Under-investment in Profitable Technologies when Experimenting is Risky: Evidence from a Migration Experiment in Bangladesh

Gharad Bryan, London School of Economics
Shyamal Chowdhury, University of Sydney
A. Mushfiq Mobarak, Yale University*

*Contact Author: ahmed.mobarak@yale.edu, 203-432-5787

April 1, 2011
Preliminary and Incomplete

Abstract

The rural north-western districts of Bangladesh, home to 10 million people, experience a pre-harvest seasonal famine, locally known as Monga, with disturbing regularity. Inspired by the observations that wages are higher, jobs are more plentiful in nearby urban areas than in the monga-prone region, and that there are no official restrictions on mobility, we provide monetary incentives in 100 study villages to encourage people to seasonally migrate out in search of employment. We employ a randomized intervention design to study the patterns of responsiveness to our incentives using a 1900 household sample, which illuminates the current constraints to seasonal out-migration. The randomization also allows us to cleanly estimate the (causal) returns to migration in terms of household expenditures, savings and earnings, and caloric intake.

The propensity to seasonally out-migrate is *very* responsive to a 600-800 Taka cash/credit incentive, raising the migration rate to 57%, relative to 34% in a set of control villages. Simply providing information on job availability and wages at destination has no effect on the migration rate. Comparing the characteristics if migrants in the treatment and control areas, the incentives appear to induce people who otherwise feel less comfortable migrating – i.e. those without job leads or social network presence at the destination, and new migrants. Households closer to subsistence are less likely to migrate, but are much more responsive to the incentive.

The migrant experience was very successful on average: monthly consumption increased by at least Tk 300 ($4) per person per month, or Tk 1050 ($15) per household per month due to the induced migration. Earnings, savings and remittances from each migration episode are several multiples of the initial investment (the incentive, or the round-trip cost of moving). Most strikingly, migration rate in the treatment areas remained significantly higher in the treatment areas (47% vs. 35%) a year after the program, even after all incentives are removed. Induced migration in one year increases the propensity to re-migrate by 40 percentage points. There is evidence of learning among induced migrants, such as a greater propensity to re-migrate when the initial migration episode was more successful, greater growth in savings/earnings per day in treatment group, and larger consumption effect in 2009 among re-migrants.

These results are consistent with a model of a migration-based poverty trap, where individuals are uncertain about their own prospects at the destination, and particularly worried about a bad outcome (e.g. undertaking the costs of moving, but then not finding a job) during a period in which their family is under the threat of famine. The uncertainty associated with migration prevents households from investing, even when the expected returns are positive. Our intervention insures households against the bad outcome, thereby allowing them to invest, learn about their private returns to the investment, and for those with positive realizations, re-migrate the next period even in the absence of the intervention. Many of the empirical results are also consistent with a simple model of a liquidity/credit constraint preventing migration.

These results suggest that a grant program, or a credit program with limited liability (which amounts to insuring households against the possibility of a bad outcome) is likely to be welfare enhancing, and can be an effective policy response against the threat of localized seasonal famines such as Monga. More broadly, providing small grants or credit that enable households to search for jobs, and leads to a better spatial and seasonal matching between potential employers and employees may be a useful complement to the much-more-popular microcredit programs that aim to create new entrepreneurs and new businesses.

Underinvestment in Profitable Technologies when Experimenation is Costly: Evidence from a Migration Experiment In Bangladesh

Gharad Bryan, Shyamal Chowdhury and A. Mushfiq Mobarak

April 1, 2011

1 Introduction

The causes and consequences of internal migration have received comparatively little attention in the economics and policy literatures, even though it is much more common than the heavily-studied issue of international migration (Borjas 1999, Rosenzweig 2005, Yang 2008, Hanson 2009). There were 240 times as many internal migrants in China in 2001 as there were international migrants (Ping 2003), and 4.3 million people migrated internally in the 5 years leading up to the 1999 Vietnam census compared to only 300,000 international migrants (Ahn et al, 2003). The development effects of internal migration are arguably more profound: the majority of households in rural Rajasthan migrate seasonally in search of employment, and use this as the primary vehicle for diversifying sources of income (Banerjee and Duflo 2006).

This paper studies the causes and consequences of internal seasonal migration in north-western Bangladesh, in a region where over 5 million people below the poverty line must cope with a pre-harvest seasonal famine, known locally as Monga, almost every year (The Daily Star 2010). This seasonal famine is emblematic of widespread pre-harvest "lean" or "hungry" seasons throughout South Asia and Sub-Saharan Africa documented in a number of studies,[1] where imperfect consumption smoothing forces households below poverty during parts of the year. Inspired by the observation that nearby urban areas offer better wage and employment op-

[1]Seasonal poverty or hungry seasons have been documented in Ethiopia (Dercon and Krishnan 2000, who show that poverty and malnourishment increase 27% during the lean season), Malawi (Brune et al 2010), Madagascar (Dostie et al 2002, who estimate that 1 million people fall below poverty before the rice harvest), Kenya (Swift 1989, who distinguishes between years that people died versus years of less severe shortage), Thailand (Paxson 1993), India (Chaudhuri and Paxson 2002) and China (Jalan and Ravallion 1999).

portunities during the lean season, we provide small grant and loan incentives (of $8.50 or 600 Taka) in 100 study villages to encourage people to seasonally migrate out in search of employment. The random assignment of incentives allows us to generate among the first experimental estimates of the effects of migration, and internal migration in particular.Estimating the returns to migration is the subject of a very large literature, but one that has been hampered by difficult selection issues (Akee 2010, Grogger and Hanson 2010, McKenzie et al 2010). The internal seasonal migration we encourage is a commonly used means to cope with seasonal poverty and natural disasters across Asia, Africa and Latin America,[2] and our estimates of the effects of migration thus have substantial external relevance.

We estimate the causal effects of being induced to migrate away during the lean season to be very large.These estimates add to an emerging literature that documents very high rates of return to small capital investments in developing countries (Udry and Anagol 200x, De Mel, McKenzie and Woodruff 200x and 200x and Duflo, Kremer and Robinson 2009), and bolsters the case made by Clemens and Pritchett (20xx), Rosenzweig and xx (200x) and Gibson and McKenzie (2010) that offering migration opportunities improves welfare by much more than any other development intervention in health, education or agriculture that has been studied. Migration induced by our intervention increases food and non-food consumption of the migrants family members remaining at the origin by 30-35%, and improves their caloric intake by 700 calories per person per day. On an initial investment of about $6-$8 (the average round-trip cost to a destination), migrants earn $110 on average during the lean season and save about half of that, which are suggestive of a very high rate of return on investment.Most strikingly, these induced migrants continue to re-migrate at a higher rate compared to a control group a year later, even after the inducement for migration is removed.These large positive returns, consumption effects and preferences revealed by the voluntary re-migration beg one very important question: Why didn't these people already engage in such a highly profitable investment?

To understand this puzzling behavior, we propose a model in which experimenting with a new technology or behavior such as migrating to a new destination - is risky, and rational households do not migrate in the face of uncertainty about their prospects at the destination even when they expect monetary returns to the activity to be positive. This results in a poverty trap in which households do not migrate out of fear of an unlikely, but devastatingly negative outcome where they undertake the cost of moving but return hungry after not finding employment during a period when their family is under the threat of famine. Inducing the inaugural

[2]Prior attempts use controls for observables (Adams 1998), selection correction methods (Barham and Boucher 1998; Acosta et al 2007), instrumental variables methods (BenYishay 2010, McKenzie and Rapoport 2007, Brown and Leeves 2007, Yang 2008) and natural policy experiments (Clemens 2010, Gibson et al 2010) to answer this question.

migration by insuring against this devastating outcome which our grant or loan with implied limited liability managed to do can lead to long-run benefits where households either learn how well their skills fare at the destination, or improve their future prospects by allowing employers to learn about them. This simple model can explain the high take-up rate for the intervention, the large positive income and consumption effects, and the greater re-migration in a future period among treatment households even after the inducement is removed.

