Bermuda Journal of Academic Research, 1:1, 24-36, (2026)

OPEN ACCESS www.bermudajournal.org

Family Comes First: How Agricultural Households Operating Under Market Failures Absorb a Labour Supply Shock

Ana Van Der Ree

University of Oxford

Keywords: *Agricultural household model, separation property, market failures, labour supply shock, Covid-19, Ethiopia, LSMS-ISA, rural labour markets*

Abstract

I use the agricultural household model's separation property as a framework to analyse how agricultural households absorb a large shock in a context where markets function poorly. Under complete markets, farm production decisions should depend solely on prices and technology, and are separable from household characteristics like the household's labour endowment. I show that Covid-19 caused a large increase to households' labour endowments in rural Ethiopian farm households. I use this finding to examine whether the labour supply shock was so large that imperfect labour markets were unable to fully absorb it, causing households to increase total labour utilisation on the farm in response. The results show that agricultural households were operating under market failures before and after the shock. I find no evidence that the labour input misallocation consequences of market failures were statistically larger after Covid-19. I test the propositions by estimating first-differences specifications for two distinct panels of farm households in rural Ethiopia, a pre-Covid panel (2011–2015) and a Covid panel (2018–2021) using LSMS-ISA data.

Corresponding author: Ana Van Der Ree *E-mail address:* anavdr@gmail.com
Received: Fall 2025 **Published:** Spring 2026

1. Introduction

Poor households constitute a large proportion of the population in low-income countries. Across Sub-Saharan Africa, many of the rural poor are reliant on farming for their livelihoods and consumption needs (Suri & Udry 2022). In these settings, labour is the household's primary endowment. A sudden shock to labour supply, therefore, prompts households to re-optimize their labour allocation decisions across tightly linked production, consumption and labour markets. Understanding how agricultural households absorb a shock to labour supply provides an important insight to the functioning of markets in rural settings.

Covid-19 introduced a sudden disruption to mobility and migration patterns, economic activity and health attitudes. Consequently, many rural households retained members at home that may have otherwise migrated away, or received returning members back to the home. This represents a large shift outwards of the labour supply curve in rural labour markets. If rural labour markets function imperfectly and wages do not adjust fully, this results in an above-clearing market wage. Wage 'stickiness' has potentially important consequences for employment and output, with a wage set above market-clearing level leading to higher unemployment, in theory.[1] Kaur (2019) finds that nominal wages do not adjust downwards in response to (negative demand) shocks in a setting for casual agricultural labour. Downwards wage rigidity causes employment reductions borne by poorer individuals—landless and small landholders—and when rationed out of the market, landholders increase labour supply to their own farms (Kaur 2019). This suggests there is vast potential for the labour supply shock from Covid-19 leading to "disguised unemployment" in agricultural settings if wages are rigid. In this paper I use the agricultural household model as a framework for a set of labour market institutions through which to analyse how rural Ethiopian households absorb a large positive shock to labour supply.

The agricultural sector is an interesting context to study labour (mis)allocation because of evidence that low agricultural productivity and high share of agricultural employment can explain—at least in part—a large fraction of cross-country dispersion in aggregate productivity (Gollin & Udry 2021; Restuccia et al. 2008). Therefore, understanding the labour allocation decisions of the agricultural household is also key to understanding what drives productivity growth. Since the demand for agricultural labour is highly seasonal and many African countries display low population density, agricultural labour markets may be highly constrained, making it difficult for farms to source labour at the right times (Suri & Udry 2022). On the other hand, farmers may be "poor but efficient" and learn not to repeatedly misallocate their (few) resources at their own peril (Schultz 1964). This would suggest that labour markets

[1] While this is a prediction of classic labour demand-supply interactions, evidence for it is mixed in the literature. There is little direct evidence that wage rigidity actually affects employment in the labour market in any setting, with Card (1990) as a notable exception.

function smoothly for the most part. The completeness of labour markets will determine how households absorb a labour supply shock, thus motivating my investigation based on a test for market completeness.

Labour markets in rural settings are often interconnected with markets for consumption, output and credit, as small family-run farms produce agricultural goods both for consumption and sale. The agricultural household model captures this interdependence (Singh et al. 1986). Assuming markets are complete and the household is a price taker in all markets, farm production decisions are treated as if the utility-maximising household operates a profit-maximising firm. This implies input choices only depend on prices and farm technology. The model yields a testable prediction of separation: consumption decisions are made in the second stage after any household income from farm profits has been pinned-down in the first stage. As the model's prediction relies on the assumption that markets are complete, rejection of the separability hypothesis is evidence for markets failures. This has prompted a literature which tests for complete markets using the separation hypothesis (Benjamin 1992; Dillon & Barrett 2017; Dillon et al. 2019; LaFave & Thomas 2016; Udry 1996).

Tests of the agricultural household model's assumption are well rehearsed in the literature. I aim to use the model's prediction as a framework to analyse how agricultural households absorb a large shock in a context where markets function poorly.

I employ a test for separation based on labour demand on the farm, as follows. If the family farm behaves as a profit-maximising firm in complete and competitive markets, then optimal production decisions—such as input utilisation—do not depend on any household characteristics—such as labour endowment. I regress the demand for farm labour on household labour endowment to test if a relationship exists. If labour endowments affect the farm's labour input decisions, resources are allocated sub-optimally; households are operating under market failures.

My first proposition is that Covid-19 caused a large positive shock to household labour endowment because mobility restrictions and business closures prompted members to stay or return home. My second proposition is that the outwards shift of the labour supply curve was so large that the shock was not fully absorbed, so households increased total labour utilisation on the farm in response. This is a failure to separate household characteristics from production choices. My third proposition is that before the shock, markets were already functioning poorly and the large supply shift led to a deterioration in the consequences of market failures. My fourth proposition is that more 'sophisticated' households were better able to absorb the shock by placing household members in the formal wage market.

I test the propositions on two panels of households—before and after Covid-19—from rural Ethiopia using 10 years of agricultural household surveys spanning from 2011–2022. Ethiopia's Living Standards Measurement Surveys and Integrated Surveys on Agriculture (LSMS-ISA) consists of an initial panel of households that was tracked from 2011–2015 and a different refreshed panel tracked from 2018–2021. Longitudinal data is key for the test to ensure results are not contaminated by unobserved heterogeneity between households. Better management practices or soil quality, for instance, are related to farm labour utilisation and can be misattributed to changes in household labour endowment if these unobserved characteristics are not controlled for. Using variation within the household across time improves the identification of the test because such heterogeneity is held constant. A limitation of Ethiopia's LSMS-ISA data is that the same households were not observed continuously before and after the Covid shock, which means observed differences may partly reflect cohort differences. But it also has a key advantage: it enables the baseline surveys (2011 and 2018) to represent two groups that face similar labour-market conditions and life-cycle stages. I explore this further in the data description.

