RISK AND THE EVOLUTION OF INEQUALITY IN CHINA IN AN ERA OF GLOBALIZATION

ETHAN LIGON

Abstract. Recent increases in urban income inequality in China are mirrored in increases in increasing inequality in consumption expenditures. This connection between changes in the distribution of income and consumption expenditures could be entirely due to differences in preferences (in which case households' intertemporal marginal rates of substitution would all be equated after every history), or could be due to imperfections in the markets for credit and insurance which would ordinarily serve to equate these intertemporal marginal rates of substitution. In this paper we presume that market imperfections drive changes in the distribution of expenditures, and use data on expenditures from repeated cross-sections of urban households in China to estimate a Markov transition function for shares of expenditures over the period 1985–2001. We then use this estimated function to compute the welfare losses due to risk over this period, and to predict the future trajectory of inequality from 2001 through 2025.

1. Introduction

The welfare loss due to risk faced by households at a point in time is intimately related to changes in inequality in expenditures. In particular, risk-averse households with time separable preferences will tend to prefer to smooth shocks to income over time, so that even entirely transitory shocks to income will tend to have a permanent effect on future consumption expenditures. Thus, the same shocks to income which makes next period's consumption uncertain will also determine the household's position in next period's distribution of expenditures.

In this paper we exploit this link by using data on the evolution of expenditure inequality to estimate both household risk preferences and the welfare loss due to risk actually borne by urban Chinese households over the period 1985–2001, an era during which China's economy has undergone dramatic reforms and experienced remarkable growth. Others have noted that increases in urban inequality imply that the rising tide of the aggregate Chinese economy has not lifted all boats

Date: August 18, 2004.

equally (Kahn and Riskin, 2001). Here we note that because households may change their position in the wealth (and expenditure) distribution, merely looking at changes in inequality will understate the displacement and (*ex ante*) welfare loss experienced by risk-averse households facing dramatic economic change.

Models having complete markets á la Arrow-Debreu yield fully Pareto efficient outcomes; in such a model any changes in inequality must be preferred by all market participants, and so yield little in terms of interesting policy implications. Complete market models which feature Gorman aggregable preferences yield the very strong prediction that the distribution of consumption across households is invariant. More interesting are models in which some friction prevents allocations from being fully Pareto optimal, and which have enough dynamic structure to yield interesting predictions regarding the evolution of the distribution of consumption.

Here we assume that all households have similar preferences, and constant relative risk aversion (Arrow, 1964). We further assume that all households have access to credit markets on equal terms, and that households exploit these credit markets to smooth their consumption over time, á la the permanent income hypothesis. Beyond this, we make no notably restrictive assumptions. We allow quite arbitrary forms of technology and shocks, and avoid the problem of measuring asset returns. Though this framework is quite general in several dimensions, we will show that conditional on the distribution of production shocks the model yields rather sharp predictions regarding the evolution of the distribution of resources across households. In particular, the model gives us the law of motion governing the inverse Lorenz curves which describe inequality in the economy.

The law of motion for inverse Lorenz curves allows us to make predictions about the sequence of Lorenz curves we would expect to observe, conditional on household risk preferences, rates of aggregate economic growth, and on the distribution of unforecastable shocks facing households in different years, at different wealth levels, and in different occupations. By comparing realized and predicted Lorenz curves, we can estimate these preferences and distributions. This same procedure yields a Markov transition function mapping shares of consumption today in to a probability distribution over possible shares tomorrow, and we use this object to calculate the risk borne by differently situated urban Chinese households in different years, and relate this risk to measures of globalization during this period.

The key to the empirical strategy of this paper involves exploiting the restrictions placed on data by Euler equations to make statements about the evolution of inequality. Related literature includes Deaton and Paxson (1994), who derive a martingale property from the consumption Euler equation and use several long panels of household level expenditure data to argue that within-cohort inequality in industrialized countries is increasing over time, and Storesletten et al. (2002) who use household panel data on expenditures from the U.S. and a more completely specified general equilibrium model to estimate a law of motion for the distribution of consumption. The central idea of those papers is to exploit intertemporal restrictions to estimate the law of motion for individual households' consumption growth, and then in effect to integrate over households to infer what the law of motion is for the distribution of consumption *across* households. The present paper reverses these last two steps—we derive equations which impose intertemporal restrictions on individual households' consumption growth, but then integrate over these equations to obtain restrictions on the the law of motion for the distribution of consumption across households before taking these restrictions to the data. The cost of the procedure followed in this paper is that one can't exploit all the information which would be available from the trajectories of consumption for many different individual households—the (closely related) benefit is that can get by without panel data, using instead only a relatively limited set of data obtainable from repeated cross-sectional surveys of household expenditures, of the sort that many countries conduct in order to, e.g., compute consumer price indices.

