Neighborly Business: Labor Exchange Networks and Wage Markets in Rural Thailand

Aisha Salih*[1]

[1]Aix-Marseille School of Economics, Marseille, France

October 18, 2026

Abstract

Why does unpaid reciprocal labor persist after rural wage markets develop? To answer this question, we study the allocation of household labor following surprise increases in crop returns. Using monthly bilateral paid and unpaid labor transactions from 16 Thai villages, we build a Bartik-type shift-share measure of crop-return innovations using unexpected changes in world prices and pre-set village crop specialization, interacted with each crop's harvest timing. An unexpected increase in crop prices near harvest time reallocates both labor exchange and hiring toward established partners relative to previously untied potential providers. The findings show that established relationships help mobilize labor when timely agricultural work becomes more valuable and provide new evidence on the role of relational infrastructure in organizing paid work in rural labor markets.

Keywords: Labor markets, informal institutions, networks, relational contracts, Thailand
JEL Codes: J43, O12, O17, D85

*aisha.SALIH@univ-amu.fr, 5–9 bd. Maurice Bourdet, 13001 Marseille

1 Introduction

Unpaid reciprocal farm labor or labor exchange is widely documented in Sub-Saharan Africa, South and Southeast Asia, Latin America, and appears in anthropological accounts from North America. Predominantly used for lumpy agricultural tasks, this practice can take the form of bilateral or small-group exchange either with balanced reciprocity or large festive work parties rewarded with food and drink where expected reciprocity is looser (Olken and Singhal, 2011; Erasmus, 1956; Arciniegas et al., 2025; Gilligan, 2004). As labor exchange remains common even where a cash-wage market is active, a natural question is why households exchange labor instead of settling every transaction in cash. One reason is that reciprocal exchange defers the cost of settlement; much like contract-farming arrangements that shift insurance premiums from planting to harvest (Casaburi and Willis, 2018), exchange allows a household to obtain labor today and repay in kind later, avoiding the cash outlay that a spot wage requires. When cash is scarce or credit markets are thin, this deferral is valuable (Rosenzweig, 1988; Fafchamps, 1993). Furthermore, social norms and repeated interaction sustain cooperation even without formal enforcement (Kranton, 1996; Parnwell, 2005). Economic accounts of labor exchange emphasize factor-market imperfections, returns to teamwork when tasks are lumpy, and monitoring advantages when workers and employers know one another (Gilligan, 2004; Kranton, 1996). A related literature emphasizes the calendar dimension of rural labor demand: seasonal concentration in agricultural tasks can generate low-work periods outside the main crop season, especially when few activities provide counter-cyclical labor demand (de Janvry et al., 2022). These functions make labor exchange a potentially important part of production even where an active cash-wage market already exists.

A common account of rural commercialization nevertheless treats reciprocal work and wage labor as successive institutions. Ethnographic evidence describes festive work parties giving way to stricter bilateral exchange and, eventually, to wages as cash markets deepen (Fujimoto, 2013; Hossain, 1997; Phelinas, 1994). Yet exchange and hiring often coexist, sometimes within the same pair of households. This coexistence matters because a labor relationship and the contract used for a particular task are distinct. A trusted exchange partner may remain an unpaid collaborator, become a paid worker, or alternate between the two arrangements. Recent evidence from Tanzania, for example, shows that labor-exchange teams can facilitate paid employment by internalizing monitoring (Arciniegas et al., 2025). Wage-market development may therefore build on reciprocal relationships rather than simply displace them.

Much has been shown in the literature on how rural wages, employment, and migration respond to agricultural productivity and labor-demand shocks (Jayachandran, 2006; Kaur, 2019; Breza and Kaur, 2025). Less is known about the contractual margin within persisting relationships. Existing ties may disappear, remain reciprocal arrangements, or provide channels for paid work. We study this relational margin. When unexpected crop returns rise, do households expand unpaid exchange with established partners, pay those same partners, or reallocate transactions toward providers with whom they had no prior labor relationship?

This question links two views of market development. In one, deeper wage and credit markets

replace informal exchange as cash permits transactions beyond a small circle of known partners. In the other, markets remain embedded in relationships because identity carries information about reliability, effort, and future interaction.

Partner identity is especially relevant in agriculture. Production tasks are local, time sensitive, and difficult to contract completely. A household may know the prevailing daily wage yet still prefer a worker whose effort and availability it has observed. Conversely, a provider may accept deferred reciprocal settlement only from a household expected to return the obligation. The set of feasible contracts is therefore relationship specific. Aggregate hired-labor totals conceal network dynamics and whether adjustment occurs through new matches, established wage relationships, or partners with experience under both cash and non-cash settlement.

The distinction also changes how structural transformation is measured. A falling share of exchange can reflect the disappearance of reciprocal relationships, a lower frequency of unpaid transactions within surviving relationships, or movement by the same pairs between exchange and wages. These paths may have different implications for access or resilience. If exchange depends on landholding and prior interaction, it can exclude households without the assets needed to reciprocate. If established ties facilitate paid work, their disappearance can also remove a channel into employment. Relationship-level records are therefore needed to separate these margins.

We answer this question using bilateral transaction records from the Townsend Thai Monthly Survey, which follows households in 16 rural villages from 1998 to 2014 (Townsend, 2011, 2016). Rural Thai labor exchange, known locally as *long khaek* or *ao raeng*, traditionally organizes neighbors to work together or in sequence across farms without immediate cash payment (Shigetomi, 1998; Parnwell, 2005). Among 686 households, 61.5 percent participate in labor exchange and 70.3 percent in paid labor, either as a receiver or provider. Nearly half participate in both. Rice and other field crops account for about 69 percent of recorded exchange episodes, and exchange is strongly seasonal.

The bilateral records distinguish a relationship from the contract used within it. We classify each pair using its labor transactions during the first 24 waves, from September 1998 through August 2000. These baseline categories—baseline-exchange-only, baseline-hire-only, baseline-dual, and baseline-untied—are fixed regressors. They describe the pair's history before the estimation period and are distinct from the contracts observed in a given wave. In each monthly survey wave, we pair every household with every other potential provider currently living in the same village, whether or not they transact, excluding those that have left the village. This structure compares established partners with available households that were not used.

We construct a Bartik-type crop-return shock by combining unexpected world-price movements with the crops in which each village specializes. Crop weights are fixed at baseline, and an unexpected shock is the difference between realized price growth and its forecast based on earlier price changes. For each household interview, we use price information set before the household's reported transactions occurred. The shock therefore captures changes in crop returns that were not predictable from recent price movements or normal seasonality. To test sensitivity to harvest

timing, we calculate a crop-specific harvest calendar from monthly records of reported harvest. For a given village and crop, the calendar measures how frequently households in other villages in the same region report harvesting that crop in each month of the year and is held fixed throughout the analysis. We then test whether the Bartik shock has a larger effect when the crops important to the focal village are in their usual harvest months. The focal village's own harvest reports are excluded from the calendar.

