Stuck at the Bottom: Caste-Based Discrimination, Relative Poverty and the Mechanisms of Intergenerational Opportunity Traps[1]

Sudeshna Maitra
Dept. of Economics,
York University
4700 Keele Street,
Toronto, Canada

This Version: May 2026

Email: smaitra.econ@gmail.com (preferred)
maitra@yorku.ca

[1] I acknowledge funding and support from UNU-WIDER, as part of their project titled *Inequality over the life course*. An earlier version of this paper was co-authored with Sudarsana Kundu, of Gender at Work and Nayi Disha (Hyderabad, India). I owe Ms. Kundu my sincere thanks for her insights on marriage market calibrations and assessments of social status standards in urban India, which feature in this work. Thanks also to Anika Tabassum Amreen, Gulnur Arazli and Abrar Iqbal at York University for their research assistance on this project.

Stuck at the Bottom: Caste-Based Discrimination, Relative Poverty and the Mechanisms of Intergenerational Opportunity Traps

Sudeshna Maitra

York University

May 2026

ABSTRACT

Caste-based labour market discrimination can take two economically distinct forms — differential access to high-wage employment and differential wages for the same type of employment . Each form perpetuates intergenerational inequality through fundamentally different channels and require fundamentally different policy responses. We develop an overlapping-generations model of class formation to separate these channels analytically, and calibrate it to urban India using the Periodic Labour Force Survey (PLFS) 2025. Theoretical experiments show that access barriers worsen intergenerational immobility and distort class composition toward the relatively poor, while wage penalties compress within-class consumption for disadvantaged households without altering class membership – findings that have direct implications for which policy instruments to deploy and which metrics to use in evaluating them. Calibrated to contemporary urban India, we find no absolute poverty trap at 2025 PLFS-estimated parameter values: the relevant form of deprivation is positional rather than absolute. Our simulations further show that while both channels are simultaneously active in contemporary India (the simulated economy resembles the access-barrier experiment in its pattern of intergenerational immobility and the wage-penalty experiment in its pattern of within-class consumption intensity), wage penalties currently impose the larger loss of opportunities for disadvantaged households. We propose the Stuck-at-the-Bottom Index, a steady-state measure of class persistence decomposed by caste, to identify which channel is driving immobility and to guide the selection and evaluation of appropriate policy tools.

1 Introduction

How do systemic modes of discrimination – such as a caste system, or racial bias – impact the transmission of intergenerational opportunities in a developing country? What are the *mechanisms* by which discriminatory norms in labour and marriage markets affect such opportunities? How could we measure such opportunities over time – using nationally representative cross-sectional data – to derive their impact on relative poverty and intergenerational mobility over time? How could such analysis guide anti-discriminatory policy making?

The above questions are especially important in the context of India, where caste-based economic inequality has proved remarkably persistent despite several decades of affirmative action policy. Reservations in public employment and higher education have been in place since the 1950s; yet caste-based earnings gaps and hiring discrimination against SC/ST workers in the private sector remain large and well-documented (Madheswaran and Attewell, 2007; Thorat and Attewell, 2007; Deshpande, 2011). Likewise, intergenerational mobility in India has remained persistently low, with lower caste households disproportionately unlikely to escape their class origins across generations (Hnatkovska et al, 2012; Asher et al., 2024).

While reducing caste-based discrimination remains an important objective of Indian policymakers, the dominant policy discourse tends to conflate two structurally distinct forms of labour market discrimination: differential access to high-wage employment (access barriers) and differential wages for the same type of employment (wage penalties). Most empirical work (see earlier references) measures discrimination through earnings gaps, which combine both channels. Most policy tools — reservations, most prominently — target only the access channel.

The most rigorous recent evidence on intergenerational mobility in India (Asher et al, 2024) also illustrates the conceptual difficulty around disentangling "access" and "wage penalty" effects. Asher et al find that affirmative action has substantially improved Scheduled Caste education rank mobility — consistent with the "access" channel being responsive to policy. But education rank measures the access channel alone, and is orthogonal to the "wage penalty" channel. Two equally educated workers of different caste, who face different wages would be indistinguishable in education rank-based data. If within-education-group wage disparities are large then (education) rank convergence overstates the reduction in total discrimination. Estimates of remaining discrimination based on such measures will be systematically too optimistic.

In this paper, we adopt a novel approach using "theoretical experiments", to understand the impact of specific sources of economic opportunities – "access" versus "wage premiums/penalties" – on class formation, long-run poverty status and intergenerational mobility. Our experiments, calibrated by estimates from the Periodic Labour Force Survey (2025), provide evidence that both channels are active in contemporary urban

India, and that wage penalties currently impose the larger loss of opportunities for the disadvantaged caste.

Our approach is rooted in an overlapping-generations (OLG) model of income generation and household choice, following Maitra (2024). In this model, households accumulate durable goods over generations through two interdependent markets: (1) a labour market where (unobserved) productivity and education type (high or low) determine whether they receive a high-wage or a low-wage job, and (2) a marriage market where accumulated durables signal social status and thereby determine the partner's wage contribution (high-wage or low-wage). Note that the marriage market is a key feature for tracking caste-based inter-generational mobility, since caste status in Indian society is traditionally derived by birth, and marriage/mating is governed by norms of caste endogamy (Deshpande, 2011; Banerjee et al, 2013).

Maitra (2024) shows that, in steady state, classes emerge endogenously – through a wealth-begets-wealth (or inter-generational persistence) effect – in the distribution of optimal household durable consumption across communities. The formulation of the OLG model permits the analysis of two separate forms of discrimination in the labour market: access-based discrimination, where low-caste households face lower probabilities of earning the high wage for any productivity-education type; and wage-penalties' based discrimination, where low-caste households earn lower wages conditional on job type. The model is thus the vehicle for showing – in a fully dynamic setting with feedback through the marriage market – that these two channels operate through fundamentally different mechanisms.

Conditional on a given set of OLG model parameters, we derive the transition matrix governing evolving household income and choices, and solve for the steady state "population" distribution of durable consumption. We then draw and analyze a synthetic dataset from the population distribution, much as data is drawn in household surveys. This approach allows us to tweak various model specifications (viz. its structure and parameters) and test how the tweaks change household choices observed in the synthetic data – a process that is akin to running "experiments" by altering the "theoretical" engine of the OLG data-generating process.

To quantify intergenerational persistence in class membership, we propose a Stuck-in-class$-j$ Index, defined as the proportion of households in class j in period$-(t-1)$ who remain in class j in period t, evaluated at steady state. The Stuck-in-class$-\hat{j}$ Index for the lowest (relatively poorest) class $\hat{j}$ is called the Stuck-at-the-Bottom (S-B) Index. The 'Stuck' indexes are motivated by the literature on transition-matrix-based mobility measurement (Shorrocks, 1978); however, they are not axiomatic mobility indexes. Instead they represent targeted diagnostics, viz. steady-state class-persistence indicators, decomposed by caste, that directly measure the impact of different types of discrimination on intergenerational immobility.

We conduct 3 theoretical experiments to demonstrate the mechanisms modeled in our framework, and analyze the impact of discrimination on relative poverty and its persistence. In our first experiment, we assume a just society where all model parameters are identical for households of both castes (i.e. there is no discrimination). This case provides a benchmark against which to compare the impact of systemic discrimination. In Experiment 2 – unjust society "P" – we assume that L-caste households face lower probabilities of receiving high-wage employment in the labour market than households who are H-caste. Finally, in Experiment 3 – unjust society "W" – we assume that L-caste households receive lower wages than H-caste households for performing the same jobs. There is ample evidence of both these forms of caste-based discrimination in Indian society (Deshpande, 2011; Banerjee and Knight, 1985; Vaid, 2014). Experiments 2 and 3 allow us to isolate the effects of each – employment access barriers (P) versus wage penalties (W).

In Sections 3.1-3.4, we report results from running theoretical experiments 1-3 using arbitrary values of model parameters, and assuming 2 caste types – high (H) and low (L). We find that the P channel (access barriers, tracked by comparing Society P with J) expands the lower class (or the "subsistence" class) and raises the L-caste share among the relatively poor, while leaving within-class living standards unchanged. However, the W channel (wage penalties, tracked by comparing Society W to J) leaves class sizes and caste composition of classes unchanged in the bottom classes, while reducing durable consumption within these classes for L-caste households. Furthermore, channel W also reduces the share of L-caste households in the upper class by diminishing their capacity to signal high status in the marriage market. These (theoretical) findings suggest that the two channels P and W are not observationally equivalent. They affect different features of the class distribution and require different metrics to detect and address.

In Section 4, we calibrate the model parameters to contemporary urban India, using PLFS 2025 data. We now simulate 4 societies – J, P, W (as before) and A. Society A represents contemporary urban India with *both* channels of discrimination active at the same time. Three findings stand out in the calibrated simulations. First, contemporary urban "pseudo" India (Society A) exhibits no absolute poverty trap, i.e. the subsistence class does not emerge at parameter values calibrated to PLFS data. The society now features two new forms of positional vulnerability: a volatile "Low-High Oscillator" class and a low-mobility/low-durables "Middle" class, both disproportionately populated by L-caste households. Second, the calibrated pseudo-urban Indian economy behaves like Society P for intergenerational immobility and like Society W for within-class consumption intensity, confirming that both channels are simultaneously active in A. Third, wage discrimination (the W channel) is currently the more severe channel as measured by the caste-based opportunities' gap in the labour market, defined as the ex-ante labour-market earnings gap in Societies P and W. The opportunities' gap induced by wage penalties (over a generation: Rs. 71.39 lakhs for Low-caste versus Rs. 98.83 lakhs for High-caste) clearly exceeds that induced by access barriers (Rs. 80.69 lakhs for Low-caste), a finding with clear implications for the ordering of policy priorities.

The findings from the calibrated experiments reported above offer several pointers for framing and evaluating anti-discrimination policy. The experiments demonstrate that the two forms of caste-based discrimination – access barriers (P) and wage penalties (W) – perpetuate intergenerational inequality through fundamentally different channels, and require fundamentally different policy responses. The more stringent of the two forms (measured by the opportunities' gap) should determine which tools to prioritize at any time. For contemporary urban India, we find a more urgent need for policy that affects wage penalties. Having selected and deployed specific policy tools, the evaluation of their impact should focus on the outcomes they *do* influence. Thus, policies targeting employment access (reservations, job training) address the P channel and should be evaluated with relative measures, e.g. the L-caste share of the classes and the Stuck-at-the-Bottom index. Policies targeting the wage gap (minimum wages, pay equity) address the W channel and should be evaluated with within-class consumption measures. Using the wrong metric to evaluate a policy tool could produce systematic understatement of effective interventions, a point of practical importance in India's ongoing policy debate around caste-based affirmative action.

The primary motivation – and contribution – of this paper is, then, to develop a practical framework, informed by theory and data, for use by policymakers in formulating anti-poverty and anti-discrimination policy. A growing literature on intergenerational mobility has established that initial inequality in endowments generates persistent inequality across generations through well-defined mechanisms — such as, the transmission of human capital and social status (Loury, 1981; Galor and Zeira, 1993; Corak, 2013) — but these models abstract away from categorical group identity as an independent source of persistent disadvantage. The literature on caste discrimination in India has documented precisely this – viz. that 'caste status' independently predicts earnings, occupational attainment and hiring outcomes, even after conditioning on education and other productivity characteristics (Banerjee and Knight, 1985; Thorat and Attewell, 2007; Madheswaran and Attewell, 2007; Deshpande, 2011; Hnatkovska et al 2012). However, the latter studies work largely in cross-section or reduced form, and do not model how labour market discrimination compounds through the marriage market into intergenerational class-persistence. Overlapping generations models of class formation (Loury, 1981; Galor and Zeira, 1993; Maitra, 2024) provide the machinery to bridge both observations. But no existing model embeds categorical discrimination with the analytical separation needed to distinguish "access barriers" from "wage penalties", in tracking their structurally different effects on class dynamics. This paper attempts to bridge that gap.