Firstly, I show the panels of households are similar at baseline across a range of household and farm characteristics, confirming they are comparable for investigation. This includes that average household labour endowment is almost exactly the same at baseline. Secondly, I show that households exposed to Covid-19 experienced a larger increase in labour endowment on average between the baseline and follow-up, supporting my first proposition. This is primarily driven by a greater retention of ageing children in the household. Thirdly, I find evidence to reject complete markets both before and after Covid-19. Households increase labour utilisation on the farm in response to increases in their labour endowment, consistent with my second proposition that households were unable to fully absorb the shock. The elasticity almost doubles for households exposed to Covid-19 compared to the earlier panel, suggesting market failures led to greater labour misallocation consequences for farm households who experienced the shock. When testing this formally in a pooled sample including all households and a binary indicator if the household was part of the Covid panel, the Covid-differential in the elasticity fails to reach statistical significance. This is qualitatively consistent but quantitatively inconsistent with the third proposition. I explore whether more sophisticated households with experience in the formal wage market absorb the shock better by using their market experience to adjust on that margin. I find no significant evidence for heterogeneity in household responses.

The rest of the paper proceeds as follows. In Section 2, I discuss the implication of complete markets (Subsection 2.1), the labour effects of Covid-19 in Ethiopia (Subsection 2.2) and the literature on the separation test (Subsection 2.3). The Agricultural Household Model is presented in Section 3. Section 4 outlines the identification strategy and Section 5 summarises the data. The key results are discussed in Section 6 followed by a heterogeneity analysis in Section 7. I conclude in Section 8.

2. Background

2.1. Complete markets

A market is complete when arbitrary units of the good can be bought at a specific price and there are no information asymmetries or transaction costs. Evidence for incomplete markets suggests agricultural households are operating under market failures and are constrained in achieving profit-maximisation. This may be contrasted with other explanations for their behaviour—such as operating in non-competitive markets, entirely missing markets, or behavioural constraints such as social norms—which would require different policy (Dillon & Barrett 2017). Markets can be incomplete for several reasons: asymmetric information leading to family members being used on the farm more than hired workers; imperfect substitutability of family and hired labour; heterogeneity in the price paid to farmers and the market price of the good; rationing from limits on the number of hours family members can work outside the household; weak farm contract enforcement.

It is not possible with a general test for separation to identify the sources of market failures. Specifically, non-separation is evidence for market failures in at least two factor or output markets. LaFave and Thomas (2016) point out this is not always a disadvantage due to the nature of rural settings having inter-linked markets. In fact, no single binding constraint explains low African agricultural productivity and it is likely that a combination of market failures operate in multiple markets simultaneously (Suri & Udry 2022). Therefore, while I implement a test based on labour utilisation on the farm, it is not strictly evidence for labour markets failures. If credit markets function poorly, credit constraints can limit a household's ability to hire labour even if labour markets are complete.[2] However, because of the nature of the Covid shock causing a large shift in labour endowments, I direct my attention to labour markets as this is primarily where the shock will be absorbed.

2.2. The Covid-19 Pandemic in Ethiopia

The Covid-19 pandemic hit Ethiopia hard, ranking it 6th in Africa for most total infections.[3] Even if the health shock was modest (arguably, it was not), the economic shock was large. The government imposed a six-month long State of Emergency in 2020 which banned public meetings, restricted travel and forced school closures (Ayele & Fessha 2021). In a comprehensive policy brief, D. Harris et al. (2021) find that employment and incomes declined across the board during the pandemic and, while they bounced back by 2021, much of this represented a move to lower-quality and temporary jobs. They observe that casual labourers were less likely to find jobs and recover their incomes compared to small business owners even long after the onset of the pandemic. 'Young Lives' also documents a shift to working in the agricultural sector and towards self-employment (Ford et al. 2021). Part of this may be explained by lay-offs in factories and export-oriented production. Garment worker lay-offs in the Hawassa Industrial Park—one of several major public industrial infrastructure projects—resulted in many workers having to migrate out of the city, which could be reflective of a broader return of industrial workers to rural areas (D. Harris et al. 2021). Although schools re-opened, by the end of 2020 around two-fifths of 19-year-olds still in education had not engaged in any form of learning since school closures, with those from the poorest households and rural areas affected most (D. Harris et al. 2021).

The descriptive findings together support the proposition that the pandemic caused a positive labour supply shock to agricultural households. The moderate shift towards agricultural "low-quality" work may also have disproportionately affected agricultural households by further constraining rural labour markets. Nolte et al. (2022) find that Ethiopian respondents state the reduced availability of hired labour as an important factor in preventing the pursuit of 'normal' agricultural activities after Covid-19, suggesting agricultural labour markets exhibited market failures.

2.3. Previous Literature

The agricultural household model was first formally developed in the foundational work by Singh et al. (1986) while Bardhan and Udry (1999) brought it to the dynamic setting. The core analytical concept of integrating consumption and production decisions was influenced by Sen (1962)'s and Chayanov (1991)'s observations that a family farm's production volume is likely related to its family size. In a seminal paper, Benjamin (1992) found no evidence for incomplete markets in rural Indonesia by estimating a cross-sectional relationship between farm labour demand and household composition. A source of concern in the cross-section is the potential bias arising from unobserved variation across households in the sample. LaFave and Thomas (2016) return to the same setting two decades later with extremely granular panel data to address the potential bias. They exploit within-household variation to rule-out fixed unobserved heterogeneity between households driving the results and come to the opposite conclusion of Benjamin (1992).

Udry (1996) brings the test to the African context to bridge the knowledge gap on characterisations of African rural markets. He uses plot-level panel data from Burkina Faso and Kenya to test for separation and rejects market completeness in both settings. The lack of rich data on plot quality limits Udry (1996) as he cannot rule-out unobserved variation in plot characteristics as confounders. This concern is echoed by more recent findings in Gollin and Udry (2021) which point out that unobserved heterogeneity and measurement error can lead to misallocation being overstated by three-fold. This signals to a broader point that the margin of productive inefficiency in Africa is possibly inflated in the literature, which motivates this investigation.

Dillon and Barrett (2017) use cross-sections of LSMS-ISA data to find evidence for market failures, estimating elasticities of household labour endowment to farm labour demand within the range of 0.32–0.53 for Sub-Saharan African countries including Ethiopia. A few years later, Dillon et al. (2019) return to the exploration by exploiting LSMS-ISA panels and, similarly, find no evidence for separation. Dillon et al. (2019) use the first two waves of Ethiopia's LSMS-ISA surveys. To my knowledge, the separation property has not been previously used to investigate how farm households operating under market failures adapt their farm labour allocations in response to a large labour supply shock.

3. Theoretical Framework

3.1. A dynamic model of the agricultural household

The agricultural household model captures the interdependence of consumption and production decisions by illustrating a profit-maximising family-run farm (firm) that operates within a household maximisation problem (Singh et al. 1986). This replicates the characteristic common in many developing countries in which households produce crops both for consumption and market sale. The model shows that when the household is a price taker for all goods which it consumes and produces and markets are complete, optimal farm inputs are determined independently of leisure and consumption choices. Then, given the income derived from production, consumption decisions are made. Therefore, decision-making is sequential.

In this simplified version of the model the household gets utility from consumption of market and agricultural goods (denoted for simplicity as one good, c_t) and leisure, l_t. The infinitely lived household maximises the expected present discounted value of current and future utility:

$$\max \ \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t U(c_t, l_t)\right] \tag{1}$$

The household can produce the good/crop on its farm using inputs land, A_t, and labour, L_t, according to its production

[2] Without a functioning credit market, the household's shadow value of cash (used to pay hired labour) is higher than the interest rate. The households will hire less labour than optimal.

