The Welfare Effects of Encouraging Rural-Urban Migration

David Lagakos
UCSD and NBER

Mushfiq Mobarak
Yale University and NBER

Michael E. Waugh
New York University and NBER

January 2020

ABSTRACT

In this paper we study the welfare effects of encouraging rural-urban migration in the developing world. To do so, we build a dynamic incomplete-markets model of migration that allows for sorting on permanent comparative advantage, idiosyncratic productivity shocks and migration risk, plus migration disutility that depends on past migration experience. We estimate the model to replicate the results of a field experiment that subsidized migration in rural Bangladesh, leading to significant increases in seasonal migration rates and consumption for induced migrants. To match the experimental evidence, the model requires that migration subsidies are more likely to induce migration from those with relatively low productivity and asset levels, and that the non-monetary disutility of migrating is substantial. We conclude that the welfare effects of migration subsidies arise through better insurance for vulnerable rural households rather than by relaxing credit constraints for those with high urban productivity but who are stuck in rural areas.

Email: lagakos@ucsd.edu, ahmed.mobarak@yale.edu, mwaugh@stern.nyu.edu. 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, Rob Townsend, Fabrizio Zilibotti, seminar participants at Barcelona, Bristol, Chicago, Edinburgh, Einaudi, Fordham, Harvard, Hong Kong, NYU, St. Andrews, Stockholm IIES, UC Berkeley, UC Irvine, UNC, USC Marshall, Washington, Yale, Zurich, and participants at numerous conferences. For outstanding research assistance we thank Elizabeth Carls, Menaal Ebrahim, Patrick Kiernan, Seungmin Lee and especially Alejandro Nakab. For financial support we thank the International Growth Centre. 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). Misallocation of factors of production across firms, sectors or regions within an economy may underlie these TFP differences.[1] One potentially large source of misallocation is an inefficient distribution of workers across space (see e.g. Bryan and Morten, 2019; McMillan and Rodrik, 2011; Vollrath, 2009). This is highlighted by the large observed gaps in productivity and wages between rural and urban workers (Gollin, Lagakos, and Waugh, 2014; Herrendorf and Schoellman, 2018; Young, 2013). Such gaps also create a development puzzle: Why do large shares of the population in many developing countries continue to live in rural areas when urban areas within those same countries offer much higher wages?

A series of field experiments in Bangladesh show that paying small travel subsidies to induce rural Bangladeshis to migrate to urban areas leads to substantial gains in income and consumption over multiple years (Akram, Chowdhury, and Mobarak, 2017; Bryan, Chowdhury, and Mobarak, 2014). One interpretation of this experimental evidence is that households are spatially misallocated, and that encouraging migration would raise productivity by releasing rural households from a "poverty trap." 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 allow them to migrate to the city, where they are permanently more productive.

In this paper, we provide a reinterpretation of these migration experiments, and then use them to better understand how migration subsidies affect the welfare of rural households. We argue that seasonal migration serves as a form of insurance that rural households would like to draw on in states of the world when their productivity and asset holdings are low. However, high costs of migration – broadly defined – keep households from migrating more frequently, leaving them more vulnerable to bad shocks than they otherwise would be. Migration subsidies help direct resources toward rural households that have experienced bad shocks and who are willing to undergo the ordeal of migration in order to bolster their low consumption levels. We show that this interpretation is quantitatively consistent with the evidence behind the migration experiments and, while our model can entertain poverty-trap scenarios as in the model of Bryan et al. (2014), it can do so only at the cost of being quantitatively inconsistent with the experimental evidence.

[1] See 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, 2016); delegation and information frictions (Akcigit, Alp, and Peters, 2016; David, Hopenhayn, and Venkateswaran, 2016); entry barriers (Yang, 2016); and other government policies (Fajgelbaum, Morales, Suárez Serrato, and Zidar, 2019; Guner and Xu, 2008).

We formalize these arguments using a dynamic incomplete-markets model of migration. The model allows for migration risk following a long tradition in development economics (e.g. Harris and Todaro, 1970). In particular, households face deterministic seasonal income fluctuations and idiosyncratic productivity shocks, both of which are endemic to developing economies. 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 Herrendorf and Schoellman, 2018; Hicks, Kleemans, Li, and Miguel, 2017; Lagakos and Waugh, 2013; Young, 2013). Households can migrate either permanently or temporarily across locations, 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 high-quality experimental data, which is a methodological innovation relative to the prior literature.[2] In particular, we replicate the results of the migration experiments described above within our model, and we use simulated method of moments to estimate the model using the experimental data. Matching these moments helps us isolate the characteristics of workers who are near the margin – meaning those that can most easily be induced to migrate – relative to those who are unlikely to migrate, those already migrating regularly or those that are permanently located in cities. The 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 that have not migrated recently.

We show that our model is consistent with a number of experimental and cross-sectional moments, including several that are not targeted, while the model of Bryan et al. (2014) is not. Unlike that model, ours predicts – as a non-targeted moment – a negligible migration response to an unconditional transfer. This is empirically confirmed in a subsequent experiment conducted on this 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. In contrast, Bryan et al. (2014)'s model predicts, counterfactually, that only those with high asset levels can bear

[2] In terms of methodology, our work follows the seminal papers by Todd and Wolpin (2006) and Kaboski and Townsend (2011) that discipline dynamic structural models using quasi-experimental evidence rather than non-experimental moments, as is 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 (2019) who draw on a structural general equilibrium model to help study the impacts of bridges built quasi-randomly.

the risk of migrating. 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 use our estimated model to quantify the welfare gains from subsidizing rural-urban migration either temporarily or permanently, in partial equilibrium and in general equilibrium. The gains from a one-time partial-equilibrium migration subsidy are about 0.4 percent in consumption-equivalent welfare on average though substantially higher for the poorest households. By comparison, the welfare gains from an alternative calibration that has migration subsidies release households from a poverty trap are five times as high. Though this alternative calibration also makes counterfactual predictions on a number of dimensions. We calculate that when migration subsidies are offered permanently, the average welfare gains are 3.6 percent in consumption equivalents in partial equilibrium and only modestly lower in general equilibrium. The gains from permanent transfers are still higher for the poorest households, but more equally spread across the income distribution than under temporary transfers, because they offer regular insurance to households against times when they are vulnerable. These welfare calculations highlight the value of migration transfers as facilitating better insurance rather than relaxing credit constraints for those with relatively high productivity.

As a frame of reference, we also compare the distributional consequences of subsidizing rural-urban migration to two alternative rural-based policies that are implemented in some parts of the developing world and part of the policy debate quite broadly: unconditional cash transfers and rural "workfare" programs, such as India's massive rural employment guarantee (NREGA). We find that the average welfare gains from the three policies (costing the same amount) are similar. Yet the conditional migration subsidies are better than alternatives at targeting the neediest households because they create an *ordeal*: generally those who have faced recent negative shocks are the ones that can be induced by the subsidy to incur the disutility of migrating to the city.

We conclude by empirically investigating the source of the migration disutility that our estimated model suggests is so important. To do so, we conduct new discrete-choice experiments in which the same experimental sample of households used to estimate the model are asked to choose between hypothetical migration options varying in wage rates, labor-income risk, housing options at destination, and frequency of visits home to see family. This exercise points to substantial disutility associated with bad housing conditions at the destination. Offering improved housing with a proper indoor latrine increases migration propensity by 17.4 percentage points. This effect size is equivalent to the effect of increasing migration wages by 21 percent. While poor housing conditions for migrants do appear to be an important barrier to migration, they do not fully explain the non-monetary utility cost in our model. We conclude that future work should continue to investigate the frictions that reduce rural-urban migration, either seasonally or permanently.

2. The Migration Experiments: A Summary

In this section, we summarize the experimental results (Akram et al., 2017; Bryan et al., 2014) 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" called the *monga* during the three-month period between planting and harvest, when farmers mostly wait for the crop to grow, and labor demand falls. Landless laborers experience a drop in wages and employment opportunities as a result, and incomes fall by an estimated 50 percent or more, on average (Khandker, 2012). To cope, some households migrate to towns and cities during the lean season in search of employment.

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, using 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

[3]The 32-village control group is comprised 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 comprised 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 and refer to them as the "the treatment group" and compare their outcomes to those of the combined control group.

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 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.

Importantly for our model, this second experiment added random variation in the proportion of the landless population in the treatment villages that were provided migration subsidy offers simultaneously. This labor-market-level variation created labor supply shocks of different magnitudes in different villages, which provides 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 working 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 later in Section 5 that this model succeeds in matching the data while the model 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.

[4]In reality net migration from rural to urban areas was likely to have been positive over this period. We abstract from positive net migration in the model since it seems unlikely to change the model's basic interpretation of the 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

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

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 or not the household is an "inexperienced migrant." The households' problem is dynamic, and households discount the future at rate $0 < \beta < 1$.[5]

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 by, for example, developing a network of friends, potential employers, or 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 Section 7). 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 formally describe more in detail below. These taste shocks are independently and identically distributed across time and options and 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 taken for largely non-economic, idiosyncratic, reasons, such as marriage. Idiosyncratic taste shocks provide a simple way to capture these (partly) non-economic factors.

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

[5] We 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 poorer households. The data suggests the opposite pattern: 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.15 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.

z in the urban area, 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. Denoting s_t as the current shock, this shock evolves according to an AR(1) process (in logs):

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

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

To allow for 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, 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

$$Y_r = A_r^i N_r^{\phi}, \tag{2}$$

where N_r are the effective labor units working in the rural area, and $0 < \phi < 1$, so that there is a decreasing marginal product of labor in the rural area, and A_r^i is rural productivity indexed

[6]The 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.

[7]The assumption of perfectly correlated shocks across regions is not an especially realistic one. We make it only made for computational tractability. In Appendix Tables A.7 and A.8 we present the results of an alternative estimation of the model in which shocks are uncorrelated across regions and over time (which is the case when $\rho = 0$). The welfare effects of conditional migration transfers 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.

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:

$$Y_u = A_u N_u, \tag{3}$$

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 are the sum across efficiency units and, thus, 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

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

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

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

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

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

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 describing 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 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 or not the household is an inexperienced migrant, x, and, thus, whether or not 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 jointly, the two aggregate states 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:

$$v(a, r, s, \nu, x, i, N_r) =$$

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

where a household chooses among staying in the rural area, seasonally moving, and permanently moving. Influencing this choice is the value function associated with each option and the household's taste shock associated with each 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 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:

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

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:

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

There are several important points to take note of 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 or not 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 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 s', 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:

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

Again, influencing this choice is the value function associated with each option and the household's taste shock associated with each 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. Households staying in the urban area 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 the rural area, as determined by the γ parameter (see equation (5)). 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 or not the household suffers the disutility.

As 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. While in the following section, we allow the data to discipline the most important determinants, 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—i.e., 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 to 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, 2016; 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 both in 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 to 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 Zylerberg (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 with low assets, or high-productivity workers with high assets, is determined by the data.

4. Model Parameterization and Quantification

We now estimate the model using simulated method of moments. We draw on two sets of moments. The first are 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

[9]For 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.

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

We draw on eight moments from the migration experiment of Bryan et al. (2014) to discipline the model. These are (i) the variance of log consumption growth from before and after the lean season in the control villages (0.19); (ii) the percent 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 those that 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 assigning some parameter values directly; 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 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 the average half-yearly inflation rate in Bangladesh, 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 productivity in the lean season to the growing season, A_{rl}/A_{rg}, to be 0.5, consis-

[10]We experiment with different values of β and R in Appendix Tables A.5 and A.6, and find that our model's predictions are not substantively altered under alternative plausible choices.

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	$0.1 \times$ 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 of a year. m_T is chosen to equal 10 percent of average rural consumption.

tent 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 such 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.