The proposed model also makes four other predictions regarding migration behavior. First, households that are close to subsistence – on whom experimentation with the new technology imposes the biggest risk – should be less likely to migrate and more likely to benefit from our intervention. Second, households that do not know someone at the destination and therefore do no have the contacts necessary for successful migration should be less likely to migrate and more likely to respond to our incentive. Third, because the poverty trap is generated by the possibility of a large negative return and that possibility can be effectively mitigated through a loan that does not have to be paid off if the migration episode was not successful, we predict that cash and credit incentives should have roughly the same effect on the migration rate. Fourth, the possibility of learning implies that those households that are most successful should be those most likely to repeat migrate and that repeat migration will be to the same location as prior migration. We find support for all of these conclusions in our data, the last one tested by using additional exogenous variation in migration location induced by conditionalities placed on migration in our experiment.

Households reluctance to engage in risky or costly experimentation can provide insight into a number of other important puzzles in growth and development. Green revolution technologies led to dramatic increases in agricultural productivity in South Asia (Evenson and Gollin 2003), but take-up and diffusion of the new technologies was surprisingly slow, partly due to low levels of experimentation and the resultant slow learning (Munshi 2004). More directly, research has linked the inability to experiment due to uninsured risk to a bias towards low risk low-return technologies that stunt long-run growth (Yesuf et al 20xx), and to reduced investments in agricultural inputs and technologies such as new high-yield variety seeds and fertilizer (Rosenzweig and Wolpin 1993a, Dercon and Christiansen 2009). Aversions to experimentation can also hinder entrepreneurship and business start-ups and growth (Hausmann and Rodrik 2003; Fischer 2009), and can force sales of productive investment assets in order to smooth consumption (Rosenzweig and Wolpin 1993b). We contribute to this literature by using experimental variation in data to build evidence on the presence of a poverty trap induced by inefficiently low levels of experimentation.

Our model and experimental results shed light on a fundamental puzzle in development economics, which is that adoption rates of highly efficacious technologies with the potential to address important development challenges - ranging from tropical diseases to low agricultural productivity to low rates of savings and investment - have remained surprisingly low (Miller and Mobarak 2010, Null et al 2010). If adoption is risky (e.g. due to risk of crop failure, or uncertainty about durability of a new stove or water purifying technology) then giving households the opportunity to experiment with the new technology by insuring against failure may be an effective marketing strategy.

Finally, from a narrow policy perspective, our experiments uncover a cost-effective response to the widespread famines that afflict the 10 million people residing in Rangpur region of Bangladesh with disturbing regularity. Such lean seasons are not phenomena unique to northwestern Bangladesh: there is a predictable pre-harvest hungry season every year in Malawi, Ethiopia, Madagascar and other countries, forcing millions to succumb to seasonal poverty. The solution we implement is inexpensive (the vast majority repay the $8.50 inducement when it is offered as a loan), it confers long-run benefits even when offered as a one-off, and it is more sustainable than subsidizing food purchases on an ongoing basis, which is the major anti-famine policy tool currently employed by the Bangladesh government (Khandker 2010; Government of Bangladesh 2005, journalistic cite). Our intervention mitigates the spatial mismatch between where people are based, and where the jobs and excess demand for labor are during the pre-harvest months. This approach may be of relevance to other countries that face geographic concentrations of poverty, such as northern Nigeria, eastern islands of Indonesia, northeast India, southeast Mexico, and inland southwest China (Jalan and Ravallion 2002). Of course one has to be careful about general equilibrium effects at the destination when such a program is scaled up, but instances of excess labor demand during particular seasons is not altogether uncommon even in developing countries. Our findings suggest that providing credit to enable households to search for jobs, and aid spatial and seasonal matching between potential employers and employees may be a useful way to augment the microcredit concept currently more narrowly focused on creating new entrepreneurs and new businesses.

The next two sections describe the context and the design of our interventions. We present results on program take-up and the effects of migration in Section 4. These findings motivate the risky experimentation model in Section 5. We present statistical tests of various implications of the model in section 6, discuss alternative explanations of the data in Section 7 and offer conclusions and policy advice in section 8.

2 The Context: Northwestern Bangladesh and the Monga Famine

Our experiments are conducted in 100 villages in two districts (Kurigram and Lalmonirhat) in the seasonal-famine prone Rangpur region of north-western Bangladesh. The Rangpur region is home to roughly 7% of the countrys population, or 9.6 million people. 57% of the regions population (or 5.3 million people) live below the poverty line.[3] In addition to the level of poverty, the Rangpur region experiences more pronounced seasonality in income and consumption, with incomes decreasing by 50-60% and expenditures on food dropping by 10-25% during the post-planting and pre-harvest season (September-November) for the main Aman rice crop (Khandker 2010). As Figure 1 indicates, the price of rice also spikes during this season, particularly in Rangpur, and thus kilos of rice consumed drops 22% even as households shift monetary expenditures towards food (Figure 2) while waiting for the Aman rice harvest. The lack of job opportunities and low wages during the pre-harvest season and the coincident increase in grain prices combines to create a situation of seasonal deprivation and famine (Mahmud and Khandker 20xx, Sen 1981)[4]. The famine occurs with disturbing regularity and thus has a name Monga. It has been described as a routine crisis (Rahman 1995), and its effects on hunger and starvation widely chronicled in the local media. Agricultural wages in the Rangpur region are already among the lowest in the country over the entire year (BBS Monthly Statistical Bulletins), and further, demand for agricultural labor plunges between planting and harvest. The resultant drastic drop in purchasing power for Rangpur households reliant on agricultural wage employment takes (or threatens to take) consumption below subsistence.

[Figure 1 about here.]

[Figure 2 about here.]

Several puzzling stylized facts about household and institutional characteristics and coping strategies motivate the design of our migration experiments. First, seasonal out-migration from the monga-prone districts appears to be low despite the absence of local non-farm employment opportunities. According to the nationally representative HIES 2005 data, only 5 percent of households in Monga prone districts receive domestic remittances, while 22 percent of all

[3]Extreme poverty rates (defined as individuals who cannot meet the 2100 calorie per day food intake even if they spend their entire incomes on food purchases only) were 25 percent nationwide, but 43 percent in the Rangpur districts. Poverty figures are based on Bangladesh Bureau of Statistics (BBS) Household Income and expenditure survey 2005 (HIES 2005), and population figures are based on projections from the 2001 Census data.

[4]Amartya Sen (1981) notes these price spikes and wage plunges as important causes of the 1974 famine in Bangladesh, and that the greater Rangpur districts were among the most severely affected by this famine.

Bangladeshi households do. Remittances generally under-predict out-migration rates, but the size and direction of this gap (lower migration out of the poorest district) are still puzzling. It is more common for agricultural laborers in other regions to migrate in search of higher wages and employment opportunities, and this is known to be one primary mechanism by which households diversify income sources in India (Banerjee and Duflo 2006).

Second, inter-regional variation in income and poverty between Rangpur and the rest of the Bangladesh have been shown to be much larger than the inter-seasonal variation within Rangpur (Khandker 2010). This suggests smoothing strategies that take advantage of inter-regional variation (i.e. migration) rather than inter-seasonal variation (e.g. savings, credit) may hold greater promise. Moreover, an in-depth case-study of the Monga phenomenon (Zug 2006) explicitly notes that there are off-farm employment opportunities in rickshaw-pulling and construction in nearby urban areas during the monga season. To be sure, Zug 2006 points out that this is a risky proposition for many, as labor demand and wages drop all over rice-growing Bangladesh during that season. However, this seasonality is less pronounced than that observed in Rangpur (Khandker 2010).