2. The Data

The data requirements for the exercise described in this paper are quite modest. The main requirement is simply that have a sequence of data on consumption shares for several population quantiles—that is to say, just the data conveyed by a sequence of Lorenz curves for consumption expenditures.

2.1. **Changes in Inequality in China, 1985–2001.** At beginning of the period for which we have data, inequality in China was remarkably low.

While one still doesn't observe gross inequities in the distribution of consumption in China, over the course of 1985–2001 one *does* observe an extremely rapid rate in the rate at which inequality is increasing, even when one confines one's attention to the urban population. Presumably this extraordinarily rapid increase in inequality has been made

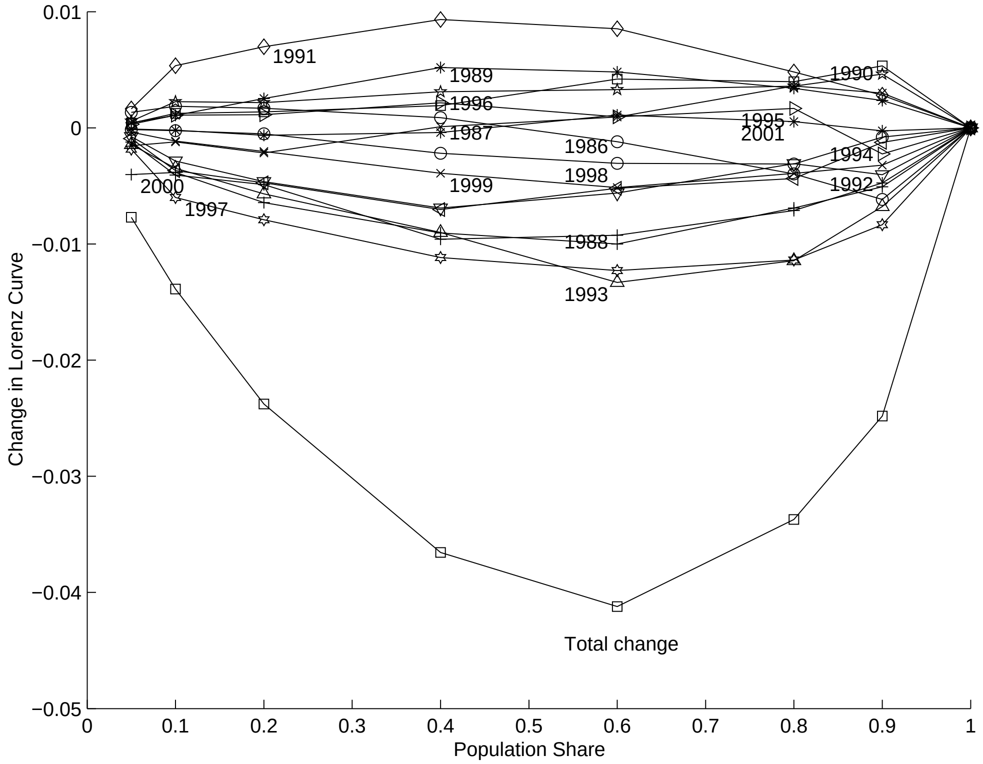

FIGURE 1. Changes in Consumption Lorenz Curves for Registered Urban Households

politically palatable by the similarly high rates of economic growth over the same period—though the worst off urban households are consuming a rapidly shrinking share of the the aggregate consumption pie, the pie itself has been growing

3. The Model

In this section we describe a model in which households can exchange debt in competitive credit markets, and derive restrictions on the the evolution of each household's share of aggregate consumption. The key assumptions we'll exploit are that households all have similar preferences featuring constant relative risk aversion, and that all households have access to credit on the same terms. Note that this latter assumption is weaker than assuming that credit markets are perfect—in particular, it may be the case that at the interest rates faced by households for some reason credit markets fail to clear.

Consider, then, an environment with n infinitely lived households. We index these households by $i = 1, 2, \dots, n$. Time is discrete, and is indexed by t. Household i derives momentary utility from consumption according to some function $u_i : \mathbb{R} \to \mathbb{R}$, and discounts future utility at a common rate $\beta \in (0, 1)$.

3.1. **Intertemporal Restrictions.** At any date t, household i can exchange claims to consumption at $t + 1$ with other households at a price $1/\rho_t$, solving the problem

$$\max_{b_{it}} u_i(c_{it} - b_{it}/\rho_t) + \beta \mathrm{E}_t \left\{ u_i(c_{it+1} + b_{it}) + \sum_{j=2}^{\infty} \beta^{j-1} u_i(c_{it+j}) \right\} \tag{1}$$

where E_t denotes the expectations operator conditional on information available at time t, b_{it} denotes the debt issued by the household at time t, and c_{it} denotes the households time t consumption expenditures.

