

EXPERIMENTAL APPROACHES IN MIGRATION STUDIES

David McKenzie, *World Bank*
Dean Yang, *University of Michigan*

Abstract

The decision of whether or not to migrate has far-reaching consequences for the lives of individuals and their families. But the very nature of this choice makes identifying the impacts of migration difficult, since it is hard to measure a credible counterfactual of what the person and their household would have been doing had migration not occurred. Migration experiments provide a clear and credible way for identifying this counterfactual, and thereby allowing causal estimation of the impacts of migration. We provide an overview and critical review of the three strands of this approach: policy experiments, natural experiments, and researcher-led field experiments. The purpose is to introduce readers to the need for this approach, give examples of where it has been applied in practice, and draw out lessons for future work in this area.

Keywords: Migration, Remittances, Experiment, Identification, Self-Selection.

Chapter prepared for the forthcoming *Handbook of Research Methods in Migration*

1. Introduction

Individuals and households decide whether or not to migrate - and whether or not to send remittances if they do migrate - with the outcome of these choices depending upon their skills, wealth, risk preferences, ambition, drive, family ties, and a myriad of other observable and unobservable characteristics. This self-selection of migrants poses a severe challenge for researchers attempting to ascertain the impacts of migration or remittances on individuals, families, and communities. For example, suppose we observe that children are more likely to attend school in households with a migrant than in households without a migrant. This may reflect the income effect of remittances, but could just as easily reflect that children in households with migrants have higher quality parental education, or better language skills, or that it is parents who care most about the education of their children who migrate to earn the money needed to pay for schooling costs. As a result, even if we condition on a wide array of observable characteristics, comparisons of migrants and non-migrants are unlikely to give convincing estimates of the impacts of migration.

Experimental approaches to migration studies aim to overcome this difficulty by exploiting situations where the reason one household engages in migration or remits and another does not is truly the result of random chance. This may occur as a result of policy experiments, such as visa lotteries; through natural experiments whereby “nature” provides the source of exogenous variation; and through researcher-led field experiments which are explicitly designed to test specific theories of constraints to migration or remittance behavior. The purpose of this chapter is to introduce readers to the need for this approach, give examples of where it has been applied in practice, and draw out lessons for future work in this area. We begin with a short discussion to illustrate the perils and challenges of trying to estimate the causal impacts of migration or remittances using non-experimental approaches, and then discuss the different experimental approaches, before concluding with lessons for future work.

2. The challenge of assessing the causal impact of migration¹

A large part of the development literature in migration attempts to answer questions of the form “what is the effect of migration or remittances on outcome Y?”. One branch of this focuses on the migrants themselves, and is interested in how migration changes their incomes, health, stress levels, and life opportunities. A second branch focuses on remaining household members and communities in the sending areas, and is interested in the impact of having a household member migrate or of receiving remittances on the education and health of children, on levels of entrepreneurship and labor supply of adult members, and on poverty and inequality levels in the village. A common approach to answering such questions is to use survey data on migrants and non-migrants, and attempt to control for differences between them in a linear regression framework. We will set this out for the case of estimating the impact of migration, but the challenge is analogous for estimating the impact of remittances. For example, researchers may attempt to estimate an equation of the form:

$$Y_i = \alpha + \beta \text{Migrant}_i + \delta' X_i + \varepsilon_i \quad (1)$$

¹ See McKenzie (2005) and Gibson et al. (2010) for more detail on these challenges.

Where $Migrant_i$ is a dummy variable which takes the value 1 if the individual is a migrant, and zero if he or she is not, Y_i is the outcome of interest (e.g. individual income), and X_i are a set of observed characteristics of the individual which are presumed not to have changed with migration (e.g. age, sex and education level, location of birth, ethnicity, religion, parental education, etc.).² Then in order for the linear regression estimate of β to give the causal impact of migration on the outcome of interest, we require that:

$$E(Migrant_i \varepsilon_i | X_i) = 0 \quad (2)$$

That is, we require the unobserved determinants of the outcome of interest (income in our example) to be uncorrelated with whether or not an individual migrates once we have conditioned on the observable characteristics of these individuals. But in the absence of experimental variation in migration, this assumption is unlikely to hold. Indeed, the seminal migration selectivity model of Borjas (1987) has migrants deciding whether or not to migrate in part on the basis of the ε_i they would expect to have at home versus abroad. It is easy to think of a whole range of typically unmeasured variables, such as entrepreneurial prowess, ambition, language proficiency, and health status which would both affect whether or not someone migrates, and also directly affect their income or other outcome of interest. Likewise liquidity constraints will likely determine both the pattern of self-selection of migrants (McKenzie and Rapoport, 2010), as well as the range of job opportunities and consumption-smoothing opportunities that individuals will have at home. As a result, equation (2) will almost always be violated in practice, so that linear regression on equation (1) will result in biased estimates of the impact of migration.

Equation (2) therefore says that we can only estimate the causal impact of migration if the only reason one person migrates and another does not is random (conditional on observed characteristics). That is, ideally we would randomly choose some people to migrate and others not to, and then by comparing these two groups, get the impact of migration. This is precisely what the experimental approach attempts to do.

3. Examples of the Experimental Approach

3.1. Policy Experiments

Several countries use visa lotteries to choose among numerous applicants desiring to immigrate through a particular migration category that has a fixed quota. The most famous of these is the United States Diversity Visa Lottery (commonly known as the Green Card Lottery), which each year makes available 50,000 visas, to be drawn randomly among eligible applications from countries with low rates of immigration to the United States. For the 2010 lottery, over 13.6 million qualified entries were received, with 102,800 applicants drawn as winners, under the assumption that half of these would migrate.³ Whilst this is the most well-known example of a migration lottery, it has not been used for research yet.