[3] When ranking in terms of most infections per population, Ethiopia falls to the 37th in Africa.

technology:

$$F(L_t, A_t) \tag{2}$$

Labour used on the farm can be sourced from family labour, L_t^F, or hired labour from the marketplace, L_t^H. Land inputs can come from the farm's own land, A_t^F, or rented-out land, A_t^M:

$$L_t = L_t^F + L_t^H \tag{3}$$

$$A_t = A_t^F + A_t^M \tag{4}$$

Although in practice other inputs such as fertilizer will be used and the farm will produce various crops, these features are abstracted away to simplify the analysis and do not change the model's predictions. The critical part of the model is that households face complete markets: they take as given the price of the good, p_t, land rental price, a_t, and wage for market work, w_t. Importantly, this implies the market wage is the shadow wage for farm work, which is therefore the same shadow wage for hired and family labour. This underpins the assumption that hired and family labour are perfect substitutes. The model also abstracts away any market frictions, such as matching frictions in hiring, time lags between seeking and obtaining a job, labour supply shortages, or limited availability of jobs. This has implications for the assumptions regarding how Covid-19 affected labour markets, which I unpack in light of my results in Section 6.

The household's time endowment, E_t^L, is divided between working on the farm, L_t^F, working in the market, L_t^M, and leisure, l_t. Similarly, the household's land endowment, E_t^A, can be used for their own farming activity, A_t^F, or rented out to other households, A_t^M:

$$E_t^L = L_t^F + L_t^M + l_t \tag{5}$$

$$E_t^A = A_t^F + A_t^M \tag{6}$$

Because this is a dynamic model, farmers can smooth consumption across time by borrowing or saving in complete credit markets with interest rate r_t. This comes together in the inter-temporal budget constraint, where Ω_t is the household's wealth at the beginning of period t:

$$\Omega_{t+1} = (1+r_t)\Big[\Omega_t + w_t(L_t^F + L_t^M) + \{p_t F(L_t, A_t) - w_t L_t - a_t A_t\} - p_t c_t\Big] \tag{7}$$

Wealth in the next period, Ω_{t+1}, is equal to the interest earned on wealth in the current period net of wage earnings from labour work $w_t(L_t^F + L_t^M)$, farm profits (total revenue less total costs: $p_t F(L_t, A_t) - w_t L_t - a_t A_t$), and consumption expenditure ($p_t c_t$). Substituting (5) we obtain an expression that includes l_t:

$$\Omega_{t+1} = (1+r_t)\Big[\Omega_t + w_t(E_t^L - l_t) + (p_t F(L_t, A_t) - w_t L_t - a_t A_t) - p_t c_t\Big] \tag{8}$$

With complete and competitive markets, households choose the quantities of c_t and l_t which maximise (1) subject to the budget constraint (8), resource constraints (3)–(6), and non-negativity constraints:

$$l_t,\ c_t,\ L_t^M,\ L_t^F,\ L_t^H \geq 0 \tag{9}$$

In the first stage, the household chooses inputs L_t^* and A_t^* to maximise expected farm profits. In the second stage, the household maximises utility, conditional on expected income from the farm. This sequentiality is what makes the model recursive. The solution to the problem is, therefore, characterised by:

$$\Pi^* = \max\ \mathbb{E}\left[p_t F(L_t^*, A_t^*) - w_t L_t^* - a_t A_t^*\right] \tag{10}$$

$$A_t^* = A_t(a_t, w_t, p_t) \tag{11}$$

$$L_t^* = L_t(a_t, w_t, p_t) \tag{12}$$

where total labour demand on the farm, L_t^*, and total land demand on the farm, A_t^*, depend only on the prices of inputs, crop price, the farm's productive technology and shadow prices. Crucially, any household preferences or characteristics that affect consumption decisions do not affect production choices. This leads to a neat result: while it initially appeared that the agricultural household solved a joint problem where consumption and production decisions were intertwined, with complete markets, production decisions are separable from consumption choices. The opposite is not true: consumption still depends on the amount of profits realized from production as this enters the budget constraint. If all markets but one are complete—for instance, if land cannot be traded—then separation still holds because relative prices can adjust to accommodate the non-tradable good (Feder et al. 1985).

The separation result has a testable prediction that, under complete and competitive markets, farm labour demand should not be affected by household labour endowment. Intuitively, if a household's labour endowment falls—for instance because a young male leaves to find work in the city—demand for labour on the farm should remain unaffected because family labour is substitutable with hired labour at the same shadow wage. Similarly, excess supply of family labour can be sold in the market at the same wage so should have no impact on production.

3.2. How Covid-19 affects the separation prediction

Proposition I: My first proposition is that Covid-19 caused a large positive shock to household labour supply. There are two possible mechanisms:

- *Retention:* Household members who otherwise may have chosen to leave to urban areas for employment or education delay departure due to uncertainty. For instance, young men in the household who would have migrated for better work opportunities stay in the household longer.
- *Return:* Members living outside the household return home during the pandemic. For instance, adults moving back with family for food security reasons, and youth returning from urban areas due to school or business closures.

This represents an increase in the household's labour endowment. If markets clear, additional labour will be fully absorbed by the labour market and will have no effect on input decisions. For instance, extra family labour will be sold to the local market at the shadow wage due to the perfect substitutability of labour, or hired farm labour will contract and be replaced by the new family member. If markets do not clear, the positive shock to labour endowment may not be fully absorbed. Farms may respond by expanding total labour utilisation on the farm.

Proposition II: My second proposition is that markets are not fully able to absorb the labour supply shock. The outwards shift of the labour supply curve in the agricultural labour market may be so large that wages do not fully adjust if wages are 'sticky'. Therefore, households respond to the increase in their labour endowment by employing a higher quantity of total labour on the farm. For instance, farms may add more family labour to their input mix without reducing hired labour inputs even though this need not be optimal for L_t^*. This is a failure to separate household characteristics from production choices. In other words, family may come first.

Proposition III: There are two possibilities regarding the state of market completeness: the first scenario is that before the Covid shock, markets were clearing. The increase in labour supply prevented what was previously characterised as market completeness. More probably, however, before the shock, markets were already functioning somewhat poorly. My third proposition is, therefore, that the large supply shift led to greater consequences from market failures. This is equivalent to the intensive margin of non-separation.

In addition to the Propositions, I provide examples for three other mechanisms via which Covid-19 can have affected households and markets:

Changing the household's demographic composition. Returning household members bring with them new spouses or children into the household; increased illness or death in the household; elderly individuals encouraged to move in with their family or children. Such changes in household demographics can increase care-taking burdens, particularly for women, and crowd out time availability for farm work. It should have no effect on farm labour demand if the model's assumptions hold because only technology, prices and inputs determine L_t^*. Under separation, the care burden is accounted for in the utility function $U(\cdot)$ via a fall in leisure and the household would compensate by adjusting its consumption-leisure trade-off. The adjustment to the care shock happens in the second stage after profit-maximising decisions have been pinned down in the production process. If instead, care burdens shift labour demand on the farm, then demographic factors are "leaking" into production decisions and separation is violated.

