How complete are labor markets in sub-Saharan Africa? Evidence from panel data in four countries*

Work in progress. Please do not cite or circulate. Comments welcome.

Brian Dillon† Peter Brummund‡
Germano Mwabu§

September 22, 2016

Abstract

We develop and implement new tests for the completeness of rural labor markets. Our approach builds on the seminal work of Benjamin (1992) and requires panel data with changes in household labor endowments over time. We show how the possibility of asymmetric responses to increases and decreases in labor endowments leads to tests of necessary conditions for a shortage or surplus in the labor market. The tests are based on strictly exogenous changes in the household labor endowment and allow for variation in labor market conditions across phases of the cultivation cycle. We implement the test using nationally representative panel data from Ethiopia, Malawi, Tanzania, and Uganda. Results differ across countries. In Ethiopia there is strong evidence of a labor shortage. In Malawi, the evidence is strongly in favor of a labor surplus. For Tanzania the results suggest a labor surplus at all times other than planting, when the pattern is less clear. In Uganda there is no evidence of non-separation in the panel, and therefore no labor market shortcoming. The paper shows that labor markets are not efficiently allocating labor in the study economies, with the possible exception of Uganda. But the underlying labor market imperfection varies across settings, highlighting the importance of tailoring policies to the circumstances of each country.

*We thank the Institute for the Study of Labor (IZA) and the UK Department for International Development for funding through the GLM-LIC program. Initial work on this project took place as part of the Myths and Facts in African Agriculture project of the World Bank and African Development Bank. For comments and helpful discussions we are grateful to Chris Barrett, Dave Donaldson, Andrew Foster, Rachel Heath, Tavneet Suri, and seminar participants at the Pacific Development Conference. We thank Joshua Merfeld, Reuben Mutegi, and Audrey Lawrence for excellent research assistance. All errors are the responsibility of the authors.

†University of Washington, Evans School of Public Affairs. Email: bdillon2@uw.edu.
‡University of Alabama. Email: pbrummund@cba.ua.edu
§University of Nairobi, Department of Economics. Email: gmwabu@gmail.com.

1 Introduction

The performance of labor markets has been central to the study of economic development since Lewis (1954) and Harris and Todaro (1970). Labor is a primary endowment of the world's poor households. As such, reductions in poverty are often driven by increases in the returns to the employment of labor. Labor is also a critical input to the agriculture sector, through which it contributes to the majority of production in rural areas of low income countries. Without an adequate understanding of the process that determines the returns to labor in rural areas, it is difficult to identify the appropriate set of policies to increase agricultural productivity and reduce poverty (Teal, 2011). Thus, the study of how rural labor markets operate should be a part of any poverty reduction agenda.

The goal of this paper is to shed light on the functioning of rural labor markets in four countries of sub-Saharan Africa: Ethiopia, Malawi, Tanzania, and Uganda. To do this, we develop new and intuitive tests for non-clearing labor markets. We then apply the tests to recently collected, nationally representative panel data sets. Our approach uses migration into and out of rural households to make inference about the labor market. However, we do not study migration in response to wage differentials or production shocks, in line with a large literature.[1]. Instead, the analysis here is rooted in the agricultural household model and the testable hypotheses that it implies.

The key antecedent to this paper is the seminal work of Benjamin (1992). Benjamin shows that with complete and competitive markets, farming households can fully separate the production and consumption sides of their utility maximization problem (the *separation hypothesis*). This leads to a testable prediction: if markets are complete and competitive, the household endowment

[1]See for example Gollin, Lagakos and Waugh (2014); Harris and Todaro (1970); Lewis (1954); Morten (2015).

of labor should not affect input demand on the farm. In a sample from Indonesia, Benjamin (1992) shows that labor demand on farm does not depend on the number of workers in the household. He therefore cannot reject the null hypothesis that rural markets – including the labor market – are complete and competitive.

This cross-sectional test has three key limitations. First, it does not control for time invariant household characteristics that could drive the result, such as a preference for working on one's own farm or unobserved managerial skill. Second, absent a compelling instrument for household labor endowment, it cannot account for possible endogenous adjustments to the household labor endowment in response to local labor market conditions. Third, it cannot distinguish between separation failures due to excess labor supply, to excess labor demand, or to failures in markets other than that for labor.

In this paper we expand and adapt the test in Benjamin (1992) to account for all three of these limitations. Our first extension is to bring the test to a dynamic setting using panel data. In country-level panels we observe changes over time in household labor endowments, as measured by the number of working age members. By including household fixed effects we control for unobserved household characteristics in regressions of on-farm labor demand on labor endowments and other variables. Identification is from the conditional correlation between changes in endowments and changes in labor demand. This aspect of our paper is similar to LaFave and Thomas (2016), who use panel data from the same setting as Benjamin to revisit his original hypothesis. LaFave and Thomas (2016) find that with the inclusion of household fixed effects, labor demand co-moves significantly with household labor endowments, overturning Benjamin's initial finding and indicating that rural markets in Indonesia are not complete and competitive.

Our second and more important innovation is to theoretically link asymmetries in the relationship between labor demand and labor endowments with underlying patterns of market failure. To understand these tests, we assume that changes in labor demand on farm are positively correlated with changes in the household labor endowment (which we will show to be the case for three of our study countries). It is also important to note that if the separation failure is because of non-clearing labor market, then that market must either be in excess or in shortage, but never in both simultaneously. We estimate regressions of the change in labor demand on the change in the labor endowment and prices, and we allow for the possibility of asymmetric responses to increases and decreases in labor endowments. We show how the interpretation of these regressions can, in some cases, reveal details about underlying labor market failures.

The intuition is as follows. Suppose that off-farm opportunities are limited so that there is excess supply in a non-clearing labor market, perhaps due to downwardly sticky wages (Kaur, 2014). Under very weak assumptions we show that households would not endogenously *increase* their labor endowments in response to this set of labor market conditions. Thus, any observed in-migration to the household must be exogenous to the local labor market. People come and go from rural households for many reasons – marriage, divorce, education, health shocks, loss of employment – and the exogenous changes in endowments are ideal for identification of the dynamic version of the separation hypothesis. Furthermore, because labor supply to the market is already in excess, this exogenous increase in the labor endowment leads to increased labor supply to the household farm, again under very weak assumptions. Thus, a necessary condition for a non-clearing labor market due to excess labor supply is that on-farm labor demand must increase when the labor endowment increases. If it also the case that labor demand on farm does *not* decrease when the household labor endowment decreases, this is strongly suggestive of a buyer's labor market. There are simply not enough jobs. While this is only a test of necessary, not sufficient conditions, the test relies on increases in endowments, which are exogenous under the case of excess labor.

Symmetrically, we show that if the local labor is characterized by a shortage, reductions in the household labor endowment must be exogenous to market conditions. Furthermore, such reductions lead to reductions in labor supply to the family farm, which cannot be offset with hired labor because of the prevailing market conditions. This leads to a test of a necessary condition for a labor shortage: on-farm labor demand must decrease when the household labor endowment decreases, but show little or no response to increases in the labor endowment. This pattern satisfies a necessary condition for a labor shortage.

These tests require some additional refinements that we discuss in the paper. Yet, they provide a clear framework for connecting violations of separation with underlying patterns of market failure. We implement these tests for all four study countries, separately. The data we use are from the recently collected LSMS-ISA data sets in Ethiopia, Malawi, Tanzania, and Uganda. These data sets are collected by the national statistics offices of the study countries, with technical support from the World Bank. They combine a large LSMS-style survey with a detailed set of agricultural modules and an emphasis on collecting nationally representative panel data.

In Ethiopia we find a strong positive correlation between changes in the amount of farm work and changes in the household labor endowment. However, when we allow for asymmetric responses to increases and decreases in the number of working age members, we find that this correlation is driven entirely by the decreases. When a worker moves out of the household there is a significant drop in the amount of labor-per-acre on the farm, but there is not a similar response when someone moves in. This is consistent with a seller's labor market – a shortage of workers. The pattern is the same in the harvest and pre-harvest periods, indicating that labor market conditions are similar across the phases of the agricultural cycle.

With the Malawi data we find the same initial pattern: changes in the household labor endowment are strongly correlated with changes in labor applied to the household farm. But when we split by increases and decreases we find the opposite result from Ethiopia. When a household in Malawi gains a working age member, that person is highly likely to work on the household farm. But decreases in the number of working age members do not impact the amount of labor applied to the farm. This is evidence of a buyer's labor market – too many workers for too few jobs. This pattern is stronger during the pre-harvest period than during the harvest period, but qualitatively the results are similar across the phases of the agricultural cycle.

In Tanzania the results are broadly similar to those from Malawi. During the weeding and harvest periods, the use of labor on the household farm changes significantly when the number of working age members increases, but not when it decreases. The same is true during the planting period, except there we also see a weakly statistically significant response to increases in the labor endowment, a finding that does not lend itself to any single interpretation. Overall, however, the results from Tanzania lean strongly toward a finding of excess labor and a lack of off-farm opportunities.

In Uganda, we find that changes in the number of working age members do not lead to statistically significant changes in the amount of work on the farm. This surprising result contradicts previous findings based on cross-sectional data. The implication is that the markets for labor and other agricultural inputs are working well in Uganda.

Another contribution of this paper is that we also provide a rich descriptive characterization of labor supply and demand in rural areas of the study countries. First, consistent with most observers' expectations, we document a clear spike in labor demand during both planting and harvest. However, the magnitude of the demand spike is far greater during planting than harvest, suggesting that if labor markets seize up it is most likely to be during the peak planting window (this finding is based on the Malawi data). Second, we show that average household labor supply to the household farm is much greater, and more variable throughout the year, than average supply of labor to the spot labor market (this result is based on the Malawi data). Third, in a surprising result, we show that older household members do not retire from working on the family farm. The labor force participation rate on farm is higher for 80-year-olds than it is for 20-year-olds in every survey wave. 70-year-olds work the same average number of days on farm as 40-year-olds. Fourth, we show that household labor endowments are constantly changing. The share of households experiencing a change in the number of working age members from one survey to the next ranges from 40-80%, depending on how we count children aging into the workforce. Fifth, we demonstrate that most villages have both households that shrink in size and households that grow in size, rendering it impossible to characterize the majority of villages as either source or destination villages for internal migrants. This is consistent with recent evidence in Young (2013). Finally, we show that migrants into and out of rural households tend to use their time similarly to existing household members. New arrivals do not work on the farm more or less than current members.

This paper advances the literature in a number of ways. As a field, we know surprisingly little about how labor markets function in sub-Saharan Africa, in part because of how difficult it is to develop sharp tests. Farming households are on many sides of the market. They both buy and sell labor, and may also barter labor with neighbors through systems of reciprocal exchange (Feder, 1985; Reardon, 1997; Barrett, Reardon and Webb, 2001; Teal, 2011). Additionally, labor supply to the market is partly determined by the shadow wage of labor on the family farm, which is difficult to observe (Benjamin, 1992; Jacoby, 1993; Skoufias, 1994). The seasonality and inherent riskiness of agricultural production add additional complexity. Early season shocks, or anticipation of late-season labor market bottlenecks, may dynamically alter the household's position in the labor market (Fafchamps, 1993; Kochar, 1999; Dillon, 2014). Finally, nominal wage rigidity due to hold-up problems, norms against wage reductions, or other sources of incomplete adjustment may reduce the effectiveness of the price-setting mechanism in allocating labor efficiently across households and sectors (Dreze and Mukherjee, 1989; Osmani, 1990; Kaur, 2014).

Our analytical approach combines many of the strengths of previous papers in this area. Like Benjamin (1992), our tests are based only on characteristics of the input demand equations that prevail under a null hypothesis of complete and competitive markets. In this way we avoid structural estimation and the large set of required assumptions. Also, we use panel data to control for unobserved household characteristics (LaFave and Thomas, 2016), and we theoretically link potential wage rigidities (non-clearing labor markets) with the household response to exogenous in- or out-migration (Kaur, 2014). In this sense our findings align with recent evidence of persistent gaps between rural and urban wages (Gollin, Lagakos and Waugh, 2014; McCullough, 2015), which are indicative of significant structural barriers to fully clearing labor markets.

What do our results suggest for policymakers? First, different approaches are called for in different countries. This may sound simplistic, because policies should always be tailored to local conditions. Here, however, we find results that appear to be similar across settings upon the initial analysis. There is a cross-sectional correlation between labor endowments and labor use on farm in all four countries (Dillon and Barrett, 2014) and a positive correlation in the panel for Ethiopia, Malawi, and Tanzania. Only when we allow for asymmetric adjustment do we see that the underlying labor market failure in Ethiopia is the opposite of that in Malawi and Tanzania.

For Ethiopia, where the results indicate a labor shortage, it would be beneficial to expand programs that match workers to farmers or that promote adoption of labor-saving technologies (herbicides, tractors). In Malawi and (for the most part) Tanzania there is evidence of a labor surplus. Possible policy responses to this situation are the expansion of jobs programs or government spending to increase aggregate demand.

Finally, we find little evidence that labor market conditions change from one agricultural phase to the next. The conditions that prevail during planting appear to persist through the end of the harvest, despite noticeable spikes in demand at certain times. This result is surprising, and contravenes both our expectations and assumptions that are commonplace in the literature.