² For simplicity we consider only the case of an individual-level outcome here, assuming that all migrants and non-migrants are observed. When the comparison involves households with and without migrants, a second form of selectivity is involved, since households can also choose whether all members migrate, or only some. Return migration also introduces a third form of selectivity. See Gibson et al. (2010) for discussion of this more complicated case.

³ Source: http://travel.state.gov/visa/immigrants/types/types_4574.html [accessed February 26, 2010].

Researchers have exploited smaller lottery programs. The first such set of studies considered the Pacific Access Category program for Tonga, which provides an opportunity for 250 individuals each year to move to New Zealand, with a random ballot used to select among all eligible applications received. McKenzie et al. (2010) collaborated with the New Zealand Department of Labour to draw a sample of individuals who had their names selected in this ballot, as well as a sample of those who applied but whose names were not drawn. They then surveyed the ballot winners in New Zealand, their remaining family members in Tonga, and the ballot losers in Tonga. As with the U.S. diversity visa, not all those whose names were chosen in the lottery migrated, so the authors also had to survey in Tonga the ballot winners who did not migrate.

If everyone who applied for the migration lottery ended up migrating, and no one who lost the lottery migrated, then linear regression of equation 1 (with or without the X controls) for the sample of lottery applicants, would give a consistent estimate of the causal impact of migration for people who enter the migration lottery program. However, in practice, some of those who win the lottery may not move (they may change their mind, or fail an entry requirement), whilst a few of those who lose the lottery may find other ways of migrating. In such cases, the outcome of the lottery can be used as an instrumental variable for migration, and still be used to identify the impact of migration on the outcome of interest. If the impact of migration varies across individuals (is heterogeneous), then what will be identified is the local average treatment effect (LATE). Recently there has been debate as to whether the LATE is a parameter of interest in many experiments (e.g. Deaton, 2010 and Imbens, 2010). However, in the case of a migration lottery, it is easy to argue that the parameter is giving an effect of policy interest. The LATE tells us the impact of migrating for someone who would migrate if they won the lottery, and not migrate otherwise. This is precisely what we would be interested in when assessing the impacts of such policies on development outcomes.

McKenzie et al. (2010) use this migration lottery to estimate that Tongans moving to New Zealand have a 263 percent increase in income, within the first year of moving. The migrants also benefit in terms of improved mental health (Stillman et al, 2009). The authors then use a sample of non-applicants to the lottery and a large sample of the overall population and compare the experimental estimate to what one would obtain using non-experimental methods. They find that linear regression would overstate the income gain to migration by 27 to 35 percent, which is consistent with migrants being positively selected on unobservables (the authors find positive selection on observables). Using non-experimental methods like difference-in-differences or propensity score matching reduces this overstatement a little, but still results in an overstatement of around 20 percent in the income gain from migration. The only non-experimental method that gets close to the experimental estimate is a good instrumental variable.

Migration policies typically limit which other family members can migrate along with the principal migrant, often restricting this just to the spouse and dependent children of the migrant. This policy rule can be used alongside a migration lottery to deal with a second form of selectivity – selectivity into which household members move and which remain in the home country. Gibson et al. (2009) use the combination of the policy rule and the lottery to look at the impact of migration on remaining household members in Tonga, finding evidence that remaining household members appear worse off in the short-term, with lower per capita incomes and consumption.

The above studies relied on the use of administrative data to track winners and losers in the Pacific Access Category. It requires considerable effort and enlightened policymakers in order for researchers to obtain access to such data. An alternative approach would be to attempt to locate households with lottery winners and losers in a survey in the migrant-sending country. Gibson et al. (2010) provide an example of this approach, using a representative survey of Samoan households to identify households which entered the Samoan Quota lottery, which allows 1,100 Samoans to migrate to New Zealand each year. Since Samoa's population is small, the lottery had been in place for several years, and there were 5,000-7,000 applications for the lottery each year, a random sample of households was able to identify sufficient numbers of lottery winners and losers for experimental analysis of the impact of migration through this category. In contrast to the Tongan results, the Samoan experiment finds migration to have reduced poverty and increased household incomes among remaining household members. However, there is suggestive evidence that this positive effect may be short-lived, with remittances and home production falling with time spent abroad of migrant members.

Clemens (2010) provides a final example of the use of a lottery provided by migration policy. He studies the H1-B visa, which is an admission channel for high skilled workers who wish to work in the United States. There is a cap on the number of people who can annually enter through this category, and while applications are processed on a first-come, first-served basis, in 2007 and 2008 so many applications were received on the opening days, that a lottery was used to select which applications to process. Rather than trying to get administrative data from the United States Citizenship and Immigration Services, which would be incredibly difficult, Clemens obtained personnel records from a large Indian information technology/software firm that supplied a large number of applicants to this lottery. He is then able to use these records to determine the impact of migrating on the migration, job title, and earnings of the applicants in this Indian firm.

An underlying assumption of the experimental estimates is that the migration lottery influences the outcome of interest for an individual only through that individual's migration decision. In particular, we require the Stable Unit Treatment Value Assumption (Rubin, 1986), which means that the outcome of one individual should not be affected by the lottery outcome of another individual. Thus when we consider the case of measuring the income gain from migration, we would require that the income of lottery losers is not affected by whether other people in the sample win or lose the lottery. One potential way this assumption could be violated would be if the lottery winners send remittances to the lottery loser households. This is unlikely to be much of an issue in cases where the number of winners is small relative to the overall population, and can be directly checked through surveys. A more problematic concern would be if the employment prospects of the lottery losers change as a result of the winners migrating. Whilst there is some evidence to suggest that large-scale migration can increase the wages earned by non-migrants through less competition for jobs (Mishra, 2007), this is again likely to be at most a second-order concern when the number of winners is small relative to the overall population. But it is more of a concern in cases like that studied by Clemens (2010), where the entrants studied are all from the same company, and it is therefore harder to imagine that the job opportunities available to the non-migrants are not affected by having some company members abroad.

