Migrating out of Poverty
RESEARCH PROGRAMME CONSORTIUM

The Effect of Return Migration on the Household Welfare: Evidence from Ethiopia

Marta Schoch, Julie Litchfield and Asmelash Haile Tsegay

Working Paper 66

January 2020

Acknowledgments

The authors wish to thank our large team of field supervisors and enumerators for collecting the data: Amare Dagnew, Mulat Kiros, Lemma Bekele, Gebremichael Wegayehu, Genetu Dagne, Teame Tesfay, Habtom Debes, Habtom Gebregergish, Yirgawesen Getahun, Solomon Kidanu, Beza Sherif , Teshome Zikarge, Abu Zeneba, Mebratu Bogale, Meaza Nugussie, Berehe Kiros, Line Kinfe, Daniel Mebrahtu, Gebremedihin Berhe and Melkamu Chefamo. We also thank participants at the Income and Remittances workshop in Acrra, Ghana, February 2019, for comments and discussion on early results, and participants at the Migrating out of Poverty regional workshop in Addis Ababa, Ethiopia in May 2019.

Migrating out of Poverty
University of Sussex, Arts B
Brighton BN1 9QN, United Kingdom
Email: migrationrpc@sussex.ac.uk
Web: www.migratingoutofpoverty.org
Twitter: @MigrationRPC
Facebook: /migratingoutofpoverty
Blogger: migratingoutofpoverty.blogspot.com

This is an output from a project funded by UK aid from the UK government. The views expressed do not necessarily reflect the UK government's official policies.

Abstract

This paper analyses the impact that return migration has on the welfare of households to which migrants return. We use panel data from a bespoke longitudinal survey of rural households in Ethiopia in 2014 and 2018. We outline the evaluation methods available to researchers investigating the causal impact effect of migration on household welfare, focussing our discussion on impacts on the household at origin. We show using a combination of difference-in-difference in methods with matching that households with a return migrant over the period experience increases in non-food household consumption and that return is also associated with large increases in remittances prior to return.

Executive Summary

Return migration has been an important feature of migration patterns in Ethiopia in recent years, caused in part due to a ban on international migration to Middle Eastern countries and inter-ethnic conflict within the country.

This paper explores the impact of return migration on welfare of the households to which migrant return.

Evaluating the impact of return is methodologically challenging as return is usually a selective process, with the decision to return in part influenced by household characteristics. This paper provides an overview of methods that can be used to address these challenges.

We use panel data for a sample of rural households which follows households over time from 2014 and 2018 and captures the evolving migration experiences of their members. We employ difference-in-differences estimation combined with propensity score matching at baseline.

We find that a large proportion (around a third) of international migrants have returned following deportation, and family issues are responsible for around a quarter of both internal and international migrants. A sizeable minority of migrants return because they had finished their contracts or earned enough money.

We find some evidence of selection in return. Returnees are slightly more likely to be male than migrants still away , suggesting that women are slightly less likely to return, and that returnees are also around 3 years older on average than those still away. We find higher incidence of return amongst Muslim migrants, and among those originating from Tigray.

In terms of household welfare we find that return has no impact on the food consumption, either in terms of the expenditure on food or on the dietary diversity of food consumed in the household. We do however find a statistically significant impact of return on non-food expenditure, with our estimates suggesting return increases non-food consumption by just over 20% compared to households with migrants still away. Much of this increase is driven by increased expenditure on clothing and household items such as kitchen utensils, furniture and linens.

Our analysis also shows that return is associated with a very large increase in remittances sent home by the migrant in the 12 months prior to return. This suggest that migrants are able to accumulate significant savings while away from home and that even if their return may not have been anticipated precisely, it does not necessarily follow that return migration should be viewed as an example of a failed migration.

Introduction

Return migration is an important part of many migration journeys. Around forty percent of international migrants in OECD countries will leave their host country within five years of arrival (Wahba, 2015a) while estimates based on all migration flows suggest that around twenty-five percent of migration events are returns to country of birth (Asoze and Raftery, 2019). Note that both of these figures ignore return that happens within countries. Return, which the International Organisation for Migration IOM defines as "the act or process of going back or being taken back to the point of departure" (IOM, 2019), can be planned or unplanned, voluntary or forced, assisted or independent. While return is sometimes characterised as being a sign of a failed migration attempt, to characterise return as such would be to overlook the growing body of evidence that suggests that return migrants often accumulate skills, savings and know-how while away and use these upon return.

Wahba (2015b) studied return migration to Egypt and the impact on returnee wages and occupational choice. She pays careful attention to selection into migration and into return. While she found that returnees are negatively selected among migrants (i.e. less likely to be endowed with characteristics that give them advantages in the labour and entrepreneurial markets), migrants as a group are positively selected. Her work shows that living abroad afforded migrants the opportunity to increase income and savings, and also to acquire new skills which resulted in a wage premium upon return. In addition, while return migrants may lose social capital during the period of their time abroad, they are more likely to become entrepreneurs when they return than non-migrants because of the human capital, savings and experience acquired overseas (Wahba and Zenou, 2012).

This paper analyse the effect that return migration has on the welfare of the households to which migrants return. We use panel data from the Migrating out of Poverty survey of rural households in Ethiopia carried out in 2014 and 2018. Our decision to focus on the impact of return migration is motivated by the drop in migration and the extent of return migration observed in our data over the period. The number of migrants (internal and international) roughly halved over the four year period, with the number of households with no current migrants almost doubling. We ascribe these changes to two main factors. Firstly it became increasing difficult to migrate to the Middle East following highly publicised cases of violence against Ethiopians, particularly young women, working abroad, which promoted a ban on international migration.[1] This combined with increasing inter-ethnic tensions within Ethiopia, escalating to violence at times, has led to a shift away from international migration towards internal migration, and to internal migrants choosing locations closer to home. We discuss these trends in more detail in Haile Tsegay and Litchfield (2019).

We first provide an overview of the empirical challenges that arise in evaluating the impacts of migration (leaving aside discussions at this stage about how migration is defined) and then describe some of the more common approaches to tackling these challenges. We then describe the Migrating out of Poverty data for Ethiopia and provide more details on the extent and nature of return migration in Ethiopia. We analyse the impact of return migration on household welfare by isolating households which had migrants at baseline where some of the migrants returned to the household by 2018 and employ a difference in differences approach with matching at baseline.

[1] Conflict in Yemen, a common route for Ethiopians, has also disrupted international migration.

Effects of migration: econometric challenges

Measuring the effects of migration is a challenging econometric task. Although correlation between migration and households' characteristics might be insightful, we might be interested in understanding whether migration affects some outcome of interest in a causal way. For example, we might be interested in the causal relationship between having a migrant in the household and households' welfare. When interested in the effects of migration on households' welfare the first strategy that comes to one's mind is to compare households with at least one migrant to households without. If, for example, we are interested in the effect of migration on poverty, we might proceed as follows. Using cross-sectional data we could compare the mean poverty index of households with migrants to the mean poverty index of households without migrants. Let us assume that we find that households with migrants have lower poverty index scores. We could conclude that migration reduces poverty. There are several problems with this conclusion.

First, households with migrants might have been less poor before the member(s) of the household migrated. This is not difficult to believe, as in many contexts richer households are more able to afford the costly investment of migration, e.g. foregone income, costly barriers to migration etc. Thus, one of the main concerns when looking at the effects of migration is that some households select into migration on the basis of pre-existing characteristics. This means that comparing a household with a migrant to a household without is not a good comparison as they are different in other characteristics to begin with. Often, this is referred to as selection bias. Failing to account for such bias might produce biased estimates of the relationship of interest, as shown by Wahba (2015b) discussed above. McKenzie et al. (2010) find that failing to account for such bias can lead to an overestimation of the income gains from migration of between the 20 to 82%. They find, for example, that migrants are more educated on average compared to non-migrants, with the former having on average two more years of education than the latter. This leads the authors to conclude that selection into treatment is not only due to observable characteristics, but plausibly also to unobservables, e.g. motivation or parental income might both drive up educational achievement and the probability to migrate. In line with this, Clemens and Tiongson (2017) find that analysis that do not correct for selection might produce very different results compared to a quasi-experimental analysis. In particular, they find that when not accounting for selection, migration seems to positively affect the probability that children in the origin household attend school, while this effect is not statistically significant in the quasi-experimental analysis. The authors conclude that unobservable characteristics such as aspirations for the migrant's children future might explain this result independently of the migration experience (Clemens and Tiongoson, 2017). Similarly, they find that non-migrants are less likely to work than migrants and that failing to account for selection into treatment might result in attributing such effect to the migration program (Clemens and Tiongson, 2017).

Second, it might be that factors other than migration are affecting poverty levels of households differently for households with a migrant and households without a migrant and poverty levels react to these other factors rather than to migration. Following from the example drawn from Clemens and Tiongson (2017), it might be that other factors make non-migrant members of the household less likely to work. This is likely to affect poverty independently of migration. Similarly, it is usually the case that proximity to a bigger town or a main road makes one more likely to migrate (McKenzie and Sasin, 2007). This is also likely to affect poverty level has proximity to a big city or to a main road might make it easier for the household to have access to more goods or to better trade opportunities. Another example might be the exposure to climate shocks, e.g. a drought, that might make one more likely to migrate as well as affecting negatively income. This makes it difficult to assess whether is poverty that affects the decision to migrate or vice versa.

