

The Welfare Effects of Encouraging Rural-Urban Migration

David Lagakos
BU and NBER

Mushfiq Mobarak
Yale University and NBER

Michael E. Waugh
Federal Reserve Bank of Minneapolis and NBER

October 2022

ABSTRACT

This paper studies the welfare effects of encouraging rural-urban migration in the developing world. To do so, we build and analyze a dynamic general-equilibrium model of migration that features a rich set of migration motives. We estimate the model to replicate the results of a field experiment that subsidized seasonal migration in rural Bangladesh, leading to significant increases in migration and consumption. We show that the welfare gains from migration subsidies come from providing better insurance for vulnerable rural households rather than from correcting spatial misallocation by relaxing credit constraints for those with high productivity in urban areas that are stuck in rural areas.

Email: lagakos@bu.edu, ahmed.mobarak@yale.edu, michael.e.waugh@gmail.com. For helpful comments we thank Paco Buera, Matthias Doepke, Mitch Downey, Ben Faber, Ed Glaeser, Greg Kaplan, Louis Kaplow, Sam Kortum, Melanie Morten, Paul Niehaus, Michael Peters, Natalia Ramondo, Chris Tonetti, Stephen Terry, Rob Townsend, Fabrizio Zilibotti and seminar participants at many universities, conferences, and research centers. For outstanding research assistance we thank Elizabeth Carls, Menaal Ebrahim, Patrick Kiernan, Seungmin Lee, Harrison Mitchell and Alejandro Nakab. For financial support we thank the International Growth Centre. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. All potential errors are our own.

1. Introduction

Differences in income per capita across countries are largely accounted for by differences in total-factor productivity (TFP), and within an economy, misallocation of factors of production across firms, sectors, or regions may underlie these TFP differences.¹ One potentially large source of misallocation is an inefficient distribution of workers across space, and in particular across rural and urban areas, between which gaps in wages and productivity are particularly large (Gollin, Lagakos, and Waugh, 2014; Herrendorf and Schoellman, 2018; Young, 2013).

The view that workers are misallocated across space is seemingly reinforced by a recent series of field experiments in rural Bangladesh. In these experiments, small subsidies for seasonal migration led to substantial gains in income and consumption for migrants over multiple years (Akram, Chowdhury, and Mobarak, 2017; Bryan, Chowdhury, and Mobarak, 2014). Bryan et al. (2014) articulate this view in a model in which migration is risky and households face credit constraints that limit migration. In their model, migration subsidies reduce spatial misallocation by helping rural households accumulate enough assets to migrate to the city, where they are permanently more productive. However, their model is quantitatively inconsistent with the returns to migration and seasonal migration rates in the experimental data.

In this paper, we provide a reinterpretation of these migration experiments and a normative evaluation of migration policies using a dynamic incomplete-markets model that is estimated to match the experimental data. Migration subsidies, in partial or general equilibrium, direct resources toward rural households that have experienced bad shocks and are willing to undergo the ordeal of seasonal migration in order to bolster their low consumption levels. Rural households in the model generally prefer to stay in the rural area, and migration decisions are not very constrained by lack of access to credit. Instead, the lack of access to credit and good savings opportunities keeps households vulnerable to bad shocks in the first place, which lead them to migrate seasonally. In other words, seasonal migration is more of an act of desperation than a “move to opportunity,” which is often associated with permanent migration.

¹See Restuccia and Rogerson (2008), Hsieh and Klenow (2009) and more recent work that links misallocation to financial frictions (Buera and Shin, 2013; Midrigan and Xu, 2014; Moll, 2014); heterogeneous markups (Peters, 2020); delegation and information frictions (Akcigit, Alp, and Peters, 2021; David, Hopenhayn, and Venkateswaran, 2016); entry barriers (Yang, 2021); and other government policies (Fajgelbaum, Morales, Suárez Serrato, and Zidar, 2019; Guner and Xu, 2008).

In order to understand the migration behavior observed in the experiments, the model is tailored to feature a rich set of migration motives that have been emphasized in the literature. It allows for migration risk, following a long tradition in development economics (e.g., Harris and Todaro, 1970) in which households face both idiosyncratic productivity shocks and seasonal income fluctuations. Households insure themselves through buffer-stock savings, as in Bewley (1977), Aiyagari (1994) and Huggett (1993) and following a large literature in macroeconomics (see, e.g., Heathcote, Storesletten, and Violante, 2009; Kaplan and Violante, 2010). The model also features sorting by comparative advantage in the urban and rural regions, as in Roy (1951), following a number of recent studies (including Hamory, Kleemans, Li, and Miguel, 2021; Herrendorf and Schoellman, 2018; Lagakos and Waugh, 2013; Young, 2013). Households can migrate either permanently or temporarily, as in Kennan and Walker (2011), and face both a monetary cost of migration and a non-monetary disutility from migration that depends on past migration experience.

We estimate the model using the experimental data as well as several aggregate and cross-sectional moments. Matching these moments helps us isolate the characteristics of workers who are near the margin – meaning those who can most easily be induced to migrate – relative to those who are unlikely to migrate, those already migrating regularly or those who are permanently located in cities. The estimated model implies that workers near the margin must be modestly negatively selected on productivity and assets and that the non-monetary disutility associated with migration is substantial on average for those who have not migrated recently. While each targeted moment plays a role in informing the model’s parameter values, we show that the prediction of negative selection is governed in large part by the relatively large local average treatment effect (LATE) of migration on consumption coming from the experiment relative to the smaller OLS “return” to migration coming from the cross-section.

We show that our model is consistent with a number of experimental and cross-sectional moments, including several that are not targeted. Our model predicts a negligible migration response to an unconditional transfer, for example, which is consistent with a subsequent experiment conducted on the same population. Second, both in our model and in the data, people who choose to migrate are those with lower-than-average consumption and asset levels. These observations are consistent with our interpretation that migration subsidies provide households a chance to more easily supply labor in urban areas in periods when rural opportunities are lacking. We show that the model of Bryan et al. (2014), based on credit constraints, makes counterfactual predictions on these and other dimensions.²

²In terms of methodology, our work follows the seminal papers by Todd and Wolpin (2006) and Kaboski and Townsend (2011). These discipline dynamic structural models using quasi-experimental evidence rather

We use our estimated model to quantify the welfare gains from subsidizing rural-urban migration either temporarily or permanently. The welfare gains from a one-time migration subsidy in partial equilibrium are about 0.5 percent in consumption-equivalent welfare on average and are substantially higher for the poorest households. Moreover, the welfare gains from migration subsidies are similar on average to the welfare gains arising from a one-time unconditional transfer program costing the same total amount, though they are somewhat worse for the poorest households. This suggests that migration per se is not the main source of the welfare gains from migration subsidies; rather, they come from targeting resources to needy households willing to undergo the ordeal of migration.

When migration subsidies are offered permanently to rural households with low assets in general equilibrium, the welfare gains are larger on average for poor rural households but negative for richer (mostly urban) households that finance the subsidies. The average welfare gain to rural households with low assets is now around 3 percent in consumption equivalents. Many rural households respond to the presence of permanent migration subsidies by holding fewer assets in equilibrium, and the rural population share rises. As in the case of temporary transfers, we show that most of the welfare gains arise from targeting resources to vulnerable rural households. In other words, the gains from offering permanent migration subsidies arise not through relaxing credit constraints for those stuck in rural areas, but from providing households with better insurance.

In order to quantify the importance of workers' misallocation across space, we solve the problem of a benevolent social planner in charge of making migration and consumption decisions for each household in each period. We then compare the solution to the planner's problem with the competitive equilibrium outcomes. Analytically, we show that the solution to the social planner's problem involves a novel set of conditions characterizing optimal migration decisions for each household, where the benefits of moving any individual household involve both a static component from higher current productivity and a dynamic component reflecting the future option value of having the household in the migration destination.

When we compute the social planner's solution for the estimated model, we find relatively modest welfare gains from better allocating workers across space. Moreover, the planner's

than non-experimental moments, which are most common in macroeconomics. Our paper builds on these by estimating our structural model directly using variation induced by an RCT, in which concerns about endogeneity are even less present. In this sense, our quantitative work is similar to that of Buera, Kaboski, and Shin (2014), who use a macro model to help interpret the general-equilibrium effects of unconditional asset transfer programs, Greenwood, Kircher, Santos, and Tertilt (2019), who build a general equilibrium model of the AIDS epidemic to complement the many related RCTs, and Brooks and Donovan (2020), who draw on a structural general equilibrium model to help study the impacts of bridges built quasi-randomly.

solution entails substantially *fewer* seasonal moves than the competitive equilibrium. We show that many of these gains come from curtailing seasonal moves for marginal households that migrate after receiving a bad shock in the rural area when assets are low. Overall, comparisons of the planner’s solution with the competitive equilibrium provide very little support for a simple story of spatial misallocation in which credit constraints keep workers in rural areas. Instead, such comparisons point toward a different type of misallocation, which arises from a lack of insurance possibilities.

Our paper contributes most directly to a recent literature on the economics of rural-urban migration in the developing world. Morten (2019) and Meghir, Mobarak, Mommaerts, and Morten (Forthcoming) also emphasize the links between insurance and seasonal migration, though their focus is mostly on network effects and spillovers from individual migration decisions. Neither of those studies explores on the role of migrant selection by asset levels or non-monetary migration costs, or compares market outcomes to the efficient migration levels. Bryan and Morten (2019) find a large role for non-monetary migration costs, as in the current study, though their focus is more on permanent migration. Our paper’s focus on migration subsidies complements studies of other related policies, such as Imbert and Papp (2019), who show that rural public works can reduce seasonal migration to cities, and de Janvry, Emerick, Gonzalez-Navarro, and Sadoulet (2015), who find that land titling raises out-migration from rural areas. See Lagakos (2020) for a recent overview of this growing literature.

2. The Migration Experiments: A Summary

In this section, we summarize the experimental results of Bryan et al. (2014) and Akram et al. (2017) that motivate our modeling and estimation choices. The setting for both experiments are rural, rice-growing areas in the Rangpur region of Bangladesh, home to around ten million people. Like many other agrarian societies, these areas experience a “lean season,” known in Bangladesh as the *monga*, during the three-month period between planting and harvest. During this time, farmers mostly wait for the crop to grow, and labor demand falls. As a result, landless laborers experience a drop in wages and employment opportunities and incomes fall by an estimated 50 percent or more, on average (Khandker, 2012). To cope, some households migrate to towns and cities in search of employment during the lean season.

In the first experiment, reported in Bryan et al. (2014), 19 poor households were randomly sampled from each of 100 randomly selected villages in two districts in the Rangpur region. “Poor” was defined as households with almost no land holdings (less than 50 decimals

of land) and that reported having missed meals during the previous lean season. These households fall in roughly the lower half of the asset distribution. In August 2008, 68 villages were randomly assigned to treatment and 32 to control. In the 19 households in each of the treatment villages, subsidies encouraged one household member to migrate during the lean season. There were no subsidies in the control villages.³ The travel subsidy was worth about 800 Taka (\$11.50), which is sufficient to pay for round-trip bus fare plus a few days of food and is equivalent to about seven to ten days of rural wages during the lean season.

All 1,900 sample households were surveyed in December 2008 (post-treatment) and June 2009 about their migration and consumption during the 2008 lean season. The random assignment of migration subsidies produced three important outcomes that will inform our modeling choices. First, while 36 percent of households in control villages sent a migrant during the lean season, 58 percent of households in treatment villages did so (Bryan et al, Table II). Second, if we use the randomized treatment assignment as an instrumental variable for migration, the estimated local average treatment effect (LATE) of migration on consumption was an increase of around 30 percent per household member (Table III of Bryan et al). Migrants reported taking jobs such as rickshaw driving and construction work, which raised their household incomes. The LATE is too large to be explained by treatment households simply consuming the transfer. In practice most of the subsidy was put towards bus fare. Third, the treatment and control groups were surveyed a year later, in December 2009, though neither group received any additional treatment. Re-migration rates during the 2009 lean season remained nine percentage points higher in the treatment group, and this was statistically significant (Table II of Bryan et al). Subsequent results in 2011 and 2013 show elevated, but decaying, migration rates in the treatment group.

The second experiment (Akram et al., 2017) was conducted in 2014 on a larger scale, with migration offers extended to 5,792 poor, landless households. The authors measure income at high frequency and show that the migration offers led to significant increases in income of a magnitude consistent with the consumption increases observed in 2008. The new experiment also finds repeat migration effects of that one-time transfer during 2015-16, similar to the re-migration observed in 2009. Notably, the main experimental results from 2008-09 that we use to estimate our model are consistent with this much larger experiment.

