

NBER WORKING PAPER SERIES

INCOME MOBILITY, INCOME RISK AND WELFARE

Tom Krebs
Pravin Krishna
William F. Maloney

Working Paper 23578
<http://www.nber.org/papers/w23578>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2017

We are grateful to seminar participants at Columbia, Harvard and Pennsylvania State University for many useful comments. Particular thanks to Francois Bourguignon, Francisco Ferreira, Gary Fields, and the other participants at the World Bank/Universitat Autònoma de Barcelona Workshop on Mobility for very helpful discussions. Our appreciation to Edwin Goni and Mauricio Sarrias for inspired research assistance. This work was partially supported by the Regional Studies Program of the Office of the Chief Economist for Latin America at the World Bank, the World Bank Research Support Budget and the Strategic Research Program. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Tom Krebs, Pravin Krishna, and William F. Maloney. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Income Mobility, Income Risk and Welfare
Tom Krebs, Pravin Krishna, and William F. Maloney
NBER Working Paper No. 23578
July 2017
JEL No. F0,J0,J6

ABSTRACT

This paper presents a framework for the quantitative analysis of individual income dynamics, mobility and welfare, with ex-ante identical individuals facing a stochastic income process and market incompleteness implying that they are unable to insure against persistent shocks to income. We show how the parameters of the income process can be estimated using repeated cross-sectional data with a short panel dimension, and use a simple consumption-saving model for quantitative analysis of mobility and welfare. Our empirical application, using data on individual incomes from Mexico, provides striking results. Most of measured income mobility is driven by measurement error or transitory income shocks and therefore (almost) welfare-neutral. Only a small part of measured income mobility is due to either welfare-reducing income risk or welfare-enhancing catching-up of low-income individuals with high-income individuals, both of which, nevertheless, have economically significant effects on social welfare. Strikingly, roughly half of the mobility that cannot be attributed to measurement error or transitory income shocks is driven by welfare-reducing persistent income shocks. Decomposing mobility into its fundamental components is thus crucial from the standpoint of welfare evaluation.

Tom Krebs
Department of Economics
University of Mannheim
68131 Mannheim
Germany
tkrebs@econ.uni-mannheim.de

William F. Maloney
The World Bank
Wmaloney@worldbank.org

Pravin Krishna
Johns Hopkins University
1740 Massachusetts Avenue, NW
Washington, DC 20036
and NBER
Pravin_Krishna@jhu.edu

I. Introduction

Income mobility, that is, the extent to which individuals move across different sections of the income distribution, is a central issue in a variety of public policy discussions today.¹ Income mobility is important as it is seen as indicative of the opportunities afforded by society to escape one's origins, thus allowing low-income individuals to catch-up with those with higher incomes. While such mobility may generally be seen as enhancing social welfare, it must be recognized that individual income changes, and thus mobility, may also be driven by the risk to incomes to which individuals are exposed,² which may be welfare-reducing instead.

In this paper, we develop an analytical framework for the estimation and welfare-theoretic evaluation of individual (intra-generational) income dynamics that takes into account these different (and, competing, from a welfare-theoretic perspective) drivers of income mobility.³ We begin by noting that the literature on income mobility has often focused on two important questions: the quantitative/empirical measurement of the extent and nature of the change in individual incomes and, separately, the social-welfare-theoretic evaluations of such changes.⁴ Two methodological issues have arisen in this area. First, the parametric formulations used in

¹In developing country contexts, see, for example, Hnatkovska, Lahiri and Paul (2012), who examine income mobility across different castes in India and the recent World Bank flagship publication, *Economic Mobility and the Rise of the Middle Class* (2012), which focuses on Latin America. Clearly income mobility is a concern in developed countries as well – see, for instance, the recent analysis of mobility patterns in the United States by Chetty et. al. (2014).

²On the relevance of risk in income earnings, especially for the poor in developing countries, see the discussion in Banerjee and Duflo (2011, Chapter 6) and also Collins et. al. (2009).

³Our focus here is on intra-generational or within-lifetime mobility, as opposed to the related but distinct concept of inter-generational mobility which looks at relative income gains achieved by individuals given the position in the income distribution occupied by their parents. Our focus on intra-generational mobility is driven by the fact that most longitudinal datasets on individuals do not support a study of inter-generational mobility as they do not contain information on parental income or educational status. We should note that focusing on intra-generational mobility nevertheless suffices for us to make our central point regarding the role of risk in driving mobility.

⁴For the former, see Lillard and Willis (1978), Shorrocks (1978b), Geweke, Marshal and Zarkin (1986), Conlisk (1990) and Fields and Ok (1996). For the latter, see, Atkinson (1983), Markandya (1982, 1984), Atkinson, Bourguignon and Morrison (1992), Dardononi (1993), and Gottschalk and Spolaore (2002). Additionally, the discussion over suitable social (income) mobility measures (indices), which may be used to evaluate mobility given the pattern of individual income changes in society, constitutes a very well researched area that has generated a number of important contributions in recent years. See Fields and Ok (1999) for a survey discussion.

the measurement of income changes are not easily used as inputs to the quantitative welfare-theoretic analysis, thereby constituting a problematic gap between these two literatures. Furthermore, as the literature has often pointed out, the measurement of dynamic income changes is itself confronted by (at least) the following two problems. First, income data are subject to measurement error and, second, a significant proportion of the observed income changes may be simply temporary in nature - resulting, typically, in an overestimation of the relevant mobility in income.⁵ This is also important from the perspective of welfare analysis, as measurement error has no effect on workers' welfare and transitory shocks to income are perhaps easily smoothed out, resulting in small welfare effects. In addition, welfare analysis is confronted by an additional challenge. Since individual utility is postulated as taking consumption rather than income as its argument, its direct valuation requires reliable data on individual consumption levels, which are often unavailable for developing countries. To use the more easily available data on incomes, a theoretical framework is required that translates the estimated income dynamics into consumption changes taking into account the institutional constraints individual agents face.

To highlight and explore the competing effects of welfare-enhancing “catch-up” relative to welfare-reducing risky income changes on mobility, our conceptual focus in this paper is on the income dynamics of ex-ante identical individuals. As we will discuss in detail in subsequent sections of the paper, this corresponds to an examination of income changes and mobility in “residual income”, income after conditioning for the standard determinants of income such as education and experience. Our focus on “residual income” mobility allows us to make our central point concerning the potentially welfare-reducing impact of income mobility when it is driven by risky income changes, without having to engage a different and difficult question regarding the welfare-theoretic evaluation of mobility between individuals who are not ex-ante identical: For instance, how much convergence of incomes is optimal, from a welfare perspective, between individuals with different initial education levels.⁶

⁵See, for example, Lillard and Willis (1978) Solon (2001), Luttmer (2002), Fields et al (2003), Glewwe (2004), and Antman and McKenzie (2007).

⁶We note additionally, that residual income is a significant portion of income in our data (as indicated in Table 2) and that measured residual income mobility is, nevertheless, not much different from unconditional income mobility in our data.

The analytical framework we develop to study income mobility provides a close link between the welfare theory and the empirical methodology used in the measurement of the income dynamics, thereby helping to bridge the gap between these literatures. At the same time, this framework overcomes many of the methodological problems that we have just discussed. Our approach consists of two steps. First, we follow a large empirical literature on individual income dynamics and postulate a stochastic income process that is highly parameterized, but sufficiently elaborate to distinguish between changes in income resulting from trend growth and other predictable factors and changes in income that are unpredictable.⁷ The unpredictable part of income change, in turn, has two components, one first degree autoregressive (AR(1)) component reflecting persistent shocks to income and another component that is i.i.d and captures transitory shocks and measurement error in the income data.⁸ We show how income mobility, measured in relation to the correlation of incomes over time,⁹ relates to the various parameters of the underlying income process. Further, we show how the parameters of the income process can be estimated using an econometric approach that exploits both the longitudinal and repeated cross-sectional features of income data, and apply our estimation strategy to individual income data from Mexico (Sections III-V).¹⁰ Our econometric approach is particularly suited for the application to Mexico and other developing countries, where data sets with a long panel dimension do not exist, but repeated cross-sectional observations with a very short panel dimension (rotating panel) are available.

⁷See, for example, Baker and Solon (2003) and Meghir and Pistaferri (2011) for a detailed discussion of the literature.