The first order conditions associated with the household's problem of debt-issuance at time t indicate that the household will consume c_{it} at t if the the usual Euler equation

$$u_i'(c_{it}) = \beta \rho_t \mathrm{E}_t u_i'(c_{it+1}) \tag{2}$$

is satisfied.

It's convenient to restrict our attention to the case in which utility functions exhibit constant relative risk aversion, so that

$$u_i(c_{it}) = \frac{c_{it}^{1-\gamma} - 1}{1 - \gamma} \tag{3}$$

where γ is the coefficient of relative risk aversion. In this case, (2) implies that

$$[\beta \rho_t]^{-1} = \mathrm{E}_t \left(\frac{c_{it+1}}{c_{it}} \right)^{-\gamma}$$

for all $i = 1, \dots, n$ and all t. As a consequence, we have

$$\mathrm{E}_t \left(\frac{c_{it+1}}{c_{it}} \right)^{-\gamma} = \frac{1}{n} \sum_{j=1}^{n} \mathrm{E}_t \left(\frac{c_{jt+1}}{c_{jt}} \right)^{-\gamma}. \tag{4}$$

We can interpret this as a prediction that with the ability to freely exchange debt all households will have the same expected growth in their marginal utilities of consumption.

Because we want to understand the links between intertemporal restrictions on consumption such as (4) and the evolution of inequality, we'll translate (4) into a statement about shares of consumption expenditures. Let $\bar{c}_t = \sum_{i=1}^n c_{it}$ denote aggregate consumption expenditures at t, and $\sigma_{it} = c_{it}/\bar{c}_t$ denote i's share of expenditures at t. Then

$$\mathrm{E}_t \left\{ \left[\left(\frac{\sigma_{it+1}}{\sigma_{it}} \right)^{-\gamma} - \frac{1}{n} \sum_{j=1}^n \left(\frac{\sigma_{jt+1}}{\sigma_{jt}} \right)^{-\gamma} \right] \left(\frac{\bar{c}_{t+1}}{\bar{c}_t} \right)^{-\gamma} \right\} = 0.$$

Let

$$\epsilon_{it+1} \equiv \left[\left(\frac{\sigma_{it+1}}{\sigma_{it}} \right)^{-\gamma} - \frac{1}{n} \sum_{j=1}^n \left(\frac{\sigma_{jt+1}}{\sigma_{jt}} \right)^{-\gamma} \right] \left(\frac{\bar{c}_{t+1}}{\bar{c}_t} \right)^{-\gamma} \tag{5}$$

denote household i's time t "forecast error" relative to the average forecast error. Note from the properties of (5) that $\mathrm{E}_t \epsilon_{it+1} = 0$, as is usual when evaluating forecast errors from Euler equations. However, as Chamberlain (1984) points out, in the usal analysis there may be an aggregate shock which induces correlation across households' forecast errors in the cross-section, so that there's no guarantee that realized forecast errors at $t+1$ will in fact average to zero. We've avoided this problem here by eliminating ρ_t; for us $\frac{1}{n}\sum_{j=1}^n \epsilon_{jt} = 0$ by construction.

3.2. **Risk.** The uncertainty facing any individual household with consumption share σ_{it} at time t which may reduce its utility at time $t+1$, then, can be summarized by three random variables. First of these just has to do with variation in the growth of aggregate consumption, $g_{t+1} \equiv \bar{c}_{t+1}/\bar{c}_t$. Second are household-specific surprises which may change the household's share of aggregate expenditures, ϵ_{it}. Third and finally, the aggregate of surprises facing all other households may change the distribution of resources, $\eta_{t+1} \equiv \frac{1}{n}\sum_{j=1}^n \left(\frac{\sigma_{jt+1}}{\sigma_{jt}}\right)^{-\gamma}$.

Define the idiosyncratic risk borne by the household at time t to be the *ex ante* loss in expected utility at $t+1$ due solely to variation in the purely idiosyncratic shock ϵ_{it+1}, or

$$R_{it} \equiv u_i(\bar{c}_{t+1}\sigma_{it}) - \mathrm{E}[u_i(c_{it+1}) | I_t, \eta_{t+1}, g_{t+1}] \tag{6}$$

where I_t denotes the information set at time t. Here the first term is the utility the household would obtain at $t+1$ if the household's share of expenditures was unchanged (as would be the case if no household faced any idiosyncratic risk) and the household knew in advance wat aggregate consumption would be in $t+1$. The second term is the utility the household would *expect* if it somehow knew in advance what the realization of all the relevant aggregate random variables would be, so that it remained ignorant only of the idiosyncratic shocks it would experience in the first period.