The rest of the paper proceeds as follows. In Section 2 we develop the theoretical framework and associated empirical tests. Section 3 describes the data sets and sample, as well as a number of descriptive statistics that clarify the setting and rationalize some of the choices we make in the empirical implementation. In Section 4 we present our main results and a series of extensions and robustness checks. The paper concludes with a discussion in Section 5.

2 Theory

We begin this section with a brief description of the static agricultural household model and the test of complete and competitive markets that it implies. We then extend the model to a dynamic setting, explicitly allowing for endogenous adjustments to household labor endowment in response to local labor market conditions. The extended model implies three tests that together distinguish between market failures due to a surplus of labor, a shortage of labor, and a shortcoming in some other related market. This section ends with a description of the empirical specifications that we bring to the data.

2.1 The static agricultural household model

We first develop a static model with pre-determined endowments of land and labor.[2] Suppose that a farming household is endowed with $\bar{L}$ units of labor, which it divides between leisure L^l, work on the household farm L^h, and supply of labor to the market, L^m. The household has preferences over consumption C and leisure L^l, represented by the strictly increasing and concave utility function $U(C, L^l)$. Household members produce a single food commodity for sale or consumption using strictly increasing, concave production technology $F(L, X, A)$, where L represents total labor application, X is a vector of non-labor inputs, and A represents land inputs. Total output is $y = F(L, X, A)\epsilon$, where ϵ represents an exogenous production shock encompassing pest pressure and agro-climatic factors. The household can hire labor on the market, represented by L^d. Let p^x be the vector of non-labor input prices and w be the market wage rate. We normalize the price of the output to 1.

If the markets for labor, credit and other inputs are complete and competitive, the household's utility maximization problem takes the following form:

$$\max_{L^l, L^h, L^d, L^m, X, C} \quad \mathbb{E}_\epsilon \left[U(C, L^l \mid \epsilon, \overline{L}, A) \right] \tag{1}$$

$$\text{subject to:} \quad C - wL^m \leq F(L, X, A)\epsilon - wL^d - p^x X \tag{2}$$

$$L = L^h + L^d \tag{3}$$

$$\bar{L} = L^h + L^m + L^l \tag{4}$$

$$L^l, L^h, L^d, L^m, X, C \geq 0 \tag{5}$$

[2]This section is closely based on Benjamin (1992), which provides many key results for the single-period case. Our theoretical innovations relate to the dynamic model and the possible endogeneity of adjustments to the labor endowment.

where the utility function is conditioned on the endowments and the stochastic output shock. The equality in (2) will hold at the solution. Under current assumptions, the model is recursive, and the consumption and production sides of the household problem can be solved separately. Household members first choose L and X to maximize expected farm profit, which is on the right-hand side of (2). They then maximize utility, conditional on farm income. The solution is characterized by the following:

$$\Pi^*(w, p^x \,|\, \epsilon, \overline{L}, A) = \max\, pF(L^*, X^*, A)\epsilon - wL^{d^*} - p^x X^* \tag{6}$$

$$L^* = L(w, p^x \,|\, \epsilon, \overline{L}, A) \tag{7}$$

$$X^* = X(w, p^x \,|\, \epsilon, \overline{L}, A) \tag{8}$$

where equation (6) is the profit function, and equations (7)-(8) are the input demand functions. The complete markets assumption imposes a testable restriction on the input demand functions, namely, that labor demand is not a function of the household labor endowment, $\overline{L}$. If markets are complete and competitive, equation (7) depends only on prices, endowments, and the exogenous shock ϵ. If there are failures in at least two relevant markets, then relative prices fail to adjust to their competitive levels, separation fails, and labor demand on farm is related to the household labor endowment. There are no insurance markets in the settings we study, so the lack of a market to sell or buy risk is the first market failure. The question is whether there are one or more additional market failures that lead to non-separation.

There are two primary ways that labor market failure could induce non-separation in a static context. The first is through excess labor supply in the local market, possibly due to nominal wages that are sticky downwards (Benjamin, 1992; Hart, 1983; Osmani, 1990; Kaur, 2014). This is experienced by the average household as a binding constraint on the amount of labor that its members can supply to the market. If a marginal household member increases the household labor endowment more than the household demand for leisure, then households with more workers supply more work to the family farm. This is detectable in the data as positive conditional correlation between labor endowments and labor demand.

The other case stems from the opposite problem in the market – a shortage of workers – but has the same reduced form effect. It is often speculated that farmers cannot find sufficient workers during periods of peak labor demand, i.e., planting and harvest. In this case, larger households will be less constrained than smaller households, on average, leading once again to positive correlation between labor endowments and labor demand. For this situation to hold, something must prevent wages from adjusting upwards to clear the market. One possibility is a social norm against hard bargaining with one's employers. Such a norm could hold in equilibrium if those supplying labor know that they might be on the hiring side in future periods. But it may also be the case that neither information nor workers can move fast enough to clear the local labor market when signaling is costly and the period of peak labor demand is short.

The details for the above two cases have been extensively worked out in Benjamin (1992). Benjamin shows in a sample from Indonesia that household labor endowments do not significantly affect on-farm labor demand, and therefore does not reject the null hypothesis of complete markets. Dillon and Barrett (2014) apply the same test to cross-sectional data from five countries in sub-Saharan Africa, using a single wave of the panel data used in this paper. They strongly reject the null hypothesis of complete and competitive markets for all countries.

As mentioned above, this cross-sectional test has three key limitations. First, the static result does not control for relevant time invariant household characteristics that could drive the result, such as a preference for working on one's own farm or unobserved managerial skill. Second, absent a compelling instrument for household labor endowment, there is no way to control for possibly endogenous adjustments to the household labor endowment. Third, the cross-sectional test cannot distinguish between separation failures due to excess labor supply, to excess labor demand, or to failures in other markets. In the following section we develop an approach that uses panel data to address these three issues.

2.2 The dynamic model with endogenous labor endowments

With panel data, the test of complete markets from the previous subsection can be extended to the dynamic context (LaFave and Thomas, 2016). The inclusion of household fixed effects controls for relevant time invariant household characteristics, such as a preference for working on one's own farm, or unobserved managerial skill. Identification in a panel setting is from inter-annual changes in labor endowments. Such changes are common, because of in-migration, out-migration, death, or children growing into young adults who can work on the farm.[3] Deaths and aging are arguably exogenous to the labor market. Likewise, there are causes of migration that are not driven by short-term labor supply or demand, such as marriage, widowhood, or return from secondary school or university. However, it is possible that some changes due to migration are endogenous to local labor market conditions. Households experiencing a shortage of workers may recruit new members for the next season. Alternatively, households facing a shortage of off-farm work opportunities may release members to migrate elsewhere. This potential endogeneity of the labor endowment is a problem for both cross-sectional and panel data, but it becomes especially salient when households

[3]Note that we do not mention exiting the workforce due to old age as a possible source of changes to the endowment. We show in the next section that aging out of the workforce is surprisingly uncommon in sub-Saharan Africa.

are surveyed for multiple periods and researchers observe the changes taking place.

In this subsection we extend the static model to a dynamic setting in a way that incorporates the possibility of endogenous changes to household labor endowment. Our goal is to formally characterize the circumstances under which the *sign* of the change in household labor endowment is informative for understanding local labor market conditions. Based on the results in Dillon and Barrett (2014) we proceed under the assumption that for the study countries, separation can be rejected in the cross-section. That is, in regressions of total labor demand on household labor endowment and prices, the coefficient on the household labor endowment is positive and statistically significant.

Consider a household solving the problem of the previous section annually, in T consecutive years that are indexed by t. We assume that utility is additively separable across time and that the discount rate is 0, for simplicity. This allows us to consider each period's problem in isolation while still incorporating the critical aspect of the dynamics, which is that the household's endowment of labor can change between years. The labor endowment in period t, $\overline{L}_t$, is equal to net migration, M_t, plus the lagged labor endowment: $\overline{L}_t = \overline{L}_{t-1} + M_t$. Net migration is the sum of changes due to factors exogenous to local labor market conditions, M_t^x, and that from endogenous adjustment in response to a shortage or excess of labor, M_t^e. That is, $M_t = M_t^x + M_t^e$. Adjustments to the household labor endowment due to aging or death are part of M_t^x. We make one further adjustment to the static model, which is that the level of cultivated acreage, A_t, is now a choice variable, with price p^A.

Endogenous adjustments to household labor endowment may be costly. Let c^o (c^i) represent the non-negative household-specific cost of sending (receiving) a single out-migrant (in-migrant). These migration costs reflect only the direct costs of adjusting the labor endowment, due to factors such as travel or the fixed costs of accommodating a new household member. Indirect costs or benefits that accrue to the household because of a change in the number of members are incorporated directly into the labor endowment and the utility function (through its conditionality on $\overline{L}$). For simplicity we assume that costs c^o and c^i are linear in the number of migrants.

The timing within each year is as follows. Prior to beginning cultivation, the household learns M_t^x. Then the household chooses whether to further adjust its labor endowment with additional migration, which is tantamount to choosing M_t^e.[4] Finally, the household makes cultivation choices with its new level of labor endowment, $\overline{L}_{t+1}$, in accordance with the model of Section 2.1.[5]

We next discuss the two sets of labor market circumstances that could lead to non-separation and endogenous adjustment: shortages of workers (excess labor demand) and constraints on off-farm labor (excess labor supply), and failures in a market other than that for labor.

2.2.1 Limits to off-farm labor

Suppose that in period $t-1$ there is a limit to the amount of labor that household members can supply off-farm, represented by H. This could be due to downward stickiness in wages that results in excess labor supply in the market. This limit will bind if the household's marginal rate of substitution between leisure and consumption is lower than the market wage, w_{t-1}, after supplying H units to the market and making optimal on-farm labor decisions. That is, the limit binds if desired labor supply by household members exceeds $H + L^*(w_{t-1}, p_{t-1}^x, p_{t-1}^a \,|\, \epsilon_{t-1}, \overline{L}_{t-1})$. If that is the case, the household supplies additional work on farm until the marginal rate of substitution equals the

[4]Seasonal migration between cultivation periods is not included in M_t^e, unless it persists into the cultivation period in response to local labor market conditions. However, migration of this kind is less common in Tanzania than in other settings, e.g. Morten (2016).

[5]In a set of robustness checks we will allow for a variation in timing that accommodates adjustments to the labor endowment after cultivation has started, though the data are only partially informative on this point. See Section 4.2.

ratio of expected returns between consumption and leisure. This implicitly defines a shadow wage for the household, w^s, that is below the market wage: $U_2(C^*_{t-1}, L^*_{t-1} \,|\, \epsilon_{t-1}, \overline{L}_{t-1}) = w^s_{t-1} < w_{t-1}$.

Under these circumstances, what options are available for M^e_t? With two additional assumptions we arrive at the key result. First, we assume that offers to potential in- or out-migrants must be incentive compatible, with migrants responding to the relative wage offers. Second, we assume that migrants' effects on the demand for leisure are never greater in magnitude than their effects on labor supply. That is, households with $w^s_{t-1} < w_{t-1}$ cannot recruit an in-migrant who makes leisure so enjoyable for everyone that desired labor supply by entire household falls.

Given these assumptions, a household facing a labor demand shortage would not want to, and could not attract an in-migrant to work for wage $w^s_{t-1} < w_{t-1}$. This rules out $M^e_t > 0$. In the face of a binding constraint on off-farm work, the household will choose between $M^e_t = 0$ and $M^e_t < 0$ by evaluating the expected utility of each. This is a discrete choice problem because $M^e_t < 0$ represents out-migration by a household member. Focusing on the most likely scenario in which a full adult (someone over the age of 15) is the potential out-migrant, the household chooses $M^e_t = -1$ instead of $M^e_t = 0$ iff.:

$$
\begin{aligned}
&\mathbb{E}_{\epsilon_t}[U(C^*_t, L^*_t \,|\, \epsilon_t, \overline{L}_{t-1} + M^x_t - 1)] \;\geq\; \mathbb{E}_{\epsilon_t}[U(C^*_t, L^*_t \,|\, \epsilon_t, \overline{L}_{t-1} + M^x_t)] \qquad (9)\\
\Rightarrow &\mathbb{E}_{\epsilon_t}\left[U\left(L^*_t, X^*_t, A^*_t)\epsilon - w_t L^{d*}_t - p^x_t X^*_t - p^a_t A^*_t + w_t H - c^i,\; \overline{L}_{t-1} + M^x_t - 1 - L^{h*}_t - H\right)\right] \geq \\
&\qquad \mathbb{E}_{\epsilon_t}\left[U\left(F(L^*_t, X^*_t, A^*_t)\epsilon - w_t L^{d*}_t - p^x_t X^*_t - p^a_t A^*_t + w_t H,\; \overline{L}_{t-1} + M^x_t - L^{h*}_t - H\right)\right]
\end{aligned}
$$

Whether or not the above inequality is satisfied is not clear *a priori*. We could extend the model by abandoning the unitary household framework and trying to model individual-level utility for all current and potential household members in all possible migration states. However, this would be no more informative than the current set-up, because it would make a series of predictions that are not empirically distinguishable. For our purposes, the relevant findings so far are that households with a labor surplus will not recruit new members to solve their labor market problems, and that whether they release a member for out-migration because of labor market conditions depends on a number of unobserved household-specific factors.[6]

It still remains for us to determine how a household facing a labor demand constraint responds to exogenous changes in the labor endowment, M_t^x. The answer proceeds immediately from the results in Benjamin (1992). If M_t^x is positive – i.e., if there is net in-migration for reasons unrelated to labor market clearing – then the household labor supply will increase unless the new member raises the demand for leisure by more than the supply of work. We have assumed away that possibility. If M_t^x is negative, the opposite pertains, and the household reduces its labor supply in equilibrium. However, because $M_t^e < 0$ is also possible, the researcher does not know whether a reduction in the labor endowment is driven by an exogenous or endogenous adjustment.