In general, the problem we often face when assessing the causal relationship between migration and the outcome of interest is the challenging task of establishing a good counterfactual. The question we want to answer is "what would have happened to this household if it didn't have a migrant?" Given that we can only observe the unit of analysis, i.e. the household, with or without a migrant, we need to use statistical methods to estimate an appropriate counterfactual.

In the next section we will discuss how this is usually done in the econometric literature on migration but first we outline the key concepts from the experimental literature as an aid to understanding different approaches. In general, we will see that the causal inference of the effect of migration on the outcomes of interest can be evaluated using experimental or non-experimental methods. The key feature of the experimental approach is that we assume the researcher is able to randomly assign treatment (e.g. a household has a migrant), to some households, labelling them the treatment group, and not to other households, which become the control group.

Figure 1 shows a graphical visualisation of evaluating the treatment effect of, say, free school meals, on test scores of children of school age. We can see that test scores, in this hypothetical example, have increased for both treatment and control groups but that the test scores of children receiving free school meals have increased by more than those of children in the control group. By observing the same children before and after treatment, we can compare the mean outcomes between the two groups and net out the treatment effect from other possible confounding factors.

Figure 1 Example of impact of treatment effect on test scores before and after intervention

Before

After

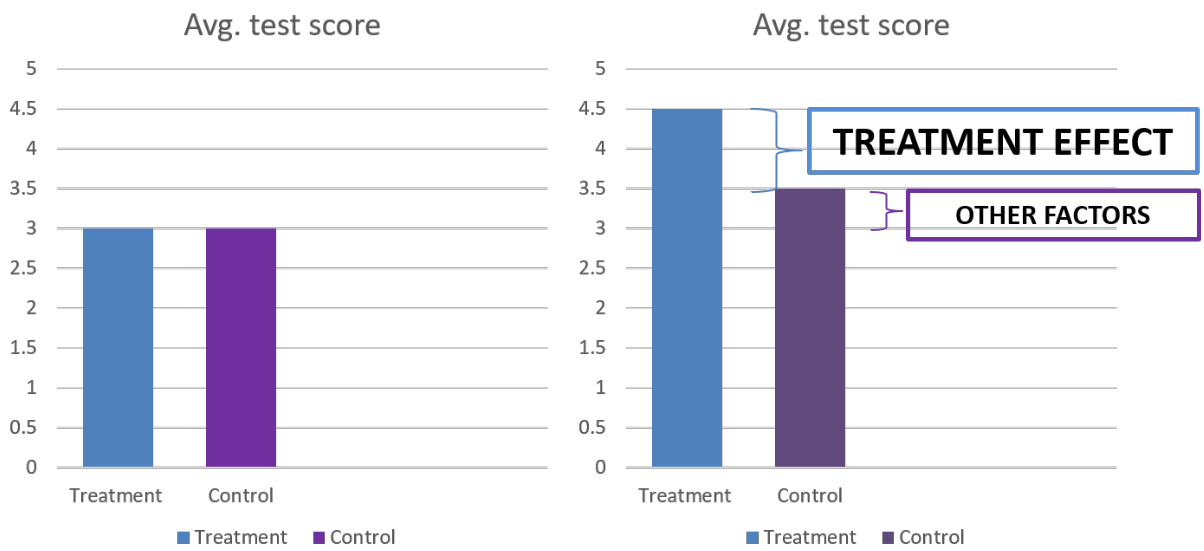

For the estimate of the treatment effect to be unbiased, the experimental setting needs to satisfy certain assumptions. First, the treatment and the control group must be comparable in terms of observable and unobservable characteristics before the treatment. Failing to satisfy such assumption might not account for selection or omitted variable bias. This is also true if the treatment and control group have different pre-treatment trends, i.e. they behave differently already before the treatment. In this case it becomes important to control for pre-treatment trends for the two groups and make sure that the only difference post-treatment is due to the treatment.

However, while random assignment of a treatment such as free school meals, vaccinations or training may be feasible, it is hard to conceive of migration being randomly assigned. There are of course migration-related interventions that can be analysed using an experimental approach: random allocation of visas for example, hence migration researchers will usually need to seek non-experimental solutions. The experimental setting however does provide a useful way to conceptualise the problems of drawing causal inferences about the impact of migration on household welfare. In practice, this means that using instrumental variables, propensity score matching or exogenous shocks we can seek to identify a valid control group that differs from the treatment group only because of the treatment, i.e. having a return migrant in the household.

Common practises in evaluating the effect of migration.

There are some examples of migration research which have been able to adopt an experimental approach. McKenzie et al (2010) used a lottery program aimed at selecting Tongan individuals to move to New Zealand for one year. By surveying all the individuals who applied for this program, the researchers could then make a meaningful comparison of households that were selected to migrate and those who were not. Had they compared those who migrated to all other households, then the comparison group in this hypothetical example would have included households who would have not wanted to migrate to begin with, and hence did not apply to the program, as well as households which had applied to the program. This is an example of establishing a valid control group or counterfactual.[2] Other examples include the work by Gibson et al (2011) and Bryan et al (2014).

Designing or exploiting natural experiments such as these is not always possible so researchers often have to find a quasi-experimental approach. We outline some significant examples below.

Exogenous shocks

One way of evaluating the effect of migration on the outcomes of interest is to use exogenous shocks to migration or remittances. Yang (2008) provides a good example of how to use exogenous shocks to address endogeneity and reverse causality in the context of migration. The paper exploits an exogenous shock to exchange rates to estimate the effect of remittances on household investment and human capital accumulation. The identification strategy relies on the Asian financial crisis being an exogenous shock to the value of remittances sent from Filipino migrants back home. The exogenous shock causes a change in the value of remittances between two periods before and after the 1997 financial crisis. Weighting for the share of migrants living in different countries, an exogenous variable which is exogenous to the decision to migrate, and also varies across the sample of households with migrants, can be constructed and used to evaluate the effect of remittances on the outcomes of interest.

Instrumental variables

Another common way of dealing with the challenges of establishing causality is to use an instrumental variable approach. A good instrumental variable is a variable that affects the outcome of interest only via migration or remittances, and not directly. Some authors have used historical

[2] See also McKenzie and Yang, 2012.

trends of migration to predict present trends as these capture the existence of migration corridors and migrant networks (Binzel and Assaad 2011, Card, 1990, Hildebrand and McKenzie 2005, Karymshakov *et al.* 2017, Munshi 2003, Mackenzie and Rappoport 2007, Vadean *et al.* 2017). Other studies have used geographical variables such as distance to roads, railroads, capital cities or borders, reflecting the cost of migration (for example, Black *et al.* 2015, Demirgüç-Kunt *et al.* 2011). Some studies make use of historical data that capture "push" factors such as local wages, incomes or unemployment rates (for example, Amuedo-Dorantes and Pozo 2010).

The fundamental assumption for an instrumental variable to be valid is that it most both be *relevant* and satisfy what is known as *exclusion* restriction. This means we are looking for a variable that is correlated with the migration phenomena we are analysing (relevance) but not directly related to the outcome we are interested in. for example, historical migration flows are usually a good predictor of current migration flows; but do not necessarily have an influence on current welfare, other than through their indirect influence via current migration. It is easy to see why in practice it can be difficult to find a valid instrument.

Panel data, fixed effects and difference-in-differences models

One of the challenges in analysing migration is that migrants, or their households, possess certain characteristics that make them different from the non-migrant population. These might be characteristics that can be measured, such as education, gender, household wealth, but may also be unobservable, such as attitudes to risk. Evidence of the importance of selection bias in migration decisions (McKenzie et al. 2010) has led to a greater use of household panel data (Clemens and Tiongson, 2017). Gibson and McKenzie (2014) used follow-up surveys to the work of McKenzie et al (2010) to estimate the effect of migration on household outcomes. The advantage of using panel data is that we can rule out that the difference between households with and without migrants is due to any of the unobservable household characteristics as long as these do not change over time.

As mentioned in the previous sections, the most common challenge when doing research on the effects of migration is the selection bias that determines who decides to migrate. In fact, even if an exogenous shock or an instrumental variable were to identify a change in the probability to migrate, we might still be concerned that who decides to migrate or who receives remittances is not randomly drawn from the population. One way to deal with selection bias, especially when we are concerned that unobservable factors might affect the decision to migrate, is to use fixed effects models. This is only possible when we can observe the same unit of analysis more than once, for example when using household panel data. The literature on migration is making more and more use of panel data, Gibson and McKenzie (2014) use the following model to estimate the effect of migration on household economic outcomes:

$$Y_{h,t} = \alpha_h + \beta Migration_{h,t} + Year_t + \varepsilon_{h,t} \tag{1}$$

where Y is the welfare outcome of household *h* at time *t*, α_h captures household fixed effects. Fixed effects introduce a different intercept for each household and so control for time-invariant unobservable characteristics that might affect the decision to migrate. Everything else being constant, if between the two times we observe a household, their welfare increases, we can rule out the possibility that this because of some unobservable time-invariant characteristic, e.g. the risk aversion of the head of household. Of course, this does not solve the problem of isolating the migration decision from time-varying characteristics, but it helps to get a closer estimate of an unbiased estimate. Household fixed effects models are being used more and more often in the literature. For example De Brauw et al. (2018) and Beegle et al. (2011) estimate the following model:

$$\Delta lnC_{h,t+1,t} = \alpha_h + \beta Migration_{h,t+1} + \gamma X_{h,t} + \varepsilon_{h,t} \quad (2)$$

where $lnC_{h,t+1,t} = lnC_{h,t+1} - lnC_{h,t}$, i.e. the percentage change in consumption between period *t+1* and *t*, α_h are again households fixed effects, and *X* is a vector of household characteristics measured at the initial period.