⁸Examples of transitory shocks are fluctuations in temporary layoffs, overtime labor supply, bonuses, lottery prizes, and bequests. Permanent innovations are associated with, for instance, promotions, layoffs that lower the value of the worker's skill set, severe health shocks and accumulations of high quality experience.

⁹Specifically, we use a quite basic and familiar measure, the Hart Index, which is the complement of the correlation between the logarithm of incomes over times (see Hart (1981) and Shorrocks (1993)). As Fields and Ok (1996) discuss, however, the literature has recently made important advances in studying the "multi-faceted concept" of mobility and a number of different theoretical measures, each capturing a different aspect of mobility have been introduced. We have no contribution to make to this discussion and simply use the Hart Index as our basic measure of mobility.

¹⁰See Bourguignon, Goh and Kim (2006) for an interesting exercise that compares results on poverty vulnerability (the propensity to move into poverty) obtained using panel data on incomes with those obtained from repeated cross-sections instead and finds that model parameters recovered from pseudo-panels approximate reasonably well those estimated directly from a true panel.

In a second step, we follow the large literature on consumption-saving models to provide a link between income dynamics, consumption, and welfare.¹¹ To translate stochastic income into consumption and welfare, we develop a dynamic general equilibrium model with incomplete markets in which the consumption/saving choice of heterogeneous workers in the presence of income variability is explicitly modeled. As is well known, general versions of such models are difficult to solve, and most work in the literature has therefore been computational. In contrast, we rely upon an extended version of the incomplete-markets model developed and analyzed by Constantinides and Duffie (1996) and Krebs (2007) that is highly tractable, but still rich enough to allow for tight links between the econometric framework and the welfare-theoretic model. This approach allows us to provide a clear analytical and quantitative discussion of these interrelated concepts, and specifically the role of income variability. We discuss in detail how different determinants of measured income mobility may have quite different implications for welfare. Specifically, we show that the auto-correlation coefficient of the AR(1) process (the catching-up parameter) measures “good mobility” in the sense that a reduction in this parameter increases both mobility and welfare. In contrast, social welfare is (almost) unaffected by measurement error or transitory income shocks even though mobility increases with the variance of the i.i.d. component of labor income. Finally, the variance of persistent income shocks (income risk) increases mobility, but decreases social welfare. This implies that two societies with the same initial distribution of income and the same level of measured income mobility and aggregate growth may experience quite different social welfare changes depending upon the different combinations of the underlying income parameters. The welfare expressions that we derive are also functions of the income parameters. Estimates of these parameters, as discussed above, thus allow for a quantitative analysis of individual and social welfare.

Two aspects of our analysis make our framework particularly suitable for application in developing country contexts. First, our theoretical framework makes specific assumptions regarding the (non-) availability of institutions to insurance against variations in *permanent*

¹¹See Deaton (1991) for the basic partial equilibrium model and Huggett (1993) and Aiyagari (1994) for the seminal applications to general equilibrium analysis. See Heathcote, Storesletten, and Violante (2009) for a recent survey.

income. Agents are allowed to save, but not allowed to borrow. No other insurance schemes, such as those that may be organized by a government, exist. We believe this characterization to be closer to that of developing economies, but this analysis would be relevant in any contexts where such borrowing constraints bind and where the relevant social insurance schemes are absent. Second, our estimation framework for estimating transitory and persistent income risk, confronts, through the use of suitable moment conditions, a common feature of developing country data sets on individual incomes, i.e., the availability of repeated cross-sections of the data with a large number of observations, but with a short panel dimension.

We present a quantitative implementation of our framework that underscores the importance of decomposing income dynamics into its components. Specifically, an application using data on individual incomes from Mexico yields striking results. Most of measured income mobility is driven by measurement error or transitory income shocks and therefore (almost) welfare-neutral, and only a small part of measured income mobility is due to either welfare-reducing income risk or welfare-enhancing catching-up of low-income individuals with high-income individuals. However, despite the small mobility effects, (idiosyncratic) persistent income risk has significant negative effects on social welfare – eliminating or insuring it would generate welfare gains that are equivalent to an increase in lifetime consumption by about 10 percent even if workers are only moderately risk-averse (log-utility).¹² Eliminating the catch-up of low income individuals with high income individuals yields a loss in social welfare of similar magnitude. Decomposing mobility into its fundamental components is thus seen to be crucial from the standpoint of welfare evaluation.

In sum, in this paper we make two contributions. First, we present a tractable framework that provides a transparent link between income dynamics, mobility, and welfare. Second, we apply our approach to individual income data from Mexico and show the importance of decomposing income mobility into its fundamental components.

¹²In comparison, for the same preference parameters, Lucas (2003) computes welfare cost of aggregate consumption fluctuations in the US that are two orders of magnitude smaller. Thus, even though our estimates of persistent income risk seem small when measured mobility is the yardstick, their welfare effect is large indeed.

We conclude the introduction with a general remark. In this paper, we use a tractable model to provide a link between income dynamics and consumption and welfare. The simplicity of our approach has the advantage of clarifying the basic channels that we like to emphasize in this paper. However, it has the disadvantage of leaving out some additional channels that are potentially important. For example, we use an exchange economy so that any effect of income risk on physical capital accumulation (Aiyagari, 1994) or human capital accumulation (Krebs, 2003) are ruled out by assumption. Further, we focus on households with little financial wealth implying very strong effects of permanent labor income shocks on consumption, a result that finds support in the data for households with little financial wealth (Blundell, Pistaferri, and Peston 2008). In contrast, for wealthy households the evidence indicates that the effect of permanent income shocks on consumption is less than one-to-one (Blundell, Pistaferri, and Peston 2008). Moreover, endogenous labor supply and payments from family members can provide consumption insurance, in particular in the context of developing countries. Finally, we have adopted the time-honored assumptions of time-additive expected utility preferences.¹³ Gottschalk and Spolaore (2002) analyze mobility and welfare in a setting with more general preferences that do not obey the independence axiom. Extending our analysis to models that allow for non-expected utility preferences and additional insurance channels are important topics for future research.

II. Income and Mobility

II.1. Income Process

Consider a large number of workers indexed by i . For notational ease, we focus on one cohort of workers who enter the labor market for the first time in period $t = 0$ so that $t = 0, 1, \dots$ stands for both calendar time and age (experience) of the worker. Let y_{it} stand for the labor income of worker i in period t . Following a longstanding tradition in microeconomics, we postulate that the log of y_{it} is given by the sum of returns to observable

¹³We also assume that the one-period utility function is logarithmic, but our analysis can easily be extended to the general case of CRRA-utility functions.

worker characteristics, and "residual" component: ¹⁴

$$\ln y_{it} = \lambda_x \cdot x_{it} + v_{it} \quad (1)$$

where x_{it} denotes the vector of observable worker characteristics, λ_x denotes the corresponding vector of coefficients and v_{it} denotes residual income and is itself given by $v_{it} = \omega_{it} + \eta_{it} + \mu$ where ω_{it} is a persistent component, η_{it} is transitory and μ denotes the mean of income.

The persistent component, ω_{it} , follows an AR(1) process

$$\omega_{i,t+1} = \rho\omega_{it} + \epsilon_{i,t+1} , \quad (2)$$

where ρ is a parameter measuring the persistence of shocks. The term ϵ denotes a stochastic innovation to labor income, which we assume to be i.i.d. over time and across individuals. We further assume that the transitory component of labor income, η_{it} , is i.i.d. over time and across individuals. Moreover, η_{it} and $\epsilon_{i,t+n}$ are uncorrelated for all t and n . All random variables are normally distributed so that labor income is log-normally distributed. More specifically, we assume that $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$, $\eta_{it} \sim N(0, \sigma_\eta^2)$, and $\omega_{i0} \sim N(0, \sigma_{\omega_0}^2)$.

Importantly, note that for expositional simplicity, we do not allow for individual means in our specification (1) above, but simply impose that the mean of the income process μ is the same for all individuals. However, theoretical expressions that we derive below can also be derived by allowing for individual means, μ_i in (1). As we will discuss later (in Section V), we are able to easily accommodate the presence of such individual "fixed effects" in our empirical implementation.