³The 32-village control group comprises of a pure control (16 villages) and an information treatment (16 villages in which general information about migration possibilities were offered, but without any travel subsidy), which looks indistinguishable from the control group in terms of the migration response. The 68-village treatment group is composed of travel subsidies in the form of a grant (37 villages) or a zero-interest loan (31 villages). The grant and loan treatments produced very similar outcomes, so for simplicity, we combine them, refer to them as the “the treatment group,” and compare their outcomes with those of the combined control group.

For the purposes of our model, it is important to note that this second experiment added random variation in the proportion of the landless population in the treatment villages that was provided migration subsidy offers simultaneously. This labor-market-level variation created labor supply shocks of different magnitudes in different villages, which provide an experimental estimate of the village wage response to out-migration. Later, we use this estimate to inform the general-equilibrium effect of emigration on the rural labor market in our model.

3. Model of Migration to Interpret the Experiments

We turn now to a model of migration that is designed to be able to match the rich experimental evidence outlined above as well as salient cross-sectional facts, such as the rural-urban average wage gap. We focus on a stationary distribution of the model, in which the fraction of workers in each region and other aggregate variables remain constant in each period, as does the distribution of workers by state.⁴ The model is flexible enough to incorporate several broad determinants of rural-urban migration, including worker sorting on permanent comparative advantage, credit and saving frictions, temporary and seasonal shocks to rural and urban productivity and migration costs of a monetary and non-monetary nature. We show in Section 4.9 that this model succeeds in matching the data, whereas that of Bryan et al. (2014) fails. Moreover, while this model is flexible enough to entertain the ideas of Bryan et al. (2014), we show that it offers an entirely different interpretation of the experimental data.

3.1. Economic Environment

Preferences. Households are infinitely lived and maximize expected discounted utility. Their period utility function over consumption, c_t , is given by

$$u(c_t) = \frac{1}{1 - \alpha} c_t^{1-\alpha} \bar{u}^{x_t}, \quad (1)$$

where α is the coefficient of relative risk aversion; \bar{u} captures the non-monetary costs of migration; and $x_t \in \{0, 1\}$ is an indicator variable representing whether the household is an “inexperienced migrant.” The households’ problem is dynamic, and households discount the future at rate $0 < \beta < 1$.⁵

⁴In reality, net migration from rural to urban areas was positive over this period, though we abstract from this in the model, since it seems unlikely to change the model’s basic interpretation of the data.

⁵We model the disutility as multiplicative, rather than additive, because it is more flexible with respect to wealth effects in migration decisions. An additive urban disutility builds in a smaller disincentive for wealthier households to migrate, relative to that of poorer households. The data suggest the opposite pattern:

Inexperienced migrants experience disutility \bar{u} if they locate in the urban area in period t , whereas experienced migrants experience no such disutility. After each period in the urban area, inexperienced migrants become experienced with probability $1 - \lambda$. This is meant to capture any way in which rural-urban migrants become accustomed to being in urban areas—for example, by developing a network of friends or potential employers, or adapting to housing conditions. Experienced migrants can become inexperienced again after returning to the rural area. In each period in the rural area, the probability that an experienced migrant will become inexperienced again is $1 - \pi$. The motivation behind these modeling choices is twofold. First, we want to model the fact that migrants dislike certain aspects of migrating to an urban area (see the discussion in Appendix Section C). Second, we also want to model the idea that one’s utility from a location improves as one becomes accustomed to living there.

Migration decisions are also subject to additive idiosyncratic taste shocks, which we describe formally below. These taste shocks are independently and identically distributed across time and migration options and are drawn from a Type-1 extreme value distribution with scale parameter σ_ν , following the quantitative migration literature (see, e.g., Caliendo, Dvorkin, and Parro (2019)). The rationale for migration taste shocks is that many migration decisions are undertaken for (largely) non-economic reasons, such as marriage.

Endowments. Households supply one unit of labor inelastically, with efficiency units that vary across time and across locations, as in Roy (1951). Households differ in permanent productivity in the urban area z , which is drawn from a Pareto distribution: $z \sim 1 - z^{-\theta}$, where $z \geq 1$ and the shape parameter θ controls the variance of urban productivity. Here, a lower θ implies more variability in urban productivity. Households are identical in rural permanent productivity, and this value is normalized to one. Thus, the vector $\{1, z\}$ describes a household’s permanent productivity in the rural and urban areas.⁶

Households experience idiosyncratic transitory shocks to their endowments. Let s_t denote the current shock. This shock evolves according to an AR(1) process (in logs):

$$\log s_{t+1} = \rho \log s_t + \epsilon_{t+1} \quad \text{with} \quad \epsilon_{t+1} \sim \mathcal{N}(0, \sigma_s),$$

poorer households are more likely to migrate than wealthier households. Thus, from a pragmatic point of view, a multiplicative disutility allows the model more flexibility in fitting the data, rather than imposing a predetermined pattern between migration and wealth. Appendix Table A.1 reports the best fit of the data in an alternative model with additive disutility and highlights its difficulties in matching migration rates in the control and treatment groups and the consumption response to migration subsidies.

⁶The assumption of one-sided selection is supported by the empirical observation that we see very low variance in the level of consumption in rural areas. Moreover, this assumption eases the computational burden, allowing us to introduce transitory shocks and behavioral responses to them.

where ρ is the autocorrelation parameter and σ_s is the standard deviation of the shocks.

To allow this shock to have a differential impact on earnings (and risk) across locations, we assume that the household-specific, transitory component on efficiency units is s for the rural area and s^γ for the urban area. Thus, the vector $\{s, zs^\gamma\}$ describes a household's endowments (both permanent and transitory) for the rural and urban areas.

The parameter γ governs differential risk across locations. In particular, if $\gamma > 1$, this formulation will imply that shocks have a larger impact on incomes in the urban area than in the rural area. Hence, if $\gamma > 1$, the urban area will be riskier than the rural area. The benefit of this modeling choice is that it allows us to reduce the dimensionality of the state space to focus on just one shock (versus multiple shock processes across locations). At the same time, it parsimoniously captures the old idea in development economics that differential risk in urban and rural areas may be a deterrent to migration, as well as a source of urban-rural average income differences (Harris and Todaro, 1970).⁷

Production. There is one homogeneous good produced in both locations by competitive producers. Locations differ in the technologies they operate. The rural technology is

$$Y_r = A_r^i N_r^\phi, \quad (2)$$

where N_r are the effective labor units working in the rural area; $0 < \phi < 1$, so that there is a decreasing marginal product of labor in the rural area; and A_r^i is rural productivity indexed by season i . Seasonality is modeled with the rural area experiencing deterministic, seasonal fluctuations. Specifically, rural productivity can take on one of two values: $i \in \{g, \ell\}$ with $A_r^g > A_r^\ell$, where if current rural productivity is A_r^g , then the economy deterministically transits to productivity state A_r^ℓ in the next period, and so forth. Superscript g is for "growing" season, and superscript ℓ is for "lean" season.

The urban technology is given by

$$Y_u = A_u N_u, \quad (3)$$

where A_u captures urban productivity and N_u is the effective labor units supplied by households working in the urban area. Notice that N_u and N_r do not sum to one but

⁷The assumption of perfectly correlated shocks across regions is not an especially realistic one. We make it only for computational tractability. In Appendix Table A.2 we compute the welfare effects of migration subsidies in a model in which shocks are uncorrelated across regions and over time (which is the case when $\rho = 0$). The welfare effects are not that different in that model from the benchmark case to follow. This suggests that the welfare gains from a model with imperfectly correlated shocks would also be similar to those of the benchmark estimation.

are the sum across efficiency units. Thus, they depend on the shock realizations and the pattern of worker sorting across sectors.

Wages. In season i , with N_r workers in the rural area, wages per efficiency unit are

$$\omega_{r,i}(N_r) = A_r^i \phi N_r^{\phi-1} \quad \text{and} \quad \omega_u = A_u. \quad (4)$$

Agents working in a particular location receive wages that are the product of (4) and the number of their efficiency units (in both permanent and transitory terms). Thus, the labor income that a household with permanent state $\{1, z\}$ and transitory state s receives for working in location i is:

$$w_r(s, i) = s\omega_{r,i} \quad \text{and} \quad w_u(z, s) = z s^\gamma \omega_u, \quad (5)$$

which depends on the product of a household's permanent and transitory productivity and wages per efficiency unit in (4).

Location Options. Households have choices about where to reside and work. Those in the rural area have three options. First, they can work in the rural area. Second, they can pay a fixed cost m_T and work in the urban area for one period and return to the rural area in the next period. This is (temporary) *seasonal* migration in the model: a one-period working spell in the urban area by a rural household. Third, the household can pay a higher fixed cost $m_P > m_T$ and work in the urban area for the indefinite future. This is *permanent* migration, which enables the household to live and work in the urban area for more than one period.⁸

Households residing in the urban area have similar options. They can work in the urban area, or they can pay a fixed cost m_P and work in the rural area for the indefinite future. The latter option allows for rural-to-urban and then urban-to-rural moves as a household's comparative advantage, experience, and asset holdings change over time.

Asset Choices. Households can accumulate a non-state-contingent asset, a , with a gross rate of return, R . Asset holdings are restricted to be non-negative, and thus there is no borrowing. Furthermore, we assume that R is exogenous.

3.2. Optimization

Before we describe the value functions of a household, it is important to have a complete accounting of the state space. The state variables for a household can be divided into

⁸The higher cost of a permanent move reflects, at least in part, the greater expenses of moving all of a household's property rather than just a part of it, as required for a temporary move.

objects that are permanent, transitory, endogenous and aggregate.

- **Permanent productivity state.** Each household is endowed with z efficiency units in the urban area and one efficiency unit in the rural area.
- **Transitory productivity state.** Each household faces transitory productivity, s .
- **Transitory moving shock.** Each household is subject to an i.i.d. moving shock, ν .
- **Endogenous state variables.** There are three endogenous (individual) state variables. The first is the household's asset holdings, a . The second is a composite variable that describes the household's location and migration status. The possible states are rural (r), seasonal-migrant ($seas$, i.e., living in the rural area but working in the urban area for one period), and urban (u). The third is whether the household is an inexperienced migrant, x , and thus whether it suffers disutility \bar{u} from locating in the urban area.
- **Aggregate state variables.** There are two aggregate state variables: the season, $i \in \{g, \ell\}$, and the number of workers in the rural area, N_r . The season determines the current and future productivity in the rural area, and the two aggregate states jointly determine the current wage per efficiency unit as in equation (4).

We begin with the problem of a rural household. Because z is time-invariant for each household, we omit it as a state variable from the formulation of the household's problem below.

Rural Households. A rural household with productivity z solves the following problem:

$$v(a, r, s, \nu, x, i, N_r) = \max \left\{ v(a, r, s, x, i, N_r | \text{stay}) + \nu^{\text{stay}}, v(a, r, s, x, i, N_r | \text{seas}) + \nu^{\text{seas}}, v(a, r, s, x, i, N_r | \text{perm}) + \nu^{\text{perm}} \right\}, \quad (6)$$

where a household chooses among staying in the rural area, seasonally moving, and permanently moving. The value function associated with each option and the household's taste shock associated with each choice influence this choice. Here we will follow the quantitative migration literature (see, e.g., Caliendo et al. (2019)) and assume that the taste shocks are independently and identically distributed across time and drawn from a Type-1 extreme value distribution with scale parameter σ_ν . This distributional assumption implies

that the probability of staying in the rural location is

$$P(a, r, s, x, i, N_r | \text{stay}) = \frac{\exp\{\sigma_\nu^{-1}v(a, r, s, x, i, N_r | \text{stay})\}}{\sum_{j_r} \exp\{\sigma_\nu^{-1}v(a, r, s, x, i, N_r, | j_r)\}},$$

where the sum across j_r 's are the different choices for rural households. Here the scale parameter shows up and modulates the strength of the preference shock in determining the move. For example, as σ_ν goes to infinity, then only the shock matters for the moving choice, and the probability of each individual choice is simply one over the number of choices.