Increasing labour market frictions. Infection risk or increased scepticism about Covid-19 increases transaction costs of hiring labour onto the farm; travel restrictions reduce supply of hired labour; lower availability of non-agricultural jobs; information frictions from difficulty of knowing who is available or eligible to be hired onto the farm. The model technically assumes there are no frictions in the labour market. In practice, frictions are likely to exist in rural settings. Covid-19 also most likely intensified frictions. This may limit the suitability of the model's assumptions for a shock context like Covid-19, which I explore further in light of the results in Section 6. Frictions may contribute to the "family comes first" effect by making it more difficult to hire labour or more desirable to use family labour on the farm.

Increasing credit constraints. Tighter household budgets or credit, for instance due to lower remittances from urban areas, make hiring costly even if labour is available; increased perception of risk by local banks or moneylenders lowers loan availability or increases interest rates. The various market failures may strengthen one another's impact on farm labour utilisation.

I do not explicitly test for the effects of changing demographic composition, market frictions or credit constraints on farm labour demand in my identification. However, they serve as an important backdrop to the results and can shed light on additional mechanisms operating in the household.[4]

4. Identification

4.1. Separation test

I test for separation in the agricultural household model through the household's demand for labour in (12). Following Dillon et al. (2019), I estimate the following empirical specifications for labour utilized by the farm household h at time t:

$$L_{ht} = \beta_0 + \beta_1 E_{ht} + \gamma A_{ht} + \gamma_2\, demog_{ht} + \eta_t + \kappa_{ea} + \epsilon_{ht} \tag{13}$$

$$\Delta L_{ht} = \beta_0 + \beta_1 \Delta E_{ht} + \gamma \Delta A_{ht} + \gamma_2 \Delta demog_{ht} + \epsilon_{ht} \tag{14}$$

Equation (13) is a pooled cross-section model where the variation is across households in the sample and equation (14) is a household fixed effects model in first differences where the variation is within-households over time. The latter is the main specification of interest. The variables L_{ht}, E_{ht} and A_{ht} enter as logs, following the literature (Dillon et al. 2019; LaFave & Thomas 2016).

L_{ht} is the total number of person days of labour used on the farm h in period t and E_{ht} is the household's labour endowment in period t. Labour demand may be related to other farm characteristics, such as land size, which is captured by the cultivated acreage of the household in each period, A_{ht}. Farm labour demand may be affected by other household demographics, such as the share of working-age or elderly adults. $demog_{ht}$ accounts for this as a vector of controls including the share of prime-age males, prime-age females, elderly males and elderly females in each household h in each period t, following Dillon et al. (2019).[5]

η_t represents time fixed effects at the survey wave/year level. This absorbs any wave-specific shocks—such as a bad harvest year, national policy or general macro trends in labour use—which influence labour demand across all households. In equation (13) η_t enters as levels and in (14) it drops out because all time-fixed effects are absorbed by first differencing. κ_{ea} are village fixed effects that absorb any time-invariant village-specific characteristics, such as elevation or distance to urban cities, in order to compare households within the same village (enumeration area, EA) by holding this heterogeneity constant. Village fixed-effects disappear in first-differencing.[6] ϵ_{ht} is an error term. We might expect the error terms of households in the same village to be correlated if there are village-level shocks that affect labour demand. To tackle the intra-cluster correlation, I allow errors in each EA-wave combination to be correlated with one another but independent of other EA-wave cells by clustering standard errors at the EA-wave level.

Under the null that separation holds and markets clear, household labour endowment E_{ht} will have no impact on farm

[4] I do not explore heterogeneous responses to the labour shock across gender lines due to time constraints. There is evidence that gender-differentiated plots matter in Western Africa, while such differentiated allocations are not as marked in Eastern Africa. An interesting addition would be to explore whether return of males compared to females affects farm labour demand. Dillon et al. (2019) find no such evidence in the same context of agricultural Ethiopia.

[5] Shares are calculated as the number of individuals in that category divided by the total number of individuals in the household. Prime-age is defined as 15–60 years old, elderly is defined as 61 years and older.

[6] Time-varying village spillover effects may be a concern. There are possible general equilibrium effects if many households in a village experience return migration, such as wage depression or lower hiring from other villages. While I do not explicitly address this, it remains as a valid consideration to keep in mind for the results.

labour demand and β_1 will be zero. However, identification in the pooled cross-section is contaminated by unobserved household-specific characteristics such as managerial skills or preference for working on one's own farm. For instance, a more productive farmer will demand more labour on her farm.[7] These characteristics tend to be time-invariant or slow-moving across time. As these characteristics are unobserved, in (13) they will be captured by ϵ_{ht} and bias the estimate if they are correlated with labour endowment. One solution is to add these characteristics into the specification as in Benjamin (1992). However, as they are unobservable and difficult to measure, this has often been a struggle in the literature that uses cross-sectional data. Specification (14) addresses this by first-differencing out all time-invariant heterogeneity in each household, allowing identification to come solely from within-household changes in labour endowment across time. This is the key advantage of using panel data.

Including the vector $demog_{ht}$ helps isolate the effect of labour endowment itself by holding constant the underlying demographic composition of the household. The number of children or elderly individuals in the household may impact farm labour demand in practice if separation fails. Therefore, by omitting this vector from the specification, it risks attributing the effects that actually come from household structure (i.e. who is available within the household) to labour endowment (i.e. how many workers are available within the household). Separating these effects leads to a cleaner test on ΔE_{ht} with a clear mechanism.

4.2. Identifying the effect of the Covid-19 shock

To examine how households absorbed the Covid-19 shock to labour endowment I estimate the specifications in two different time periods:

(i) **Panel I:** a pre-Covid period representing counterfactual market conditions,
(ii) **Panel II:** a Covid period representing market conditions after the shock.

Comparing the estimated elasticity across the panels suggests whether the shock amplified non-separable behaviour. This approach provides a comparative test across the two regimes revealing how households absorbed the shock—however, it does not constitute a formal differences-in-differences design as the same households are not observed continuously across all four waves. A larger elasticity in the Covid period suggests greater separation failure.

However, the comparison is treated cautiously because Covid-19 may have altered market conditions—such as by introducing market frictions—in ways that are important for the cleanness of the separation test. I discuss this in Section 6. The two panels, of course, differ in more ways than only Covid exposure—other shocks or cohort effects could also drive differences in the estimated slopes. I examine differences in the baseline data for each of the panels in the next section and find that the data suggests the two panels are comparable across a range of farm and household characteristics.

[7] A more skilled manager extracts more output from each worker-day (better task allocation, fewer mistakes) so that each additional worker has a higher marginal product. Under profit-maximization, the manager demands labour up to the point where the marginal product of labour equals its (shadow) wage cost. Thus, a more productive farmer will demand a higher quantity of farm labour, ceteris paribus.

5. Data

5.1. Overview of the data

I test the model's predictions using four waves of comprehensive household data from Ethiopia's LSMS-ISA, surveys conducted by Ethiopia's Central Statistics Agency and the World Bank. They collect extremely rich survey data on over 3000 nationally representative rural households on a wide range of topics necessary for the separation test including farm production, consumption, labour and individual demographics. The granularity of the agricultural variables and ability to track individuals across waves are key advantages of this dataset.