In addition to the lotteries discussed, the U.S. has had a lottery for Cubans in the mid 1990s⁴, New Zealand's Pacific Access Category also has small lotteries for Kiribati and Tuvalu, and New Zealand also used a lottery to allocate places in its Family Quota and Refugee Family Quota Categories in the early 2000s.⁵ Currently only a few studies to date have utilized migration policy experiments. However, given the massive excess demand for migration into many countries worldwide, a lottery system for choosing which applications to process provides one fair and equitable mechanism for countries to process such applications, and we see this as a rich area for both researchers and policymakers to work on in the future.

3.2. Natural Experiments

In addition to exogenous variation generated by government policy, identification of causal effects in migration research can also take advantage of other sources of exogenous variation in right-hand-side variables of interest, or so-called *natural experiments*. We review here recent studies that take advantage of two types of natural experiments: exchange rate shocks experienced by migrants, and weather shocks to which migrants' origin households are exposed in the home country.

Impacts of changes in migrants' economic conditions

A question of general interest in migration studies is "What is the impact of changes in a migrant's economic conditions on outcomes in their origin household?" Such questions help reveal the extent to which changes in migration host countries (such as economic conditions, exchange rates, job opportunities, restrictions on legal work, etc.) affect a migrant's willingness and ability to send resources home, as well as the ways in which migrant resources are used by recipient households. Impacts can be mediated by several channels, such as remittances sent home, the stock of savings held by migrants overseas, or return migration decisions.

A central difficulty in answering this type of question is that migrant earnings or migrant economic conditions more generally are in general not randomly allocated, so that any observed relationship between migrant economic conditions and household outcomes may simply reflect the influence of unobserved third factors. For example, more ambitious households could have migrants who work for higher wages or in destinations with more attractive work opportunities, and also have higher entrepreneurial investment levels in the origin household. Alternately, households that recently experienced an adverse shock to existing investments (say, the failure of a small business) might send members overseas to make up lost income, and when migration decisions are made under duress migrants may accept going to less-attractive destinations. In sum, simply observing a statistical correlation between migrant economic conditions overseas and outcomes in migrant origin households does not imply that migrant economic conditions cause the origin household outcomes in question.

An experimental approach to establishing the impact of migrant economic opportunities on household outcomes could start by identifying a set of households that already had one or

⁴ <http://havana.usint.gov/media/pdfs/lottery.pdf> [accessed March 1, 2010].

⁵ These categories have now been replaced, but New Zealand still uses a random ballot to fill residual places in its Refugee Family Support Category which provides a means for refugees to sponsor parents, adult siblings or grandparents into New Zealand.

more members working overseas, assigning each migrant a randomly-sized economic shock, and then examining the relationship between changes in household outcomes and the size of the shock dealt to the household's migrants.

Yang and Martinez (2005) and Yang (2008b) take advantage of a real-world natural experiment that is analogous to the experiment just described. Many households in the Philippines have one or more members working overseas. These overseas Filipinos work in dozens of foreign countries, many of which experienced sudden changes in exchange rates due to the 1997 Asian financial crisis. Crucially for the analysis, the changes were unexpected and varied in magnitude across overseas Filipinos' locations. The net result was large variation in the size of the exchange rate shock experienced by migrants across source households. Between the year ending July 1997 and the year ending October 1998, the US dollar and currencies in the main Middle Eastern destinations of Filipino workers rose 50% in value against the Philippine peso. Over the same time period, by contrast, the currencies of Taiwan, Singapore, and Japan rose by only 26%, 29%, and 32%, while those of Malaysia and Korea actually fell slightly (by 1% and 4%, respectively) against the peso.

Taking advantage of this variation in the size of migrant exchange rate shocks, these papers examine the impact of the shocks on changes in outcomes in migrants' origin households, using detailed panel household survey data from before and after the Asian financial crisis.

Yang (2008b) shows that these exogenous increases in migrant resources are used primarily for investment in origin households, rather than for current consumption. Households experiencing more favorable exchange rate shocks see greater increases in child schooling and reductions child labor (for children aged 10-17). They also raise their non-consumption expenditures in several areas likely to be investment-related (in particular in educational expenditures), and show enhanced entrepreneurship participation in entrepreneurial activities. Households raise hours worked in self-employment, and become more likely to start relatively capital-intensive household enterprises (transportation/communication services and manufacturing). By contrast, there is no large or statistically significant effect of the exchange rate shocks on current household consumption. Yang and Martinez (2005) extends the analysis and shows that these positive migrant exchange rate shocks also lead these households to be more likely to exit poverty status.

Aside from impacts on migrant origin households, it is also of interest to examine how migrant return decisions changed in response to the exchange rate shocks accompanying the Asian financial crisis. In research on migration decision-making, a current debate is whether durations of migrants' stays overseas are determined primarily by straightforward life-cycle considerations, as opposed to being driven by the need to reach target-earnings levels. By "life-cycle" considerations, one means simply that households choose the length of stay overseas that balances the marginal benefit from higher savings overseas (and thus higher lifetime consumption) against the marginal utility cost of overseas work (as in Stark, Helmenstein, and Yegorov (1997) and Dustmann (2003)). On the other hand, when households face borrowing constraints and minimum investment levels, lengths of stay overseas can be determined by the amount of time needed to accumulate a "target-earnings" level, as in Piore (1979) and Mesnard (2004).

Distinguishing between the two alternative motivations for return migration is important, because the return decisions of "life-cycle" migrants and "target-earners" can respond very differently to changes in overseas economic conditions. For "life-cycle" migrants, improved

economic conditions in host countries—say, increased wages—can lead to longer overseas stays (as long as substitution effects dominate any income effects). For “target-earners,” on the other hand, improved economic conditions should lead to shorter overseas stays, as migrants reach their earnings goals more quickly.