Conditional on staying in the rural area, the household's value function is

$$v(a, r, s, x, i, N_r | \text{stay}) = \max_{a' \in \mathcal{A}} \left\{ u(Ra + w_r(s, i, N_r) - a') + \beta \mathbb{E}[v(a', r, s', \nu', x', i', N'_r)] \right\}, \quad (7)$$

which means that the household solves only a consumption-savings problem. The asset holdings must respect the borrowing constraint and thus must lie in the set \mathcal{A} . Given asset choices, a household's consumption equals the gross return on current asset holdings, Ra , plus labor income from working in the rural area, $w_r(z, s, i)$, minus future asset holdings. Next period's state variables are the new asset holdings, location in the rural area, the transitory productivity shock, the experience level, the subsequent season, and the aggregate rural efficiency units in the next period. The expectation operator is defined over two uncertain outcomes: the transitory shocks and the change in experience. Recall that if the household is experienced, it stays that way with probability π and becomes inexperienced with probability $1 - \pi$; if the household is inexperienced, then it stays inexperienced.

The value function associated with a permanent move is

$$v(a, r, s, x, i, N_r | \text{perm}) = \max_{a' \in \mathcal{A}} \left\{ u(Ra + w_r(z, s, i, N_r) - a' - m_p) + \beta \mathbb{E}[v(a', u, s', \nu', x', i', N'_r)] \right\}.$$

While this function is similar to the staying value function, there are several points of difference. First, the agent must pay m_p to make the permanent move, and this costs resources. Second, the continuation value function denotes that the household's location changes from the rural to the urban area.

The value function associated with a seasonal move is

$$v(a, r, s, x, i, N_r | \text{seas}) = \max_{a' \in \mathcal{A}} \left\{ u(Ra + w_r(s, i, N_r) - a' - m_T) + \beta \mathbb{E}[v(a', \text{seas}, s', x', i', N'_r)] \right\}. \quad (8)$$

If a household decides to move seasonally, it pays the moving cost m_T and works in the urban area in the next period. The key distinction between the permanent move and the seasonal move is that the seasonal move is for just one period. Hence, the location state variable is $seas$ and not u , as this indicates that the household is going to work in the urban area and return in the next period. The value function associated with a seasonal move while in the urban area is

$$v(a, seas, s, x, i, N_r) = \max_{a' \in \mathcal{A}} \left[u(Ra + w_u(z, s) - a')\bar{u}^x + \beta \mathbb{E}[v(a', r, s', \nu', x', i', N'_r)] \right]. \quad (9)$$

There are several important points to note in (9). First, this household has only one choice: how to adjust its asset holdings. By the definition of a seasonal move, the household works in the urban area for one period and then returns to the rural area. Second, note how the disutility from living in the urban area appears (i.e., the presence of \bar{u}). Moreover, the state variable of a household's experience x determines whether the disutility is experienced.

Equations (8) and (9) illustrate the forces that shape the decision to move seasonally and, in turn, our inferences from the experimental and survey results. Generally, the choice to move seasonally will relate to a household's comparative earnings advantage in the urban area relative to that in the rural area. However, several forces may lead a household with a permanent comparative advantage in the city not to move. First, the urban disutility may prevent the household from moving, even though its comparative advantage in the urban area is expected to be high. Second, there is risk associated with the move. A household does not know the realization of the transitory shock before moving, and hence there is a chance that the income realization in the urban area will not be favorable. Third, the household may have limited assets that simply make a move infeasible or not sufficient to insure against a bad outcome in the urban area.

Urban Households. Urban households face problems similar to those described above, though they choose between just two options: staying or making a permanent move. For a household with productivity level z , the problem is

$$v(a, u, s, \nu, x, N_r, i) = \max \left\{ v(a, u, s, x, N_r, i \mid \text{stay}) + \nu^{\text{stay}}, v(a, u, s, x, N_r, i \mid \text{perm}) + \nu^{\text{perm}} \right\}. \quad (10)$$

Again the value function associated with each option and the household's taste shock associated with each choice influence this choice. These taste shocks are independently and identically distributed across time and distributed Type 1 extreme value distribution with the same scale parameter σ_ν .

The value functions for urban households are analogous to those of the rural households, so we omit them for brevity. However, these households have several key differences from those staying in the rural area. First, their wage depends on their permanent productivity level, z , and not on the season or number of aggregate efficiency units in the rural areas. Moreover, the transitory productivity shocks may have more or less volatility relative to those of the rural area, as determined by the γ parameter (see equation (33)). Third, the disutility from living in the urban area appears (i.e., the presence of \bar{u}), and the state variable of a household's experience, x , determines whether the household suffers it.

As in the case with rural households, expectations are over the transitory shock and the change in experience. However, as these households are in the urban area, inexperienced households stay that way in the next period with probability λ and become experienced with probability $1 - \lambda$. Experienced households retain their experience. Urban households must pay m_p to make a permanent move back to the rural area. Furthermore, the continuation value function denotes the household's location changes from the urban to the rural area. After a permanent move to the rural area, experienced households keep their experience with probability π and lose it with probability $1 - \pi$.

3.3. Discussion: Determinants of Migration and Location Choice

The model allows for a rich set of determinants of migration and of location choice more generally. Although we allow the data to discipline the most important determinants in the following section, it is worth discussing them informally here first.

One clear determinant of migration in the model is the season. Since the growing season has higher productivity than the lean season, rural households will be more likely to migrate (seasonally or permanently) to the urban area in the lean season, all else equal. The permanent urban productivity level, z , is another important determinant of migration. All else equal, agents with higher values of z will have stronger incentives to locate in the urban area. The migration disutility, \bar{u} , is also an unambiguous deterrent to migration. The higher is \bar{u} , the less likely it is that inexperienced households will locate in the urban area. Furthermore, those with migration experience are more likely to migrate, as these households face no disutility of locating in the urban area. Finally, both effects—permanent comparative advantage and experience—will interact, as households with a stronger comparative advantage in the urban area are more likely to migrate and hence are more likely to be experienced at migrating.

The probabilities of gaining or losing experience, λ and π , mostly affect the extent of repeat migration. When experience is easy to obtain and hard to lose—that is, λ is low and π

is high—a subsidy to migration will induce inexperienced rural-urban migrants to repeat migrate (or to stay in the urban area) for many periods in the future. For rural households induced to migrate seasonally, the lower is π , the less likely they will be to migrate in subsequent periods, since experience is lost at a faster rate.

The transitory shock, s , and asset level, a , have ambiguous effects on migration and location choice. For concreteness, suppose that shocks are persistent, so that households with a high shock today are more likely to receive a high shock one period hence. Consider first the case that $\gamma > 1$, so that transitory shocks are more volatile in the urban area. In this case, rural households may be more likely to migrate to the urban area after receiving a good shock. The asset holdings also play a role in this case. High values of assets allow for insurance, which may mean that households migrate in this case only when their assets are sufficiently high. If this is the case, subsidizing migration may induce these high-productivity households to migrate and realize large consumption gains due to a better allocation of their urban-specific productivity. This is just how the model of Bryan et al. (2014) works, as we discuss further below. It is worth emphasizing that this case is more likely to occur the lower is the return to saving, R , since for higher savings rates, workers can self-finance and save their way out of these credit constraints (see, e.g., Donovan, 2021; Midrigan and Xu, 2014; Moll, 2014).

Consider next the case that $\gamma < 1$, so that shocks are more volatile in the rural area than in the urban area. In this case, rural households may be more likely to migrate when they have bad shocks than when they have good shocks. Since migration is costly in both monetary terms and non-monetary disutility, households may migrate only when they are sufficiently unproductive and when their assets are too low for them to insure themselves against their current low productivity. In this case, subsidizing migration may induce these low-productivity households to migrate and realize large consumption gains to avoid bad outcomes in the rural area and reap benefits of higher average productivity in the urban area. This case is related to the findings of Gröger and Zylberberg (2016) and Kleemans (2015), who find evidence that workers use migration as a coping mechanism after bad shocks.⁹ In practice, whether induced migrants tend to be low-productivity workers with low assets or high-productivity workers with high assets is determined by the data.

⁹For the case of international migration, Bazzi (2017) finds that credit constraints limit emigration from poorer rural areas in Indonesia, though in more developed rural areas, those with higher permanent income shocks are less likely to migrate.

4. Model Estimation

We now estimate the model using simulated method of moments. We draw on two sets of moments. The first is the migration experiments described in Section 2. The second is a large nationally representative household survey from Bangladesh. Taken together, these moments help the model to jointly fit key aggregate facts from the Bangladeshi economy relevant for understanding rural-urban migration, plus the household responses to migration incentives, which are well identified through the experimental evidence.

4.1. Data and Targeted Moments

To discipline the model, we draw on eight moments from the migration experiment of Bryan et al. (2014). These are (i) the variance of log consumption growth from before and after the lean season in the control villages (0.19); (ii) the percentage of control households with no liquid assets (47 percent); (iii) the seasonal migration rate in the control villages (36 percent); (iv) the OLS “return” to migration in the control villages (10 percent); (v) the seasonal migration rate in the treatment villages minus that of the control villages (22 percentage points); (vi) the same difference but in year 2 (9 percentage points); (vii) the IV return to migration (LATE) in terms of consumption (30 percent); and (viii) the probability of repeat migration for individuals in the control villages (0.68).

We take three moments from a large nationally representative household survey called the 2010 Household Income and Expenditure Survey, which surveyed 12,240 households. These moments are (i) the fraction of households residing in rural areas (62 percent); (ii) the ratio of urban to rural average wages (1.89); and (iii) the variance of log wages in the urban area (0.56). To construct the wage variance, we restrict attention to wage earners, since the data on wage earnings are likely to be more reliable than the data on self-employed income or farm income. We also restrict attention to males aged 20 and older that work “full time” (which we define as having worked at least 5 months in the last year, for at least 15 days in the last month, and for an average 5 or more hours per day). We compute the wage as monthly earnings in the main occupation divided by weekly hours multiplied by four.

4.2. Directly Chosen Parameters

We begin by directly assigning some parameter values; Table 1 provides a summary. These are parameters that are related directly between the model and the data or are difficult to identify from the data. We choose the time period to be half a year, which allows for seasonal migration and seasonal variation in rural productivity, which are important features in the experimental data. We set the risk-aversion parameter, α , to be two, which is within

Table 1: Pre-Assigned Parameters

	α	β	R	A_{rl}/A_{rg}	m_T	m_p	ϕ
Value	2.0	0.95	0.95	0.50	$0.1 \times \text{rural cons.}$	$2 \times m_T$	0.91

Note: the table reports the values of the 8 pre-assigned parameters in the model. A period is defined to be half a year; m_T is chosen to equal 10 percent of average rural consumption.

the range of commonly chosen values in the macroeconomics literature. We choose the discount factor, β , to be 0.95. The return on assets, R , is set to 0.95 to capture Bangladesh’s average half-yearly inflation rate, which is around five percent. This choice is consistent with the asset composition of households’ balance sheets in our experimental sample, which is primarily cash.¹⁰

We set the ratio of productivity in the lean season to that in the growing season, A_{rl}/A_{rg} , to be 0.5, consistent with estimates by Khandker (2012). The seasonal moving cost, m_T , is set at ten percent of average rural consumption. This is approximately the seasonal migration cost (round-trip bus fare plus a few days of food during travel) reported in Bryan et al. (2014). We set the permanent migration cost, m_p , high enough so that gross flows across regions are negligible, which is what is observed in the eight years of tracking in the Bangladeshi data. We find that our results are not substantially affected by this parameter value. We note that permanent migration is an important feature in the model because it allows individuals to sort endogenously into the rural and urban areas based on permanent comparative advantage (and the other other state variables.)

Finally, we set the elasticity of output with respect to labor, ϕ , to be 0.91, following the estimates of the effects of large-scale migration subsidies of Akram et al. (2017). They observe that rural wages rise more in villages that randomly receive more migration subsidies (and have more out-migration). Our choice of ϕ replicates their elasticity of a 2.2 percent increase in rural wages for every ten percent increase in seasonal migration. Akram et al. (2017) also document that the consumption increase from travel subsidies is entirely due to migration income earned by the migrant (with no change in labor supply of other household members), and our modeling choices reflect that fact.

4.3. Parameters to Estimate

We estimate the remaining eleven parameters of our model. The first nine are part of the model and defined above. They are (i) θ , the shape parameter controlling the urban indi-

¹⁰We experiment with different values of β and R in Appendix Tables A.3 and A.4, and we find that our model’s predictions are not substantively altered under alternative plausible choices.

vidual productivity distribution; (ii) \bar{u} , the disutility of migration; (iii) λ , the probability of remaining inexperienced after a move to the urban area; (iv) π , the probability of remaining experienced following a return to the rural area; (v) γ , the relative volatility of the urban area; (vi) A_u , the urban aggregate productivity level; (vii) σ_s , the standard deviation of stochastic shocks; (viii) ρ , the autocorrelation of urban shocks; and (ix) σ_ν , the standard deviation of idiosyncratic taste shocks.