In a world with complete markets and the assumed CES preferences we work with here, it's easy to establish that each households' share of aggregate consumption will remain constant, eliminating all idiosyncratic risk. Thus, we can interpret the first term of (6) as the utility the household would obtain if no households bore any idiosyncratic risk less the expected utility of consumption when the household does bear this idiosyncratic component of risk. It's trivial to establish that this *cardinal* measure of risk is uniquely consistent (up to an linear transformation of u_i) with the notion of increasing risk defined by Rothschild and Stiglitz (1970). Because our measure of idiosyncratic risk is denominated in utils, it's straightforward to construct a variety of useful measures of the welfare loss associated with this risk.

Let $\Psi_t(\sigma_{it+1}|\sigma_{it}, x_{it}, g_{t+1}, \eta_{t+1})$ denote the time t Markov transition function for the household's share σ given household characteristics x_{it} and knowledge of the aggregate quantities g_{t+1} and η_{t+1}. Then the expression for household i's time t idiosyncratic risk, as defined above, may be written

$$R_{it} = u_i(\bar{c}_{t+1}\sigma_{it}) - \int u_i(\bar{c}_{t+1}\sigma')d\Psi_t(\sigma'|\sigma_{it}, x_{it}, g_{t+1}, \eta_{t+1}).$$

Note that idiosyncratic risk can depend on both household characteristics x_{it} as well as on the household's current position in the consumption distribution, σ_{it}. Let the distribution of characteristics x of households having a share σ of aggregate consumption time t be given by $G_t(x|\sigma)$. Then to calculate *average* idiosyncratic risk of households with share σ we integrate out the characteristics x, obtaining the marginal Markov transition function $\tilde{\Psi}_t(\sigma'|\sigma, g_{t+1}, \eta_{t+1}) = \int \left[\int^{\sigma'} d\Psi_t\right] dG_t(x|\sigma)$.

Let the distribution of σ at t be given by $\Gamma_t(\sigma)$ (this is the inverse of the Lorenz curve). Average idiosyncratic risk is then given by

$$R_t = \int u_i(\bar{c}_{t+1}\sigma)d\Gamma_t(\sigma) - \int u_i(\bar{c}_{t+1}\sigma')d\tilde{\Psi}_t(\sigma'|\sigma, g_{t+1}, \eta_{t+1})d\Gamma_t(\sigma).$$

3.3. **Distribution.** The Markov transition function $\tilde{\Psi}$ which is critical for calculating average risk in the population is also critical for understanding how the distribution of resources changes over time. In particular, the inverse Lorenz curves $\{\Gamma_t\}$ satisfy a law of motion

$$\Gamma_{t+1}(\hat{\sigma}) = \int_{\{\sigma'<\hat{\sigma}\}} d\tilde{\Psi}_t(\sigma'|\sigma, g_{t+1}, \eta_{t+1})d\Gamma_t(\sigma). \tag{7}$$

Accordingly, knowledge of the transition functions $\tilde{\Psi}_t$ suffices to characterize both average risk as well as the evolution of inequality in the population.

3.4. **Forecast Errors and Markov Transitions.** Recall that an individual household's uncertainty depends only on relative forecast errors ϵ_{it+1}, η_{t+1}, and on g_{t+1}. In particular, we can use (5) to express $\tilde{\Psi}_t$ in terms of the distribution of relative forecast errors in the population. Let ϵ_{it+1} have the cumulative probability distribution $F_t(\epsilon|\sigma_{it}, x_{it})$. Then note from (5) that we have

$$\sigma_{it+1} = \sigma_{it}\left(g_{t+1}^{\gamma}\epsilon_{it+1} + \eta_{t+1}\right)^{-1/\gamma}, \tag{8}$$

so that

$$\tilde{\Psi}_t(\sigma'|\sigma, g_{t+1}, \eta_{t+1}) = \int\int_{\left\{\epsilon > \frac{(\sigma/\sigma')^{\gamma}-\eta_{t+1}}{g_{t+1}^{\gamma}}\right\}} dF_t(\epsilon|\sigma, x)dG_t(x|\sigma). \tag{9}$$

Let $\tilde{F}_t(\epsilon|\sigma) = \int dF_t(\epsilon|\sigma, x)dG_t(x|\sigma)$ denote the marginal distribution of relative forecast errors ϵ for households having consumption share σ. Then, because time $t+1$ shares must integrate to one, we have the adding up restriction

$$\int \sigma[\eta_{t+1} + \epsilon g_{t+1}^{\gamma}]^{-1/\gamma} d\tilde{F}_t(\epsilon|\sigma)d\Gamma_t(\sigma) = 1, \tag{10}$$

which pins down the value of the $\{\eta_t\}$ in terms of the remaining objects in (10). Accordingly, given knowledge of the distributions $\{(F_t, G_t)\}$, the sequence of realized aggregate consumptions to pin down $\{g_t\}$, and the risk aversion parameter γ, we can completely describe the evolution of inequality and the distribution of risk in the population.