2 The Theoretical Framework

The theoretical framework we will use is the same as the one presented in Maitra (2024), but with the additional assumption of the existence of a caste system – characterized by caste-based discriminatory norms in the labour market and a social norm of caste endogamy in the marriage market. Section 2.1 below is drawn largely from Maitra (2024). In Section 2.4, I extend the framework to include a caste system.

2.1 The Model without Caste (from Maitra, 2024)

The decision making process of each household is modeled using an overlapping generations framework in which every household is defined by three generations – 0 (child), 1 (parent) and 2 (grandparent). In any period t, the earners and decision-makers in a household are the parents (generation 1), while children and grandparents (generations 0 and 2) are dependents. Parents choose the level of education of their children and the amount of durables they wish to purchase, ensuring, first, that a level of subsistence consumption, C (> 0), is met. C is the theoretical counterpart of the absolute poverty line typically used to measure poverty. The common household utility of members in any period t is given by

$$U(B_t) = C + B_t \tag{1}$$

where B_t is the total value of durables present in the household in t.

Note that B_t represents total accumulated durables, viz. durables purchased by period$-t$ parents in period t as well as those acquired by period$-t$ grandparents in period $(t-1)$, i.e.

$$B_t = b_t + b_{t-1} \tag{2}$$

where b_k indicates the durables purchased in period k by period$-k$ parents. (1) and (2) indicate that when period$-t$ grandparents pass so do the durables they accumulated when they were parents (in period $(t-1)$), viz. durables depreciate over time.

Household income in any period is the sum of incomes of two parents: one who was born and raised in the household in question and the other that married into the household. The income of the parent born in the household is low (w_L) or high (w_H) depending on two factors: (1) whether that parent is of low or high productivity (α_L and α_H, respectively) and (2) whether they have high or low education (e_H and e_L, respectively). The productivity level α_{t-1} of the generation of parents in period t is determined randomly at the time of their birth (in $(t-1)$) and is unobservable. We assume that productivity is α_L with probability q_L (α_H otherwise). Likewise, the education level of the period$-t$ parents is determined by the amount invested in it by *their* parents when they were children, i.e. in $(t-1)$. We denote the wage of the parent raised in the household by w_1 and refer to it as the household's 'labour income'.

The income of the parent who marries into the household is assumed to depend on the social standing of the household, which determines marriage market success. We assume that marriages are arranged and that households with higher social standing – as measured by the value of durables observed to be in use in that household (B_t) – attract partners with a higher wage. In particular, we assume that a household that has B durables in any period will attract a partner with high wage w_H with probability $\Phi_S(B)$, where $\Phi_S(B)$ is the cumulative distribution function of a normal distribution $N(\beta, \sigma^2)$. The latter assumption has the following interpretation (see Figure 1).

In any period, there is a certain level of durables ownership, β, that is generally acknowledged to mark households of high social standing. The 'skepticism' around this common belief is represented by σ^2. Higher accumulated durable ownership B increases the probability of attracting a partner with a high wage, with the rate of increase in probability highest around the level $(\beta - \epsilon)$ ($\epsilon > 0$, small). Moreover, the higher the skepticism (σ^2) regarding the common social standard β, the lower the increased probability of acquiring a high-wage partner at most levels of accumulated durables B, around β. An example of a society with low skepticism (or low σ^2) would be one where everyone agrees unanimously on the connection between durables and social standing, such as might be likely in small, close-knit communities in rural settings. Higher skepticism could occur in more anonymous communities such as might exist in urban settings. We will henceforth refer to $\Phi_S(B)$ as the signal function under beliefs $S = (\beta, \sigma^2)$; where β denotes the social standard and σ^2 denotes the skepticism regarding β. Further, we denote the wage of the parental partner by w_2 and call it the household's 'marriage market income'.

Household income in period t can be written as $I_t(e_{t-1}, \alpha_{t-1}, B_{t-1})$, where e_{t-1} is the education level of period$-t$ parents, α_{t-1} is their (random, unobserved) productivity level and B_{t-1} is the total number of durables in the household (signaling its social standing) when period$-t$ parents were matched in the marriage market. In particular,

$$I_t(e_{t-1}, \alpha_{t-1}, B_{t-1}) = w_{1t}(e_{t-1}, \alpha_{t-1}) + w_{2t}(B_{t-1}) \quad (3)$$

where w_{1t} (labour income in period t) is w_H with probability $p(e_{t-1}, \alpha_{t-1})$, w_L otherwise, and w_{2t} (marriage market income in t) is w_H with probability $\Phi_S(B_{t-1})$ (w_L, otherwise); $S = (\beta, \sigma^2)$. We assume

$$p(e_L, \alpha_L) = p_1 \tag{4}$$

$$p(e_L, \alpha_H) = p_2 \tag{5}$$

$$p(e_H, \alpha_L) = p_3 \tag{6}$$

$$p(e_H, \alpha_H) = p_4 \tag{7}$$

$$0 < p_1 < p_2 < p_3 < p_4 < 1 \tag{8}$$

where condition (8) – along with our additional assumptions that p_1, p_2 are small and p_3, p_4 are large – indicate that education is *effective* in generating high labour market income.

Household expenses (E_t) in any period consist of three components: (1) the subsistence consumption level C that must be met, (2) the investment in education of the generation of children in that period, and (3) the expenditure on durables:

$$E_t = C + c(e_t) + b_t \tag{9}$$

where $c(e_t)$ represents the parental generation's spending on education and b_t is the spending on durables in period t. We assume that there are two possible levels of education e_t – high (e_H) and low (e_L) – and that the cost $c(e_t)$ of providing the same are E (> 0) and 0, respectively. We also assume that there are no savings opportunities, so the residual household income after spending C and $c(e_t)$ is used to purchase durable goods.

The timing of income-realization and decision-making is as follows. At the beginning of any period t, parents find themselves with (realized) income I_t based on the education level and productivity of one parent(α_{t-1}, e_{t-1}), and the wage of the other parent determined by the household's durables level B_{t-1}. Given income I_t, parents choose their children's education level e_t and the amount to spend on durables b_t in order to maximize their own expected lifetime utility $[U(B_t) + \delta E_t U(B_{t+1})]$, where δ (ϵ $(0, 1)$) is the discount factor. At the end of period t, the current generation of children (with education e_t) enters the labor market and earns a wage based on their productivity draw and the education e_t they have received. In addition, the total value of durables ($B_t = b_{t-1} + b_t$) in the household in period t determines the wage of their partner by arranged marriage: w_H with probability $\Phi_S(b_{t-1} + b_t)$, w_L otherwise. The sum of own wage and partner's wage determines the household income of the parental generation I_{t+1} in the next period.[1]

The optimization problem of the parental generation in period t can be written as:

$$\underset{(e_t, b_t)}{Max}\ U(b_{t-1} + b_t) + \delta E_t U(b_t + b_{t+1}) \tag{10}$$

$$subject\ to$$

$$c(e_t) + b_t \leq I_t(e_{t-1}, \alpha_{t-1}, b_{t-1} + b_{t-2}) - C \tag{11}$$

$$E_t b_{t+1} = E_t[I_{t+1}(e_t, \alpha_t, b_t + b_{t-1}) - C - c(e_{t+1})] \tag{12}$$

$$b_t \geq 0 \tag{13}$$

where all durable expenditures in the continuous interval $[0, Max\{0, I_t(e_{t-1}, \alpha_{t-1}, b_{t-1} + b_{t-2}) - C\}]$ are *feasible* choices for the household.

Notice how the period$-t$ decision variables (e_t, b_t) impact the decision makers' (parents') lifetime utility. The spending on children's education e_t represents a trade-off between current and future consumption, since it involves an expenditure now that increases income (potentially) in the future $[(11) - (13)]$. However, the spending on current durables b_t improves consumption now as well as in the future since, (i) it increases direct consumption utility in both periods (10) and, (ii) it also increases the potential of higher income (hence, consumption) in the future (12).

It is is easy to see that solving the optimization exercise in $(10) - (13)$ under the model assumptions reduces to ascertaining which of the two education levels (e_H or e_L) generates a higher expected lifetime utility for the decision maker, conditional on their realized income (I_t) and their inherited durable stock (b_{t-1}). The residual income after spending on this optimal education level and subsistence consumption is assigned to durables. We can show that households choose the higher education level (e_H) if the following condition holds:[2]

$$(1 + \delta)E + \delta(w_H - w_L)[q_L(p_1 - p_3) + (1 - q_L)(p_2 - p_4) + (\tilde{p} - \tilde{p}_e)] \tag{14}$$

$$-\delta E[Pr(e_{t+1} = E/e_t = e_L) - Pr(e_{t+1} = E/e_t = e_H)] < 0$$

[1]Notice the intergenerational impact of household decisions – the income of parents in any period t depends on the investment in their education by their parents (viz. $(t-1)-$parents) and the accumulated durables they have inherited from their parents (b_{t-1}) and grandparents (b_{t-2}).

[2]See Maitra (2024).

where $\tilde{p} = \Phi_S(b_{t-1} + I_t - C)$ and $\tilde{p}_e = \Phi_S(b_{t-1} + I_t - C - E)$ are the values of the signal function under $S = (\beta, \sigma^2)$ when e_L (with cost 0) or e_H (with cost E) is chosen, respectively.

The model outlined in (1) – (14) describes a stochastic process $\{b_{t-1}, b_t\}$, driven by the following parameters:

$$(w_H, w_L, p_1, p_2, p_3, p_4, \beta, \sigma^2, q_L, \delta, E, C)$$

In this model, total household income in any period t could take one of 3 possible values – $(2w_L)$, $(w_L + w_H)$ or $(2w_H)$. For each of the 3 values of household income, there are 2 possible (optimal) choices for durable expenditure (b_t), corresponding to whether education level e_L or e_H is chosen in t (as per condition (14)). Thus, in any period t, optimal durable expenditure b_t could take one of 6 $(= 2 \times 3)$ possible values. Furthermore, for each of the 6 possible values of b_t, there are (similarly) 6 possible values of b_{t-1}. These comprise 36 (i.e. 6^2) "states" that $\{b_{t-1}, b_t\}$ can pass through in any period t. Thus, the transition matrix P that governs the movement from (b_{t-1}, b_t) to (b_t, b_{t+1}) is of order (36×36).

Let us make the additional simplifying assumption that $C = 2w_L$, which implies that households with the lowest income level $(2w_L)$ can barely afford to pay for subsistence consumption; hence they always choose $e_t = e_L$. This reduces the possible values that b_t can take, to 5 (instead of 6). This leads to 5^2 $(= 25)$ possible "states" and a transition matrix of order (25×25).