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 already above. They are: (i) θ, the shape parameter controlling the urban individual 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 of 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:

$$\hat{g}_{c,i} = g_{c,i} + \upsilon_{r,i}, \tag{11}$$

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

$$\log \hat{y}_i = \log y_i + \log \upsilon_{u,i}, \tag{12}$$

where y_i is observed income of household i; y_i is actual income; and $\log \upsilon_{u,i}$ 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

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

such that simulated moments from the model match up with moments in the data. This is analogous to the generalized method of moments estimation, but we do not have closed-form representations of model moments. Thus, we solve the model and construct moments from simulated data. We compute asymptotic standard errors for each parameter as described in Appendix C. 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 percent 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 appropriate 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, they conducted their

Table 2: Moments Targeted in the Estimation

Moments	Data	Model
Control: Variance of rural log consumption growth	0.19 (0.03)	0.19
Control: Percent of rural households with no liquid assets	47 (1.13)	48
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)	70
Treatment: Seasonal migration relative to control	22 (2.39)	21
Treatment: Seasonal migration relative to control in year 2	9 (2.44)	4
Treatment: Consumption increase of induced migrants (LATE)	30 (9.67)	29
Urban-Rural wage gap	1.89 (0.18)	1.89
Percent 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.

baseline survey prior to the lean season; thus, we follow the same timing in the baseline sample selection and measurement for our model. Furthermore, they 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 or not 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 close to it, while one (the repeat migration rate) is somewhat lower in the model than in the data. Figure 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. As the figure shows, the model also does well in other years, capturing the declining pattern present in the data. By five years after the experiment, the difference in migration rates between the two groups is positive, but small in magnitude, in the model, at two percent, and statistically insignificant in the data.

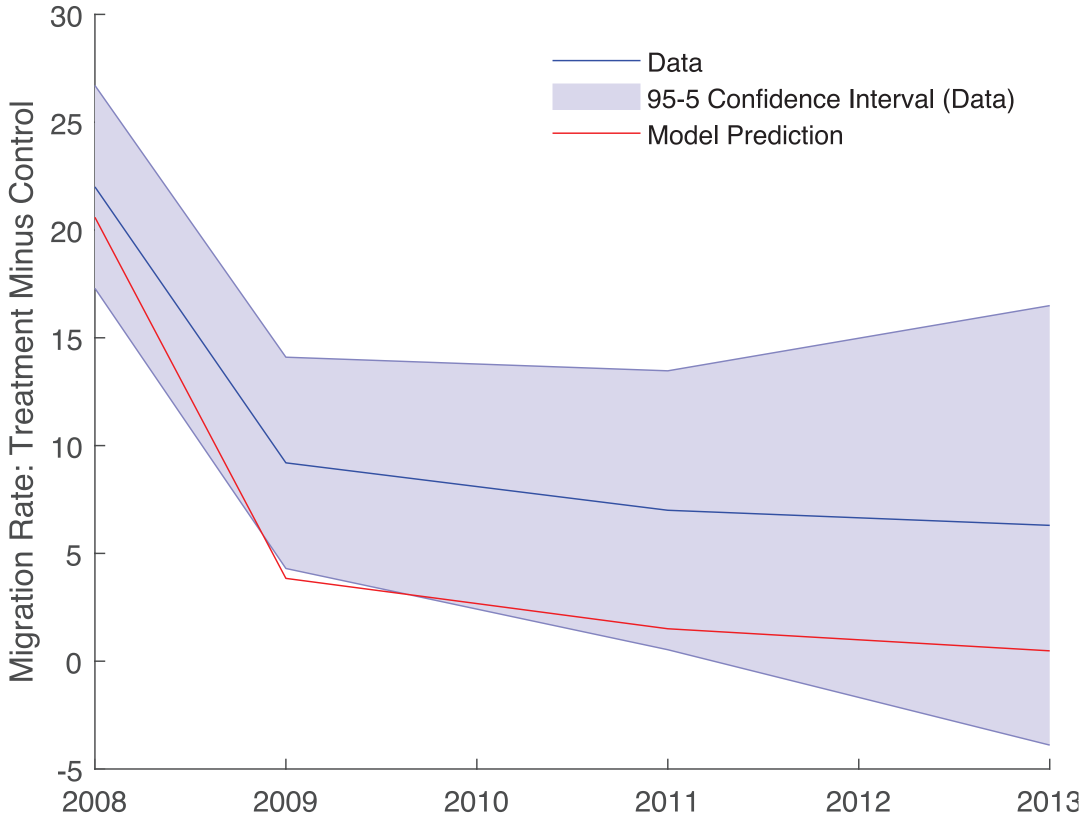

Figure 1: Difference in Migration Rates in Treatment and Control Groups

Table 3 shows the estimated parameter values and their standard errors. 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 parameter controlling permanent differences in ability, θ, is quite low, at 0.54. 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 1.51 for $\bar{u}$ implies that experiencing disutility of

Table 3: Estimated Parameters and Standard Errors

$\frac{1}{\theta}$	$\bar{u}$	λ	π	γ	A_u	σ_s	ρ	σ_ν	σ^2_{rc}	σ^2_{ui}
0.54	1.51	0.67	0.63	0.57	1.55	1.28	0.74	0.11	0.15	0.15
(0.002)	(0.004)	(0.054)	(0.028)	(0.002)	(0.022)	(0.866)	(0.217)	(0.010)	(0.005)	(0.004)

Note: The table reports the values of the 11 jointly-estimated parameters and their standard errors.

migration is equivalent, in a static sense, to cutting consumption by 33 percent. In Section 7, we use new survey data to help point to specific reasons why households suffer disutility of temporary migration. Next, the probability of remaining inexperienced, λ, as well as the probability of remaining experienced, π, are substantially less than one, with around one third of experienced households losing their experience each half-year, and only one third of inexperienced households gaining experience after a move. Finally, standard errors are in general quite small. The exceptions are the standard deviation and autocorrelations of the transitory shocks, σ_s and ρ, which are estimated fairly imprecisely.

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 these outcomes. Appendix D provides a more formal, quantitative discussion of which data moments identify which parameters.

We focus on rural households leading into the lean season since most migration occurs then. Figure 2 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 relative to low z households is the experimental migration response to the treatment.

This observation highlights an important implication about who migration subsidies may affect. Households sort themselves into rural and urban areas largely on the basis of permanent

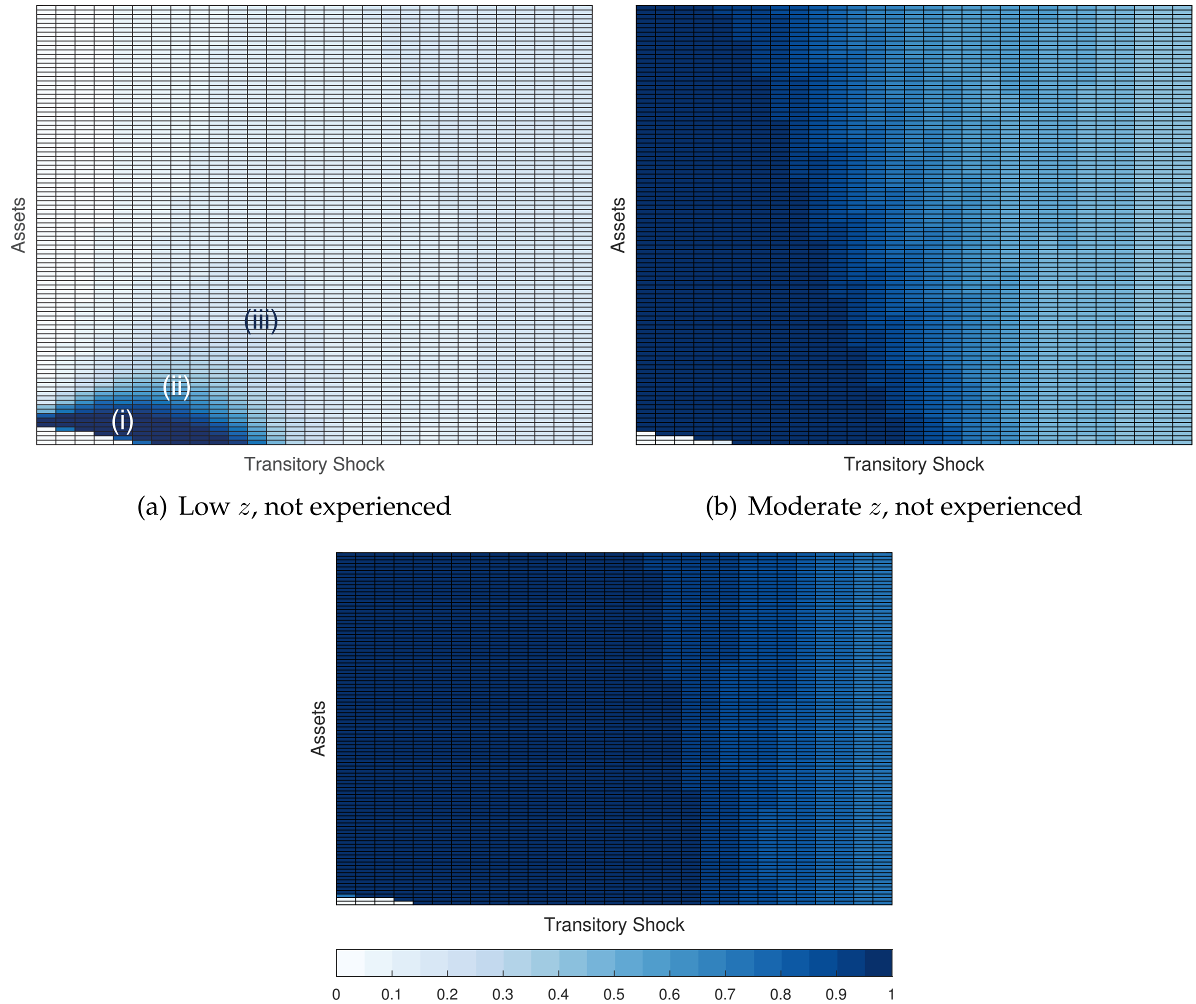

(a) Low z, not experienced

(b) Moderate z, not experienced

(c) Moderate z, experienced

Figure 2: Migration Probabilities for Select Household Types

comparative advantage. Thus, the set of households living in the rural area in any given period are most likely 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 a important deterrent to migration. Figure 2, panels (b) and (c), contrast the moving policy for the same z but different experience levels and, hence, 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.

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 is falling 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 can not 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 2 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. At first glance this may seem surprising, if one's expectation is that credit constraints are the primary reason that households do not migrate. If migration costs are high and the credit constraint binds, then the migration region would originate from the northeast corner in Figure 2, 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 moments 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 their consumption, specifically their transitory shock and asset holdings. Thus, the fact that the OLS coefficient is less than IV coefficient pushes our model to accommodate the idea households with low assets and low transitory shocks are more likely to migrate as in Figure 2.

This push is achieved through a large $\bar{u}$ parameter (as described above), but 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 asset households 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 due 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 how migration rates vary by initial consumption level speak 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).[11]

Migration rates by initial consumption level. Figure 5 plots the migration rates by quintile of the initial rural consumption level. Panel (a) reports migration rates from the data, and panel (b) reports the value in the estimated model (which are not targeted). As the Figure shows, migration rates are lowest in the poorest quintile but otherwise fairly similar across the quintiles. This suggests that migration in the data is, if anything, more likely for those with lower consumption levels to begin with. Our model gets this prediction largely correct.

Asset holdings by migration status. Table 5, Panel B, reports the average asset holdings of migrants and non-migrant households in the control villages of the experiment, relative to their average monthly consumption. As the Table shows, assets relative to income are modestly lower for the migrants than for the non-migrants in the control villages. The model also gets this feature right, though with a level of asset holdings that is a bit too low relative to the data.

Effects of unconditional transfers. Table 5, Panel C, reports the migration response of an unconditional transfer in the data, which was conducted along with the original migration experiments. Households did respond positively to an unconditional transfer, but to a much smaller extent than the conditional transfers. Though 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

[11] In Appendix B we formalize a notation of misallocation in a simplified version of the model in which we can analytically characterize the social planner's solution and compare it to the market allocation.