4. Empirics

How can we go about using data on the evolution of Lorenz curves to estimate the risk borne by differently situated households? Given our maintained assumption of equal access to credit markets and some initial distribution of consumption shares Γ_0, (7) allows us to trace out changes in the distribution over time given knowledge of the Markov transition functions $\{\tilde{\Psi}_t\}$ and of the sequence $\{g_t, \eta_t\}$. However, each $\{\Psi_t\}$ must be consistent with the law of motion for shares (8), while the unknown sequence $\{\eta_t\}$ is determined by the adding-up restriction (10). As a consequence, the extent of our ignorance regarding $\{\tilde{\Psi}_t\}$ amounts to ignorance regarding the risk-aversion parameter γ, and the marginal distributions of relative forecast errors in each period, $\{\tilde{F}_t\}$.

Though we don't begin with knowledge of the distributions of errors $\{\tilde{F}_t\}$, the first moment of each $\tilde{F}_t$ must be equal to zero by (2), while the support of the distribution at t must be a subset of $[\eta_{t+1}/g_{t+1}^{\gamma}, \infty)$. After examining the empirical distribution of estimated relative forecast errors for a small panel of urban Chinese households[1], it appears that this empirical distribution at t adequately represented by what Johnson and Kotz (1970) call the "three-parameter log-normal distribution," with $\log(\epsilon - \theta_{t+1})$ distributed $N(\mu_t(\sigma), v_t^2(\sigma))$. Of the three parameters $(\theta_t, \mu_t(\sigma), v_t(\sigma))$ only two are free, with $\theta_t = -\eta_t/g_t^{\gamma}$ (since shares must all lie in the $(0, 1)$ interval) and $\mu_t = \log(-\theta_t) - v_t^2(\sigma)/2$ (since the expected value of ϵ must be zero).

With these restrictions on the distribution of relative forecast errors, the only things which remain for us to infer from data are the coefficient of relative risk aversion γ and the scale parameters $\{v_t(\sigma)\}$. In practice we only work with a finite number (say n) of share values, and have only a finite number of periods (T) of data on the distribution of consumption. We impose a log-linear structure on these scale parameters, assuming that for every year $t = 1, \dots, T$ and share $\sigma \in \{\sigma_1, \dots, \sigma_n\}$ there exists an ℓ-vector of observable variables x_{it} which determines the scale parameters via

$$\log v_t(\sigma_i) = \delta' x_{it}. \tag{11}$$

for some ℓ-vector δ. This assumption allows us to estimate a set of ℓ parameters which may be presumed to be smaller than Tn, and guarantees that estimated values of $v_t(\sigma)$ will be positive, as they must be to be interpretable as the standard deviation of a normally distributed variable.

As a consequence of the foregoing, we're left with the problem of estimating $\ell + 1$ parameters $b_0 = (\gamma, \delta')$. We have data on the share of consumption expenditures for population quantiles $(x_1, x_2, \dots, x_m)$ for each of $T + 1$ years. We use these data on consumption expenditures to approximate the Lorenz curves $\{L_t(x)\}_{t=0}^T$ of expenditure shares. We fix an initial guess of our parameters b. Noting that $L_t = \Gamma_t^{-1}$, conditional on this guess we use the law of motion (7) (along with the adding-up restriction (10), and (11)) to predict a sequence of Lorenz curves $\{\hat{L}_t(x|b)\}_{t=1}^T$. We compute a simple measure of distance between

[1]We have data on monthly expenditures for the 2002 round of the National Bureau of Statistics Urban Household Survey for the province of Shandong.

FIGURE 2. Growth rate necessary to compensate for risk, by quantiles. Each line indicates the minimum rate of growth necessary to compensate for the risk faced by the average household at the consumption share quantile indicated at the far right. For example, the line labelled '5' gives the rate of growth in each year necessary to compensate the average household *ex ante* at the 5 per cent quantile for the risk borne by that household.

the predicted and actual Lorenz curves

$$d(b) = \sum_{t=1}^{T} \sum_{i=1}^{m} \left(L_t(x_i) - \hat{L}_t(x_i|b) \right)^2, \tag{12}$$

and then use a simplex minimization routine to find the value $\hat{b} = \operatorname{argmin}_b d(b)$.