Equations (1) and (2) together imply that:

$$v_{it} = \rho^t \omega_{i0} + \sum_{n=0}^{t-1} \rho^{t-n-1} \epsilon_{i,n+1} + \eta_{it} + \mu . \quad (3)$$

Thus, labor income in period t is determined by initial condition, ω_0 , and stochastic changes, the latter being represented by the transitory shocks, η , and permanent shocks, ϵ . From (3)

¹⁴See, for example, Baker and Solon (2003) and Meghir and Pistaferri (2011) for a detailed discussion of the literature. In contrast to some papers in the literature (Gottschalk and Moffitt, 1994, and Carroll and Samwick, 1997), we do not impose the random walk restriction on the persistent component of labor income.

and our assumptions about ϵ , η , and ω_0 it follows that expected labor income is $E[\ln y_{it}] = \mu$ and labor income uncertainty before ω_{i0} is known is given by

$$var[v_{it}] = \begin{cases} \rho^{2t} \sigma_{\omega_0}^2 + \sigma_{\eta}^2 + \frac{1-\rho^{2t}}{1-\rho^2} \sigma_{\epsilon}^2 & \text{if } \rho \neq 1 \\ \sigma_{\omega_0}^2 + \sigma_{\eta}^2 + t \sigma_{\epsilon}^2 & \text{if } \rho = 1 \end{cases} . \quad (4)$$

As we have mentioned earlier, our study examines income mobility of ex-ante identical individuals within their lifetimes (i.e., intra-generational income mobility).¹⁵ From (2), the parameter ρ measures persistency of income and thus $(1 - \rho)$ measures the extent to which individuals with low levels of income “initially” will catch up with individuals with high initial income. In our context, the “initial” period corresponds to the time of entry into the work force after the completion of formal education. Since labor income may vary initially for equivalent individuals, catching-up in this context measures the extent to which individuals with initially low incomes catch up to those with initially high incomes.¹⁶

II.2. Mobility

As noted in the introduction, our empirical measure of income mobility between 0 and t , which we denote by m_t , is the Hart index, defined as the complement of the correlation in residual incomes at 0 and t (see Shorrocks (1993)):

$$\begin{aligned} m_t &= 1 - corr(v_{i0}, v_{it}) \\ &= 1 - \frac{cov(v_{i0}, v_{it})}{\sigma_{v_{i0}} \cdot \sigma_{v_{it}}} , \end{aligned} \quad (5)$$

where we have used the notation $\sigma_{v_{i0}} = \sqrt{var(v_{i0})}$ and $\sigma_{v_{it}} = \sqrt{var(v_{it})}$. Using our income specification from the previous section, we find the following expression for the covariance:

¹⁵For recent work on intra-generational mobility, see Antman and McKenzie (2007), Cuesta and Pizzolitto (2010), Dang et. al. (2011), and Cruces et. al (2011).

¹⁶In the terminology of the growth literature, it measures convergence. To see this, suppose $\rho < 1$. In this case, we have convergence towards the “steady state”: $E[\ln y_{it} | \omega_{i0}] \rightarrow \mu$. Let $\Delta_0 = \ln y_{i0} - \bar{d}$ be the initial distance from the steady state and $\Delta_t = \ln y_{it} - \bar{d}$ be the distance in period t . We can then define the time, T , it takes to get halfway towards the steady state, which is simply the solution to $\Delta_T / \Delta_0 = 1/2$. Using the expression for Δ_T and Δ_0 , it is straightforward to see that T is increasing in ρ for $\rho < 1$, that is, an increase in ρ reduces the speed of convergence.

$$\begin{aligned}
cov(v_{i0}, v_{it}) &= cov(\omega_{i0} + \eta_{i0}, \rho^t \omega_{i0} + \sum_{n=0}^{t-1} \rho^{t-n-1} \epsilon_{i,n+1} + \eta_{it} + \mu) \\
&= \rho^t \sigma_{\omega_0}^2
\end{aligned} \tag{6}$$

Using (3) and (6), we find the following expression for income mobility:¹⁷

$$m_t = \begin{cases} 1 - \frac{\rho^t \sigma_{\omega_0}^2}{\sqrt{\sigma_{\omega_0}^2 + \sigma_{\eta}^2} \sqrt{\rho^{2t} \sigma_{\omega_0}^2 + \sigma_{\eta}^2 + \frac{1-\rho^{2t}}{1-\rho^2} \sigma_{\epsilon}^2}} & \text{if } \rho \neq 1 \\ 1 - \frac{\sigma_{\omega_0}^2}{\sqrt{\sigma_{\omega_0}^2 + \sigma_{\eta}^2} \sqrt{\sigma_{\omega_0}^2 + \sigma_{\eta}^2 + t \sigma_{\epsilon}^2}} & \text{if } \rho = 1 \end{cases} . \tag{7}$$

Equation (7) defines residual income mobility as a function of the parameters of interest, σ_{ϵ}^2 , σ_{η}^2 , and ρ . It is straightforward to show that mobility is increasing in the volatility parameters σ_{ϵ}^2 and σ_{η}^2 . This is intuitive as an increase in the variance of income shocks increases the variability of individual incomes, lowering the correlation between incomes across time, thus increasing mobility.

Importantly, income mobility is decreasing in ρ :

$$\frac{\partial m_t}{\partial \sigma_{\epsilon}^2} > 0 \quad , \quad \frac{\partial m_t}{\partial \sigma_{\eta}^2} > 0 \quad , \quad \frac{\partial m_t}{\partial \rho} < 0 . \tag{8}$$

Intuitively, any increase in ρ increases income persistence, reducing the catching-up effect and therefore reducing mobility.

III. Econometric Implementation

The discussion in the preceding sections has described how the different parameters of the income process ($\sigma_{\omega_0}^2$, σ_{ϵ}^2 , σ_{η}^2 and ρ) affect mobility. To get to a quantitative assessment of these linkages, we turn next to the methodology and data used to estimate these parameters.

¹⁷For $\rho < 1$, the ω -process has a stationary distribution. If we choose as initial distribution this stationary distribution, the ω -process becomes stationary with $\sigma_{\omega_t}^2 = \sigma_{\omega_0}^2 = \sigma_{\epsilon}^2 / (1 - \rho^2)$. In this case the mobility expression (7) reduces to $m_t = 1 - \rho^t / (1 + \sigma_{\eta}^2 / \sigma_{\omega_0}^2)$.

III.1. Estimation

We continue to assume that log labor income, $\ln y_{it}$, is specified as in (1) but with some elaboration to account for age cohorts as follows:

$$\begin{aligned}\ln y_{it} &= \lambda_t + \lambda_x \cdot x_{it} + \sum_z \lambda_z \delta(z_{it}) + v_{it} \\ v_{it} &= \mu + \omega_{it} + \eta_{it}\end{aligned}\tag{1'}$$

where, again, x'_{it} is vector of observable individual characteristics beyond age (education, education², gender), λ_t is a constant that varies by calendar time period (thus absorbing the effects of macroeconomic factors such as aggregate productivity growth and aggregate economic fluctuations on income), λ_x is a vector of coefficients for the vector of worker characteristics x' , and $\delta(z_{it})$ are age-dummies, with λ_z being the corresponding coefficients.

Equation (1') resembles a typical Mincer specification for labor income for which the residual, v_{it} , is the sum of two unobserved stochastic components, ω_{it} and η_{it} . As in Carroll and Samwick (1997), we first use equation (1') to estimate the residuals v_{it} and then use these estimated residuals to estimate, in a second step, the parameters of interest. As noted above, this implies, importantly, that our mobility measure relates to residual income v rather than unconditional income y .

For notational simplicity, assume that all individuals i “are born” in period $t = 0$, so that t and z simultaneously stand for age of the individual and calendar time. Equations (1) and (2) which describe our labor income process imply that the change in residual income variance with age is given by:

$$\begin{aligned}\text{Var}[v_{iz}] &= \text{var}[(\omega_{iz} + \eta_{iz})] \\ &= \sigma_\eta^2 + \rho^{2z} \sigma_{\omega_0}^2 + \frac{1 - \rho^{2z}}{1 - \rho^2} \sigma_\epsilon^2\end{aligned}\tag{4'}$$

(4') links the changes in cross sectional residual income variances over for any age cohort z with our parameters of interest. Unfortunately, however, (4') is not sufficient to separately

identify $\sigma_{\omega_0}^2$ and σ_{ϵ}^2 since, as can be seen from the expression on the right hand side, both evolve at the same rate with z . We therefore also use the covariance restriction,

$$\begin{aligned} cov(v_{iz}, v_{i,z+1}) &= cov((\omega_{iz} + \eta_{iz}), (\omega_{i,z+1} + \eta_{i,z+1})) \\ &= \rho^{2z+1} \sigma_{\omega_0}^2 + \frac{1 - \rho^{2z}}{1 - \rho^2} \rho \sigma_{\epsilon}^2 \end{aligned} \quad (6')$$

to achieve identification of all four parameters. Notice that (4') requires, on the left hand side, estimates of the cross-sectional variance of residual income for each age group z , while (6') requires that we use the panel dimension of our data set to estimate the covariances in individuals' residual incomes v_{iz} over time. Thus, our estimation strategy exploits both the panel dimension and the repeated cross sections available in the data set. As in Carroll and Samwick (1997), we use residual income data at the individual level to obtain unbiased estimators of the terms on the left hand side of (4') and (6'). Specifically, v_{iz}^2 and $v_{iz}v_{i,z+1}$ serve as individual level "observations" of the variance and covariance terms on the left hand sides of (4') and (6') respectively.