This brings us to one of the key theoretical insights of the paper. Under the assumption of a binding constraint on off-farm labor, we can treat *increases* in the labor endowment as exogenous. Furthermore, in expectation, those increases will lead to increases in labor supplied to the household farm. Decreases in the labor endowment could be due to exogenous or endogenous factors, making inference more difficult. This generates a testable, necessary condition for non-separation due to a constraint on off-farm labor: labor supply must respond to increases in the labor endowment, which are necessarily exogenous. If increases in the labor endowment do not lead to increases in labor demand on farm, we can reject the hypothesis of a binding constraint on off-farm labor.

[6] In Section 3.2 we will examine whether migrants ever appear to be recruited because they have agricultural-specific skills.

2.2.2 Shortages of workers

The alternative form of a non-clearing labor market comes from a shortage of workers at the prevailing wage. This is most likely to be a problem during peak periods of labor demand. It is straightforward to see that this presents the opposite set of predictions from the previous subsection.

Suppose that in period $t-1$ there is a limit, D, on the amount of agricultural labor that the household can hire in the market. The limit binds if the marginal rate of substitution between leisure and consumption is greater than w_{t-1} at the optimal level on-farm labor. In this case, desired labor demand by household members is less than $L^*(w_{t-1}, p^x_{t-1}, p^a_{t-1} \,|\, \epsilon_{t-1}, \overline{L}_{t-1}) - D$. Once again the household supplies work on farm until the marginal rate of substitution equals the ratio of expected returns between consumption and leisure. This implicitly defines a shadow wage for the household, w^s, that is above the market wage: $U_2(C^*_{t-1}, L^*_{t-1} \,|\, \epsilon_{t-1}, \overline{L}_{t-1}) = w^s_{t-1} > w_{t-1}$.

Invoking the same relatively innocuous assumptions as the previous section, it is straightforward that a household facing a labor demand shortage would not release an out-migrant as a response to labor market conditions. This rules out $M^e_t < 0$. Once again taking M^x_t as given, the household chooses $M^e_t = 1$ instead of $M^e_t = 0$ iff.:

$$\mathbb{E}_{\epsilon_t}[U(C^*_t, L^*_t \,|\, \epsilon_t, \overline{L}_{t-1} + M^x_t + 1)] \;\geq\; \mathbb{E}_{\epsilon_t}[U(C^*_t, L^*_t \,|\, \epsilon_t, \overline{L}_{t-1} + M^x_t)] \tag{10}$$

$$\Rightarrow \mathbb{E}_{\epsilon_t}\left[U\left(F(L^*_t, X^*_t, A^*_t)\epsilon - w_t D - p^x_t X^*_t - p^a_t A^*_t + w_t L^{m*} - c^o,\ \overline{L}_{t-1} + M^x_t + 1 - L^{h*}_t - L^{m*}_t\right)\right] \geq$$

$$\mathbb{E}_{\epsilon_t}\left[U\left(F(L^*_t, X^*_t, A^*_t)\epsilon - w_t D - p^x_t X^*_t - p^a_t A^*_t + w_t L^{m*},\ \overline{L}_{t-1} + M^x_t - L^{h*}_t - L^{m*}_t\right)\right]$$

We cannot determine whether the above inequality holds or not. This leaves open the question of whether increases in the labor endowment are exogenous or endogenous to labor market conditions. However, we can be certain that households facing a labor shortage will not release out-migrants as a response to the non-clearing labor market. Reductions in the labor endowment must be exogenous.

This leads to the second key theoretical prediction of the paper. Under the assumption of a binding shortage of workers, we can treat *decreases* in the labor endowment as exogenous. Furthermore, in expectation, such decreases lead to reductions in labor demand on the household farm. This generates a testable, necessary condition for non-separation due to a shortage of workers: labor demand must fall when the labor endowment falls, and such a change is necessarily exogenous. If decreases in the labor endowment do not lead to decreases in labor demand on farm, we can reject the hypothesis of a binding labor shortage.

2.3 Empirical framework

The discussion above can be synthesized as a set of predictions about the dynamic relationship between labor demand and changes in household labor endowment. First, if there is excess labor, labor demand on farm will respond to increases in the labor endowment, but not to decreases. Second, if there is a labor shortage, the opposite is true. And third, if the labor market frictions are caused by credit constraints or other non-labor factors, labor demand on farm responds to both increases and decreases in household labor endowment. These predictions can be taken directly to the data by extending Benjamin's empirical framework to the panel setting, and by allowing the change in household labor endowment to affect labor demand asymmetrically. That is, decreases in household labor endowment are allowed to have a separate impact on labor demand than do increases in household labor endowment.

There is also a within-season, inter-temporal dimension to the farming household's problem, driven by the natural sequencing of agricultural activities. Farmers in an area typically plant during the same few weeks, when the rainy season begins, in order to promote seed germination and optimal uptake of nutrients. Other forces can induce rapid changes in labor demand and supply, such as the appearance of pests, the opening of a narrow harvesting window, the regulated start of a marketing period, or a drought. These shocks are highly spatially correlated, because households in a single area face similar agro-climatic and marketing conditions. Often, demand for labor increases just as supply decreases, because household members become less willing to supply work off-farm. Wages may fail to adjust quickly enough to clear the market. The implication is that the labor market may clear during some weeks or months of the cultivation period, but not during others. To allow for this possibility, we allow for differences in labor market conditions across agricultural activities, as much as is allowed by the data[7].

To control for household fixed effects we estimate the equation of interest in first differences. This specification lends itself naturally to the additional step of allowing asymmetric responses to the change in the labor endowment. Our first main estimating equation is as follows:

$$\Delta \log L_{fht} = \alpha + \beta_1 \Delta \log \overline{L}_{ht} + \beta_2 \mathbf{F} + \beta_3 \{\mathbf{F} \times \Delta \log \overline{L}_{ht}\} + \gamma_1 \Delta \log A_{ht} + \gamma_2 \Delta dem_{ht} + \nu_t + \epsilon_{fht} \quad (11)$$

where $\Delta \log L_{fht}$ is the change in log labor demand for activity f of household h between periods t and $t-1$. This dependent variable measures total labor demand during the relevant activity phase, including both household labor and hired labor. $\Delta \log \overline{L}_{ht}$ is the change in log labor endowment between periods t and $t-1$. $\mathbf{F}$ is a vector of dummy variables for agricultural activities: cultivation, planting, weeding, cultivation (which is weeding + fertilizer application), and harvest. The included

[7] More details about the different agricultural activities captured in different countries are provided in Section 3

activities differ between countries, depending on the degree of disaggregation in the agricultural labor questionnaire. $\Delta \log A_{ht}$ is the change in area cultivated by the household, Δdem_{ht} represents the change in the demographic composition of the household, ν_t are period fixed effects (which only apply to countries with more than two waves of data available), and ϵ_{fht} represents any idiosyncratic shocks which affect the change in labor demand for activity f of household h between periods t and $t-1$.

This specification overcomes the first limitation of the static specification, by controlling for time invariant household characteristics that may affect how households choose their labor demand. The coefficients of interest are β_1 and β_3, capturing how labor demand for a particular activity responds to changes in household labor endowment. We use the earliest agriculture phase, which is planting or all pre-harvest activities depending on the country, as the excluded category.

Estimates of specification (11) provide insight into whether changes in household labor endowment are correlated with changes in labor use on farm. But they do not account for the possible endogeneity of changes in endowments, and do not admit an interpretation the reveals the structure of underlying market failures. To allow the household to respond differently to increases and decreases in the labor endowment, we expand specification (11) by separating the change in household labor endowment variable, $\Delta \log \overline{L}_{ht}$ into three categories, households that increased their labor endowment $\Delta^+ \log \overline{L}_{ht}$, households that decreased their labor endowment $\Delta^- \log \overline{L}_{ht}$, and households which did not change their labor endowment. The third category is the omitted reference group. Substituting these two new terms into our previous specification, we get:

$$\begin{aligned} \Delta \log L_{fht} \quad &= \quad \alpha + \beta_1 \Delta^+ \log \overline{L}_{ht} + \beta_2 \Delta^- \log \overline{L}_{ht} + \beta_3 \mathbf{F} + \beta_4 \{\mathbf{F} \times \Delta^+ \log \overline{F}_{ht}\} \\ &\quad + \beta_5 \{\mathbf{F} \times \Delta^- \log \overline{L}_{ht}\} + \gamma_1 \Delta \log A_{ht} + \gamma_2 \Delta dem_{ht} + \nu_t + \epsilon_{fht} \end{aligned} \tag{12}$$

Now the coefficients of interest are, β_1, β_2, β_4, and β_5, as they directly apply the three predictions described above. More specifically, if β_1 is positive and statistically significant, and β_2 is not significant, this is consistent with excess labor during the period relevant for the omitted agricultural activity. Conversely, if β_2 is positive and significant, while β_1 is not, this is consistent with a labor shortage. Similarly, if $\beta_1 + \beta_4$ is significant, but not $\beta_2 + \beta_5$, then the labor market is characterized by excess labor during the period represented by the relevant member of $\mathbf{F}$. If $\beta_2 + \beta_5$ is positive and significant, while $\beta_1 + \beta_4$ is not, this is consistent with a labot shortage.

Things are less clear if both coefficients for a particular agricultural activity are statistically significant. If both β_1 and β_2 are significant, this could be evidence of a failure in some other market (e.g. that for credit). However, another possibility is that there is indeed a failure in the labor market, but that households are successfully adjusting their labor endowments between periods. In such a case, one direction of changes would be exogenous, the other a mix of exogenous and endogenous changes. There is no way to distinguish between these possibilities within the framework developed so far.

3 Data and descriptive patterns

3.1 Data and sample

To test the predictions of the above model we use panel data from the Living Standards Measurement Study and Integrated Surveys on Agriculture (LSMS-ISA) project. These surveys are comprehensive household and agricultural surveys, conducted by the statistics offices of participating countries with cooperation and guidance of the World Bank. The data are nationally representative, span a wide range of topics, and are reasonably comparable across countries.

We focus our attention on farming households in Ethiopia, Malawi, Tanzania, and Uganda. A household is considered a farming household if it reports cultivation of a positive amount of land during the seasons under study. Table 1 describes the various waves of data we have for each country.

Table 1: Data Coverage in LSMS-ISA

Country	Number of Waves	Years
Ethiopia	2	2011-12, 2013-14
Malawi	2	2010-11, 2013
Tanzania	3	2008-09, 2010-11, 2012-13
Uganda	4	2005-06, 2009-10, 2010-11, 2011-12

Table 2 provides summary statistics for each country. Observations are captured at the household-wave level. Table 2 shows that Tanzania households have the largest households, and that Malawian households have the least amount of labor demanded for their agricultural activities. Ethiopian households have the largest amount of labor demanded. The four countries have relatively similar shares of prime-age adults in their households (between 47% and 50%), though Ethiopia households have larger shares of elderly adults, 15% compared to 8, 8, and 10% for the other countries. Households in Malawi have the least amount of acres under cultivation, and also the youngest household heads. Household heads in Ethiopia have obtained the least amount of education.

Two details about the data apply to all of the countries. First, households are included in the sample if they report non-zero acres cultivated and non-zero labor demand for at least half of the waves available for that country. Secondly, for the acres owned and cultivated variables, we first use the area as reported by respondent, and supplement with surveyor recorded GPS area if reported area was missing.

While the data for each of the countries studied are largely comparable, there are some differences. For Ethiopia, the data identifies children who worked on the farm as being from the household, hired, or working as free/exchange labor. The Ethiopia data also separately records agricultural labor as either being pre-harvest or harvest related. As for Malawi, the data does not separately record outside workers who may have worked as free/exchange laborers. While household labor is categorized into preparation, weeding, and harvest activities for Malawi, the hired labor is only categorized as pre-harvest and harvest labor. Therefore, we combine the household preparation and weeding activities into a pre-harvest category to match the hired laborers.

For the Tanzania data, children hired from outside the household are only recorded in Waves 2 and 3. All three waves of the Tanzania data record labor for preparation, weeding, and harvest activities. Waves 2 and 3 also record fertilizing labor, but we exclude that as it appears to not be included in any of the Wave 1 categories. For the Uganda data, children working on the farm from the household are only captured in Wave 1, whereas children hired from outside the household are recorded in all four waves. The Wave 1 data for Uganda records labor separately by agricultural activity, but Waves 2-4 only record total labor, so we combine all the Wave 1 activities together for consistency. Wave 4 also separately records free/exchange labor, which we include in the total labor demand as it appears to have been counted as hired labor in the previous waves.