We find a reallocation of transactions toward established partners when a positive crop-return innovation arrives near harvest. The harvest-timing coefficient is not itself the total effect; it measures how the response to a return innovation changes as harvest intensity rises. Because the average and harvest-timed exposures are generated by the same price innovations, a return innovation changes both jointly. We therefore evaluate their combined effect at the village's peak harvest month. On this measure, hiring rises about as much as exchange once each response is compared with its baseline rate. The point estimate is larger for baseline-dual pairs than for baseline-exchange-only pairs, but the difference is not statistically precise. Within baseline-dual pairs, both mutual and one-way exchange increase, consistent with a combination of contemporaneous coordination and reciprocity settled across time. The overall pattern is a broad activation of existing labor ties rather than a shift from cash settlement to unpaid exchange.

The paper makes two main contributions. First, it brings a relationship-level perspective to rural labor-market adjustment. By observing both receivers and providers of a transaction, we distinguish changes in total hiring from reallocation across partners and contract forms—a margin not captured by household-level wage or employment outcomes. Second, we contribute to work on labor exchange and informal institutions by showing that reciprocity remains responsive within relationships that also support paid work. Market and non-market contracts are therefore not confined to separate relationships. This relational margin is relevant for evaluating shifts in transportation, labor-market matching institutions, and mechanization, each of which can modify access to or demand for seasonal workers. These impacts may be reflected in the selection of contracting partners and in contract structure even when village wages remain sticky due to broader market-clearing forces or other wage-setting frictions.

The rest of the paper proceeds as follows. Section 2 describes the setting, data, and labor-exchange institution. Section 3 states the mechanisms and empirical predictions. Section 4 presents the dyadic specification and the external return shifter. Section 5 reports the estimates and validity checks. Section 6 concludes.

2 Data and Institutional Setting

This section describes the Townsend Thai Monthly Survey and the construction of the dyadic sample used in the analysis. We first discuss the study villages and explain how the survey modules are combined. We then define labor relationships and contract outcomes at the household-pair level. The final two subsections document the role of labor exchange in village production and describe

how exchange relationships evolved as wage and credit markets expanded. These institutional patterns motivate the distinction between a labor relationship and the contract used for a particular transaction.

2.1 Sample

The Townsend Thai Monthly Survey follows households in 16 rural villages in the provinces of Lopburi and Chachoengsao in the Central region and Buriram and Sisaket in the Northeast (Townsend, 2011, 2016). The four provinces are shown in Appendix Figure A.1, and their baseline crop composition is reported in Appendix Table A.2. Approximately 45 households are interviewed in each village. Exploratory fieldwork for the Townsend Project began in the early 1990s, followed by a large baseline survey in April and May 1997. After Thailand floated the baht in July 1997, resulting in a financial crisis, the project returned to these communities in 1998 and established the monthly panel used here (Townsend et al., 2013; Townsend, 2016). Beginning in August 1998, the survey collected detailed information on household labor, farming, transactions, income, assets, and financial activity. The monthly frequency is particularly useful for this study because agricultural tasks and labor transactions vary substantially within the year.

The study provinces span distinct agricultural environments. Lopburi and Chachoengsao are closer to Bangkok and have relatively diversified commercial agriculture, while Buriram and Sisaket are predominantly rice-growing provinces in the Northeast. Rice accounts for more than four fifths of recorded baseline output in Chachoengsao, Buriram, and Sisaket, whereas corn dominates in Lopburi. The empirical strategy uses these initial differences in crop specialization to construct village exposure to subsequent changes in world crop prices. The five crops included in the exposure measure are rice, corn, soybeans, groundnuts, and sorghum[1].

The analysis combines several survey modules, each of which serves a different purpose. Household activity records measure person-days in own-farm and paid work. Farm modules record the crops households grow, when they harvest, and how much they produce. Transaction modules identify the counterparty to a labor transaction and distinguish paid labor from exchange labor. Household and village rosters indicate whether both parties were resident in the village at the time of the transaction. We use the transaction records to construct the dyadic outcomes, the activity records to describe the agricultural labor calendar, the farm records to measure baseline crop specialization, and the rosters to define the set of potential partners in each wave. As a result, the descriptive statistics may be reported at the household-month or transaction level, while the main outcome is measured for a directed household pair in a given survey wave.

Panel A of Table 1 describes the 686 households observed during the first 24 survey waves (baseline). Nearly all households hold agricultural land. Farm activities account for about two thirds of recorded productive work, and wage work accounts for most of the remainder. Thus, the study villages had active cash labor markets from the beginning of the panel: households cultivated

[1] Aquaculture is also important in parts of Chachoengsao, but it is excluded for the absence of a comparable monthly world-price series.

their own plots, hired workers, supplied wage labor, and participated in labor exchange.

Table 1: Baseline households and the strict-recall dyadic sample

	Mean	Std. dev.		
Panel A. Baseline household characteristics, raw waves 1–24				
Households	686			
Household members	4.2	1.8		
Landholding household (%)	93.4			
Landholdings (hectares)	3.9	3.6		
Farm share of productive work (%)	67.4			
Wage share of productive work (%)	22.1			
Panel B. Interview-aligned directed pair-wave sample				
	Directed	Share of	Pair-waves with transaction (%)	
Baseline pair type	pairs	pairs	Any hired	Any exchange
Baseline-untied	22,746	0.891	0.10	0.15
Baseline-exchange-only	1,100	0.043	0.58	3.75
Baseline-dual	446	0.017	1.68	3.10
Baseline-hire-only	1,242	0.049	0.70	0.52
Directed pair-waves				2,629,888
Receivers; villages				641; 16

Notes: Panel A reports 1998–2000 baseline characteristics. Panel B uses the presence-aware directed pair-wave sample for which the receiver has a valid recall interval and a strict pre-recall world-price innovation. Both endpoints are observed residents in the village-wave. Baseline pair types use undirected labor transactions in raw waves 1–24. The effective treatment period is September 2000–February 2010.

The period also saw a substantial expansion in formal credit. In 2001–02, the Million Baht Village Fund established a village bank in every Thai village, including all sixteen villages in the monthly survey. Previous work using these data finds that the program increased credit, consumption, agricultural investment, and wages (Kaboski and Townsend, 2012). We treat the fund as part of the institutional setting since the household-by-wave fixed effects in the main specification absorb contemporaneous household conditions associated with its introduction. Reciprocal and paid labor thus continued to coexist during a period in which access to credit was increasing.

2.2 Transaction network and variable construction

The main analysis uses bilateral labor transactions within the sampled villages. We describe a transaction as a directed flow from household j, the labor provider, to household h, the receiver. Exchange labor consists of unpaid work supplied by one household to another, generally with the expectation that labor or other help will be returned. Hired labor refers to work settled in cash; we use the terms hired labor and paid labor interchangeably.