We emphasize here that equations (4') and (6') taken together enable separate identification of the variance of initial incomes, $\sigma_{\omega_0}^2$, and the variance of persistent income shocks, σ_{ϵ}^2 . Specifically, the system allows for initial income differences (determined possibly by heterogeneous but fixed and unobserved individual characteristics) and for persistent shocks to income whose evolution over time (as contrasted with fixed individual characteristics which do not change, by definition) enables the separate identification of these parameters. Initial draws of persistent characteristics ω_{i0} will be important in determining income inequality, but income mobility will nevertheless be determined by the magnitude of shocks to income (as indicated in (7)).

We estimate our system of two equations ((4') and (6')) using a simultaneous, non-linear, seemingly unrelated regressions model (NLSUR) (as described in Gallant, 1975 and Amemiya, 1983). This permits the estimation of the two non-linear equations, with the cross-equation restrictions implied by the common parameters, simultaneously and achieves additional estimation efficiency by combining information from both equations (Davidson &

MacKinnon, 2004).¹⁸

IV. Data

Using the estimation methodology described in the preceding section, we estimate income mobility parameters using individual income data from Mexico. Specifically, the individual income data are taken from the Encuesta Nacional de Empleo Urbano (ENEU, Mexican National Urban Employment Survey) which was conducted by the Instituto Nacional de Estadística, Geografía e Informática (INEGI, National Institute of Statistics, Geography and Information), the primary statistical agency in Mexico, and the Secretaría del Trabajo y Previsión Social (STPS, Secretariat of Labor and Social Security), Mexico's Labor Ministry.

Until recently, the ENEU was the primary survey instrument for collecting earnings and employment data in Mexico. The survey is sampled to be representative geographically and by social strata (see INEGI 2000). The basic sampling unit is the dwelling. Demographic information is collected on the household (households) occupying each dwelling. Subsequently, an employment questionnaire is administered for each individual aged 12 and above in the household on position in the household, level of education (years of schooling), age and sex as well as standard measures related to participation in the labor market: occupation, hours worked, employment conditions, search and earnings.

The ENEU is constructed as a rotating panel, where households are surveyed every quarter for a total of five quarters.¹⁹ ENEU, in its modern form, has employed a consistent survey instrument from 1987 to 2004; it is thus one of very few long-running surveys with a

¹⁸See also Davidson and MacKinnon (2004) for a thorough discussion of the asymptotic equivalence between estimates obtained using a non-linear-least-squares methodology and the generalized method of moments.

¹⁹In each round of the rotating panel, the questionnaire records absent members, adds any new members who have joined the household, and records any changes in schooling that have taken place. If none of the original group of household members is found to be living in the dwelling unit in the follow-up survey, the household is recorded as a new household. The interviewers do not track households that move, so they leave the panel. Rates of attrition are comparable to other developing countries (See Antman and McKenzie, 2007).

panel dimension in the developing world. In our study, we are able to use this 18 year span comprising a total of 72 quarters of data, with, as we have discussed, households appearing in the survey for five quarters before they are dropped.²⁰ Worker earnings include overall earnings in the individual’s principal occupation from fixed salary payments, hourly or daily wages, piece-meal work, commissions, tips and self employment earnings.

We note that while the ENEU survey records employment information on all members of the household above 12 years old, for younger workers employment is generally transient and time is often divided among schooling, unpaid support to the household and paid work. Similarly, much later in life, work again becomes more transient. In our analysis, we focus on individuals between the ages 20 and 65.

V. Results

As discussed in the previous section, our estimation methodology proceeds in two steps. As in Carroll and Samwick (1997), we first use individual data to estimate a Mincer earnings regression. In a second step, the residuals from the Mincer regression are used to estimate income mobility parameters using (4’) and (6’). Table 2 reports the estimates from the first stage earnings regression using the ENEU data described in the preceding section. Our estimates are consistent with earlier findings in the literature. Specifically, earnings increase, but at a decreasing rate, with education. Further, earnings increase with potential experience (age) up until the age of 44 after which they decrease again. Males appear to earn 31 percent more than women, conditional on the other covariates.²¹

We use next the residuals from the earnings regression, v_{it} , to construct individual level

²⁰Since 2004, the ENEU has been replaced by the Encuesta Nacional de Ocupacion y Empleo (ENOE, Survey of Occupation and Employment) in 2005. Unfortunately, however, the ENOE instrument differs from ENEU in important ways that make it impossible to match the surveys with confidence.

²¹For robustness we have also run alternate earnings specifications, allowing for both more and less temporal variation, by allowing all parameters to vary in each time period, and separately by constraining even the constant to be invariant across periods (unlike in the specification reported on in Table 2, which includes year fixed effects). The results do not change appreciably.

“observations” of income variances v_{it}^2 and covariances $v_{it}v_{i,t+1}$,²² that are to be used on the left hand side of equations (4') and (6') to estimate the income mobility parameters. The age profile of the constructed variance and covariance measures are indicated in Figures 1 and 2, which are generated by regressing the two variables respectively on a complete set of age and time dummies and then plotting the former against age (see Deaton and Paxson, 1994, for a similar exercise). Consistent with equations (4') and (6'), the accumulation of persistent shocks, σ_ϵ^2 , as age increases, gives both relationships an upward slope, albeit at rates differing by a factor of ρ .

Estimation results from the joint estimation of (4') and (6'), as described in the previous section, yield the parameter estimates listed in Table 3. The first column presents the results using the full sample, while the second column provides results obtained using data from just those households that enter the sample in the first quarter of each year. Our estimates of the income mobility parameters are also in line with those obtained previously in the literature. The autoregressive component, ρ , is estimated to be 0.977, which suggests that persistent shocks to income experienced by any individual i will indeed last a long time.

We note here that the estimate of ρ that we have obtained as being close to one is *not* driven by the exclusion of individuals mean effects in the income specification (1). If we allow individual fixed effects, μ_i , in income as follows: $v_{it} = \omega_{it} + \eta_{it} + \mu_i$, the expressions (4') and (6') are modified in the following fashion:

$$Var[v_{iz}] = \sigma_\mu^2 + \sigma_\eta^2 + \rho^{2z}\sigma_{\omega_0}^2 + \frac{1 - \rho^{2z}}{1 - \rho^2}\sigma_\epsilon^2 \quad (4'')$$

$$cov(v_{iz}, v_{i,z+1}) = \sigma_\mu^2 + \rho^{2z+1}\sigma_{\omega_0}^2 + \frac{1 - \rho^{2z}}{1 - \rho^2}\rho\sigma_\epsilon^2 \quad (6'')$$

where σ_μ^2 denotes the variance of individual fixed effects, μ_i . In estimating the framework above, however, we were unable to obtain separate estimates for σ_μ^2 and $\sigma_{\omega_0}^2$ due to the apparent collinearity between these two – as would be implied by a value of ρ close to 1. Thus, allowing the initial shocks to dissipate at a different rate than the fixed effects does not

²²Note that v_{t+1} denotes individual i 's residual one year (four quarters) after t

yield a different estimate of ρ . Separately, the system above may be estimated by allowing for different autoregressive coefficients ρ_ω and ρ_ϵ to be attached to $\sigma_{\omega_0}^2$ and σ_ϵ^2 respectively – thereby allowing the initial draws of persistent income to dissipate at a different rate than the later set of persistent shocks, ϵ . Specifically, this allows for a different estimate of ρ to be obtained from the rate of dissipation of later shocks (while also allowing for individual fixed effects). Estimating this system yields estimates of $\rho_\omega \approx 1$ and $\rho_\epsilon \approx 1$, affirming once again a value of $\rho \approx 1$.