Empirically, attempts to distinguish between the two alternatives typically examine the correlation between return migration and migrants’ overseas earnings. The evidence has been inconclusive. Borjas (1989) finds among the foreign-born in the US that higher earnings are associated with less return migration. By contrast, Dustmann (2003) documents, among immigrants in Germany, that higher migrant wages (instrumented by parental education) are associated with more return migration (shorter overseas stays). Constant and Massey (2002) find no statistically significant relationship between earnings and migrant returns in the same German dataset, although migrants who are unemployed or marginally employed are more likely to return.

A key methodological concern with existing empirical work on this topic is that the independent variable of interest—foreign earnings—is not randomly assigned across migrants, so any observed relationship between foreign earnings and return migration may simply be caused by unobserved third factors. For example, a finding that migrants with higher earnings have shorter lengths of stay overseas need not imply that higher earnings cause shorter migration durations. Rather, higher-wage migrants could simply have other characteristics that make early return attractive (such as better job prospects at home, or stronger family ties).

Yang (2006) exploits the exchange rate shocks experienced by Filipino overseas migrants, making possible a causal estimate of the effect of migrant economic conditions on return migration. In so doing, it also sheds light on the relative importance of life-cycle versus target-earnings explanations for return migration. Overall, the paper finds that more favorable exchange rate shocks lead to fewer migrant returns, which supports the “life-cycle” explanation for return migration. A positive exchange rate shock raises the marginal benefit of staying overseas (by raising the domestic-currency value of foreign wages), and leads to less return migration on the margin. However, the paper also finds that even though life-cycle considerations seem to dominate on the whole, migrants from a subset of households appear to be target-earners. In households with intermediate values of the foreign wage index, the exchange rate shocks lead to increases in variables associated with household investment, such as vehicle or real estate purchases and entrepreneurial income. These results are consistent with the theoretical prediction that the migrants most likely to be target-earners are those in the middle of the foreign wage distribution: positive exchange rate shocks make target-earners more likely to return home and to invest (because they become more likely to have reached the minimum investment threshold).

The general methodology used in these studies on Philippine migrants – examining the impact of an economic shock experienced by overseas migrants on remittances and the outcomes of family members left behind – can potentially be applied in a variety of different contexts. Studies using a similar methodology can be useful to ascertain whether the results in the Philippine case extend to other contexts, or, if not, what might account for the differences in impacts.

The key requirements for such a study are: 1) an origin country whose migrants are destined to a wide variety of overseas destinations, 2) large and heterogeneous economic shocks in destination areas, 3) data on migrant locations before the shocks, and 4) data on migrant and

origin household outcomes after the shocks. Many situations satisfy elements 1) and 2): for example, migrants from India and the other countries of South Asia also are destined for a wide variety of overseas destinations, and regional or global country-level economic shocks (such as the 2008-09 global financial crisis) are often heterogeneous in magnitude across migrant destinations. Migrants from specific countries in Latin America are often destined for a variety of locations across the United States, and hence it can be possible to exploit state-level (and perhaps occupation- or industry-specific) economic shocks experienced by migrants to achieve identification. Antman (2010) provides one example of this approach.

A likely hindrance to future research along these lines among other migrant populations is that there are fewer situations where the requisite survey data (requirements 3 and 4) are available. The Philippine case is unusual, in that the National Statistics Office of the Philippines administers a linked set of high-quality surveys to a nationally-representative household sample that includes a detailed module on migration which is administered if the household reports having one or more members overseas. Importantly, the migration module (called the Survey on Overseas Filipinos) includes questions on migration history that allows a researcher to track migration episodes up to 5 years in the past. Such a module turns out to be crucial for identifying households that had migrants in specific shock-exposed locations *prior to* the shock, because location after the shock could be endogenous and therefore introduce bias in estimation. A key implication for surveys of migrant sending households in countries where migrants tend to only go to a one or a few main destinations is to collect information not only on which country the migrant is in, but also the city or region to provide more scope for identifying local shocks.

Remittance responses to conditions in migrant origin areas

An important potential benefit from international migration is that remittances may serve as insurance, rising in the wake of negative shocks in migrants' home countries. Rural households in many developing countries are highly exposed to weather risk, experiencing storms, flooding, and droughts with great frequency. Households therefore should benefit greatly from access to formal and informal insurance that alleviates their most important sources of weather risk. Potential benefits include the ability to maintain nutritional, health, and educational investments, to adopt new production technologies, and to start new entrepreneurial activities that weather risk made previously unattractive. A large literature has examined the mechanisms through which households cope with risk in developing countries, but until recently the insurance role of remittances has not been investigated.

Yang and Choi (2007) and Yang (2008a) explore whether migrant remittances serve as insurance in the wake of negative weather shocks. This is a mechanism for coping with shocks *ex post* on which previous micro-level studies have not focused. At the international level, it is commonly posited that remittance flows from overseas buffer economic shocks in the migrants' home countries (for example, Ratha 2003), but there have been relatively few empirical tests of this claim with micro-level household data. Mishra (2005) examines remittances in 13 Caribbean countries from 1980 to 2002 and finds that every 1 percent decrease in GDP is associated with a 3 percent increase in remittances two years later. Related research on the role of internal (domestic) migration in pooling risk within extended families includes Lucas and Stark (1985), Rosenzweig and Stark (1989), and Paulson (2003).