Let $\theta_1, \theta_2, ..., \theta_{25}$ denote the 25 possible "states" or values that the process $\{b_{t-1}, b_t\}$ can pass through in any period t. Note that each θ_i $(i = 1, 2, ..., 25)$ has, associated with it, an amount of total durables $(b_{t-1i} + b_{ti})$ observed in a household in state θ_i in period t.[3] Furthermore, let $x_t = (x_{1t}, x_{2t}, ..., x_{25t})$ denote the proportions of households in states $\theta_1, \theta_2, ..., \theta_{25}$ respectively, in the population in period t $(0 \leq x_{it} \leq 1, \sum_{i=1}^{25} x_{it} = 1, i = 1, 2..., 25)$. Thus, households' transition through various states of durables expenditure can be written as:

$$x_t P = x_{t+1} \tag{15}$$

Moreover, the steady state distribution of durables expenditures over states $(\theta_1, \theta_2, ..., \theta_{25})$, denoted by $x^* = (x_1^*, x_2^*, ..., x_{25}^*)$, will satisfy the condition:

$$x^* P = x^* \tag{16}$$

[3]Clearly, $(b_{t-1} + b_t)$ is not unique across the 25 states, since any household that has (b_0, b_1) will be observed to own the same amount of total durables as a household with (b_1, b_0). These 2 households would, however, have very different transition probabilities to other possible states since the durables accumulated by grandparents play a role in determining the optimality of education (condition (14)).

It can be shown that the stochastic process described above converges to a steady state (Tsokos, 1972).

The model (in $(1)-(14)$) serves to power an economically meaningful data-generating process (DGP) that generates the income (and durables') distributions in the economy in each period. We will focus on *steady state equilibrium* (16) in this economy, viz. one where the same distribution of states (durables) persists over time, and where households' expectations about future opportunities are fulfilled in each period.

2.2 Measuring Intergenerational Mobility: The "Stuck-at-the-Bottom" Index

We measure inter-generational mobility in the above model by focusing on the income "history" of households, in steady state equilibrium. Recall that every state of the transition matrix maps to a history of income draws (I_t, I_{t-1}) in two consecutive generations, where I_t and I_{t-1} can each be *low* $(= 2w_L)$, *middle* $(= w_L + w_H)$ or *high* $(= 2w_H)$. Hence, the steady-state proportion of households yield a distribution of "income histories," captured by 'classes'. We assign classes, as follows, to households based on their income history (I_t, I_{t-1}).

Class	Steady-State (I_t, I_{t-1})
Lower (L)	(low, low)
Lower-Middle (LM)	$(low, middle)$ or $(middle, low)$
Middle (M)	$(middle, middle)$
Upper-Middle (UM)	$(middle, high)$ or $(high, middle)$
Upper (U)	$(high, high)$
Low-High Oscillator (LHO)	$(low, high)$ or $(high, low)$

We define inter-generational income persistence of as the persistence of income level j (= *low*, *middle*, *high*) in two consecutive generations. In particular, we define the "Stuck-at-the-Bottom" Index as follows.

For each income level j of I_t (j = *low*, *medium*, *high*) the "Stuck-in-j" index is the proportion (in steady state) of income-j households in period $(t-1)$ who also draw income j in period t. Thus the "Stuck-in-j" indices map to the Lower (L), Middle (M) and Upper (U) classes defined in the table above.

The "Stuck-at-the-Bottom" index is the "Stuck-in-$\hat{j}$" index for the lowest level $\hat{j}$ observed in steady state. Households in class $\hat{j}$ represents the positionally disadvantaged group, or the group in "relative poverty." In addition, the class that draws the $'low'$ income $(2w_L)$ in consecutive periods is defined as the "subsistence class," viz. the class that is persistently in absolute poverty, in steady state.

A Stuck-in-j index of value 0 indicates perfect intergenerational mobility out of class j; a value of 1 indicates perfect immobility ("stuck" in class j). In the experiments conducted below, we report the class structure that emerges in the simulated steady state equilibrium, along with the "Stuck" indices described above.

The next subsection describes how simulated data on steady-state household durable expenditures is generated from the model presented in Section 2.1 above.

2.3 Example of an economy without systemic discrimination

Example 1

$w_H = 100, w_L = 10, p_1 = 0.1, p_2 = 0.3, p_3 = 0.4, p_4 = 0.8,$

$\beta = 150, \sigma = 110, q_L = 0.5, \delta = 0.5, E = 11.5, C = 20.$

Figure 2(a) plots the derived steady state distribution for Example 1 – the $x-$axis showing $(b_t + b_{t-1})$, the total accumulated durable spending observed in any time t; and the $y-$axis showing the proportion of households that would be observed in steady state at each level of total durable spending.

Figure 2(a) shows the theoretical steady state distribution – or the data-generating process (DGP) – that corresponds to the parameters in Example 1. It is possible now to draw samples from a population that is distributed as in Figure 2(a). For example, Figure 2(b) plots the histograms of durable expenditures observed in two samples of 1000 observations, drawn independently from the theoretical steady state distribution (the DGP) in Figure 2(a).

Figure 2(c) corresponds to an economy comprising 1000 communities (such as in Example 1). The parameter values representing each community is from uniform distributions over specified ranges (see figure). We draw a sample of 1000 households from each community and pool these observations to generate a synthetic dataset ($N = 1,000,000$) representing the economy (of 1000 communities). Figure 2(c) is the histogram of total durable expenditures in the pooled synthetic dataset.

Note that, in the synthetic dataset, the two-generational incomes of households (I_t, I_{t-1}), and hence their 'class' is known. This allows us to track class structure – thence intergenerational mobility – in the current population. Moreover, by changing the values of model parameters, we can track changes in class structure generated by different economic channels in the model. This is the very exercise we undertake in this paper.

In the next sub-section, we introduce a caste system in the theoretical model.

2.4 The Model with Systemic Discrimination (A Caste System)

Assume now that households can be characterized as high (H-caste) or low (L-caste) caste families. Due to caste endogamy, households marry only other households of the same caste. This ensures that the caste identity of households is determined "at birth", i.e. it does not change over time.

Recall that the economic and social conditions in a community are encapsulated in the model parameters $\{w_H, w_L, p_1, p_2, p_3, p_4, \beta, \sigma, q_L, \delta, E, C\}$. In a "just society" these parameters would necessarily be the same for all households, regardless of categorical group identities such as caste status. The very first theoretical "experiment" we will conduct in Section 3 is, therefore, to track the process of class formation – especially the size and characteristics of the "lowest" class – in a just society J. Our findings for the just society will serve as a benchmark against which we can compare the effects of systemic discrimination, in subsequent experiments described below.

We will assume that labour market discrimination can take place in 2 fundamental forms, consistent with widespread evidence of caste-based discrimination in India (Deshpande, 2011; Banerjee and Knight, 1985), as well as the larger literature on racial and gender-based discrimination (Becker, 1971; Darity, 2005, Phelps, 1972). The first form – employment-access barriers – operates through the probability parameters (p_1, p_2, p_3, p_4). Specifically, we will simulate a society – call it Unjust Society P – where L-caste households face lower probabilities (p_1, p_2, p_3, p_4) of receiving high-wage employment at each level of productivity and education, than H-caste households. The second form of discrimination – in what we will call "Unjust Society W" – operates as wage penalties, wherein L-caste households face lower wages (w_H, w_L) than H-caste households, for the same jobs. Our model permits us to separate the impacts of these two effects on the lowest emergent class, thence on relative poverty.

In Section 3 below, we run 3 experiments – (1) Just Society J (2) Unjust Society P (3) Unjust Society W – albeit, with arbitrary parameter values, to compare the size and characteristics of the lowest class obtained in unjust societies versus the just society.

3 Theoretical Experiments

Suppose that our economy/society is comprised of 1000 communities. We assume that each community j $(j = 1, 2, ..., 1000)$ consists of 30% H-caste households and 70% L-caste households. The marriage market in the society is characterized by caste endogamy whereby households only marry into the same caste that they are born into. This ensures that the caste identity of households remain the same over time and generations.

Also, caste-specific wages earned in the labour market (which may be subject to discrimination) are repeated in the marriage market, since the spouse too must belong to the same caste. Finally, we assume that the signal function (Φ_S), as well as the social standard (β) and skepticism (σ) parameters are the same for households of all castes. This assumption is motivated by the fact that all households (irrespective of caste) belong to the same society and have access to the same signals of social status as each other (presumably from consuming the same TV shows, movies and other aspects of observable culture).

The parameters of H-caste households in community j ($j = 1, 2, ..., 1000$) are drawn from uniform distributions over the following (arbitrary) ranges. Note that the set of realized parameter values (drawn from the uniform distribution), remain the same for H-caste households in *all* the experiments.

Parameters

$w_H \in (80, 120), w_L \in (5, 20), p_1 \in (0, 0.5), p_2 \in (p_1, 0.5), p_3 \in (0.5, 1), p_4 \in (p_3, 1),$

$\beta \in (500, 3500), \sigma \in (100, 500), q_L = 0.5, \delta = 0.5, E \in (10, 100), C = 2w_L.$

The specific assumptions around parameters for L-caste households are listed below:

1. In Experiment 1 (Just Society J), we assume that L-caste households in each community j have the same values of ALL parameters as H-caste households from community j.

2. In Experiment 2 (Unjust Society P), we assume that the L-caste households in each community j have lower probabilities ($p_i, i = 1, 2, 3, 4$) of securing high-wage (w_H) employment. In particular, for low-caste households in any community j, $p_i^{L-caste,j} = x^j p_i^{H-caste,j}$ ($i = 1, 2, 3, 4$), where $x^j \sim U(0.70, 0.75)$. All parameters other than probabilities are the same across both castes.

3. In Experiment 3 (Unjust Society W), we assume that in every community j, L-caste households receive lower wages (w_H and w_L) than the H-caste households. All parameters other than probabilities are the same across both castes.

4. To ensure comparability between Experiments 2 and 3, we assume that in both these cases, the ex-ante expected labour market income for L-caste households is the same.[4] Clearly, in unjust societies, the ex-ante expected labour market income for L-caste households is lower than in the just society. Per our assumption, this (lower) expected income is driven by lower probabilities (access barriers) in Experiment 2, and by lower wages (wage penalties) in Experiment 3.

[4] The ex-ante expected labour market income is the income expected *before* (unobserved) productivity is drawn and education level is chosen. It is equal to $[q_L\{p_1w_H+(1-p_1)w_L+p_3w_H+(1-p_3)w_L\}+(1-q_L)\{p_2w_H+(1-p_2)w_L+p_4w_H+(1-p_4)w_L\}]$.

As per the steps of our methodology, we start by deriving the steady state distribution of total durable expenditures – the data-generating processes (DGP) – for each caste, in each community j ($j = 1, 2, ..., 1000$). Then we draw a caste-stratified sample of households from each community – 300 from the H-caste DGP and 700 from the L-caste DGP – such that we have synthetic samples of 1000 households from each community. Finally, we combine the observations from each community to form our "pooled" synthetic dataset of 1,000,000 households. In the results for each experiment, we will present analyses of the pooled synthetic dataset of households for the experiment in question.

3.1 Experiment 1: A Just Society "J"

Results for a just society are presented in Tables 1(i)-(ii) (columns 1-3). The first observation of note is that class-based inequality occurs even in a just society characterized by a complete absence of any form of systemic discrimination. The specific size and characteristics of the classes reflect the nature (and overall abundance/scarcity) of economic and social opportunities that are available in the economy/society, as encapsulated in the DGPs of the constituent communities. But the relative success of some households over others are also influenced by the uncertainties ("luck") inherent in the process of converting opportunities into income/durables over time. Therefore, the society in question is "just" in the sense of equal *access* to opportunities of all households – as captured in model parameters. The birth-characteristics of households is immaterial in the process of *who* moves ahead or falls behind in a systematic manner – hence we see that the caste composition of *each* class (Table 1, Panel A, Column 3) is the same as that in the population (70:30). However, *equal access* to opportunities does not translate into *equal realization* of outcomes; hence we see class formation in realized total durable expenditures. Our model plays a useful role in characterizing this distinction – between *access* to and *realization* of opportunities – in a measurable way; thus providing a basis for estimating "injustice" by deviation from the outcomes of a just society.