The final two parameters govern the extent of measurement error in income and consumption, which we want to allow for, since income and consumption data at the micro level are clearly measured with noise. Hence, we do not want to force the model to ascribe all the income and consumption variance to permanent or temporary shocks rather than to measurement error. In particular, we assume that rural consumption growth (which we observe using the experimental data) satisfies

$$\hat{g}_{c,i} = g_{c,i} + v_{r,i}, \quad (11)$$

where $\hat{g}_{c,i}$ is observed consumption growth of household i ; $g_{c,i}$ is actual consumption growth; and $v_{r,i}$ is measurement error, which we assume is normally distributed with mean zero and variance $\sigma_{r,c}$. Urban income, in turn, satisfies

$$\log \hat{y}_i = \log y_i + \log v_{u,i}, \quad (12)$$

where y_i is observed income of household i ; y_i is actual income; and $\log v_{u,i}$ is measurement error in income, which we assume is normally distributed with mean zero and variance $\sigma_{u,i}$.

4.4. Estimation by Simulated Method of Moments

We estimate these eleven parameters of the model using simulated method of moments. The basic idea is to pick the parameter vector

$$\Theta = \{\theta, \bar{u}, \lambda, \pi, \gamma, A_u, \sigma_s, \rho, \sigma_\nu, \sigma_{r,c}, \sigma_{u,i}\} \quad (13)$$

to minimize the distance between the simulated moments and the data. The eleven data moments from which we estimate the parameters are listed in Table 2 and are divided into two basic groups: the experimental moments (top eight) and the cross-sectional survey moments (bottom three). We construct the simulated moments in the following way. For the cross-sectional moments, we solve the household's problem and construct the stationary distribution of households. From the stationary distribution, we compute the urban-rural

wage gap, the percentage of households that permanently live in the rural area, and the variance of log income in the urban area.

A novel feature of our estimation procedure is that we replicate the experiment to be targeted directly in our model. We implement this procedure in the following way. First, we present the model households with a one-time, unanticipated seasonal migration opportunity without the monetary cost m_T and compute their optimal responses. We do this in partial equilibrium, which is an appropriate choice given the relatively small number of experiment participants in each village and the relatively small number of villages in the experiment. We then randomly sample rural households from the model's stationary distribution, consistent with the sample selection criteria in the migration experiments discussed in Section 2. Specifically, Bryan et al. (2014) conducted their baseline survey before the lean season; thus, we follow the same timing in the baseline sample selection and measurement for our model. Furthermore, Bryan et al. (2014) selected households that were relatively poor to start with, and we implement this in the model by selecting rural households that are in the bottom half of the asset distribution for rural residents.

Given the appropriate sample of households and their optimal policies if treated or not, we compute the moments from the control and treatment groups described above. To compute the OLS return to migration in the control group, we regress the consumption of the model's rural households in the lean season on an indicator variable for whether the household migrated in the lean season. To compute the LATE of migration on consumption, we use data from both the treatment and control groups in the model and run an IV regression in which consumption is regressed on (instrumented) migration, with migration instrumented by assignment to the treatment group in the first stage.

4.5. Estimation Results

Table 2 presents the data and the model moments. In general, the model's predicted moments are quite similar to its counterparts in the data. Ten of the eleven moments are matched exactly, or almost exactly it, while one (the repeat migration rate) is somewhat lower in the model than in the data.¹¹

Table 3 shows the estimated parameter values and their bootstrapped confidence intervals. While the next two sections discuss the economic implications and identification of these parameter values, several features of Table 3 are worth pointing out. First, the shape pa-

¹¹Figure A.1 plots the difference in migration rates between the treatment and control groups in the model and data in 2008 (the year of the experiment in the model) and for five subsequent years. By five years after the experiment, the difference in migration rates between the two groups is positive but statistically insignificant.

Table 2: Moments Targeted in the Estimation

Moments	Data	Model
Control: Variance of rural log consumption growth	0.19 (0.03)	0.19
Control: Percentage of rural households with no liquid assets	47 (1.13)	47
Control: Seasonal migration rate	36 (2.64)	36
Control: Consumption increase of migrants (OLS)	10 (4.47)	10
Control: Repeat migration rate	68 (0.46)	71
Treatment: Seasonal migration relative to control	22 (2.39)	21
Treatment: Seasonal migration relative to control in year 2	9 (2.44)	5
Treatment: Consumption increase of induced migrants (LATE)	30 (9.67)	29
Urban-Rural wage gap	1.89 (0.18)	1.89
Percentage in rural area	62 (1.36)	60
Variance of log urban wages	0.56 (0.06)	0.56

Note: The table reports the moments targeted using simulated method of moments, their values in the data and in the model, and the standard errors of the empirical moments.

parameter controlling permanent differences in ability, θ , is quite low, at 0.55. This implies that there is large variation in permanent productivity in the urban area. Second, the urban relative risk parameter, γ , is less than one, implying that shocks in the urban area are less volatile than those in the rural area. Third, the disutility of migrating, \bar{u} , is sizable (and positive, since the level of household utility is negative). In terms of magnitudes, the value of 1.55 for \bar{u} implies that experiencing disutility of migration is equivalent, in a static sense, to cutting consumption by 33 percent. This is consistent with the large seasonal migration costs estimated by Imbert and Papp (2020) in India. Next, the probability of remaining inexperienced, λ , and that of remaining experienced, π , are substantially less than one, with around one-third of experienced households losing their experience each half-year, and only around one-third of inexperienced households gaining experience after a move.

The bootstrapped confidence intervals are reasonably tight for $1/\theta$, A_{ur} , and ρ . The migra-

Table 3: Estimated Parameters and Confidence Intervals

$\frac{1}{\theta}$	\bar{u}	λ	π	γ	A_u	σ_s	ρ	σ_ν
0.55	1.54	0.66	0.60	0.52	1.55	1.30	0.75	0.13
(0.42,0.60)	(1.34,1.90)	(0.40,0.90)	(0.23,0.88)	(0.15,0.85)	(1.40,1.70)	(0.80,2.25)	(0.59,0.87)	(0.01,0.21)

Note: The table reports the values of nine jointly estimated parameters and their bootstrapped 95-percent confidence intervals.

tion disutility, \bar{u} , and urban relative risk parameter, γ , have somewhat wider confidence intervals but are always consistent with a moving cost substantially above one and an urban relative risk parameter below one. The variance of the transitory shock and that of migration taste shocks are even wider, though the latter is small relative to other studies in the literature employing a migration taste shock. The probabilities of remaining experienced and inexperienced, π and λ , have wide confidence intervals and rule out only fairly extreme cases near zero or one.¹²

4.6. Who Migrates and Why?

In this section, we discuss how the estimated model’s policy functions for location choice depend on permanent productivity, asset holdings, the transitory shock, and migration experience. In the discussion of these outcomes, we discuss how the data informs them.

We focus on rural households leading into the lean season, since most migration occurs then. Figure 1 plots the migration probabilities in the estimated model for select rural households with different levels of urban productivity and migration experience as a function of their transitory shocks and asset holdings. The x -axis represents the transitory productivity shock, and the y -axis is the asset holdings of the household. For each transitory shock and possible asset value, the color represents the probability of migrating, with darker colors meaning higher probabilities.

Higher urban productivity leads to more migration. Panels (a) and (b) contrast the moving policy for low- z households and moderate- z households. The dark blue migration region in the southwest portion of the panels is larger for moderate- z than for low- z households. This means that those with a stronger comparative advantage are more likely to seasonally migrate to the urban area. Intuitively, the key data moment determining how many moderate- z households there are, compared with the number of low- z households, is the experimental migration response to the treatment.

¹²For expositional purposes, Table 3 omits the values of measurement error terms, σ_{rc} and σ_{ui} . These are each estimated to be 0.15, with bootstrapped confidence intervals of (0.11, 0.22) and (0.00, 0.29).

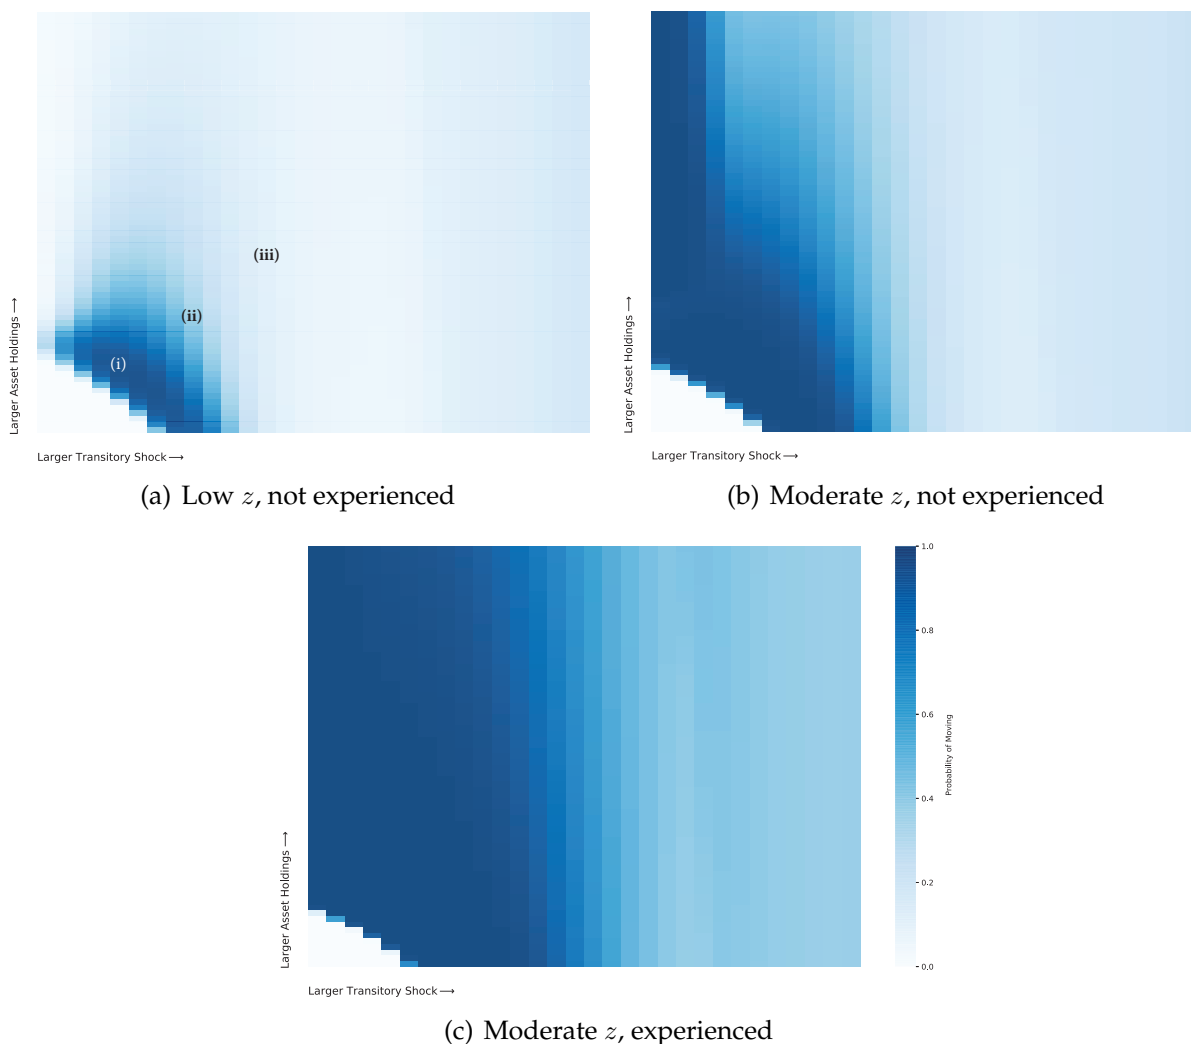


Figure 1: Migration Probabilities for Select Household Types

This observation highlights an important implication about whom migration subsidies may affect. Households sort themselves into rural and urban areas largely on the basis of permanent comparative advantage. Thus, the set of households living in the rural area in any given period consists most likely of those with relatively low z values. Households with higher z are more likely to be in the city. As a policy, migration subsidies will be offered to those with low comparative advantage in the urban area.