Yang and Choi (2007) and Yang (2008a) emphasize credible identification of the effect of negative shocks on international remittances. Existing studies of the impact of household

income on remittance receipts use cross-sectional data, and so are subject to potentially severe biases in directions that are not obvious a priori. Reverse causation is a major concern: productive investments funded by migrant remittances can raise household income, leading to positive correlations between household income and remittances. Alternately, remittances may reduce households' need to find alternative income sources, leading to a negative relationship between remittances and domestic-source income. Even if reverse causation from remittances to income in migrants' source households was not a problem, it would be difficult to separate the cross-sectional relationship between income and remittances from the influence of unobserved third factors affecting both income and remittances (for example, the entrepreneurial spirit of household members).

Yang and Choi (2007) resolve these identification problems by focusing on income changes for the migrant-sending family due to shocks—changes in local rainfall—that are credibly exogenous, so that bias due to reverse causation is not a concern. Among households in the Philippines with members who are overseas migrants, they find that changes in income from domestic sources lead to changes in remittances in the opposite direction of the income change: remittances fall when income rises, and remittances rise when income falls. In such households, the amount of insurance is large: roughly 60% of exogenous declines in income are replaced by remittance inflows from overseas. In a similar vein, Clarke and Wallsten (2004) find, using panel data from Jamaica, that remittances from overseas replaced 25% of damages from Hurricane Gilbert in 1992. Yang (2008a) examines the impact of hurricanes on international financial flows using country-level panel data and finds that, for the poorest developing countries, hurricane damage leads to large inflows of migrants' remittances, amounting to 20% of experienced damages. Strikingly, the remittance response to hurricanes for these countries is large in magnitude: roughly one-quarter as large as the response of foreign aid.

General thoughts on natural experiments

As these examples show, natural experiments offer the potential to provide a credible means of helping answer many important questions in migration studies. However, there are often clear limits to the types of questions such natural experiments can answer (see also Rosenzweig and Wolpin, 2000) and one needs to be cautious in interpreting the results. For example, work such as that by Yang (2008b) directly looks at the impact of the exchange rate shocks on outcomes for migrant sending households. This does not reveal the average impact of remittances or of migration, but rather the response of these sending households to temporary shocks in the earnings their migrants earn abroad. A finding that households save or invest much of the additional amount remitted as a result should therefore not be used to infer that remittances in general are used for largely productive purposes – economic theory tells us we should expect households to save or invest more in response to temporary income shocks and temporary increases in remittances caused by exchange rate fluctuations than they would from regular income or their usual level of remittances.

This is not a concern for the Yang (2008b) study in which the question of interest is indeed the response of migrant households to temporary shocks in the conditions facing their migrants abroad. But it is a more of a concern for studies which attempt to use these natural experiments to generate an instrumental variable for migration or remittances. When these natural experiments are used to instrument migration, the impact identified is the local average treatment effect for households affected by the instrument. For example, Antman (2010)

instruments father's migration with economic conditions in the main U.S. destinations in a study of the impact of paternal migration on child schooling. Assuming the instrument meets the other criteria required, the effect identified is only the impact of parental migration for children whose father's migration decisions are affected by temporary economic shocks in the U.S. destinations. Given the large income differences between Mexico and the U.S., the set of households likely to have their migration decisions change as a result of these temporary shocks may be small, and thus the impact identified not be one that applies to much of the overall migrant population.

An important warning about natural experiments is that they must be scrutinized carefully before they are used as instruments in instrumental variables (IV) estimation. When an exogenous source of variation in economic conditions is identified (e.g., weather, exchange rates), it is often tempting to take the next step and use the shock as an instrumental variable. The concern is with the validity of the IV exclusion restriction, namely the requirement that the instrument only affect the 2nd-stage variable of interest via the endogenous right-hand-side variable of interest (which is being instrumented).⁶ It is actually quite rare for exogenous shocks to satisfy the exclusion restriction, because there are usually a number of different channels through which the shock can affect the 2nd-stage outcome of interest. When this is the case, instrumenting for just one of several channels with the shock will generally lead to biased estimates. That said, it is generally acceptable to examine the "reduced form" effect of the shock (e.g., in a regression of the outcome of interest directly on the shock variables), and to interpret the effect of the shock as operating through multiple potential channels. For example, Yang (2008b) examines only the reduced form effect of the exchange rate shock on the dependent variables of interest precisely because the exchange rate shocks could operate through at least two channels: through remittances sent home as well through the Philippine-peso value of savings and other assets held overseas. In this case using the exchange rate shock as an instrument for remittances would have led to biased estimates. This bias would probably have been in an upward direction because any effects of the exchange rate operating via changes in the value of unremitted overseas savings would have been "loaded" onto the coefficient on remittances.

3.3. Field Experiments

While governments should use them more often, policy experiments are rare. Natural experiments, while valuable and revealing when they occur, are difficult to find. When policy or natural experiments do not exist, a large set of questions can be answered via randomized control trials or field experiments. Well-designed field experiments can help us understand not only the impact of a particular program or intervention, but can also shed light on underlying causal mechanisms or test particular theories. While field experiments have become increasingly common in development economics research, they have only just begun to be attempted in research on migration.

In this section we provide overviews of a handful of recent or ongoing field experiments on migration: studies of savings among migrants in the U.S. and research on barriers to migration (for internal migrations in Bangladesh and international migrants from the Philippines).

⁶ For a useful overview of this issue in instrumental variables estimation, see Angrist and Pischke (2009), section 4.1.

Studies of migrant savings

While remittances bring numerous benefits to households in developing countries, to date we know very little about how migrants make their remittance-sending decisions. In particular, it is unknown whether migrants desire greater control over how family members back home use the remittances they receive. This question is relevant not only for migration studies but also for the large and active literature in development economics on intra-household resource allocation. What's more, a better understanding of these questions could have substantial impact on public policy, by suggesting policies to further stimulate remittance flows and potentially channel them towards more productive uses in migrant source countries.