5. Results

In this section we present results for different specifications of the variance structure, the general form of which is given by (11). We

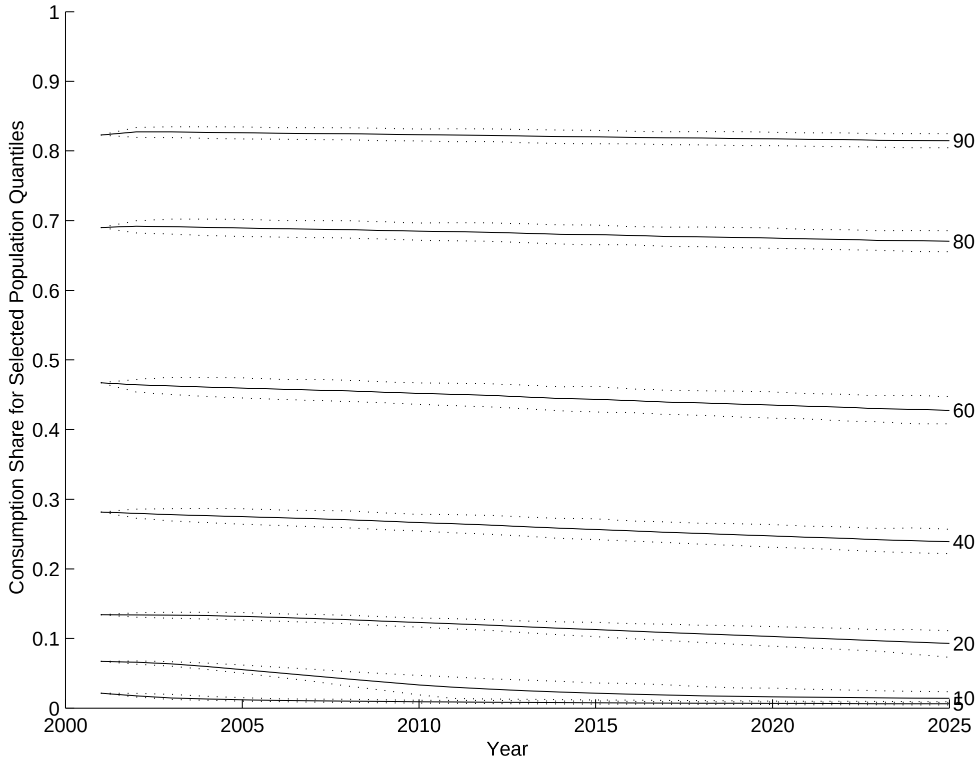

FIGURE 3. Predicted evolution of consumption shares $L(x)$ for selected population quantiles x.

begin with (adapting some language from Amemiya (1984) a simple 'error components' structure, permitting the log of the variance of the relative forecast error to depend on the sum of a year-specific constant and a quantile-specific constant, so that

$$\log v_{it} = \alpha_i + \nu_t.$$

Here v_{it} is the standard error of the relative forecast shock for a household in the ith quantile of the consumption share distribution in year y; in practice we divide this distribution into 17 different quantiles, but for the sake of identifying these parameters we constrain the top quantile to have $\alpha_{17} = 0$.

Table 1 presents the fitted parameters given this 'error components' variance structure. The leftmost panel of the table shows parameters which vary across years. Estimates of the normalizing constants $\{\eta_t\}$ appear in the first column of this panel, while the "year effects" part of the variance structure, $\{\nu_t\}$ appears in the second column. Recall from our earlier discussion the interpretation of η_t as a measure of the

Year	η_t	ν_t	Quantile	α_i	γ	R^2
1986	1.0635	-1.2862	0.0139	-3.5888	0.7229	0.4723
1987	1.1295	-1.1886	0.0278	0.0009		
1988	1.2037	-1.1411	0.0539	-0.0004		
1989	1.0884	-1.2590	0.0800	3.8415		
1990	1.0463	-1.1884	0.1142	-0.6764		
1991	1.1009	-1.2512	0.1485	-0.5182		
1992	1.1228	-1.1603	0.2280	-0.0510		
1993	1.2187	-1.1499	0.3075	-0.2359		
1994	1.2660	-1.1199	0.4008	1.5517		
1995	1.2004	-1.2411	0.4940	0.0261		
1996	1.0926	-1.2052	0.6030	0.3490		
1997	1.0868	-1.1026	0.7120	-0.0000		
1998	1.0682	-1.2444	0.7756	0.0001		
1999	1.1003	-1.2407	0.8393	-0.1229		
2000	1.1205	-1.1694	0.9197	0.7127		
2001	1.1064	-1.3337				

TABLE 1. Parameter Estimates assuming log-normal relative forecast errors and variance specification I.

aggregate uncertainty at time t—this specification gives us a simple way to check the model, since ν_t provides a direct measure of aggregate uncertainty. In this case, the correlation between the two measures is 0.47, consistent with our expectations.

The primary virtue of this 'error components' specification of the structure of the variance of relative forecast errors has to do with the simplicity of interpreting estimates of α_i and ν_t. In particular, for years in which ν_t is relatively large, the entire population faces greater risk than usual. At the same time the specification allows for variation in uncertainty by wealth (consumption share); the average households in a quantile for which α_i is negative faces less uncertainty in an average year than do the very wealthiest households, while the average household in a quantile with α_i greater than zero faces more.