Table 4: 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 those that migrate (do send a migrant).

higher migration. Thus, it is fair to say that the model predicts a smaller effect of an unconditional transfer on migration than a conditional transfer, as in the data, and an effect that is small overall, and within the confidence intervals of the model's prediction.

Variances of consumption growth by migration status. Table 4 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, at 0.18, for those that migrate. The model is similar, with 0.18 for the stayers and 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 than for 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.

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

We now compare our model to the model proposed by Bryan et al. (2014) and show that their model is quantitatively inconsistent with the experimental evidence, and, as a result, 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 which 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. Unlike in our model, there is no disutility of migration and no temporary productivity shocks. Learning in their model is permanent: no worker knows their productivity in the urban area until after migrating, and then they learn it forever.

To compare their model's prediction to 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 to the data and to 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 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 households up 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 5 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 local average treatment effect (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 in both the data and in the current model. As we discussed above, getting a smaller OLS coefficient than IV coefficient is a key determinant of whether model households that choose to migrate are negatively or positively selected on income and assets. The model of Bryan et al gets the OLS coefficient counterfactually high, at 57 percent compared to 10 percent in the data. It also gets the IV coefficient too high, at 52 percent compared to 30 percent in the data. Perhaps most importantly, it gets the OLS efficient counterfactually higher than the IV coefficient.

Figure 5 provides an alternative way to see how migration is determined in the two models and in the data. All three subfigures plot migration rates by consumption quintile in the control group (light blue), the treatment group (medium blue), and the simple difference between the treatment and control groups (dark blue). Panel (a) represents the data, panel (b) the predictions of the model and panel (c) the predictions of the model of Bryan et al. In the data, one can see that migration rates are higher in the treatment than in the control group across all consumption quintiles. The differences are similar across quintiles, and somewhat larger in the lowest quintile.

Table 5: 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: Assets Relative to Average Monthly Consumption

	Migrants	Non-Migrants
Data	0.7	0.8
Model	0.4	0.5
Model of Bryan et al (2014)	1.1	0.5

Panel C: Effects of an Unconditional Transfer on Migration

	Control	Treatment
Data	34	44
Model	36	37
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 average asset holdings by migration status, expressed relative to average monthly consumption. Panel C reports the migration responses to an unconditional cash transfer in the control and treatment villages.

Subfigures (b) and (c) of Figure 5 show the stark differences between the predictions of two models. In the current model, migration rates are higher in the treatment than in the control across all five quintiles, and somewhat larger in the lowest quintile. This parallels the data closely. In the model of Bryan et al, migration rates are identical in the highest two quintiles, and counterfactually high at one-hundred percent. Thus, it is those with the highest consumption levels that migrate, as these households are following a simple rule of migration once assets are sufficiently high (at which point consumption levels are high as well). The reason the treatment has an effect for the other quintiles is that the migration subsidy pushes other households up over the threshold. However, as one can plainly see in panel (a) of Figure 5, this is not an empirically accurate depiction of how households make migration decisions in the data.

Panel B of Table 5 reports the average asset holdings of migrants and non-migrants from the control group of the experimental data, and the corresponding predictions from the two models. Asset holdings are expressed as a ratio of average monthly consumption (across the en-

Figure 3: Migration Rate by Consumption Quintile

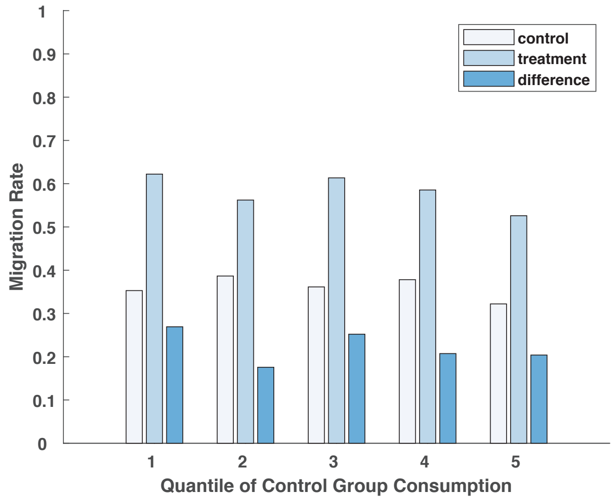

(a) **Data**

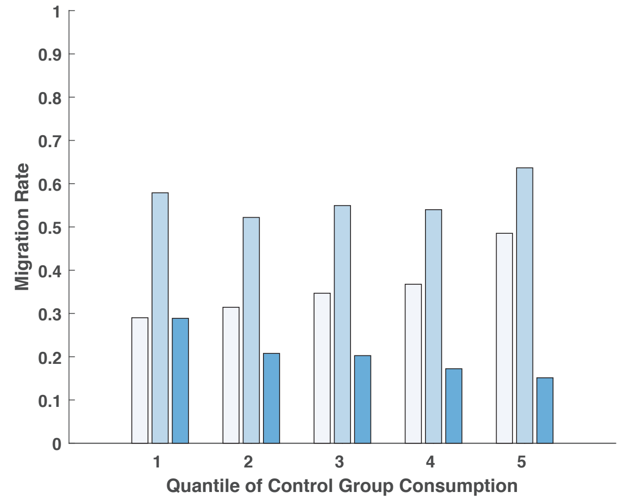

(b) **Baseline Model**

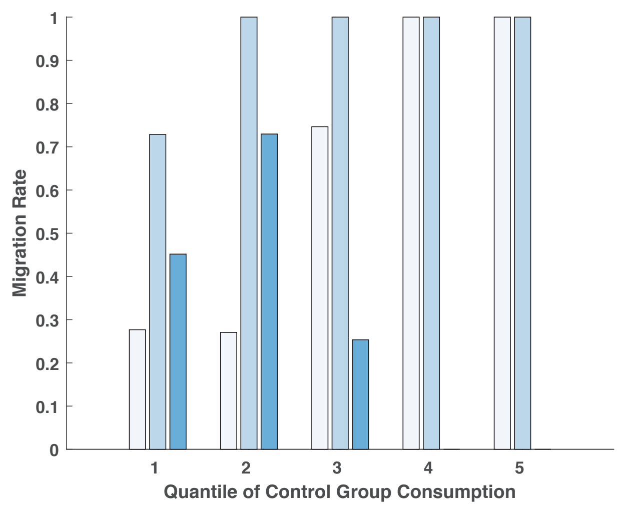

(c) **Model of Bryan et al (2014)**

tire control group). In the data, migrants have somewhat lower reported asset holdings (0.7 months of average consumption) than non-migrants (0.8 months). The same is true in the current model, with 0.4 months of assets for migrants compared to 0.5 months for non-migrants. Recall that these moments are not targeted in our estimation. The model of Bryan et al predicts that migrants have far higher asset levels (1.1 months) than non-migrants (0.5 months). This highlights yet again the counterfactual nature of migration decisions in the Bryan et al model, in which households migrate once they have sufficiently large buffer stocks of assets.

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 C of Table 5 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.

The bottom three rows of Panel C report the predicted effects of an unconditional transfer in the current model and the model of Bryan et al. Our model predicts that unconditional transfers induce a negligible increase in migration of around one percent. Again, this moment is not targeted. The model of Bryan et al predicts a counterfactually large increase in migration rates of 22 percentage points. This substantial increase in migration again reflects the constraints on migration in that model, which is 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 need to make migration worthwhile.

What about subsistence and permanent learning about migration ability, which are the other two features of the model of Bryan et al? Surely subsistence constraints are a feature of life in village economies in the developing world, given the low income levels there. Yet, as we show in Appendix Table A.13, 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 our model is not substantially changed once we add subsistence constraints. As such, we find it hard to see subsistence constraints as a central component in the interpretation of the experimental evidence in question.

As for permanent learning about migration ability, the model of Bryan et al makes predictions that are counterfactual here as well. Repeat migration was a key experimental finding, since a significant fraction of households in the treatment villages responded to the migration subsidies even a year after the treatment concluded. In the data, the probability of a household sending a migrant in period t conditional on having sent one in period $t-1$ is 0.68, while the probability of sending a migrant in t conditional on *not* having sent one in $t-1$ is 0.26 (see Appendix Table A.14). In our model, these probabilities are 0.70 and 0.14. In the model of Bryan et al, these probabilities in the second period (the first in which the migration probabilities can actually be calculated) are 0.52 and 0.74, counterfactually higher for those that didn't migrate earlier, and in the long run exactly half the population migrates (since they know their types by then). These counterfactual predictions for repeat migration suggests that the role of permanent learning is not at the heart of the experimental results in question.

To summarize, the model in the current paper succeeds in matching the experimental evidence of Bryan et al. (2014), whereas the model explored in that paper does not. Not by coincidence, the current model offers a substantially different interpretation of the experimental evidence than theirs. In our model, credit constraints and migration risk play very little role in the migration decisions of rural households. Instead, the households that can be induced to migrate are those that have a modest urban productivity level and below-average savings, which is driven in large part by the low return on savings, R. The fact that migration is mostly among those with low assets and income levels pushes away from a world where households forgo migration due to migration risk and lack of savings. Subsistence constraints, while certainly a feature of village life in the developing world, are a central part of the interpretation of the experimental evidence at hand. Nor is learning about permanent migration ability, which predicts counterfactually high persistence in migration rates. It is perhaps not surprising then that the welfare implications of a model with the features of Bryan et al (2014) differ substantially from those of our model, as we show in the following section.

6. 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 consumption-equivalent metric used in macroeconomics since Lucas (1985) and extensively thereafter. This welfare metric computes the percent increase in consumption, p, that makes the household indifferent between a p-percent consumption increase in perpetuity and being offered the conditional migration subsidy.

We compute the welfare gains of three types of conditional migration subsidies. The first is a one-time transfer in partial equilibrium, as in the migration experiments Bryan et al. (2014). The second is a permanent subsidy that offers rural households a conditional migration transfer each period conditional on having low assets, in partial equilibrium, meaning that wages do not change. The third is a permanent subsidy but in general equilibrium, so that rural wages adjust depending on how much migration is induced in equilibrium.

We simulate the permanent effects of the policies as follows. In the estimated economy, we record the level of assets that satisfies the sampling criteria used by Bryan et al. (2014); call this a^*. Recall that we compute a^* so that only the bottom 50 percent of the rural asset distribution are offered the migration subsidy. In the case of permanent subsidies, we offer a migration subsidy every lean season to every household with assets below a^*. This assures that the permanent transfers are targeting households in the same way as in the one-time transfers. In the general-equilibrium counterfactual, we offer a permanent migration subsidy and in addition

Table 6: Welfare Effects of Conditional Migration Subsidies

		One-Time Partial Equilibrium		Permanent Partial Equilibrium		Permanent General Equilibrium	
		Welfare	Migr. Rate	Welfare	Migr. Rate	Welfare	Migr. Rate
Income Quintile	1	1.05	85	4.79	92	4.41	91
	2	0.40	62	3.80	79	3.44	76
	3	0.26	53	3.47	70	3.09	67
	4	0.16	43	3.15	60	2.81	57
	5	0.11	39	2.69	51	2.38	48
Average		0.39	57	3.58	71	3.22	68
Fraction Rural			60		66		65
Fraction w. Experience			23		42		41

Note: The table reports the (lifetime) consumption-equivalent welfare gains from conditional migration transfers. The rows are for different income quintiles, with 1 being the poorest quintile and 5 being the richest. The first two columns are for a one-time transfer in partial equilibrium. The next two columns are for a permanent conditional migration transfer program offered each period indefinitely, assuming a fixed wage rate. The final two columns are for a permanent transfer adjusting wages to clear labor markets.

re-solve for the rural wage that is consistent with Equation (4).