The disutility of the urban area is an important deterrent to migration. Panels (b) and (c) of Figure 1 contrast the moving policy for the same z but different experience levels. Hence, these households face a disutility to migrating or not. The dark blue migration region is larger for the experienced household than for the inexperienced one. This illustrates the point that in the estimated model, we infer an important role for a non-monetary disutility

of the urban area in shaping the migration choice.¹³

Several features of the data combine to push our model to infer a substantial non-monetary disutility of the urban area. At the most basic level, this is about the overall level of migration (not the experimental response which informs the distribution of ability). In the lean season, productivity falls by fifty percent, yet many do not migrate. One way for our model to accommodate this observation is to have a large disutility of migration.

The alternative explanation is that households simply cannot afford to migrate. The important observation is that whatever force prevents migration must be consistent with the consumption response and the pattern of selection implied by it. As we discuss next, in both the data and the model, migrants are negatively selected on assets and income. And a large disutility of migration (to get the overall migration flows correct) is consistent with the pattern of selection in the data; a credit constraint story is not.

Households with low assets and low transitory shocks are more likely to migrate. All three panels of Figure 1 highlight how households with low assets and low transitory shocks are more likely to migrate. This point is seen by noting that in all cases, the dark blue migration region always originates out of the southwest corner. If one's expectation is that credit constraints are the primary reason that households do not migrate, this may seem surprising. If migration costs are high and the credit constraint binds, then the migration region would originate from the northeast corner in Figure 1, because this is where the constraint would be alleviated. In fact, credit constraints prevent migration for a very small part of our parameter space, seen as the white region in the very lower left corners of all three panels in the figure.

The key behind this result is the observation that the OLS coefficient from consumption regressed on migration is smaller than the IV (LATE) coefficient. Because OLS is smaller relative to IV, the implication is that migrants are *negatively selected* on the determinants of consumption. In our model, a household's state variables are the determinants of its consumption — specifically, its transitory shock and asset holdings. Thus, the fact that the OLS coefficient is less than the IV coefficient pushes our model to accommodate the idea

¹³To empirically investigate the source of the substantial migration disutility in our estimated model, we conducted some additional discrete-choice experiments on the same set of households used to estimate the model. We found evidence that potential migration opportunities involving poor housing conditions at destination were substantially less desirable in wage terms, all else equal. In particular, offering improved housing with a proper indoor latrine increases reported migration propensity by 17.4 percentage points. This effect size is equivalent to the effect of increasing migration wages by 21 percent. Thus, the lack of a good housing option after migrating, which is likely quite common for potential migrants in this setting, is one concrete example of the migration disutility in our model. See Appendix C for a complete description of these discrete-choice experiments and our empirical findings.

that households with low assets and low transitory shocks are more likely to migrate as in Figure 1.

This push is achieved through a large \bar{u} parameter (as described above), but it is also facilitated by the inference that rural area is very risky $\gamma < 1$ and that households have a low return on their savings and hence difficulties self-insuring. That is, rural households generally prefer rural areas but sometimes find themselves in periods of low productivity and low assets, particularly in the lean season. Because self-insurance through savings is difficult, these households use seasonal migration as a form of insurance allowing them to temporarily raise their incomes and thus smooth their consumption over time. This latter observation is essentially a spatial analog to Pijoan-Mas (2006), in which households with low productivity or low assets increase their labor supply to self-insure.

This observation has an important implication for how migration subsidies may improve welfare in our model. They serve to channel resources toward the rural households that are unproductive and vulnerable owing to a lack of assets that can be used for self-insurance. In other words, the migration subsidy provides households an insurance opportunity to more easily supply their labor in the urban location.

4.7. Non-Targeted Moments

How does the model fare in predicting non-targeted moments? We answer this question by examining several features of the data: migration rates by initial consumption level, asset holdings by migration status, variances of consumption growth by migration status, and the migration effects of unconditional cash transfers.

We focus on these non-targeted moments for the following reasons. Looking at how migration rates vary by initial consumption level speaks to households' heterogeneous responses to the treatment. Asset holdings by migration status tell us about the extent to which migration decisions are driven by buffer-stock savings strategies versus strategies in which migration serves as insurance when productivity and assets are low. The variance of consumption growth and the effects of unconditional transfers speak to the importance of potential spatial misallocation from credit constraints and migration risk, as in Harris and Todaro (1970) and the model of Bryan et al. (2014).

Migration rates by initial consumption level. Table 4 reports the migration rates in the control group by the levels of consumption and assets. Panel A exhibits migration rates from the data, and panel B reports the value in the estimated model (which are not targeted). For simplicity, we consider consumption only above and below the median and assets above and below the 800 Taka — approximately the size of the transfer. As

Table 4: Migration Rate by Asset and Consumption Level

Panel A: Data			
		Assets	
		≤ 800 Taka	> 800 Taka
Consumption	Below Median	40	29
	Above Median	36	31

Panel B: Model			
		Assets	
		≤ 800 Taka	> 800 Taka
Consumption	Below Median	41	28
	Above Median	41	38

Note: The table reports migration rates in the control group by consumption level and by level of assets, defined as cash plus bank account balances. Panel A shows the migration rate in data. Panel B reports the rates from the estimated model. Both are expressed in percentage points.

the table shows, migration rates are higher for people with lower consumption and asset levels. For the households with below-median consumption, those with less than 800 Taka in assets had a migration rate of 40 percent, compared to a migration rate of 29 percent for those with more than 800 Taka (p -value = 0.09). For those with above median consumption, migration rates were 36 percent and 31 percent for the two groups (p -value = 0.45). The model gets this prediction largely correct, with migration rates of 41 percent and 28 percent for households with below-median consumption, and migration rates of 41 and 38 percent for those with above-median consumption.

Effects of unconditional transfers. Table 7, Panel B, reports the migration response in the data of an unconditional transfer, which was conducted along with the original migration experiments. Households did respond positively to an unconditional transfer but to a much smaller extent than they did to the conditional transfers. However, the confidence interval is quite large and comfortably includes zero: the p -value of a test of the null hypothesis that the unconditional transfer has no effect is 0.24. In the model, the effect of an unconditional transfer is one percent higher migration. Thus, it is fair to say that the model predicts that an unconditional transfer has a smaller effect on migration than a conditional transfer, as in the data, and an effect that is small overall and within the confidence

Table 5: Variance of Log Consumption Growth

	Control Group		Treatment Group		
	Stay	Migrate	Stay	Migrate	
Data	0.15	0.18	Data	0.16	0.19
Model	0.18	0.19	Model	0.17	0.19

Note: The table reports the variance of log consumption growth from before the lean season to afterwards. The left panel is for the control group, and the right panel is for the treatment group. The columns represent the set of households that stay (do not send a migrant) versus that of those that migrate (do send a migrant).

intervals of the model’s prediction.

Variations of consumption growth by migration status. Table 5 lists the variance of log consumption growth for households that stay and those that migrate, in both the data and the model. In the data, the control group (upper panel) has log consumption growth variance of 0.15 for stayers and marginally higher variance, of 0.18 for those that migrate. The model is similar, with 0.18 for the stayers and a marginally higher at 0.19 for the migrants. The treatment group (lower panel) in the data is somewhat similar to the control group, and again the model matches the similar but marginally higher log consumption variance of the migrants.

It is worth discussing how our model correctly predicts higher consumption growth variance for migrants compared with that of non-migrants, even though it features higher transitory shock variance in the rural area ($\gamma < 1$). The reason is as follows. The model’s prediction is that households with relatively low transitory shocks and asset levels do more temporary migration, all else equal. In the estimated model, these temporary migrants see large gains in income and hence consumption, since they are largely “hand-to-mouth.” This tends to increase consumption growth variance for migrants. In the aggregate, this force leads to larger consumption growth variance for migrants, even though migrants face lower income risk at the individual level.

4.8. Matching the Model to the Data

In this subsection, we discuss how the main features of the model are necessary in order to match the moments in question. To do so, our approach is to re-estimate the model several times under specific parameter restrictions and then to report the best fit of the targeted moments in each case. We report several such re-estimations in Table 6. This approach generally results in a worse fit of the targeted moments since the model has fewer degrees

of freedom once parameters are restricted ex-ante. The value of the exercise comes from seeing exactly how the model fails once particular parameters are restricted.¹⁴

The first two columns reproduce the targeted moments from the data and the estimates from the baseline model. The third column illustrates the model's best fit taking $\theta \rightarrow \infty$ as a numerical limit, implying that heterogeneity in permanent productivity across workers is negligible. In this case the model greatly over-predicts seasonal migration rates in the control group and the OLS return to migration, while under-predicting the LATE return to migration and repeat migration. This highlights how worker heterogeneity in the model is necessary for matching observational and experimental migration rates as well as returns to migration.

The fourth column of Table 6 reports the best fit of the model when $\bar{u} = 1$ — that is, when migration disutility is shut off. This version of the model substantially over-predicts seasonal migration rates in the control group, and under-predicts migration rates in the treatment group. These shortcomings highlight how it is hard to generate so little seasonal migration, and such a large response of migration to the subsidies, without migration disutility. Moreover, note that leaving off migration disutility forces the model to under-predict both the OLS return to migration and LATE of migration on consumption. In other words, matching both the observational and experimental returns to migration requires some non-monetary migration disutility.

The next column explores the model's fit when $\lambda = 0$, meaning that no individual can retain migration experience. In this case, the model generates no repeat migration in the treatment group of the experiment. This illustrates the (intuitive) conclusion that the model's persistence of experience is important for matching repeat migration rates in the data. Moreover, since the model does fairly well in matching the other moments, we conclude that dynamics of experience are not crucial outside of repeat migration patterns.

The last column of Table 6 explores the case when γ is restricted to equal one, meaning that transitory shocks are equally volatile in rural and urban areas for potential migrants. The big failure in this case is that the OLS coefficient of migration on consumption is far too high, and counterfactually similar to the LATE of migration on consumption. The model's inability to match these two measured returns to migration signals the importance of differential shock volatility, and the resulting patterns of sorting, in matching the larger experimental effect of migration on consumption than the simple OLS return to migration.

¹⁴A related, but distinct, approach is to compute the elasticity of each targeted moment to each parameter, starting from the estimated model. Table A.5 reports these elasticities, which suggest broadly similar insights about which of the model's parameters are most relevant for each predicted moment.

Table 6: Model Moments When Model Is Re-Estimated under Parameter Restrictions

Moments	Data	Baseline	$\theta = \infty$	$\bar{u} = 1$	$\lambda = 0$	$\gamma = 1$
<u>Control</u>						
Percentage with no liquid assets	47	47	46	47	48	46
Seasonal migration rate	36	36	55	56	36	45
Consumption gain, migrants (OLS)	10	10	18	6	8	24
Control: Repeat migration rate	68	71	56	55	67	66
<u>Experiment</u>						
Seasonal migration rate	22	21	16	13	27	20
Migration, year 2	9	5	2	2	0	4
Consumption gain, induced migrants (LATE)	30	29	24	25	29	29
LATE - OLS	20	20	6	19	21	5
<u>Aggregate</u>						
Urban-Rural wage gap	1.89	1.89	1.89	1.91	1.90	1.90
Percentage in rural	62	60	61	65	63	62

Note: The table reports the values of eight moments of the model when the model is re-estimated under restrictions on model parameters. The first column contains the moments in the data, and the second contains the moments from the estimated model. The remaining columns present the model's moments when estimated to match the same moments but assuming that one particular parameter is restricted ex-ante. Specifically, these columns focus on the cases when $\theta = \infty$, meaning no permanent worker heterogeneity; when $\bar{u} = 1$, meaning no migration disutility; when $\lambda = 0$, meaning no chance of retaining migration experience; and when $\gamma = 1$, meaning transitory shocks are equally volatile in both regions.

4.9. Comparison to Model of Bryan et al. (2014)

We now compare our model with the one proposed by Bryan et al. (2014) and show that their model is quantitatively inconsistent with the experimental evidence. As a result, it leads to an inaccurate interpretation of the experiment. To recap, the model of Bryan et al. (2014) has three main features: migration risk, a credit constraint that prevents households from borrowing to migrate, and individual learning about urban productivity. The model starts from pre-specified (i.e., non-equilibrium) initial conditions and offers a conditional migration transfer to all model households in the rural area. Rural households differ in permanent urban ability and their stock of assets, which they can accumulate through buffer stock savings. In contrast to our model, there is no disutility of migration, and there are no temporary productivity shocks. Learning in their model is permanent: no worker knows his or her productivity in the urban area until after migrating, and then the worker learns it forever.