Finally, both government and large NGO monga-mitigation efforts have concentrated on direct subsidy programs like free or highly-subsidized grain distribution (e.g. "Vulnerable Group Feeding,"), or food-for-work and targeted microcredit programs. These programs are expensive, and the stringent micro-credit repayment schedule may itself keep households from engaging in profitable migration (Shonchoy 2010). There are structural reasons associated with rice production seasonality for the seasonal unemployment in Rangpur, and thus encouraging seasonal migration towards where jobs are appears to be a sensible complementary policy to experiment with.

3 Design of Interventions and Experiment

The two districts where the project is conducted (Lalmonirhat and Kurigram) represent the agro-ecological zones that regularly witness the monga famine (World Bank Umar cite). We randomly selected 100 villages in these two districts and first conducted a village census in each location in June 2008. Next we randomly selected 19 households in each village from the set of households that reported (a) that they owned less than 50 decimals of land, and (b) that a household member was forced to miss meals during the prior (2007) monga season. We conducted a baseline survey of these 1900 households during the pre-monga season in July 2008.

In August 2008 the researchers randomly allocated the 100 villages into four groups: Cash, Credit, Information and Control using a pure random number generator in Stata. These treatments were subsequently implemented in collaboration with PKSF[5] and their NGO partner organizations with field presence in the two chosen districts. We trained the NGOs on the implementation procedure in August 2008, and they implemented the interventions during the 2008 Monga season starting in September. 16 of the 100 study villages (consisting of 304 sample households) were randomly assigned to form a control group. A further 16 villages (consisting of another 304 sample households) were placed in a job information only treatment. These households were given information on types of jobs available in four pre-selected destinations, the likelihood of getting such a job and approximate wages associated with each type of job and destination. 703 households in 37 randomly selected villages were offered cash of Taka 600 (∼US$8.50) at the origin conditional on migration, and an additional bonus of Taka 200 (∼US$3) if the migrant reported to us at the destination during a specified time period. We also provided exactly the same information about jobs and wages to this group as in the information-only treatment. Taka 600 covers a little more than the average round-trip cost of safe travel from the two origin districts to the 4 nearby towns on which we provided job and wage information. The 589 households in the final set of 31 villages were offered the same information and the same Tk 600 + Tk 200 incentive to migrate, but in the form of a zero-interest loan to be paid back at the end of the monga season (in December) rather than a cash grant.

In the 68 villages where we provided monetary incentives for people to seasonally out-migrate (37 cash + 31 credit villages), we sometimes randomly assigned additional conditionalities to different households within the village. For example, 12 of the 19 households were required to migrate in groups, and specific group members were assigned in half the case. Also, about half the households were required to migrate to a specific destination. We will not directly analyze the effects of such conditionalities in this paper, but these conditionalities were helpful in creating random within-village variation in specific households propensity to migrate. This random variation will be useful for some of the empirical analysis presented later. Figure 8 provides an overview of the randomized intervention conditions.

[5] PKSF (Palli Karma Sahayak Foundation) is an apex micro-credit funding and capacity building organizations in Bangladesh. It is a company not for profit set up by the Government of Bangladesh in 1990, for poverty alleviation through the provision of micro-credit through its Partner Organisations (POs).

4 Returns to Seasonal Migration

In this section we report simple "program evaluation" results on take-up of the treatment, the effects of seasonal migration on the households consumption and welfare at the origin, income and savings of the migrant at the destination, and the propensity to re-migrate in 2009 after incentives are removed. After establishing the large, positive returns to seasonal migration, in subsequent sections we turn to the question of why households were not engaging in this profitable behavior to begin with.

[Table 1 about here.]

Table 1 examines program take-ups for two incentivized groups – cash and credit put together, and out-migration rates between incentivized groups and not-incentivized groups – control and information put together. Against a migration rate of 58% for the incentivized group (after six months of incentives offered), the migration rate for the not-incentivized group stood at 36%. The statistical analysis conducted in the various panels of Table 1 show that the effect of the cash and credit incentives are highly statistically significant and that providing information has essentially a zero effect on migration propensity. There is also no difference between providing cash and providing credit – households appear to react very similarly to either incentive, we therefore combine the impact of these two treatments for much of our analysis. We return to provide a positive explanation of this fact in Section 5.

Table 2 shows the effects of migration on consumption expenditures and other expenditures amongst remaining household members during the monga season. Consumption is an ideal measure of the benefit of migration, aggregating as it does the impact on the family of migrating. Using such a measure is important as it overcomes the problems associated with measuring the true costs and benefits of technology adoption that are highlighted, for example, in Foster and Rosenzweig (2010). The effects are calculated from a regression where the choice to migrate is instrumented with whether or not a household was randomly placed in the incentive group. These estimates therefore show the impact of migration on those households that were induced to migrate by our intervention (that is the LATE).

[Table 2 about here.]

Migration of a household member during the monga season has substantial impact on the remaining household members well-being. Compared to non-migrant households, per capita food, non-food, and caloric intake among (induced) migrant households increase by 30% to

35%, and their monthly consumption expenditure increase by at least $4 per capita ($15 per household) due to induced migration. There are changes towards higher quality diets as food consumption shifts towards meat and child education expenditures increase among migrant households. There is also a statistically significant increase in non-food expenditures on clothing and shoes, and expenditures of health for male members of households.

[Table 3 about here.]

[Table 4 about here.]

The positive consumption effects on the remaining household members come from remittance that migrant members send from their earnings and saving. Table 3 (which should be compared to Table 4) breaks down the effects. The incentivized migrants earn $110 on average during the lean season and save about half of that, and the average savings plus remittance is about a dollar a day. Non-incentivized migrants earn more per episode, save and remit more per day, which indicates our induced migrants earning potentials were less than the average. However, on an incentive of about $6-$8, the induced migrants earn $110 on average during the lean season and save and remit more than half of that, which is a very high rate of return on per dollar incentive.However, not all migrants were successful in finding jobs at destinations, earnings and remitting to their households at origin. About 20% of the migrants were unsuccessful which had consequence on remigration possibilities in subsequent year.

Our one-time incentive to migrate was offered in 2008 prior to the lean season. As shown in Table 1, in 2009 47% of the incentivized groups re-migrated even after the incentive was taken away. The comparable figure for the non-incentivized group was 37% that shows almost no change on year-to-year basis in migration rate in the control and information groups. We find that if a household was induced to migrate in 2008, it doubles the chances (35-45 percentage point effect) that it will migrate in 2009.

5 A Model of Risky Experimentation

In the previous section we documented three facts regarding our intervention. First, a large number of households were motivated to migrate in response to a relatively small incentive (600 Taka). Second, there was a positive average return to migration, indicating that households were not migrating despite a positive expected profit. Third, a large portion of the households that were incentivized to migrate in year one continued to send a seasonal migrant in year two and continued to earn a high return to this activity. We consider these to be important observations, they document a situation in which a small transfer has a meaningful and lasting impact on the behavior and well being of poor households.

In this section we propose a (very) simple model of a poverty trap, and argue that it can explain the three main findings from the previous section. The basic model is a much stripped down version of the classic models of Kihlstrom and Laffont (1979) and Banerjee and Newman (1991) and can be easily extended along the lines of Banerjee (2004) to encompass a more traditional view of poverty traps. To this extent we see our results as lending support to the basic proposition that poverty traps can exist. We then use the model to determine additional predictions that should hold in the data. In section 7 we discuss alternative theories that can also generate similar facts and argue that our interpretation is to be preferred on several dimensions.