Turning our attention to differences in the uncertainty faced by households across the distribution, consider the second panel of Table 1. Here we see that the households in the bottom quantile who collectively consume 1.4 per cent of the aggregate face the least uncertainty, with a quantile "fixed effect" of -3.59. However, households in the 5–8 per cent quantile face the most, with an estimated quantile fixed effect of 3.84. Eight of the consumption share quantiles bear more uncertainty than does the topmost quantile, while seven bear less.

At this point let us pause a moment to be careful about what is meant by "uncertainty" above. Differences in the parameters $\{\nu_t\}$ across time or $\{\alpha_i\}$ across quantiles are really just related to the standard deviation of the relative forecast errors ϵ_{it}. These relative forecast errors have numerous desirable properties, but don't have a straight-forward interpretation either in terms of the welfare costs of uncertainty or in terms of variation in quantities which might be observable. In particular, the distribution of ϵ_{it} depends both on household risk-preferences (here γ) and on the distribution of consumption growth, making it critical to estimate γ and the variance structure of the forecast errors simultaneously. In the present case, the estimated value of the coefficient of relative risk aversion is 0.723. This is on the low end of the range of estimates of this parameter in the micro-econometric literature, but doesn't seem obviously wrong.

To get a sense of the magnitude of the risk facing individual households, we use the parameters reported in Table 1 to estimate the risk facing the average household at selected consumption-share quantiles in Figure 2. To construct this figure we've started with estimates of the measure of risk given by (6), but rather than reporting the welfare loss due to uncertainty in utils (which may be difficult to interpret), we've computed the growth rate of aggregate consumption expenditures which would be just enough to compensate households for the risk they bear. In particular, for a household at consumption-share quantile σ we find the growth rates $g_t(\sigma)$ such that

$$u(\bar{c}_{t+1}\mathrm{E}[\sigma_{t+1}|\sigma_t, \eta_{t+1}, \bar{c}_{t+1}]) = \mathrm{E}[u(\bar{c}_{t+1}\sigma_{t+1}g_{t+1}(\sigma_t))|\sigma_t, \eta_{t+1}, \bar{c}_{t+1}].$$

In the present case, because we assume CES utility functions future aggregate consumption $\bar{c}_{t+1}$ cancels out of this equation. By substituting in our estimates of the parameters $\{\eta_t\}$ and $\{\tilde{\Psi}_t\}$ a simple line-search algorithm can be used to find the compensating growth rates.

These compensating growth-rates are shown for selected consumption-share quantiles in Figure 2. Note first that the growth we're referring to is the rate of growth in aggregate consumption expenditures for urban households—this quantity grew at an average annual rate of 12 per cent over the period 1985–2001. Using the quantity $g_t(\sigma)$ as our measure of the welfare loss of uncertainty, the poorest (displayed) quantile of households is much the worst off—from 1985 to 1986 these households would have needed urban expenditures to have grown by nearly 14 per cent before they would have preferred the status quo to stagnation and an "iron rice bowl." Setting the poorest households aside, risk does increase in a monotone way, with households at the 92 per cent quantile requiring compensation which never exceeds five per cent. Thus, were we to graph it, this measure of risk would display a "U" pattern, with the poorest households bearing a great deal of risk, low income households bearing the least, and risk gradually increasing with consumption shares throughout the rest of the distribution.

We now turn our attention from risk to inequality. Recall that we're now able to construct estimates of the Markov transition functions $\{\tilde{\Psi}_t\}$. If these functions were invariant across time, it would be a trivial matter to calculate the future evolution of the distribution of consumption for as many periods as we choose, simply by using an estimate of some initial distribution Γ_0, and then applying (7) iteratively to trace out future distributions.

Of course, matters are not quite so simple. Instead, we have estimates of $\tilde{\Psi}_t$ for 16 different values of t, from 1985–2000, and while it's a simple matter to trace out the predicted trajectory over the course of this sample period, this tells us little about future inequality. We adopt the following simple strategy. Given our collection of 16 different estimated transition functions, we simply assume that these functions are representative of the kinds of transition functions which may be realized in the future. Thus, to estimate the evolution of the distribution of consumption over τ periods we simply make τ random draws (with replacement) from the collection of transition functions $\{\tilde{\Psi}_t\}$. Starting with the actual distribution of consumption shares in 2001, we substitute these draws sequentially into (7); inverting the resulting function Γ_t yields an estimate of the Lorenz curve L_t. Then we use these τ equations to calculate one possible sample trajectory of the Lorenz curves, which we denote by $\{\hat{L}_t^1\}_{t=1}^{\tau}$. We repeat this procedure many times, so that we have a bootstrapped sample of m possible trajectories for the Lorenz curve over time, or $\mathcal{L} = \{\{\hat{L}_t^i\}_{t=1}^{\tau}\}_{i=1}^{m}$.