Table 7: Comparison to Model of Bryan et al. (2014)

Panel A: Effect of Migration on Consumption		
	OLS	IV (LATE)
Data	10	30
Model	10	29
Model of Bryan et al. (2014)	57	52

Panel B: Effects of an Unconditional Transfer on Migration		
	Control	Treatment
Data	34	44
Model	36	35
Model of Bryan et al. (2014)	66	88

Note: This table reports moments of the experimental data and the predictions for the same moments in the current model and the model of Bryan et al. (2014). Panel A reports the values of the OLS and IV (LATE) returns to migration on consumption per capita, expressed in percentage points. Panel B reports the migration responses to an unconditional cash transfer in the control and treatment villages.

To compare their model’s prediction with the data, we assume a common CRRA parameter of two, which is the same value used in the current model and in the middle of the range of values explored by Bryan et al. (2014). We compare our model’s predictions with the data and their model, starting from their preferred initial conditions. A key difference between the two models’ predictions is in the constraints that hold back migration, which determine how migrants are selected in equilibrium. In the model of Bryan et al. (2014), many households would like to migrate, but they lack the credit or savings to do so. Their decision rule for migration is to migrate once disposable income reaches a certain threshold. A migration subsidy in that model induces migration by pushing up households over the threshold. In contrast, in our model, households wait until their prospects in the rural area are sufficiently bad for them to migrate, which means that households with lower income and asset levels migrate.

As one way to illustrate this, Panel A of Table 7 reports the effect of migration on consumption in the data and models, measured two different ways. The first way is the simple OLS coefficient from a regression of consumption on whether the household sent a seasonal migrant. The second is the LATE of migration on consumption, measured using an IV regression with migration instrumented using assignment to the treatment group. As the first two rows of Panel A show, the OLS coefficient on migration is substantially

lower than the IV coefficient both in the data and in the current model. As we discussed above, getting an OLS coefficient smaller than the IV coefficient is a key determinant of whether model households that choose to migrate are negatively or positively selected on income and assets. Bryan et al.'s (2014) model has a counterfactually high coefficient: 57 percent, compared with 10 percent in the data. Its IV coefficient is also too high: 52 percent, compared with 30 percent in the data. Perhaps most importantly, its OLS coefficient is counterfactually higher than the IV coefficient.

Another way of comparing the migration incentives in the two models is by considering the effects of an *unconditional* transfer that Bryan et al. (2014) conducted as a component of the original experiment, on a smaller subset of villages. Panel B of Table 7 reports the effects of this unconditional transfer in the data and in the two models. The top row reports the seasonal migration rates in the control and treatment groups, their simple difference, and the standard error of the difference. In the data, the migration rates were 34 percent in the control villages and 44 percent in the villages with the unconditional transfers, for a difference of 10 percentage points. The standard error of this difference is 6.5 percent, however, and the p -value is 0.24, meaning that this estimated effect is statistically insignificantly different from zero at any conventional significance level.

Our model predicts that unconditional transfers induce a negligible decreased in migration of around one percent. The model of Bryan et al. (2014) predicts a counterfactually large increase in migration rates of 22 percentage points. This substantial increase again reflects the constraints on migration in that model — namely, that households cannot save or borrow — and the migration decisions that involve migration once assets are sufficiently high. The unconditional transfer helps the agents in that model get above the threshold asset level needed to make migration worthwhile.¹⁵

5. The Welfare Effects of Conditional Migration Subsidies

Given that the model does well in matching the salient features of the data, we feel confident in using it to measure the welfare implications of encouraging rural-urban migration through migration subsidies. To do so, we compute welfare as the percent increase in

¹⁵The model of Bryan et al. also emphasizes subsistence consumption and permanent learning about migration ability, though in hindsight, neither of these features seems central to the interpretation of the experimental evidence in question. As we show in Appendix Tables A.6, the predictions of the model of Bryan et al. are similarly counterfactual with and without subsistence constraints. Furthermore, the interpretation of the experimental evidence reached by the current model is not substantially changed once we add subsistence constraints (see Appendix Tables A.7 and A.8). As for permanent learning about migration ability, in Appendix Table A.9 we show that the model of Bryan et al. makes counterfactually high predictions for repeat migration rates.

consumption, p , that makes the household indifferent between a p -percent consumption increase in perpetuity and being offered the conditional migration subsidy.

5.1. One-Time Migration Subsidies in Partial Equilibrium

We begin by computing the welfare gains of a one-time migration subsidy for rural households with low assets. We treat the migration subsidies as unanticipated, available only to the bottom half of the rural asset distribution, and in partial equilibrium, meaning that rural wages are held fixed. We assume, as in the experiments, that the funding for the transfers comes from outside the economy. As a frame of reference, we also compute the welfare gains of a one-time unconditional transfer to this same population. To compare apples to apples, we choose the size of the transfers to each household so that the total cost of the transfers equals the total cost of the migration subsidies.

Table 8: Welfare Effects of One-Time Migration Subsidies

		Migration Subsidy		Migration Subsidy		Unconditional Transfer	
		Migration Endogenous		Migration Policy Fixed		Migration Endogenous	
		Welfare	Migr. Rate	Welfare	Migr. Rate	Welfare	Migr. Rate
Income Quintile	1	1.17	85	0.77	48	1.05	45
	2	0.45	63	0.31	38	0.56	37
	3	0.29	52	0.20	34	0.40	33
	4	0.20	46	0.15	31	0.32	31
	5	0.12	40	0.10	31	0.20	31
<u>Average</u>							
Rural & Low Assets		0.44	57	0.30	36	0.51	35
All Rural		0.22	41	0.15	31	0.25	30

Note: The first two columns report the lifetime consumption-equivalent welfare gains and migration rates for rural assets with low assets from one-time conditional migration subsidies. The next two columns report the same when the migration policies are held fixed for every agent. The final two columns report the welfare gains and migration rates from a one-time unconditional transfer costing the same total amount as the migration subsidies. The rows are for different income quintiles of the rural households eligible for the subsidy, with 1 being the poorest and 5 being the richest. All three experiments are in partial equilibrium, meaning that the rural wage is held fixed, and without financing the subsidies in equilibrium.

Table 8 reports the welfare gains from the migration subsidy and unconditional transfer. For expositional purposes, we report the average welfare gain and seasonal migration rate by income quintile among rural households eligible for the subsidy (with one being the lowest). The first two columns of Table 8 focus on migration subsidies with endogenous

migration responses, which represent the migration subsidies described in the field experiments of Section 2. Several features of these two columns are worth noting. First, the welfare gains for the one-time migration subsidy are modest on average, at 0.44 percent consumption equivalents. Second, gains are higher for the poorer quintiles and reach as high as 1.17 percent consumption equivalents for the poorest quintile. Third, migration rates are higher for the poorer quintiles, reaching as high as 85 percent for the poorest quintile. The average migration rate of 57 percent corresponds to the migration rate of the treatment group of the original migration experiment.

The middle two columns of Table 8 report the welfare effects of the conditional migration transfers where we make the transfers to all households willing to migrate given the transfer but counterfactually hold the migration policies of each household fixed at pre-transfer levels. The purpose of this exercise is to illustrate how many of the gains occur just from making the transfers to the households willing to undergo the ordeal of seasonal migration in order to obtain the transfer. The welfare gains from this exercise are 0.30 consumption equivalents on average and rise up to 0.77 consumption equivalents for the poorest quintile. This implies that many of the welfare gains from the migration subsidies — around two thirds of them — arise just from making the transfers to the households willing to take them. Notice that the actual migration rates in this exercise are identical to the equilibrium without transfers, highlighting that none of the welfare gains here arise from actually migrating.

The last two columns report the welfare effects of making unconditional transfers to the same set of households, where the total amount of the transfers equals the total cost of the migration subsidies in the previous two exercises. Notice that this implies a smaller transfer to each household, since every eligible household gets the unconditional transfer, whereas only 57 percent of eligible households take up the migration subsidy. Overall, the unconditional transfers lead to modestly higher average welfare gains than the migration subsidies, at 0.51 percent consumption equivalents. The poorest households gain less, since they get a smaller transfer, and the richest gain more, signaling that the transfers are less precisely targeted to needier households. Interestingly, migration rates are far lower under the unconditional transfers, meaning that few eligible households actually want to migrate when given the choice not to do so. This is consistent with the evidence (of Section 4.9) on how an unconditional transfer had a modest effect on migration.

5.2. Permanent Migration Subsidies in General Equilibrium

We now consider the welfare gains from permanently offering migration subsidies to rural households with low assets, using the same asset availability criteria as for the temporary transfers above. We simulate offering these policies in general equilibrium — meaning that we allow rural wages and population shares to adjust — and financing the transfers through taxation of labor income on all households. We report the welfare gains on average for four subgroups of the population: rural households with low assets (i.e., those eligible for the subsidies), all rural households, all urban households, and all households.¹⁶

Table 9 reports the welfare gains for these population subgroups and other key statistics. We first report the welfare gains of the permanent migration subsidies abstracting from taxation, in the first column of the table. Rural households with no assets gain the most, at 2.06 percent consumption equivalents. However, all rural households gain nearly as much, at 1.62 percent, suggesting that the availability of migration subsidies in the future is still valuable even for households not migrating in the current period. Urban households have a modest gain from occasionally accessing the subsidies in some future period, and all households gain by just under 1 percent in consumption equivalents.

The second column of Table 9 adds back in the financing of subsidies but holds migration decisions fixed ex-post. This setting is not meant to capture an actual subsidy one could implement in practice. Instead, it illustrates the welfare gains just from targeting transfers to households willing to undergo the ordeal of migration in exchange for the subsidy. The welfare gains in this case are only somewhat lower than those in the first case, at 1.62 percent. In other words, most of the welfare gains arise from targeting transfers rather than from changes in income after migrating. Urban households now lose -0.29 percent in welfare equivalents, since they pay taxes but benefit little from the subsidies themselves.

The third column of Table 9 reports the full welfare effects of permanent migration subsidies in general equilibrium. The welfare gains are now 2.26 percent for the poor rural households with low assets. Urban households now lose even more, at -1.26 percent, while all households gain 0.80 percent on average. Interestingly, when the migration transfers are permanently offered, household location and saving decisions change substantially. In particular, the share of households living in the rural area rises from 60 to 66 percent, and the share of rural households having low assets (i.e., below the original asset cutoff) rises from

¹⁶A limitation of this analysis is that the experimental evidence is not well suited to measure impacts on urban wages. To the extent that capital, land and other non-labor factors of production are fixed in urban areas, large changes in urban shares of the population could depress urban wages, all else equal. If there are agglomeration externalities, larger urban populations could result in positive externalities on the wages of urban residents.

Table 9: Welfare Effects of Permanent Migration Subsidies

	Migration Fixed No Taxation	Migration Fixed Tax Financed	Migration Endogenous Tax Financed (G.E.)
Rural & Low Assets	2.06	1.62	2.26
All Rural	1.62	1.19	1.85
All Urban	0.15	-0.29	-1.26
All Households	1.03	0.59	0.80
% in Rural Area	60	60	66
% of Rural Seasonally Migrating	31	31	56
% of Rural with Low Assets	50	50	74
Tax Rate (% of labor income)	0	0.4	1.3

Note: The first column reports the effects of permanently offering conditional migration subsidies to rural households with sufficiently low assets, as in the migration experiments, but with migration policies held fixed and without any taxation to pay for the transfers. The second column is the same, but the migration subsidies are financed through labor taxation. The third column allows migration to be endogenous and finances the transfers through labor taxation.

50 to 72 percent. This raises the overall take-up of migration subsidies, which increases the labor tax rate required from 0.4 percent (in the case of fixed migration decisions) to 1.3 percent.

The welfare calculations of this section also highlight which behavioral and equilibrium responses are important considerations when considering permanently subsidizing migration, rather than just subsidizing it one time. Both the large increases in seasonal migration rates and the overall increase in rural population shares play primary roles in driving the overall effects of a permanent subsidy scheme. General equilibrium impacts working through wages are modest, as rural wages rise by only around 3 percent in lean periods. The overall increase in tax burden is not enormous but not trivial either, especially since most low-income countries have limited capacity to raise tax revenues. An increased number of rural households with low asset levels means that households are substituting self-insurance through asset holdings for the public insurance offered by migration subsidies. This again suggests that insurance motives are behind most seasonal migration decisions.