Consider a household that lives for an infinite number of periods and in each period decides between staying at home and earning a certain income of y and sending a migrant and earning an income of $y + b$ with probability $\mu(b)$ and $y + g$ with probability $\mu(g)$. We assume that $g > b$ and we also assume that after one migration episode the household becomes informed about whether the return is g or b. The best way to interpret the return in this context is that successful migration (g) requires connections at the destination. Before the first migration episode the household is unsure whether it will be able to find a connection, however, once the connection is in place it is permanent.

We assume that the household is a discounted expected utility maximizer with utility function u (that obeys all the usual assumptions), discount factor δ and that b and g are such that:

$$\begin{aligned} &u(y+g) > u(y); \textit{and} \\ &u(y+b) < u(y). \end{aligned}$$

Given these assumptions the return to migrating is

$$V^m = \frac{\mu(g)u(y+g) + \mu(b)((1-\delta)u(y+b) + \delta u(y))}{1-\delta},$$

while the return to staying at home is

$$V^h = \frac{u(y)}{1-\delta}.$$

Figure 3 shows a pretty typical plot of V^m and V^h as a function of baseline income y where V^m is in pink and V^h in blue. While the two curves need not cross, if they do it is usually true that they cross only once with V^m being optimal above some threshold income y^*.[6] It is much more likely that the curves will cross like this if the outcome b is sufficiently bad so that $u(y+b)$ is in the very steep part of the utility function. In the example used to generate Figure 3 this requires that $y+b$ be close to 0, the point at which the derivative of the ln utility function becomes infinite.

Given that the conditions for a single crossing are met, the model generates a poverty trap. Households with income $y < y^*$ never migrate, they consume their initial income y forever. On the other hand, households that have income of $y \geq y^*$ migrate and learn their migration status. If they are able to find a connection they earn $y+g$ forever after and if they are not able to find a connection they return to consuming y. The long run distribution of income is therefore as depicted in Figure 4. Thus we have a very simple poverty trap, those with initial incomes below a certain level stagnate at that level while those with starting incomes above that point are able to take advantage of migration and push their long run income up by g.

[Figure 3 about here.]

[Figure 4 about here.]

Of course, in reality incomes vary across years and those that are close in average income to the cutoff y^* will eventually jump from the low income equilibrium to the high income equilibrium. Nevertheless, poverty trap models such as this one provide a role for policy in helping those household whose income is sufficiently low that it is unlikely that they will move between the equilibria. Further, if the return g is large enough, a policy that helps households make the transition can lead to a modest transformation in people's lives.

Given this background, it seems reasonable to assume that a policy maker may wish to assist households in making the jump between equilibria, and it is natural to ask how this is best achieved. Figure 5 compares three policies. The first panel shows the baseline income threshold y^*, the second panel considers a once off transfer, the third an incentive similar to our "cash" treatment and the fourth an "insurance" program that provides a transfer only if the household migrates *and* realizes state b. This fourth panel considers a policy that is similar to a limited liability loan and is formally similar to our "loan" treatment. Two points should be noted. First, because the transfer increases both V^m *and* V^h it has a much more limited impact on y^*. The alternative incentive treatment is much more effective in encouraging migration. Second, because the poverty trap is driven almost entirely by the possibility of a near subsistence outcome in state b, the insurance treatment realizes most of the benefits of the incentive policy. The insurance policy is also much more cost effective; the incentive must be paid with certainty,

[6] Conditions for the crossing can be found in Banerjee (2004) for example.

while the insurance only pays out with probability $\mu(b)$. To translate this into our practical setting, the loan was repaid in roughly 85% of cases and the total cost of the loan treatment (excluding operating costs) is therefore only 15% of the cost of the incentive treatment.

[Figure 5 about here.]

These two observations translate into two important implications for our experiment. First, we can quantify what it means for the impact of the incentive to be "large". Given the parameters used in Figure 5 the transfer reduces y^* from 0.275 to 0.259. On the other hand, the incentive decreases y^* to 0.241, that is, the same size incentive has roughly twice the impact.[7] Another way to see this is to note that an incentive of this size has the same impact as an income transfer of over twice the size. We, therefore, expect our incentive scheme to have a large impact relative to yearly income variation that is roughly the same size, and this gives scope for an incentive policy to be cost effective in moving households between equilibria. Second, because the insurance policy realizes the lions share of the gain from the incentive policy, we predict that empirically our "cash" and "loan" treatments will have a similar size impact on the migration rate.

How general are these results and when would we expect them to hold? The situation considered in Figure 5 has a particular feature. Household income y^* is very low, in that it falls in the very steep part of the utility function used for the example. This implies that the transfer policy has a relatively large impact on V^h and therefore that the effect of the incentive is large relative to the effect of the transfer. If, however, y^* is high relative to subsistence then this will no longer be the case. Figure 6 illustrates. The left panel shows how a transfer will affect the curves V^m and V^h, while the right panel illustrates the impact of an incentive. Because income in this example is relatively far from the subsistence point (which is 0 given the ln utility function used) the transfer has only a small impact on V^h and as a result the incentive and transfer policies have a similar impact. In such a context it seems unlikely that our incentive would have much of an impact as it is small relative to yearly variation in income and therefore most of the households that are far from subsistence that would be affected by our policy will already have moved to the high income equilibrium. This reasoning combined with the fact that the overall impact of our incentive on V^m will be largest when y is small (due to concavity of the utility function) leads us to predict that our treatments will have a larger effect on those households that are close to subsistence.

[Figure 6 about here.]

[7] This assumes that households are uniformly distributed with respect to income.

Our model is able to generate such a stark poverty trap because we assume that there is learning about the state of the world after one migration episode. We make this assumption for a simple reason. In our experiment a large portion of those households that were induced to migrate in round 1 continue to migrate in round 2 (roughly 50%). Put simply it is difficult to generate this degree of remigration in a model without learning. Again, this is because the earnings from migration operate similar to a transfer policy and the payoff to migration would therefore have to be large relative to the yearly variation in income. This seems unlikely because the incentive offered (600 Taka) is small relative to yearly variation in income and it seems unlikely that the return to that 600 Taka investment is more than 100%. The inclusion of a learning effect also generates two additional implications. First, if we interpret learning as finding out if you can establish a network connection, and then keeping that connection we should see that the incentive has most of its impact on those that do not already have network connections at a destination. Second, we should see location specific learning. That is, households that migrate to a particular location (for exogenous reasons) should continue to migrate to that location.

We plan to test the four empirical predictions of the model in the next section, so we collect them here for reference:

Predictions 1. *If our three main results are driven by a poverty trap model as presented in this section then:*

1. *Cash and Credit treatments should have a similar impact on migration rates;*
2. *The incentive should have a larger impact on households that are close to subsistence;*
3. *The incentive should have a larger impact on households that do not have network connections at the destination; and*
4. *Households should exhibit location specific learning.*

The model and these four predictions also help us to understand the type of situation in which we would expect incentive and insurance policies to lead to the sorts of long term effects observed in our experiment. We should look for situations in which there is household specific learning, where households are close to subsistence and where there is a reasonable chance of a negative outcome from investment. Prediction 2 also provides an answer to the puzzle which motivated this work – i.e. why does Rangpur have such a low migration rate relative to the rest of Bangladesh? Our model suggests that this is *because* households in Rangpur are close to subsistence at the time when it makes most sense to migrate.

6 Additional Tests of the Model

In this section we provide evidence that is consistent with the for predictions from our model.