Ashraf, Aycinena, Martinez, and Yang (2010), henceforth AAMY, address some of these questions via a randomized controlled trial among migrants from El Salvador who are living and working in the Washington, D.C. metro area. The research aims to shed light on the extent to which migrants' lack of direct control over the use of remittances affects remittance flows, and on the impact of new financial products that could increase migrant control.

In particular, AAMY focus on improving the ability of migrants to ensure that remittances are deposited and accumulated in savings accounts in the home country. In survey data collected as part of the study, Washington, D.C.-based migrants from El Salvador report that they would like recipient households to save 21.2% of remittance receipts, while recipient households prefer to save only 2.6% of receipts. Migrants often intend the savings to be for the use of the recipient household in the future, but such savings also can be intended for the migrant's future use. In the latter case, migrants may send their own funds to be saved in El Salvador because they perceive savings held in the U.S. as relatively insecure (particularly for undocumented migrants who fear deportation and loss of their assets).

AAMY designed a field experiment that offered new facilities for Salvadoran migrants to directly channel some fraction of their remittances into savings accounts in El Salvador. Savings facilities were offered in conjunction with Banco Agricola, El Salvador's largest bank. To isolate the importance of migrant control over savings, AAMY test demand for different products that offer migrants varying levels of control. For example, they investigate differential demand for savings accounts that must be solely in the name of a remittance recipient in El Salvador, versus accounts that are either jointly owned with the migrant or for which the migrant is the sole owner.

The impact evaluation uses a randomized treatment-control methodology. Migrants in the study are randomly assigned across treatment conditions, and so comparisons across the various treatment conditions reveal the causal impact of offering migrant control on the outcomes of interest (which include savings account take-up, savings balances, and remittances). The intervention studied is unusual among development economics field experiments in that it is conducted among migrants who are located in a developed country, while several primary outcomes of interest (savings) are those of individuals who remain behind in a developing country. Data on activity at our partner bank are available from the bank's administrative records. Baseline and follow-up surveys administered to both migrants in the U.S. and their corresponding remittance-receiving households in El Salvador provide data on a broader set of other outcomes.

AAMY's results provide evidence that a desire for control over remittance uses—in particular, control over the extent to which remittances are saved in formal savings accounts—is

quantitatively large and has an important influence on financial decision making by migrants. Across the experimental conditions in the sample, migrants were much more likely to open savings accounts when offered the option of greater control over the accounts. What's more, offering greater migrant control over El Salvador-based savings accounts led to higher savings accumulation in El Salvador.

A related randomized experiment on savings among immigrants was conducted by Chin, Karkoviata, and Wilcox (2010). This study examines the impact of providing Mexican immigrants in the US with assistance obtaining a form of ID (a *matricula consular*) that can be used as identification when opening a US bank account. Study participants were made aware of a collaborating US bank that had an ongoing savings promotion among Hispanic immigrants, but the *matriculas consulares* in principle could have been used at any number of US banks. Impacts of the treatment were assessed in an in-person follow-up survey. Assignment to the treatment is found to lead to increased opening of US bank accounts, higher savings in the US, and reduced remittances to Mexico. Among migrants who report they have "no control" over how remittances are used in Mexico, the abovementioned effects are larger, and there is also a large, positive and statistically significant treatment effect on migrant earnings.

Taken together with AAMY, the Chin, Karkoviata and Wilcox (2010) study reinforces the conclusion that migrants have a variety of types of demand for savings facilities. There is demand for savings in the US, as well as demand for savings in the country of origin, and providing access to appropriate savings devices can have large impacts on savings. What's more, both studies underline the importance of migrant control over savings accounts in facilitating savings accumulation. We view such studies as just the tip of the proverbial iceberg. There is likely to be great potential for analogous future studies that partner with institutions to offer a variety of financial services to immigrants. Products that have yet to be investigated include credit, insurance, and direct payment facilities targeted towards the needs of migrants and their origin households.

Identifying barriers to migration

While international and internal migration flows are large in magnitude, even greater numbers of individuals do *not* migrate, even in the face of substantial wage differentials between less- and more-developed areas. There are likely to be large number of potential migrants who are deterred from migrating by a variety of barriers, such as imperfect information on migrant wages and job conditions, imperfect information on one's own affinity for or returns from migrant work, lack of information on job-seeking procedures, and credit constraints (when migration or job search involves non-negligible fixed costs).

At the moment we know little as researchers about the relative importance of these various potential barriers to migration. Credible evidence on the importance of migration barriers has important policy implications as well. A number of developing countries – most prominently, the Philippines – have enacted policies intended to facilitate and regulate international migration and view such policies as integral components of their overall economic development strategies. If there is a desire to promote migration, it is crucial to understand which barriers are operative and the impact of interventions that are aimed at reducing these barriers. While no study has been completed so far, randomized control trials in Bangladesh and the Philippines are currently underway and seek to shed light on the relative importance of several potential barriers to migration.

Bryan, Chowdhury, and Mobarak (2010) are currently analyzing the results of an ongoing randomized field experiment in the northwestern region of Rangpur in Bangladesh. A relatively impoverished area, Rangpur experiences annual famines that lead to seasonal declines in household income and consumption. A key coping strategy for households in the face of the famine is internal labor migration to other parts of Bangladesh that are less- or unaffected by the famine. The experiment involves 100 Rangpur villages that were randomly allocated to the following experimental conditions: a control group; a treatment group offered information on jobs available, typical wages, and the likelihood of finding migrant work in a set of migration destinations; and other treatment conditions that offered cash or credit to cover the initial fixed costs of migration. The experiment was implemented in 2008 and the endline survey of households in the 100 migrant-origin villages were implemented in 2009. Preliminary results are revealing, suggesting that the information treatments had no effect but the cash and credit interventions had substantial effects on migration both in the year they were offered as well as in the next year's famine season (when the cash/credit were no longer offered by the research project). Treatments that had effects on migration also led to substantial increases in consumption in migrant households. Should they hold up, the results provide evidence of the importance of credit constraints as a migration barrier, and also – intriguingly – suggest that policies providing a small incentive to migrate in an initial period can have persistent effects in future periods even after the incentives are removed.