Now, for any population quantile x we can compute a "mean" trajectory by computing

$$\bar{L}_t(x) = \frac{1}{m}\sum_{i=1}^{m} L_t^i(x),$$

or characterize the distribution of possible trajectories by simply working with the bootstrap sample $\mathcal{L}$.

Figure 3 shows values of $\bar{L}_t(x)$ for selected values of x[2] (the solid lines) along with 80 per cent confidence intervals, for predicted trajectories beginning in 2001 and running through 2025. The figure has several notable features. First, note that the confidence intervals are very tight, relative to the variation across population quantiles. This is a reflection of a fact already noted above—differences across households are much more pronounced than differences across time. The very small variation in our estimated "time effects" $\{\nu_t\}$ and normalizing constants $\{\eta_t\}$ mean that in fact our estimated transition functions don't change very much over time at all; as a consequence it doesn't matter very much what actual sequence of transition functions we draw in our bootstrap exercise.

Thus emboldened, we henceforth refer to the evolution of $\bar{L}_t$. Our estimated model predicts that inequality will continue to increase in China through 2025, but at a relatively slow rate. However, the bottom ten per cent of the population will, by then, consume a *much* smaller share in 2025 (2 per cent) than at present (6.5 per cent). Neglecting the welfare costs of risk discussed above, to keep the *level* of consumption constant for this poorest 10 per cent of the population, aggregate urban consumption must grow at an average rate of about five per cent to compensate this part of the population whose share is sharply declining. To compensate for risk, of course, much higher growth rates would have to be sustained.

6. Conclusion

China's economy has changed dramatically over the last two decades, but household level data to understand the effects of China's growth and openint to the outside world are very difficult to come by—data from China's National Bureau of Statistics either have very limited coverage or are very aggregated.

In this paper we make a silk purse of a sow's ear by using aggregate data on the distribution of consumption expenditures across (registered) urban households to construct a sequence of Lorenz curves, and then use intertemporal restrictions on individual households' consumption expenditures implied by optimizing behavior by risk-averse households to derive the restrictions on the evolution of these Lorenz curves implied by theory. The evolution of the Lorenz curve turns out to depend on just two kinds of objects: household utility functions,

[2]These are those available in the China Statistical Yearbooks, and are equal to $(0.5, 0.1, 0.2, 0.4, 0.6, 0.80, 0.9)$

and the distribution of "relative forecast errors" for intertemporally optimizing households.

To pin down household utility, we assume that household preferences exhibit constant relative risk aversion. To pin down the distribution of relative forecast errors we use data from a small subset urban households and construct the empirical probability distribution for these relative forecast errors from these data, up to a set of parameters.

For any estimate of the coefficient of relative risk aversion and parameters governing the distribution of relative forecast errors, we are able to predict a sequence of future Lorenz distributions. We compare this predicted trajectory with the actual sequence of distributions realized between 1986–2001, choosing our preference and distributional parameters so as to minimize a measure of the distance between these sequences of Lorenz curves.

We present two major empirical findings. First, the risk (*ex ante* welfare loss due to variation in future consumption) borne by households depends much more on households' resources than it does on the year—even though there are enormous changes in China's aggregate economy over this period, idiosyncratic risk is much more important than any aggregate shock in determining household welfare and in determining evolution of inequality over time.

Second, our estimates of the law of motion governing the Lorenz curves for urban China allow us to make predictions about future consumption inequality. Looking at the entire distribution, we predict that most of the increase in inequality between 1985 and 2025 has already occurred; however, we also predict that the share of consumption accruing to the poorest decile of these households will continue to fall at a relatively rapid rate, lowering the share of consumption for these households from 6.5 per cent in 2001 to only two per cent in 2025.

References

Amemiya, T. (1984). Tobit models: A survey. *Journal of Econometrics 24*(1–2), 3–61.

Arrow, K. J. (1964). The role of securities in the optimal allocation of risk bearing. *Review of Economic Studies 31*, 91–96.

Chamberlain, G. (1984). Panel data. In Z. Griliches and M. D. Intriligator (Eds.), *Handbook of Econometrics*. North-Holland.

Deaton, A. and C. Paxson (1994). Intertemporal choice and inequality. *Journal of Political Economy 102*(3), 437–467.

Johnson, N. L. and S. Kotz (1970). *Continuous Univariate Distributions - 1*. New York: John Wiley & Sons.

Kahn, A. R. and C. Riskin (2001). *Inequality and Poverty in China in the Age of Globalization*. Oxford: Oxford University Press.

Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. A definition. *Journal of Economic Theory 2*, 225–243.

Storesletten, K., C. I. Telmer, and A. Yaron (2002). Consumption and risk sharing over the life cycle. Forthcoming in the Journal of Monetary Economics.

Department of Agricultural & Resource Economics, University of California, Berkeley