When we have information on the treatment and control group before and after the treatment, we can implement another evaluation method: difference-in-differences (DID). DID can be used to establish causality if we observe a treatment and a control group over time and the treatment takes place during this period. In terms of estimation a DID model will include a treatment status and a time-specific binary variable, in the case of two periods this will be a pre-post dummy variable. The crucial assumption for DID to produce credible estimates is that it must be possible to both the treatment and control group before the treatment. This allows us to check whether the two groups were comparable before the treatment took place. This assumption is often referred to as common or parallel trends, i.e. the trend over time in the outcome variable was the same for the two groups and it changes only after treatment takes place. If the parallel trend assumption is satisfied, the control group can be a good counterfactual. Gibson and McKenzie (2014) use a difference-in-differences specification to establish the effect of participating to seasonal work-related migration programs (RSE) on household income, consumption and savings:

$$Y_{h,t} = \alpha + \beta EverRSE_h + \sum_{t=2}^{4} \delta_t + \gamma RSE_{h,t} + \varepsilon_{h,t} \quad (3)$$

Where $Y_{h,t}$ is the outcome of interest for household *h* at time *t*. $EverRSE_h$ is a binary variable indicating if the household ever participated in this migration program, and δ_t are survey year dummies. γ gives the average treatment effect of participating in the migration program.

The next subsection discusses a commonly used methodology that deals with selection bias in non-experimental settings. De Brauw et al. (2018), Gibson and McKenzie (2014) Beegle et al. (2011), all use propensity score matching in combination with fixed-effects and diff-in-diff models. However, propensity score matching can be used as an estimation strategy also with cross-sectional data.

Propensity Score Matching (PSM)

It might be the case that we are interested in evaluating the effect of a treatment in a non-experimental setting. It is not always possible for the researcher to apply experimental design to select a control group. Studying the effects of migration is a common example of this limitation. Propensity score matching involves pairing treatment (households with migrants) and control (households without) units by looking at their observable characteristics. A good primer on propensity score matching application is Heinrich et al. (2010, IDB technical note) and for more formal discussion see Dehajia (2005), Dehajia and Wahba (2002), Smith and Todd (2005).

The key intuition of propensity score matching (PSM) is that we can use information from a pool of units that did not receive the treatment, but that are similar to those who did in terms of observable characteristics, to identify our counterfactual. What this means in practice is that we look for a unit in the control group that is similar to a unit in the treatment group and evaluate the effect of the treatment by taking the difference in outcomes between the two. For this to meaningful, we need to match treatment and control group according to observable characteristics that might influence the outcome and are not the treatment of interest, here migration. The propensity score is defined as a "probability that a unit in the combined sample of treated and control units receives the treatment, given a set of observed variables" (Heinrich et al., 2010). What this means is that considering the presence of selection bias in a non-experimental design, e.g. migration is not random among the treatment and the control group because richer households are more likely to have a migrant (be treated), propensity score matching will help to identify a control group that is as close as possible to the treatment group given a set of observable characteristics.

The two crucial conditions for propensity score matching are that:

1) we can observe characteristics that affect the probability of being treated other than the treatment effect (*conditional independence assumption*);
2) we can observe the characteristic of interest in the control group. This second point means that we need to have non-zero observations for the covariates of interest in the control group, e.g. income level data needs to be observable for both the treated and the control groups (*common support* or *overlap condition*).

Following from these two assumptions, data availability defines how implementable propensity score matching techniques might be. The researcher should make sure that:

a. Data has been drawn from the same source for the treated and control groups
b. Missing values are handled in the same way for the treatment and control groups
c. The number of observations in the treatment and control group must be large enough

Once these assumptions are satisfied, we can look at the practicalities of using PSM for statistical analysis. Recall that the method involves pairing households. There are different algorithms that can be used to do this pairing. The most common matching technique is called *nearest neighbour covariate matching*. This consists of matching a unit in the treatment group with the closest unit in the control group for a given characteristic. Heinrich et al. (2010) provide a good example of a simple matching procedure (see table 1-2, pp. 19-20). Matching on a single characteristic is straightforward and intuitive. However, in most empirical analysis we face a problem of multidimensionality. Let's say we observe that larger households are more likely to have a migrant. Matching on this dimension would be easy, we should just match a household with say five members and a migrant to a unit with five (or six, if we are thinking about pre-migration household size) members but no migrant.

How does this change if we also observe that male-headed households are more likely to have a migrant? Or if one dimension predicts migration in the opposite direction of the other? The solution to multidimensionality is to compute a *propensity score,* i.e. the probability of receiving the treatment given a series of characteristics. This propensity score will allow us to have a balanced (i.e. equally distributed) distribution of characteristics (covariates) between the treatment and the control groups. More importantly, it will allow us to solve the multidimensionality problem and to match treatment and control units based on the propensity score itself.

The following steps outline the implementation of PSM techniques in practice:

a. Estimating the propensity score

To calculate the propensity score we need to define a model of the type below, known as the participation model:

$$Treat_{h,t} = \beta_0 + \beta_1 Age_{hh} + \beta_2 Gender_{hh} + \beta_3 Education_{hh} + \beta_4 Village_{r,t} + \cdots + \varepsilon_{h,t} \qquad (4)$$

Where $Treat_{h,t}$ is the probability of receiving the treatment (for example, that a household has a least one migrant, or receives remittances) conditional on a set of characteristics. Given that the outcome variable is a binary variable, researchers usually estimate equation 3 with a probit or logit model. The crucial question in estimating the propensity score is what characteristics (or covariates) to include in the regression. In general, we would want to include all the characteristics that determine participation in the treatment group. In the stylised example above we included age, gender and education of the head of household, but we could also include other important characteristics that might determine migration: location of the household, year of the survey and distance from the capital/main road. It is also common to include higher order polynomials of relevant covariates. Notice that deciding on a participation model is one of the crucial steps of this evaluation technique. The model should include only covariates that are relevant in predicting the treatment, thus it is advisable to start with a parsimonious model and include only variables that have a statistically significant relationship with the treatment status.

Once the propensity scores have been estimated, the second step is to match treatment and control units to be compared.

b. Matching treated and control units using the estimated propensity scores

There are several statistical packages that can be used to match propensity scores in the treatment group to those in the control group once the propensity score has been calculated. More robust results will test whether similar effects can be obtained using different matching techniques (see Heinrich et al. 2010 for a detailed explanation). Most statistical software can be used to perform the matching and generate estimations.[3]

c. Evaluate effect of treatment

If the matching was successful, this simply consists in taking the difference of the average value of the outcome variable between the treatment and control group. In fact, once the matching has made the treatment and control group comparable, the evaluation of the result is as in an experimental setting. We estimate this using a fixed effects model

$$consumption_{h,t} = \beta_0 + \beta_1 Treated_h \times Wave_t + \beta_2 Wave_t + \beta_3 Treated_h + +\varepsilon_{h,t} \qquad (5)$$

where the coefficient β_1 captures the effect on consumption of being treated.

Does return migration improve welfare in Ethiopia?

In this section we provide an example of a way to analyse the relationship between having a migrant in the household and household welfare. We focus on a specific definition of migration, i.e. whether the household has a return migrant between two waves of data. Thus, in this example we define the

[3] In STATA **psmatch2** includes different options for matching (nearest neighbour, caliper matching, radius matching) and graphing options (**psgraph**). It also includes a command for covariate balancing tests (**pstest**).

treatment status using a binary variable equal to one if the household has at least one migrant returning home by wave two.

Choosing our control group is also an important step. First, this depends on the relationship of interest. Second, we want to make sure that the treatment and control groups are as comparable as possible at baseline. Thus, we choose to use as control group households that have at least one migrant at baseline and at follow-up and do not have migrants returning to the household. The reason for this is the following. Let us assume we were to use as a control group households with no migrants in either wave. On the one hand, these households might be very different compared to households with at least one migrant at baseline. On the other hand, interpreting the results and their economic significance might prove less meaningful. We would essentially be looking at the effect of a member of the household leaving and returning to the household between the two rounds of the survey. In contrast, using as a control group households with at least one migrant in both waves, but no returning household members, will allow us to have fairly comparable groups at baseline (both have at least one migrant away) and to capture the effect of interest, i.e. having at least one member of the household returning. This also captures an important part of the migration patterns we have observed in Ethiopia.

We might expect that having a migrant returning to the household impacts welfare in different ways. On the one hand, it might have a negative effect on household consumption due to the drop in remittances and goods that are no longer sent home by the migrant. It might also cause per-capita consumption to decrease if the return of the migrant simply translates in an increase in household size. On the other hand, this effect might be positive if the migrant returns and brings back savings accumulated during the period away. The migrant returning to the household might contribute to agricultural production and increase earned income. Moreover, a stream of literature has looked at the transmission of norms and practises that the migrant has learned while abroad, e.g. better health and contraception practices that might improve living standards (Hildebrand and McKenzie, 2005).

Before we set out our method in detail we provide some context for the research as well as a description of our data.

The Ethiopian context and Migrating out of Poverty data for Ethiopia

Ethiopia has experienced rapid development in recent years, with significant investments in infrastructure and substantial growth in the services sector, and average growth of over 10 per cent per year between 2004 and 2018. Yet Ethiopia remains a low income country, with per capita GDP (in 2010 USD) of around $1800 in 2018.[4] Employment opportunities, particularly for young people are limited and thus migration is often seen as a way to explore new opportunities.