A field experiment seeking to shed light on barriers to international migration is being implemented by Beam, McKenzie, and Yang (2010) in Sorsogon province, the Philippines. International labor migration from the Philippines is very large in magnitude: the Philippine Overseas Employment Agency, with which all overseas labor contracts of Filipinos must be officially registered, has recorded over one million new contracts per year since 2006. Roughly one-third of this number are “new hires” or first deployments overseas, and the remaining two-thirds are “rehires” or new work contracts for workers who are already overseas or who have previously been overseas. Within the country, international labor migration rates are highest in areas closest to major cities like Manila (the capital) and Cebu. An open question is why individuals in some outlying provinces – such as Sorsogon, which is more than 12 hours by bus from the capital – typically have substantially lower rates of international labor migration despite facing larger income gaps between home and abroad.

The baseline survey and intervention for the Sorsogon experiment was completed in mid-2010, with roughly 5,000 households in the sample. Randomization was at the household level. The experimental conditions were as follows: a control group; a group offered information on typical wages in common overseas work destinations, on procedures for applying for overseas work, and on the typical fixed costs involved in overseas labor migration; and a group offered assistance in applying for overseas work (in addition to the information offered to the previous treatment group). Households in the “assistance” treatment were offered access to and assistance with uploading information on themselves into a job-seekers’ website that recruitment agencies in Manila could then use to search for suitable candidates to fill overseas job openings. These interventions are intended to test the relative importance of various types of information and transactions costs in explaining the low incidence of international labor migration from outlying areas of the Philippines.

Among the sample of individuals are enrolled in the job-seekers’ website database, a follow-on randomization will be implemented, in collaboration with a local microfinance

institution, intended to shed light on the importance of credit constraints as a migration barrier. Prospective migrants in the database will be randomly allocated into the following groups: a group offered a small loan to cover the costs of job search (mainly costs of travel to and lodging in Manila to attend job interviews); a group offered a larger loan to cover costs of travel overseas once a job offer is obtained; a group offered both types of loans; and a control group offered no loan products. Estimated effects of these loan treatments should reveal the relative importance of credit constraints at two different stages of the international labor migration process (initial job search vs. overseas travel once a job offer is obtained).

4. Conclusions and Lessons for Future Work

The decision of whether or not to migrate has far-reaching consequences for the lives of individuals and their families. But the very nature of this choice makes identifying the impacts of migration difficult, since it is hard to measure a credible counterfactual of what the person and their household would have been doing had migration not occurred. Migration experiments provide a clear and credible way for identifying this counterfactual, and thereby allowing causal estimation of the impacts of migration. Yet to date there have been relatively few such experiments, and we believe there are large gains from policymakers and researchers using experiments more frequently.

On the policy side, Governments could use experiments more as a way of learning about the effectiveness of their policy initiatives. We have seen examples where migration lotteries have been used as a fair and equitable way of deciding among excess demand for quota-constrained immigration categories. A second natural place for such experiments is in piloting the introduction of new policies. For example, seasonal worker programs (such as those recently introduced by New Zealand and Australia) are seen as a way to enable less-educated poor households to reap some of the gains possible with international migration. Yet there is debate about the optimal way to select such workers, and about the extent to which it will actually benefit sending communities. Randomly selecting among eligible workers in the pilot phase, and experimenting with different recruitment mechanisms would provide a way for the design of such policies to be fine-tuned, alleviate potential concerns about political favoritism determining which individuals and villages participate, and provide a means for the development impacts to be identified. On the sending country side, several developing countries have shown interest in providing pre-migration orientation seminars for potential migrants, in a similar way to the Philippines. But there is little evidence as to the effectiveness of such programs, or as to which components really matter. Before introducing such programs on a large scale, Governments could experiment with offering different content to different groups of migrants, and measure which is most effective.

One argument Governments might muster against randomly choosing among applicants for a given migration quota is a belief that they can get higher quality migrants by intensive screening of all applications. This argument is likely to be more important for policies to admit skilled migrants. Points systems for migration provide one such approach to screening, in which prospective migrants are scored on a basis of marketable skills and desirable characteristics, with only those individuals scoring above a certain threshold eligible to migrate. For example, Australia's points system in scores applicants out of 170 based on their age, English skills, occupation, work experience, Australian qualifications, and other characteristics. The system had

pass marks of 100 and 120 for different visa categories in 2010.⁷ Whilst one can imagine the Government being reasonably confident that someone with a score of 170 is a more desirable immigrant than someone with a score of 60, there is likely to be much more uncertainty about whether people getting a score of 95 are all that different from those getting a score of 105. Governments could therefore consider randomizing among excess applications within some range around the pass mark, in order to learn more about the impacts of such policies. This would be analogous to the approach used in some microfinance experiments, which have randomized access to credit for marginal applicants (Karlan and Zinman, 2010).

References

Angrist, Joshua and Jorn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton: Princeton University Press, 2009.

Antman, Francisca (2010) "The Intergenerational Effects of Paternal Migration on Schooling and Work: What Can We Learn from Children's Time Allocations?", Mimeo. University of Colorado, Boulder.

Beam, Emily, David McKenzie, and Dean Yang, "Financial and Informational Barriers to Migration: A Field Experiment in the Philippines," ongoing study, University of Michigan and World Bank, 2010.

Borjas, George (1987) "Self-selection and the earnings of migrants", *American Economic Review* 77(4): 531-53.