Two further aspects of these policy experiments are worth noticing. The first is that when the subsidies are made permanent, the households will update their expectations accordingly. This will affect the households' location choices, their willingness to seasonally migrate and their accumulation of experience. The second is that the wage changes in general equilibrium affect those that choose not to migrate each period. Note that this can also cut the other way if more people decide to locate in the rural area, since wage rates in the good season will fall, hurting rural households.

Table 6 reports the welfare gains of the three migration subsidies in question. For all three subsidy types we report the average welfare gain and seasonal migration rate by consumption quintile in the rural area (with one being the lowest). We also report the average overall welfare gains for rural households, the percent of households living in the rural area (to see how permanent migration is affected) and the percent of rural households that have experience in migrating in the new equilibrium.

Several features stand out in Table 6. First, the welfare gains for the one-time transfer in partial equilibrium are modest on average, at 0.39 percent consumption equivalents, though much larger for the poorest quintile, who gain 1.05 percent. Second, for the permanent transfers, general equilibrium affects work against the welfare gains, but by relatively small margins. This can be seen by comparing the results for the latter two columns, which show an average welfare gain of 3.58 percent in partial equilibrium and 3.22 percent in general equilibrium. There are several forces behind these results. First, because seasonal migration increases by a large amount, labor becomes scarcer in the rural area and rural wages rise by 3 percent as a result. This clearly benefits those who do not move in the Monga season. However, this force works also in the opposite direction during the good season. Because of the change in policy, more households relocate and live in the rural area. As Table 6 shows, the rural population actually increases from 60 percent to 65 percent in response to the permanent migration subsidies. This population increase lowers rural wages in the good season by about one percent.

With that said, the welfare gains from permanent transfers are large on average. There are several reasons why this is the case. The first is the permanence of the transfers, and the associated insurance provided by subsidized migration for those with sufficiently low assets. This can be seen by noting that for households in the top income quintile, the one-time transfer delivers a tiny benefit of 0.11 percent consumption equivalents, whereas the permanent transfer deliver a much larger gain of 2.38 percent. Intuitively, the permanent migration subsidies do not help the richest quartile much in the present, but do so in future states of the world in which their rural productivity and asset levels are low (and hence migration is valuable).

A second reason why the welfare gains are so much larger for permanent migration subsidies is that they result in a large endogenous change in migration experience. As a result, many households do not face the disutility associated with migration. As Table 6 shows, the fraction of the rural population with experience almost doubles from the one-time case to permanent case in general equilibrium, from 23 percent to 41 percent. The reason is that in the new stationary distribution, a lower monetary cost of migration allows more households to migrate more frequently, which leads more of them to be experienced in the new stationary equilibrium. This dynamic complimentarily between experience and the subsidies leads even more households to migrate seasonally, and to do so with substantially larger utility gains than if they were inexperienced at migrating.

The welfare calculations of this section also highlight which inefficiencies are most important in our model, and which are less important. Pecuniary externalities on rural wages are less important. Even the substantial changes in rural migration rates and rural residency induced by permanent migration subsidies lead to a relatively modest rural wage increase of just 3 percent. Market incompleteness and the lack of better savings technologies are more important inefficiencies. Once rural households are offered the possibility of subsidized migration when their assets are sufficiently low, even the richest households can expect substantial welfare benefits. To the extent that the migration disutility represents an inefficiency, it is also central to interpreting the model's welfare implications, as it plays a key role in the welfare benefits from any given migration episode and the fraction of households that migrate in any given period.

One can get some additional insight into how the welfare gains arise in the model by revisiting Figure 2. Panel (a) of the figure shows the policy functions for a household with a low level of z and no migration experience, in the control group. The policy functions for the treatment group are not depicted (for expositional purposes) but would expand the migration regions up and to the right. In the figure, household (i) is inframarginal and will make a temporary move whether or not it is offered a conditional migration transfer. Household (ii) is on the margin and is induced to migrate by a conditional transfer, but would otherwise stay in the rural area. What household (i) and (ii) have in common is that they are willing to go through with the ordeal of migration since marginal utility of consumption is so high for them. Household (iii) will not migrate even when offered a transfer. Given the high level of assets and the high shock, this household prefers the rural area even with the transfer.

Who gains the most from the conditional migration transfers? Perhaps surprisingly, it is household (i), the inframarginal household. This household has low levels of assets and a bad shock, so has a very low level of consumption. Marginal utility is relatively high for this household, so the transfer leads to a relatively large increase in its welfare. Household (ii) also gains a lot, but by somewhat less than household (i), since household (ii) has a higher level of consumption before the transfer. It is true that this household changed its behavior as a result of the experiment, but that is not the key driver of welfare gains, it is the channeling of funds to the vulnerable households who are the only ones willing to go through with the ordeal of migration. Household (iii) doesn't take up the conditional transfer, and so its welfare does not change at all due to the intervention.

6.1. Welfare Under Alternative Scenarios

One may wonder why the consumption-equivalent welfare gains from migration subsidies are not higher if the workers induced to migrate raise their consumption by 30 percent on average. To answer this question we consider alternative parameterizations of the model, each time changing exactly one parameter value and re-computing the welfare gains from the migration transfers. The overall goal in this section is to further illustrate how our model allows for an interpretation of the experiments of Section 2 based on spatial misallocation, with credit constraints and migration risk driving migration outcomes. As we show below, such an interpretation would give rise to substantially larger welfare gains from migration subsidies than found in this paper, but at the cost of making counterfactual predictions about important aspects of the experimental data.

Table 6.1 summarizes the model's welfare predictions under these alternatives. The first row

Table 7: Welfare Gains Under Alternative Parameterizations

	Average Welfare Gains	LATE (Cons.)	OLS (Cons.)	Treatment Effect (Migration)	Seasonal Migration Control
Data	-	30	10	22	36
Benchmark calibration	0.39	29	10	21	36
+ Higher urban risk	0.12	27	51	10	16
+ No migration disutility	0.51	9	29	28	55
+ Higher urban TFP	1.29	33	51	15	84
+ Higher migration cost	1.98	16	34	62	36

Note: This table reports the average welfare gains implied by the model, the LATE and OLS effects of migration on consumption, seasonal migration in the control group, and the treatment effect on migration implied by the model for each specific calibration. Row 1 shows the data. Row 2 is the benchmark calibration that results from the simulated method of moments. Row 3 ("+ Higher urban risk") changes the parameter shaping the urban relative shock by setting $\gamma = 1.5$. Row 4 ("+ No migration disutility") further removes the disutility of migration by setting $\bar{u} = 1$. Row 5 ("+ Higher urban TFP") further doubles the level of urban TFP of 3 (instead of $A_u = 1.5$). Row 6 ("+ Higher migration cost") sets p_T to be 50 percent of rural consumption so that the model matches seasonal migrant rates in the control group.

reproduces the main experimental moments on which we will focus, and the second row reports the model's predictions for the same moments plus the average welfare gain from the migration transfers. The third row raises γ from the estimated value of 0.57 up to 1.5, meaning that shocks are now relatively larger in the urban area. By itself this leads welfare gains to fall to 0.12 percent consumption equivalent, the OLS coefficient of consumption on migration to rise to a counterfactual 51 percent, and the treatment effect on migration to fall to a counterfactually low level of 10 percent. The fourth row sets $\bar{u} = 1$, which means there is no disutility from migration. Welfare gains raise substantially to 0.51 percent, but the LATE falls to a counterfactually low level of 9 percent, and the migration rate in the control group rises to 55 percent, well above the data. Clearly though, the lower value of $\bar{u}$ is an important driver of the model's welfare gains. The fifth column doubles A_u, the urban productivity, to a value of 3. The welfare gains now increase further to 1.29 percent, while other moments remain counterfactual, in particular the seasonal migration rate, which is now an implausible 84 percent.

To lower the migration rate back to a level similar to the data, the last row increases the migration cost up to $m = 0.19$, which is the value that matches the 36 percent migration rate in the control group again. This change also raises the amount of the migration transfer, by construction, since the migration subsidies are intended to cover the migration cost and actually induce migration. Under this parameterization, the welfare gains from the transfers rise to 1.98 percent, or five times what they are in the benchmark calibration. The source of the welfare gains now become relaxing credit constraints, which keep risk-averse migrants from reaching a much more productive urban area, in the spirit of the model of Bryan et al. (2014). Yet the data do not support such an interpretation. As one example, the LATE effect of consumption on migration is counterfactually lower than the OLS coefficient, pointing to inaccurate sorting patterns for migrants. As another example, the treatment effect of migration is far too large, pointing to the counterfactually large migration costs in this calibration of the model.

6.2. Alternative Rural-Based Policies

It is useful to compare the welfare generated from migration subsidies to other methods that policy makers often use to address rural poverty. Unconditional cash transfers (UCTs) are one such policy tool. Another common place-based policy utilized in developing countries are rural "workfare" programs that provide employment guarantees in rural areas. For example, India's enormous NREGA program provides funding for rural workers to work in public projects in rural areas. These policies are explicitly tied to rural areas, and, thus, discourage rural-urban migration (Imbert and Papp, 2016).

The only fully experimental RCT-based evaluation of a rural workfare program finds no significant benefits and even negative spillovers on non-beneficiaries (see Beegle, Galasso, and Goldberg (2017) for Malawi). In contrast, Imbert and Papp (2015) report some positive benefits from India's program. Thus, we simulate the effects of a rural workfare program in our model as transfers to rural households conditional on those workers remaining in the rural area for that period. The goal is to capture the general spirit of rural workfare programs without tying our exercise to particular policy details in specific countries. To conduct a budget-neutral comparison with the migration subsidy, we set the total expenditure on workfare transfers to be equal to the conditional migration subsidies.

Table 8 compares the welfare gains from migration transfers, unconditional transfers, and rural workfare policies costing exactly the same amount. Overall, the average welfare gains are very similar in the three programs, at 0.39 for migration transfers, 0.44 for unconditional transfers and 0.41 for the rural workfare programs. The migration transfers are better for the poorest households than the other two policies but worse for the richest households. The poorest quintile of households gain 1.05 percent under the migration transfers compared to 0.92 percent for the unconditional transfers and 0.78 percent for the rural workfare policies. The reason is that the conditional migration subsidies are best at targeting the neediest households is that they create an ordeal: only those households with the worst options in the rural area are induced by the subsidy to incur the disutility of migration. This allows more funds to be directed to these neediest households and away from the least needy (who decline the conditions of the

Table 8: Migration Transfers, Unconditional Transfers and Rural Workfare Policies

		Migration Transfers		Unconditional Transfers		Rural Workfare	
		Welfare	Migr. Rate	Welfare	Migr. Rate	Welfare	Migr. Rate
Income Quintile	1	1.05	85	0.92	46	0.78	28
	2	0.40	62	0.49	37	0.46	25
	3	0.26	53	0.36	35	0.35	23
	4	0.16	43	0.27	30	0.26	22
	5	0.11	39	0.19	30	0.18	24
Average		0.39	57	0.44	36	0.41	25

Note: The table reports the lifetime consumption-equivalent welfare gains (in percentage points) from the conditional migration transfers relative to an unconditional transfer program costing and rural workfare programs costing the same total amount.

migration transfer) while keeping the total cost of the program the same.

7. Empirical Evidence on the Source of Migration Disutility

Our model infers that many rural residents experience significant non-monetary disutility from migration, and this plays an important role in our interpretation of the evidence of Section 2. This is also crucial for our welfare calculations, in that some of the large consumption gains from migration are offset by this disutility. Therefore, we explore whether the large disutility is plausible, and what the source of that disutility might be. To do so, we collect new survey data from Bryan et al.'s (2014) experimental sample of migrants on their preferences for specific migration attributes. This allows us to characterize exactly what this disutility may represent, for the exact same sample of households that we used to estimate our model.