3.2 Experiment 2: Unjust Society "P" (differential probabilities or access barriers)

Tables 1(i)-(ii) (columns 4-6) report the results from Experiment 2 – an unjust society where L-caste households have lower probabilities of receiving high-wage employment than H-caste households. Notice that relative poverty – the size of the lower class – is higher (55.67%) in Society P than in the just society (46.72%). Moreover, the share of L-caste households (74.87%) in the lower class is also higher in P than in the overall population (70%). Finally, note that the mean durable expenditure of each class in Table 1(ii) (columns 5 versus 6) is quite similar to those in J, even though it is less in P for the entire sample (46.34 in P vs. 58.24 in J). This effect is due to the expansion of the lower class in P (that has an excess L-caste households) than in the just society.

The results above suggest that differential probabilities of high-wage employment affect relative poverty with the (excess) burden of relative poverty falling disproportionately on L-caste households (viz. those with the employment-access disadvantage). Notice, however, that the mean durable expenditure level by *class* is similar in P versus J, so the L-caste households in the lower class face the same intensity of absolute poverty in Society P than in Society J.[5] However, compared with J, all L-caste households in P have lower durable expenditures, on average, than H-caste households, due to their higher presence in the lower class in P.

3.3 Experiment 3: Unjust Society "W" (differential wages or wage penalties)

Now consider results from Experiment 3 – an unjust society where L-caste households have lower wages (a wage penalty) than H-caste households – reported in Tables 1(i)-(ii) (columns 7-9). Notice that in Society W, the size of the classes (especially the lower class and hence, relative poverty) are very similar to those in the just society J. Moreover, the proportion of L-caste households in the lower and middle classes are close to that in the population (70%). However, we now find that the levels of total durable expenditure in each class (and in the entire sample) are much lower than in society J. For the lower class, durable expenditures are 0 throughout (since they cannot be negative). But recall that L-caste households earn a lower w_L here than in the just society. Hence, their levels of subsistence consumption ($2w_L^{L-caste}$) are lower in Society W than in Society J. Thus, the intensity of absolute poverty among L-caste households is higher in W than in J. Due to the caste-based wage differentials, in fact, L-caste households even in the middle and upper classes have lower consumption (denoted by durable expenditures) than their H-caste counterparts.

3.4 Lessons from Theoretical Experiments 1-3

The experimental results reported above provide some interesting insights on characterizing class dynamics in steady state, and their implications for anti-discrimination policy.

Notice, first, the important distinction between "type" of jobs ('high-wage' vs 'low-wage') and the wages that a job actually pays (w_H and w_L). For example, consider two households A and B who each earn the same wage w; however, w comes from a low-wage job for A and from a high-wage job for B. While their

[5]Recall that the theoretical equivalent of the absolute poverty line is the level of subsistence consumption C,assumed to be equal to the lowest income level possible in each period, viz. $2w_L$. Since the level of $w_L$is the same in J and P, the lower class face the same level of subsistence consumption in both these societies.

absolute well-being is the same, A and B probably belong to different classes (denoting relative well-being) because of the different "type" of jobs they hold within their communities. Thus, class membership in the (synthetic) data contain information on "relative" well-being in the context of the entire economy.

Second, having understood the difference between job "type" and job "wages", it is easy to see that class mobility occurs when the probabilities of obtaining the high-type jobs changes, or are different (as is the case for L-caste households). This is why we find a greater proportion (74.87%) of L-caste households in the lower class in Society P, compared with J (70%). Note that households in the lower class do not have a lower level of absolute well-being in P than in J. However, since there are excess L-caste households (relative to the population) in the lower class than in J, the headcount ratio of absolute poverty among L-caste households will be larger in P.[6]

In Society W, we notice that class sizes and caste compositions are very much the same as in J. However, the absolute well-being of households in each class is lower because the L-caste households in each class are earning less than their H-class counterparts. In other words, the intensity of (absolute) poverty – *how poor* are the poor? – is higher in the lower class in W than in J. Let us now turn to the specific finding that the upper class households have a much lower component (47.95%) of L-caste households than in J. How do lower wages affect class membership in the upper class (which depends on the probabilities of earning high-wage jobs)? Note that to be in the upper class, households must make draws of w_H both in the labour and marriage markets (and in 2 consecutive periods). If they receive a low value of w_H, L-caste households are unable to accumulate sufficient durables to signal high-status in the marriage market (given the same social standard for all households). Thus, the probability of obtaining w_H in the marriage market (necessary for being in the upper class) is diminished for the caste that gets paid lower absolute wages w_H!

In summary, access barriers that affect probabilities of employment into high-wage job "types" addresses the issue of relative poverty, whereas wage penalties for the same job "types" addresses the intensity of absolute well-being (thence, absolute poverty). Hence, policy that seeks to remove discriminatory effects must choose its target – (1) reducing the size of L-caste households in relative poverty (or increasing upward mobility from the lowest class), or (2) reducing the intensity of absolute poverty among L-caste households in the lowest class. If (1) is the target of policy, the probability parameters are clearly what need to be addressed, e.g. by job training programs[7]But if (2) – reducing the intensity of poverty – is the target of policy, then wage equalization – especially of the lower wages w_H – is key to success (e.g. minimum wage

[6]Since the absolute poverty line (in the model) is $C = 2w_L$, and there are more L-caste households in the lower class (earning $2w_L$ in each period) in society P than in J.

[7]Also by encouraging high education among L-caste households. Note that cost of high effective education E is a model parameter and we can run similar experiments with differential E across castes.

policy). Conversely – and importantly – the above analysis also directs us to the appropriate *metrics* to use to monitor and evaluate policy changes. For example, if our policy tool pertains to the probability parameters, then the appropriate metric to evaluate its efficacy would be the size of L-caste households in the lower class; and not the intensity of poverty among L-caste households in the lower class (which would be unaffected by this particular policy!). This point is sometimes missed in policy discourse around anti-discriminatory policy owing to the difficulty of disentangling policy impact on relative and absolute measures. Our framework offers a way to disentangle these effects, by conceptualizing the mechanisms that drive the persistence of relative and absolute poverty in a complex dynamic setting.

4 Calibration of Parameters

The experiments reported above were run based on arbitrary values of parameters, and intended to demonstrate the propagation mechanisms of two distinct channels of discrimination, P and W. Now, we calibrate model parameters using Indian survey data, to simulate a "pseudo" Indian urban economy.

4.1 Labour Market parameters $(w_H, w_L, p_{1,}p_2, p_3, p_4, E)$ and PLFS (2025) estimates

We calibrate labour market parameters $(w_H, w_L, p_{1,}p_2, p_3, p_4)$ using estimates from the Period Labour Force Survey (PLFS), 2025. We group observations in the survey by regions of India, viz. North, Central, East, North-East, South and West. Workers from households identifying as Scheduled Caste, Schedule Tribe or Other Backward Classes are defined to be of 'Low' caste (henceforth denoted L-caste); the rest are identified as 'High' caste (henceforth denoted H-caste) households. Also, workers are identified as having 'high education' if they have completed a graduate (or higher degree), else they are classified as having 'low' education. We focus on monthly earnings of workers that include pay reported from self-employment, salaried or casual work (per primary mode of employment reported).

Tables 4(a)-(d) provide a summary of statistics from PLFS 2025, viz. caste and population shares, and monthly earnings by caste and education level. Two broad patterns that emerge from this data are that (1) H-caste workers have higher earnings throughout the entire earnings' distribution[8], and (2) Workers with graduate education (of either caste) earn more than workers without. These patterns are consistent with the assumptions of the OLG model, hence we proceed to calibrate the lower and upper bounds of model parameters $(p^j_{1,}p^j_2, p^j_3, p^j_4)$ as follows $(j = H - caste, L - caste)$:

[8]This is violated in the North-East region, which, albeit has a small population share of 1.9%.

1. Suppose w_H lies between the 85th and 90th percentile of earnings in the national ("All India") distribution. w_L lies between the 15th and 20th percentile of earnings in the national distribution (Table 4(b), Panel A). These values of (w_H, w_L) establish the national standard for 'high' and 'low' types of jobs.

2. Consider the distribution of earnings of workers, by each group of caste and education level (Table 4(c)-(d), Panels B and C). We assume that high productivity workers in each caste-education group have earnings between the 55th and 60th percentiles of the relevant earnings distribution; and low productivity workers earn between the 40th and 45th percentiles of the relevant earnings distribution. Now, for each of these bounds, we derive the probability (conditional on being employed) at which the expected income (given w_H and w_L defined in (1) above) is equal to these particular values.[9] We multiply the conditional probabilities by the employment rate (Tables 4(c)-(d)) to obtain unconditional probabilities of receiving the high wage. This exercise yields bounds for the unconditional probabilities p_i^j – the probability that worker i of caste j earns w_H, where $i = 1, ..., 4$ represents the (unobserved) productivity level and education level of the worker[10] and j represents the worker's caste.

3. Let the (realized) high wage earned by caste$-j$ workers (w_H^j) lie between the 85th and 90th percentile of earnings in the caste-specific earnings distribution (Table 4(c), Panels B and C). In the experiments simulated below, we solve for transition matrices corresponding to parameters $(w_H^j, w_L^j, p_1^j, p_2^j, p_3^j, p_4^j)$ for caste j.[11]

4. We convert all values of monthly earnings (as reported from the PLFS survey) into generational earnings expressed in Rupees (lakhs), 2025.

Finally, as detailed in Data Appendix A, the lifetime cost E of acquiring a graduate education is obtained from estimates from National Sample Survey, 2017-18. The lower and higher bound is set to be the education cost estimates for Tier-2 and Tier-1 cities in India, respectively. We multiply these estimates by a factor of 1.39 to express as 2025 Rupees, estimated from inflation reports provided by RBI.

[9]Suppose X denotes the expected earnings (calibrated by various percentiles of the caste-education earnings distribution). Then the probability q at which X is the expected earnings is $q = \frac{(X - w_L)}{(w_H - w_L)}$,where $w_H$and $w_L$are calibrated from the national earnings-distribution (see also Hsieh et al., 2019, for a similar identification approach in the context of racial and gender barriers to occupational access).

[10]Recall, from equations $(4) - (8)$ in the OLG model, $i = 1$ for L-productivity/L-education, $i = 2$ for H-productivity/L-education, $i = 3$ for L-productivity/H-education and $i = 4$ for H-productivity/H-education workers.

[11]Note that the probabilities $(p_1^j, p_2^j, p_3^j, p_4^j)$ of earning the high wage are based on caste$-j$ workers' expected earnings relative to *national* opportunities (w_H, w_L), and not caste-specific opportunities (w_H^j, w_L^j). This distinction is deliberate: it allows the probability of securing a 'high'-type job – which reflects economy-wide access to opportunities – to be calibrated and interpreted separately from the wage level that is actually realized (w_H^j, w_L^j). The two are therefore analytically separable in the calibration, and permit the interpretation of 'classes' as relative positions in the national society. However, the distinction does imply that in the simulations, realized incomes (at probabilities $p_i^j$and earnings(w_H^j, w_L^j)) may deviate from the percentile values reported in the PLFS summary tables.