The surveys were conducted in 2011, 2013, 2015, 2018, 2021.[8] The panel of households was refreshed in the 2018 wave to ensure the sample was representative for Ethiopia's eleven regions and for rural and urban areas. Therefore, the 2018–2021 households represent a new panel, not a follow-up of the 2011–2015 sample.

Ideally a differences-in-differences design would follow the same households through all waves, so that pre- versus post-Covid comparisons come entirely from within-household changes. Using data from the two different panels might seem like a disadvantage because changes observed between panels could in principle arise from cohort or survey differences. But it also has a key advantage: it enables the baseline surveys (2011 and 2018) to represent two groups that face similar labour-market conditions and life-cycle stages. If the same individuals had been tracked from 2011 to 2021, the 18–24 year-olds in the first wave would be 28–34 years old in the fifth wave, which encapsulates an entirely different life-cycle labour decision.[9] In contrast, by comparing differences between 2011–2015 to differences between 2018–2021, the two panels each include (i) the baseline survey for their sample, and (ii) span a relatively similar interval of three or four years.[10] Comparison is corroborated by the feature that both panels are nationally representative and households are selected from the same regional states.

Since the separation test relies heavily on changes in household size, it is important that selective attrition does not contaminate inferences. For the pre-Covid panel, 91.2 percent of the rural households from 2011 were re-surveyed in 2015, representing relatively low attrition. The Covid panel suffers from a larger response rate problem: the Tigray conflict in northern Ethiopia between 2020–2022 resulted in the Tigray region being excluded from the 2021 survey. This meant the 398 rural households (or 54 rural enumeration areas) from the Tigray region were dropped in the Covid-panel. To minimise the potential influence of this discrepancy on the results, I run all key regressions of the earlier panel both with and without Tigray households and find that the results remain unaffected (see Addendum A1).

5.2. Variables and survey details

The questionnaires begin with a household roster including basic demographic data (age, sex) on the individuals living in the household. I use this list of household members to create a variable for the labour endowment of household h in year t,

[8] The surveys are conducted over the course of 9–12 months starting in September, so technically the waves are 2011–2012, 2013–2014, 2015–2016, 2018–2019, 2021–2022. For simplicity, I will refer to the waves as the year in which the surveying began.

[9] For instance, households from the 2011 sample that have children of the relevant age for this investigation would have moved beyond that phase by 2021 and been replaced by characteristics that were not relevant for the test in 2021 but were in 2011.

[10] The 2021 wave was likely originally planned for 2020 following the consistent 2-year increment in the LSMS surveys but delayed due to the pandemic. Therefore, the interval between waves 4–5 is three years.

E_{ht}, by counting the number of working-age individuals. I had to choose an age cut-off at which someone enters or exits the working-age. Following Dillon et al. (2019), I include all adults (including senior citizens) in the count for E_{ht} and allow children to gradually age into E_{ht} using an equivalence scale: 10-year olds and younger do not count, 11-year-olds count as 0.2 adults in the workforce, 12-year-olds as 0.4 adults, and so on where a 15-year-old counts as one full adult. The results are robust to different specifications of labour endowment including without the phase-in of children. I also use the household roster to construct the vector of demographic controls, $demog_{ht}$, which is the share of prime-age (15–60 years-old) and elderly (61+) males and females. To construct the cultivated acreage, A_{ht}, I use the area reported by plot holders in the farm or, in cases where area was not recorded, I use the GPS-measured area from the survey.

The agricultural survey modules are granularly measured at the holder-parcel-field-crop level. There is evidence that within-household allocation of labour inputs across plots, including across gender lines, matters for misallocation and productivity (Duflo & Udry 2004; Goldstein & Udry 2008; Kilic et al. 2015; Udry 1996). Nevertheless, while acknowledging this literature, I aggregate all variables to the household-level because this is the relevant unit for my analysis; variation at the plot level is not necessarily informative for my predictions about farm behaviour in response to the Covid-19 shock. It is also difficult to link disaggregated plots across survey years as these units are sometimes noisily measured, so aggregating also smooths these measurement errors (Gollin & Udry 2021).

Because the analysis explores the household's demand for labour on the farm, my sample consists of only land-holding households that cultivated in at least one survey year—i.e. excluding households that display no demand for labour on their farm ever.

Accurate measurement of labour used on the farm is crucial for this research. The surveys distinguish between two farming seasons—cultivation (also known as post-planting) and harvest—and collect comprehensive data on family and hired labour in each season on the farm.[11] Surveys conducted in informal settings have struggled with accurate measures of labour supply: Gollin and Udry (2021) emphasize the prevalence of measurement error in similar settings exploring misallocation in African agriculture.

While LSMS-ISA surveys measure the quantity of hired workers, the survey structure may overlook nuances in how labour is exchanged in Ethiopian informal settings. In an analysis of his Ethiopian fieldwork, Aspen (1993) discusses how financially-unremunerated exchange labour for agricultural tasks like weeding and harvesting is common in rural Ethiopia, known as *däbo* or *wänfäl*. This involves a mutual exchange where the host serves food and beer in exchange for labour, and participation is expected to be reciprocal. Importantly, exchange labour is not solely a production decision: participation is also motivated by the desire to remain invested in the social network and feel a sense of loyalty. The LSMS-ISA do not clarify whether mutual exchange labour is included in the respondent's count of hired labour. I find that between 13–16% of individuals in the samples said they provided free exchange labour to, on average, 4–5 other households. This indicates exchange labour is prevalent. More detailed data tracking households that received exchange labour, or the number of person-days exchanged, is lacking.

5.3. Summary statistics

Table 1 presents summary statistics of the data for each year. Only households tracked between survey waves are included so that the baseline characteristics represent the samples used in the empirical tests. Tigray households are excluded in both waves of Panel II.

The most important statistic of Table 1 lies in the first row: in both panels the average household has an endowment of roughly 3 working age adults at baseline (with significant variation across households). This is one measure that validates the comparability between panels for this investigation. We also see in 2021, average household labour endowment is considerably larger than at baseline, which is suggestive of Proposition I.

The pattern is robust to different measures of household labour endowment that exclude children. Average shares of demographic groups in the household are balanced across panels. Farm characteristics are reported in the table. In baseline years the average farm cultivates roughly 6 acres of land, of which around two-thirds are owned. What remains consistent across the panels is that the vast majority (80–90 percent) of labour is provided by family members. There is an active labour market with 35% of households hiring labour on their farms.

5.4. Shock to labour endowment

There are two additional features necessary in the data to motivate the paper: exogenous variation of household labour endowment is required for identification and a large (positive) shock to household labour supply from Covid-19 is the premise for Proposition I.

Table 2 reports average changes in household labour endowment between the baseline and follow-up years. The statistics are highly consistent with Proposition I that Covid-19 caused a positive shock to labour supply. Average ΔE in 2018–2021 is much larger and positive (0.64) compared to in 2011–2015 (0.09).

All the flows of working-age individuals in Panel II compared to Panel I suggest that Covid-19 caused an expansion in the average household's labour endowment: average move-outs of the household are notably attenuated in Panel II; average move-ins rise in Panel II—possibly consistent with return migration; and the labour endowment boost from retaining ageing children is markedly stronger in Panel II. There is identification of the key regressor as Table 2 shows that 75% of households in Panel I and 70% in Panel II experienced any net ΔE from one survey to the next. I argue the variation is primarily exogenous, especially in Panel II because of the nature of Covid-19 being an extreme shock.