The first prediction of our model has already been verified. The migration rate in the cash treatment is 58.96% while in the credit treatment it is 56.8% and the difference is not statistically significant. This is consistent with the claim that the cash and credit treatments should have similar impacts on the migration rate, with the credit treatment being slightly less effective.

The second prediction of the model is that those who are close to subsistence should be less likely to migrate and that the incentive should have the greatest impact on those households that are close to subsistence. We use the portion of household expenditure that goes toward food as a measure of how close the household is to subsistence. Figure 7 provides kernel density plots of the distribution of this index among households that migrate and those that do not migrate. The left hand panel shows the plot among the control villages. It is clear in these diagrams that households for that are close to subsistence are far less likely to migrate. The right panel shows the same plot amongst the treatment villages. In these villages there is no such discrepancy – essentially the migrant and non-migrant households look identical on this measure in the treatment villages. These two plots are, therefore, consistent with the model.

[Figure 7 about here.]

The third implication of our model is that the incentive should have a larger impact on those households that do not have connections at the destination and therefore have more to learn – or equivalently, have not yet learned whether they are g or b households. Table 5 provides evidence on this conjecture. The table breaks migration episodes within the first round data down with the first migration episode presumably being the one most affected by our incentive. The table shows that the incentive led to an increase in the number of households migrating that do not know someone at the destintion and that do not already have a job lead at the destination. These observations are again consistent with our model.

[Table 5 about here.]

Finally, our model predicts learning. We make two predictions here. First, households that are successful in their first year should be more liklely to migrate in the second year and this effect should be more pronounced among the incentivized migrants. Second, one of our treatment requires households to send their migrant to a particular location. Our model predicts that remigration should be disproportionately back to this same location.

Table 6 provides evidence in favor of the first question. Among the incentivized villages, households that remigrate are those that were successful in the first round – i.e. those that have higher savings and earnings in the first round. In contrast there is a much weaker different in the impact of success among non-treatment villages. This is consistent with the model and the idea that the treatment households have something to learn, while the non-treatment villages do not.

[Table 6 about here.]

Table 7 provides evidence on the second conjecture. It shows results of a multinomial logit and confirms that households that were successful in a particular location were likely to return to that location.

[Table 7 about here.]

7 Alternatives

In this section we discuss alternative explanations for our empirical results. The most plausible alternative explanation is that, during the lean season, households face a strong liquidity constraint – they simply do not have the 600 Taka required to pay for a bus ticket and therefore cannot migrate. Relieving the liquidity constraint in one year could lead to ongoing migration either through learning as discussed above or because increased income relaxes the liquidity constraint again the following year.[8] Such an explanation can account for the majority of our facts. A large portion of the population may be liquidity constrained in any given year leading to the large effect of the incentive, the return to migration can be positive and yet people do not migrate because they lack can, the impact of cash and credit treatments will be the same and, to the extent that subsistence is correlated with the liquidity constraint, we would expect the incentive to have a greater impact on those close to subsistence.

Nevertheless we do not thing a liquidity constraint can explain our results. Given the strong repeat impact on the second year, and the small size of the transfer, a liquidity constraint story would not hold over time. That is, we would expect all households to have received a sufficiently strong positive income shock in the past to have overcome the liquidity constraint and moved to the rich equilibrium prior to our intervention. For example, year on year average absolute deviation in *weekly* expenditure in our sample is 307 Taka between rounds one and two

[8] It should be noted that some form of credit constraint is required for the poverty trap model above to work. The liquidity constraint discussed here is a more severe form of credit constraint.

and 368 Taka between rounds two and three. The standard deviation of absolute deviation in consumption is 635 and 508 Taka respectively. While part of this is no doubt driven by measurement error, these figures show what seems obvious, season on season variation in incomes likely swamps the 600 Taka of our incentive. This criticism applies to all mechanisms that treat the incentive as a one time transfer. Further to this basic criticism, the lean season in this region lasts for around 3 months. It is hard to see how any household that has sufficient funds, either as storage or recurring wages, to survive this period would be unable to find 600 Taka. This may be a very costly choice, but it will be possible. Liquidity constraints therefore seem to be counter factual.

The above discussion applies equally to any mechanism that does not emphasize the incentive effect of our treatments. This observation, however, raises a perhaps troubling point for our preferred model. The basic argument against 600 Taka having a large effect in the form of a transfer is that yearly variation in income is larger. Therefore, for our incentive to work it must be the case that many households would prefer to keep a once of transfer of 600 Taka rather than spend it one migration. But this implies that the welfare impact of our intervention is bounded above by the benefit of giving each household 600 Taka to use as they please *and* the households not migrating. This point is of course obvious in a fully rational model because households always prefer to have money without constraints. We leave it to the reader to determine whether they believe that it is reasonable to assume this. Alternative models that posit less rationality, either in the form of excessive impatience (δ low), non-rational expectations ($\mu(b)$ too low) or some reason why the household is consistently on the poverty line despite variation in income (for example hyperbolic disounting á la Laibson (1997) and Duflo et al. (2009)) can make some room for policy makers to believe that the welfare gains of this intervention are larger than the purely rational model would imply.

8 Conclusions

We conducted a randomized experiment in which we incentivized households in Rangpur to send a seasonal migrant to an urban area. The main results show that a small incentive led to a large increase in the number of seasonal migrants, that the migration was successful in the sense that it led to an average increase in consumption of around 35% and that households given the incentive in one year continued to be more likely to migrate in the following year.

We argue that the results are consistent with a simple (rational) model of a poverty trap where households that are close to subsistence are unable to learn whether migration is successful due to a small possibility that migrating will turn out badly, leaving households consumption below subsistence.

From a policy perspective, the results support a policy of providing micro-insurance that mitigates the potential downside of experimentation. Our analysis also suggests that such a policy will be particularly beneficial where households are close to subsistence, risk averse and could benefit from a technology that requires individual level experimentation in order to determine profitability.

[Table 8 about here.]

References

Banerjee, A., "The Two Poverties," in Stefan Dercon, ed., *Insurance against poverty*, Oxford Scholarship Online Monographs, 2004, pp. 59–76.

Banerjee, A.V. and A.F. Newman, "Risk-bearing and the theory of income distribution," *The Review of Economic Studies*, 1991, *58* (2), 211.

Duflo, E., M. Kremer, and J. Robinson, "Nudging farmers to use fertilizer: theory and experimental evidence from Kenya," *NBER Working Paper*, 2009.

Foster, A.D. and M.R. Rosenzweig, "Microeconomics of technology adoption," *Annu. Rev. Econ.*, 2010, *2* (1), 395–424.

Kihlstrom, R.E. and J.J. Laffont, "A general equilibrium entrepreneurial theory of firm formation based on risk aversion," *The Journal of Political Economy*, 1979, *87* (4), 719–748.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 477 1997, *113* (2), 443.