A separate issue concerns our estimation assumption of a common ρ over different workers, especially those of different ages. To assess the restrictiveness of this assumption, we allow ρ to vary by age groups - for instance, by dividing workers into three groups - those under 30, those between 30 and 50 and, finally, those over 50 years of age. However, we are unable to reject that the estimates of ρ for the different age groups are different from each other – thus mitigating this concern.

The estimated variance of transitory shocks to income, $\sigma_\eta^2 = 0.202$, is significantly larger than the variance of persistent shocks to income, $\sigma_\epsilon^2 = 0.015$. This is not surprising given that the η -term in our specification also captures measurement error in income, which we expect to be quite large in our data set.²³ Finally, the estimated variance in initial incomes $\sigma_{\omega_0}^2 = 0.104$. The R^2 goodness-of-fit statistic is estimated to be 0.13.²⁴ As the results in the second column indicate, the estimates are not appreciably different with the restricted sample of households who enter the survey in just the first quarter of each year.

Given our estimates of the income parameters (i.e., $\rho = 0.977$, $\sigma_{\omega_0}^2 = 0.104$, $\sigma_\eta^2 = 0.202$ and $\sigma_\epsilon^2 = 0.015$), we can use expressions (7) to analyze mobility patterns. In particular, we can compute how much the individual parameters contribute to overall mobility. Plugging in our estimates of the parameters characterizing the income process into (7), we obtain estimates of mobility in residual income which we report in Table 4. Specifically, mobility in residual income over a 1 year period is calculated to be 0.67. For 10 years, calculated mobility

²³See Antman and McKenzie (2007) for a discussion of measurement error and mobility using this data.

²⁴ R^2 for equations (4') and (6') is estimated to be 0.15 and 0.16 respectively.

increases to 0.76 and for 25 years, it increases to 0.84. The reasons behind the surprisingly high one-year mobility level, and relatively modest increases thereafter, become clearer in the next rows which set to zero each of the key parameters and calculate the resulting change in mobility. Notice, first, that 1-year mobility falls by a full 90 percent if we set $\sigma_\eta^2 = 0$ – no transitory shocks or measurement error. As we have noted earlier, measurement error should not enter welfare calculations and individuals can often smooth transitory shocks through own savings so that their welfare impact is limited. By contrast, “bad mobility” σ_ϵ^2 due to risk and “good” mobility due to convergence, ρ , account for roughly 1 percent each across one year.²⁵

The relative impact of these parameters clearly changes as we increase the span over which we are measuring mobility. At 25 years, setting transitory shocks to zero reduces mobility by a still large, but much reduced 23 percent (as transitory shocks are, by definition, transitory and mobility over this duration is driven to a greater extent by the cumulative effect of persistent shocks experienced by individuals over this period). By contrast, mobility due to persistent risk accounts for 7.4 percent and mobility due to convergence accounts for 8.6 percent. Having identified which parameters have the largest influence on measured mobility, we now turn to their relative contribution to welfare.

VI. Welfare Analysis

The voluminous literature on consumption and saving with individual income risk and incomplete insurance markets has generated a number of insights.²⁶ One important insight is that workers can effectively self-insure against transitory income shocks through borrowing or own saving, and that the effect of these shocks on equilibrium prices and quantities are relatively small.²⁷ A second important insight of this literature is that very persistent or fully permanent income shocks have substantial effects on consumption and welfare even if

²⁵Note that since mobility is highly non-linear in its underlying parameters, measured mobility does not decompose additively into its component parts.

²⁶See, for example, Heathcote, Storesletten, and Violante (2009) for a recent survey.

²⁷See, for example, Aiyagari (1994) for quantitative work and Levine and Zame (2002) for a theoretical argument.

individual households have own savings, but no or only limited access to insurance markets. Indeed, when labor income is the main source of income and labor income shocks are highly persistent, we would expect that consumption responds (almost) one-for-one to labor income shocks. This point has been made more formally Constantinides and Duffie (1996) using dynamic general equilibrium exchange models with incomplete markets. Constantinides and Duffie (1996) only consider the case in which income follows a random walk ($\rho = 1$), but Krebs (2007) also analyzes an extension with $\rho < 1$ and costs of financial intermediation that introduce a spread between the borrowing rate and the lending rate. In this section, we discuss the main ideas and results of the model analyzed in Krebs (2007).

VI.1. Consumption

The model features long-lived, risk-averse ex-ante identical workers with homothetic preferences who make consumption/saving choices in the face of uninsurable income shocks. Workers' preferences over consumption plans, $\{c_{it}\}$, allow for a time-additive expected utility representation with one-period utility function of the CRRA-type, where in this paper we confine attention to the log-utility case (degree of relative risk aversion of 1):

$$U(\{c_{it}\}|\omega_{i0}) = E \left[\sum_{t=0}^{\infty} \beta^t \ln c_{it} | \omega_{i0} \right] . \quad (9)$$

Workers maximize expected lifetime utility subject to a sequential budget constraint that allows them to transfer wealth across periods through saving (or borrowing). The model is an exchange economy with endogenous interest rate (general equilibrium). Since workers are ex-ante identical, we do not distinguish between income y_{it} and residual income v_{it} in this discussion.

In order to apply the equilibrium characterization result of Krebs (2007), we need to introduce three modification of the labor income process (1). First, we abstract from ex-ante heterogeneity and time-effects. For simplicity, we set $\mu = 0$ so that the mean of labor income (aggregate labor income) is normalized to one (see below). Second, measurement error should not enter into the worker's budget constraint, and the part of η that represents measurement error should therefore be omitted. Further, as we have argued before, the part

of η that is due to true income shocks is expected to have only small effects on equilibrium consumption and welfare. To simplify the analysis, we neglect these small effects of transitory income shocks and set $\ln y_{it} = \omega_{it}$, where $\{\omega_{it}\}$ is an AR(1) process as in specification (2). Third, the distribution of the innovation term, ϵ , and the distribution of initial income, ω_0 , include a mean-adjustment: $\epsilon \sim N(-\sigma_\epsilon^2/2, \sigma_\epsilon^2)$ and $\omega_0 \sim N(-\sigma_{\omega_0}^2/2, \sigma_{\omega_0}^2)$. This adjustment is necessary to ensure that σ_ϵ^2 and $\sigma_{\omega_0}^2$ can be interpreted as uncertainty parameters (see below).²⁸

Our specification of the labor income process implies that

$$\begin{aligned} E[y_{i,t+1}|I_t] &= y_{it}^\rho & (10) \\ \text{var}[y_{i,t+1}|I_t] &= e^{\sigma_\epsilon^2} - 1 \\ E[y_{i0}] &= 1 \\ \text{var}[y_0] &= e^{\sigma_{\omega_0}^2} \end{aligned}$$

where I_t denote the information available at time t . Thus, increases in either σ_ϵ or σ_{ω_0} increase the variance of labor income without any change in the (conditional) mean – they lead to a mean-preserving spread. In other words, the two parameters measure risk/uncertainty.²⁹

If $\rho = 1$ and labor income follows a random walk, then the equilibrium interest rate will adjust so that individual workers will optimally decide to set consumption equal to labor income (see Constantinides and Duffie (1996) and Krebs (2007) for details). The argument is outlined in the Appendix. If ρ is not equal to one, but not too far away from one, then a sufficiently large difference in the borrowing and lending rate (cost of financial

²⁸The main part of the analysis in Krebs (2007) deals with the random walk case, but the Appendix discusses the extension to labor income shocks that are not fully permanent. The labor income process specified in the Appendix of Krebs (2007) is equivalent to an AR(1) process with an innovation term that has finite support, which rules out the case of a normal distribution. One way to apply the results of Krebs (2007) to the present analysis is to truncate all normal distributions at an arbitrarily large point, and to think of all equilibrium results as approximate results for which the approximation error can be made arbitrarily small.

²⁹The n -period ahead variances, $\text{var}[y_{i,t+n}|I_t]$, in general depend on σ_ϵ^2 for $n \geq 2$ if $\rho < 1$. We can correct for these “higher-order” effects without essentially changing the main results of the paper. More precisely, a modified version of the welfare formula (11), which adjusts for the change in mean income, yields quantitative results that are very close to the results reported here. Details are available on request.

intermediation) will ensure that in equilibrium households still choose to set consumption equals labor income (see the Appendix of Krebs (2007) for details). In short, in equilibrium we have $c_{it} = y_{it}$, that is, consumption and labor income move one-for-one.