Estimates of migration incidence are a little patchy. According to the United Nations Population division (UNPD), more than 1.2 million Ethiopian live abroad,[5] but some estimates reach as high as three million, including undocumented and irregular migrants. The top three international destinations of Ethiopian migrants are the USA, Saudi Arabia, and Israel. The Middle East generally is a popular destination, particularly for women who are in demand as care and domestic workers. The Ministry of Labour and Social Affairs (MoLSA) estimates that around 460,000 Ethiopians have legally

[4] World Development Indicators.
[5] UNDP(2017).

migrated to the Middle East, mainly Saudi Arabia, Kuwait and Dubai between September 2008 and August 2013. Internal migration estimates come from the country's Labour Force Survey: the 2013 survey suggest that internal migration[6] in Ethiopia remains limited and accounted for only 6.5% of adult population. In rural areas in particular, mobility is limited, with a mere 3.5 percent of adults moving out of their original zone of residence between 2008 and 2013 (the five years preceding the 2013 LFS. Migrants account for a higher share of the population in urban areas: in 2013, 17 percent of urban dwellers were recent migrants (i.e. had come to the city in the five years up to 2013), and the LFS also suggests a shift from rural to rural migration to rural to urban migration.[7]

The Migrating out of Poverty data for Ethiopia is a two-round panel survey, of approximately 1200 households survey in September-October 2014 and again in September-October 2018, in the four big regions (Oromia, Amhara, SNNP and Tigray) and 9 dominant migrant-sending districts in the country (see Figure 2).[8]

Figure 2: Sampled Districts and woredas

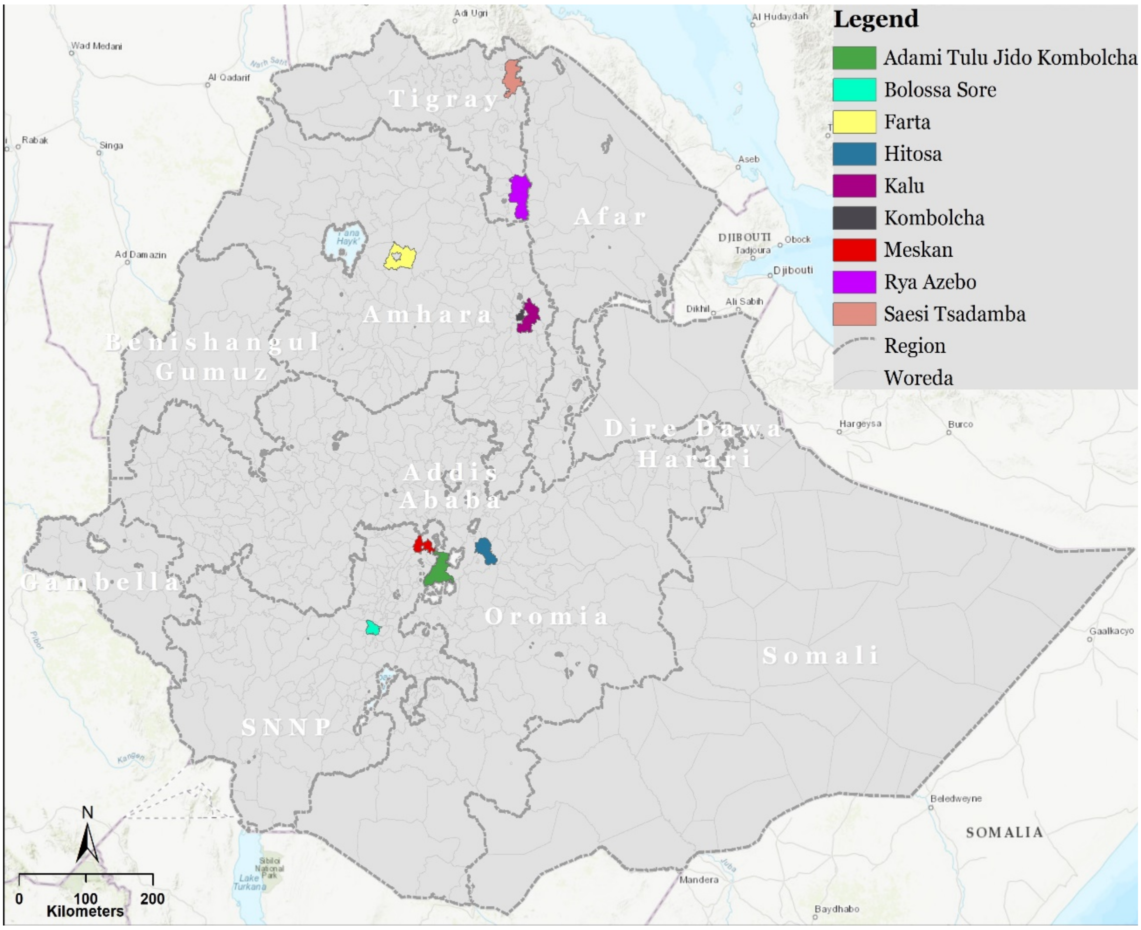

Source: Haile Tsegay and Litchfield (2019)

[6] Note however that the scale of internal migration will be underestimated as the LFS only picks up a change in *zone* of residence. Movements within zones, e.g. from one *woreda* to another, will not be considered as internal migration
[7] See Haile Tsegay and Litchfield (2019) for a fuller discussion.

[8] Migrating out of Poverty data for Ghana, Ethiopia, Zimbabwe, Indonesia and Bangladesh is available at www.migratingoutofpoverty.org

Migrants are defined as former resident members of the household who have moved out of the *kebele* within the last 10 years and have been away for a period of at least 3 months. The definition thus excludes very local and very short-term movements of people. We minimised attrition rates to around 1% by asking households to consent in 2014 to a follow up survey and by keeping detailed contact information, tracking households which had relocated to a neighbouring *kebele*. Households in 2014, our baseline survey, were selected randomly from two strata: a rapid listing exercise identified households in each *kebele* with and without current migrants, and samples were drawn from these in a roughly 2:1 ratio

Our data reveal substantial return migration between 2014 and 2018 of both internal and international migrants. In 2014 roughly a third of our households had no migrants (by definition given our sampling approach), with another third having internal migrants, and another third having either international migrants or a mixture of both internal and international migrants. In 2018 we see the share of households without migrants grow to around 60%, and it is this scale of return that motivates the analysis in this paper.

Figure 3: Sankey graph showing changes in household migration status between 2014 and 2018.

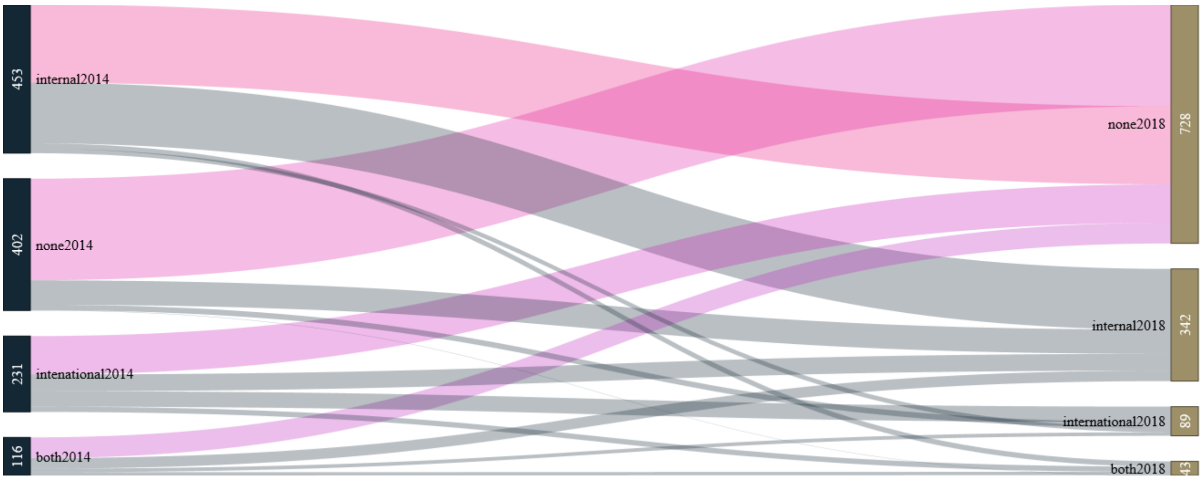

Source: authors' calculations from Migrating out of Poverty data

During this period, Ethiopia has experienced considerable and rising intra-ethnic conflict, escalating to violence in 2018. At the same time, migration of Ethiopians to the Middle East has come under scrutiny because of a number of highly publicised incidents of violence against migrants, particularly women. This has led to a suspension of migration between Ethiopia and a number of Gulf States, with agreements to resume migration yet to be implemented. We suggest that both of these factors contribute to return migration. Furthermore shifts in destinations of internal migrants, to *woredas* and zones closer to home than in previous years suggest a caution among potential migrants to move into areas where they are ethnic minorities.

Our data provides a sample of 300 return migrants: around 70% are returning internal migrants, and all the returning international migrants are form the Middle East. The reasons provided by the return migrants for their decisions are varied, from family issues and home-sickness to deportation (for the international migrants) and end of contracts. Few explicitly state that fears about violence were a key factor in their decision, but this may reflect the sensitivity of the situation. Our field work in 2018 took place during some of the worst violence Ethiopia has seen. A sizeable group returned home simply because they had either earned enough money or their work contract had finished.