Both endpoints are observed when the receiver and provider belong to the sampled village roster. This feature makes it possible to study the choice of partner as well as the choice of contract, although it also narrows the scope of the analysis. Providers outside the sampled roster account for 45.0 percent of hired transactions and 36.4 percent of exchange transactions. The esti-

mates therefore describe how households allocate transactions within the sampled village network. Broader household labor use is measured in person-days in a separate module, which generally does not identify the provider. Because the two modules differ in coverage and recall conventions, we use the transaction archive for relational outcomes and the household module for descriptive and mechanism exercises.

We classify relationships using transactions reported during the first 24 waves, from September 1998 through August 2000. A pair is connected at baseline if either household supplied labor to the other. We then assign each undirected pair to one of four mutually exclusive categories. A baseline-exchange-only pair used exchange but no paid labor, a baseline-hire-only pair used paid labor but no exchange, and a baseline-dual pair used both contracts. All remaining pairs are classified as baseline-untied. These labels describe the pair's history in waves 1–24. They are held fixed as regressors and are distinct from the contracts used in any later wave. For example, baseline-dual status records experience with both forms of settlement; it carries no treatment interpretation and does not necessarily remain the pair's contract type.

The outcome panel begins in wave 25. In each village and wave, we pair every eligible receiver with every other sampled household that was resident in the village and could therefore have supplied labor. Transactions that do not occur are coded as zeros. Requiring both endpoints to be resident avoids treating a household that has left the village as an available provider. We also require the receiver to have a recall interval between 15 and 60 days, because the length of this interval determines which price information was available before the reported transactions. These restrictions produce an effective estimation period from September 2000 through February 2010. Panel B of Table 1 reports the number of pairs and transaction rates in this sample.

The unit of observation is a directed pair-wave (h, j, w). We define *any exchange* as an indicator that provider j supplied unpaid labor to receiver h during the receiver's recall period. *Any hired labor* indicates that h paid j for labor over the same period. These outcomes overlap: a pair using both contracts is coded one for each. To distinguish contract mix, we define three mutually exclusive outcomes: exchange only, hire only, and both contracts. Exactly one of these outcomes equals one when a transaction occurs; if all three equal zero, the pair had no recorded labor transaction during the recall interval. By construction, any exchange is the sum of exchange only plus both contracts, and any hired labor is the sum of hire only plus both contracts.

Interview dates vary within a survey wave. The median absolute difference between receiver and provider interview dates is six days, and a substantial share of interviews are more than one week apart. We anchor the outcome to the receiver because the transaction record describes labor supplied to that household over its own recall interval. This dating convention is also used to assign the crop-return innovation in Section 4.2.

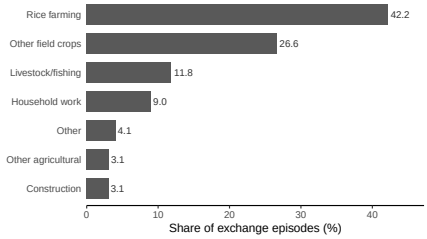
The resulting risk set is sparse because it retains all potential providers, including those with whom the receiver does not transact. This feature is necessary for the empirical comparison. A transaction with an established partner must be evaluated relative to other households the receiver could have used in the same wave. On baseline-dual pairs, exchange occurs in 3.10 percent of

pair-waves and hiring in 1.68 percent. The corresponding rates among baseline-untied pairs are 0.15 and 0.10 percent. Baseline status remains strongly predictive after August 2000, although some baseline ties become inactive and some previously untied pairs begin to transact. Keeping relationship type fixed allows us to ask how a stock of relationships observed before the treatment period may influence subsequent adjustment.

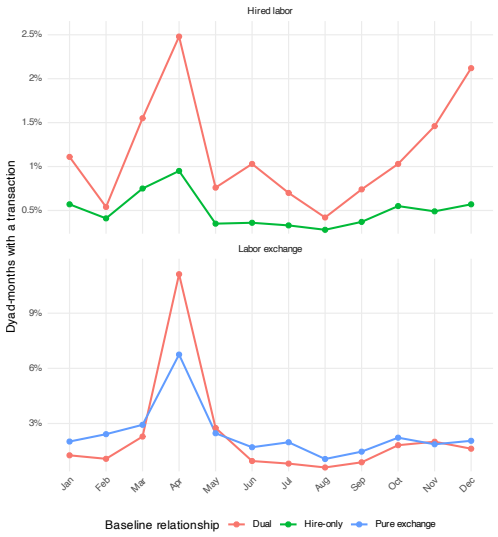
2.3 Labor exchange in the data

We next use the survey records to describe how labor exchange operates in these villages. The purpose is to establish the types of tasks for which exchange is used, how exchange varies over the agricultural calendar, and how quickly reciprocal obligations are settled. These patterns help interpret the contract responses estimated below.

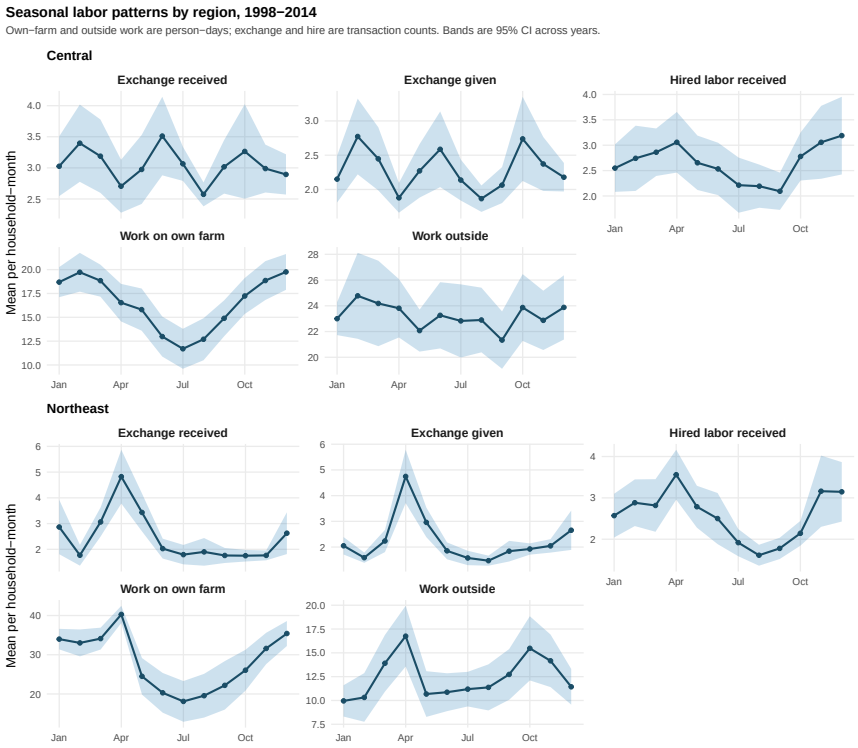
The available evidence points to an informal practice organized among neighbors. In the 1997 key-informant survey, village heads estimated that an average of 26 percent of agricultural labor was supplied through exchange, while reporting almost no membership in formal labor-exchange groups (Appendix Table A.3). Figure 1 provides a more detailed view. Panel A shows that rice and other field crops account for about 69 percent of recorded exchange episodes. As Panel B shows, these transactions are highly seasonal, especially among pairs with a prior exchange relationship. Nearly one third of sampled exchange transactions occur in April, around land preparation and the beginning of the rainy-season crop calendar. Although exchange is uncommon in an average month, households that report using it receive about 5.7 person-days of exchange labor.