VI.2. Mobility and Welfare

Using $c_{it} = y_{it} = \omega_{it}$ and the income specification discussed above, we can evaluate the expected lifetime utility (9) of an individual with initial income ω_{i0} . Taking the expectation over ω_{i0} yields social welfare, W , where we assume that each individual household is assigned equal weight in the social welfare function. In other words, social welfare is the expected lifetime utility from an *ex ante* point of view when the initial condition, ω_0 , is not yet known (veil of ignorance). More formally, we have

$$\begin{aligned}
 W &= E \left[\sum_{t=0}^{\infty} \beta^t \ln c_{it} \right] \\
 &= E \left[E \left[\sum_{t=0}^{\infty} \beta^t \ln c_{it} \mid \omega_{i0} \right] \right] \\
 &= E \left[-\frac{\beta}{(1-\beta)(1-\beta\rho)} \frac{\sigma_{\epsilon}^2}{2} + \frac{1}{1-\beta\rho} \omega_{i0} \right] \\
 &= -\frac{\beta}{(1-\beta)(1-\beta\rho)} \frac{\sigma_{\epsilon}^2}{2} - \frac{1}{1-\beta\rho} \frac{\sigma_{\omega_0}^2}{2}
 \end{aligned} \tag{11}$$

The formula (11) shows how social welfare depends on the various income parameters and the preference parameter β . In particular, (11) shows that an increase in uncertainty, either about initial conditions or about future labor market conditions, will reduce social welfare. Further, an increase in ρ increases uncertainty about lifetime income, and therefore reduces welfare:

$$\frac{\partial W}{\partial \sigma_{\omega_0}^2} < 0 \quad , \quad \frac{\partial W}{\partial \sigma_{\epsilon}^2} < 0 \quad , \quad \frac{\partial W}{\partial \rho} < 0 \tag{12}$$

In order to express welfare changes in economically meaningful units, we calculate the corresponding change in consumption in each period and possible future state that is necessary to compensate the worker for the change in uncertainty. For example, suppose we compare

two economies, one with income parameters $(\sigma_{\omega_0}^2, \sigma_\epsilon^2, \rho)$ and one with income parameters $(\hat{\sigma}_{\omega_0}^2, \hat{\sigma}_\epsilon^2, \hat{\rho})$. We then define the consumption-equivalent welfare change, Δ , of moving from $(\sigma_{\omega_0}^2, \sigma_\epsilon^2, \rho)$ to $(\hat{\sigma}_{\omega_0}^2, \hat{\sigma}_\epsilon^2, \hat{\rho})$ as

$$E \left[\sum_{t=0}^{\infty} \beta^t \ln (c_{it}(1 + \Delta)) \right] = E \left[\sum_{t=0}^{\infty} \beta^t \ln \hat{c}_{it} \right], \quad (13)$$

where c is consumption in the first economy and \hat{c} is consumption in the second economy. Using the definition (13) and the welfare formula (11), we find:

$$\begin{aligned} \ln(1 + \Delta) = & \frac{\beta}{(1 - \beta\hat{\rho})} \frac{\hat{\sigma}_\epsilon^2}{2} + \frac{(1 - \beta)}{(1 - \beta\hat{\rho})} \frac{\hat{\sigma}_{\omega_0}^2}{2} \\ & - \frac{\beta}{(1 - \beta\rho)} \frac{\sigma_\epsilon^2}{2} - \frac{(1 - \beta)}{1 - \beta\rho} \frac{\sigma_{\omega_0}^2}{2} \end{aligned} \quad (14)$$

As mentioned before, measurement error and transitory shocks have (almost) no effect on welfare. In contrast, the effect of the other two mobility parameters, σ_ϵ and ρ , turn out to be quite substantial. For example, based on the welfare formula (14) and an annual discount factor of $\beta = 0.96$, a value that is standard in the macro-economic literature (for example, Cooley and Prescott, 1995), we find that removing all “bad mobility”, setting $\sigma_\epsilon^2 = 0$, leads to a welfare gain of about 12 percent of lifetime consumption. Using the same discount factor, the welfare cost of removing all “good mobility”, setting $\rho = 1$, is equal to 8 percent of lifetime consumption, again a significant welfare effect. Finally, removing both “good” and “bad” mobility at the same time, setting $\sigma_\epsilon^2 = 0$ and $\rho = 1$, leads to a net welfare gain of about 10 percent of lifetime consumption. The last result shows that the welfare formula (14) is highly non-linear and that the positive welfare effect of catching-up, $\rho < 1$, is closely linked to the presence or absence of persistent income shocks, ϵ . Calculations with other values of β yield similar results as indicated in Table 5.

In sum, the application of our general framework to Mexico provides striking results. The parameter that accounts for the largest part of measured mobility, σ_η , has (almost) no effect on welfare, and the two parameters that have large effects on welfare, σ_ϵ and ρ , have only a modest contribution to measured mobility, and least over small time durations. Clearly,

our welfare results depend on the choice of preference parameters, namely the degree of risk aversion and the degree of impatience (discounting). However, by using a logarithmic utility function we have already chosen a relatively low degree of (relative) risk aversion, namely one, and any increase in the degree of risk aversion would only increase the welfare effects. Further, lowering the discount factor β will lower the welfare effects, but for a wide range of values of β the welfare effects remain substantial and the ranking of the different parameters remains the same (see Table 5). In short, our welfare results are valid for a wide range of preference parameters.

VII. Conclusions

This paper develops an analytically tractable framework linking individual income dynamics, social mobility and welfare. This analytical framework that we develop has the merit that the links between different determinants of income mobility and social welfare are drawn out in a simple and transparent manner – allowing for a clearer analytical and quantitative discussion of these interrelated concepts than has generally been possible in the past. In particular, we discuss in detail how different determinants of measured income mobility (shocks to income, and convergence forces, for instance) may have quite different implications for welfare. This implies that two societies with the same initial distribution of income and the same level of measured income mobility may be characterized by quite different levels of social welfare. Decomposing the determinants of mobility is thus shown to be crucial from the standpoint of welfare evaluation.

An important strength of the proposed framework is its empirical implementability. The quantitative evaluation of mobility and welfare in our context entails the estimation of income process parameters may be achieved using combined cross sectional and longitudinal data on individual incomes and relatively straightforward econometric techniques. The results from Mexico are striking. Most of measured mobility is estimated to be driven by transitory shocks to income and is therefore (almost) welfare neutral. Only a small part of mobility (i.e., mobility in permanent income) is driven by either social-welfare-reducing

persistent income shocks or welfare-enhancing catching-up of low-income individuals with high-income individuals. Importantly, roughly half of the of the mobility that is not driven by measurement error or transitory income shocks, can be attributed to welfare-reducing shocks to permanent income. Decomposing mobility into its fundamental components is thus crucial from the standpoint of welfare evaluation.

Appendix

Here we outline the proof that $c_{it} = y_{it}$ is an equilibrium choice. Details can be found in Krebs (2007) and Krebs, Krishna, and Maloney (2010).

Consider the sequential budget constraint

$$a_{i,t+1} = (1 + r)a_{it} + y_{it} \quad \forall t \quad (A1)$$

where a_{it} is asset holding (financial wealth) of individual i in period t and r is the risk-free interest rate. The first-order conditions (consumption Euler equation) associated with the worker problem of maximizing expected lifetime utility (11) subject to the sequential budget constraint read:

$$1 = \beta(1 + r) E \left[\left(\frac{c_{it}}{c_{i,t+1}} \right) \mid I_{it} \right] , \quad (A2)$$

where I_t is the information available to individual i in period t . A straightforward but lengthy argument shows that the first-order conditions (A2) in conjunction with the transversality condition are sufficient conditions for the solution of the worker's utility maximization problem, and that the transversality condition is equal to the no-Ponzi-scheme condition.