Table 1 Main reason to return by destination		
Reason reported by return migrant	**Location of Migrant Prior to Return**	
	Middle East	Internal
Marital problems	0.00%	6.80%
Homesick	4.35%	4.37%
Earned enough money	4.35%	3.40%
Sick family member	7.61%	11.17%
Contract ended	7.61%	13.59%
Could not find work	1.09%	5.83%
Deported	30.43%	1.46%
To get married	2.17%	5.34%
Family issues	25.00%	27.18%
Permit not renewed	2.17%	1.46%
Migrant became ill	6.52%	5.34%
Concerns about violence	6.52%	4.37%
Other	2.17%	9.71%
Total	100%	100%

We also compare our sample or returnees to those migrants who are still away from home in 2018. Table 2 show that returnees are older, on average, than those still away, more likely to be male, slightly more likely to be married, to have children and more likely to be a returning head or spouse than a child of the household head. Notably they are twice as likely to be illiterate than those remaining away. This brief sketch suggests that return migrants are different from those migrants who remain away, providing further support to the possibility that return is not a random process, and that this needs to be addressed in our analysis.

Table 21 Individual level characteristics of migrants and returnees in wave two.		
	Migrant	Returnee
Age	24.09	27.12
Male	0.457	0.532
HH Head’s child	0.848	0.656
Married	0.309	0.343

Illiterate	0.0694	0.156
Has kids	0.223	0.404
Orthodox	0.582	0.473
Muslim	0.342	0.493
Tigray	0.228	0.299
Amhara	0.386	0.319
Oromiya	0.182	0.186
SNNP	0.204	0.196
Observations	764	301

Treatment and control groups

Our sample is comprised of 410 households. We define our treatment group as households which experience the return of at least one migrant between waves, and the control group as having migrants in both waves, with no returnees. Table 3 shows information on the treated and control group by wave. We have 155 households that had at least one migrant in wave one and at least one returnee in wave two. Of these, 72 households also have at least one migrant at follow up, and 83 only have at least one returnee but no other member of the household away[9]. Our control group is composed of 255 households who have at least one migrant in both waves, but none returning in wave two.

Table 3. Sample by Treatment And Control Status

	2014		2018	
	With migrants	Without migrants	With migrants	Without migrants
Treated	155	0[a]	72	83
Control	255	0[a]	255	0[a]

Note: [a] These cells are empty by definition as at baseline we only consider households with migrants at the time of the survey and our control group is those who have no return migrants in 2018.

The outcome variables are derived from a detailed section in the Migrating out of Poverty dataset on household expenditure of food and non-food goods and services, including own-produced food and gifts. We construct three measures of welfare: monthly food, monthly non-food consumption and monthly total consumption. In the non-food consumption section, each household is asked to report the amount spent in the last month for thirteen non-food items, and the amount spent in the last year for another fourteen non-food items consumed less regularly (clothes, furniture, ceremonial expenses, education, healthcare). For each household, we construct a non-food consumption value to be the natural logarithm of the sum spent for all these items in a period of one month.

[9] We explore the heterogeneity among the treatment group below.

Constructing food consumption values proved less straightforward as for each of the twenty-five items listed in the household questionnaire, each household could report to have purchased, produced or received as a gift said item. Thus, while quantity consumed in the past week is available for all items in the list, information on the amount spent for each item is only available for purchased ones. We construct a median price per unit at the *kebele* level for each of the items and use this as an indication of amount spent per unit. We then construct the food consumption variable by taking the natural logarithm of the sum spent by each household in a month. Total monthly household consumption is defined as the sum of these two measures. We use both household and per capita consumption values in our analysis.

All values of food and non-food consumption are in real terms and in 2014 Birr. We adjust prices using regional consumer price indices, constructed separately for food and non-food items and made available by the Central Statistical Agency of Ethiopia.

In addition we explore some of the components of consumption in a little more detail, namely consumption of animal products, which generally indicates the household is better off, and expenditure on specific non-food items.

Table 4 Summary statistics monthly consumption: household level and per capita

	2014		2018	
	Control	Treated	Control	Treated
Monthly total consumption	2099.8	2024.2	2174.1	2222.5
Monthly food consumption	1432.2	1388.5	1457.4	1358.3
Monthly non-food consumption	667.6	635.7	716.7	864.3
Monthly total consumption pc	412.7	426.5	337.6	347.8
Monthly food consumption pc	283.2	299.4	226.9	213.1
Monthly non-food consumption pc	129.4	127.1	110.6	134.7
Proportion of HHs which consumed animal products in last 7 days	0.48	0.49	0.39	0.40
Yearly non-food consumption				
- Clothing items	17255	16398	18443	21160
- Health & education	9347	8694	10299	11463
- Household items	11054	11385	11615	12109
Remittances received in previous 12 months	3496	2757	2519	5196
Observations	405	405	405	405

Table 4 provides some context for our sample. We see that our sample is, on average, poor: food share of total expenditure is between 60 and 70%, which is higher than the country average, suggesting our rural *kebeles* are drawn from the lower part of the welfare distribution.[10] We also see that welfare rankings and trends over time are very sensitive to how we measure welfare. Generally, the measures of household level consumption variables show an improvement over time, particularly for non-food consumption, but the per capita figure generally suggests the opposite. This suggests that household size is increasing by proportionately more than consumption levels. We believe this may reflect the particular characteristics of land use and land rights in rural Ethiopia, whereby land usufruct rights are granted to households with conditions that land is used productively, and of the strong gender norms surrounding farm labour, with women often precluded from some agricultural activities, for example tasks involving the use of animal labour. The majority of current migrants in our sample are male (see Haile Tsegay and Litchfield, 2019) and it is plausible that with declining availability of household farm labour, households replace family labour with hired, live-in labour.

Comparing how our treatment and control groups fare over the period, we see a reversal in the rankings over time. Initially, households in our control group (those households which have migrants in both waves of the survey and no return migrants) appear to be better off than the treatment group, but this ranking switches by 2018, and is particularly strong for non-food consumption, which grows by almost 40% for the treatment group, compared to less than 10% for the control group.[11] the rise in non-food consumption is notably for clothing items, but not limited to clothing with expenditure of health, education and general household items all increasing at a higher rate for the treatment group.

One possible reason for this is the difference in remittances received by households in the twelve months prior to the survey. Those household with returning migrants experienced an increases in remittances over time, while those with continued migration experience and no returnees see a fall. We explore the relevance of these factors in our analysis below.

In order to analyse whether there is any impact of return migration on household welfare we first explore to what extent return migration is a random phenomenon, whether characteristics of the household make them more likely to have a return migrant, and whether these characteristics might be associated with higher or lower levels of welfare.

Table 5 below shows some key characteristics of the households by treatment and control stats of households in each wave. We do indeed see that at baseline, the treatment households are more likely to be unemployed, female-headed, have older heads and heads with less education than those households which do experience return over the period. All of these factors are likely to be associated with lower welfare levels, and motivates our approach to match households at baseline to control for differences in observables between the two groups.

[10] Hassen et al (2016) estimate the average food share in Ethiopia in 2011 to be around 48% and on a declining trend.

[11] And per capita non-food consumption drops for the control group, while rises for the treatment group.

Table 5 Summary statistics of household characteristics				
	2014		2018	
	Control	Treated	Control	Treated
Household size	5.267	5.221	6.780	6.768
HH Male	0.717	0.636	0.752	0.729
HH Age	55.34	57.36	57.51	56.63
HH Married	0.757	0.695	0.792	0.729
HH no school	0.558	0.669	0.616	0.542
HH Employed	0.785	0.740	0.708	0.781
HH has kids	0.689	0.578	0.560	0.600
Orthodox	0.594	0.539	0.536	0.568
Catholic	0	0	0.004	0
Protestant	0.0637	0.0195	0.0840	0.0129
Muslim	0.343	0.442	0.376	0.419
Tigray	0.251	0.357	0.240	0.361
Amhara	0.323	0.279	0.324	0.290
Oromiya	0.195	0.195	0.208	0.181
SNNP	0.231	0.169	0.228	0.168
Observations	405	405	405	405

Results

Table 6 show the results of our analysis, which is the effect on household consumption of having a return migrant, estimated using equation (5) and restricting the analysis to a comparable sample matched on observable characteristics. $Treated_h$ in equation (5) is a binary variable equal to one if the household has at least one migrant away at baseline and at least one returning migrant at follow-up and equal to zero if it has at least one migrant away in both periods and none returning.

The results show an increase in overall consumption, both per capita and household, and increases in non-food consumption, with a small decline in food consumption. Only the increase in non-food consumption is statistically significant. The results tell us that having a return migrant increases non-food consumption by 22-23 percent, but has no statistically significant effect on either food or on total consumption. We ran the same model but without matching and obtained a larger estimate, of around 30 percent, suggesting that not controlling for the non-randomness of return migration would over-state the effect of return on household welfare.[12]

[12] We also try an alternative method to propensity score matching, inverse probability weighting. Our results are shown in the Appendix but ae qualitatively the same as those we find with psm.

