frac{(\rho+\mu)Z_1 + \mu Z_0 - \rho(\rho+2\mu)\frac{Z_1}{Z_0}\left[\frac{(\rho+\mu)x_0 - \mu x_1}{k_0}\right]^{1-b}}{(\rho+\mu)Z_0 + \mu Z_1 - \rho(\rho+2\mu)\left[\frac{(\rho+\mu)x_0 - \mu x_1}{k_0}\right]^{1-b}} \\
&< \frac{(\rho+\mu)Z_1 + \mu Z_0 - \rho(\rho+2\mu)\left[\frac{(\rho+\mu)x_0 - \mu x_1}{k_0}\right]^{1-b}}{(\rho+\mu)Z_0 + \mu Z_1 - \rho(\rho+2\mu)\left[\frac{(\rho+\mu)x_0 - \mu x_1}{k_0}\right]^{1-b}} \\
&< \frac{(\rho+\mu)Z_1 + \mu Z_0}{(\rho+\mu)Z_0 + \mu Z_1}
\end{aligned}$$

where the last inequality uses the fact that both numerator and denominator are positive since $\frac{dx_j}{ds} < 0$. Hence, any point $(x_0, x_1) > (0, 0)$ for which (4.2) holds with equality is consistent with optimal behavior if it satisfies

$$\frac{dx_1}{dx_0} < \frac{(\rho + \mu) Z_1 + \mu Z_0}{(\rho + \mu) Z_0 + \mu Z_1}$$

i.e. a path $(x_0 (s), x_1 (s))$ which satisfies the first order conditions and cuts through the line defined by (4.2) must come from a region where (4.2) is violated. But since

$$\lim_{s \to \overline{s}} \frac{x_1 (\overline{s})}{x_0 (\overline{s})} = \frac{(\rho + \mu) Z_1 + \mu Z_0}{(\rho + \mu) Z_0 + \mu Z_1}$$

this requires that there exists an $s' \leq \overline{s}$ such that

$$\frac{x_1 (s')}{x_0 (s')} = \frac{(\rho + \mu) Z_1 + \mu Z_0}{(\rho + \mu) Z_0 + \mu Z_1}$$

and

$$\frac{dx_1 (s')}{dx_0 (s')} > \frac{(\rho + \mu) Z_1 + \mu Z_0}{(\rho + \mu) Z_0 + \mu Z_1}$$

We have just established that this cannot be the case for any $s' < \overline{s}$, while continuity establishes this cannot be the case for $s' = \overline{s}$. This contradiction establishes $\left.\frac{ds_{it}}{dt}\right|_{z_t = Z_0} > \left.\frac{ds_{it}}{dt}\right|_{z_t = Z_1}$. ■

Proof of Proposition 4: To show that $s_{it} \geq s_{jt}$, recall from the proof of Proposition 3 that when aggregate productivity is equal to Z_k, the change in skills is given by

$$\left.\frac{ds_t}{dt}\right|_{z_t = Z_k} = a \left(\frac{b}{1 - b} \right)^b \left[\frac{\rho x_k (s) + \mu [x_k (s) - x_{-k} (s)]}{Z_k} \right]^b$$

Since $x_k (s)$ and $x_{-k} (s)$ represent the value of the option to learn, which is weakly increasing in a, it follows that $\left.\frac{ds_t}{dt}\right|_{z_t = Z_k}$ is strictly increasing in a for both values of k. Suppose there exists a date $t \geq 0$ and a state of productivity $k \in \{0, 1\}$ such that $s_{it} < s_{jt}$. Then since $s_{i0} = s_{j0}$, there must exist a date $t' \geq 0$ and a state of productivity k such that $\left.\frac{ds_{it'}}{dt}\right|_{z_t = Z_k} < \left.\frac{ds_{jt'}}{dt}\right|_{z_t = Z_k}$, which contradicts what we just established. Since $\left.\frac{ds_t}{dt}\right|_{z_t = Z_k}$ is strictly increasing in a whenever $s_t < \overline{s}$, it follows that $s_{it} = s_{jt}$ if $s_{it} = \overline{s}$.

Next, we establish that an individual still reaches $\overline{s}$ in finite time, i.e. there exists a time T such that for any realization z_t where $t \in [0, \infty)$, $s_t = \overline{s}$ for $t \geq T$. We do this by proving that if either $j = 0$ or $j = 1$, the individual would reach $\overline{s}$ in finite time. It follows from this that for any realization z_t, the individual would reach $\overline{s}$ in finite time.