We assume that assets are in zero aggregate net supply: $E[a_{it}] = 0$. Suppose all workers begin life with no financial wealth, $a_{i0} = 0$. Clearly, $c_{it} = y_{it}$ and $a_{i,t+1} = 0$ solve the sequential budget constraint. Suppose further that income follows a logarithmic random walk, $\rho = 1$. Substituting $c_{it} = y_{it}$ into the first-order conditions (A2) and using $y_{i,t+1} = e^{\epsilon_{i,t+1}} y_{it}$ shows that (A2) holds if the interest rate is given by:

$$r = \beta^{-1} e^{-\sigma^2} - 1 , \quad (A3)$$

Further, a straightforward argument shows that the transversality condition is satisfied. Thus, the choice $c_{it} = y_{it}$ is individually optimal, that is, it solves the workers' utility maximization problem. It also satisfies the asset market clearing condition. Hence, $c_{it} = y_{it}$ and $a_{i,t+1} = 0$ conjunction with the interest rate (A3) constitute an equilibrium.

If $\rho < 1$, an argument similar to the above argument shows that $c_{it} = y_{it}$ and $a_{i,t+1} = 0$ solve the corresponding first-order conditions and transversality condition for an arbitrary but bounded set of income realizations if i) there is a spread between the borrowing rate and the lending rate (cost of financial intermediation) and ii) the spread is large enough. Krebs (2007) provides an argument along those lines. Krebs, Krishna, and Maloney (2010) rule out borrowing (credit constraints), which is equivalent to the limit of an infinite borrowing rate. Clearly, the interest rate spread necessary to decentralize the no-trade equilibrium depends on the set of income realization chosen.

References


- Aiyagari, R., 1994, "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109: 659-84.
- Amemiya, T., 1983, "Nonlinear Regression Models" *Handbook of Econometrics*, Amsterdam, The Netherlands 1:333-389.
- Antman, F. and McKenzie, D., 2007, "Earnings Mobility and Measurement Error: A Pseudo-Panel Approach," *Economic Development and Cultural Change*, 56(1): 125-162.
- Atkinson, A. B., 1983, "The Measurement of Economic Mobility," in *Social Justice and Public Policy*, MIT Press, Cambridge, MA.
- Atkinson, A., B., Bourguignon, F., and Morrisson, C., 1992, "Empirical Studies of Earnings Mobility," in Lesourne, J., and Sonnenschein, H., eds., *Fundamentals of Pure and Applied Economics*, Volume 52, Harwood Academic Publishers, Philadelphia.
- Attanasio, O., and C. Binelli (forthcoming)" Mexico in the 1990s: the Main Cross-Sectional Facts," in *Review of Economic Dynamics*.
- Baker, M., and Solon, G., 2003, "Earnings Dynamics and Inequality among Canadian Men, 1976-1992: Evidence from Longitudinal Income Tax Records," *Journal of Labor Economics*, 21: 289-321.
- Banerjee, A., and Duflo, E., *Poor Economics*, Public Affairs: NY
- Benabou, R., and Ok. E. A., 2001a, "Social Mobility and the Demand for Redistribution: The POUM Hypothesis," *Quarterly Journal of Economics*, 116: 447-487.
- Benabou, R., and Ok, E. A., 2001b, "Mobility as Progressivity: Ranking Income Processes According to Equality and Opportunity." *NBER Working Paper Series*, Number 8431.
- Bourguignon, F., Goh, C. and Kim, D., 2006, "Estimating Individual Vulnerability to Poverty with Pseudo-Panel Data" in Morgan, S. L., Grusky, D. B., and Fields G. S., eds., *Mobility and Inequality Frontiers of Research in Sociology and Economics*, Stanford University Press. Stanford.
- Blundell, R., Pissaferrri, I. Peston (2008) "Consumption Inequality and Partial Insurance," *American Economic Review* 99: 1887-1921.
- Carroll, C., and Samwick, 1997, "The nature of precautionary wealth," *Journal of Monetary Economics*, Elsevier, vol. 40(1), pages 41-71

- Chetty, R., Hendren, N., Kline, P., and Saez, E., 2014, "Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States", NBER Working Paper No. 19843
- Collins, D., Murdoch, J., Rutherford, S., and Ruthven, O., 2009, *Portfolios of the Poor*, Princeton University Press, Princeton, NJ
- Conlisk, J., 1990, "Monotone Mobility Matrices," *Journal of Mathematical Sociology*, 15:173-191.
- Constantinides, G., and Duffie, D., 1996, "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy* 104: 219-240.
- Cruces, G., P. Lanjouw, L. Lucchetti, E. Perova, R. Vakis, and M. Viollaz, 2011, "Intra-generational mobility and repeated cross-sections: a three-country validation exercise," Policy Research Working Paper Series 5916, The World Bank.
- Cuesta, J., H. and G. Pizzolitto, 2011, "Using Pseudo Panels To Measure Income Mobility In Latin America," *Review of Income and Wealth*, 57:2 p.224-246.
- Dang, Hai-Anh, P. Lanjouw, J. Luoto, D. McKenzie 2011 "Using Repeated Cross Sections to Explore Movements in and out of Poverty," Policy Research Working Paper 5550.
- Danziger, S., and Gottschalk, P., 1998, "Family Income Mobility: How Much Is There, and Has It Changed?" in Auerbach, J., and Belous, R., eds., *The Inequality Paradox*, Washington D.C.: National Policy Association.
- Dardanoni, V., 1993, "Measuring Social Mobility," *Journal of Economic Theory*, 61: 372-94.
- Davidson, R and J. G. MacKinnon, 2004, *Econometric Theory and Methods* Oxford University Press, New York.
- Daeton, A., 1991, "Saving and Liquidity Constraints," *Econometrica* 59: 1221-1248.
- Deaton, A. and Paxson, C., 1994, "Intertemporal Choice and Inequality," *Journal of Political Economy*, 102(3): 437-467.
- Fields, Gary S. and Ok, E. A., 1996, "The Meaning and Measurement of Income Mobility," *Journal of Economic Theory*, 71: 349-377.
- Fields, Gary S. and Ok, E. A., 1999, "The Measurement of Income Mobility: An Introduction to the Literature," in Silber, J., ed., *Handbook of Inequality Measurement*, Dordrecht, Germany: Kluwer Academic Publishers.

- Fields, Gary S., Cichello, P., Feije, S., Menendez, M., and Newhouse D., 2003 “For Richer or Poorer? Evidence from Indonesia, South Africa, Spain and Venezuela” *Journal of Economic Inequality*, 1: 67-99.
- Friedman, M. 1962. *Capitalism and Freedom*, Princeton, NJ: Princeton University.
- Gallant, A. R. 1975, “Seemingly Unrelated Nonlinear Regressions”’ *Journal of Econometrics*, 3(1): 35-50.
- Geweke, J., Marshall, R., and Zarkin, G., 1986, “Mobility Indices in Continuous Time Markov Chains,” *Econometrica*, 54(6): 1407-1423.
- Glewwe, P., 2004, “How Much of Observed Economic Mobility is Measurement Error? A Method to Remove Measurement Error, with an Application to Vietnam”, —itMimeo, World Bank, Washington DC.
- Gottschalk, P., and Moffit, R., 1994, “The Growth of Earnings Instability in the U.S. Labor Market, *Brookings Papers on Economic Activity*, 2: 217-72.
- Gottschalk, P., and Spolaore, E., 2002, “On the Evaluation of Economic Mobility,” *Review of Economic Studies*, 69(1): 191-208.
- Hart P.E., 1981, “The Statics and Dynamics of income Distributions: A Survey,” in Klevmarket, N. A., Lybeck, J. A., and Tieto, C., eds. *The Statics and Dynamics of Income*.
- Heathcote, J., Storesletten, K., and G. Violante, 2009, “Quantitative Macroeconomics with Heterogeneous Households,” *Annual Review of Economics* 1: 319-354.
- Hnatkovska, Viktoria, Amartya Lahiri, and Sourabh Paul, 2012, “Castes and Labor Mobility.” *American Economic Journal: Applied Economics*, 4(2): 274-307
- Huggett, M., 1993 “The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies,” *Journal of Economic Dynamics and Control* 17: 953-969.
- Instituto Nacional de Estadística y Geografía (INEGI), 2000, “Documento Metodológico de la Encuesta Nacional de Empleo Urbana,” Mexico.
- Krebs, T., 2003, “Human Capital Risk and Economic Growth,” *Quarterly Journal of Economics* 118: 709-744.
- Krebs, T., 2007, “Job Displacement Risk and the Cost of Business Cycles,” *American Economic Review* 97: 664-686.