List of Figures

Figure 1: Seasonality of Rice Prices

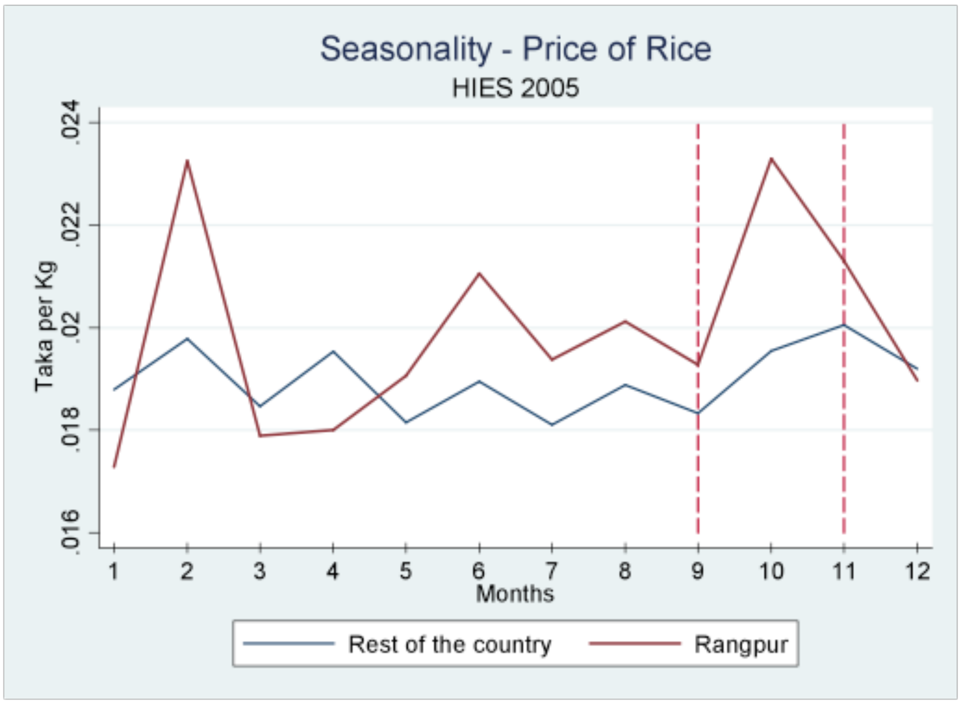

Figure 2: Seasonality of Food Expenditure

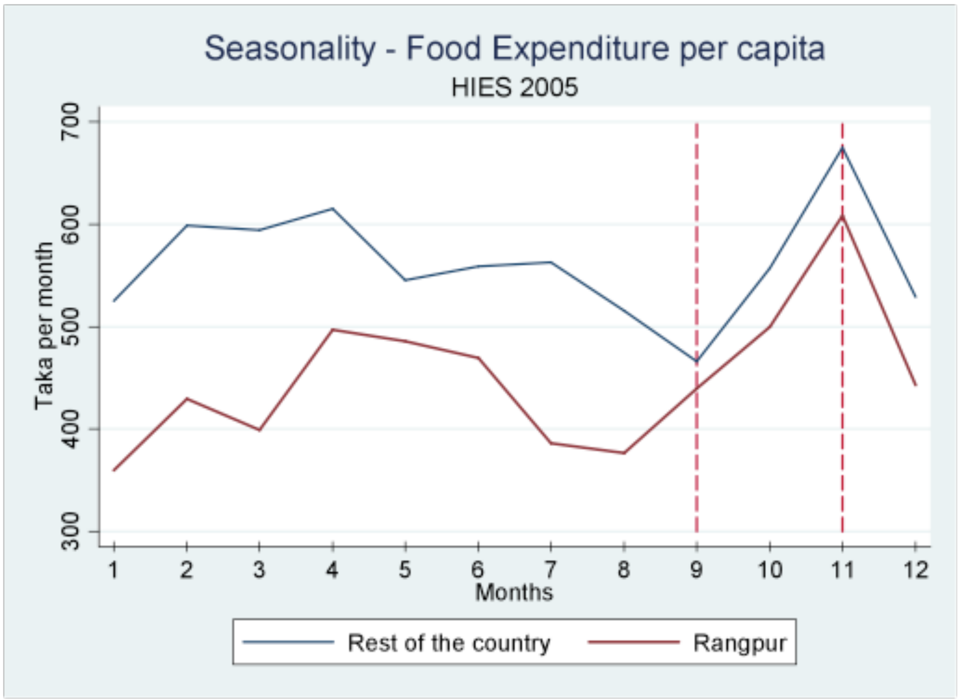

Figure 3: The return to migration as a function of income

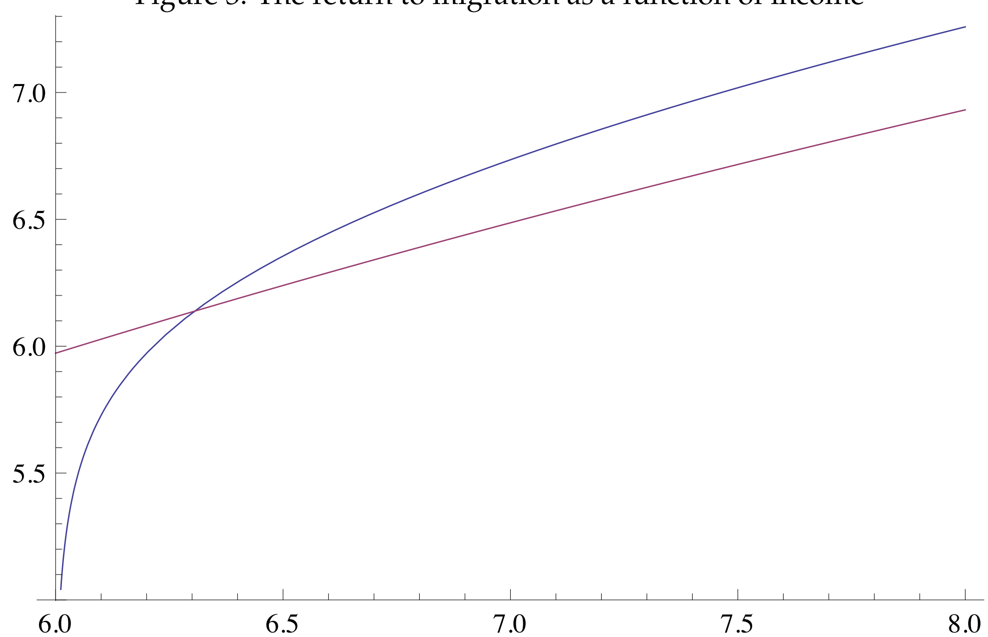

$u = \ln$, $\mu(g) = 0.7$, $b = -6$, $g = 3$ and $\delta = 0.7$.

Figure 4: Long run income distribution

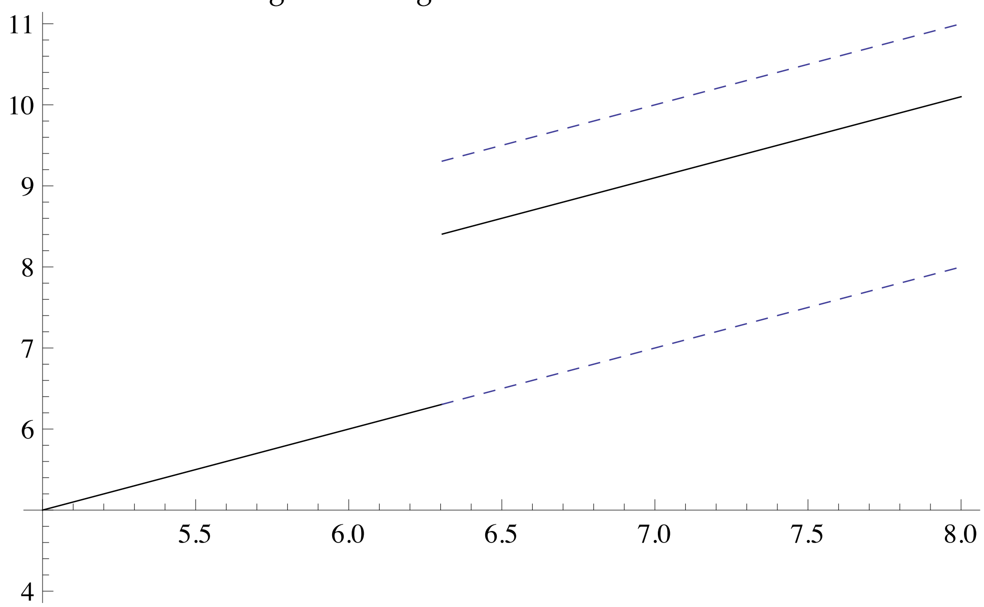

Parameters as in Figure 3. Dashed lines represent possible outcomes and solid lines average outcomes.