4.2 Marriage market parameters (β, σ^2)

The social standard β is calibrated using a basket of household durables (e.g. car, air conditioners, large refrigerators, washing machine, smart TV, laptops, smartphones etc.), the ownership of which is indicative of a "high" social-status household in contemporary urban India. The basket is ascertained (by quality and quantity) for Tier-1 and Tier-2 cities – for upper and lower bounds – and priced at 2025 prices. The procedure is detailed in Data Appendix A. Finally, we assume that the skepticism parameter σ^2 lies between 2.5 and 3.5.

4.3 Other parameters (q_L, δ)

We calibrate the probability q_L of drawing (unobserved) low productivity to be equal to the proportion of Tier-2 cities in urban India (0.85). The discount factor is set to be 0.5.

Having calibrated the regional upper and lower bounds for all parameters, we now draw 100 communities – with L-caste and H-caste sub-communities – within each region with parameters drawn from a uniform distribution in the appropriate calibrated range. We solve for the transition matrix and steady state durables' distributions in each of these sub-communities. Since, each sub-community is characterized by region r (=1,...,6) and caste j (= L-caste, H-caste)), we then draw households from each sub-community in proportion to the relative population size of caste j in region r and region j in the country as a whole. The pooled observations from each sub-community yield the simulated data set representing contemporary urban Indian "pseudo" society.

5 Simulation Results

As in Sections 3.1-3.3, we report simulated findings from experiments J, P, W below. In addition, our calibrations permit us to simulate a model of contemporary urban "pseudo" India, which we denote as Society A. Our simulations show that contemporary urban "pseudo" Indian society has no subsistence (lower) class, but an U class that is 11.3 percentage points smaller than it would be in the absence of caste discrimination. In experiments J, P and W below, we parse out the mechanisms that explain this aggregate finding.

5.1 Contemporary Indian Urban "Pseudo" Society – Society A

The Contemporary Indian Urban "Pseudo" Society (call it Society A) comprises households that face opportunities specific to their caste in their community/ region. Households are drawn from communities in proportion to regional and caste representation in their communities, and pooled to comprise our simulated data set. The simulated data are summarized in Tables 2-3.

The foremost finding for Society A is the disappearance of the lowest class L – or the subsistence class – in the class profile. There is a small percentage (4.88%) of (potential) single-period subsistence earners – in terms of the Low-High Oscillators (class LHO) – but there are no households who are stuck in the L class, viz. who receive subsistence-level incomes over consecutive generations. Society A is characterized by an M class of size 29.12% and an upper (U) class of size 21.48%. In addition, there emerges a "new middle" class – the "upper middle (UM) class" of size 44.52% – of households that oscillate between the middle and high incomes across generations. Mean durable consumption is lowest for class LHO and increases across classes M, UM and U.

Two types of relative deprivation emerge in the urban "pseudo" Indian society A. The first type of relative deprivation occurs through a high volatility of intergenerational income – the Low-High Oscillator (LHO) class. The LHO-class (albeit, only 4% in size) is characterized by high intergenerational mobility – the generations oscillate between low and high income – but the households therein are "poorest" by durable consumption in any period (suggesting that these households come from communities with very low w_H as well as w_L). LHO also has a 0% rate of high-education acquisition (Table 3, Panel B) in sharp contrast with an overall economy-wide education attainment rate of 90.03% .

The second type of relative deprivation – the M-class – occurs through low intergenerational mobility, as these households remain "stuck" at mid-level incomes across generations. That the mean durables of the M-class are the second-lowest of all the levels across the classes indicate that M-class households represent the class that is "stuck at the bottom" in this society in the long run. Therefore, in the remaining analysis, we will use the "Stuck-at-the-Bottom" index to reference the "Stuck-in-M" index.

Moreover, in Society A, we see that L-caste households are more likely than H-caste households to face either type of relative deprivation. The LHO-class struggling with high volatility comprise L-caste households alone (100%). The M-class that struggles with intergenerational immobility is 74.37% L-caste households, a number higher than the overall proportion of L-caste households (64.25%) in the entire population.

The higher burden of relative deprivation on L-caste households is also supported by the Stuck-in-M index, which measures the proportion of period$-(t-1)$ households in M-class who are also in M-class in period$-t$, in steady state. In Society A, the Stuck-in-M index for Society A is 61.7% for L-caste households, compared with 47.5% of H-caste households.

Recall that the parameters used to simulate income-generation in Society A vary by caste and region.

Hence, communities of caste j and community c (identified by geographical region) are generated using parameters $(w_H^{j,c}, w_L^{j,c}, p_{1,}^{j,c} p_2^{j,c}, p_3^{j,c}, p_4^{j,c}, \beta^c, \sigma^{2,c}, q_L, \delta, E^c)$ calibrated as described above. The mean ex-ante expected labour-market earnings (reported in Table 2, Panel B) provide a measure of the difference in opportunities inherent in the differential parameters governing income-generation in the different caste-communities within regions.[12] Clearly, in Society A, the mean ex-ante labour market earnings are higher for H-caste than for L-caste households (98.3 vs 59.14), since the former face better opportunities represented by higher probability parameters and higher wages.

Next, we simulate a just society, where all castes have equal opportunities (as expressed in parameters and ex-ante expected labour-market outcomes).

5.2 Just Society J

Just Society J is characterized by equal opportunities for all households, regardless of their caste. To simulate a Just Society J, we assume that households of all castes experience opportunities represented by the parameters of the H-caste households in Society A.

The simulation results (Tables 2-3) show, once again, the absence of a persistent subsistence class, L, in Society J. However, the LHO-class of low-high oscillators (that appear in Society A) disappear completely as well. Instead we have 3 classes in Society J – an M-class (of size 20.34%), an UM-class (46.86%) and an U-class (32.80%). While higher classes enjoy higher durable consumption, there is no distinction in class-specific durable levels by caste. Likewise, the caste composition of each class is almost the same as in the overall population (~ 64%), as predicted by Theoretical Experiment J (Section 3.1); and the Stuck-in-M (or, Stuck-at-the-Bottom) index is the same across the castes. Notice that compared with just society J, contemporary Indian "pseudo" urban society A has an U class that is 11 percentage points smaller (21.48% vs 32.80%). This difference is attributable to the altered opportunities in A – their direct effects in the labour market and indirect effects through the marriage markets – generated by caste-based norms.

Society J now constitutes the benchmark against which we will compare the impact of differential access to opportunities, on class formation and relative deprivation by caste.

5.3 Unjust Society P (differential probabilities or access barriers)

In Unjust Society P, we assume that L-caste households receive the same low and high wages as H-caste households. However, L-caste households face lower probabilities of earning the high wage than H-caste

[12] See footnote 4.

households,. In other words, H-caste households in any community c, face opportunities encapsulated in $(w_H^{Hcaste,c}, w_L^{Hcaste,c}, p_{1,}^{Hcaste,c} p_2^{Hcaste,c}, p_3^{Hcaste,c}, p_4^{Hcaste,c}, \beta^c, \sigma^{2,c}, q_L, \delta, E^c)$, whereas L-caste households access opportunities $(w_H^{Hcaste,c}, w_L^{Hcaste,c}, p_{1,}^{Lcaste,c} p_2^{Lcaste,c}, p_3^{Lcaste,c}, p_4^{Lcaste,c}, \beta^c, \sigma^{2,c}, q_L, \delta, E^c)$. We could interpret P as a society where there are barriers to access for L-caste households to the high 'type' of jobs, but when hired to such jobs households of all castes received the same wage. The ex-ante expected labour market earnings reported for P in Table 2 (Panel B) provide a numerical measure of the difference of opportunities represented by differential probabilities, viz. 80.69 for L-caste and 98.83 for H-caste households, respectively.

The simulation results (Tables 2-3) show the existence of 4 classes – LHO (a very small share of 0.06%), M (31.64%), UM (46.27%) and U (32.80%). Compared with Society J, we see an appearance of the volatile LHO class, an expansion of M-class and a shrinkage of U-class. The mean durables by class are, however, identical in J and P. These patterns are expected since the income levels (defined by high and low wages) are the same in J and P (hence same mean durables). However, access barriers that affect 65% of the total population (the L-caste households) leads to an expansion of the M-class (the stuck-at-the-bottom class). We also see that L-caste households comprise a much higher proportion of the M-class (76.41%) and a lower proportion of the U-class (47.91%) in P than in J. Consistent with all of these patterns, the Stuck-in-M index is also higher for L-caste households in P than in J. Interestingly, we see that the Stuck-in-M index is similar in Societies A and P, suggesting that the effect of intergenerational immobility of L-caste households in A works largely through the P channel of access barriers.

5.4 Unjust Society W (differential wages or wage penalties)

In Society W, we assume that L-caste households face the same probabilities of earning the high-wage as H-caste households. However, L-caste households face lower wages (both high and low) than H-caste households. In other words, H-caste households in any community c, face opportunities (as before) encapsulated in $(w_H^{Hcaste,c}, w_L^{Hcaste,c}, p_{1,}^{Hcaste,c} p_2^{Hcaste,c}, p_3^{Hcaste,c}, p_4^{Hcaste,c}, \beta^c, \sigma^{2,c}, q_L, \delta, E^c)$, whereas L-caste households access opportunities $(w_H^{Lcaste,c}, w_L^{Lcaste,c}, p_{1,}^{Hcaste,c} p_2^{Hcaste,c}, p_3^{Hcaste,c}, p_4^{Haste,c}, \beta^c, \sigma^{2,c}, q_L, \delta, E^c)$. In Society W, therefore, all castes have equal access to high 'type' jobs. However, L-caste households receive a lower wage than H-caste households, for the same jobs. The ex-ante expected labour market earnings reported for W in Table 2 (Panel B) provide a numerical measure of the difference of opportunities represented by differential probabilities, viz. 71.39 for L-caste and 98.83 for H-caste households, respectively. Thus, L-caste households face a lower "level" of opportunities in W than in P – a feature of the contemporary labour market opportunities' differential (PLFS, 2025).[13]

The simulation results (Tables 2-3) reveal 4 classes – LHO (a small share of 0.32%; larger than in P), M (20.37%), UM (46.55%) and U (32.76%). Note that the sizes of M, UM and U are almost identical to that in J, as is the proportion of L-caste households in each class, and the Stuck-in-M index for each caste. However, the mean durables of each class is lower than in J. In the M class, specifically, the gap between L-caste and H-caste mean durables exhibits a gap of 59.4 (120.26 vs. 179.66; Table 3, Panel A, Society W). In J and P, this gap is essentially zero (180.33 vs 179.66 in J; 180.46 vs 179.66 in P). This illustrates that W generates a sizable within-class caste gap – mirroring A – while P leaves within-class living standards completely unchanged. The effect of discrimination on poverty intensity in A works largely through the W channel of wage penalties.

5.5 Societies J/P/W versus Society A

The following insights emerge from comparing the simulated findings for Societies J/P/W with each other and with those of Society A (the contemporary, urban "pseudo" Indian society).

First, the negative impact on overall opportunities' – the caste-based "opportunities" gap – imposed by wage penalties (W) is larger than that of access barriers (P), in contemporary urban India. We see this in the ex-ante expected labour market earnings in each society (Table 2, Panel B). In J, both castes have ex-ante earnings of 98.83. This value remains the same for H-caste households in all societies (by construction). But for L-caste societies, ex-ante earnings is lowest in A (59.14), followed by (in ascending order) W (71.39) and then P (80.69).