Conducting field experiments that vary a number of non-monetary attributes of the migration experience (such as quality of living conditions, wages, risk, family separation) would be practically challenging and prohibitively expensive, so our approach in this section is to conduct discrete-choice experiments (DCE) on the migrant sample. The DCE presented respondents with a series of hypothetical scenarios in which we randomly varied a few key attributes associated with one of two migration options. The surveys presented respondents with (hypothetical) options for the fall 2015 lean season and asked them to indicate which migration choice they would make. The attributes we presented under each option randomly varied the probability of finding employment in the city, the wage if employed, how frequently the migrant could return to visit family (to minimize separation), and access to a hygienic latrine in their residence at the migration destination, which is a useful proxy for the quality of housing amenities that the migrant would experience in the city.

There is a reasonable concern that, in DCEs, people's responses to hypothetical questions may not accurately reflect their real-world behavior. They may, for example, express an interest in migrating in response to a hypothetical question, even though they may be more hesitant if the actual choice ever presented itself. We are therefore careful not to make any inferences about people's overall migration propensity using this exercise. Instead, in analyzing people's responses to the hypothetical scenarios, we infer the *relative weights* people place on quality of living conditions relative to wages or concerns about family separation.

Each respondent was asked to choose one of two migration options or a third, "opt-out" no-migration option. The experimental setup for the hypothetical options was created to mimic the circumstances under which the equivalent decision would be made in the real world (Ryan and Skatun, 2004). In the example shown in Appendix Figure A.1, for example, both options feature a 33 percent chance of employment. Choice #1 offers a lower wage if employed but better amenities (more regular family contact and a hygienic latrine in the residence) compared to Choice #2.

We conducted these DCEs on a sample of 2,714 respondents, presenting each respondent with seven different choice sets for which the values of attributes are varied. We used the Choice Experiment tools in JMP12 (built on SAS) to generate algorithms that pick values for the attributes under each migration option in each choice problem in such a way that the power of the experiment is maximized. We observed a total of 18,998 choices, but to eliminate any bias stemming from recent induced migration experience, we used only choices made by respondents who resided in the control villages in the experiments. We estimated a multinomial logit model of migration choice as a function of the offered attributes of each location, using the remaining 3,462 observations.

Table 7 presents the predicted probabilities and estimated marginal effects from this multinomial logit regression. We report the marginal effects of improving each attribute associated with option #2.[12] The middle two data columns of Table 7 show the predicted probabilities (PP) and marginal effects (ME) on the propensity to migrate to destination #2. The first and last two data columns show the PP and ME on destination #1 and "No Migration" when the characteristics of destination #2 are varied.

The four attributes for each destination that we specified in our surveys are as follows. The first is the probability of employment, with three possible values that were randomly varied

[12]We set all attributes associated with option #1 at their least attractive values, and those associated with option #2 at median values. The rationale for this is to effectively create only two relevant choices for the potential migrant: either migrate to destination 2, or stay at home. This binary choice most closely resembles the decisions made by agents in our model. Recall that we model a binary migration choice.

Table 9: Estimated Marginal Effects on Migration

	Migration Opp. #1		Migration Opp. #2		No Migration	
	PP	ME	PP	ME	PP	ME
33% Prob. Employment	0.112***	0.000	0.587***	0.000	0.301***	0.000
	(0.019)	(.)	(0.056)	(.)	(0.061)	(.)
66% Prob. Employment	0.075***	-0.037***	0.716***	0.129***	0.209***	-0.092***
	(0.013)	(0.010)	(0.047)	(0.031)	(0.047)	(0.032)
100% Prob. Employment	0.045***	-0.067***	0.794***	0.207***	0.160***	-0.141***
	(0.009)	(0.013)	(0.036)	(0.037)	(0.035)	(0.040)
Family visit once in 60 days	0.074***	0.000	0.717***	0.000	0.209***	0.000
	(0.014)	(.)	(0.044)	(.)	(0.044)	(.)
Family visit twice in 60 days	0.075***	0.001	0.716***	-0.001	0.209***	0.001
	(0.013)	(0.007)	(0.047)	(0.025)	(0.047)	(0.024)
Family visit 4 times in 60 days	0.063***	-0.011	0.723***	0.005	0.214***	0.005
	(0.012)	(0.007)	(0.053)	(0.030)	(0.054)	(0.030)
No Latrine in residence	0.075***	0.000	0.716***	0.000	0.209***	0.000
	(0.013)	(.)	(0.047)	(.)	(0.047)	(.)
Pucca Latrine in residence	0.026***	-0.049***	0.906***	0.190***	0.068***	-0.141***
	(0.005)	(0.009)	(0.022)	(0.031)	(0.020)	(0.033)
Daily Wage (Taka), Opp # 2		-0.001***		0.003***		-0.002***
		(0.000)		(0.000)		(0.000352)
Observations	3462	3462	3462	3462	3462	3462

Note: The PP columns represent the predicted probabilities of migrating at each given condition, and the ME columns represent the marginal effects of changing migration conditions in each category. Both are measured while fixing the conditions of migration Choice #1 at the worst values, and fixing the conditions of migration Choice #2 at the median values. The sample includes only households from the control group.

across the choice scenarios: 33%, 66% and 100%, which is meant to capture labor-income risk. The second is the daily wage, which could take one of five possible values running from 200 to 340 Taka per day. The third is living conditions in the city, which had two categories: either a *pucca* (hygienic) latrine in the residence, or no latrine. This is a context-relevant proxy for the overall quality of housing. The fourth is the extent of family separation, which had three possible categories: the ability to go back and visit family once, twice or four times during the monga. The daily wage is modeled as a continuous variable in the multinomial logit, while the other attributes are modeled as categorical variables.

Table 7 shows that an increase in employment probability at destination 2 from 33 percent to 66 percent or 100 percent (holding destination #1 characteristics fixed) increases the propensity to migrate to destination #2 by 19.7 and 20.7 percentage points. Labor-income risk is, therefore, a quantitatively important deterrent to migration. The next three rows show that the frequency of family visits has a negligible (and statistically insignificant) effect on migration choices.

In stark contrast, having a latrine in one's residence increases the probability of choosing destination # 2 by 19 percentage points. Housing conditions at the destination therefore appear to be an important determinant of migration choices. Finally, the probability of migrating to destination #2 increases by 0.3 percentage points for every additional Taka in daily wage that is offered, meaning a 15 percentage point increase in the migration probability for an extra 50 Taka in income. Thus, having a better housing option is similar to an additional 63 Taka per day in wages, which amounts to a 32 percent increase over the base value of 200 Taka per day that we used in our hypothetical DCE scenarios (and corresponds to roughly the average wages earned by migrants in the city). To the extent that rural-urban migrants generally face poor urban housing options (proxied by a lack of access to convenient latrines, which is a realistic worry in the slums of South Asian cities), this represents a large non-monetary cost of migration and a substantial offsetting force to the higher wages earned by migrants. The large migration disutility that our model infers from people's actual migration and re-migration behavior does appear to be validated in the DCEs when these (potential) migrants are asked to explicitly consider the non-monetary dimensions of the migration experience.

The contrast between the weight that potential migrants place on urban housing conditions versus their relative inattention to the length of family separation is notable. For short-run seasonal migration, frequency of family visits appears less important than housing quality. What makes this contrast interesting from a policy perspective is that concerns about housing conditions can be more easily addressed through policy compared to concerns about family separation. The large welfare gains for the poor and the disutility parameter that we estimate from our model, coupled with these DCE results, suggest that governments may want to improve urban slum housing conditions, as a way to raise the welfare gains from migration to cities.

8. 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 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, using 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 rural-urban migration thus come about largely from providing better insurance opportunities for rural households in periods where they are vulnerable. Future research should explore the consequences of encouraging internal migration in other countries and settings, as well as the interactions between urban infrastructure and the welfare gains from migration. Scaled-up versions of any migration encouragement program may introduce other complexities such as strategic interactions in people's travel decisions due to their pre-existing connections. Future work will need to grapple with these complexities of scale.

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, which 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 (2016): "Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries," Unpublished Working Paper, University of Chicago.

AKRAM, A. A., S. CHOWDHURY, AND A. M. MOBARAK (2017): "General Equilibrium Effects of Emigration on Rural Labor Markets," Unpublished Working Paper, Yale University.

ANDREWS, I., M. GENTZKOW, AND J. M. SHAPIRO (2017): "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *Quarterly Journal of Economics*, 132, 1553–1592.

BAZZI, S. (2017): "Wealth Heterogeneity and the Income Elasticity of Migration," *American Economic Journal: Applied Economics*, 9, 219–255.

BEEGLE, K., E. GALASSO, AND J. GOLDBERG (2017): "Direct and indirect effects of Malawi's public works program on food security," *Journal of Development Economics*, 128, 1–23.

BEWLEY, T. (1977): "The Permanent Income Hypothesis: A Theoretical Formulation," *Journal of Economic Theory*, 16, 252–292.

BROOKS, W. AND K. DONOVAN (2019): "Eliminating Uncertainty in Market Access: The Impact of New Bridges in Rural Nicaragua," Unpublished Working Paper, Yale University.

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.

DAVILA, J., J. H. HONG, P. KRUSELL, AND J.-V. RIOS-RULL (2012): "Constrained Efficiency in the Neoclassical Growth Model With Uninsurable Idiosyncratic Shocks," *Econometrica*, 80, 2431–2467.

DONOVAN, K. (2016): "Agricultural Risk, Intermediate Inputs, and Cross-Country Productivity Differences," Unpublished Working Paper, University of Notre Dame.

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.

GOURIEROUX, C. AND A. MONTFORT (1996): *Simulation-Based Econometric Methods*, Oxford: Oxford University Press.

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.

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 Heterogenous 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.

Hicks, J. H., M. Kleemans, N. Y. Li, and E. Miguel (2017): "Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata," NBER Working Paper No. 23253.

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 (2015): "Labor Market Effects of Social Programs: Evidence from India's Employment Guarantee," *American Economic Journal: Applied Economics*, 7, 233–263.

——— (2016): "Short-term Migration, Rural Public Works and Urban Labor Markets: Evidence from India," Department of Economics, University of Warwick.

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. and M. E. Waugh (2013): "Selection, Agriculture, and Cross-Country Productivity Differences," *The American Economic Review*, 103, 948–980.

Lucas, R. E. (1985): "Models of Business Cycles," in *Lectures on Economic Growth*, Oxford: Basil Blackwell.

McMillan, M. S. and D. Rodrik (2011): "Globalization, Structural Change and Productivity Growth," NBER Working Paper No. 17143.

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.

Peters, M. (2016): "Heterogeneous Mark-Ups, Growth and Endogenous Misallocation," Unpublished Working Paper, Yale University.

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.

Ryan, M. and D. Skatun (2004): "Modelling Non-demanders in Choice Experiments," *Health Economics*, 13, 397–402.

Terry, S. (2019): "The Macro Impact of Short-Termism," Unpublished Working Paper, Boston University.

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.

Vollrath, D. (2009): "How Important are Dual Economy Effects for Aggregate Productivity?" *Journal of Development Economics*, 88, 325–334.

Yang, M.-J. (2016): "Micro-level Misallocation and Entry Selection," Unpublished Working Paper, University of Washington.

Young, A. (2013): "Inequality, the Urban-Rural Gap, and Migration," *The Quarterly Journal of Economics*, 128, 1727–1785.