Table 6 The effect of having a returning migrant on consumption: household fixed effects and matching						
	(1)	(2)	(3)	(4)	(5)	(6)
	Household			Per-capita		
Consumption	Total	Food	Non-food	Total	Food	Non-food
Wave	0.0242	-0.000754	-0.00642	-0.259***	-0.282***	-0.296***
	(0.0537)	(0.0548)	(0.0802)	(0.0512)	(0.0540)	(0.0753)
Treat	-0.00632	0.00240	-0.0388	-0.0200	-0.0113	-0.0525
	(0.0537)	(0.0548)	(0.0802)	(0.0511)	(0.0538)	(0.0750)
Treated x Wave	**0.0508**	**-0.0121**	**0.223****	**0.0511**	**-0.0145**	**0.229****
	(0.0759)	**(0.0775)**	**(0.113)**	**(0.0723)**	**(0.0763)**	**(0.106)**
Observations	766	766	766	763	763	763
R-squared	0.003	0.000	0.012	0.053	0.071	0.023

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Mechanisms between return and consumption

We explore a number of avenues in order to better understand these results.

Although expenditure on food does not appear to increases we explore whether having a return migrant improves the quality of the household diet. In column 1 of table 7 we estimate whether treated households become more likely to consume animal products (milk, eggs, cheese and meat). We do not find an effect on diet quality.

We then investigate whether the positive effect of having a returning migrant on non-food consumption is concentrated in certain expenditure categories. In column 2, we estimate the effect of having a returning migrant on expenditure on health and education for members of the household. This is the natural logarithm of the monthly expenditure of the household in the following items as listed in the questionnaire: education (fees, books, uniforms) and health (consultation fees, medicines and medical supplies). We find a positive but not statistically significant effect on this outcome. We then look at whether the household spends more on household items: kitchen equipment, linens, furniture, lamps/torches (column 4). We find that having a returning migrant increases the expenditure in this category by 39 percent. We also find a 43.5 percent increase in expenditure on clothing, both for adults and for boys and girls aged less than 18 (column 5).

Finally we also explore the role of remittances. Return migrants often send remittances prior to return, transferring assets and savings home in preparation for their visit. We measure remittances for households in the control group as the sum of remittances received in the last 12 months by each migrant away (and take the natural log). For households in the treatment group, we do the same exercise for wave one, while for wave two we also add any money received by the household in the previous 12 months by the returning migrant. Our data suggests that households with return migrants received substantially more remittances in the twelve months prior to the survey, and this is compared to the control group of households which still have migrants away from home. We find that households with a returning migrant experience a 135 percent increase in remittances received.

Table 7 The effect of having a returning migrant on food and non-food categories, and remittances					
	(1)	(2)	(3)	(4)	(5)
VARIABLES	Consumed Animal Prod last week	Non-food: Health &Edu	Non-food: household	Non-food: clothing	Remittances last 12 months
Treated	0.0254	-0.0559	-0.0779	-0.0854	0.215
	(0.0522)	(0.102)	(0.133)	(0.107)	(0.367)
Wave	-0.0702*	-0.0201	-0.313**	-0.119	-0.874***
	(0.0419)	(0.0901)	(0.129)	(0.0918)	(0.331)
Treated x Wave	**-0.0137**	**0.160**	**0.392****	**0.435*****	**1.350****
	(0.0685)	**(0.136)**	**(0.194)**	**(0.135)**	**(0.558)**
Observations	766	697	420	729	766
R-squared	0.006	0.002	0.016	0.015	0.021

Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

Heterogeneity by household type

This section revisits the definition of treatment and control, noting that the treatment group of households with returning migrants may continue to have migrants away from home. Households with returning migrants can be split into two subsamples: households that have no migrants by wave two, and households that still have at least one migrant. The effect of having a returning migrant may vary between with the latter potentially continuing to receive remittances.

We use a triple difference estimator to evaluate the heterogeneous effect of having at least one returning migrant and at least one migrant away using the following model:

$$\begin{aligned} consumption_{h,t} &= \beta_0 + \beta_1 Treated_h \times Wave_t + \beta_2 Wave_t + \beta_3 Treated_h \\ &+ \beta_4 Treated_h \times Wave_t \times Migrant_{h,2} + \beta_5 Migrant_{h,2} + \varepsilon_{h,t} \end{aligned}$$

β_4 measures the effect of having at least one migrant away by wave two conditional on being a treated household, i.e. with at least one returning migrant, compared to being treated but not having any migrants away. We opt for this methodology to increase the power of our estimation. We use only observations on the common support when estimating this model.

Table 8 shows the summary statistics for the outcome variables by these two subgroups. While we do observe an increase in total monthly consumption between the two waves, this is mainly driven by the increase in non-food consumption in both per household and per capita terms. The largest difference is in the amount of remittances received in the second period. In particular, the amount of money received by returning migrants in the previous 12 years is more than six times larger than what received as remittances in wave one while both groups had migrants away. This suggests that returning migrants bring back home a large sum of savings after their migration experience.

Table 8 Summary statistics for treated households: by heterogeneity whether with/without migrants away by wave two.

	Wave 1		Wave 2	
	Treated		Treated	
	No migrants in wave 2	1 or more migrants in wave 2	No migrants in wave 2	1 or more migrants in wave 2
Monthly tot consumption	1937.1	2123.5	2057.5	2412.8
Monthly food consumption	1325.6	1460.2	1288.7	1438.4
Monthly non-food consumption	611.5	663.3	768.8	974.4
Tot consumption pc	407.1	448.5	363.2	330.0
Monthly food consumption pc	281.6	319.6	227.8	196.1
Monthly non-food consumption pc	125.5	129.0	135.4	134.0
Total yearly remittances	2862.7	2670.1	4795.8	5657.6
Remittances from returning migrant (last 12 months)	-	-	9710.9	9370.1
Observations	83	72	83	72

Table 9 Heterogeneity between households with at least one returning migrant and with at least one migrant away

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Household			Per-capita			
	Total consumption	Food consumption	Non-food consumption	Total consumption	Food consumption	Non-food consumption	Remittances
Treated	-0.103	-0.104	-0.103	-0.0627	-0.0644	-0.0637	0.232
	(0.0751)	(0.0821)	(0.0956)	(0.0674)	(0.0765)	(0.0882)	(0.448)

Wave	0.00889	-0.00678	-0.0278	-0.272***	-0.289***	-0.304***	-0.874***
	(0.0475)	(0.0500)	(0.0681)	(0.0438)	(0.0476)	(0.0641)	(0.332)
Treated x Wave	0.0843	0.0316	0.248**	0.172**	0.119	0.336***	0.971
	(0.0871)	(0.0968)	(0.120)	(0.0800)	(0.0901)	(0.115)	(0.723)
Migrant	0.152*	0.179**	0.113	0.0808	0.106	0.0477	-0.0366
	(0.0858)	(0.0897)	(0.120)	(0.0853)	(0.0936)	(0.109)	(0.559)
Treated x Wave x Migrant	**0.0327**	**-0.0409**	**0.101**	**-0.181***	**-0.253****	**-0.118**	**0.817**
	(0.103)	**(0.111)**	**(0.149)**	**(0.0980)**	**(0.111)**	**(0.136)**	**(0.892)**
Constant	7.502***	7.116***	6.211***	5.931***	5.546***	4.635***	4.061***
	(0.0382)	(0.0402)	(0.0512)	(0.0312)	(0.0338)	(0.0454)	(0.240)
N	766	766	766	762	762	762	766
R-squared	0.014	0.009	0.019	0.058	0.076	0.029	0.022

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

The results show that households with at least one returning migrant and with at least one migrant away experience a decrease of 25 percent in food-consumption per capita relative to households with at least one returning migrant in wave 2 and no migrants still away. This drives a decrease of 18 percent in total consumption per capita. The differences in household consumption are not statistically different from zero. In terms of total remittances, households with at least one migrant away receive higher remittances, in line with what expected, but this result is not statistically different from zero.

This worsening of per capita consumption mirrors the results of the main analysis, where we find that households with returning migrants have higher consumption than households with migrants still away. This might suggest a worsening of the migration experience for households with migrants away in both waves due to a drop in international migration, the uncertainties caused by inter-ethnic conflict or to a decline in the value of remittances received, which the returning migrant is not capable to offset. In addition, this result might suggest a short-term negative effect of migration on per capita consumption. It might be that the household sees a previous migrant coming back between the two waves, but a different member of the household migrating by wave two[13]. This finding is be in line with other literature (Gibson et al, 2011) that migration can have short-term

[13] Our measures of migration are defined at the household level. It might be that the household has at least a returning migrant by wave two and still a migrant away in the same period. However, this does not mean that the same member of the household was away in both periods. In particular, it might be that the as one member returns another member migrates. It is also possible that the household had two migrants away in wave one and only one returns by wave two. The effect of these different migration patterns is plausibly different on consumption levels.

negative effects on food consumption if it decreases agricultural labour productivity, disrupts the decision-making process in the household or changes the household gender and age composition[14]. Having a returning migrant might not be enough to offset these negative effects if, for example, the returning member is older or has had to return earlier than anticipated. But note, this is speculative on our part given the small sample we are working with here.

We are unable to disentangle these mechanisms in our analysis, both because of our measure of migration at the household level and because of the limited sample size when delving into sub-group analysis. However, presenting this heterogeneity analysis shows that migration is a complex phenomenon and understanding its effects on household welfare depends on different dimensions.