- Kuznets, S., 1955, "Economic Growth and Income Inequality," *American Economic Review*, 45(1): 1-28.
- Levine, D. and Zame, W., 2002, "Does Market Incompleteness Matter?," *Econometrica* 71: 1695-1725.
- Lillard, Lee A. and Willis, R. J., 1978 "Dynamic Aspects of Earnings Mobility" *Econometrica*, 46: 985-1012.
- Lucas, R., 2003, "Macroeconomic Priorities," *American Economic Review*, 93: 1-14.
- Luttmer, E., 2002, "Measuring Economic Mobility and Inequality: Disentangling Real Events from Noisy Data," *Mimeo*, World Bank, Washington DC.
- Markandya, A., 1982, "Intergenerational Exchange Mobility and Economic Welfare," *European Economic Review*, 17: 307-324.
- Markandya, A., 1984, "The Welfare Measurement of Changes in Economic Mobility," *Economica*, 51: 457-471.
- Meghir, C., and L. Pistaferri, 2011, "Earnings, Consumption, and Life-Cycle Choices," *Handbook of Labor Economics*
- Shorrocks, A. F., 1978a, "Income Inequality and Income Mobility," *Journal of Economic Theory*, 19: 376-393.
- Shorrocks, A. F., 1978b, "The Measurement of Mobility," *Econometrica*, 46(5): 1038-1024.
- Shorrocks, A., 1993, "On the Hart measure of Income Mobility," in M. Casson and J. Creedy, (eds.), *Concentration and Inequality: Essays in Honour of Peter Hart* (Edward Elgar, Aldershot, UK).
- Solon, G., 2001, "Mobility Within and Between Generations" in Welch, F., ed., *The Causes and Consequences of Increasing Inequality*, Chicago: University of Chicago Press.

Figure 1: Variance of Unpredicted Part of Earnings vs. Age (1987-2003)



Note: Variance is the coefficient on age from a regression of the Mincer residual squared on age and year dummies. Estimates from Mexican Urban Employment Survey using individuals age 20-65. 5% confidence intervals

Figure 2: Covariance of Unpredicted Part of Earnings across 5 Quarters vs. Age (1987-2003)



Note: Covariance is the coefficient on age from a regression of the covariance of the Mincer residual in quarter 1 vs. quarter 5 on age and year dummies. Estimates from Mexican Urban Employment Survey using individuals age 20-65. 5% confidence intervals.

Table 1: Summary Statistics: 1987-2003

	<i>Mean</i>	<i>Std Deviation</i>	<i>Min</i>	<i>Max</i>
<i>Age</i>	36.271	10.626	20	65
<i>Schooling</i>	10.624	5.460	0	22
<i>Sex</i>	0.737	0.440	0	1

Note: Based on the Mexican Monthly Urban Employment Survey, 1987-2003 using individuals between 20 and 65 years of age. Age and schooling in years.

Table 2: Mincer Regression

	<i>Coef</i>	\overline{Sd}	<i>t</i>	<i>p</i> > <i>t</i>	
Cons	3.699	0.009	422.450	0.000	
Sex	0.310	0.002	191.140	0.000	
Sch	0.077	0.001	143.160	0.000	
Sch ²	-0.001	0.000	-45.380	0.000	
Age					
21	0.044	0.005	8.730	0.000	
22	0.088	0.005	17.540	0.000	
23	0.127	0.005	25.740	0.000	
24	0.173	0.005	34.570	0.000	
25	0.208	0.005	41.530	0.000	
26	0.242	0.005	47.780	0.000	
27	0.268	0.005	52.450	0.000	
28	0.288	0.005	56.510	0.000	
29	0.309	0.005	60.240	0.000	
30	0.328	0.005	64.350	0.000	
31	0.348	0.005	66.740	0.000	
32	0.360	0.005	68.540	0.000	
33	0.370	0.005	71.270	0.000	
34	0.382	0.005	71.890	0.000	
35	0.389	0.005	73.360	0.000	
36	0.391	0.005	73.160	0.000	
37	0.407	0.005	75.150	0.000	
38	0.422	0.005	77.940	0.000	
39	0.421	0.005	76.910	0.000	
40	0.426	0.006	77.270	0.000	
41	0.442	0.006	76.630	0.000	
42	0.451	0.006	77.750	0.000	
43	0.448	0.006	76.700	0.000	
44	0.459	0.006	74.430	0.000	
45	0.455	0.006	74.150	0.000	
46	0.450	0.006	70.690	0.000	
47	0.452	0.007	66.900	0.000	
48	0.441	0.007	64.660	0.000	
49	0.430	0.007	61.210	0.000	
50	0.434	0.007	60.680	0.000	
51	0.431	0.008	56.920	0.000	
52	0.430	0.008	54.200	0.000	
53	0.423	0.008	52.360	0.000	
54	0.420	0.009	48.510	0.000	
55		0.398	0.009	44.750	0.000
56		0.400	0.009	42.730	0.000
57		0.393	0.010	39.600	0.000
58		0.367	0.011	34.870	0.000
59		0.356	0.011	32.000	0.000
60		0.322	0.011	28.640	0.000
61		0.307	0.012	24.860	0.000
62		0.302	0.014	22.000	0.000
63		0.286	0.015	19.640	0.000
64		0.309	0.016	19.460	0.000
65		0.247	0.016	15.360	0.000
<i>Year and wave dummies</i>			Yes		
N			782179		
<i>R</i> ² Adj			0.595		

Note: Regression of log income on sex, age as a dummy variable, schooling, schooling square and a year time specific dummy and a dummy for whether the data correspond to the first period or the fifth. Data are pooled across all years. Based on the Mexican Monthly Urban Employment Survey, 1987-2003, using individuals between 20 and 65 years of age.

Table 3: Estimation of Mobility Parameters

=		
	<i>Full Sample</i>	<i>Restricted</i>
ρ	0.977*** (0.0019)	0.976*** (0.0037)
σ_ω^2	0.104*** (0.0038)	0.104*** (0.0068)
σ_ϵ^2	0.015*** (0.0008)	0.016*** (0.0016)
σ_η^2	0.203*** (0.0039)	0.217*** (0.0073)
<i>Time Dummies</i>	Yes	Yes
R^2	0.13	0.13
N	387460	99570

*Note: Estimation using Non-linear SUR estimation. Dependent variables: Eq 1 variance, Eq 2 covariance. Variance calculated as the square of the residual of the mincer regression. Covariance as the covariance of the residual in the first quarter observed with that of the fifth quarter. ρ represents the autoregressive coefficient or convergence parameter. σ_ω^2 represents the variance of the initial distribution of income. σ_ϵ^2 represents the variance of permanent shocks. σ_η^2 represents the variance of the transitory or measurement error component of income. A complete and separate set of time dummies is included in each equation. Estimates using the Mexican Monthly Urban Employment Survey, 1987-2003, using individuals between 20 and 65 years of age. Column 1 uses all observations. Column 2 just those beginning Q1 of each year. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.*

Table 4: Mobility Analysis

	Time Span (years)		
	1	10	25
Computed Mobility	0.674	0.763	0.846
% Δ if:			
$\rho = 0$	-0.77	-5.2	-8.6
$\sigma_\epsilon^2=0$	-1.2	-6.6	-7.4
$\sigma_\eta^2=0$	-89.7	-45.7	-23.3

Note: Table shows the percentage decline in residual income mobility as component parameters are individually set to zero relative to actual mobility calculated from equation (7) using parameters estimated in Table 3 based on the Mexican Monthly Urban Employment Survey, 1987-2003. ρ represents the autoregressive coefficient or convergence parameter. σ_ϵ^2 represents the variance of permanent shocks. σ_η^2 represents the variance of the transitory or measurement error component of income. Mobility is calculated by using estimated income parameters in (7) across a span, t , of 1, 10 and 25 years.

Table 5: Welfare Analysis

	$\sigma_\epsilon^2 = 0$	$\bar{\rho} = 1$	$\sigma_\epsilon^2 = 0$ and $\rho = 1$
% Change			
$\beta = 0.96$	12.56	-8.04	10.5
$\beta = 0.95$	10.64	-5.82	8.91
$\beta = 0.94$	9.21	-4.45	7.72
$\beta = 0.90$	5.87	-2.05	4.93

Note: Table shows the percentage change in welfare calculated measured as a percent of lifetime consumption as σ_ϵ^2 , the variance of permanent shocks, is set to 0 (no income risk) and ρ , the convergence parameter, is set to one (no convergence). β is the annual discount factor. Welfare is calculated using equations (11) and (14) and the estimated values in Table 3 using the Mexican Monthly Urban Employment Survey 1987-2003.