Appendix (for Online Publication)

A. Appendix Tables and Figures

Table A.1: Variance Covariance Matrix of Cross-Sectional Moments (V_c)

	Urban-Rural wage gap	Percent rural	Var of log urban wage
Urban-Rural wage gap	0.032606		
Percent in rural	-0.000514	0.000186	
Var of log urban wage	0.008279	-0.000083	0.003475

Table A.2: Variance-Covariance Matrix of Experimental Moments (V_e)

	Con: Var log cons Δ rural	Control: P(hh no a)	Control: Seasonal migrants	Cons OLS	Migration, T-C	Migration, T-C, year 2	Cons, LATE	P(09/08)
Cont: V(log rur Δcon)	0.00128584							
Cont: P(hh no a)	0.00001115	0.00012712						
Cont: Mig Rate	-0.00003059	0.00000817	0.00069678					
Cons. OLS	0.00019928	0.00001947	-0.00000131	0.00199811				
Mig, T-C	0.00002116	-0.00001500	-0.00036251	-0.00001064	0.00057068			
Mig, T-C , Y2	0.00004212	-0.00000289	-0.00017265	0.00001570	0.00020858	0.00059693		
Cons, LATE	0.00000031	-0.00003335	0.00034881	-0.00037858	-0.00031222	-0.00019477	0.00934914	
P(mig09/Mig08)	-0.00007174	-0.00001168	0.00001397	-0.00025705	0.00000863	-0.00037700	0.00004016	0.00208298

Table A.3: Derivatives of Moments with Respect to Each Parameter (M)

	θ	$\bar{u}$	λ	π	γ	A_u	σ_s	ρ	var	σ_{rc}	σ_{ui}
Con: Var(log cons Δ rural)	0.01	0.00	0.01	0.00	0.00	0.00	-0.05	-0.03	0.00	1.00	0.00
Con: P(households no assets)	-0.20	-1.75	0.01	-0.05	-2.96	-0.41	-3.97	7.82	-0.07	0.00	0.00
Con: Seasonal migrants	-2.30	-1.90	-0.66	1.02	-1.12	0.28	-0.34	0.26	1.77	0.00	0.00
Consumption, OLS	0.71	0.23	-0.08	0.24	0.52	-0.27	-0.33	0.84	-0.02	0.00	0.00
Migration, Treat-Control	-0.22	-0.12	0.26	-0.35	-0.16	0.18	-0.25	0.49	-0.61	0.00	0.00
Migration, Treat-Control, year 2	0.03	-0.12	-0.08	0.23	-0.37	0.03	-0.07	0.24	-0.31	0.00	0.00
Consumption, LATE	-0.34	-0.09	-0.01	0.18	-0.34	-0.06	-0.59	1.01	0.06	0.00	0.00
P(mig09/mig08)	0.22	-0.54	-0.16	0.05	0.26	-0.06	-0.69	1.56	-0.98	0.00	0.00
Urban-Rural wage gap	4.95	0.06	0.45	-0.04	0.44	-0.25	-0.46	1.02	-0.73	0.00	0.00
Percent in rural	0.92	-0.18	-0.05	0.18	0.09	-0.49	-0.12	0.35	0.13	0.00	0.00
Variance of log wages in urban	1.47	-0.03	-0.09	0.06	0.26	0.04	0.16	-0.28	-0.01	0.00	1.00

Table A.4: Andrews et al (2017) Matrix of Sensitivity of Parameters to Moments

	V(log c Δ rural)	Con: P(hh no a)	Con:Seas. migrants	Cons, OLS	Mig., T-C	Mig., T-C, y2	Cons, LATE	P(mig09 /mig08)	wage gap	% rural	Var log w_u
θ	0.0	-0.1	0.0	0.3	0.6	0.1	0.2	-0.1	-0.2	0.2	0.0
$\bar{u}$	0.0	-0.1	0.4	-0.1	1.1	-0.4	-0.6	0.2	0.1	0.8	0.0
λ	0.0	0.4	-0.8	-3.0	-5.9	-1.0	-0.6	2.3	-0.0	-1.6	0.0
π	0.0	0.4	-0.5	-1.7	-1.8	-1.9	-1.4	0.8	-0.10	-0.3	0.0
γ	0.0	0.2	-0.2	-0.8	-0.6	1.0	-0.2	-0.3	-0.0	0.1	0.0
A_u	0.0	-0.1	-0.4	-1.0	0.0	-0.3	1.1	0.2	-0.4	2.5	0.0
σ_s	0.0	-0.6	-0.8	-6.1	-4.4	-4.1	8.3	2.8	0.8	-0.2	0.0
ρ	0.0	-0.4	-0.4	-3.5	-2.2	-1.8	4.1	1.4	0.4	0.2	0.0
σ_ν	0.0	-0.2	-0.2	-0.8	-0.1	0.6	0.9	0.6	-0.0	0.3	0.0
σ_{rc}	-1.0	-0.1	-0.0	-0.4	-0.2	-0.2	0.5	0.2	0.1	0.0	0.00
σ_{ui}	0.0	0.0	-0.0	-0.4	-1.0	-0.1	-0.5	0.3	0.3	-0.4	-1.0

Table A.5: Targeted Moments in Data and Model with Different R values

Moments	Data	Model R=0.93	Model R=0.95	Model R=0.97
Control: Variance of log consumption growth in rural	0.19	0.19	0.19	0.19
Control: Percent of rural households with no liquid assets	47	60	48	38
Control: Seasonal migrants	36	38	36	35
Control: Consumption increase of migrants (OLS)	10	9	10	11
Treatment: Seasonal migration relative to control	22	20	21	21
Treatment: Seasonal migration relative to control in year 2	9	4	4	4
Treatment: Consumption of induced migrants (LATE)	30	30	29	27
Control: Probability of repeat migration	68	71	70	70
Urban-Rural wage gap	1.89	1.88	1.89	1.90
Percent in rural	62	59	60	60
Variance of log wages in urban	0.56	0.56	0.56	0.56

Note: The table reports the main moments of the paper for alternative values of R. The estimated model has $R = 0.95$. The model is not re-estimated in the cases of $R = 0.93$ and $R = 0.97$.

Table A.6: Targeted Moments in Data and Model with Different β values

Moments	Data	Model β=0.93	Model β=0.95	Model β=0.97
Control: Variance of log consumption growth in rural	0.19	0.19	0.19	0.19
Control: Percent of rural households with no liquid assets	47	59	48	38
Control: Seasonal migrants	36	37	36	36
Control: Consumption increase of migrants (OLS)	10	8	10	12
Treatment: Seasonal migration relative to control	22	20	21	20
Treatment: Seasonal migration relative to control in year 2	9	4	4	4
Treatment: Consumption of induced migrants (LATE)	30	29	29	28
Control: Probability of repeat migration	68	71	70	71
Urban-Rural wage gap	1.89	1.90	1.89	1.87
Percent in rural	62	60	60	59
Variance of log wages in urban	0.56	0.56	0.56	0.56

Note: The table reports the main moments of the paper for alternative values of $\beta = 0.95$. The estimated model has $\beta = 0.95$. The model is not re-estimated in the cases of $\beta = 0.93$ and $\beta = 0.97$.

Table A.7: Targeted Moments in Data and Models with no $\bar{u}$ and $\rho = 0$

Moments	Data	Model Full	Model $\bar{u} = 1$	Model $\rho = 0$
Control: Variance of log consumption growth in rural	0.19	0.19	0.19	0.28
Control: Percent of rural households with no liquid assets	47	48	48	2
Control: Seasonal migrants	36	36	55	34
Control: Consumption increase of migrants (OLS)	10	10	-7	17
Treatment: Seasonal migration relative to control	22	21	10	22
Treatment: Seasonal migration relative to control in year 2	9	4	0	4
Treatment: Consumption of induced migrants (LATE)	30	29	23	21
Control: Probability of repeat migration	68	70	56	71
Urban-Rural wage gap	1.89	1.89	1.86	1.88
Percent in rural	62	60	73	57
Variance of log wages in urban	0.56	0.56	0.65	1.54

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

Table A.8: Welfare Under Alternative Models with no $\bar{u}$ and $\rho = 0$

		Full Model		$\bar{u} = 1$		$\rho = 0$	
		Welfare	Migr. Rate	Welfare	Migr. Rate	Welfare	Migr. Rate
Income Quintile	1	1.0	85	1.5	87	1.3	64
	2	0.4	62	0.7	74	0.8	61
	3	0.2	53	0.4	64	0.5	57
	4	0.1	43	0.3	54	0.3	52
	5	0.1	39	0.3	49	0.1	48
Average		0.4	57	0.6	65	0.6	56

Note: The table reports the (lifetime) consumption-equivalent welfare gains from the conditional migration transfers relative to an unconditional transfer program costing the same total amount and to a rural workfare program that costs the same amount. The numbers in the table are the average percent increase in consumption each period that would make the households indifferent between the consumption increase and the transfers, and the seasonal migration rates, by quintile of the rural income distribution.

Table A.9: Targeted Moments in Data and Model with Subsistence

Moments	Data	Model	
		Full Calibration	Full Calibration w/ Subsistence
Control: Variance of log consumption growth in rural	0.19	0.19	0.23
Control: Percent of rural households with no liquid assets	47	48	0
Control: Seasonal migrants	36	36	76
Control: Consumption increase of migrants (OLS)	10	10	5
Treatment: Seasonal migration relative to control	22	21	14
Treatment: Seasonal migration relative to control in year 2	9	4	4
Treatment: Cons of induced migrants relative to control (LATE)	30	29	46
Control: Probability of repeat migration	68	70	80
Urban-Rural wage gap	1.89	1.89	1.66
Percent in rural	62	60	56
Variance of log wages in urban	0.56	0.56	0.64

Note: The table reports the moments targeted using simulated method of moments and their values in the data and in the model. The final calibration reports the moments when a subsistence consumption constraint is added and set to equal 25 percent of average rural consumption in the lean season.

Table A.10: Welfare Effects of Migration Subsidies with Subsistence

		Benchmark Model		Subsistence	
		Welfare	Migr. Rate	Welfare	Migr. Rate
Income Decile	1	1.5	93	2.0	92
	2	0.6	76	1.0	82
	3	0.5	68	0.8	77
	4	0.3	57	0.6	71
	5	0.3	54	0.5	67
	6	0.2	51	0.4	61
	7	0.2	46	0.3	56
	8	0.2	42	0.3	53
	9	0.1	41	0.2	52
	10	0.1	38	0.3	45
Average		0.4	57	0.6	65

Note: The table reports the (lifetime) consumption-equivalent welfare gains from the conditional migration transfers. The numbers in the table are the average percent increase in consumption each period that would make the households indifferent between the consumption increase and the transfers, and the seasonal migration rates, by declie of the rural income distribution. The first column is the benchmark model, and the second is a model with a subsistence constraint equal to 25 percent of average rural consumption in the lean season.

Table A.11: Targeted Moments in Data and Model and Migration Costs

Moments	Data	Model: Full Cal $m_p = 2 * m_t$	Model: Full Cal $m_p = m_t$
Control: Variance of log consumption growth in rural	0.19	0.19	0.18
Control: Percent of rural households with no liquid assets	47	48	45
Control: Seasonal migrants	36	36	32
Control: Consumption increase of migrants (OLS)	10	10	11
Treatment: Seasonal migration relative to control	22	21	23
Treatment: Seasonal migration relative to control in year 2	9	4	4
Treatment: Cons of induced migrants relative to control (LATE)	30	29	24
Control: Probability of repeat migration	68	70	60
Urban-Rural wage gap	1.89	1.89	1.79
Percent in rural	62	60	59
Variance of log wages in urban	0.56	0.56	0.59

Note: The table reports the moments targeted using simulated method of moments and their values in the data and in the model in the benchmark calibration and under alternative assumptions about migration costs.

Table A.12: Welfare Under Alternative Assumptions About Migration Costs

		Benchmark Model		$m_p = m_t$	
		Welfare	Migr. Rate	Welfare	Migr. Rate
Income Quintile	1	1.05	85	0.91	81
	2	0.40	62	0.36	60
	3	0.26	54	0.22	50
	4	0.16	43	0.15	43
	5	0.11	40	0.10	39
Average		0.39	57	0.35	55

Note: The table reports the (lifetime) consumption-equivalent welfare gains from migration transfers by income quartile under alternative assumptions about migration costs.