Another issue with the Uganda data is that the survey only provides details for up to three workers on each plot. The total amount of days worked on each plot is recorded, but if more than three family members worked on a plot, we are not able to identify who they are. This means that we are not able to separate out children from adults for household workers, and need to treat all household labor equally.

Table 2: Summary Statistics for LSMS-ISA Data

	(1) Ethiopia mean/sd	(2) Malawi mean/sd	(3) Tanzania mean/sd	(4) Uganda mean/sd
Log of Labor demand (Person-Days)	4.54	4.14	4.24	4.17
	(1.94)	(1.12)	(2.34)	(2.02)
Household Size	2.72	2.78	3.06	2.90
	(1.27)	(1.37)	(1.67)	(1.59)
Prime male share	0.23	0.23	0.24	0.23
	(0.17)	(0.17)	(0.19)	(0.20)
Prime female share	0.26	0.25	0.26	0.24
	(0.17)	(0.16)	(0.17)	(0.17)
Elderly male share	0.08	0.03	0.04	0.04
	(0.13)	(0.12)	(0.13)	(0.12)
Elderly female share	0.07	0.05	0.06	0.04
	(0.15)	(0.15)	(0.17)	(0.14)
Acres Cultivated	7.15	2.04	5.89	4.40
	(150.19)	(4.38)	(14.33)	(35.15)
Acres Owned	5.98	1.81	5.47	3.78
	(149.46)	(4.37)	(14.38)	(34.49)
Age of head (years)	45.53	44.43	49.73	46.58
	(14.97)	(16.20)	(15.48)	(15.28)
Education of head (years)	1.59	5.57	5.06	4.50
	(2.85)	(4.31)	(3.29)	(3.36)
Expenditure per Capita	4,384	60,958	531,866	48,935
	(189,631)	(87,057)	(480,372)	(57,030)
Number of Obs.	5,695	4,824	6,115	8,661

Notes: Standard deviations are in parentheses. Expenditure figures are expressed in the local currency.

Table 3: Summary of Labor Demand for LSMS-ISA Data

	(1) Ethiopia mean/sd	(2) Malawi mean/sd	(3) Tanzania mean/sd	(4) Uganda mean/sd
Cultivation	229.04	97.40		
	(497.19)	(95.27)		
Harvest	83.07	21.56	48.52	
	(150.14)	(35.93)	(73.89)	
Preparation			61.21	
			(74.50)	
Weeding			55.84	
			(63.54)	
All Farm Activities				131.33
				(137.41)
Number of Obs.	5,695	4,824	6,115	8,661

Notes: Standard deviations are in parentheses. Labor Demand is measured in person-days.

3.2 Descriptive patterns

In this subsection we present seven descriptive patterns that are informative for the empirical analysis to follow. Some of these descriptive findings are presented to clarify the setting. Others provide justification for our empirical approach.

3.2.1 Intra-annual time path of labor demand

Because weather conditions are highly spatially correlated, and the timing of farming activities is connected to the weather, a natural concern is that labor markets may seize up during peak periods of demand. Short-term breakdowns in the labor market are even more likely if those who supply labor are farmers themselves, so that there is a negative labor supply shock at the same time as a positive labor demand shock. To examine the demand side of this issue we estimate kernel regressions of plot-level labor demand across time. We are only able to do this for the 2012-2013 data from Malawi. The other survey waves include activity-specific measures of labor demand (to varying degrees), but lack sufficient information for us to match activities to time periods.

Figure 1 shows region-specific kernel regressions of plot-level labor demand against time, for the 2012-2013 data from Malawi. The dependent variables are quantity demanded of household labor on the left, and quantity demanded of hired labor on the right. The scale of on-farm labor provided by household members dwarfs that of hired workers. The significant peak in labor demand in late 2012 represents the planting season. Planting is typically timed to coincide with or immediately follow the onset of the rains, and begins earlier in the South region than in the Central or the North. A second period of higher labor demand occurs around harvest. However, the spike at the harvest is not as pronounced as that at planting, and in the South there is no increase in labor

Figure 1: Region-specific time path of plot-level labor demand, Malawi, 2012-2013

All variables used to construct this figure were reported at the household-plot level. Respondents reported the month in which planting ended, and whether planting ended in the first half or second half of the month. To estimate the beginning of the planting period, we took the maximum number of weeks that any household member worked for land preparation or planting on the plot, and rounded to the nearest half-month, letting 4 weeks equal a month. If household planting labor on the plot was reported to be zero (7% of plots), we assigned the plot the shortest possible planting duration, 0.5 months. Households reported the month in which harvesting began, and the month in which harvesting ended. We estimated the harvest duration as 0.5*(difference+1), where difference = end month - start month. This set the harvest period duration equal to its expected value, given information on the start and end month only. We then randomly assigned plots to begin the harvest at the beginning or middle of the reported start month. Based on these estimates, a small number of plots (<3%) reported an end of planting period that occurred after the beginning of the harvest period. While this is possible on multi-cropped plots, we dropped these plots because of suspected misreporting. We then assigned the end of the planting period to be the beginning of the cultivation (weeding and fertilizing) period, and the beginning of the harvest period to be the end of the cultivation period. Household labor supply to the plot was reported separately for planting, cultivation, and harvest. We evenly divided the total person-days for each period between the half-months in that period. Hired labor supply to the plot was reported as an aggregate figure for planting and cultivation, with harvest reported separately. We evenly divided the planting and cultivation labor between the half-months in those two periods, and assigned harvest labor to that period in the same manner as the household harvest labor. Note that if hired labor follows a pattern similar to household labor, with more intensive application at planting, then this approach will underestimate planting labor and overestimate cultivation labor. Child labor is excluded from these estimates as it was not collected for the hired workers, but this is a negligible fraction of total labor. We assume that a full work day is 6 hours long. The figure shown is a local polynomial regression of labor demand (in person-days per half-month) on time, using an Epanechnikov kernel.

demand at harvest. This reflects the fact that for many crops farmers have more leeway with the timing of the harvest than they do with planting.[8]

The pattern in Figure 1 is specific to Malawi. The degree of intra-annual variation may be different in the other study countries. Nevertheless, these figures underscore the importance of allowing for short-term, intra-annual variation in the tightness of the labor market when we test for non-separation. While we do not observe the week or month of labor demand in any surveys other than the one shown here, we do observe (to varying degrees) the breakdown of labor demand

[8]A limitation of the data underlying Figure 1 is that it is specific to a particular cultivation cycle. If work on the previous or subsequent cultivation cycle overlaps the months depicted for some households, then the figure underestimates labor demand near the beginning and end of the shown period. This is unlikely to be a problem for the middle 8-10 months shown in the figure, which is our main concern.

across activities such as planting, weeding, or harvesting. In the empirical analysis we allow for heterogeneity across these activities.

3.2.2 Intra-annual time path of labor supply

It is also instructive to examine the time path of labor supplied to one's own farm and to the market. Presumably, if the quantity of labor demanded on farms follows the seasonal pattern shown in Figure 1, then the quantity of labor supplied should show a similar pattern. The 2010-2011 Malawi survey is the only one with sufficient detail for us to examine the time path of labor supply over the year. The sample in that year included 11,744 households, with interviews spaced uniformly across a calendar year. The survey included questions about the number of hours each household member spent on particular activities over the previous seven days. To focus on the supply side of the agricultural labor market, we use the responses for time spent working on one's own farm and time spent supplying *ganyu* labor, or casual farm labor, in a local spot market.

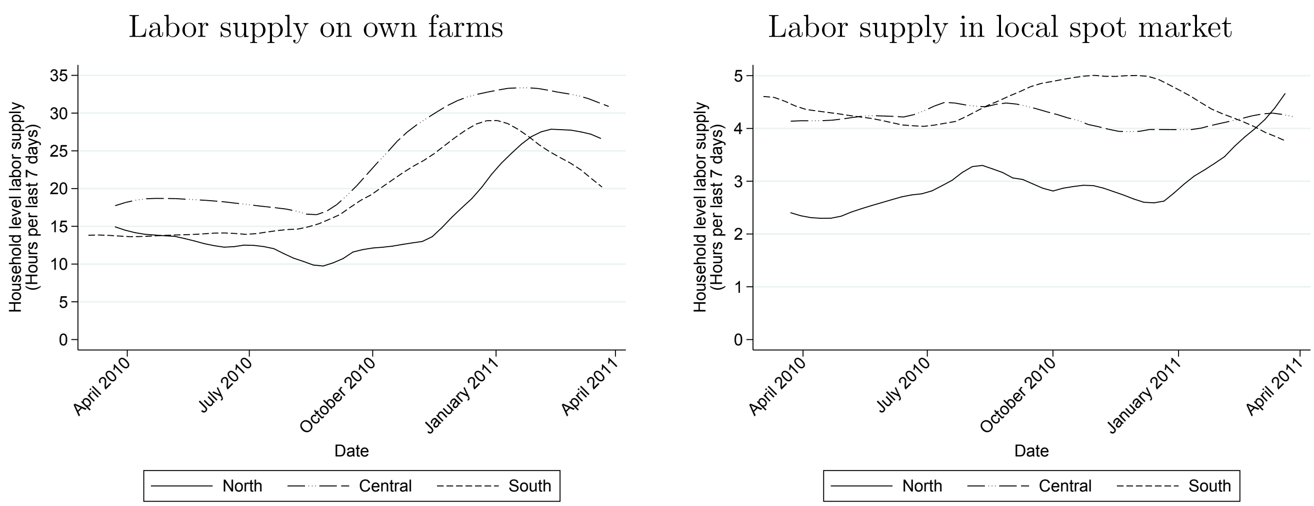

Figure 2: Region-specific time path of household-level labor supply, Malawi, 2010-2011

Results show kernel regressions of household-level labor supply by men or women against interview date. Sample sizes are 2,212 households in North, 4,053 households in Central, 5,479 households in South. Marked dates are for the first day of the month.

Figure 2 shows the time path of labor supply on own farms and in the agricultural labor market. The plots show kernel regressions of the number of hours spent on the activity over the last seven days against the interview date. There are numerous takeaways from these figures. First, average labor supply to own farms is much greater than that to the market, at all times of year. Second, intra-annual variation in labor supply to the market is less pronounced than that to own farms. The latter increases by 100-300% from peak to trough, while the former ranges from near-constant to at most a 50% increase from peak to trough. These two patterns are suggestive of non-separation: workers are ready and willing to supply additional labor on their own farm when needed, but do not (or cannot) exhibit a similar response in the labor market. Third, the general pattern of intra-annual dynamics matches the labor demand side. The peak in own-farm labor supply occurs around the planting period. The period of highest labor supply is earlier in the South than in the North, as in Figure 1. The peak for the Central region is slightly earlier than in the labor demand case.

Taken together with the labor demand findings from the previous subsection, the time paths in Figure 2 underscore the importance of allowing for different labor market conditions at different times of year. Because we cannot match the labor data in most surveys to specific dates, we instead accommodate this concern by allowing for heterogeneity in separation across farming activities, which roughly align with time periods.

3.2.3 Labor force participation on the extensive and intensive margins

What constitutes the "labor endowment" of a farming household in sub-Saharan Africa? Because the analysis in this paper hinges on our definition of the labor endowment, we need a clear answer to this question. The main challenge relates to choosing the age cutoffs at which someone enters or exits the labor force. The survey data do not include information on desired or intended labor supply. Some children and seniors supply labor to the market. Many of those who do not work in the market still participate in household agriculture.