6. The Social Planner's Allocation

This section formulates and solves the efficient allocation in our model economy. Up to this point, we have measured the welfare gains from particular policy interventions aimed

at increasing seasonal rural-urban migration. However, these experiments leave open the normative question of how much households *should* be migrating in an efficient allocation. In other words, how would a benevolent social planner choose migration rates in this setting, and how would that compare to actual migration outcomes?

6.1. The Social Welfare Function

To study these normative questions, we take a stand on the social welfare function. We focus on a planner with Pareto weights $\psi(z)$ that vary with the permanent productivity of an individual. The idea is to (potentially) restrict how much redistribution there is across households of different permanent ability levels while having the planner treat households with the same z identically and in a utilitarian manner. We define the social welfare function as:

$$\mathcal{W}^{SP} = \sum_{t=0}^{\infty} \sum_{j'} \sum_j \int_z \int_s \int_x \int_{\nu} \beta^t \psi(z) u_{j',j}(c_j(z, s, x, t), x, \nu) \lambda_j(z, s, x, \nu, t) d\nu dx ds dz. \quad (14)$$

Here social welfare is the weighted average value of household utility across locations j and going to location j' , productivity states z and s , experience x , and preference shocks ν . The average is computed with respect to the Pareto weights, $\psi(z)$, and the measure of households $\lambda_j(z, s, x, \nu, t)$ with those shock states, experience levels, and preference shocks at all dates t . Utility depends directly upon the consumption allocation $c_j(z, s, x, t)$, but also directly on the location j through the migration disutility \bar{u} and the idiosyncratic preference shocks.

We formulate the social planner's problem so that her choice variables are the consumption allocations and migration probabilities for all households. To cast the problem in terms of migration probabilities, we integrate out the preference shocks conditional on a set of migration probabilities for each household state. These migration probabilities prescribe an assignment of those households with the largest relative preference shock to migrate or not. So given set of states j, z, s, x, t , utility is

$$u(c_j(z, s, x, t), x) + E[\nu \mid \{\mu_{j',j}(z, s, x, t)\}_{j'}], \quad (15)$$

where $\mu_{j',j}(z, s, x, t)$ is the migration probability going from location j to location j' and then $E[\nu \mid \{\mu_{j',j}(z, s, x, t)\}_{j'}]$ is the expected value of the preference shock conditional on the migration probabilities. So, for example, if the Planner dictates that all people migrate from location j to location j' , then this value is the unconditional mean of a Type-1 extreme

value shock. Now we can re-write the social welfare function as

$$\sum_{t=0}^{\infty} \sum_j \int_z \int_s \int_x \beta^t \psi(z) \left\{ u(c_j(z, s, x, t), x) + E[\nu \mid \{\mu_{j',j}(z, s, x, t)\}_{j'}]} \right\} \lambda_j(z, s, x, t) dx ds dz. \quad (16)$$

6.2. The Law of Motion and Feasibility

The planning problem maximizes (16) subject to the law of motion describing how the population evolves across states and locations and then how many resources are available — that is, feasibility. Below, we describe each of these aspects of the environment.

Law of Motion. The law of motion describing how the measure of households evolves across states and locations is

$$\begin{aligned} \lambda_j(z, s', x', t+1) &= \int_s \int_x \mu_{j,j}(z, s, x, t) \pi(s', s) \varphi(x', x, j) \lambda_j(z, s, x, t) dx ds \\ &+ \sum_{j' \neq j} \int_s \int_x \mu_{j,j'}(z, s, x, t) \pi(s', s) \varphi(x', x, j') \lambda_{j'}(z, s, x, t) dx ds, \end{aligned} \quad (17)$$

where $\pi(s', s)$ describes the transition across transitory states and $\varphi(x', x, j')$ describes the transition in experience. This equation says that given the current distribution $\lambda_{j'}(z, s, x, t)$ in location j' , the measure of households $\lambda_j(z, s', x', t+1)$ reflects the migration probabilities of households in each location (the μ 's), how their productivity evolves over time (π 's), and how their experience $\varphi(x', x, j)$ evolves.

Labor Supply, Aggregate Production, and the Resource Constraint. Given a distribution of households, the effective labor units in the urban and rural area are

$$\begin{aligned} N_{u,t} &= \sum_{j=[\text{urban,seas}]} \int_z \int_s \int_x z s^\gamma \lambda_j(z, s, x, t) dx ds dz, \quad \text{and} \\ N_{r,t} &= \int_z \int_s \int_x s \lambda_{\text{rural}}(z, s, x, t) dx ds dz. \end{aligned} \quad (18)$$

The urban area includes the seasonal and permanent urban workforce. Aggregate production of the final good is

$$Y_t = A_u N_{u,t} + A_{r,t} (N_{r,t})^\phi. \quad (19)$$

Combining the amount of resources available in (19) with the consumption and moving decisions, we have the following resource constraint:

$$Y_t \geq \sum_j \int_z \int_s \int_x c_j(z, s, x, t) \lambda_j(z, s, x, t) dx ds dz + \sum_j \sum_{j'} \int_z \int_s \int_x m_{j',j} \mu_{j',j}(z, s, x, t) \lambda_j(z, s, x, t) dx ds dz, \quad (20)$$

which says that production must be greater than or equal to consumption plus the moving costs associated with households migrating across locations. Here we compactly sum across all j' and j location pairs making use of the fact that the moving cost for staying in a location is zero; that is, $m_{j,j} = 0$.

6.3. The Planner's Problem and the Efficient Allocation

The Social Planner's Problem is the following:

$$\max_{c_j, \mu_{j',j}} \sum_{t=0}^{\infty} \sum_j \int_z \int_s \int_x \beta^t \psi(z) \left\{ u(c_j(z, s, x, t), x) + E[\nu \mid \{\mu_{j',j}(z, s, x, t)\}_{j'}]} \right\} \lambda_j(z, s, x, t) dx ds dz \quad (21)$$

subject to (17), (18), (19), (20) and an initial condition, $\lambda_j(z, s, x, 0)$. Here the planner is choosing consumption allocations and migration probabilities for each state and date and these allocations must respect the production technology, feasibility, the laws of motion and an initial condition. Finally, note that the planner only considers allocations that depend on the current state, not the entire history.

In spite of the seemingly complicated problem facing the social planner, we show that the solution to the planner's problem has an intuitive characterization which we formalize as Proposition 1 below.

Proposition 1 (Efficient Consumption and Migration) *The consumption and migration rates solving the social planner's problem are as follows. Consumption allocations for each type z equate the marginal utility of consumption in all locations, productivity and experience states for each date:*

$$u'(z, t) = u'(c_j(z, s, x, t)) = u'(c_{\bar{j}}(z, s', x', t)) \quad \forall j, s, x, \quad (22)$$

$$\text{and for all types } z \text{ and } z': \quad \frac{u'(z, t)}{u'(z', t)} = \frac{\psi(z')}{\psi(z)}. \quad (23)$$

Migration probabilities satisfy:

$$\mu_{j'j}(z, s, x, t) = \exp\left(\frac{-u'(z, t) m_{j',j} + \beta \mathbb{E}[\chi_{3j'}(t+1)|z, s, x]}{\sigma}\right) \Big/ \sum_{j'} \exp\left(\frac{-u'(z, t) m_{j',j} + \beta \mathbb{E}[\chi_{3j'}(t+1)|z, s, x]}{\sigma}\right), \quad (24)$$

with the multiplier $\chi_{3j'}(z, s, x, t+1)$ satisfying the following recursive relationship

$$\chi_{3j'}(z, s, x, t+1) = u_{j'}(z, s, x, t+1) + u'(z, t+1)\kappa_{j'}(z, s, x, t+1) + \beta \mathbb{E}[\chi_3(t+2)|z, s, x], \quad (25)$$

where

$$\kappa_{j'}(z, s', x', t+1) = mpl_{j'}(z, s', t+1) - c_{j'}(z, s', x', t+1) - \sum_{j''} m_{j'',j'} \mu_{j'',j'}(z, s', x', t+1). \quad (26)$$

Proposition 1 has a lot of content to unpack. First, consumption allocations equalize the marginal utility of consumption each permanent productivity type z across all locations, productivity, and experience states for each date t . This is a standard optimality condition ensuring that there is no way to better allocation consumption across states for a given z type. Second, the ratio of marginal utilities for two different permanent productivity types z and z' are inversely proportional to the weights assigned to them by the planner. This allows the planner to potentially allocate more consumption to agents with higher permanent productivity, for example, as in Negishi (1960).

Note that these optimality conditions do not imply that the planner equates the level of consumption across households with a given permanent productivity z . In our model, a household's experience and the disutility of being in the urban area affect the households' marginal utility of consumption. Mechanically, this necessitates that the planner compensates non-experienced rural-urban migrants with higher consumption to equate the marginal utility of consumption across all households of a particular z type.

Two observations follow from this one. First, when moving households across space, the planner must factor in the differential in the consumption cost associated with moving an inexperienced rural household to the urban area, not just the moving cost. Second, and more broadly, gaps in consumption (either in the cross-section or when households move, as in the migration experiments) arise in an efficient allocation because of migration disutility. In other words, consumption gaps do not necessarily imply the allocation is inefficient or that there is scope for gains to reallocate households across space.

While the migration probabilities (24) take on a familiar exp sum form that is facilitated by

the Type 1 extreme value shocks, the way the migration probabilities connect with fundamentals in (25) and (26) is novel.

The destination-specific component of the migration probabilities reflects two components that capture the social cost and benefits of a move to that destination. The first is the cost of a move which shows up directly but is evaluated at the marginal utility of consumption $-u'(z, t) m_{j',j}$. So if the migration cost from j to j' is high, then fewer people move. The second component is the expected value of the multiplier for a given destination, $\chi_{3j'}(z, t + 1)$, with the expectation taken across the shock states and experience. This multiplier takes on a recursive formulation in (25), where the multiplier today reflects current utility, the κ term (which we explain below), and the expected discounted value of the multiplier tomorrow, which takes into account all possible moving options out of location j' .

This term in (25) describes, on net, how much a household receives from the planner versus what is produced in return. What it receives is the first term in (25): utility evaluated at the optimal consumption allocation and the expected preference shock. Then the κ term in (26) reflects the net resources produced by the household; that is, it is the marginal product of labor net of consumption received and net expected moving costs—all evaluated at the marginal utility of consumption.

All of this gives rise to intuitive features of the allocation. The planner moves households expected to have a high marginal product of labor in location j . The planner is unlikely to move a household if it requires relatively more consumption in that location (which would arise when the household is inexperienced). The planner also takes into consideration the moving cost associated with that location, including the term $\sum_{j''} m_{j'',j'} \mu_{j'',j'}(z, s', x', t + 1)$, which is a migration-weighted average of the moving costs. This term thus indexes how hard it is to move a household out of that location in the future. If a location features high migration costs, the planner moves fewer people in.

How does this allocation relate to the competitive equilibrium? We know that with incomplete markets, the marginal utility of consumption in the competitive equilibrium allocation will not be properly equated across households of a given z level. Nor will marginal utilities between households of different z values necessarily be proportional to the planner's weights. What is less obvious is how the planner will locate households across space and how her seasonal migration decisions will compare to those of the competitive equilibrium. We explore these issues next.

Table 10: Competitive Equilibrium vs. Social Planner's Solution

	(1)	(2)	(3)	(4)
	Competitive	Planner,	Planner,	Difference
	Equilibrium	Migration Fixed	Migration Endogenous	(3) - (2)
Average Welfare Gain (%)	-	46.6	48.5	1.9
% in Rural Area	60	60	54	-6
% Seasonally Migrating	31	31	27	-4
Rural-Urban Wage Gap	1.88	1.88	1.70	-0.18

Note: The first column reproduces features of the competitive-equilibrium outcome. The second column reports the welfare gains from moving to an allocation with full consumption insurance, but holding migration policies fixed, and the features of this allocation. The third column reports the welfare gains from moving to the efficient allocation, with full consumption insurance and allowing for endogenous migration decisions.

6.4. The Efficient Allocation versus the Competitive Allocation

This section computes the efficient allocation using the same parameters as in Section 4 and compares the results to the competitive allocation. As in our competitive equilibrium, we focus on stationary equilibria that give rise to time-invariant consumption allocations and migration probabilities. Here we present the planner's problem and summarize key features of its solution; Appendix B presents the planner's solution in full detail.