Table A.13: Seasonal Migration Rates

	Control	Treatment	Difference
Data	36	58	22***
			(2.39)
Model	36	57	21
Model of Bryan et al (2014) (initial conditions)	66	97	31
Model of Bryan et al (2014) (+ no subsistence)	83	98	15
Model of Bryan et al (2014) (long run)	50	50	0

Note: This table reports the seasonal migration rates in the control and treatment villages of Bryan et al (2014) expressed in percentage points, and the standard error and statistical significance of the difference, where ***,** and * mean significance at the 1-percent, 5-percent and 10-percent levels. The next four rows present the same statistics in the current model, the model of Bryan et al (2014) under the initial conditions (t=0), without the subsistence constraint, and in the long run ($t \geq 10$) with the subsistence constraint.

Table A.14: Repeat Migration Patterns

	Pr(Migrate$_t$ | Migrate$_{t-1}$)	Pr(Migrate$_t$ | Not Migrate$_{t-1}$)
Data	0.68	0.26
Model	0.70	0.14
Model of Bryan et al (2014)	0.52	0.74

Note: This table reports patterns of repeat migration, measured by probabilities of migration conditional on migration the previous year, and on no migration in the previous year. The first row reports the conditional migration probabilities in the experiment of Bryan et al (2014). The second and third rows present the same probabilities in the current model and in the model of Bryan et al (2014).

Table A.15: Alternative Estimation with Additive Migration Disutility

Moments	Data	Benchmark	Additive
Control: Variance of rural log consumption growth	0.19	0.19	0.19
	(0.03)		
Control: Percent of rural households with no liquid assets	47	48	50
	(1.13)		
Control: Seasonal migration rate	36	36	45
	(2.64)		
Control: Consumption increase of migrants (OLS)	10	10	5
	(4.47)		
Control: Probability of repeat migration	0.68	0.70	0.60
	(0.46)		
Treatment: Seasonal migration relative to control	22	21	30
	(2.39)		
Treatment: Seasonal migration relative to control in year 2	9	4	-1
	(2.44)		
Treatment: Consumption increase of induced migrants (LATE)	30	29	14
	(9.67)		
Urban-Rural wage gap	1.89	1.89	1.89
	(0.18)		
Percent in rural area	62	60	67
	(1.36)		
Variance of log urban wages	0.56	0.56	0.56
	(0.06)		

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

Figure A.1: Sample Migration Opportunity

			S.1.C.2
Given the attributes below, which option do you choose? Please evaluate each new pair of migration options independent of the ones you saw earlier.			
	Choice #1: Migration	**Choice #2: Migration**	**Choice #3: No Migration**
Chance of Employment	33%	33%	N/A
Daily Wage (Taka)	270	340	Wage at Home in November
Latrine Facility during Migration	Pucca Latrine in Residence	Walk to Open Defecate or Public Pay Toilet	N/A
Family Contact	See Family Every Month	See Family Every 2 Month	N/A
s16bq2_1 Your Choice *(Tick Single Box)*	☐	☐	☐

B. Planner's Problem in Simplified Version of Model

In this section we present a simplified version of the model to illustrate how workers can be misallocated across space due to credit constraints and migration risk. To do so we will solve both the competitive equilibrium allocation and a constrained planner's problem that gives rise to a particular constrained efficient allocation, in the spirit of Davila, Hong, Krusell, and Rios-Rull (2012). We focus on the potential inefficiency coming from the inability of households with low asset holdings to insure themselves again migration risk, as posited by Bryan et al (2014). We abstract from the issue of redistribution across low and high productivity households, which a more general planner's problem would take into consideration, since this is not the focus of this paper.

In this simple model there are two periods and two locations: urban and rural. For simplicity, in period one the households only work, and in period two they only consume. Households all begin in the rural area and either migrate or not to the urban area for work. If they migrate, they pay a cost m. Households have log utility over consumption, and supply labor inelastically. Households differ in their exogenous asset holdings, a, and their urban productivity, z, where $z \geq 0$ and $a \geq 0$ for all households. We assume that there is unit measure of households distributed over (z, a) space, and we refer to a household's type as its (z, a) value.

Households are all identical in the rural area and supply one efficiency unit if they work in the rural area. If they migrate to the urban area they supply their z efficiency units of labor, or they supply zero efficiency units, as we describe below. Households can save at gross interest rate R but are credit constrained and cannot borrow at any interest rate. The production functions in the rural and urban areas are given by:

$$Y_u = A_u N_u \quad \text{and} \quad Y_r = A_r N_r \tag{14}$$

where N_u is the total number of efficiency units supplied in the urban area, N_r is the total number of workers in the rural area, and A_u and A_r are the exogenous productivities in two locations. Markets are competitive, and hence wages per efficiency unit are given by $w_u = A_u$ and $w_r = A_r$.

Migration Risk and Borrowing Constraints

As in the full model, we assume that migration is risky. If a household decide to move, then there are two possible states, s, good ($s = g$) and bad ($s = b$). In the good state, which occurs with probability p, agents get income $z \cdot w_u$. In the bad state, which realizes with probability $1 - p$, the household's efficiency units become worthless, and households cannot supply any labor. In this case the household must consume from savings net of migration costs.

Competitive Equilibrium. We begin by characterizing the outcome of the competitive equilibrium. Letting $M \in \{0, 1\}$ be the endogenous moving decision, then the household's problem can be written as:

$$\max_{c,a',M} \mathbb{E}[\log(c)]$$

s.t.

$$c = \begin{cases} A_r + aR & \text{if } M = 0 \text{ or } a < m \\ A_u z + (a-m)R & \text{if } M = 1 \text{ and } a \geq m \text{ and } s = g \\ (a-m)R & \text{if } M = 1 \text{ and } a \geq m \text{ and } s = b. \end{cases}$$

Note that for households to migrate, their level of assets has to be larger than the moving cost. But, more importantly, the expected utility from migrating has to be larger than the utility coming from consumption of the certain income from the rural area. The reason is that if the second constraint is satisfied, then the first will be too. To see this, note that when households have exactly enough assets to pay the moving costs, $a = m$, then the bad state would leave them with zero consumption and hence negative infinity utility. Thus, it must be true that $a > m$ for agents to possibly prefer migrating than not migrating.

One can show that households are indifferent between migrating and staying when z and a satisfy

$$A_r + aR = [A_u z + (a-m)R]^p [(a-m)R]^{(1-p)}. \tag{15}$$

When a or z is higher, then the household strictly prefers to migrate. Otherwise the household strictly prefers to remain in the rural area.

Constrained Efficient Allocation. We next characterize a constrained planner's problem. In particular we allow the planner to redistribute resources within fixed (z, a) types but not across z or a types. The idea is that the planner may wish to pool resources and encourage more migration by providing better insurance against bad migration outcomes. For all agents with ability z, the planner solves:

Figure B.1: Migration and Misallocation with Migration Risk

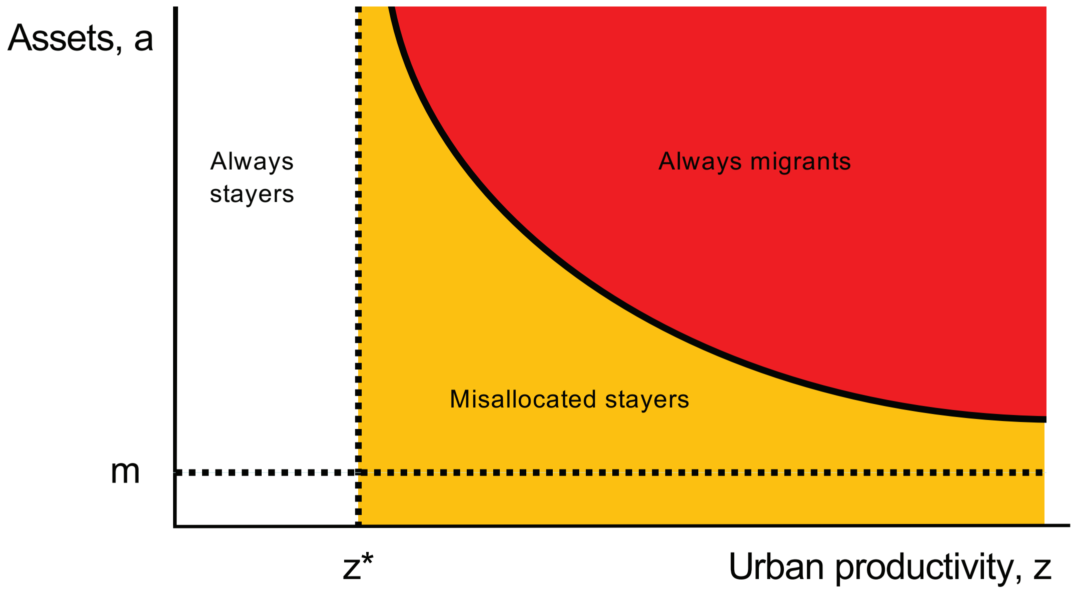

$$\max_{c,M} \log(c)$$
$$\text{s.t.} \quad c = \begin{cases} A_r + aR & \text{if } M = 0 \\ pA_u z + (a-m)R & \text{if } M = 1. \end{cases}$$

One can show that the planner does not care about assets, a, when making migration decisions. All households with z above a cutoff will migrate, and the planner will give all migrating agents the same consumption stream whether or not they have a bad migration outcome. The cutoff is determined by setting the income in the rural area equal to the expected income from migrating to the urban area:

$$A_r + aR = pA_u z + (a-m)R. \tag{16}$$

This implies that the cutoff is $z^* = (A_r - mR)/pA_u$. The consumption of each (z, a) type is determined by whether they migrate or not, which depends on whether their z is above or below the cutoff, their average income, and their asset endowment.

Figure B.1 illustrates the migration decisions in the competitive equilibrium and under the planner's solution. The red area in the figure represents the households that migrate in the competitive equilibrium, with the black curve representing the households who are indifferent between migrating or not, as in equation (15). The orange area plus the red area represents the set of households that migrate in the planner's solution, which are all those with urban productivity above the cutoff z^*. Thus, the red area represents those that always migrate, while the orange area are those that migrate only in the planner's allocation. As such, the orange area represents the households that *misallocated*, staying in the competitive equilibrium but migrating in the constrained efficient allocation. The white area are those that are always stayers, remaining in the rural area either in the market outcome or under the planner.

The outcomes illustrated in Figure B.1 illustrate how an asset transfer could reduce misallocation. Among households with z above z^*, those with higher asset levels are more likely to migrate. The misallocated households are those with relatively low assets. As such, in this economic environment, like in the model explored by Bryan et al (2014), an asset transfer to low-asset households could help reduce misallocation. Reducing misallocation can raise welfare, since the planner's allocation yields higher expected utility for any household type (z, a) than the competitive equilibrium allocation.

Pecuniary Externality

We next characterize a second exercise to analyze the role of pecuniary externalities in our model. In particular we assume that there is no migration risk and we further assume that there are decreasing returns to scale in the rural area. We assume that in the urban area productivity per efficiency units is given by A_u. Therefore, the production functions in the rural and urban areas are given by:

$$Y_u = A_u N_u \quad \text{and} \quad Y_r = A_r N_r^{\phi} \tag{17}$$

where N_u is the total number of efficiency units supplied in the urban area, N_r is the total number of workers in the rural area, and A_u and A_r are the exogenous productivities in two locations, and $\phi < 1$. Markets are competitive, and hence wages per efficiency unit are given by $w_u = A_u$ and $w_r = A_r \phi N_r^{\phi-1}$. We do not take into account the profits coming from the rural area, which we assume (for simplicity) belongs to some agent outside this economy.

Given that there are decreasing returns to scale we need to assume a particular distribution for z and a to get the wage in the rural area in equilibrium. Thus, we assume that z and a are iid and uniformly distributed with $a \sim U[0, \bar{a}]$, $z \sim U[0, \bar{z}]$.