Figure 3: Labor force participation on the extensive and intensive margins, by country and activity

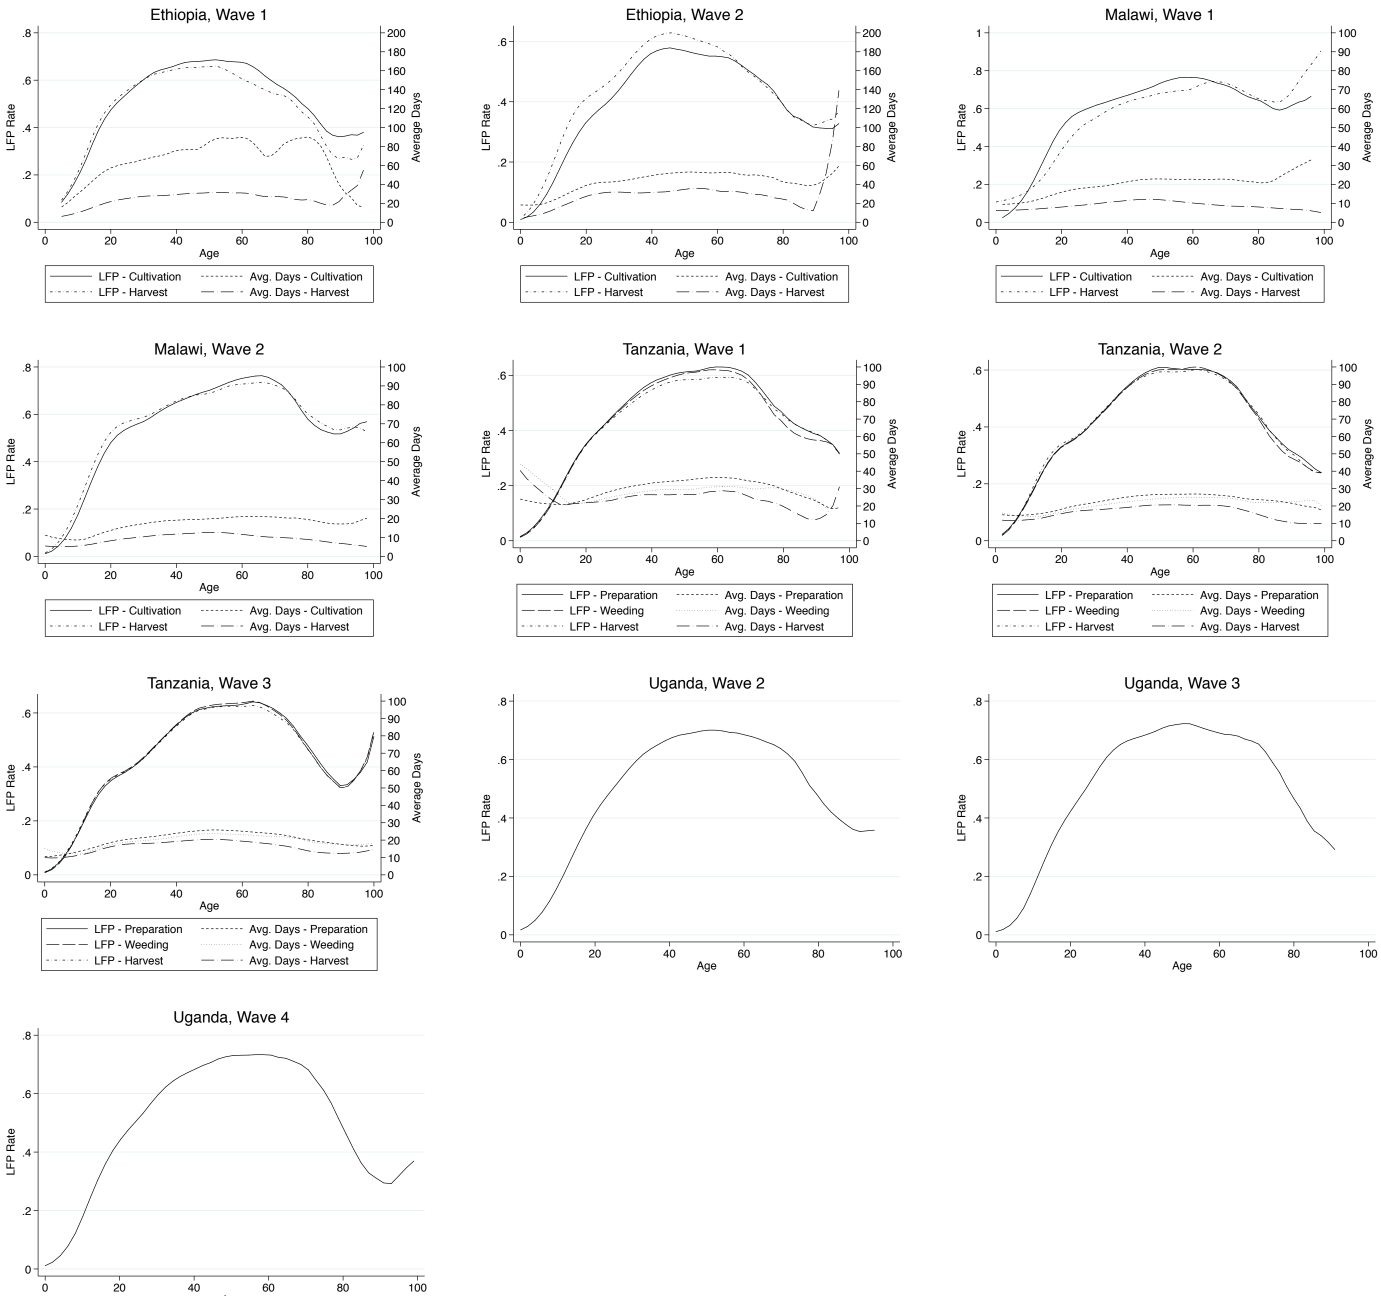

Authors' calculations from LSMS-ISA data. Uganda data do not include a breakdown by activity and do not allow for differentiation in work days at the individual level.

To guide this decision we look at labor force participation by age across our data sets. Figure

3 shows kernel regressions of the labor force participation rate on the household farm (LFP - the extensive margin) and the average number of days worked by those who are working (the intensive margin), plotted against age. Separate plots are shown for each farming activity. In all figures the scales are set so that the LFP regression lines appear above those for average days worked. The scale for LFP is on the left and the scale for average number of days worked is on the right.

There are again three key patterns in these figures. First, to our surprise, older people do significant work on farms. In all figures, the LFP rate for 80-year-olds is higher than that for 20-year-olds. This finding is based on admittedly small samples of people at or around age 80.[9] Confidence intervals at the upper end of the age distributions are wide. Yet, the extensive margin drop-off in work between ages 60 and 80 is more gradual than the increase that occurs during youth. Furthermore, there is little variation across the age distribution in the average days worked by those who supply labor to the farm. Across countries and activities there is essentially no difference in average days worked between 40-year-olds and 70-year-olds.

Second, there is little meaningful variation between farming activities, on either margin. On the extensive margin there are no differences that merit comment. On the intensive margin, the only notable finding is that workers supply fewer average days of work for the harvest than for other activities, a pattern that holds across ages and countries.

Third, the most rapid changes in labor force participation occur between ages 10 and 20. Youth in Malawi work more than in the other countries, reflecting the overall higher levels of labor force participation in that country. However, for all countries the rate of growth of the LFP rate increases during childhood and only begins to decline during late teenage years. Clearly, though

[9] The upper tails of the age distributions are as follows: ET wave 1, 2.0% are over age 70, 0.6% over age 80; ET wave 2, 2.1% and 0.5%; MW wave 1, 2.0% and 0.6%; MW wave 2, 1.9% and 0.6%; TZ wave 1, 2.6% and 0.6%; TZ wave 2, 2.7% and 0.9%; TZ wave 3, 2.8% and 1.0%.

it would be incorrect for us to treat all children as part of the labor endowment, it would also be incorrect to exclude teenagers entirely.

These observations guide our definition of the household labor endowment. Based on the first point, we do not assume that anyone "ages out" of the labor endowment. All adult household members are counted as potential workers. At the other end of the age distribution, we allow children to gradually age into the workforce as the data allows[10]. From age 11-15 we assume a linear adult equivalence scale: 11-year-olds count as 0.2 adults in the labor endowment, 12-year-olds as 0.4, and so on. As a robustness check we also use a binary cut-off at age 15.

3.2.4 Between-year changes in household labor endowments

With a definition of the labor endowment ($\overline{L}_t$) in hand, we next examine the degree of inter-annual variation in these endowments. This variation is the basis of our identification strategy. Changes in labor endowments can occur for three reasons: new people move in, previous household members move out or pass away, or children age into the workforce.

Table 4 shows the changes in the household labor endowment between survey waves. The upper panel is based on our preferred measure of "aging in", using a linear equivalence scale from ages 11 to 15. The lower panel is based on a discrete cut-off of the labor endowment at age 15. All entries are the sample means of household-level variables. The units for the labor endowment statistics are adult equivalents. The first row in each panel shows the average net change in the number of household members, without consideration of age. The second row gives the average net change in the size of the labor endowment. The next three rows decompose labor endowment changes into move-ins, move-outs (which includes deaths), and aging into the workforce. The final

[10]The ages of children working on the farm are not captured for Ethiopia, hired labor in Malawi, hired labor in waves 2 and 3 for Tanzania, and for Uganda. In these cases, children's labor supply is discounted by 50%.

Table 4: Inter-annual changes in number of members and labor endowment, household level

	Ethiopia	Malawi	Tanzania		Uganda		
Change between waves:	1 & 2	1 & 2	1 & 2	2 & 3	1 & 2	2 & 3	3 & 4
Method 1: children grow into labor endowment gradually from age 11-15							
Δ Number of members	0.04	0.43	0.35	-0.04	0.40	-0.31	0.03
Δ Labor endowment	0.05	0.35	0.25	0.04	0.28	-0.15	0.05
Δ Labor endowment: move-ins	0.21	0.36	0.31	0.27	0.76	0.21	0.30
Δ Labor endowment: move-outs	-0.43	-0.41	-0.34	-0.52	-1.04	-0.56	-0.46
Δ Labor endowment: aging children	0.27	0.40	0.28	0.29	0.56	0.19	0.20
Any net Δ in labor endowment (=1)	0.69	0.71	0.68	0.71	0.80	0.71	0.71
Increase in labor endowment (=1)	0.43	0.53	0.50	0.45	0.51	0.41	0.47
Decrease in labor endowment (=1)	0.26	0.18	0.18	0.26	0.30	0.31	0.24
Method 2: children enter labor endowment at age 15							
Δ Number of members	0.04	0.43	0.35	-0.04	0.40	-0.31	0.03
Δ Labor endowment	0.04	0.29	0.25	0.02	0.22	-0.15	0.05
Δ Labor endowment: move-ins	0.19	0.32	0.29	0.25	0.64	0.18	0.26
Δ Labor endowment: move-outs	-0.38	-0.38	-0.32	-0.48	-0.92	-0.50	-0.40
Δ Labor endowment: aging children	0.23	0.35	0.28	0.26	0.50	0.17	0.20
Any net Δ in labor endowment (=1)	0.42	0.47	0.44	0.45	0.59	0.41	0.42
Increase in labor endowment (=1)	0.23	0.34	0.31	0.25	0.36	0.17	0.24
Decrease in labor endowment (=1)	0.18	0.13	0.13	0.20	0.23	0.24	0.17

Notes: Authors' calculations from LSMS-ISA data. In the upper panel, children between ages 11 and 15 are counted as 1-0.2(15-age) working age adults. In the lower panel, children are not counted as workers until they reach age 15, when they are counted as one full worker. Δ Labor endowment is the sum of the three categories immediately below.

three rows shows the proportion of households experiencing any change in labor endowment, a positive change, or a negative change, respectively. The only noteworthy difference between the upper and lower panels is in these final three statistics. A higher proportion of households experience a change in labor endowment when children age into the labor endowment gradually, because more households have one or more members in the age range 11-14 than have a child who turns 15 between waves.

Focusing on the upper panel of the table, the most important takeaway is that approximately 70-80% of households experience a net change in labor endowment from one survey to the next.

The majority of surveyed households contribute to identify the effects of interest. Looking at the decomposition of changes to the labor endowment, the average reduction due to move-outs is greater than the average increase due to move-ins. However, after accounting for children aging into the workforce, the average net change in labor endowment is positive in all but one survey wave. The average changes in labor endowments are not large. The largest is for Malawi, at 0.35 adult equivalents. The smallest is in Uganda between waves 2 and 3, at -0.15 adult equivalents. The last two rows in the top panel show that roughly twice as many households experience a net increase as experience a net decrease.

3.2.5 Spatial distribution of net changes in labor endowments

We have just established that most households experience between-survey changes in labor endowments. We next examine the spatial distribution of these changes. This helps us understand the prevalence of arguably exogenous variation in labor endowments over time. In Section 2 we formalized the intuitive notion that only one direction of adjustment can be endogenous to a non-clearing labor market: increases in endowments are exogenous to surplus labor, and decreases are exogenous to labor deficits. If we see that households are not clustered geographically based on the sign of labor endowment changes, this is suggestive support for the widespread presence of at last some changes that are exogenous to local labor market conditions.