Figure 5: The effect of policy on cutoff for migration

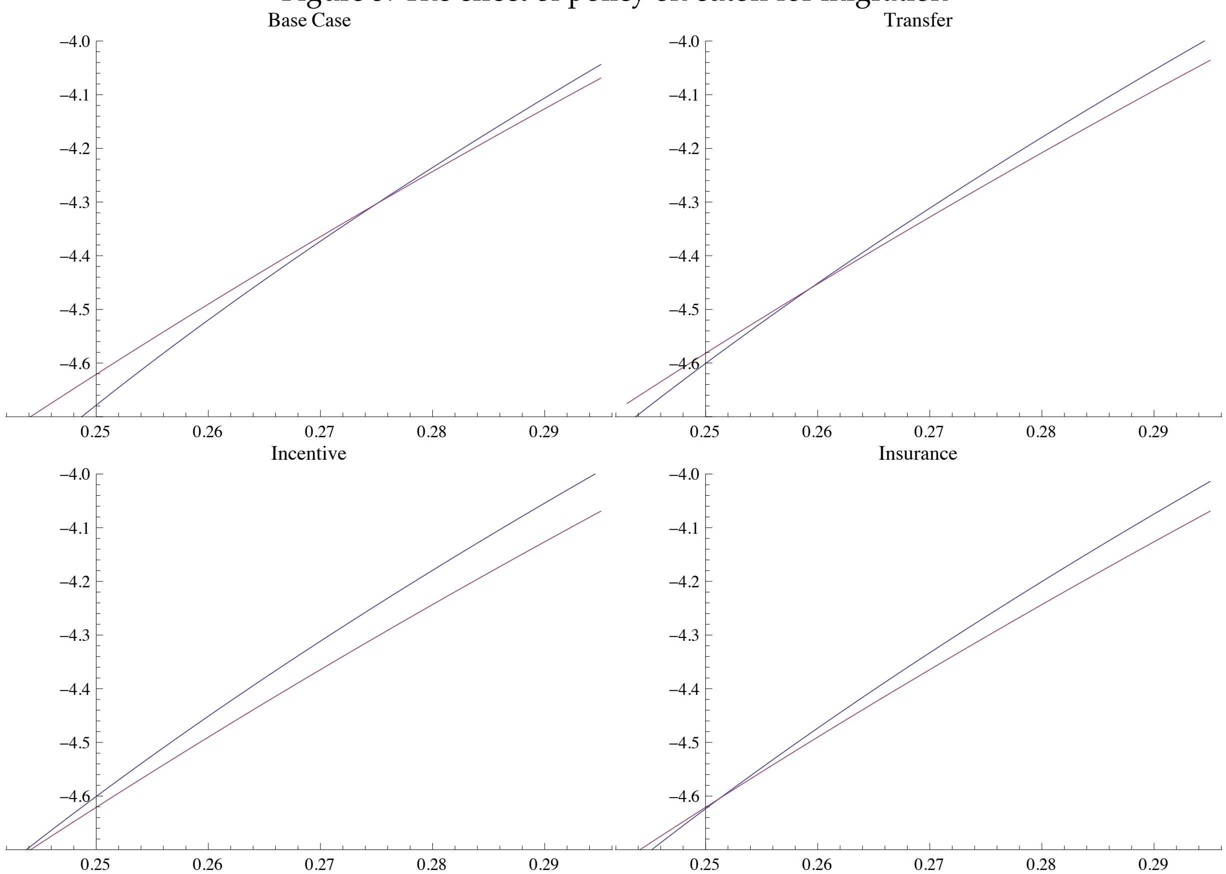

$u = \ln$, $\mu(g) = 0.7$, $b = -0.2$, $g = 0.05$ and $\delta = 0.7$.

Figure 6: The effect of policy on cutoff for migration with income above subsistence

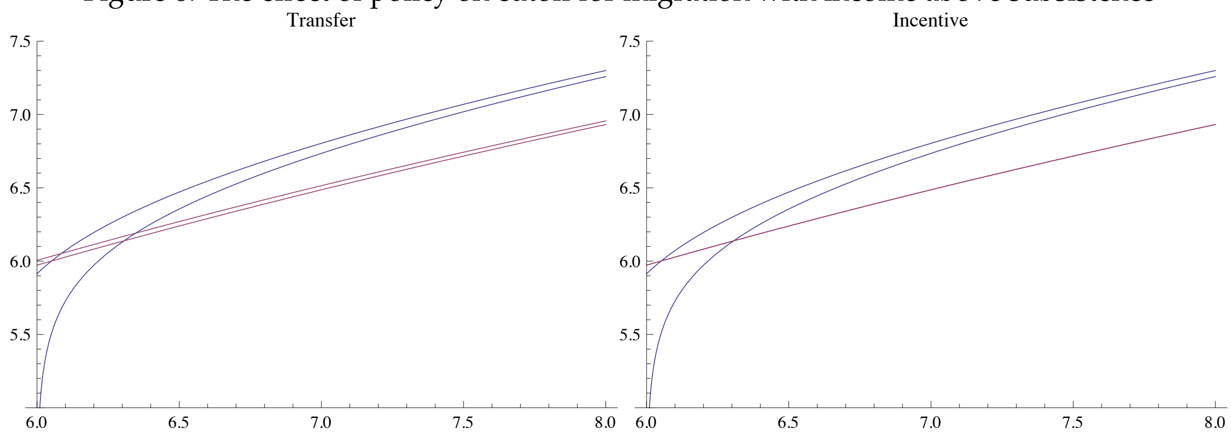

$u = \ln$, $\mu(g) = 0.7$, $b = -6$, $g = 3$ and $\delta = 0.7$.

Figure 7: Kernel Density Plots of Subsistence Index by Treatment Status

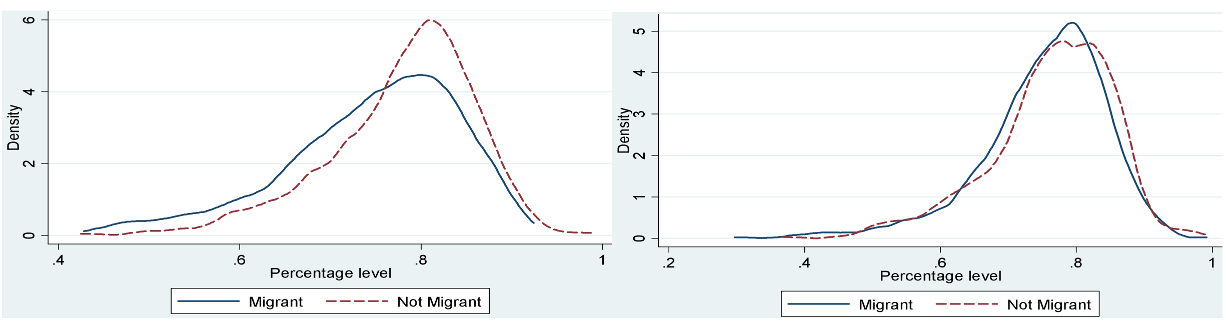

List of Tables

Table 1: Program Takeup

	Offer Accepted	Kept Money	Migration Rate
Cash	71.88%	48.26%	58.96%
Credit	52.98%	34.21%	56.80%
Info	35.14%	.	35.93%
Control	.	.	35.97%
	Incentivized	**Not Incentivized**	**P-Value**
Migration Rate	58% (0.014)	36% (0.0196)	0.00
Remigration Rate	47% (0.014)	37% (.020)	0.00

Offer accepted refers to household that accepted the offer and received the money. Some of the household returned the money given. Kept Money refers to the individuals who accepted the offer and kept the money from the treatment. The P-value is obtained from the difference between migration rates of incentivized and non incentivized households, this refers to all household in Cash, Credit, Info and Control, whether they kept the money or not.