6. Empirical Results

6.1. Elasticity of farm labour demand to household size

This section presents the estimated results for the model of demand for farm labour. Table 3 presents the estimates of β_1 which, under the null of separation, should be zero.

6.1.1. Cross-section model and first differences model

The first column in Table 3 reports estimates of β_1 pooling all households and treating the sample as a cross-section as in specification (13), separately for each panel. In both panels, separability is strongly rejected in the cross-section. The elasticity of farm labour demand to household labour endowment is around 0.4.

[11] The household labour questionnaire asks the total number of weeks each household member worked, the average days per week and the average hours per day. The hired labour questionnaire asks the total number of men/women/children hired and the total number of days they were hired. These units are aggregated to create person-day variables for hired and household labour.

Table 1. Summary statistics

	Panel I: Pre Shock		Panel II: Shock	
Variable	**2011**	**2015**	**2018**	**2021**
Household characteristics				
Labour endowment, with kids	3.02 (1.41)	3.03 (1.45)	2.91 (1.30)	3.24 (1.53)
Labour endowment, no kids	2.72 (1.27)	2.72 (1.32)	2.65 (1.17)	2.95 (1.40)
Age of head (years)	44.47 (15.21)	47.55 (15.41)	44.33 (14.98)	46.54 (14.72)
Education of head (years)	1.57 (2.82)	1.84 (3.21)	2.28 (3.60)	2.53 (3.86)
Months away (hh average)	0.16 (0.49)	0.11 (0.48)	0.34 (0.96)	0.31 (0.81)
Share of prime age males	0.23 (0.17)	0.24 (0.19)	0.26 (0.19)	0.26 (0.18)
Share of prime age females	0.26 (0.16)	0.27 (0.19)	0.28 (0.17)	0.28 (0.18)
Share of elderly males	0.03 (0.10)	0.04 (0.13)	0.03 (0.11)	0.03 (0.10)
Share of elderly females	0.03 (0.12)	0.05 (0.18)	0.04 (0.14)	0.04 (0.14)
Share of children	0.43 (0.23)	0.48 (0.33)	0.40 (0.24)	0.45 (0.34)
Farm characteristics				
Land cultivated (acres)	6.10 (34.02)	5.11 (36.63)	6.03 (38.79)	9.19 (63.84)
Land owned (acres)	4.22 (17.89)	3.61 (16.18)	4.00 (18.58)	3.15 (14.03)
Farm demand characteristics (person-days)				
Total Labour Demand	246.92 (513.01)	230.29 (752.19)	159.04 (194.94)	130.84 (158.68)
Family supplied labour	223.76 (427.08)	179.70 (232.77)	134.21 (160.41)	110.70 (138.49)
Hired labour	23.16 (268.28)	50.59 (702.58)	24.83 (76.67)	20.13 (63.96)
Labour demanded for…				
Cultivation	175.89 (469.02)	141.92 (661.19)	100.90 (143.13)	66.81 (92.23)
of which hired	17.19 (264.07)	37.80 (630.14)	14.94 (54.54)	8.57 (33.70)
Harvest	71.03 (107.86)	88.37 (253.72)	58.14 (80.22)	64.03 (91.65)
of which hired	5.97 (34.03)	12.79 (223.10)	9.89 (42.23)	11.57 (47.94)
Observations	2777	3130	1688	1612

Notes: Standard deviations are in parentheses. "Labour endowment with kids" uses adult equivalence scale for children aged 11–15. "Prime" demographic groups are those aged 15–60. Total labour demand is the sum of hired and family labour in both cultivation and harvest seasons.

Table 2. Inter-annual changes in number of members and labour endowment

Changes between waves:	Panel I: Pre Shock 2011 & 2015	Panel II: Shock 2018 & 2021
Δ Number of members	−0.299	0.590
Δ Labour endowment (E)	0.091	0.642
ΔE from ageing children	−0.114	0.471
ΔE from move-ins	0.185	0.281
ΔE from move-outs	−0.628	−0.110
Any net Δ in E (=1)	0.753	0.704
Increase in E (=1)	0.455	0.650
Decrease in E (=1)	0.298	0.054

Notes: All entries are household-level averages that phase ageing children into the labour endowment except for the first row which counts each member as one person. Children aged 11–14 count as (Age −10) × 0.2 in the labour endowment. Δ Labour endowment (E) is the sum of the three categories above.

The key results are reported in Column (2) of Table 3 for the main specification of interest (14). In the counterfactual pre-shock sample, separation is rejected. This suggests market failures were already operating before the effect of the Covid shock on households. In the Covid-19 panel, separation is also rejected. The Covid sample estimate is larger and more significant than the counterfactual estimate, roughly doubling from 0.33 to 0.66. Households that experienced the positive shock to their labour supply increased labour inputs on their farm more than households that did not experience the shock; increased farm labour utilisation helped partly absorb the larger household labour supply.

Comparing the cross-sectional and fixed effects estimates within each panel, in Panel I the elasticity falls by a quarter when accounting for time-invariant household factors. This indicates unobserved characteristics are important for explaining non-separation in the 2011–2015 panel, but cannot explain it alone as I still find evidence for non-separation when including household fixed effects. However, the cross-sectional bias in Panel II is in the opposite direction: the elasticity increases by almost half. I suggest this is because the large sweeping shock to labour endowment weakens the correlation between the unobserved household qualities and household size; there is now a stronger driver of household size.

6.1.2. Farm demand for family labour and hired labour

Columns (3) and (4) in Table 3 disaggregate the dependent variable into demand for family labour and hired labour, respectively. The insignificant estimated elasticities in column (4) are not evidence in favour of complete markets; column (2) is the key variable of interest for the separation test. Rather, statistically insignificant elasticities of hired labour demand suggest the farm is not substituting efficiently in response to increases in its family size. Furthermore, farms expand family labour on the farm in response to larger household labour endowments, and do so significantly more in 2018–2021 (looking at Column (3), elasticities are 0.38 vs 0.70). This is consistent with households operating under labour market failures and absorbing the shock by expanding family labour utilisation on the farm (Proposition II).

Columns (3) and (4) indicate that households do not shed hired labour when family labour rises; if anything, they hold on to (or even slightly increase) hired labour alongside their own labour. This could be due to informal insurance agreements with other households in the village to help each other out with financial hardship from Covid-19. There is evidence in the literature that social norms play a potentially big role in governing rural labour markets (Breza et al. 2019). This motivates the intuition that "family comes first": when households get more internal labour, they use it—rather than shedding external workers—and expand total farm production to absorb that labour.

6.1.3. Confounding time-varying shocks

First-differencing eliminates time-invariant household heterogeneity, but time-varying shocks could still confound the estimates. A shock that may confound the estimates is government support to households. Ethiopia's Productive Safety Net Program (PSNP) is a large government program that provides food and cash transfers and remunerated public work to help relieve food insecurity of the rural poor. I find the results in Table 3 are robust to the exclusion of households that participated in the PSNP. Another concern is that Covid-19 impacted ΔFarm Labour directly, not just through ΔE, for instance by lowering crop prices. In practice, I attempt to mitigate this concern by controlling for cultivated acreage and household demographic composition.