Figure 3.2.5 shows location-level scatter plots of the proportion of households with a net increase (vertical axis) and a net decrease (horizontal axis) in the labor endowment, by survey wave. The points are defined at the enumeration area level for all countries. By construction, all points lie in the triangle formed by the origin and the value of 1 on each axis. The clear pattern in all surveys is of clustering in the interior of that triangle. This indicates that in the large majority

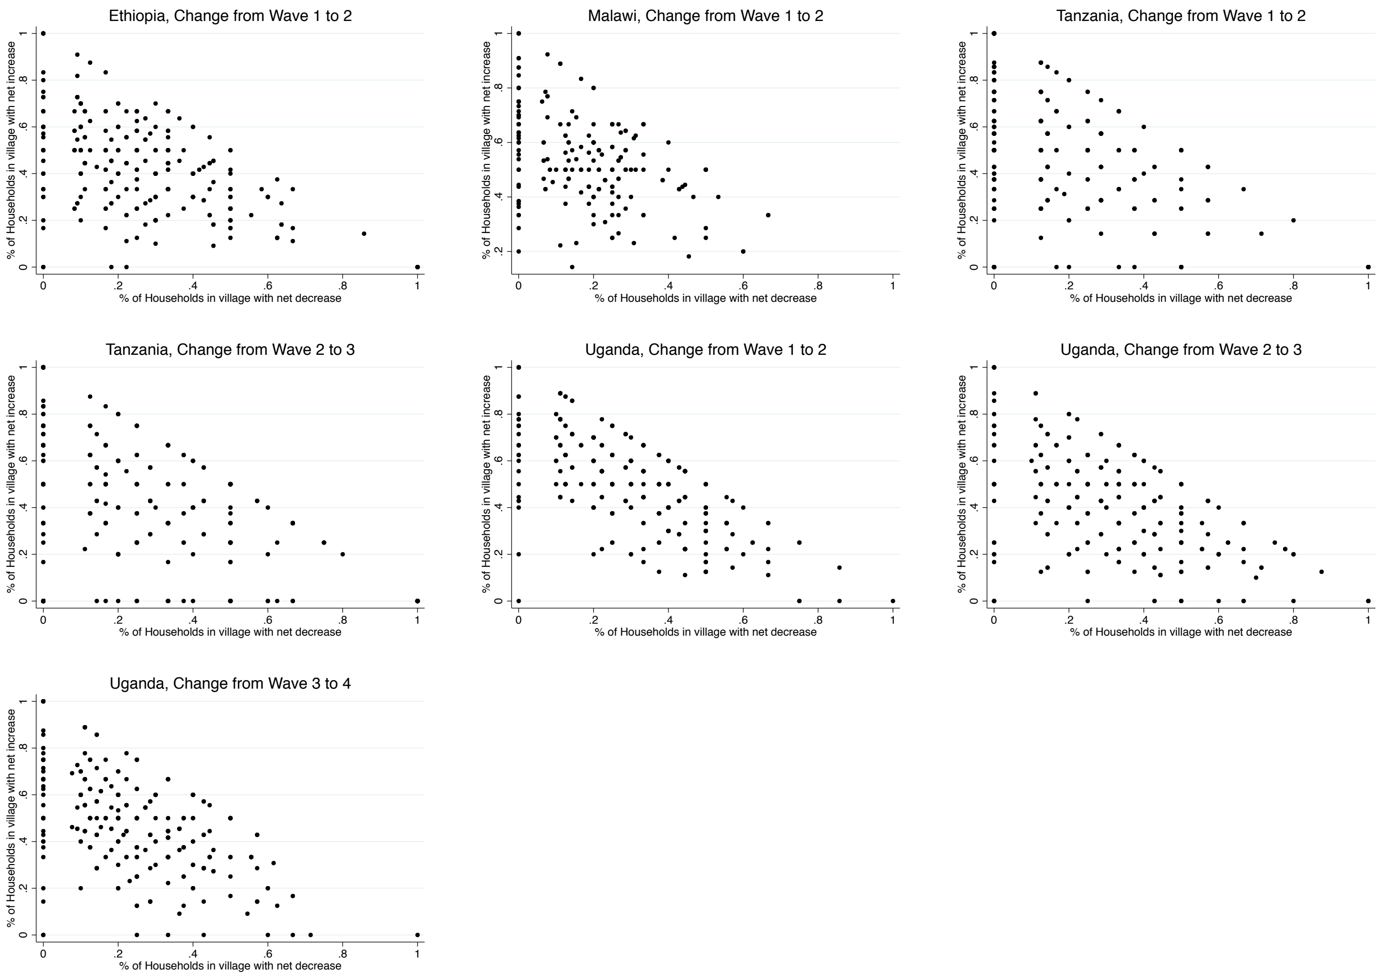

Figure 4: Distribution of net labor endowment changes, village level

Notes: Authors' calculations from LSMS-ISA data. Each dot represents a village or enumeration area. These are of roughly equivalent size in most surveys. Alternative figures using circles to represent the size of each village show a similar pattern, and are available upon request.

of locations there are some households that experience a net increase in labor endowment, and some that experience a net decrease. In the empirical analysis to follow we will lack statistical power to allow for spatial heterogeneity by interacting the change in endowment with location fixed effects. So it is comforting that positive and negative changes in endowments are widely spread. We are not asserting that there are no endogenous adjustments, only that there are many exogenous adjustments as well, and that they are geographically disbursed within each study country.

3.2.6 Descriptive statistics for migrants and stayers

What tasks do in- or out-migrants do, relative to others in the household? One indication of endogenous adjustments to the labor endowment, particularly endogenous increases, would come from an unusually substantial role for migrants in agricultural activities. Imagine a similar test for domestic workers in urban areas. These household members are recruited specifically to do housework. Thus, we would expect to find in summary statistics that they do substantially more housework than permanent members. While not a perfect test of skill- or task-specific recruitment, the general pattern is still informative.

In this subsection we examine the labor supply to the farm by migrants and non-migrants. We also look at the summary statistics for migrants and stayers in order to characterize the population that underlies our identification strategy.

Table 5 shows migrant and stayer characteristics for Ethiopia. Columns 1 and 2 differentiate between out-migrants and in-migrants. The table uses data from two waves. Column 1 uses data from Wave 1, before the individual moved-out, and column 2 only uses data from Wave 2, after the migrant moved-in. Column 3 takes an average of the time varying variables (e.g. age, average days worked, etc.) across Waves 1 and 2. Essentially none of the between-column differences are statistically significant. There are two notable takeaways. First, the large majority of both in- and out-migrants are family members. The average migrant is in his or her late 20s, with balance across genders. It is likely that many of these changes in the labor endowment are due to marriage. Second, in-migrants and out-migrants do not do substantially more or less work on farm than stayers. Rather, they seem to fit into the working life of the household when they are a part of it, working on the family farm much like everyone else.

Table 5: Summary statistics for migrants and stayers: Ethiopia

	(1) Moved-Out mean/sd	(2) Moved-In mean/sd	(3) Stayed mean/sd
Age	26.25	29.05	22.28
	(15.36)	(16.73)	(18.44)
% Male	0.51	0.45	0.51
	(0.50)	(0.50)	(0.50)
% Child	0.60	0.34	0.57
	(0.49)	(0.47)	(0.50)
% Other Family Member	0.35	0.57	0.43
	(0.48)	(0.50)	(0.49)
% Not a Family Member	0.06	0.07	0.00
	(0.23)	(0.26)	(0.04)
% Worked Cultivation	0.54	0.35	0.42
	(0.50)	(0.48)	(0.45)
Avg. Days Worked Cultivation	66.11	43.32	59.81
	(117.96)	(52.97)	(98.43)
% Worked Harvest	0.55	0.50	0.42
	(0.50)	(0.50)	(0.44)
Avg. Days Worked Harvest	23.10	27.22	37.02
	(29.32)	(37.02)	(39.30)
Number of Obs.	1,084	544	13,180

Notes: Authors' calculations from LSMS-ISA data.

In Table 6 we show the same statistics for Malawi. The characteristics of migrants and stayers are remarkably similar to those for Ethiopia. The only noteworthy difference between the countries is that an even smaller percentage of migration in Malawi is by non-family members. Once again the differences between the three columns are not statistically significant.

In Table 7 we show the same set of summary statistics, for Tanzania. The first three columns display the results for changes between Waves 1 &2, and the last three columns for changes between Waves 2 & 3. Once again, the third and sixth columns show the characteristics of stayers pooled across both waves. For Tanzania the most important patterns from Ethiopia and Malawi are maintained. It is notable, and somewhat surprising, that more women than men migrate in Tanzania. This could reflect between-country variation in customs governing the residency location

Table 6: Summary statistics for migrants and stayers: Malawi

	(1) Moved-Out mean/sd	(2) Moved-In mean/sd	(3) Stayed mean/sd
Age	28.37	29.25	22.28
	(16.04)	(15.21)	(18.60)
% Male	0.50	0.51	0.49
	(0.50)	(0.50)	(0.50)
% Child	0.59	0.31	0.58
	(0.49)	(0.46)	(0.49)
% Other Family Member	0.40	0.65	0.41
	(0.49)	(0.48)	(0.49)
% Not a Family Member	0.01	0.02	0.00
	(0.10)	(0.13)	(0.03)
% Worked Cultivation	0.55	0.48	0.45
	(0.50)	(0.50)	(0.45)
Avg. Days Worked Cultivation	19.06	16.78	19.49
	(19.75)	(15.33)	(14.58)
% Worked Harvest	0.41	0.55	0.45
	(0.49)	(0.50)	(0.42)
Avg. Days Worked Harvest	9.80	7.43	14.64
	(16.06)	(10.76)	(12.05)
Number of Obs.	841	703	10,004

Notes: Authors' calculations from LSMS-ISA data.

of couples after marriage. However, the finding that both in- and out-migrants work on the farm to a similar extent as everyone else is as true in Tanzania as in Ethiopia and Malawi.

Finally, in Tables 8, 9, and 10, we show the similar statistics for Uganda. However, in Uganda, the data does identify the specific person who worked on each plot in Wave 1, so column 1 of Table 8 can not show what percent over out-migrants worked on the farm. Also, none of the waves of the Uganda data captures how much each person works on each plot, so each of the tables here are unable to show the average days worked for each type of migrant. Despite these data limitations, the tables show that migrants in Uganda fit a similar pattern as the other countries, and do not appear to be much different than other household members.

Table 7: Summary statistics for migrants and stayers: Tanzania

	Wave 1-2			Wave 2-3		
	(1)	(2)	(3)	(4)	(5)	(6)
	Moved-Out	Moved-In	Stayed	Moved-Out	Moved-In	Stayed
	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd
Age	28.44	28.71	24.02	29.60	29.77	24.32
	(17.10)	(15.74)	(20.05)	(17.72)	(16.68)	(20.34)
% Male	0.37	0.38	0.49	0.41	0.32	0.50
	(0.48)	(0.48)	(0.50)	(0.49)	(0.47)	(0.50)
% Child	0.44	0.28	0.48	0.40	0.27	0.49
	(0.50)	(0.45)	(0.50)	(0.49)	(0.45)	(0.50)
% Other Family Member	0.51	0.65	0.52	0.55	0.62	0.51
	(0.50)	(0.48)	(0.50)	(0.50)	(0.48)	(0.50)
% Not a Family Member	0.04	0.07	0.00	0.05	0.10	0.01
	(0.21)	(0.25)	(0.05)	(0.21)	(0.30)	(0.08)
% Worked Cultivation	0.49	0.39	0.44	0.47	0.41	0.44
	(0.50)	(0.49)	(0.44)	(0.50)	(0.49)	(0.44)
Avg. Days Worked Cultivation	27.18	23.38	28.10	22.79	23.70	24.07
	(35.87)	(19.68)	(25.82)	(20.06)	(21.91)	(17.79)
% Worked Harvest	0.47	0.39	0.41	0.48	0.44	0.44
	(0.50)	(0.49)	(0.49)	(0.50)	(0.50)	(0.50)
Avg. Days Worked Harvest	25.13	19.26	25.07	18.44	20.00	17.44
	(32.78)	(18.88)	(40.53)	(22.88)	(21.40)	(19.82)
Number of Obs.	655	598	10,088	1,001	508	10,302

Notes: Authors' calculations from LSMS-ISA data.

Table 8: Summary statistics for migrants and stayers: Uganda Waves 1-2

	(1)	(2)	(3)
	Moved-Out	Moved-In	Stayed
	mean/sd	mean/sd	mean/sd
Age	25.67	27.45	22.58
	(14.33)	(14.46)	(18.17)
% Male	0.48	0.44	0.50
	(0.50)	(0.50)	(0.50)
% Child	0.36	0.22	0.47
	(0.48)	(0.42)	(0.50)
% Other Family Member	0.41	0.54	0.44
	(0.49)	(0.50)	(0.50)
% Not a Family Member	0.03	0.08	0.00
	(0.18)	(0.27)	(0.05)
% Worked on Farm		0.29	0.40
		(0.45)	(0.49)
Number of Obs.	2,392	1,666	10,630

Notes: Authors' calculations from LSMS-ISA data.

Table 9: Summary statistics for migrants and stayers: Uganda Waves 2-3

	(1) Moved-Out mean/sd	(2) Moved-In mean/sd	(3) Stayed mean/sd
Age	25.29	26.35	21.32
	(12.61)	(13.29)	(18.43)
% Male	0.50	0.38	0.50
	(0.50)	(0.48)	(0.50)
% Child	0.33	0.23	0.50
	(0.47)	(0.42)	(0.50)
% Other Family Member	0.33	0.57	0.48
	(0.47)	(0.50)	(0.50)
% Not a Family Member	0.07	0.12	0.01
	(0.25)	(0.33)	(0.09)
% Worked on Farm	0.26	0.13	0.35
	(0.44)	(0.34)	(0.48)
Number of Obs.	1,322	482	13,733

Notes: Authors' calculations from LSMS-ISA data.

Table 10: Summary statistics for migrants and stayers: Uganda Waves 3-4

	(1) Moved-Out mean/sd	(2) Moved-In mean/sd	(3) Stayed mean/sd
Age	25.56	26.95	21.69
	(14.20)	(14.71)	(18.81)
% Male	0.46	0.45	0.49
	(0.50)	(0.50)	(0.50)
% Child	0.38	0.32	0.51
	(0.48)	(0.47)	(0.50)
% Other Family Member	0.44	0.52	0.48
	(0.50)	(0.50)	(0.50)
% Not a Family Member	0.07	0.09	0.01
	(0.25)	(0.28)	(0.08)
% Worked on Farm	0.30	0.28	0.38
	(0.46)	(0.45)	(0.48)
Number of Obs.	1,014	651	12,693

Notes: Authors' calculations from LSMS-ISA data.