(a) Tasks in exchange episodes



(b) Transactions by month and baseline relationship



(c) Household labor patterns by region

Figure 1: Where and when labor exchange occurs. Panel A reports the task composition of exchange episodes, 1998–2005. Panel B reports the percentage of pair-waves with exchange or paid labor by calendar month and baseline relationship. Panel C reports mean household labor use by region, 1998–2014.

Panel C places these transactions in the broader labor calendar. In the Central provinces, own-

farm work, exchange, and hiring peak in March and April, with a second concentration of own-farm work and hiring in November and December. In the Northeast, own-farm work is high around the main wet-season rice harvest and again in April, while exchange and hiring are concentrated between April and June. Labor demand thus tends to rise at times when potential providers may also be occupied on their own farms. Across the sample, recorded productive work is 23 percent higher in January–April than in July–October, and farm work accounts for most of this difference. Off-farm work changes little between these periods, suggesting that other activities provide limited smoothing of the agricultural calendar (de Janvry et al., 2022). Appendix B reports the full labor and wage calendars. It also compares survey labor with satellite vegetation, which rises after the early-year labor peak as planting gives way to crop growth. The vegetation measure is used only to validate timing and does not enter the crop-return shock.

The transaction records indicate that labor exchange involves reciprocal obligations, although these obligations are seldom settled through an equal and immediate return of labor. In 1998–2002, only 5 percent of pair-years record equal flows in both directions. Over the same period, 24.8 percent of directed episodes have a reverse flow in the same month, and 39.3 percent have a reverse flow during the following twelve months. Both measures decline over the panel: same-month mutual exchange falls to 21.0 percent in 2003–2008 and 12.7 percent in 2009–2014, while reverse flows during the following year fall to 33.8 and 22.2 percent. Return labor is therefore observed less often within a fixed period as the panel progresses. This pattern is consistent with obligations taking longer to settle, although weaker reciprocity and reduced transaction-module coverage may also contribute to the decline. Repayment may occur after the observed window, may be incomplete, or may involve forms of help that are not recorded in the labor module. A bilateral flow may also form part of a larger work party in which obligations are sustained socially. The dyadic records identify the two households that transact but do not capture the full set of obligations within the wider group (Erasmus, 1956; Parnwell, 2005).

2.4 Baseline relationship types and structural change

The types of households that participate in exchange help clarify what this institution can accomplish and who can access it. Panel A of Table 2 groups households according to participation in sampled exchange and hiring during 1998–99. A household using both contracts may exchange with one partner and hire another; this household classification is distinct from the pair categories fixed over raw waves 1–24. More than 95 percent of households using exchange only or both contracts at baseline hold land, compared with 83 percent of households without a sampled labor tie. Exchange users are therefore overwhelmingly cultivators, although they are not the wealthiest households in the sample. The groups differ more clearly in their use of labor. Households that participate in both exchange and hiring record 26.6 own-farm person-days per month, compared with 16.4 among households using exchange only at baseline, 12.1 among households using hiring only at baseline, and 10.2 among households without a sampled labor tie. Households using both contracts at baseline are also more likely to have a borrowing or lending relationship.

These household comparisons associate the use of both contracts at baseline with active cultivation and repeated interaction, with wealth playing a limited role. At the pair level, baseline-dual relationships tend to join cultivators of similar size and are more likely than baseline-hire-only relationships to overlap with kinship, gifts, asset transfers, or loans. Most labor ties, however, are not simultaneously credit or gift ties. Baseline-dual status may therefore capture features of the match—such as durability, trust, or a broader range of feasible settlements—that are not summarized by farm size or observable social links. Section 5.4 examines how much of the result is explained by the frequency of baseline transactions.

The close association between exchange and landholding also limits its potential reach. Because each household retains the output from its own plot, exchanging labor shares work without pooling the returns to land. Households without land or with little labor available to reciprocate may consequently find it difficult to enter these arrangements, a pattern also documented for labor-exchange groups in Indonesia by Gilligan (2004).

Panel B follows exchange over the life of the panel. The share of non-wage days supplied through exchange falls from 15.2 percent in 1998–2002 to approximately 7 percent in the two later periods. Return flows within twelve months and mutual exchange within the same wave also become less common. Part of this decline may reflect changes in survey coverage, since fewer transaction modules were completed in later years (Appendix Table A.4). We therefore use these trends as institutional background and leave the causes of the long-run decline open.

The broad pattern is consistent with earlier accounts of Thai agriculture. Phelinas (1994) describes family labor, exchange labor, and hired labor as coexisting among cultivators, with exchange already declining in rice-growing areas undergoing technological change. Related work on Asian rice systems documents a gradual move toward daily wages and piece rates as wages, migration, and mechanization restructure production (Hossain, 1997). However, the process is not uniform. (Kharel et al., 2023) find that, in some migrant-sending areas, labor scarcity has led households to rely more heavily on reciprocal labor.

The pair-level records reveal another feature of this transition. Pairs that exchanged during the baseline period often hired one another later, and some baseline-hire-only pairs subsequently exchanged. These unconditional transitions show that a relationship can support several forms of settlement. The relationship may persist while the contract changes from one task or period to the next. Appendix C provides additional evidence on partner characteristics and on the overlap of labor relationships with kinship, gifts, and credit.

Table 2: Labor relationships at baseline and over the panel

	No sampled labor contract	Exchange only	Hiring only	Both contracts
Panel A. Household contract use in 1998–99				
Households	123	81	141	341
Landholding household (%)	82.9	95.1	92.2	97.4
Own-farm person-days/month	10.2	16.4	12.1	26.6
Outside paid person-days/month	16.5	16.3	13.8	12.4
Wealth index (standardized)	−0.24	−0.35	0.30	−0.24
Any borrowing or lending tie (%)	15.4	27.2	31.2	46.9
Panel B. Recorded exchange over time				
	1998–2002	2003–2008	2009–2014	
Exchange share of non-wage days (%)	15.2	7.0	7.1	
Household-waves with exchange (%)	2.1	1.8	0.7	
Return flow within 12 months (%)	39.3	33.8	22.2	
Same-wave mutual flow (%)	24.8	21.0	12.7	

Notes: Household categories use sampled labor transactions during 1998–99, as receiver or provider. A household using both contracts participates in exchange with any sampled partner and hiring with any sampled partner; these need not be the same partner. Exchange-only and hiring-only households use only the respective contract, and remaining households have no sampled labor transaction in this window. These categories differ from the fixed pair-level baseline categories in the main regressions, which use raw waves 1–24. Panel B combines the household labor module and named transaction records. The later decline partly reflects transaction-module thinning, documented in Appendix A.