Competitive Equilibrium. We begin by characterizing the outcome of the competitive equilibrium. Again letting $M \in \{0, 1\}$ be the moving decision, then the household's problem is defined as:

$$\max_{M\in\{0,1\}} c$$

s.t.

$$c = \begin{cases} A_r \phi N_r^{\phi-1} + aR & \text{if } M = 0 \text{ or } a < m \\ A_u z + (a-m)R & \text{if } M = 1 \text{ and } a > m. \end{cases}$$

This means that all households with $a < m$ cannot migrate not matter what z is. We assume that agents withe larger z values migrate first, which implies that the marginal z agent is given by:

$$U\left(A_r \phi N_r^{\phi-1} + aR\right) = U(A_u z - mR + aR).$$

Making use of the uniform distribution, the marginal z is given by:

$$A_r \phi \left(\frac{1}{\frac{m}{\bar{a}} + \frac{\bar{a}-m}{\bar{a}} \frac{z*}{\bar{z}}} \right)^{1-\phi} = A_u z^* - mR.$$

Note that marginal household only cares about income in each location and therefore she does not internalize the effect of migrating on the rural wage for all the rest of the agents.

Constrained Efficient Allocation. We next characterize a constrained planner's problem. In particular we allow the planner to lend resources to households with low levels of assets so they can migrate and furthermore the planner internalizes the wage decrease as agents migrate out of the rural area. For all z agents, the planner solves:

$$\max_{M} \int_i c_i di$$

$$\text{s.t.} \quad c_i = \begin{cases} A_r \phi N_r^{\phi-1} + a_i R & \text{if Mi=0.} \\ A_u z_i + (a_i - m)R & \text{if Mi=1} \end{cases}$$

$$N_r = \int_{(z,a)} M_i di$$

Therefore, in equilibrium the marginal household has a and z given by:

$$U\left(A_r \phi N_r^{\phi-1} + aR\right) = U(A_u z - mR + aR) + \int_{(z,a)\in R} U'(c_i) w'(N_t) dij$$

Figure B.2 illustrates the migration decisions that arise in the competitive equilibrium and un-

Figure B.2: Migration and Misallocation with a Pecuniary Externality

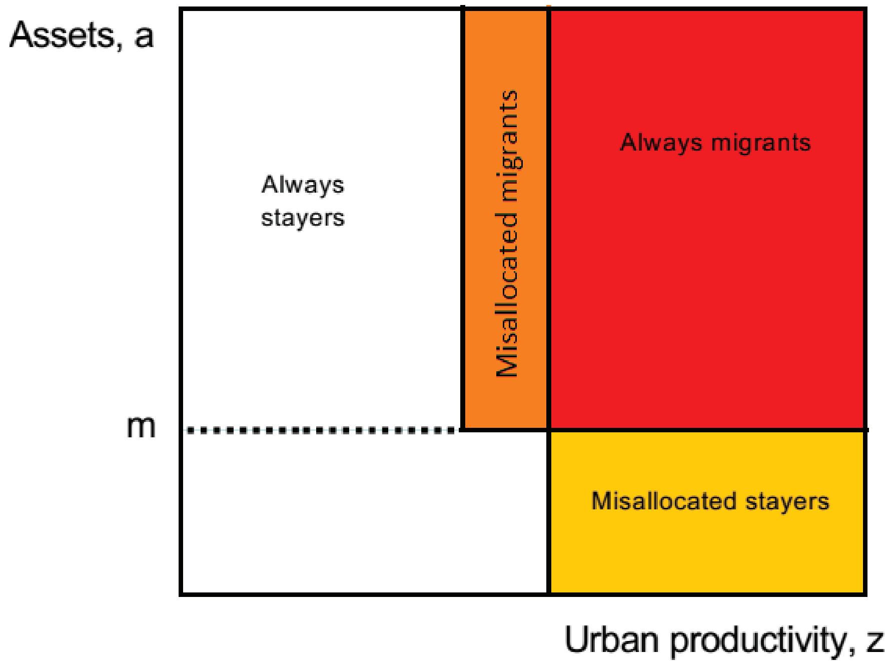

der the solution to the planner's problem. As before, agents with assets below m can be misallocated if their urban productivity is high enough. This is illustrated by the yellow region on the bottom right of the figure. Now there is a possibility of misallocated migrants as well. These are for agents with assets above m but sufficiently low productivity, illustrated by the orange region in the center of the figure. These misallocated agents arise because the planner allocates more agents to the urban area than does the competitive equilibrium allocation. The agents take only their own wages into consideration when deciding to migrate. The planner takes into consideration the wages of each agent plus the effects of each agents migration choice on the wages of the others.

C. Estimation and Quantitative Analysis

As explained in the text, we estimate the model using simulated method of moments. Following Gourieroux and Montfort (1996) and Terry (2019), one can write the asymptotic variance-covariance matrix of the estimated parameters as:

$$\Sigma = \left[M'WM\right]^{-1} M'W\Delta V\Delta WM \left[M'WM\right]^{-1},$$

where M is the Jacobian matrix representing the derivatives of each moment with respect to each parameter, W and Δ are weighting matrices, and V is the asymptotic variance-covariance matrix of the empirical moments targeted. We compute the M matrix using numerical differentiation taking the average changes in each moment to parameter changes of 0.75 percent, 1 percent and 1.25-percent. We set the matrix W to be the identity matrix, which puts each weight on all moments. We also set Δ to be the identity matrix, which puts each weight on the cross-sectional and experimental moments.

The asymptotic variance-covariance matrix of the targeted moments, V, is comprised of two parts: V_c, the variance-covariance matrix of our cross-sectional moments, and V_e, the variance-covariance matrix of the experimental moments. We compute V_c using the Rao-Wu (1998) rescaling bootstrap with replacement, which corrects for bias due to a finite number of primary sampling units per strata. For our experimental moments, our estimands are the true value of these moments internal to our experimental sample (instead of the population of Bangladesh, as with our cross-section). The uncertainty in our experimental moments comes only from our random assignment of treatment status. We therefore bootstrap V_e by resampling households with replacement stratified by treatment group. The intuition here is that the distribution of outcomes observed in the treatment group is a proxy for the distribution of treatment potential outcomes across our full sample. Our bootstrap is replicating our randomization by redrawing potential outcomes for treatment and control households from this proxy distribution. We assume the cross-sectional and experimental moments are independent, therefore the full variance-covariance matrix is $V = [V_c, 0; 0, V_e]$. The exact values for V_c, V_e and M are given in Tables A.1, A.2 and A.3.

D. Identification

In this section, we further illustrate how the experimental and cross-sectional moments help identify the parameters of the model. To do so, we start with the benchmark calibration and then compute the elasticity of each targeted moment to each parameter. Table D.1 reports these elasticities. For expositional purposes, we put in bold any elasticity greater than or equal to one in absolute value. It is useful to discuss the results in Table D.1 one parameter at a time, as well as the moments that are most sensitive to the change in parameters. In Appendix Table A.4, we report also the matrix proposed by Andrews, Gentzkow, and Shapiro (2017) that computes the sensitivity of each estimated parameter value to the value for each moment chosen.

Table D.1: Elasticities of Targeted Moments to Parameters

	θ	$\bar{u}$	λ	π	γ	A_u	σ_s	ρ	σ_ν	σ_{rc}	σ_{ui}
Con: Var(log cons Δ rural)	0.0	0.0	0.0	-0.0	0.0	-0.0	-0.3	-0.1	0.0	**1.0**	0.0
Con: Pr(households no assets)	-0.2	**-5.4**	0.0	-0.1	**-3.4**	**-1.3**	**-10.6**	**12.1**	-0.0	0.0	0.0
Con: Seasonal migrants	**-3.4**	**-7.9**	**-1.2**	**1.8**	**-1.8**	**1.2**	**-1.2**	0.5	0.5	0.0	0.0
Consumption, OLS	**3.9**	**3.6**	-0.5	**1.5**	**3.0**	**-4.2**	**-4.3**	**6.4**	-0.0	0.0	0.0
Migration, Treat-Control	-0.6	-0.9	0.9	**-1.1**	-0.4	**1.4**	**-1.5**	**1.8**	-0.3	0.0	0.0
Migration, Treat-Control, year 2	0.5	**-5.0**	**-1.5**	**3.8**	**-5.5**	**1.1**	**-2.5**	**4.9**	-0.9	0.0	0.0
Consumption, LATE	-0.6	-0.5	-0.0	0.4	-0.7	-0.3	**-2.7**	**2.6**	0.0	0.0	0.0
Control: Pr(repeat migration)	0.2	**-1.2**	-0.2	0.0	0.2	-0.1	**-1.3**	**1.7**	-0.2	0.0	0.0
Urban-Rural wage gap	**1.4**	0.1	0.2	-0.0	0.1	-0.2	-0.3	0.4	-0.0	0.0	0.0
Percent in rural area	0.8	-0.5	-0.1	0.2	0.1	**-1.3**	-0.3	0.4	0.0	0.0	0.0
Variance of log wages in urban	**1.4**	-0.1	-0.1	0.1	0.3	0.1	0.4	-0.4	0.0	0.0	**1.0**

Note: This table reports the elasticities of each targeted moment to each parameter, computed as the percent increase in each moment to a one percent increase in each parameter, starting from the estimated parameters of the model. Elasticities greater than or equal to one in absolute value are printed in bold.

Permanent productivity in the urban area, θ. The migration rate in the control group and OLS return to migration are most sensitive to θ. The intuition here is that θ controls how many households potentially have a comparative advantage in the urban area – i.e., how many "marginal households" there are. In the context of Figure 2, θ controls how many moderate z households there are relative to low z households. Changing θ therefore has a substantial impact on migration rates, as well as the productivity gains for those choosing to migrate.

Disutility in the urban area, $\bar{u}$, and the dynamics of experience, λ and π. The migration disutility term, $\bar{u}$, has the biggest impact on migration rates. Intuitively, if migrating to the urban area involves substantial disutility, then households will not migrate in general, and they will also not migrate in response to a migration subsidy. The comparison of panels (b) and (c) of

Figure 2 illustrates this point. The dynamics of experience are pinned down by migration rates and, especially, by the subsequent re-migration response in the years after the initial treatment. If the probability of staying inexperienced, λ, is larger, fewer households will want to migrate to begin with and fewer will end up repeat migrating. If the probability of remaining experienced, π, is larger, then migration rates in the control group and repeat migration rates will be larger.

Urban relative volatility, γ. The moments most sensitive to this parameter are the OLS return to migration, the probability of holding no assets and the probability of repeat migration. The link to the OLS return to migration implies that γ is being identified in part from the extent to which migrants are positively or negatively selected on transitory shocks. When γ is smaller, households induced to migrate are more likely to be those with lower transitory shocks and few asset holdings. If these households migrate, therefore, they are more likely to be those with the relatively lowest productivity draws in the urban areas. This, in turn, leads to a lower OLS coefficient of consumption migration, since those deciding to migrate have relatively lower consumption levels. The probability of having no assets and probability of repeat migration inform γ as higher urban risk means less migration and less use of migration as insurance.

Urban productivity, A_u and transitory shock process, ρ and σ_s. Urban productivity affects urban wages and, hence, migration rates and the percent of households that permanently locate in the rural and urban areas. It also effects the OLS returns to migration and probability of holding few assets. The asset distribution is informative about these parameters, as are migration rates and the OLS returns to migration. In the estimated model, the main effect of higher σ_s is to cause households to hold more assets and migrate less frequently due to the increased migration risk. The effect of a higher ρ is almost the exact opposite, as more autocorrelated shocks leave households more vulnerable to shocks but cause them to migrate more when assets and productivity are low.

Variance of idiosyncratic shocks and measurement error: σ_ν, σ_{rc} and σ_{ui}. The variance of idiosyncratic shocks, σ_ν, plays a minor role in all moments except the probabilities of repeat migration and migration in the control group. The reason is that idiosyncratic shocks affect the utility of migration directly, and for reasons unrelated to past migration. The latter two parameters inform only the measurement error moments by construction.