The descriptive statistics presented in this subsection provide a detailed overview of the relationships between demographics, labor, and migration. Sections 3.2.1 and 3.2.2 show the time path of labor demand and labor supply for Malawi. The substantial intra-annual fluctuations highlight the importance of allowing for possible heterogeneity in labor market conditions across phases of cultivation. In Section 3.2.3 we found that older household members do sufficient work to be counted as part of the household labor endowment. Adolescents and teenagers age quickly into the workforce. Table 4 in Section 3.2.4 shows the patterns of changes in labor endowments. Most households experience some change in endowment from one period to the next. Changes due to moving in, moving out, and aging are all commonplace. In Section 3.2.5 we confirm that there is no clear spatial pattern to the move-ins and move-outs, suggesting that our identifying variation is not driven by macro-level migration patterns related to local labor market conditions. Lastly, in Section 3.2.6 we see that in-migrants and out-migrants are primarily relatives of the household head who do similar levels of farm work as other members of the household. This is at least suggestive evidence that migration is not driven by the recruitment of members who are designated for farm work.

4 Results

In this section we present our empirical findings. The first subsection includes our main results, followed by a subsection with robustness checks.

4.1 Main results

In this subsection we first present the results of OLS regressions based on the empirical specifications in Section 2.3. For our main specifications we define both the household labor endowment and agricultural labor demand using an adult equivalence scale that treats everyone over age 15 as 1 worker, and children aged 11-14 as $(Age - 10) \times 0.2$ workers. One of the robustness checks in the next subsection uses a different definition of adult equivalents.

Table 11: Panel regressions of change in agricultural labor demand on change in labor endowment

Dependent variable: Δ Log of on-farm labor demand	Ethiopia (1)	Ethiopia (2)	Malawi (3)	Malawi (4)	Tanzania (5)	Tanzania (6)	Uganda (7)	Uganda (8)
Δ Labor endowment	0.631***		0.769***		0.614***		0.078	
	(0.198)		(0.172)		(0.203)		(0.084)	
Δ LE if increase		0.241		1.079***		0.680**		0.229
		(0.318)		(0.242)		(0.311)		(0.146)
Δ LE if decrease		0.933***		0.280		0.556*		-0.059
		(0.312)		(0.321)		(0.326)		(0.126)
Harvest (=1)	-0.190***	-0.200***	-1.441***	-1.371***	-0.372***	-0.403***		
	(0.032)	(0.044)	(0.031)	(0.042)	(0.032)	(0.041)		
Harvest $\times$ Δ LE	-0.017		-0.272***		-0.048			
	(0.124)		(0.096)		(0.102)			
Harvest $\times$ Δ LE if increase		0.038		-0.526***		0.115		
		(0.237)		(0.152)		(0.171)		
Harvest $\times$ Δ LE if decrease		-0.061		0.130		-0.200		
		(0.179)		(0.162)		(0.169)		
Weeding (=1)					0.070**	0.032		
					(0.028)	(0.036)		
Weeding $\times$ Δ LE					0.137			
					(0.093)			
Weeding $\times$ Δ LE if increase						0.338**		
						(0.160)		
Weeding $\times$ Δ LE if decrease						-0.050		
						(0.157)		
Observations	5688	5688	4204	4204	12069	12069	5695	5695
R^2	0.03	0.03	0.11	0.11	0.06	0.06	0.26	0.26
Adjusted R^2	0.02	0.02	0.11	0.11	0.06	0.06	0.26	0.26
F: weeding $\times$ change = 0 (p-val)					0.000			
F: weeding $\times$ increase = 0 (p-val)						0.000		
F: weeding $\times$ decrease = 0 (p-val)						0.112		
F: harvest $\times$ change = 0 (p-val)	0.002		0.009		0.004			
F: harvest $\times$ increase = 0 (p-val)		0.421		0.039		0.007		
F: harvest $\times$ decrease = 0 (p-val)		0.005		0.258		0.254		

Notes: Authors' calculations from LSMS-ISA data. Regressions account for household fixed effects by differencing at the household-year level. Standard errors in parentheses. Standard errors clustered at household level. Significance: *** 0.01, ** 0.05, * 0.1. Dependent variable is the inter-annual change in log of total labor demand. All regressions include controls for changes in log of cultivated acreage, changes in demographic shares, and year fixed effects. The excluded farming activity is all non-harvest work for Ethiopia; all non-harvest work for Malawi; planting and land preparation for Tanzania.

Columns 1, 3, 5, and 7 of Table 11 show the baseline estimates for each country, based on equation (11). These models assume symmetric effects from increases and decreases in the labor endowment. All specifications allow for heterogeneity across phases of cultivation, except for those from Uganda, where the data does not separate agricultural labor by phase. The first coefficient listed is that for the change in labor endowment between waves during the reference (excluded) phase. The reference phase is all pre-harvest activity in Ethiopia and Malawi, and planting / land preparation in Tanzania. Results for the non-excluded phases are found by summing the first coefficient with the interaction terms "Harvest $\times$ Δ Labor endowment" or "Weeding $\times$ Δ Labor endowment", as appropriate. The lowest panel of the table reports the results of F-tests for the significance of the combined effects in the non-excluded phases.

The baseline findings are positive and statistically significant for all countries other than Uganda. The levels are remarkably consistent across countries and agricultural activities. In Ethiopia the elasticity of labor demand with respect to the labor endowment is 0.63 pre-harvest and 0.61 during the harvest. For Malawi, elasticities are 0.77 pre-harvest and 0.50 in the harvest. In Tanzania, the elasticities are 0.61 during planting, 0.75 during weeding, and 0.57 during harvest. Only in Uganda is the relationship between the labor endowment and labor demand statistically insignificant and small in magnitude. However, as discussed in Section 3.1, the Uganda survey instrument censors the measurement of household farm labor in a way that introduces measurement error in labor demand that is correlated with the labor endowment. Thus, even though our identification is off of the changes in the labor endowment, we hesitate to make any definitive claims for Uganda.

The baseline findings in Table 11 suggest that the cross-sectional results in Dillon and Barrett (2014), which show a similarly strong correlation between labor endowments and labor demand, are not driven by household-level characteristics that vary little from year to year, such as managerial skill or preferences for working on one's own farm. In both a static and dynamic setting we can reject the null hypothesis of complete markets for Ethiopia, Malawi, and Tanzania. This pattern of findings differs from that in Indonesia. There, the cross-sectional finding of complete markets in Benjamin (1992) was over-turned by the panel estimation in LaFave and Thomas (2016).

The results discussed so far do not provide any evidence about the possible nature of the market failures that underlie non-separation. For that, we turn to the extended results in columns 2, 4, 6, and 8, which are based on specification (12). These models allow for possible asymmetric responses of labor demand to increases and decreases in the labor endowment. Recall the intuition from Section 2. If reductions in the labor endowment do not lead to statistically significant reductions in labor demand, we can reject the hypothesis of a labor shortage. Conversely, if increases in the labor endowment do not lead to statistically significant increases in labor demand, we can reject the hypothesis of a labor surplus. If labor demand responds to both increases and decreases in the labor endowment, then we can be less certain about the interpretation. Symmetrical responses could be due to failures outside the labor market, or to successful endogenous adjustment of labor endowments to labor market shortcomings.

Once we allow for heterogeneous responses to increase and decreases in the endowment, the pattern of findings is no longer the same across countries. Consider first the results for Ethiopia, in column 2. There, we see that the elasticity of labor demand with respect to the labor endowment is only significant for decreases in the endowment. This finding is the same in the pre-harvest and harvest phases. Clearly, the combined effect from column 1 is driven entirely by decreases in labor endowments. It is notable that the elasticity when the endowment falls is near unity, indicating that the departure of a worker from the household leads to a nearly proportional reduction in the amount of labor applied to the farm. For Ethiopia, this is consistent with a model of rural labor markets with insufficient workers.

For Malawi, the results are the opposite. In column 4 of Table 11 we see that the elasticity of labor demand to changes in the labor endowment is only significant when the labor endowment increases. In the pre-harvest period the value of the elasticity is essentially unity. In the harvest, it is estimated to be 0.55. Reductions in the labor endowment do not lead to statistically significant reductions in labor demand. This is strong evidence in favor of rural labor markets with insufficient employment opportunities, possibly due to downwardly sticky wages.

The results in column 6 indicate that for Tanzania the situation is less consistent across agricultural activities. During all three phases, labor demand responds to increases in the endowment. The elasticities of labor demand with respect to increases in the labor endowment are 0.68, 1.018, and 0.795 during the planting, weeding, and harvesting periods, respectively. During planting, the reference phase, labor demand also responds to decreases in the labor endowment, although the effect is smaller than that for increases and only marginally significant. In the weeding and harvest periods, decreases are not statistically significant. This intra-annual variation is suggestive of a market in which conditions vary, albeit slightly, over the course of the season. During planting, we cannot distinguish between types of labor market failures, because the results are consistent with excess or surplus labor accompanied by successful endogenous adjustment. Planting period results are also consistent with the presence of failures outside of the labor market. In the weeding and harvest periods we can reject the necessary conditions for a labor shortage, and results are consistent with non-clearing due to excess labor supply. One interpretation of this pattern is that households maintain their labor endowment to satisfy needs at planting time, when there is a correlated positive labor demand shock for a short time. They are then left with more workers than necessary at later points in the season.

Finally, for Uganda we note that the insignificant pooled effect is unchanged when we split by increases and decreases. With these data, labor demand does not respond to any type of change in the labor endowment.

4.2 Robustness and extensions

In this subsection we estimate three additional classes of models to test robustness and to extend the inference of the previous section. We first describe the three sets of additional tests, then present the results by country.

In a first set of robustness checks we exclude children under the age of 15 from the definitions of labor demand and the labor endowment. As discussed in Section 3.2.3, there is no single best way to allow the labor endowment to change as children age into the workforce. We use a binary cut-off at age 15 as an alternative to the gradual aging-in that we used for the main results.

The second set of robustness checks relates to how we deal with changes in cultivated acreage. Households can, and do, adjust the acreage under cultivation from one season to the next. This raises two issues. First, these adjustments are potentially endogenous to changes in the labor endowment, and thereby to labor market conditions. This is less concerning when identification stems from endowment changes that are exogenous to labor market conditions, as it does for one direction of adjustments. A second concern is that because desired acreage is likely a function of the labor endowment, underlying land market failures could drive the results. To deal with this concern, we instrument for changes in cultivated acreage using changes in owned acreage.

In a third and final set of extensions we modify our definition of the labor endowment to account for time spent away from home. In our main estimates we defined the labor endowment based on current household members or recent departures, as of the date of the survey. However, each survey also asks for how many of the last 12 months each person was living in the household. We use these data to re-estimate the labor endowment using the share of months spent at the household as weights. This approach has the potential to detect endogenous adjustments, if households recruit or send migrants during the cultivation season in response to labor market conditions. The main drawback is that we do not know the timing of entries and exits. By smoothing the time spent away across all months, even for cases where all endowment changes occurred before beginning cultivation, we risk introducing a new source of measurement error that is negatively correlated with real variation in the endowment.

We implement these additional tests only for specification (12), which allows for asymmetric effects. In the following subsections we discuss the results of these extensions separately for each country.

4.2.1 Robustness and extensions: Ethiopia

Table 12 shows results of the above extensions for Ethiopia. In columns 1 and 2 we see that neither modification of the labor endowment measure nor use of an instrument for cultivated acreage have a meaningful impact on results. The elasticity of labor demand with respect to the labor endowment is near unity, but only for decreases in demand. The results for the pre-harvest and harvest periods are the same. These findings are supportive of the conclusion that labor shortages rather than labor surpluses drive non-separation in rural Ethiopia.

The results of the third set of extensions, weighting household membership by time spent away, provide some additional nuance. In column 3 of Table 12 the effect of decreases in the

Table 12: Robustness checks and extensions: Ethiopia

Dependent variable: Δ Log of labor demand	(1)	(2)	(3)
	No Kids	IV Acreage	Time Away
Δ Labor endowment if increase	0.443	0.253	0.179***
	(0.294)	(0.319)	(0.047)
Δ Labor endowment if decrease	0.995***	0.946***	0.491***
	(0.296)	(0.312)	(0.089)
Harvest (=1)	-0.189***	-0.200***	-0.214***
	(0.041)	(0.044)	(0.035)
Harvest $\times$ Δ Labor endowment if increase	-0.101	0.038	0.010
	(0.203)	(0.237)	(0.024)
Harvest $\times$ Δ Labor endowment if decrease	0.018	-0.061	-0.143
	(0.169)	(0.179)	(0.096)
Observations	5688	5688	5688
R^2	0.029	0.026	0.036
Adjusted R^2	0.027	0.025	0.034
F: harvest $\times$ increase = 0 (p-val)	0.28	0.40	0.00
F: harvest $\times$ decrease = 0 (p-val)	0.00	0.00	0.00

Notes: Authors' calculations from LSMS-ISA data. Regressions account for household fixed effects by differencing at the household-year level. Standard errors in parentheses. Standard errors clustered at household level. Significance: *** 0.01, ** 0.05, * 0.1. Dependent variable is the inter-annual change in log of total labor demand. All regressions include controls for changes in log of cultivated acreage, changes in demographic shares, and year fixed effects. The excluded farming activity is all non-harvest work.

endowment is attenuated by nearly 50%, and the effect of increases is now positive and statistically significant, with an elasticity of 0.179. This finding is consistent with various interpretations. One is that this is a spurious result, due to measurement error that is correlated with labor demand. However, it may also be that this finding sheds light on the process of endogenous adjustment. As discussed in Section 2, successful recruitment of temporary members by otherwise labor-constrained households is one reason that we might see correlation between demand and endowments in both directions. This interpretation seems plausible, in combination with the strong and consistent results in favor of a labor shortage from all other specifications.