6.2. Did the estimated elasticity increase post-Covid?

Taking stock, I have found that Covid-19 increased labour supply in the household which is consistent with Proposition I. Next, I found evidence for Proposition II that households were unable to fully absorb the shock, which resulted in household characteristics leaking into input allocation decisions. This raises the important question of whether the labour supply shock led to greater consequences from market failures, statistically speaking. Therefore, Proposition III is focused on the intensive margin of market failures.

In Table 4 I pooled both panels of households into one sample and interacted all right-hand-side regressors with a Post-Covid dummy variable equal to 1 if the household was surveyed in the 2018–2021 panel. This allows a formal test of whether the separation is statistically weaker after the Covid shock or if the difference in estimates is due to sampling fluctuation.

Although the estimate of the interaction term (0.327) is economically large—implying the elasticity roughly doubled in the Covid panel—it is measured imprecisely (standard error of 0.25) and fails to reach statistical significance. Therefore, while Table 3's estimates suggested a substantially bigger farm input response to ΔE during Covid, Table 4 shows the difference is not statistically distinguishable from sampling variation. I conclude that the evidence is qualitatively consistent but quantitatively inconsistent with Proposition III.

Between the follow-up survey of Panel I to the follow-up of Panel II (2015 to 2021), Ethiopia's rural markets likely became increasingly integrated and evolved, following the general trend of agricultural markets deepening over the last quarter century (Suri & Udry 2022). Against this backdrop of deepening markets, the persistent finding of non-separation—especially a potential intensification during the Covid-19 period—becomes more notable. It suggests that Covid-19 likely re-introduced or exacerbated imperfections that had otherwise been diminishing over time.

7. Heterogeneity in household sophistication

How well the system of agricultural labour markets absorb the shock to household labour endowment may depend on how 'sophisticated' the household is. Sophisticated households in the rural setting are households with higher wealth or consumption

Table 4. Estimation of Covid-19 impact on labour endowment elasticity

	Δ Farm labour
ΔE	0.329**
	(0.131)
$\Delta E\times$ Covid-panel	0.327
	(0.252)
Covid-panel	0.131
	(0.141)
Controls	Yes
Observations	4,350

Notes: Standard errors in parentheses, clustered at the EA-wave level. Covid takes value 1 for waves 4–5 and 0 otherwise. Controls include Δ ln(acres) and demographic shares. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

levels, higher human capital, participation in the wage market, formal salaried jobs or more land. The Covid-19 shock may have affected them differently because these households could better smooth the consequences of a labour supply shock. I put forward Proposition IV that sophisticated households absorb the Covid shock better than non-sophisticated households by placing more family members into the formal wage market.

Table 5 presents results for the heterogeneous responses by households in the wage market along two measures of sophistication: previous experience in the wage market and higher human capital. The statistically significant estimates of the first interaction term (2nd row of each panel) show that more sophisticated households have more household members working in the wage market. Previous market employment is more important than having an educated household head for getting members into formal work.

The "experience" benefit that sophisticated households exhibit disappears with Covid. The triple interaction terms are negative for both measures of sophistication and larger in magnitude than the "experience" benefits. This suggests the Covid shock was structural and affected households across the spectrum, even sophisticated households with pre-established networks in the market. Taken together, the estimates of Table 5 are inconsistent with Proposition IV. Sophisticated households were not better able to absorb the shock, suggesting all households faced market

Table 5. Heterogeneous sophistication of the household

	Δ Number of members employed in wage market
A. Employment in wage market at baseline	
ΔE	−0.005
	(0.013)
$\Delta E\times$ Had market employment	0.351***
	(0.132)
$\Delta E\times$ Covid-panel	0.161
	(0.125)
$\Delta E\times$ Had market empl. × Covid-panel	−0.445*
	(0.246)
B. Educated household head at baseline	
ΔE	−0.005
	(0.021)
$\Delta E\times$ Educated	0.116**
	(0.052)
$\Delta E\times$ Covid-panel	0.243*
	(0.132)
$\Delta E\times$ Educated × Covid-panel	−0.226
	(0.175)

Notes: Pooled sample of both panels with a binary indicator "Covid-panel" equal to 1 if the household was part of the 2018–21 panel. Standard errors in parentheses, clustered at the enumeration area for each wave. All regressions include controls for changes in log of cultivated acreage and changes in demographic shares. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3. Effect of household labour endowment on farm labour demand

	A. Pooled	**B. Household FE in first differences**		
Dep. variable:	Farm labour	Δ Farm labour	Δ Family labour	Δ Hired labour
	(1)	(2)	(3)	(4)
Panel I: Pre Shock				
Labour endow. (E)	0.412***			
	(0.057)			
ΔE		0.329**	0.381***	−0.188
		(0.131)	(0.133)	(0.181)
N. Households	5,561	2,777	2,777	2,777
Panel II: Shock				
Labour endow. (E)	0.463***			
	(0.083)			
ΔE		0.657***	0.700***	0.254
		(0.215)	(0.210)	(0.368)
N. Households	3,262	1,574	1,574	1,574

Notes: Standard errors in parentheses, clustered at the enumeration area for each wave. All regressions include controls for changes in log of cultivated acreage, changes in demographic shares, and wave fixed effects. Labour endowment E and all dependent variables are measured in logs. All dependent variables are measured in person-days. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

failures in the same degree and the Covid-19 shock was structural. I find this result is robust to using other proxies for household sophistication such as household consumption expenditure per capita and PSNP participation (a proxy for lower 'sophistication'), as shown in Addendum A2.

8. Conclusion

The separation test is a well-established procedure for testing the completeness of rural markets (Dillon et al. 2019). I use it as a framework to analyse how the system of rural markets and agricultural households absorb a shock. I find that rural Ethiopian farm households (i) felt the impact of Covid-19 through an increase in their household size; (ii) were operating in a set of markets exhibiting market failures; (iii) this meant the additional family labour supply could not be fully absorbed by labour markets; (iv) households responded by increasing labour utilisation on the farm, representing a failure to separate production input allocations from consumption characteristics; (v) this effect was structural and affected households at all levels of sophistication.

Separation is a 'portmanteau test' for completeness in at least two factor or output markets (LaFave & Thomas 2016). While it is not technically possible to identify which markets exhibit failures, I take the results as suggestive of imperfections in labour markets because of the nature of Covid-19 as a large unprecedented disruption to labour supply. Furthermore, even if labour markets are complete but separation fails because of other market failures (such as credit), this still represents misallocation on the farm which is important in itself. It highlights that the set of rural labour market institutions are inadequate to help buffer the misallocation and efficiency consequences from a large shock to agricultural households. As these households tend to rely on agriculture for livelihoods and consumption needs, this outcome is relevant for their day-to-day wellbeing.

While the Covid-19 shock represented an increase in the labour supply in agricultural rural markets—potentially causing a temporary labour surplus—it is important to distinguish which constraints are afflicting the labour market. Further research may explore whether it is demand constraints including lack of available jobs, hiring frictions or discrimination; supply constraints of workers not having the right skill-set or employability; or intermediation constraints including matching frictions, wage rigidity, information asymmetries, skills signalling difficulties or lack of trust. Effective policy prescription requires an understanding which constraints are operating in the market. It is also relevant to policy-makers to understand whether the misallocation consequences from the Covid-19 shock persist today.