3 Conceptual Framework

This section develops a simple framework for interpreting how established labor relationships respond to changes in crop returns. The main distinction is between the relationship linking two households and the contract used for a particular task. An established pair may exchange labor, use a cash payment, or move between the two forms of settlement over time. We first discuss why exchange can remain useful when a wage market is active and then derive the empirical patterns associated with different forms of adjustment.

Repeated interaction can make both paid and unpaid transactions easier to sustain. When formal enforcement is weak, a prior relationship provides information about a worker's effort and reliability and about the likelihood that an exchange obligation will be honored (Kranton, 1996; Munshi, 2003; Jackson, 2008; Hull, 2007; Parnwell, 2005). The same information may be useful when the worker is paid. A household facing a time-sensitive task can assemble familiar workers more quickly, spend less effort screening them, and rely on past experience when monitoring the work. A known provider may therefore be valuable under either form of settlement.

Exchange has two additional features that can be useful in agriculture. First, it allows settlement to be deferred. A household obtains labor when the task must be completed and returns labor later, avoiding an immediate cash outlay. Second, it can help organize tasks that require several workers over a short period. Neighboring cultivators may work across fields in sequence or assemble a common work party, thereby coordinating labor even when their own production calendars overlap. These arrangements are particularly useful when tasks are difficult to monitor or when showing up at the right time matters for output (Gilligan, 2004; Kranton, 1996).

To see how these considerations affect contract choice, consider a household h seeking labor from provider j. Under a wage contract, the transaction is settled in cash at or near the time of work. Under exchange, the labor supplied by j creates a claim that h can repay with labor or other assistance, either immediately or at a later date. The return to obtaining an additional unit of labor depends on the value of the crop and on the timing of the task. At the same time, the provider's opportunity cost depends on her own farm work and off-farm opportunities. Hiring is feasible when the parties agree on compensation and the receiver can finance the payment. Exchange is feasible when the expected value of future assistance, discounted for delay and the risk of non-repayment, is sufficient to compensate the provider.

An increase in crop returns during a period of concentrated agricultural activity affects both sides of this transaction. Timely labor becomes more valuable to the receiver, while potential providers may also become busier on their own farms. A prior relationship can reduce the cost of finding and supervising labor, but exchange additionally requires the parties to coordinate work across fields or to carry an obligation forward. The relative response of exchange and hiring is therefore ambiguous. It depends on the receiver's liquidity, the value of information about the provider, the flexibility with which reciprocal obligations can be settled, and the extent to which the two households' labor needs coincide.

This ambiguity has implications for access and welfare. Exchange may allow a cultivating

household to obtain labor without an immediate cash payment, but participation usually requires land or labor with which to reciprocate. Hiring permits workers who cannot reciprocate to supply labor, although it requires the receiver to have cash or credit. Established ties can reduce transaction costs for the households involved while limiting opportunities for potential providers outside the network. The empirical analysis identifies whether adjustment occurs through existing relationships and which contracts are used within those relationships. A welfare ranking would additionally require information on output, effort, and forgone opportunities.

3.1 Empirical predictions

An unexpected increase in crop returns raises the value of completing agricultural work on time. Because each transaction requires a receiver to seek labor and a provider to supply it, the estimates capture the equilibrium response of both parties. Interacting the price innovation with the harvest calendar allows us to ask whether this response becomes stronger in months when the crops grown in the village are usually harvested. We organize the discussion around three mechanisms which harvest timing can help separate.

First, liquidity and deferred settlement. A positive price surprise at harvest relaxes cash constraints, because output is about to be sold. If this mechanism dominated, harvest-timed shocks should favor hiring over exchange. The exchange response should instead be stronger outside harvest, when cash is scarce and deferred reciprocal settlement is more valuable.

Second, coordination and the timeliness of lumpy tasks. Exchange is used to assemble several workers over a short period. If this mechanism is important, exchange should respond mainly at harvest, and a larger share of the response should take the form of same-wave mutual exchange. Conditional on a transaction occurring, days per exchange episode should also rise if households lengthen joint work or assemble larger work parties.

Third, information and monitoring with known partners. Prior interaction reduces search and supervision costs under either contract. Both exchange and hiring should then rise with established partners, and in similar proportion to their baseline rates. Days per transaction should rise for both contracts if activation also deepens the use of each pair. This is the broad-activation prediction.

These mechanisms are not mutually exclusive, but they differ in the relative size of the harvest-timed exchange and hiring responses and in whether adjustment occurs on the extensive or the intensive margin. Comparing the two contract-specific responses is therefore necessary for determining whether adjustment favors one form of settlement. A finding that exchange does not fall relative to hiring at harvest is evidence against a pure liquidity explanation.

The direction of exchange within the pair provides additional information about settlement. Mutual exchange within the same wave is consistent with households coordinating work across their fields over a short period. A one-way flow can represent either the creation or the repayment of an obligation outside the current recall window. When both households need labor at the same time, repayment may be delayed; when their production calendars differ, the pair may be able to exchange labor at separate dates. We therefore report both components without treating either as

a complete record of reciprocal obligations.

Finally, the baseline categories may capture the depth of the relationship rather than the types of contracts it previously supported. We examine this possibility by replacing the categories with continuous measures of baseline transaction frequency and dual-contract use. We also compare established pairs with previously untied providers, allowing adjustment to occur through new matches. These exercises can show whether prior relationships affect the allocation of transactions, although they cannot separately identify the roles of liquidity, monitoring, coordination, and insurance. Empirically, we measure the reallocation of labor transactions within the sampled village network following an unexpected change in crop returns.

4 Empirical Strategy

This section describes the empirical strategy. We begin with the dyadic estimating equation and explain the variation that remains after the fixed effects are included. We then construct the crop-return innovation from world prices and predetermined village crop shares and introduce the harvest calendar used to time this exposure. The section concludes with the robustness specifications and inference procedures.