Combining the evolution of $x_j(s)$ and s_t yields the following law of motion for x_j as a function of t conditional on $z_t = Z_j$:

$$\begin{aligned}
\left.\frac{dx_j}{dt}\right|_{z_t=Z_j} &= \left.\frac{dx_j}{ds}\cdot\frac{ds_t}{dt}\right|_{z_t=Z_j} \\
&= \left(\left[\frac{y_j}{k_j}\right]^{1-b} - \frac{(\rho+\mu)Z_j+\mu Z_{-j}}{\rho(\rho+2\mu)}\right) a\left(\frac{b}{1-b}\right)^b\left[\frac{y_j}{Z_j}\right]^b \\
&= \frac{a}{k_j}\left(\frac{b}{1-b}\frac{k_j}{Z_j}\right)^b y_j - \left(\frac{b}{1-b}\frac{1}{Z_j}\right)^b \frac{a(\rho+\mu)Z_j+a\mu Z_{-j}}{\rho(\rho+2\mu)} y_j^b \\
&\equiv A_j y_j - B_{j1} y_j^b
\end{aligned}$$

where $y_j \equiv (\rho+\mu)x_j(s) - \mu x_{-j}(s)$ and the coefficients A_j and B_j are positive. Differentiating y with respect to time yields

$$\begin{aligned}
\left.\frac{dy_j}{dt}\right|_{z_t=Z_j} &= \rho\left.\frac{dx_j}{dt}\right|_{z_t=Z_j} + \mu\left(\left.\frac{dx_j}{dt}\right|_{z_t=Z_j} - \left.\frac{dx_{-j}}{dt}\right|_{z_t=Z_j}\right) \\
&= \rho\left[A_j y_j - B_j y_j^b\right] + \mu\left(\left.\frac{dx_j}{dt}\right|_{z_t=Z_j} - \left.\frac{dx_{-j}}{dt}\right|_{z_t=Z_j}\right)
\end{aligned}$$

Note that except for the last term, the dynamics of y_j are similar to those as for x_j in the case with no aggregate fluctuations. For that dynamical system, we already established that y_j hits 0 in finite time. Turning to the additional component, we have

$$\left.\frac{dx_j}{dt}\right|_{z_t=Z_j} - \left.\frac{dx_{-j}}{dt}\right|_{z_t=Z_j} = \left(\frac{dx_j}{ds} - \frac{dx_{-j}}{ds}\right)\cdot\left.\frac{ds}{dt}\right|_{z_t=Z_j}$$

From Proposition 3, we know that as $s \to \overline{s}$,

$$\frac{dx_1/ds}{dx_0/ds} \to \frac{(\rho+\mu)Z_1+\mu Z_0}{(\rho+\mu)Z_0+\mu Z_1} > 1$$

Hence, there exists an $\varepsilon > 0$ such that if $s < \overline{s} - \varepsilon$, the laws of motion for y when $z = Z_1$ satisfy

$$\left.\frac{dy_1}{dt}\right|_{z_t=Z_1} < \rho\left[A_1 y_1 - B_1 y_1^b\right]$$

This insures y_1 will hit 0 in finite time, so that if $z_t = Z_1$ for all t, the individual would reach $\overline{s}$ in finite time.

Next, we know that

$$\begin{aligned}\frac{dx_j/dt|_{z_t=Z_j}}{dx_j/dt|_{z_t=Z_{-j}}} &= \frac{ds_j/dt|_{z_t=Z_j}}{ds_j/dt|_{z_t=Z_{-j}}} \\ &= \left[\frac{Z_{-j}}{Z_j}\frac{(\rho+\mu)\,x_j(s)-\mu x_{-j}(s)}{(\rho+\mu)\,x_j(s)-\mu x_{-j}(s)}\right]^b\end{aligned}$$

and so

$$\lim_{s\to\overline{s}}\frac{dx_j/dt|_{z_t=Z_j}}{dx_j/dt|_{z_t=Z_{-j}}} \to \left[\frac{Z_{-j}}{Z_j}\frac{Z_j}{Z_{-j}}\right]^b = 1$$

This insures that if $z_t = Z_0$ for all t, the individual would reach $\overline{s}$ in finite time. Otherwise, it must be the case that as $s \to \overline{s}$, $x_j \to 0$ more rapidly when $z_t = Z_1$ than when $z_t = Z_0$. But this would imply

$$\lim_{s\to\overline{s}}\frac{dx_j/dt|_{z_t=Z_1}}{dx_j/dt|_{z_t=Z_0}} > 1$$

which is a contradiction.

Next, consider two individuals i and j where $a_i > a_j$. Since $s_{it} > s_{jt}$ as long as $s_0 < s_{it} < \overline{s}$, it follows that i reaches $\overline{s}$ before j. At this point, wage inequality is clearly decreasing, since

$$\frac{d}{dt}\left(\frac{w_i}{w_j}\right)\bigg|_{z_t=Z_k} = \overline{s}\frac{d}{dt}\left(\frac{1}{s_j}\right)\bigg|_{z_t=Z_k} < 0$$

By continuity, there exists an $\varepsilon > 0$ such that both derivatives are negative whenever $s_{it} < \overline{s}-\varepsilon$, which establishes the proof. An analogous argument that relies on Proposition 3 establishes the final part of the Proposition.