Second, the calibrated urban "pseudo Indian economy (Society A) behaves like Society P for intergenerational immobility and like Society W for within-class consumption intensity, confirming that both channels are simultaneously active in the data. Notice that societies A and P – and, separately, Societies J and W – have similar values for (1) class sizes, (2) caste-composition of classes, and (3) values of Stuck-in-M Indexes for L-caste and H-caste households. Note, also, the similarity of class-specific mean durables across J and P, and the lower value of durables in A and W versus J/P. In particular, notice the large caste gap in mean durables in W and A (179.66 for H-caste and ~120 for L-caste), a gap that is almost zero for J and P (179.66 for H-caste and ~180 for L-caste). Hence, as before, we can conclude that channel P impacts relative measures and class composition (viz. how large the relatively poor? how are L-caste/H-caste households distributed

[13] Note that in the theoretical experiments, we assumed that ex-ante expected earnings were the same for P and W. The goal of that exercise was to isolate and compare (in theory) the mechanisms operating in P and W. The PLFS-based calibrations permit us to measure the difference in opportunities inherent in earnings' distribution. This has important implications for optimal policy selection (Section 5.6).

among the relatively poor?) whereas channel W impacts the intensity of relative poverty (viz. how poor are the relatively poor? How poor are the disadvantaged caste?).

Third, caste hierarchy is reproduced and sustained from both ends of the class distribution, i.e. not only are L-caste households more likely to be stuck at the bottom, H-caste households are more likely to be stuck at the top. We see this from the fact that the Stuck-in-U index (Table 2, Panel B) is consistently higher for H-Caste (0.573) than for the L-caste across in A, P and W (0.375 in A, 0.411 in P, 0.584 in W). Hence, policy that seeks to minimize caste disadvantage in households' relative positions must ensure not only that L-caste households successfully migrate to U class, but also that they remain there in the long-run (at the same rate as H-caste households).

Finally, notice that attainment of graduate education is near-universal in J/P/W (Table 3, Panel B), a finding that is not mirrored in PLFS data (27.4% for all of India, Table 4(a)). It is likely that the cost of education E used in our simulations – based on 2017-18 levels adjusted by inflation – may be too low an estimate of graduate education costs in 2025. Indeed, upon running the simulations with higher (imposed) values of E, we find that (1) graduation rates fall across both castes, albeit more for L-caste households, and (2) the size of the LHO class (that have a 0% rate of graduate-education attainment) increases. However, we continue to see the absence of a subsistence-class in the society even for values of E as high as 3 times those used in the reported simulations. These combined findings characterize a society transitioning from a subsistence economy to one with better overall – albeit precarious – income-generation opportunities. LHO households are households for whom education does not pay at current (calibrated) E values – higher values of E push more households toward non-investment in education, expanding the size of the LHO class. This makes the LHO class a cautionary illustration of what high education costs could produce more broadly — not a lower (subsistence) class in the traditional sense, but a volatile, non-investing (in education) class whose poverty reproduces through income swings rather than persistent low income.

To the extent that the size of the LHO class is a measure of the "precariousness" (or volatility) of existing opportunities, it appears that reducing education costs E (conditional on other labour-market parameters) may be an effective means of reducing such volatility in the long run. However, the true effectiveness of $E-$reduction in mitigating inter-generational precariousness depends on the values of all other model parameters – the "composite opportunities" – which jointly determine the process of class formation. Our ongoing research seeks to estimate more accurate/current education-cost parameters for simulations, and also to establish the precise (theoretical) range of precariousness-inducing parameter levels that jointly generate volatility in intergenerational income-generation.

5.6 Policy Implications

The simulation results demonstrate that caste-based discrimination could take two economically distinct forms — differential probabilities (barriers to accessing high-wage employment) and differential wages (caste-based wage penalties for the same 'type' of employment). Each form of discrimination perpetuates intergenerational inequality through fundamentally different channels, requiring fundamentally different policy responses.

As demonstrated in Sections 5.1-5.5, the calibrated urban "pseudo Indian economy A behaves like Society P for intergenerational immobility and like Society W for within-class consumption intensity. Hence, both channels are simultaneously active in the data and both channels require policy attention. However, the calibrated results also indicate that wage discrimination is currently the more severe form of discrimination in urban India as measured by the caste-based opportunities' gap (ex-ante labour market earnings, Table 2, Panel B). This finding indicates that wage-based intervention is more urgent under present conditions, even though both channels of interventions are ultimately necessary to eradicate all forms systemic discrimination evident in the data.

The most "effective" policy in any circumstance would, therefore, involve an iteration of the following steps: (1) an identification of the more stringent form of discrimination in society – to determine urgency and policy timing, (2) a selection of a *specific* policy tool that targets that area of discrimination, and (3) an evaluation of policy outcomes by focusing on the *relevant* metrics impacted by that selected tool.

Policy tools that reduce the gap in employment access (such as reservations for employment or at educational institutions) can be expected to lower relative poverty, reduce the proportion of L-caste households among the relatively poor, and also reduce intergenerational immobility (stuck-at-the-bottom index). Policy tools that narrow the caste-based wage penalty (e.g. minimum wages) for the same jobs will achieve a reduction in poverty intensity among the relatively poor. The specific tool selected should be guided by the more stringent source of discrimination that requires urgent attention.

Conversely, from the viewpoint of policy evaluation, it is critical that the efficacy of the selected policy tool be measured using the appropriate metrics that will be impacted by that tool. For probability-based policies, the appropriate metrics to track include the Stuck-at-the-Bottom (here, the Stuck-in-M) index and the L-caste shared of each class. For wage-based policies, the key metric is the caste-based gap in durable expenditures within each class. Finally, as with all policy that targets dynamic, long term impact, an iteration of steps (1) - (3) is recommended for practical implementation, in order to track changing opportunities and adapt policy tools to evolving socio-economic conditions.

6 Conclusion

We demonstrate – using a novel methodology of theoretical experiments based on an OLG model – the distinction between the "access barriers" (P) channel and the "wage penalty" (W) channel of discrimination on class formation, relative poverty and intergenerational mobility. Specifically, we define a "Stuck-at-the-Bottom" Index that measures the intergenerational immobility of households in relative poverty, by caste. Calibrating our OLG model parameters from PLFS 2025, we find evidence that both P and W channels are active in contemporary urban India, with the caste-based gap in opportunities being larger for channel W. Hence, wage-based policy measures that affect the W channel may be more urgently required for alleviation of caste-based discrimination in India.

Our analysis offers insights on how to select, formulate and evaluate policy tools that best impact the targeted mode of discrimination at any point of time. In particular, it generates the important insight that caste hierarchy reproduces itself simultaneously from the lowest and the uppermost ends of the class distribution: L-caste households are more likely to be stuck at the bottom and H-caste households are more likely to be stuck at the top. Any policy that successfully moves L-caste households into the upper class must also address the structural disadvantage in maintaining that position — otherwise the intergenerational gain dissipates in the long run.

We conclude by noting that while the theoretical assumptions and calibration in this paper are specific to India, the conceptual framework applies to any economy characterized by categorical inequality across groups and a labour market with both access and wage discrimination. The methodology presented here is amenable to the addition of complexity in the theoretical framework, as well as to more granular, sophisticated approaches for calibration. In this regard, the paper's central methodological contribution is portable across contexts.

References

[1] Asher, S., Novosad, P., & Rafkin, C. (2024). Intergenerational mobility in India: New measures and estimates across time and social groups. *American Economic Journal: Applied Economics*, 16(2), 66–98.

[2] Banerjee, A., Duflo, E., Ghatak, M., & Lafortune, J. (2013). Marry for what? Caste and mate selection in modern India. *American Economic Journal: Microeconomics*, 5(2), 33–72.

[3] Banerjee, B., & Knight, J. B. (1985). Caste discrimination in the Indian urban labour market. *Journal of Development Economics*, 17(3), 277–307.

[4] Becker, G. S. (1971). *The Economics of Discrimination* (2nd ed.). University of Chicago Press.

[5] Corak, M. (2013). Income inequality, equality of opportunity, and intergenerational mobility. *Journal of Economic Perspectives*, 27(3), 79–102.

[6] Darity, W. (2005). Stratification economics: The role of intergroup inequality. *Journal of Economics and Finance*, 29(2), 144–153.

[7] Deshpande, A. (2011). *The Grammar of Caste: Economic Discrimination in Contemporary India*. Oxford University Press.

[8] Galor, O., & Zeira, J. (1993). Income distribution and macroeconomics. *Review of Economic Studies*, 60(1), 35–52.

[9] Hnatkovska, V., Lahiri, A., & Paul, S. B. (2012). Castes and labor mobility. *American Economic Journal: Applied Economics*, 4(2), 274–307.

[10] Hsieh, C.-T., Hurst, E., Jones, C. I., & Klenow, P. J. (2019). The allocation of talent and U.S. economic growth. *Econometrica*, 87(5), 1439–1474.

[11] Loury, G. C. (1981). Intergenerational transfers and the distribution of earnings. *Econometrica*, 49(4), 843–867.

[12] Madheswaran, S., & Attewell, P. (2007). Caste discrimination in the Indian urban labour market: Evidence from the National Sample Survey. *Economic and Political Weekly*, 42(41), 4146–4153.

[13] Maitra, S. (2024). On the theory and measurement of relative poverty using durable ownership data. *Journal of Economic Behavior and Organization*, 225, 153-169. Available at: `https://www.sudeshnamaitra.com/Maitra_relative_poverty_Sep2023.pdf`

[14] Phelps, E. S. (1972). The statistical theory of racism and sexism. *American Economic Review*, 62(4), 659–661.

[15] Shorrocks, A. F. (1978). The measurement of mobility. *Econometrica*, 46(5), 1013–1024.

[16] Thorat, S., & Attewell, P. (2007). The legacy of social exclusion: A correspondence study of job discrimination in India. *Economic and Political Weekly*, 42(41), 4141–4145.

[17] Tsokos, C. P. (1972). Probability Distributions: An Introduction to Probability Theory with Applications. Duxbury Press, Belmont California.

[18] Vaid, D. (2014). Caste in contemporary India: Flexibility and persistence. *Annual Review of Sociology*, 40(1), 391–410.

Figure 1: The Signal Function

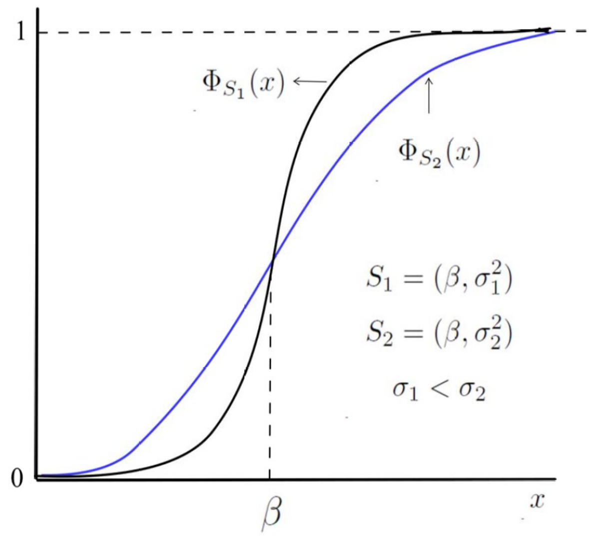

Figure 1: The Signal Function, normal c.d.f. $\Phi_S(x)$

Assumption: $\Phi_S(X)$ (where S = (β, σ)) represents the probability that the marriageable generation in a household with durables β will find a partner that earns the high wage, when skepticism around the belief β is given by σ.