4.1 Dyadic estimating equation

Let h index the potential receiver, j the potential provider, v the village, and w the survey wave. Let $m(w)$ denote the calendar month. For baseline pair type $r \in \{\text{exchange-only}, \text{dual}, \text{hire-only}\}$, D^r_{hj} indicates that undirected category. Baseline-untied pairs are the omitted group. We estimate

$$\begin{aligned} Y_{hjvw} = &\sum_r \beta_r \left(S_{hvw} D^r_{hj}\right) + \sum_r \theta_r \left(H_{hvw} D^r_{hj}\right) \\ &+ \alpha_{hj,m(w)} + \gamma_{hw} + \delta_{jw} + u_{hjvw}, \end{aligned} \tag{1}$$

where Y_{hjvw} is an indicator for any exchange, any hired labor, or one of the mutually exclusive within-wave outcomes defined in Section 2.2. The two exposure terms are

$$S_{hvw} = \frac{1}{0.1} \sum_k s^0_{vk} \varepsilon_{k,\tau(h,w)}, \qquad H_{hvw} = \frac{1}{0.1} \sum_k s^0_{vk} c^{-v}_{vk,m(w)} \varepsilon_{k,\tau(h,w)}, \tag{2}$$

with crop shares, price innovations, and the leave-one-village-out harvest calendar constructed in Sections 4.2 and 4.3. The pair-type-specific part of Equation (1) can then be rewritten as

$$\beta_r S_{hvw} + \theta_r H_{hvw} = \frac{1}{0.1} \sum_k s^0_{vk} \left(\beta_r + \theta_r c^{-v}_{vk,m(w)}\right) \varepsilon_{k,\tau(h,w)}. \tag{3}$$

The effect of a price innovation in crop k is therefore $\beta_r + \theta_r c^{-v}_{vk,m(w)}$. The harvest calendar is standardized to have mean zero and unit variance within village and crop, so β_r is the response

in a month of average harvest intensity and θ_r is the increment per standard deviation of harvest intensity. Because H_{hvw} is the same price innovation reweighted by harvest intensity rather than an independent shock, a return innovation changes S_{hvw} and H_{hvw} jointly. We therefore report the combined effect, $\beta_r + \theta_r c_v^{\text{peak}}$, evaluated at the village's peak harvest month and scaled to a one-standard-deviation movement in S_{hvw}. For each village, the peak month maximizes the baseline-crop-share-weighted average of the crop-specific standardized harvest calendars, so a village growing several crops has a peak defined by its full baseline crop basket. We average these village peak values over the eligible directed pair-waves, giving greater weight to villages that contribute more observations; the resulting pair-wave-weighted peak-calendar value is 2.11. All reported coefficients and standard errors are scaled by $\text{SD}(S_{hvw}) = 0.581$, so that the average-month response is $0.581\beta_r$, the harvest-timing increment is $0.581\theta_r$, and the peak-harvest total is $0.581(\beta_r + \theta_r c^{\text{peak}})$. This common rescaling leaves the p-values unchanged. As a transparency check, a binned specification replaces H_{hvw} with the interaction of S_{hvw} and an indicator for above-average harvest intensity.

The fixed effects make the identifying comparison quite demanding. The directed-pair-by-month effect, $\alpha_{hj,m(w)}$, absorbs the pair's usual propensity to transact in each month of the year. It therefore accounts for stable relationships, recurring agricultural seasons, and pair-specific seasonal patterns. The receiver-by-wave effect, γ_{hw}, absorbs all conditions shared by the potential providers of a given receiver at the time of the survey, including the level of the crop-return innovation and the receiver's total adjustment in labor use. The provider-by-wave effect, δ_{jw}, absorbs the provider's circumstances in that wave, such as migration, own-farm work, and opportunities to work for other households.

Because the innovation is common to all potential providers of a receiver in a given wave, its level is absorbed by γ_{hw}. Likewise, the main effects of baseline relationship type are absorbed by the pair-by-month fixed effects. The coefficients β_r and θ_r are identified from the interaction between the receiver's exposure and the predetermined type of each potential relationship. In practical terms, the regression tests whether a crop return innovation changes the probability of transacting with an established provider relative to an untied provider available to the same receiver in the same wave. The pair-by-month effects further require this difference to be measured relative to the relationship's usual behavior in that part of the agricultural calendar.

The identifying assumption is that, conditional on these fixed effects, no other shock simultaneously changes the crop-return measure and reallocates transactions toward the same baseline relationship types. The design does not require the crop-return innovation to be independent of the receiver's overall labor demand, since receiver-by-wave effects absorb that response. It requires the innovation to be unrelated to omitted determinants of the relative use of established and untied providers. The construction below addresses this requirement by using external price movements, fixing crop exposure before the outcome period, and assigning price information dated before the transaction recall window.

Equation (1) is estimated on the complete set of eligible pair-waves, including the zeros for potential providers who were not used. Conditioning on observed transactions would select on the

outcome and remove the alternatives needed to study partner choice. We use a linear probability model because its interaction coefficients can be read directly in percentage points and because it accommodates the high-dimensional fixed effects in a sample with very rare outcomes. Each eligible pair-wave receives equal weight. The resulting estimand is the change in the probability that a particular provider supplies a particular receiver, which we scale by the relevant relationship-specific transaction rate when discussing economic magnitude.

The two directions of an undirected pair are treated as separate observations. A household may receive labor from a partner in one period and provide labor to the same partner in another, while baseline relationship type remains common to the pair. Standard errors are clustered by village, the level at which baseline crop exposure is assigned, allowing outcomes to be correlated across both directions and across all pairs within the same village. With sixteen villages, conventional cluster inference is fragile. The primary p-values are therefore from a restricted wild-cluster bootstrap with Rademacher weights. We use 9,999 draws for the headline coefficients and 999 draws for the remaining terms, and we base significance stars in the main results tables on these p-values. Appendix Table D.3 retains village-clustered $t(15)$ stars, as stated in its notes. Two-way clustering by village and survey wave and time-HAC standard errors are reported in Table 5.

4.2 Bartik-type crop-return shock

The innovation combines fixed village crop specialization with unexpected movements in world crop prices. For crop k and price month t,

$$\varepsilon_{kt} = \Delta \log P_{kt}^{W} - \widehat{\mathbb{E}}_{t-1}\left[\Delta \log P_{kt}^{W}\right]. \tag{4}$$

The forecast is crop-specific and estimated recursively. It includes three lags of monthly price growth and annual sine and cosine terms. For every month, the estimation window ends at $t-1$, so later prices stay out of the forecast. The World Bank Pink Sheet series are Thai 5 percent rice, maize, soybeans, groundnuts, and sorghum.

We use forecast errors because persistent and seasonal changes in prices may already be incorporated into planting decisions, inventories, and labor arrangements. The recursive procedure ensures that each forecast uses only information available before the price movement occurs. Three lags of price growth capture short-run persistence, while the annual sine and cosine terms remove regular seasonality. The procedure uses a parsimonious information set, and a trader with richer information may forecast part of ε_{kt}. Identification requires that, after the fixed effects are included, the remaining innovation be uncorrelated with omitted shocks that differentially change transactions across baseline relationship types.

The timing of the innovation follows the receiver's interview history. Let $\tau(h, w)$ denote the last complete calendar month before the receiver's previous interview month. We assign the price innovation from this month to receiver-wave (h, w). Because $\tau(h, w)$ precedes the beginning of the receiver's transaction recall interval, the assigned price information is available before every

transaction included in the outcome. This restriction matters in a monthly survey. A transaction reported at the current interview may have occurred before a price recorded later in the same month, and even the preceding calendar month may overlap the recall interval when interviews are delayed. Receivers in the same survey wave can therefore be assigned different price months according to their interview dates. We retain recall intervals between 15 and 60 days, for which this ordering is credible.