4.2.2 Robustness and extensions: Malawi

Table 13 shows the robustness and extension results for Malawi. Unlike the Ethiopia case, here we find complete consistency between the main results and those in Table 13. The correlation between labor demand and labor endowments is driven entirely by increases in the labor endowment. The elasticity is lower in the harvest period than the pre-harvest period. The lack of significant effects from decreases in demand allows us to reject the possibility of a labor shortage. Non-separation in Malawi is driven largely by a lack of off-farm work opportunities.

Table 13: Robustness checks and extensions: Malawi

Dependent variable: Δ Log of labor demand	(1) No Kids	(2) IV Acreage	(3) Time Away
Δ Labor endowment if increase	1.122***	1.053***	1.022***
	(0.225)	(0.242)	(0.246)
Δ Labor endowment if decrease	0.301	0.251	0.248
	(0.293)	(0.326)	(0.249)
Harvest (=1)	-1.407***	-1.371***	-1.414***
	(0.038)	(0.042)	(0.041)
Harvest $\times$ Δ Labor endowment if increase	-0.407***	-0.526***	-0.368**
	(0.135)	(0.152)	(0.151)
Harvest $\times$ Δ Labor endowment if decrease	0.168	0.130	-0.075
	(0.151)	(0.162)	(0.113)
Observations	4204	4204	4204
R^2	0.111	0.108	0.109
Adjusted R^2	0.109	0.106	0.107
F: harvest $\times$ increase = 0 (p-val)	0.00	0.05	0.01
F: harvest $\times$ decrease = 0 (p-val)	0.15	0.30	0.53

Notes: Authors' calculations from LSMS-ISA data. Regressions account for household fixed effects by differencing at the household-year level. Standard errors in parentheses. Standard errors clustered at household level. Significance: *** 0.01, ** 0.05, * 0.1. Dependent variable is the inter-annual change in log of total labor demand. All regressions include controls for changes in log of cultivated acreage, changes in demographic shares, and year fixed effects. The excluded farming activity is all non-harvest work.

4.2.3 Robustness and extensions: Tanzania

Results for Tanzania are shown in Table 14. Findings are somewhat less clear, in this case. In Table 11 we found that while both increases and decreases in the labor endowment were associated with significant changes in labor demand during planting, only increases mattered in other periods. With changes to the definition of the labor endowment, in columns 1 and 3 of Table 14, we see that labor demand also responds to decreases in the endowment during the weeding period, though still not in the harvest period (see the results of F-tests in the lower panel of the table). When we instrument for changes in cultivated acreage using changes in owned acreage, labor demand responds to changes in the endowment in both directions in all periods of the year.

This result does not lend itself to a single interpretation. The effect of increases in the endowment on labor demand is clearly established across all specifications. Whether or not reductions in labor demand also significantly affect labor demand is less clear. If demand co-moves with changes in either direction, this could be evidence of successful placement of some out-migrants from households in markets otherwise characterized by surplus labor. However, it may also indicate that non-separation in Tanzania is driven by failures outside the labor market. We favor the former interpretation because it is more closely aligned with the full set of results across Tables 11 and 14, but we cannot rule out the latter.

4.2.4 Robustness and extensions: Uganda

The results of robustness checks and extensions for Uganda are reported in Table 15. Here we find no cause us to amend the main findings. Labor demand responds weakly to increases in the labor endowment in column 2, and weakly to decreases in the labor endowment in column 3. Yet there

Table 14: Robustness checks and extensions: Tanzania

Dependent variable: Δ Log of labor demand	(1)	(2)	(3)
	No Kids	IV Acreage	Time Away
Δ Labor endowment if increase	0.629**	0.700**	0.544**
	(0.272)	(0.283)	(0.273)
Δ Labor endowment if decrease	0.592**	0.416*	0.573**
	(0.287)	(0.226)	(0.288)
Weeding (=1)	0.047	0.051	0.049
	(0.033)	(0.034)	(0.033)
Weeding × Δ Labor endowment if increase	0.206	0.231*	0.202
	(0.127)	(0.124)	(0.127)
Weeding × Δ Labor endowment if decrease	-0.023	0.060	-0.024
	(0.145)	(0.076)	(0.146)
Harvest (=1)	-0.397***	-0.367***	-0.397***
	(0.039)	(0.041)	(0.039)
Harvest × Δ Labor endowment if increase	0.049	0.003	0.067
	(0.144)	(0.152)	(0.145)
Harvest × Δ Labor endowment if decrease	-0.134	0.133	-0.184
	(0.148)	(0.146)	(0.154)
Observations	12069	12069	12069
R^2	0.059	0.060	0.058
Adjusted R^2	0.058	0.059	0.057
F: weeding × increase = 0 (p-val)	0.00	0.00	0.00
F: weeding × decrease = 0 (p-val)	0.04	0.03	0.05
F: harvest × increase = 0 (p-val)	0.01	0.01	0.02
F: harvest × decrease = 0 (p-val)	0.10	0.01	0.16

Notes: Authors' calculations from LSMS-ISA data. Regressions account for household fixed effects by differencing at the household-year level. Standard errors in parentheses. Standard errors clustered at household level. Significance: *** 0.01, ** 0.05, * 0.1. Dependent variable is the inter-annual change in log of total labor demand. All regressions include controls for changes in log of cultivated acreage, changes in demographic shares, and year fixed effects. The excluded farming activity is planting and land preparation.

is no clear pattern, and all results for Uganda are subject to the caveat that measurement error in labor demand may attenuate all results.

Table 15: Robustness checks and extensions: Uganda

Dependent variable: Δ Log of labor demand	(1) No Kids	(2) IV Acreage	(3) Time Away
Δ Labor endowment if increase	0.079	0.242*	0.097
	(0.123)	(0.146)	(0.151)
Δ Labor endowment if decrease	-0.099	-0.025	0.249*
	(0.106)	(0.125)	(0.143)
Observations	5695	5695	5695
R^2	0.265	0.258	0.266
Adjusted R^2	0.264	0.257	0.265

Notes: Authors' calculations from LSMS-ISA data. Regressions account for household fixed effects by differencing at the household-year level. Standard errors in parentheses. Standard errors clustered at household level. Significance: *** 0.01, ** 0.05, * 0.1. Dependent variable is the inter-annual change in log of total labor demand. All regressions include controls for changes in log of cultivated acreage, changes in demographic shares, and year fixed effects.

5 Discussion

The findings in the previous section present a picture of labor markets in three of the four study countries that are far from fulfilling their goal of optimal resource allocation. The findings in Table 11 reveal two key patterns. The first is that the results for specification (11) , when a symmetric relationship between changes in labor endowment and labor demand is assumed, show a remarkably similar pattern across Ethiopia, Malawi, and Tanzania. The second, however, is that when we allow an asymmetric relationship between increases and decreases in labor endowment, the results reveal substantial heterogeneity across the countries. These patterns highlight both the value of allowing for heterogeneity in the treatment when possible, and also the risks associated with taking the results from one country as evidence for another.

The heterogeneity revealed in Table 11 also indicates that the policy implications of our findings vary between countries. In Ethiopia, our findings suggest a shortage of labor in rural areas. Surprising as this may sound, it is consistent with other recent evidence of rising rural wages and increased use of labor-saving technologies in agriculture (Duncan et al., 2016). Policies to promote further use of labor-saving technologies (e.g. herbicides, tractors) are well-suited to these conditions. In Malawi and Tanzania, our findings are consistent with the presence of disguised unemployment and under-utilization of labor in the countryside. Under these conditions, policies to promote job creation and stimulate aggregate demand are more likely to generate efficient resource allocation than policies that promote substitution of capital for labor in agriculture. In Uganda, we find no evidence in the panel of rural market failures. The implication is that rural markets are reasonably complete and competitive in Uganda, with prices that adjust to clear the markets for agricultural inputs and outputs.

This paper improves our understanding of the relationships between demographics, labor markets, and migration. Building on the theoretical model of the agricultural household, we showed how differentiated reactions to increases and decreases in labor endowments can reveal specific shortcomings in the labor market. Our findings provide a portrait of rural markets in sub-Saharan Africa that share some common features, but that in fact differ in important ways upon deeper analysis. Three of the four study countries share some degree of labor market failure. The policies to fix these failures, however, require context-specific solutions. The analysis presented here is an important step toward reducing rural poverty by identifying the right set of policies to address labor market failures in four major sub-Saharan African economies.

References

Barrett, Christopher B, Thomas Reardon, and Patrick Webb. 2001. "Nonfarm income diversification and household livelihood strategies in rural Africa: concepts, dynamics, and policy implications." *Food Policy*, 26(4): 315–331.

Benjamin, Dwayne. 1992. "Household composition, labor markets, and labor demand: testing for separation in agricultural household models." *Econometrica: Journal of the Econometric Society*, 60(2): 287–322.

Dillon, Brian. 2014. "Risk and Resilience among Tanzanian Farmers." University of Washington.

Dillon, Brian, and Christopher B. Barrett. 2014. "Agricultural factor markets in Sub-Saharan Africa: an updated view with formal tests for market failure." World Bank.

Dreze, Jean P, and Anindita Mukherjee. 1989. "Labour contracts in rural India: Theories and evidence."

Duncan, Alan J., Fantu Nisrane Bachewe, Kindu Mekonnen, et al. 2016. "Crop residue allocation to livestock feed, soil improvement and other uses along a productivity gradient in Eastern Africa." *Agriculture, Ecosystems & Environment*, 228: 101–110.

Fafchamps, Marcel. 1993. "Sequential labor decisions under uncertainty: An estimable household model of West-African farmers." *Econometrica: Journal of the Econometric Society*, 61(5): 1173–1197.

Feder, Gershon. 1985. "The relation between farm size and farm productivity: The role of family labor, supervision and credit constraints." *Journal of Development Economics*, 18(2): 297–313.

Gollin, Douglas, David Lagakos, and Michael E Waugh. 2014. "The Agricultural Productivity Gap." *Quarterly Journal of Economics*, 129(2): 939–993.

Harris, John R, and Michael P Todaro. 1970. "Migration, unemployment and development: a two-sector analysis." *The American Economic Review*, 60(1): 126–142.

Hart, Oliver D. 1983. "Optimal labour contracts under asymmetric information: an introduction." *The Review of Economic Studies*, 50(1): 3–35.

Jacoby, Hanan G. 1993. "Shadow wages and peasant family labour supply: an econometric application to the Peruvian Sierra." *The Review of Economic Studies*, 60(4): 903–921.

Kaur, Supreet. 2014. "Nominal wage rigidity in village labor markets." National Bureau of Economic Research.

Kochar, Anjini. 1999. "Smoothing consumption by smoothing income: hours-of-work responses to idiosyncratic agricultural shocks in rural India." *Review of Economics and Statistics*, 81(1): 50–61.

LaFave, Daniel, and Duncan Thomas. 2016. "Farms, families, and markets: New evidence on agricultural labor markets." Econometrica.

Lewis, W Arthur. 1954. "Economic development with unlimited supplies of labour." *The Manchester School of Economic and Social Studies*, 22(2): 139–191.

McCullough, Ellen B. 2015. "Labor productivity and employment gaps in Sub-Saharan Africa." *World Bank Policy Research Working Paper*, , (7234).

Morten, Melanie. 2015. "Temporary Migration and Endogenous Risk Sharing in Village India." *mimeo.*

Osmani, Siddiqur Rahman. 1990. "Wage determination in rural labour markets: the theory of implicit co-operation." *Journal of Development Economics*, 34(1): 3–23.

Reardon, Thomas. 1997. "Using evidence of household income diversification to inform study of the rural nonfarm labor market in Africa." *World Development*, 25(5): 735–747.

Skoufias, Emmanuel. 1994. "Using shadow wages to estimate labor supply of agricultural households." *American Journal of Agricultural Economics*, 76(2): 215–227.

Teal, Francis. 2011. "The price of labour and understanding the causes of poverty." *Labour Economics*, 18: S7–S15.

Young, Alwyn. 2013. "Inequality, the urban-rural gap and migration." *The Quarterly Journal of Economics*, 128(4): 1717–1785.