We assume that the planner's weights take the form $\psi(z) = \exp(\alpha/z)$, where α captures the relative weight the planner puts on households with lower z values. In practice we find that our broad conclusions in this section do not depend on our choice of α (see Appendix B). Thus, in what follows we focus on the case of $\alpha = 0$, which corresponds to equal weights on all households. We measure welfare as the permanent percent increase in consumption that a household would require to be indifferent between the competitive equilibrium allocation and the planner's allocation.

Table 10 reports the results. The first column reproduces some moments from the competitive equilibrium from the baseline economy. Our main interest is in how the migration policies of the planner differ from migration outcomes under the competitive equilibrium. To this end we report, in the second column in Table 10, the welfare gains of implementing the planner's consumption allocation but fixing the migration policies to be those of the competitive equilibrium allocation. In other words, we hold the size of the pie the same as in the competitive allocation, and change only how the pie is divided up across households. Welfare gains from this second column arise from providing insurance for households of




Figure 2: Migration Probabilities: Competitive Equilibrium vs Planner’s Solution

each z type and then from redistributing from low- to high- z types. The welfare gains are large, at 46.6 percent, but not our focus, as these gains are a consequence of known sources of gains from the redistribution across permanent types and completing markets.

Our main focus is on how much the planner can improve welfare by better allocating workers across space. The third column helps answer this question by computing the full efficient allocation including endogenous location and migration choices for each household. The difference between the third column and second column thus isolate the gains from a better allocation of workers across regions relative to the competitive allocation. Interestingly, the planner chooses an allocation that induces a smaller rural population (54 percent versus 60 percent in the competitive allocation) and has less seasonal migration than in the competitive economy (27 percent versus 31 percent in the competitive allocation). Moreover, the planner chooses a smaller – but still large – wage gap (as implied by the marginal product of labor) between rural and urban areas. The incremental welfare gains from changes in migration relative to the full insurance allocation are modest, at 1.9 percent (48.5 percent minus 46.6 percent). This suggests fairly modest gains from better allocating workers across space, and that in an efficient allocation there would be less, rather than more, seasonal migration.

To better understand these results, Figure 2 plots the seasonal migration policy function

by transitory shock and asset holdings for low z households in the competitive allocation (left panel) and planner's allocation (right panel). One can see that in the competitive equilibrium, many migrants move out of "desperation." That is, there is a lot of seasonal migration by households with low transitory shocks and low assets. From the planner's perspective, these moves are socially wasteful, and they are virtually eliminated in the planner's allocation. From the planner's perspective, giving these households insurance and keeping them in the rural area is the efficient thing to do. In other words, the planner chooses lower seasonal migration rates because she eliminates these desperation moves.

The outcome from the planner's solution suggests that there is not much spatial misallocation overall. At best the planner could raise welfare by around 1.9 percent by better allocating workers across space, and even then, this would entail less seasonal migration, not more. The planner's solution suggests that the real issue in the model is the lack of insurance and how households use seasonal migration as a device to raise incomes in periods when their asset holdings and productivity levels are low. In other words, the planner gains very little from releasing labor that is stuck in rural areas but does gain from reducing moves of desperation among vulnerable rural households.

How does the planner's solution relate to the permanent migration transfers discussed in the previous section? The results of this section show that such a policy is really "second best," in that it pushes migration in the opposite direction as the planner. However, the policy is still welfare enhancing because the costs of distorting migration are small relative to the gains from providing insurance even indirectly via migration. So we see substantial gains from the permanent migration subsidies on average. This is consistent with the result from Table 10 that most of the welfare gains from moving to the planner's problem come from equalizing marginal utilities of consumption, rather than from better allocating workers across space.

7. Conclusion

This paper studies the welfare implications of subsidizing rural-urban migration in low-income countries. Cross-sectional data show that wages are much higher in urban areas than in rural areas, and recent experiments show that subsidies for seasonal migration raise the income and consumption of migrants. It is thus tempting to conclude from this evidence that many rural workers are stuck in poverty traps in which credit constraints and income risk keep them from the higher average wages of cities.

Our analysis, which uses a dynamic model of migration estimated to match this cross-sectional and experimental data, suggests that this is not the correct interpretation. Rather

than migration being deterred by risk, we argue that households use migration as a way to insure themselves against states of the world in which their productivity and asset holdings are low. The welfare gains from subsidizing seasonal rural-urban migration thus come about largely from providing better insurance opportunities for rural households in periods when they are vulnerable.

The large non-monetary costs of seasonal migration that we estimate in our model suggest scope for future research into the policies that can reduce the costs of rural-urban migration for vulnerable rural individuals. We provide suggestive evidence that the migration disutility is related in part to poor housing conditions for new migrants in urban areas (Appendix C). To the extent that policy makers can influence the living conditions for migrants arriving in urban centers, this could raise welfare by creating better opportunities for villagers to insure themselves by periodically moving to the city. Future research should look more broadly into the non-monetary and monetary barriers to seasonal migration.

In terms of methodology, our paper departs from the previous macroeconomic literature in how we discipline our model quantitatively and, in particular, in how we replicate a randomized controlled trial within a macroeconomic model. Our method of combining a dynamic incomplete-markets model with experimental data can be used more broadly to study other macroeconomic phenomena — such as savings behavior, labor market search activity, or investments in new technologies — that have been the focus of recent randomized experiments.

References

- AIYAGARI, S. R. (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 109, 659–84.
- AKCIGIT, U., H. ALP, AND M. PETERS (2021): “Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries,” *American Economic Review*, 111, 231–275.
- AKRAM, A. A., S. CHOWDHURY, AND A. M. MOBARAK (2017): “General Equilibrium Effects of Emigration on Rural Labor Markets,” Manuscript, Yale University.
- BAZZI, S. (2017): “Wealth Heterogeneity and the Income Elasticity of Migration,” *American Economic Journal: Applied Economics*, 9, 219–255.
- BEWLEY, T. (1977): “The Permanent Income Hypothesis: A Theoretical Formulation,” *Journal of Economic Theory*, 16, 252–292.

- BROOKS, W. AND K. DONOVAN (2020): "Eliminating Uncertainty in Market Access: The Impact of New Bridges in Rural Nicaragua," *Econometrica*, 88, 1965–1997.
- BRYAN, G., S. CHOWDHURY, AND A. M. MOBARAK (2014): "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh," *Econometrica*, 82, 1671–1748.
- BRYAN, G. AND M. MORTEN (2019): "The Aggregate Productivity Effects of Internal Migration: Evidence From Indonesia," *Journal of Political Economy*, 127, 2229–2268.
- BUERA, F. J., J. P. KABOSKI, AND Y. SHIN (2014): "Macro-Perspective on Asset Grants Programs: Occupational and Wealth Mobility," *American Economic Review Papers and Proceedings*, 104, 159–164.
- BUERA, F. J. AND Y. SHIN (2013): "Financial Frictions and the Persistence of History: A Quantitative Exploration," *Journal of Political Economy*, 121, 221–272.
- CALIENDO, L., M. DVORKIN, AND F. PARRO (2019): "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock," *Econometrica*, 87, 741–835.
- DAVID, J., H. HOPENHAYN, AND V. VENKATESWARAN (2016): "Information, Misallocation and Aggregate Productivity," *Quarterly Journal of Economics*, 131, 943–1005.
- DE JANVRY, A., K. EMERICK, M. GONZALEZ-NAVARRO, AND E. SADOULET (2015): "Delinking Land Rights from Land Use: Certification and Migration in Mexico," *American Economic Review*, 105, 3125–49.
- DONOVAN, K. (2021): "The Equilibrium Impact of Agricultural Risk on Intermediate Inputs and Aggregate Productivity," *Review of Economic Studies*, 88, 2275–2307.
- FAJGELBAUM, P. D., E. MORALES, J. C. SUÁREZ SERRATO, AND O. ZIDAR (2019): "State Taxes and Spatial Misallocation," *Review of Economic Studies*, 86, 333–376.
- GOLLIN, D., D. LAGAKOS, AND M. E. WAUGH (2014): "The Agricultural Productivity Gap," *Quarterly Journal of Economics*, 129, 939–993.
- GREENWOOD, J., P. KIRCHER, C. SANTOS, AND M. TERTILT (2019): "An Equilibrium Model of the African HIV/AIDS Epidemic," *Econometrica*, 87, 1081–1113.
- GRÖGER, A. AND Y. ZYLERBERG (2016): "Internal Labor Migration as a Shock Coping Strategy: Evidence from a Typhoon," *American Economic Journal: Applied Economics*, 8, 123–153.

- GUNER, NEZIH, G. V. AND Y. D. XU (2008): "Macroeconomic Implications of Size-Dependent Policies," *Review of Economic Dynamics*, 11, 721–744.
- HAMORY, J., M. KLEEMANS, N. Y. LI, AND E. MIGUEL (2021): "Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata," *Journal of the European Economic Association*, 19, 1522–1555.
- HARRIS, J. R. AND M. P. TODARO (1970): "Migration, Unemployment and Development: A Two-Sector Analysis," *American Economic Review*, 60.
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2009): "Quantitative Macroeconomics with Heterogeneous Households," *Annual Reviews in Economics*, 1, 319–354.
- HERRENDORF, B. AND T. SCHOELLMAN (2018): "Wages, Human Capital, and Barriers to Structural Transformation," *American Economic Journal: Macroeconomics*, 10, 1–23.
- HSIEH, C.-T. AND P. J. KLENOW (2009): "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124, 1403–1448.
- HUGGETT, M. (1993): "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17, 953–969.
- IMBERT, C. AND J. PAPP (2019): "Short-term Migration, Rural Public Works and Urban Labor Markets: Evidence from India," *Journal of the European Economic Association*, 18, 927–963.
- (2020): "Costs and Benefits of Rural-Urban Migration: Evidence from India," *Journal of Development Economics*, 146.
- KABOSKI, J. P. AND R. M. TOWNSEND (2011): "A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative," *Econometrica*, 79, 1357–1406.
- KAPLAN, G. AND G. L. VIOLANTE (2010): "How Much Consumption Insurance Beyond Self-Insurance?" *American Economic Journal: Macroeconomics*, 2, 53–87.
- KENNAN, J. AND J. R. WALKER (2011): "The Effect of Expected Income on Individual Migration Decisions," *Econometrica*, 79, 211–251.
- KHANDKER, S. R. (2012): "Seasonality of Income and Poverty in Bangladesh," *Journal of Development Economics*, 97, 244–256.
- KLEEMANS, M. (2015): "Migration Choice Under Risk and Liquidity Constraints," Unpublished Working Paper, University of Illinois, Urbana-Champaign.

- LAGAKOS, D. (2020): "Urban-Rural Gaps in the Developing World: Does Internal Migration Offer Opportunities?" *Journal of Economic Perspectives*, 34, 174–192.
- LAGAKOS, D. AND M. E. WAUGH (2013): "Selection, Agriculture, and Cross-Country Productivity Differences," *The American Economic Review*, 103, 948–980.
- MEGHIR, C., A. M. MOBARAK, C. MOMMAERTS, AND M. MORTEN (Forthcoming): "Migration and Informal Insurance: Evidence from a Randomized Controlled Trial and a Structural Model," *Review of Economic Studies*.
- MIDRIGAN, V. AND D. XU (2014): "Finance and Misallocation: Evidence from Plant-Level Data," *American Economic Review*, 104, 422–458.
- MOLL, B. (2014): "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?" *American Economic Review*, 104, 3186–3221.
- MORTEN, M. (2019): "Temporary Migration and Endogenous Risk Sharing in Village India," *Journal of Political Economy*, 127, 1–46.
- NEGISHI, T. (1960): "Welfare Economics and Existence of an Equilibrium for a Competitive Economy," *Metroeconomica*, 12, 92–97.
- PETERS, M. (2020): "Heterogeneous Mark-Ups, Growth and Endogenous Misallocation," *Econometrica*, 88, 2037–2073.
- PIJOAN-MAS, J. (2006): "Precautionary Savings or Working Longer Hours?" *Review of Economic Dynamics*, 9, 326 – 352.
- RESTUCCIA, D. AND R. ROGERSON (2008): "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments," *Review of Economic Dynamics*, 11, 707–720.
- ROY, A. (1951): "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 3, 135–46.
- TODD, P. E. AND K. I. WOLPIN (2006): "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility," *American Economic Review*, 96, 1384–1417.
- YANG, M.-J. (2021): "Micro-level Misallocation and Selection," *American Economic Journal: Macroeconomics*, 13, 341–368.
- YOUNG, A. (2013): "Inequality, the Urban-Rural Gap, and Migration," *Quarterly Journal of Economics*, 128, 1727–1785.