Let s^0_{vk} be crop k's share of baseline crop participation in village v, measured in raw waves 1–24 and held fixed thereafter. The exposure is

$$S_{hvw} = \frac{1}{0.10} \sum_k s^0_{vk} \varepsilon_{k,\tau(h,w)}. \tag{5}$$

Crop participation shares are calculated as the fraction of household–crop records in the baseline period associated with crop k. They are less sensitive than output shares to a small number of unusually large farms. As a robustness check, Panel D of Table 5 weights crops by baseline harvested quantities. In both cases, the shares are fixed before the outcome period so that current planting and harvest decisions cannot affect exposure. We use village rather than household shares because crop histories are missing for many potential providers and may themselves be related to household conditions. The price innovations provide the time variation in S_{hvw}, while the baseline shares determine the relative importance of each crop to the village.

The measure is an index of returns to crops that were locally important before the treatment period. Household income remains a broader object because yields, sales, domestic marketing, and local spillovers also matter. Even households that do not grow the crops in the index may respond through the labor market or the village network. Moreover, receiver-by-wave fixed effects absorb the common effect of the index on each receiver. The coefficients in Equation (1) capture the allocation of transactions across relationship types following this local exposure.

We divide the weighted innovation by 0.10, so that a unit of S_{hvw} corresponds to a 10-log-point innovation for a fully exposed crop basket. In the estimation sample, the standard deviation of S_{hvw} is 0.581 and its 10th and 90th percentiles are -0.465 and 0.628. Both positive and negative innovations therefore contribute to identification. The variation is concentrated in a small number of crops: rice accounts for 53.8 percent of the absolute identifying variation and corn for 37.2 percent, yielding an effective crop count of 2.31. Figure 2 displays the support of the measure and the contribution of each crop. Because the normalized unit is larger than a typical movement in the data, all reported effects and standard errors are scaled using the main estimation-sample standard deviation of 0.581 (0.5806547 before rounding).

4.3 Harvest timing

The value of a crop-return innovation for labor demand is likely to depend on when it occurs in the production cycle. We therefore construct a fixed crop-by-month harvest calendar from the Townsend harvest records. For each focal village, the calendar uses the frequency with which

households in other villages in the same region harvest crop k in each month. When the regional sample for a crop-month cell is too small, we use a national leave-one-village-out frequency. The resulting calendar is standardized within village and crop and held fixed across years. Let c^{-v}_{vkm} denote this standardized measure. We define the harvest-timed exposure as

$$H_{hvw} = \frac{1}{0.10} \sum_{k} s^{0}_{vk} c^{-v}_{vk,m(w)} \varepsilon_{k,\tau(h,w)}. \tag{6}$$

Thus, a positive value of H_{hvw} indicates that a positive return innovation falls disproportionately on crops usually harvested in month $m(w)$. The calendar excludes the focal village because realized harvest dates may respond to weather, planting decisions, and local labor conditions. Using other villages in the same region retains common agro-climatic timing while preventing the focal village's current outcomes from entering the regressor. Holding the calendar fixed across years likewise ensures that current shocks do not alter treatment intensity. The standard deviation of H_{hvw} in the estimation sample is 0.499.

Harvest is the relevant window for the shock because a price surprise revalues the crop that is currently being harvested. That window may not coincide with the months in which exchange labor is most frequently used. Figure 1 shows that exchange peaks in April, around land preparation and planting, whereas the rice harvest calendar that accounts for most of the identifying variation in H_{hvw} peaks in November–January (Figure 2, Panel C). The April peak describes when households supply exchange labor. It does not define when a return innovation changes the value of the standing crop.

Appendix Table D.4 provides an alternative-calendar check. Because the data do not contain crop-specific planting dates, P_{hvw} is a timing proxy rather than a measured planting exposure. It reconstructs H_{hvw} after rotating the leave-one-village-out harvest calendar backward. The land-preparation proxy shifts the harvest profile backward by eight months, mapping a December rice harvest to April. The second proxy shifts the harvest profile backward by four months, mapping a December rice harvest to August. It provides a sensitivity check for an alternative pre-harvest window rather than a measure of observed planting dates. The standalone alternative-calendar increments are generally negative, but they should not be interpreted as negative planting effects because the harvest and land-preparation proxies are negatively correlated (-0.32). The informative comparison includes both timings together. In that specification, the harvest-timing coefficient on baseline-dual exchange is 0.78 percentage points ($p_{\text{WCB}} = 0.006$), close to the main estimate, while the land-preparation term is 0.04 percentage points ($p_{\text{WCB}} = 0.431$). The April labor peak therefore does not account for the positive harvest-timed response.

Appendix Figure B.2 compares the other-village calendar with harvests reported in the excluded village. For rice and corn, which together account for 91 percent of the shifter's identifying variation, the calendar closely tracks the timing of harvests during the baseline and in later waves.

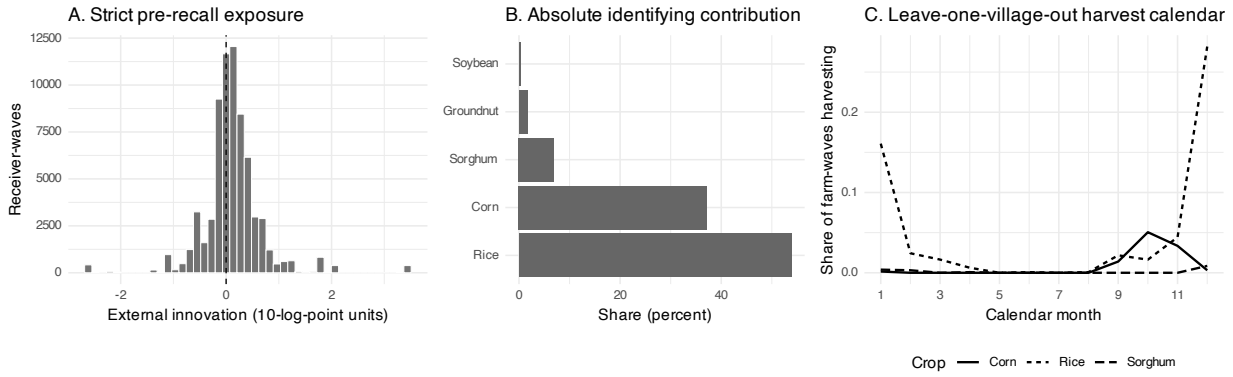


Figure 2: Support, crop composition, and harvest timing of the Bartik-type crop-return shock. Panel A shows the receiver-wave distribution of the baseline-share-weighted forecast innovation. Panel B reports each crop's normalized absolute contribution to identifying variation. Panel C reports the leave-one-village-out monthly harvest rate for the three crops with the largest relevant support.