Suri and Udry (2022) express that the literature is limited on the role that poorly functioning labour markets play in hampering agricultural productivity in Africa. I hope that my paper contributes, in some small way, to address these shortcomings and fits into the broader agenda that urges more research on labour market failures in rural settings and how they reverberate to agricultural productivity in developing economies.

Addendum

A1: Excluding Tigray in Panel I

Table 6. Effect of household labour endowment on farm labour demand (excluding Tigray in Panel I)

	A. Pooled	B. Household FE in first differences		
Dep. variable:	Farm labour	Δ Farm labour	Δ Family labour	Δ Hired labour
	(1)	(2)	(3)	(4)
Panel I: Pre Shock				
Labour endow. (E)	0.412***			
	(0.059)			
ΔE		0.306**	0.366***	−0.164
		(0.140)	(0.143)	(0.197)
N. Households	4,938	2,466	2,466	2,466
Panel II: Shock				
Labour endow. (E)	0.463***			
	(0.083)			
ΔE		0.657***	0.700***	0.254
		(0.215)	(0.210)	(0.368)
N. Households	3,262	1,574	1,574	1,574

Notes: Standard errors in parentheses, clustered at the enumeration area for each wave. All regressions include controls for changes in log of cultivated acreage, changes in demographic shares, and wave fixed effects. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

A2: Additional heterogeneity measures

Table 7. Heterogeneous sophistication of the household (additional measures)

	Δ Farm labour
A. Expenditure per capita at baseline	
ΔE	0.269*
	(0.161)
$\Delta E\times$ High expenditure pc	0.052
	(0.250)
$\Delta E\times$ Covid-panel	0.161
	(0.125)
$\Delta E\times$ High expenditure pc × Covid-panel	−0.189
	(0.805)
B. PSNP participation at baseline	
ΔE	0.351**
	(0.142)
$\Delta E\times$ Participated in PSNP	−0.199
	(0.290)
$\Delta E\times$ Covid-panel	0.294*
	(0.262)
$\Delta E\times$ Participated in PSNP × Covid-panel	0.301
	(0.556)

Notes: Pooled sample of both panels with a binary indicator "Covid-panel" equal to 1 if the household was part of the 2018–21 panel. Standard errors in parentheses, clustered at the enumeration area for each wave. All regressions include controls for changes in log of cultivated acreage and changes in demographic shares. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

References

[1] Aspen, H. (1993). Competition and Co-operation. North Ethiopian Peasant Households and their Resource base.

[2] Ayele, Z. & Fessha, Y. (2021). COVID-19 in Ethiopia: A Year in Review.

[3] Bardhan, P. K. & Udry, C. (1999). *Development Microeconomics*. Oxford University Press.

[4] Benjamin, D. (1992). Household Composition, Labor Markets, and Labor Demand: Testing for Separation in Agricultural Household Models. *Econometrica* 60(2), 287.

[5] Breza, E., Kaur, S. & Krishnaswamy, N. (2019a). Scabs: The social suppression of labour supply.

[6] Breza, E., Kaur, S. & Krishnaswamy, N. (2019b). Social Norms as a Determinant of Aggregate Labor Supply. NBER Working Paper No. w25880.

[7] Breza, E., Kaur, S. & Shamdasani, Y. (2021). Labor Rationing. *American Economic Review*.

[8] Card, D. (1990). Unexpected Inflation, Real Wages, and Employment Determination in Union Contracts. *American Economic Review* 80(4), 669–688.

[9] Chayanov, A. (1991). *The Theory of Peasant Co-operatives*. I.B. Tauris & Co Ltd.

[10] Dillon, B. & Barrett, C. B. (2017). Agricultural factor markets in Sub-Saharan Africa: An updated view with formal tests for market failure. *Food Policy* 67, 64–77.

[11] Dillon, B., Brummund, P. & Mwabu, G. (2019). Asymmetric non-separation and rural labor markets. *Journal of Development Economics* 139, 78–96.

[12] Duflo, E. & Udry, C. (2004). Intrahousehold Resource Allocation in Cote d'Ivoire: Social Norms, Separate Accounts and Consumption Choices. NBER Working Paper No. w10498.

[13] Feder, G., Just, R. E. & Zilberman, D. (1985). Adoption of Agricultural Innovations in Developing Countries: A Survey. *Economic Development and Cultural Change* 33(2), 255–298.

[14] Ford, K., Porter, C. & Pankhurst, A. (2021). A Lost Year of Learning for Girls in Ethiopia: Evidence from the Young Lives At Work Covid-19 Phone Survey.

[15] Goldstein, M. & Udry, C. (2008). The Profits of Power: Land Rights and Agricultural Investment in Ghana. *Journal of Political Economy* 116(6), 981–1022.

[16] Gollin, D. & Udry, C. (2021). Heterogeneity, Measurement Error, and Misallocation: Evidence from African Agriculture.

[17] Harris, D., Baird, S., Ford, K., Hirvonen, K., Jones, N., Kassa, M., Meyer, C., Pankhurst, A., Wieser, C. & Woldehanna, T. (2021). The Impact of COVID-19 in Ethiopia: Policy Brief.

[18] Harris, J. & Todaro, M. (1970). Migration, Unemployment and Development: A Two-Sector Analysis.

[19] Jones, M., Kondylis, F., Loeser, J. & Magruder, J. (2022). Factor Market Failures and the Adoption of Irrigation in Rwanda. *American Economic Review* 112(7), 2316–2352.

[20] Kaur, S. (2019). Nominal Wage Rigidity in Village Labor Markets. *American Economic Review* 109(10), 3585–3616.

[21] Kaur, S. (2024). Labor Markets in Developing Countries. Nobel Symposium on Development Economics.

[22] Kilic, T., Palacios-López, A. & Goldstein, M. (2015). Caught in a Productivity Trap: A Distributional Perspective on Gender Differences in Malawian Agriculture. *World Development* 70, 416–463.

[23] LaFave, D. & Thomas, D. (2016). Farms, Families, and Markets: New Evidence on Completeness of Markets in Agricultural Settings. *Econometrica* 84(5).

[24] Nolte, K., Sipangule, K. & Wendt, N. (2022). Agricultural households in times of crisis. The COVID-19 pandemic, livelihoods and land-use decisions. *Journal of Land Use Science* 17(1), 134–160.

[25] Restuccia, D., Yang, D. T. & Zhu, X. (2008). Agriculture and aggregate productivity: A quantitative cross-country analysis. *Journal of Monetary Economics* 55(2), 234–250.

[26] Schultz, T. W. (1964). Changing Relevance of Agricultural Economics. *Journal of Farm Economics* 46(5), 1004.

[27] Sen, A. (1962). An Aspect of Indian Agriculture. *Economic Weekly*.

[28] Singh, I., Squire, L. & Strauss, J. (1986). A Survey of Agricultural Household Models: Recent Findings and Policy Implications.

[29] Suri, T. & Udry, C. (2022). Agricultural Technology in Africa. *Journal of Economic Perspectives* 36(1), 33–56.

[30] Udry, C. (1996). Gender, Agricultural Production, and the Theory of the Household.

[31] World Development Indicators. The World Bank.