Borjas, George, "Immigrant and Emigrant Earnings: A Longitudinal Study," *Economic Inquiry*, 27(1), 1989, pp. 21-37.

Bryan, Gharad, Shyamal Chowdhury, and A. Mushfiq Mobarak, "The Effect of Seasonal Migration on Households during Food Shortages in Bangladesh," ongoing study, Yale University, 2010.

Chin, Aimee, Leonie Karkoviata, and Nathaniel Wilcox, "Impact of Bank Accounts on Migrant Savings and Remittances: Evidence from a Field Experiment," mimeo, University of Houston, 2010.

Clarke, George and Scott Wallsten, "Do Remittances Act Like Insurance? Evidence from a Natural Disaster in Jamaica," working paper, World Bank, 2004.

Clemens, Michael (2010) "How Visas Affect Skilled Labor: A Randomized Natural Experiment", Mimeo. Center for Global Development.

⁷ See <http://www.immi.gov.au/skilled/general-skilled-migration/points-test.htm>

Constant, Amelie, and Douglas Massey, "Return Migration by German Guestworkers: Neoclassical Versus New Economic Theories," *International Migration*, 40 (4), 2002, pp. 5-38.

Deaton, Angus (2010), "Instruments, Randomization, and Learning about Development", *Journal of Economic Literature* 48(2): 424-55.

Dustmann, Christian, "Return Migration, Wage Differentials, and the Optimal Migration Duration," *European Economic Review*, 47, 2003, pp. 353-369.

Gibson, John, David McKenzie and Steven Stillman (2009) "The impacts of international migration on remaining household members : omnibus results from a migration lottery program", *Review of Economics and Statistics*, forthcoming.

Gibson, John, David McKenzie and Steven Stillman (2010) "Accounting for Selectivity and Duration-Dependent Heterogeneity When Estimating the Impact of Emigration on Incomes and Poverty in Sending Areas", Paper presented at World Bank/AFD conference September 2009.

Imbens, Guido (2010) "Better LATE than nothing: Some comments on Deaton (2009) and Heckman and Urzua (2009)", *Journal of Economic Literature* 48(2): 399-423.

Karlan, Dean and Jonathan Zinman (2010) "Expanding Microenterprise Credit Access: Using Randomized Supply Decisions to Estimate the Impacts in Manila", Mimeo. Yale University.

Lucas, Robert E. B. and Oded Stark, "Motivations to Remit: Evidence from Botswana," *Journal of Political Economy*, 93(5), 1985, pp. 901-918.

McKenzie, David (2005) "Beyond Remittances: The effects of Migration on Mexican households", pp123-148 in C. Özden and M. Schiff (eds) *International Migration, Remittances and the Brain Drain*, The World Bank: Washington D.C. 2005

McKenzie, David, John Gibson, and Steven Stillman (2010) "How Important is Selection? Experimental Vs Non-experimental Measures of the Income Gains from Migration" *Journal of the European Economic Association*, forthcoming.

McKenzie, David and Hillel Rapoport (2010) "Self-selection patterns in Mexico-U.S. migration: the role of migrant networks", *Review of Economics and Statistics*, forthcoming.

Mishra, Prachi (2007) "Emigration and wages in source countries: Evidence from Mexico", *Journal of Development Economics* 82: 180-99.

Mishra, Prachi, "Macroeconomic Impact of Remittances in the Caribbean," working paper, International Monetary Fund, 2005.

Paulson, Anna, "Insurance Motives for Migration: Evidence from Thailand," Kellogg Graduate School of Management, Northwestern University, mimeo, 2003.

Ratha, Dilip, "Workers' Remittances: An Important and Stable Source of External Development Finance," in *Global Development Finance 2003: Striving for Stability in Development Finance*. Washington, DC: International Monetary Fund, 2003.

Rosenzweig, Mark and Oded Stark, "Consumption Smoothing, Migration, and Marriage: Evidence from Rural India," *Journal of Political Economy*, 97(4), 1989, pp. 905-926.

Rosenzweig, Mark and Kenneth Wolpin (2000) "Natural "Natural Experiments" in Economics", *Journal of Economic Literature* 38: 827-74.

Rubin, Donald (1986) "Which ifs have causal answers?", *Journal of the American Statistical Association* 81: 961-62.

Stark, Oded, Christian Helmenstein, and Yury Yegorov, "Migrants' Savings, Purchasing Power Parity, and the Optimal Duration of Migration," *International Tax and Public Finance*, 4, 1997, pp. 307-324.

Stillman, Steven, David McKenzie and John Gibson (2009) "Migration and mental health: Evidence from a natural experiment", *Journal of Health Economics* 28(3): 677-87.

Yang, Dean and Claudia Martinez A., "Remittances and Poverty in Migrants' Home Areas: Evidence from the Philippines" in Caglar Ozden and Maurice Schiff, eds., *International Migration, Remittances, and the Brain Drain*, World Bank, 2005.

Yang, Dean, "Why Do Migrants Return to Poor Countries? Evidence from Philippine Migrants' Responses to Exchange Rate Shocks," *Review of Economics and Statistics*, 88(4), 2006, pp. 715-735.

Yang, Dean and HwaJung Choi, "Are Remittances Insurance? Evidence from Rainfall Shocks in the Philippines," *World Bank Economic Review*, 21(2), 2007, pp. 219-248.

Yang, Dean, "Coping with Disaster: The Impact of Hurricanes on International Financial Flows, 1970-2002," *B.E. Journal of Economic Analysis and Policy*, 8(1), (Advances), Article 13, 2008a.

Yang, Dean, "International Migration, Remittances, and Household Investment: Evidence from Philippine Migrants' Exchange Rate Shocks," *Economic Journal*, 118, 2008b, pp. 591-630.