Table 2: Effects of Migration on Expenditure Amongst Remaining Household Members

	Sub-district F.E.		Sub-district F.E. + controls		
	OLS	IV	OLS	IV	Mean of Dependent Variable
Food Expenditures	79.16***	224.8*	92.40***	220.3*	729.2
	(18.08)	(124.2)	(18.28)	(117.9)	
Non Food Expenditures	46.04***	111.5**	55.91***	111.2**	274.4
	(8.448)	(49.54)	(8.388)	(44.41)	
Total Expenditures	124.5***	337.5**	147.4***	332.5**	1003.1
	(22.36)	(154.1)	(22.26)	(143.0)	
Total Caloric intake	231.3***	729.4***	262.7***	666.6***	2091.3
	(40.61)	(238.1)	(42.69)	(234.8)	
Total Calories from Protein	4.168***	10.51	4.889***	10.07	46.9
	(1.176)	(8.709)	(1.276)	(8.237)	
Expenditures on Meat Products	4.605	29.25	7.013*	33.60*	28.2
	(3.707)	(18.15)	(3.852)	(18.29)	
Expenditures on Milk & Eggs	1.64	-0.595	0.909	-0.938	13.1
	(1.453)	(8.067)	(1.431)	(7.975)	
Expenditures on Fish	0.564	-12.77	1.615	-8.693	74.4
	(6.552)	(48.12)	(7.133)	(43.03)	
Expenditure on Children's Education	-3.216	28.10**	-4.089*	21.05*	18.2
	(2.34)	(12.88)	(2.434)	(12.06)	
Expenditures on Clothing and Shoes	6.243***	12.72	6.928***	12.54*	38.8
	(1.53)	(8.101)	(1.528)	(7.395)	
Expenditures on Health for Female	1.565	53.37	0.204	50.80	61.5
	(6.935)	(39.36)	(7.314)	(39.78)	
Expenditures of Health for Male	18.32**	24.48	19.09**	16.42	71.4
	(8.199)	(35.28)	(8.640)	(37.70)	

Robust standard errors in parentheses, clustered by village. *** p<0.01, ** p<0.05, * p<0.1. The coefficient reported is from the migration variable that indicates if the households migrated or not. The Instruments are all the randomization treatments, such as if the individual was offered cash, credit, info, or has an assigned destination, etc. Food, Non Food and Total Expenditures is a variable that represents the amount in taka spent per person per month in each of these categories. Total Caloric intake refers to the amount of quantity consumed per individual per day. The number of individuals in the household that we used as a denominator for calculating per person expenditures is estimated for the number of individuals in the house that were present in the last seven days. This adjustment is made because some of the measures of expenditures were asked as weekly consumption. We control for household education, income proxy, percentage of food expenditure, number of adult males, number of children, credit constraints, total expenditure per capita in round 1, and expectations about monga.

Table 3: Earnings & Saving from Migration (Round 2)

	All Migrants	Incentivized	Not Incentivized	Obs
Total Savings by household	3490.5	3506.6	3434.9	951
Total Earnings by household	7777.2	7451.3	8894.4*	952
Savings per day	56.8	56.5	57.8	905
Earnings per day	99.4	96.1	111.5***	926
Remittances per day	17.8	16.2	23.3***	926
Travel Cost per Episode	222.1	222.2	221.8	953

*** p<0.01, ** p<0.05, * p<0.1, the p-values refers that incentivized and not incentivized are statistically significant different from each other. The measures for total savings, earnings and savings and earnings per day do not include outliers (Less than 20,000 for total savings and 120000 for earnings. Individuals savings per day less than 500 and individuals earnings per day less than 700. Travel cost refers to the cost of food and travel to get to the destination. Average migration duration 76 days.

Table 4: Earnings for non-migrants

Income	Only Employed	Employed & Unemployed
Job type: Daily	94.7	87.9
Job type: Salary	64.9	60.6
Non Agricultural Business Daily Profits	61.1	.

"Job Type Daily" refers to earnings from jobs that are paid on day to day basis. Salary refers to earnings from jobs that are paid on a month to month basis. For which we divided by 30 days to get daily earnings. Non Agricultural Business, are enterprises of individuals who are not dedicated to the agricultural business. They report monthly profits, hence we divided this number over 30 days to calculate daily profits. Unemployed is defined as people who were not working and looking for a job, or stop looking for a job because they could not find any work available. Unemployment is about 10%. The Last column is the average earnings across employed and unemployed workers.

Table 5: Who Migrates by Treatment Group

Table: Percentage of People that Know Someone at Destination

	Incentive	Non incentive	Diff	Std Error
First Episode	47%	65%	0.17***	0.04
Second Episode	60%	72%	0.12**	0.06
Third Episode	68%	82%	0.14	0.09
Fourth Episode	86%	88%	0.06	0.11

*** p<0.01, ** p<0.05, * p<0.1

Table: Percentage of People that had a Job Lead at Destination

	Incentive	Non incentive	Diff	Std Error
First Episode	27%	44%	0.17***	0.03
Second Episode	29%	47%	0.18**	0.06
Third Episode	36%	54%	0.18**	0.09
Fourth Episode	53%	59%	0.06	0.15

*** p<0.01, ** p<0.05, * p<0.1

Table 6: Remigration by Earnings in First Migration Episode

Earnings from Round 2	Migration in round 3: Migrant	Migration in round 3: Non Migrant	Sig.
Total Savings from migration			
Incentivized	3983.5	2671.0	***
Not Incentivized	3617.0	3294.3	
Total Earnings from migration			
Incentivized	8304.2	5948.3	***
Not Incentivized	8819.6	8638.4	*
Household Savings per day			
Incentivized	61.9	51.8	*
Not Incentivized	57.6	55.9	
Household Earnings per day			
Incentivized	101.9	86.4	
Not Incentivized	116.4	105.3	***

*** p<0.01, ** p<0.05, * p<0.1, This is for the test that compares the two horizontal numbers in the respective lines.

Table 7: Remigration by Earnings in First Migration Episode

	Effects of going to		Effects of having one	
Destination	Succesful	Unsuccesful	Succesful	Unsuccessful
Dhaka	0.257	0.109	0.008	-0.024
Bogra	0.137	0.111	0.042	0.008
Tangail	0.272	0.238	0.013	0.013
Munshigonj	0.156	0.051	0.001	0.010
Comilla	0.270	0.066	0.033	0.057

Success is defined as a the individual's migrant perception of his expected earnings.

Table 8: Structure of the Randomization

Group nature:	A. Individual		B. Assigned group				C. Self-formed group				Total
Group size:			Two		Three		Two		Three		
Destination type:	Assigned	Chosen	Assigned	Chosen	Assigned	Chosen	Assigned	Chosen	Assigned	Chosen	
Incentives:											
a) Information only											**304**
b) (a) + conditional cash transfer	133	126	66	48	54	54	66	48	60	48	**703**
c) (a)+ conditional credit	105	112	42	54	42	48	42	54	36	54	**589**
Control group											**304**
Total # of households	**238**	**238**	108	102	96	102	108	102	96	102	**1900**