Figure 2(a): Steady state distribution of household durables, derived for parameters as in Example 1 (Section 2.3)

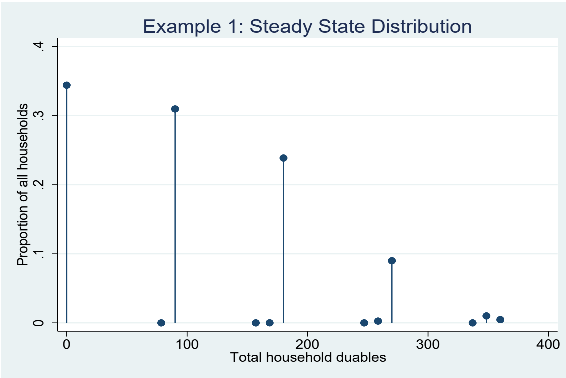

Figure 2(b): Histograms from samples independently drawn from the theoretical distribution in Figure 2(a)

Sample 1

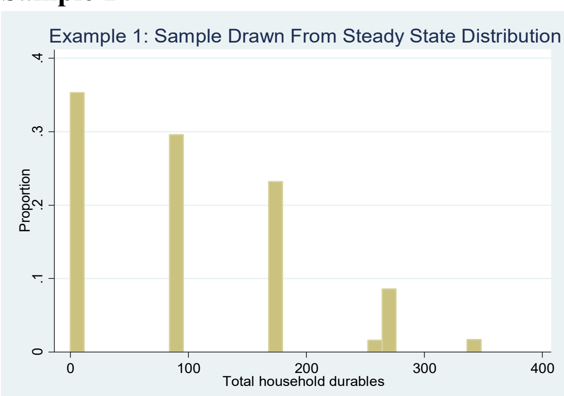

Sample 2

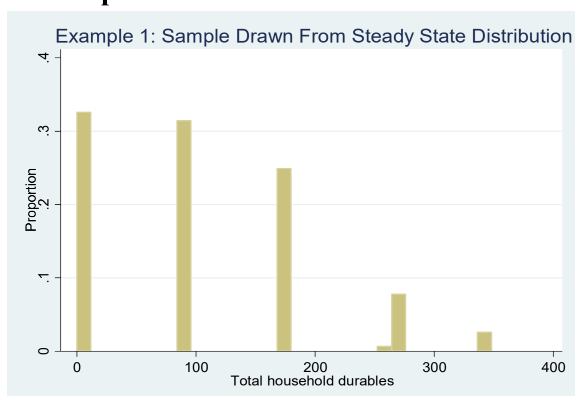

Figure 2(c): Total durable expenditure in pooled sample of 1000 communities (N=1,000,000)

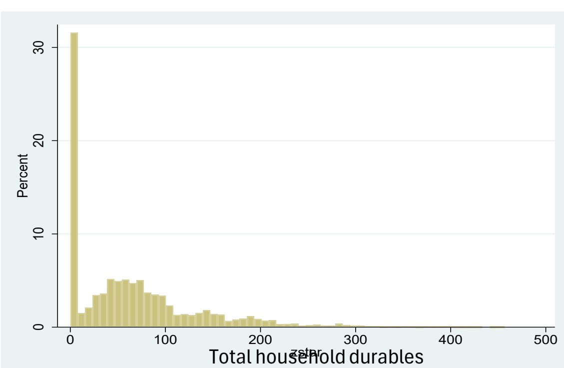

Example 2. *(Set A)* $w_H \in (80, 120), w_L \in (5, 20), p_1 \in (0, 0.5), p_2 \in (p_1, 0.5), p_3 \in (0.5, 1), p_4 \in (p_3, 1), \beta \in (100, 500), \sigma \in (100, 500), q_L = 0.5, \delta = 0.5, E \in (10, 100), C = 2w_L$.

Simulation Results

Table 1(i), Panel A: Class Characteristics in Theoretical Experiments (Section 3.1-3.4)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Just Society J			**Unjust Society P**			**Unjust Society W**		
Class	Class Size (%)	Mean Durables	% L-Caste	Class Size (%)	Mean Durables	% L-Caste	Class Size (%)	Mean Durables	% L-Caste
Lower	46.72	0.00	70.05	55.67	0.00	74.87	47.31	0.00	70.29
Lower-Middle	39.67	87.02	69.95	35.59	87.14	66.51	39.93	70.27	70.37
Low/High Oscillators	7.74	174.32	70.45	5.27	174.19	56.58	7.80	140.64	71.05
Middle	5.70	171.22	69.30	3.35	173.87	47.84	4.81	142.57	62.59
Upper-Middle	0.16	266.50	69.18	0.11	267.42	55.33	0.14	219.87	63.37
Upper	0.01	355.98	70.41	0.01	357.34	45.28	0.01	306.56	47.95
All Classes	100	58.24	70	100	46.34	70	100	46.22	70

Table 1(i), Panel B: Stuck -at-the-Bottom Indices in Theoretical Experiments (Section 3.1-3.4)

	Just Society J			**Unjust Society P**			**Unjust Society W**		
	All	H-caste	L-caste	All	H-caste	L-caste	All	H-caste	L-caste
Stuck in L	0.654	0.658	0.652	0.723	0.658	0.750	0.660	0.658	0.661
Stuck in M	0.232	0.232	0.231	0.172	0.232	0.123	0.202	0.240	0.179
Stuck in U	0.003	0.002	0.003	0.002	0.002	0.002	0.002	0.003	0.001

[1] The exante labour market earnings reflect the earnings expected *before* productivity levels and education levels are drawn/chosen -- a function of community-specific parameter values. Reported values are the mean exante expected earnings for all communities in the simulated sample.

Simulation Results

Table 1(ii): Mean Durables, by Caste, in Theoretical Experiments (Section 3.1-3.4)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Just Society J			**Unjust Society P**			**Unjust Society W**		
Class	All	H-Caste	L-Caste	All	H-Caste	L-Caste	All	H-Caste	L-Caste
Lower	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0
Lower-Middle	87.02	86.98	87.04	87.14	86.98	87.22	70.27	87.03	63.21
Low/High Oscillators	174.32	173.73	174.56	174.19	173.73	174.53	140.64	175.37	126.50
Middle	171.22	172.25	170.76	173.87	172.25	175.64	142.57	170.17	126.07
Upper-Middle	266.50	266.11	266.67	267.42	266.11	268.47	219.87	267.07	192.58
Upper	355.98	366.34	351.63	357.34	366.34	346.46	306.56	355.00	253.96
All Classes	58.24	58.33	58.20	46.34	58.33	41.20	46.22	58.24	41.07

Calibrated Simulation Results

Table 2, Panel A: Class Characteristics in "Psuedo Urban" Society, and Societies J, P, W (Section 5)

Class	**"Pseudo-Urban" Society**			**Just Society J**			**Unjust Society P**			**Unjust Society W**		
	Class Size (%)	Mean Durables	% L-Caste	Class Size (%)	Mean Durables	% L-Caste	Class Size (%)	Mean Durables	% L-Caste	Class Size (%)	Mean Durables	% L-Caste
Lower	-	-	-	-	-	-	-	-	-	-	-	-
Low/High Oscillators	4.88	118.66	100	-	-	-	0.060	181.98	100	0.32	109.42	100
Middle	29.12	136.51	74.37	20.34	180.08	63.27	31.64	180.27	76.41	20.37	142.03	63.35
Upper-Middle	44.52	216.51	62.26	46.86	273.22	64.11	46.27	273.15	63.68	46.55	214.47	63.90
Upper	21.48	308.22	46.53	32.80	365.96	64.95	22.04	365.79	47.91	32.76	285.84	64.95
All Classes	100	208.14	64.25	100	284.69	64.21	100	264.13	64.25	100	222.75	64.25

Table 2, Panel B: Stuck -at-the-Bottom Indices in "Pseudo Urban" Society, J, P, W

	"Pseudo-Urban" Society			**Just Society J**			**Unjust Society P**			**Unjust Society W**		
	All	H-caste	L-caste	All	H-caste	L-caste	All	H-caste	L-caste	All	H-caste	L-caste
Stuck in L	0	-	0	-	-	-	0	-	0	0	-	0
Stuck in M	0.573	0.475	0.617	0.470	0.475	0.466	0.583	0.475	0.627	0.471	0.475	0.468
Stuck in U	0.460	0.573	0.375	0.579	0.573	0.582	0.482	0.573	0.411	0.580	0.573	0.584
Mean Exante Lab. Mkt. Earnings[1]	-	98.83	59.14	-	98.83	98.83	-	98.83	80.69	-	98.83	71.39

[1] The exante labour market earnings reflect the earnings expected *before* productivity levels and education levels are drawn/chosen -- a function of community-specific parameter values. Reported values are the mean exante expected earnings for all communities in the simulated sample.

Calibrated Simulation Results

Table 3, Panel A: Mean Durables, by Caste, in "Pseudo Urban Society", J, P & W, (Section 5)

Class	"Pseudo-Urban" Society			Just Society J			Unjust Society P			Unjust Society W		
	All	H-Caste	L-Caste	All	H-Caste	L-Caste	All	H-Caste	L-Caste	All	H-Caste	L-Caste
Lower	-	-	-	-	-	-	-	-	-	-	-	-
Low/High Oscillators	118.66	-	118.66	-	-	-	181.98	-	181.98	109.42	-	109.42
Middle	136.51	179.66	121.64	180.08	179.66	180.33	180.27	179.66	180.46	142.03	179.66	120.26
Upper-Middle	216.51	272.39	182.64	273.22	272.39	273.68	273.15	272.39	273.59	214.47	272.39	181.75
Upper	308.22	364.58	243.46	365.96	364.58	366.71	365.79	364.58	367.10	285.84	364.58	243.34
All Classes	208.14	282.64	166.68	284.69	282.64	285.83	264.13	282.64	253.84	222.75	282.64	189.43

Table 3, Panel B: Percentage of Households with 'High' Education, by Class and Caste, "Psuedo Urban" Society, J, P, W

Class	"Pseudo-Urban" Society			Just Society J			Unjust Society P			Unjust Society W		
	All (%)	H-caste[1] (%)	L-caste[2] (%)	All (%)	H-caste[1] (%)	L-caste[2] (%)	All (%)	H-caste[1] (%)	L-caste[2] (%)	All (%)	H-caste[1] (%)	L-caste[2] (%)
Lower	-	-	-	-	-	-	-	-	-	-	-	-
Low/High Oscillators	0	-	0	-	-	-	0	-	0	0	-	0
Middle	90.76	100	87.58	100	100	100	99.81	100	99.75	100	100	100
Upper-Middle	95.04	100	92.04	100	100	100	99.90	100	99.84	99.74	100	99.60
Upper	99.11	100	98.08	100	100	100	99.97	100	99.94	99.99	100	99.99
All Classes	90.03	100	84.48	100	100	100	99.83	100	99.73	99.55	100.00	99.3

[1] Figures in this column represent the % of households of H-caste with H education, within each class.

[2] Figures in this column represent the % of households of L-caste with H education, within each class.