Conclusions

The aim of this paper was to assess the impact of return migration on welfare of the households to which migrants return. We evaluate this in comparison to households with migrants still away from home and no return migrants. We offer methodological guidelines to researchers interested in assessing the relationship between migration and household economic conditions. We began with an overview of the common problems and challenges in the literature on the effects of migration, with a focus on selection bias. It then provided an overview of econometric and statistical tools available commonly used in the literature on migration. By discussing more in detail these econometric tools, this paper aims to explain in an accessible way the main assumptions behind econometric techniques to researchers who are approaching an empirical evaluation of the effects of migration, with a focus on a non-experimental setting and contexts where data availability might be an issue.

The paper explores the impact of return migration on household welfare using difference-in-differences estimation combined with matching of treatment and control households at baseline, in line with best practice in the literature.

Returnees are asked to report the main reason for this return. The literature often cites return as being a sign of a failed migration experience, despite evidence that a large proportion of migrants return. We find that a large proportion (around a third) of international migrants have returned following deportation, and family issues are responsible for around a quarter of both internal and international migrants. However a sizeable minority of migrants return because they had finished their contracts or earned enough money. Moreover we find that remittances sent home by returnees prior to return are substantial, suggesting that returnees are able to make a significant contribution to local economic development and that their return should not be viewed as a failed venture.

We find some evidence of selection in return. Returnees are slightly more likely to be male than migrants still away , suggesting that women are slightly less likely to return, and that returnees are also around 3 years older on average than those still away. We find higher incidence of return amongst Muslim migrants, and among those originating from Tigray.

[14] Mueller et al. (2019) show that in Ethiopia is usually young productive male members who migrate. This might have negative effects in terms of agricultural labour productivity both because it reduces the household size and because of the gender-specific division of labour present in the country.

We measure welfare using household consumption, both food and non-food. We find that return has no impact on food consumption, either in terms of the expenditure on food or on the dietary diversity of food consumed by the household. We do however find a statistically significant impact of return on non-food expenditure, with our estimates suggesting return increases non-food consumption by just over 20% compared to households with migrants still away and no returnees. Exploring the composition of expenditure, we find that much of this increase is driven by increased expenditure on clothing and household items such as kitchen utensils, furniture and linens. We explore the heterogeneity of these results by whether households with return migrants also have migrants still away but the small sample size makes this analysis very tentative.

Our analysis also shows that return is associated with a very large increase in remittances sent home by the migrant in the 12 months prior to return. This suggest that migrants are able to accumulate significant savings while away from home and that even if their return may not have been anticipated precisely, it does not necessarily follow that return migration should be viewed as an example of a failed migration. We conclude therefore that return migration, far from being a sign of a failed attempt at migration, is much more nuanced, and that returning migrants have the potential to contribute towards development.

References

Amuedo-Dorantes, Catalina and Pozo, Susan, Accounting for Remittance and Migration Effects on Children's Schooling, *World Development*, 38, issue 12 (2010) p. 1747-1759.

Azose, Jonathan J. and Adrian E. Raftery "Estimation of emigration, return migration, and transit migration between all pairs of countries" Proceedings of the National Academy of Sciences Jan 2019, 116 (1) 116-122.

Beegle, Kathleen, Joachim De Weerdt, and Stefan Dercon. "Migration and economic mobility in Tanzania: Evidence from a tracking survey." *Review of Economics and Statistics* 93.3 (2011): 1010-1033.

Binzel, C., and R. Assaad. 'Egyptian Men Working Abroad: Labour Supply Responses by the Women Left Behind'. Labour Economics, 18 (2011)

Black, D. A., Sanders, S. G., Taylor, E. J., & Taylor, L. J. The Impact of the Great Migration on Mortality of African Americans: Evidence from the Deep South. *The American Economic Review*, *105*(2, 2015), 477–503.

Brauw, Alan de, Valerie Mueller, and Tassew Woldehanna. "Does Internal Migration Improve Overall Welfare in Ethiopia? †." *Journal of African Economies* 27, no. 3 (June 1, 2018): 347–65.

Bryan, G, S Chowdhury and A Mobarak , "Under-investment in a profitable technology : The case of seasonal migration in Bangladesh", *Econometrica* 82(5 2014): 1671–1748.

Card, D. "The Impact of the Mariel Boatlift on the Miami Labor-Market", Industrial & Labor Relations Review, Vol. 43 (2, 1990), 245-257.

Clemens, Michael, and Timothy Ogden. "Migration as a Strategy for Household Finance: A Research Agenda on Remittances, Payments, and Development," 2014.

Clemens, Michael, and Tiongson. "Split Decisions: Household Finance When a Policy Discontinuity Allocates Overseas Work | The Review of Economics and Statistics | MIT Press Journals." *The Review of Economics and Statistics*, July 2017.

Deaton, Angus. "Measuring poverty." Understanding poverty (2006): 3-15.

Dehejia, Rajeev. "Practical Propensity Score Matching: A Reply to Smith and Todd." *Journal of Econometrics* 125, no. 1–2 (2005): 355–364.

Dehejia, Rajeev H., and Sadek Wahba. "Propensity Score-Matching Methods for Nonexperimental Causal Studies." *Review of Economics and Statistics* 84, no. 1 (2002): 151–161.

Demirgüç-Kunt, A., Córdova, E. L., Pería, M. S. M., & Woodruff, C. Remittances and banking sector breadth and depth: Evidence from Mexico. *Journal of Development Economics*, *95*(2, 2011), 229-241.

Gibson, John, and David McKenzie. "The Development Impact of a Best Practice Seasonal Worker Policy." *Review of Economics and Statistics* 96, no. 2 (2014): 229–243.

Gibson, John, David McKenzie, and Steven Stillman. "The Impacts of International Migration on Remaining Household Members: Omnibus Results from a Migration Lottery Program." *Review of Economics and Statistics* 93, no. 4 (2011): 1297–1318.

Haile Tsegay, Asmelash , and Julie Litchfield (2019) "Changing patterns of migration and remittances in Ethiopia 2014-2018" *Migrating out of Poverty* Working Paper 59 Brighton: : School of Global Studies, University of Sussex.

Hassen, Ibrahim Worku; Dereje, Mekdim; Minten, Bart; and Hirvonen, Kalle. "Diet transformation in Africa: The case of Ethiopia" ESSP Working Paper 87 (2016). Washington, D.C. and Addis Ababa, Ethiopia: International Food Policy Research Institute (IFPRI) and Ethiopian Development Research Institute (EDRI).

Heinrich, Carolyn, Alessandro Maffioli, and Gonzalo Vazquez. "A Primer for Applying Propensity-Score Matching." Inter-American Development Bank, 2010.

Hildebrandt, Nicole, and David J. McKenzie. The effects of migration on child health in Mexico. The World Bank, 2005.

International Organisation for Migration (IOM) "Glossary on Migration", IOM 2019.

Karymshakov, Kamalbek and Burulcha Sulaimanova. "Migration impact on left-behind women's labour participation and time-use" UNU WIDER Working Paper 2017/119

McKenzie, David, and Hillel Rappoport. «Network effects and the dynamics of migration and inequality: Theory and evidence from Mexico.» Journal of Development Economics 84 (2007): 1-24.

McKenzie, David and Marcin J. Sasin Migration, remittances, poverty, and human capital: conceptual and empirical challenges. The World Bank, 2007.

McKenzie, David, Steven Stillman, and John Gibson. "How Important Is Selection? Experimental vs. Non-Experimental Measures of the Income Gains from Migration." *Journal of the European Economic Association* 8, no. 4 (2010): 913–945.

McKenzie, David and Dean Yang (2012) "Experimental Approaches in Migration Studies", in Carlos Vargas-Silva ed. Handbook of Research Methods in Migration. Edward-Elgar.

Munshi, Kaivan. «Networks in the modern economy: Mexican migrants in the U.S. labor market.» The Quarterly Journal of Economics, May 2003.

Smith, Jeffrey A., and Petra E. Todd. "Does Matching Overcome LaLonde's Critique of Nonexperimental Estimators?" *Journal of Econometrics* 125, no. 1–2 (2005): 305–353.

United Nations Population Division (UNDP) Word Population Prospects 2017

Vadean, Florin, Teresa Randazzo, and Matloob Piracha. "Remittances, Labour Supply and Activity of Household Members Left-Behind." The Journal of Development Studies 55.2 (2019): 278-293.

Wahba, Jackline, "Who benefits from return migration to developing countries?" IZA World of Labor, February 2015a.

Wahba, Jackline "Selection, selection, selection: the impact of return migration." *Journal of Population Economics*, 28 (3) (2015b): 535–563.

Wahba, J., & Zenou, Y. "Out of sight, out of mind: migration, entrepreneurship and social capital." *Regional Science and Urban Economics*, 42 no 5 (2012), 890-903.

World Development Indicators https://datacatalog.worldbank.org/dataset/world-development-indicators.

Yang, Dean. "International Migration, Remittances and Household Investment: Evidence from Philippine Migrants' Exchange Rate Shocks." *The Economic Journal* 118, no. 528 (2008): 591–630.

Appendix

1. Propensity Score Matching

We are interested in the effect of having a returning migrant in the household on household consumption. A straightforward way to do this is to compare the welfare of households in the treatment group to that of households in the control group. For this comparison to be meaningful, we want our treatment and control group to be as comparable as possible, i.e. we want them to differ only because of the difference in treatment status.

In the context of a non-experimental setting, propensity score matching can be used to balance the treatment group and the control group in a way that this comparison is achievable. Propensity score matching uses observable characteristics to match observations in the treatment group to the most comparable observations in the control group. The simplest way to do this matching when we have only one observable characteristic, say household size, is to use *nearest neighbour covariate matching*, which uses the most similar untreated unit to build the counterfactual of a given treatment unit. This, however, is reasonably not enough to compensate potential confounding factors that might affect income in the treatment and control group and not be associated with household size, but to the treatment. To deal with this multidimensionality in the characteristics that differentiate the treatment and the control group, we can estimate propensity scores.