References

[1] Barlevy, Gadi, 1999. "The Sullying Effect of Recessions" Mimeo, Northwestern University.

[2] Barlevy, Gadi, 2000. "Stabilization Policy in a Model of Endogenous Growth" Mimeo, Northwestern University.

[3] Beaudry, Paul and John DiNardo, 1991.

[4] Becker, Gary, 1975. *Human Capital; a Theoretical and Empirical Analysis, with Special Reference to Education*. Chicago: NBER.

[5] Ben-Porath, Yoram, 1967. "The Production of Human Capital and the Life Cycle of Earnings" *Journal of Political Economy*.

[6] Blank, Rebecca and Alan Blinder, 1986. "Macroeconomics, Income Distribution and Poverty" in *Fighting Poverty: What Works and What Doesn't*, edited by Sheldon Danziger and D. Weinberg, Cambridge, MA: Harvard University Press, p180-208.

[7] Blinder, Alan and Howard Esaki, 1978. "Macroeconomic Activity and Income Distribution in the Postwar United States" *Review of Economics and Statistics*, November, 60(4), p604-609.

[8] Burtless, Gary, 1990. "Earnings Inequality over the Business and Demographic Cycles" in *A Future of Lousy Jobs? The Changing Structure of U.S Wages*, edited by Gary Burtless, Washington, D.C.: Brookings Institution, p77-117.

[9] Caselli, Francesco, 1999. "Technological Revolutions" *American Economic Review*, March, 89(1), p78-102.

[10] Cole, Harold and Lee Ohanian, 1999. "The Great Depression in the United States from a Neoclassical Perspective" *Federal Reserve Bank of Minneapolis Quarterly Review*, 23(1), Winter, p2-24.

[11] Cooper, Russel and John Haltiwanger, 1993. "The Aggregate Implications of Machine Replacement: Theory and Evidence" *American Economic Review*, June, 83(3), p181-186.

[12] Cutler, David and Lawrence Katz, 1991. "Macroeconomic Performance and the Disadvantaged" *Brookings Papers on Economic Activity*, 2, p1-74.

[13] David, Paul, 1991. "Computer and Dynamo: the Modern Productivity Paradox in a Not-too-Distant Mirror" *Technology and Productivity: the Challenge for Economic Policy*. Paris: OECD.

[14] Dellas, Harris and Plutarchos Sakellaris, 1997. "On the Cyclicality of Schooling: Theory and Evidence" Mimeo, University of Maryland.

[15] Galor, Oded and Daniel Tsiddon, 1997. "Technological Progress, Mobility, and Economic Growth" *American Economic Review*, June, 87(3), p363-82.

[16] Goldin, Claudia and Robert Margo, 1992. "The Great Compression: The Wage Structure in the United States at Mid- Century" *Quarterly Journal of Economics*, February, 107(1), p1-34.

[17] Goldsmith, Selma. "Changes in the Size Distribution of Income" in *Inequality and Poverty*, edited by Edward Budd. New York: WW Norton and Company, 65-79..

[18] Greenwood, Jeremey and Mehmet Yorukoglu, 1997. "1974" *Carnegie Rochester Conference Series on Public Policy*, December, 46(2), p49-95.

[19] Kuznets, Simon, 1953. *Shares of Upper Income Groups in Income and Savings*. New York: NBER.

[20] Kuznets, Simon, 1955. "Economic Growth and Income Inequality" *American Economic Review*. May, 45(1), p1-28.

[21] Lucas, Robert, 1987. Models of Business Cycles. Oxford: Basil Blackwell.

[22] McLean, Ian, 1991. "The Distributional Impact of the Depression in the United States" in *Income Distribution in Historical Perspective*, edited by Y. Brenner, Hartmut Kaelble, and Mark Thomas, Cambridge, UK: Cambridge University Press, p189-211.

[23] Roy, 1951.

[24] Schultz, Theodore, 1975. "The Value of the Ability to Deal with Disequilibria" *Journal of Economic Literature*, September, 13(3), p827-46.

[25] Solon, Gary, Robert Barsky, and Jonathan Parker, 1994. "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?" *Quarterly Journal of Economics*, February 1994, 109(1), p1-26.

[26] Stevens, Rosemary, 1989. *In Sickness and in Wealth: American Hospitals in the Twentieth Century*, New York: Basic Books.

[27] Swanson, Eric. "Real Wage Cylicality in the PSID, 1967-1991" Mimeo, Stanford University, 1998.

[28] Yona Rubinstein and Tsiddon, Daniel 1998. "Coping with Technological Progress: the Role of Ability in Making Inequality so Persistent" Mimeo, Tel Aviv University.

[29] Yona Rubinstein and Tsiddon, Daniel 2000. "Born to be Unemployed: Unemployment and Wages over the Business Cycle" Mimeo, Tel Aviv University.

[30] Williamson, Jeffrey and Peter Lindert, 1980. *American Inequality: a Macroeconomic Perspective*, New York: Academic Press.