Table 4a: PLFS 2025 — Labour Force Characteristics, Urban India, All Workers (18–80 years)

Region	Prop. L-caste (earner sample)	Prop. Graduate (earner sample)	Regional Population Share (HH-wtd)	L-caste Share in Region (HH-wtd)	N (earners)
Panel A: All Castes					
All India	**67.5%**	**27.4%**	**100.0%**	**64.6%**	**178,594**
North	54.3%	28.1%	15.2%	56.4%	32,320
Central	72.8%	28.0%	16.1%	70.2%	40,303
East	61.0%	26.6%	13.9%	50.9%	25,026
North-East	74.5%	27.5%	1.9%	64.3%	14,856
South	83.4%	27.6%	32.5%	82.2%	40,103
West	53.3%	26.0%	20.4%	47.7%	25,986
Panel B: Low Caste					
All India	**100.0%**	**23.3%**	**100.0%**	**64.6%**	**120,508**
North	100.0%	20.9%	15.2%	56.4%	17,539
Central	100.0%	22.4%	16.1%	70.2%	29,329
East	100.0%	21.7%	13.9%	50.9%	15,274
North-East	100.0%	27.1%	1.9%	64.3%	11,074
South	100.0%	26.1%	32.5%	82.2%	33,441
West	100.0%	21.4%	20.4%	47.7%	13,851
Panel C: High Caste					
All India	**0.0%**	**35.4%**	**100.0%**	**64.6%**	**58,086**
North	0.0%	36.3%	15.2%	56.4%	14,781
Central	0.0%	42.0%	16.1%	70.2%	10,974
East	0.0%	34.2%	13.9%	50.9%	9,752
North-East	0.0%	28.6%	1.9%	64.3%	3,782
South	0.0%	35.0%	32.5%	82.2%	6,662
West	0.0%	30.9%	20.4%	47.7%	12,135

Sample includes urban workers of age 18-80 years with non-missing values of monthly earnings (from self-employment, salaried or casual work). L-caste workers come from SC/ST/OBC households. H-caste workers come from households with "social group" = "General". High education refers to workers with a completed graduate or higher degree. All statistics are unweighted, unless otherwise noted.

Table 4b: PLFS 2025 — Monthly Earnings, Urban India, All Workers (18–80 years)

Region	P15	P20	P40	P45	P50	P55	P60	P85	P90	Mean (Rs.)
Panel A: All Castes										
All India	**5,214**	**7,170**	**12,322**	**14,339**	**15,000**	**16,000**	**18,000**	**35,000**	**45,000**	**20,912**
North	5,500	8,000	14,000	15,000	15,500	18,000	20,000	40,000	50,000	22,873
Central	3,500	5,500	10,000	12,000	12,167	14,000	15,000	30,000	38,000	17,117
East	4,500	6,500	12,000	12,167	14,000	15,000	16,000	35,000	41,000	18,833
North-East	5,500	7,320	13,036	15,000	15,463	17,500	19,650	40,000	45,560	21,651
South	7,000	8,690	15,000	15,000	17,000	18,000	20,000	36,000	45,000	22,446
West	6,500	8,500	15,000	15,000	16,000	18,000	20,000	40,000	50,000	23,575
Panel B: Low Caste										
All India	**5,000**	**6,952**	**12,000**	**13,000**	**15,000**	**15,000**	**16,000**	**30,000**	**40,000**	**18,628**
North	5,000	7,000	12,000	13,688	15,000	15,000	16,000	30,000	40,000	18,789
Central	3,000	5,000	9,777	10,429	12,000	12,167	14,000	25,000	30,500	15,016
East	4,000	6,000	10,863	12,000	13,000	14,500	15,000	28,010	35,355	16,695
North-East	5,300	7,000	13,036	15,000	15,600	18,000	20,000	40,000	45,016	21,584
South	6,952	8,500	14,500	15,000	16,000	18,000	19,500	35,000	42,000	21,198
West	5,866	7,800	12,000	14,000	15,000	15,229	17,800	32,000	40,000	19,637
Panel C: High Caste										
All India	**6,000**	**8,500**	**15,000**	**16,000**	**18,000**	**20,000**	**22,000**	**45,000**	**55,000**	**25,651**
North	6,000	9,000	15,000	18,000	20,000	22,000	25,000	50,000	60,000	27,718
Central	5,000	7,500	13,500	15,000	15,500	18,000	20,000	40,000	50,000	22,729
East	5,000	7,500	12,167	14,000	15,000	16,949	18,250	40,000	50,000	22,180
North-East	6,000	8,000	13,036	15,000	15,208	16,635	18,250	39,500	47,000	21,844
South	7,821	10,000	16,126	18,000	20,000	22,000	25,000	50,000	60,000	28,709
West	8,000	10,000	16,000	18,000	20,000	21,500	25,000	45,000	57,000	28,070

Pi denotes the ith percentile of the distribution of monthly earnings from self-employed, salaried or casual work. Earnings in Rs. (2024-25). P15–P20 used as wL bounds; P40–P45 used as low-prod expected earnings; P55–P60 used as high-prod expected earnings; P85–P90 used as wH bounds. All statistics unweighted unless otherwise noted.

Table 4c: PLFS 2025 — Monthly Earnings, Urban India, Graduate Workers (Graduate Degree or Higher))

Region	P15	P20	P40	P45	P50	P55	P60	P85	P90	Mean (Rs.)	Unemp. Rate within Labour Force	N (earners)
Panel A: All Castes												
All India	**9,125**	**12,000**	**20,000**	**23,000**	**25,000**	**30,000**	**32,000**	**58,000**	**68,000**	**33,112**	**13.1%**	**54,404**
North	10,000	12,000	22,000	25,000	30,000	34,000	36,808	65,000	75,000	36,739	18.4%	9,748
Central	6,000	8,000	15,000	18,000	20,000	22,000	25,000	50,000	58,500	26,715	8.8%	12,364
East	8,000	10,000	18,400	20,000	25,000	26,500	30,000	52,000	60,040	30,234	16.0%	7,610
North-East	10,000	12,000	21,000	25,000	29,000	32,000	35,000	55,000	64,640	32,553	17.3%	4,750
South	12,000	15,000	23,000	25,000	28,000	30,000	35,000	60,000	70,000	35,689	14.4%	12,192
West	12,000	15,000	25,000	25,000	30,000	31,081	35,000	60,000	75,000	37,878	5.1%	7,740
Panel B: Low Caste												
All India	**8,690**	**10,646**	**18,400**	**20,000**	**24,000**	**25,000**	**30,000**	**52,000**	**62,000**	**30,216**	**15.0%**	**30,777**
North	8,000	10,000	18,000	20,000	22,400	25,000	30,000	60,000	70,000	31,888	23.5%	3,766
Central	5,000	7,500	15,000	15,000	18,000	20,000	22,000	45,000	55,000	24,235	9.7%	7,015
East	7,500	10,000	17,000	18,402	20,500	25,000	27,500	46,330	55,000	27,059	20.1%	3,710
North-East	10,000	12,000	21,592	25,000	30,000	32,000	35,000	53,300	61,272	32,259	18.1%	3,493
South	12,000	15,000	22,000	25,000	26,000	30,000	32,000	56,995	68,000	33,827	15.1%	9,488
West	10,000	12,000	20,000	22,000	25,000	27,000	30,000	55,000	65,000	32,026	6.1%	3,305
Panel C: High Caste												
All India	**10,000**	**12,800**	**24,000**	**25,000**	**30,000**	**32,600**	**35,000**	**64,000**	**75,000**	**36,885**	**10.5%**	**23,627**
North	10,000	15,000	25,000	30,000	35,000	37,000	40,000	70,000	80,000	39,794	14.8%	5,982
Central	8,000	10,000	18,000	20,000	23,000	25,000	30,000	55,000	63,662	29,967	7.5%	5,349
East	9,189	12,000	20,000	24,566	25,450	30,000	35,000	56,000	65,000	33,255	11.7%	3,900
North-East	9,000	12,000	20,000	25,000	28,000	31,000	35,000	56,404	68,000	33,369	15.0%	1,257
South	14,000	15,560	28,000	30,000	35,000	35,565	40,000	70,000	82,139	42,223	12.2%	2,704
West	13,530	15,500	25,480	30,000	32,000	35,000	40,000	68,000	80,000	42,240	4.4%	4,435

Sample: Urban workers who have completed a graduate degree or higher. Pi denotes the ith percentile of the distribution of monthly earnings (from self-employed, salaried or casual work). Earnings in Rs. (2024-25). All statistics are unweighted unless otherwise noted.

Table 4d: PLFS 2025 — Monthly Earnings, Urban India, Non-Graduate Workers

Region	P15	P20	P40	P45	P50	P55	P60	P85	P90	Mean (Rs.)	Unemp. Rate within LF	N (earners)
Panel A: All Castes												
All India	**4,500**	**6,000**	**10,863**	**12,000**	**13,000**	**15,000**	**15,000**	**25,000**	**30,000**	**15,568**	**3.2%**	**124,190**
North	5,000	6,518	12,000	13,000	15,000	15,000	16,000	27,000	35,000	16,884	4.6%	22,572
Central	2,856	5,000	9,000	10,000	10,863	12,000	12,500	20,000	25,000	12,869	2.6%	27,939
East	3,500	5,214	10,000	10,863	12,000	12,500	14,000	21,000	25,000	13,851	3.6%	17,416
North-East	4,256	6,000	11,000	12,167	13,400	15,000	15,280	28,682	35,000	16,526	4.5%	10,106
South	6,000	7,500	12,000	13,525	15,000	15,200	16,500	25,000	30,000	16,661	2.8%	27,911
West	5,214	7,000	12,000	13,036	15,000	15,000	16,000	28,000	35,000	17,507	1.8%	18,246
Panel B: Low Caste												
All India	**4,345**	**6,000**	**10,200**	**12,000**	**12,167**	**14,000**	**15,000**	**24,000**	**28,000**	**14,654**	**3.1%**	**89,731**
North	5,000	6,500	12,000	12,000	13,688	15,000	15,000	24,333	28,000	15,208	4.7%	13,773
Central	2,500	4,500	8,799	10,000	10,429	12,000	12,167	19,500	23,000	12,118	2.3%	22,314
East	3,476	5,214	10,000	10,500	12,000	12,167	14,000	20,000	25,000	13,370	3.8%	11,564
North-East	4,000	6,000	10,863	12,170	13,500	15,000	15,600	30,000	35,000	16,666	4.6%	7,581
South	6,000	7,500	12,000	13,036	15,000	15,000	16,000	25,000	30,000	16,196	2.8%	23,953
West	5,000	6,518	11,000	12,000	13,000	15,000	15,000	25,000	30,000	15,754	1.7%	10,546
Panel C: High Caste												
All India	**5,000**	**7,000**	**12,000**	**13,500**	**15,000**	**15,000**	**16,500**	**30,000**	**35,000**	**17,948**	**3.4%**	**34,459**
North	4,563	7,000	13,036	15,000	15,000	16,729	18,000	35,000	40,000	19,509	4.5%	8,799
Central	3,500	6,000	10,429	12,000	12,500	14,516	15,000	25,081	32,000	15,847	3.4%	5,625
East	3,872	6,000	10,000	11,500	12,000	13,000	14,400	25,000	30,000	14,800	3.3%	5,852
North-East	5,000	6,500	11,312	12,000	13,036	15,000	15,000	25,000	31,427	16,107	4.4%	2,525
South	6,000	8,000	14,000	15,000	15,304	17,381	18,000	30,000	35,800	19,476	3.3%	3,958
West	6,000	8,000	14,000	15,000	15,208	17,000	18,250	30,000	40,000	19,908	1.8%	7,700

Sample: Urban workers who have not completed a graduate degree or higher. Pi denotes the ith percentile of the distribution of monthly earnings (from self-employed, salaried or casual work). Earnings in Rs. (2024-25). All statistics are unweighted unless otherwise noted.