These scores will allow us to summarize the information in one single number that captures the probability that the household is in the treatment group given a set of characteristics. If this set of characteristics is on average very similar for an untreated and treated unit, we can match the propensity scores of these two units. The rationale behind the propensity score matching is that using observable information we can predict the probability of a household having at least one return migrant. Hence, if for the same probability, say close to 0.8[15], we find two households where one is treated and one is not, we can assume that conditional on those observable characteristics the only difference between the two units is the treatment effect. Once the matching has taken place, we can estimate the effect of the treatment to be the difference in outcome between the treated and control group.

We estimate a participation model for our data (based on equation 4 above), which models the probability that a household is in the treatment group. Table A1 shows the results of this participation model estimated using a logit regression, and baseline values of each covariate. We include in our regression only binary variables indicating head of household characteristics: age, gender, marital status, employment status, education level, religion, region or residence. We also include binary variables indicating the quintile of the total household consumption distribution to which the household belongs. [16]

---

[15] A more detailed explanation of different matching procedures can be found in Heinrich et al. (2010). In general, we might expect that the treated and untreated unit are not an exact match, i.e. both with a propensity score of 0.8. Given reduced sample size, we might want to match control and treatment unit with a propensity score as close as possible to 0.8. A whole additional discussion should be made about whether the matching happens with or without replacement, and whether we use only one control unit to match to a treatment unit or a set of units.

[16] De Brauw et al. (2018) include the following covariates in estimating their propensity scores: age, gender and education of the migrant, age and education of the household head, and proxy variables for the source household's wealth (landholdings and number of livestock units). Gibson and McKenzie (2014) include:

Table A1 Logit model estimation for the propensity scores	
VARIABLES	Treat
Household size	0.0227
	(0.0587)
HH Male	0.350
	(0.445)
HH Age	0.00413
	(0.00962)
HH Married	-1.006**
	(0.465)
HH no school	-0.367
	(0.251)
HH Employed	0.148
	(0.275)
HH has kids	0.237
	(0.269)
Orthodox	-0.448
	(0.277)
Protestant	-3.244***
	(1.133)
Amarigna	-13.27
	(511.5)
Guraghigna	-1.150
	(0.910)
Oromifa	-1.194
	(0.859)
Siltigna	-1.541

demographic variables (household size, number of adults, school-aged children, males aged 18 to 50) characteristics of males in working age (literacy test, health status, number of days worked), household's past migration experience, household baseline assets and housing infrastructure, geographical characteristics, past household's wage history. Plus the square of all of the above.

	(1.116)
Tigrigna	-0.418
	(0.826)
1st quintile	-0.493
	(0.398)
2nd quintile	0.176
	(0.350)
3rd quintile	-0.489
	(0.356)
4th quintile	0.00813
	(0.332)
Amhara	12.18
	(511.5)
Constant	0.990
	(1.059)
Pseudo R-squared	0.066
Observations	403

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

Once the model is estimated, for each unit will we be able to estimate the logarithm of the odds ratio of each propensity score defined as p/(1-p). It is common to use the log of the odds ratio to do the matching, as using the propensity score directly might not be robust to "choice-based sampling". We then need to decide which algorithm to use in the matching. The default option is usually nearest neighbour matching (this is the case for psmatch2[17] command in STATA). There are other matching algorithms: radius matching or matching on a maximum propensity score distance, and kernel matching with uses nonparametric techniques. Once the matching procedure has been performed, the researcher will evaluate the difference in mean in the outcome of interest for the treated and untreated group.

We apply propensity score matching techniques to our analysis to make the treatment and control group as comparable as possible. Figure 4 shows graphically the result of our matching. As expected the matching makes the distribution of the propensity scores for the treatment and control group much more similar. Thus, once we restrict our regression analysis to only observations with a propensity score on the common support, the two groups are comparable. In light of what we have

[17] Notice that *psmatch2* includes an option to estimate the propensity scores (and the logs odds ratio) in the same line of command that performs the matching algorithm. In this case, however, the independent variables to be included in the estimation of the propensity score need to be always specified.

explained so far, then, the estimation of the treatment effect will be unbiased. We conduct the analysis for households with a propensity score falling in the common support range [0.114-0.701] (for the control [0.00-0.732], for the treated [0.114-0.701]).

Figure 4 Distribution of the propensity scores before and after matching

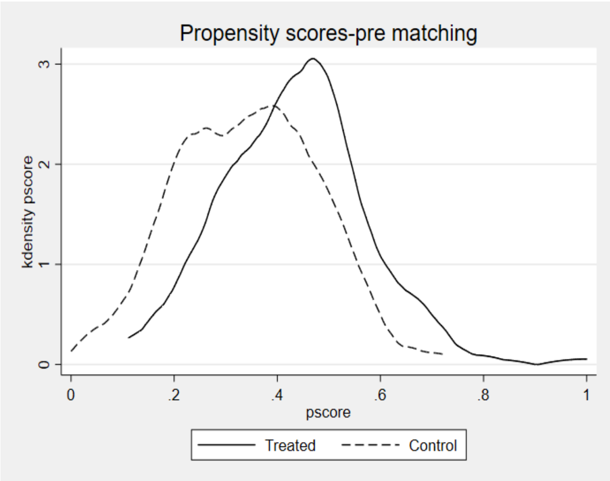

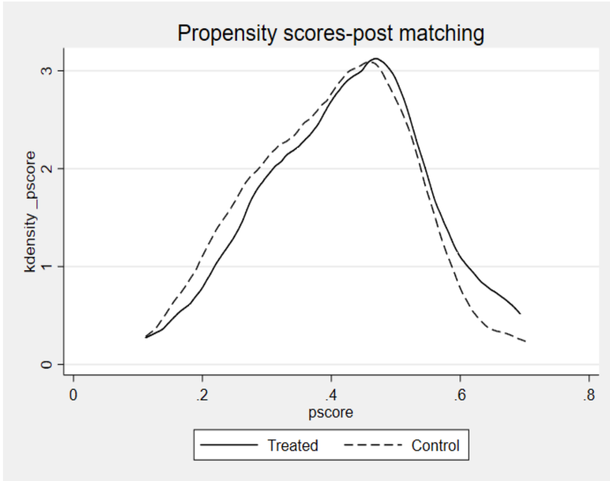

2. Robustness checks

An alternative to propensity score matching is to use inverse propensity weights in the estimation of equation (5). Table A2 shows consistent results of an increase in non-food consumption. The larger magnitude and higher statistical significance of the results arises from the inclusion of less comparable households in the treatment and control group and from the higher power due to a larger sample size. This suggests that propensity score matching techniques do a better job at reducing the impact of selection bias on the estimates.

Table A2 The effect of having at least a returning migrant in the household on consumption: inverse propensity weights

	(1)	(2)	(3)	(4)	(5)	(6)
		Household			Per-capita	
	Total	Food	Non-food	Total	Food	Non-food
Wave	-0.0567	0.0398	-0.319*	-0.0120	0.0875	-0.281*
	(0.120)	(0.125)	(0.172)	(0.114)	(0.126)	(0.152)
Treat	0.0427	0.0267	-0.0177	-0.237***	-0.250***	-0.304***
	(0.0459)	(0.0487)	(0.0656)	(0.0435)	(0.0472)	(0.0629)
Diff-in-diff	**0.0597**	**-0.0223**	**0.282*** **	**0.0272**	**-0.0578**	**0.256*** **
	(0.0721)	**(0.0761)**	**(0.107)**	**(0.0691)**	**(0.0784)**	**(0.0950)**

Observations	798	798	798	795	795	795
R-squared	0.006	0.000	0.018	0.048	0.067	0.025

Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1

About Migrating out of Poverty

Migrating out of Poverty research programme consortium is funded by the UK's Department for International Development (DFID). It focuses on the relationship between migration and poverty – especially migration within countries and regions – across Asia and Africa. The main goal of **Migrating out of Poverty** is to provide robust evidence on the drivers and impacts of migration in order to contribute to improving policies affecting the lives and well-being of impoverished migrants, their communities and their countries through a programme of innovative research, capacity building and policy engagement.

Migrating out of Poverty is coordinated by the University of Sussex and led by Research Director Dr Priya Deshingkar and Dr Robert Nurick as Executive Director. Core partners are the Centre for Migration Studies (CMS) at the University of Ghana, and the African Centre for Migration & Society (ACMS) at the University of the Witwatersrand in South Africa, the Organisation for Social Science Research in Eastern and Southern Africa (OSSREA) at Addis Ababa University, Ethiopia and L'Université Assane Seck Ziguinchor (UASZ) in Senegal. Past partners included the Refugee and Migratory Movements Research Unit (RMMRU) in Bangladesh, the Asia Research Institute (ARI) at the National University of Singapore; and the African Migration and Development Policy Centre (AMADPOC) in Kenya. Please visit the website for more information.

Migrating out of Poverty

University of Sussex, Arts B

Brighton BN1 9QN, United Kingdom

Email: migrationrpc@sussex.ac.uk

Web: www.migratingoutofpoverty.org

Twitter: @MigrationRPC

Facebook: /migratingoutofpoverty

Blog: migratingoutofpoverty.blogspot.com