The study villages are too small to affect world prices, which provides an external source of time variation. The main identifying assumption nevertheless concerns the interaction between these price innovations and baseline relationship type. After conditioning on the fixed effects, the innovations must be unrelated to other events that would independently shift transactions toward established providers. The design addresses the most immediate concerns through strict information timing, predetermined crop shares, and pair-specific seasonal effects. The future-innovation tests described below examine whether current outcomes are associated with price information that had not yet arrived.

Two limitations remain. First, the rice benchmark is a Thai export price and may itself respond to national policy or production conditions. We therefore decompose the estimates by crop and report results that exclude corn, which provides the second-largest source of variation. Second, villages specializing in rice may differ systematically from villages specializing in corn. Fixed baseline shares and pair-specific seasonality remove permanent differences in levels and regular seasonal patterns, while leave-one-village-out estimates reveal whether a single location accounts for the result. With sixteen villages, the data do not support a fully flexible response for every exposure pattern.

The estimates also incorporate equilibrium responses within the village network. A return innovation may raise a receiver's demand for labor and the opportunity cost of providers working on their own farms. Providers may then reallocate labor across receivers. Receiver-by-wave and provider-by-wave effects condition on the circumstances of both endpoints, so the estimated coefficients describe how transactions are distributed across the observed choice set. They should not be interpreted as a partial-equilibrium labor-demand elasticity.

Finally, precision is governed by the five price series and sixteen baseline village crop baskets that generate exposure. The dyadic panel improves the measurement of how a shock is allocated

across relationships, while the effective number of treatment units remains small.

4.4 Robustness specifications

We use four extensions of Equation (1) to assess the main threats to interpretation. First, we add three future values of the average innovation and three future values of the harvest-timed innovation, each interacted with baseline-dual status. These innovations were unavailable during the current recall interval, so their joint association with current transactions would indicate anticipation, residual seasonality, or another timing problem. Second, we reconstruct the two exposure measures using rice only, corn only, and all crops except corn. Third, we replace crop-participation weights with baseline harvested-quantity shares. Finally, we replace the categorical relationship indicators with the larger and smaller of the shares of baseline waves in which the pair used exchange or hiring. These measures describe activity under the two contracts and are distinct from their joint use within the same wave. The results of these exercises are reported in Section 5.4.

4.5 Inference and falsification

The main standard errors are clustered by village, since baseline crop exposure is assigned at that level. With only 16 villages, the usual cluster asymptotic approximation may be unreliable. Significance stars in the main tables are based on restricted wild-cluster bootstrap p-values. Two-way clustering by village and survey wave, a time-HAC covariance estimator with six wave lags, and leave-one-village-out ranges are reported in Table 5. These procedures address different forms of dependence, and we give greatest weight to results whose sign and magnitude are stable across them.

The future-innovation test is conducted separately for each outcome. For each contract, we jointly test the three future average innovations and three future harvest-timed innovations using an $F(6, 15)$ reference distribution. The test has limited power with sixteen exposure units, so a failure to reject is interpreted as evidence consistent with the timing restriction. A rejection indicates that the outcome should not be used to support a causal claim. We apply this criterion separately because combining exchange and hiring changes both the economic margin and the residual variation.

As a check on the first link in the proposed channel, we also examine whether the external innovations pass through to local farm-gate prices. The coefficient on baseline-share-weighted domestic price growth is 0.299, indicating that external price movements are reflected in the returns faced by local producers. Household labor and harvest outcomes are examined separately in Section 5.5 to determine whether the data also support a contemporaneous increase in physical labor demand.

5 Results

This section presents the effect of crop-return innovations on labor transactions across baseline relationships. We begin with the contract-specific estimates and the peak-harvest totals. We then

compare exchange and hiring in proportional terms and report the mutually exclusive within-wave outcomes. A second table separates mutual from one-way exchange and reports days conditional on transacting. The final two subsections report the validity checks and return to the mechanisms in Section 3.1.

5.1 Labor transactions across baseline relationships

Table 3 reports the coefficients from Equation (1) in percentage points. Each column is one regression. Columns (1) and (2) are the overlapping indicators for any exchange and any hired labor. Columns (3)–(5) are mutually exclusive: exchange only, hire only, and both contracts in the wave. In each case, the coefficient for an established pair is measured relative to a baseline-untied potential provider available to the same receiver in the same wave. All panels use the main estimation-sample standard deviation of S_{hvw}, 0.581. Panel A reports $0.581\beta_r$, the response in a month of average harvest intensity. Panel B reports $0.581\theta_r$, the additional response for a one-unit increase in the standardized harvest calendar. Panel C reports the object of interest: $0.581(\beta_r+\theta_r c^{\text{peak}})$, evaluated at the pair-wave-weighted village peak. The pair-wave-weighted peak calendar is 2.11.

For baseline-dual pairs, a one-standard-deviation positive return innovation at peak harvest raises the probability of exchange by 1.68 percentage points ($p = 0.036$). This increase is equivalent to 54 percent of their mean monthly exchange rate of 3.10 percent. The harvest-timing increment is 0.77 percentage points ($p = 0.006$) for the same one-standard-deviation movement in S_{hvw} when the standardized harvest calendar rises by one unit, indicating that the response becomes stronger as harvest intensity rises.

Baseline-exchange-only pairs also become more active at harvest. A one-standard-deviation positive return innovation at peak harvest raises exchange among those pairs by 1.08 percentage points ($p = 0.048$), or 29 percent of their mean monthly exchange rate. Although the estimated harvest-timing response is larger for baseline-dual pairs, the difference between the two pair types is not statistically significant ($p = 0.291$). We therefore find that prior exchange relationships respond, but cannot conclude that experience with both contracts produces a stronger response.

Baseline-hire-only pairs also show a suggestive exchange response. At peak harvest, a one-standard-deviation positive return innovation raises exchange by 0.26 percentage points ($p = 0.060$), equivalent to 50 percent of their 0.52 percent mean monthly exchange rate. Thus, relationships previously used only for hiring may also support subsequent exchange, although this evidence is weaker than for pairs with a history of exchange.

The average-month coefficients are smaller and, for baseline-dual pairs, imprecise. A binned specification that replaces H_{hvw} with an indicator for above-average harvest intensity produces the same qualitative pattern (Appendix Table D.3).

These coefficients describe relative allocation across providers. Receiver-by-wave effects absorb any change in the receiver's total transactions in that wave. A positive coefficient means that the crop-return innovation shifts transactions toward that type of established relationship relative to baseline-untied providers.

Table 3: Unexpected crop returns and labor transactions across baseline relationships

	(1) Any exchange	(2) Any hired	(3) Exchange only	(4) Hire only	(5) Both contracts
A. Average-month response: $S_{hvw}\times$ baseline type ($0.581\beta_r$)					
$S\times$ baseline-exchange-only	0.156	0.038**	0.142	0.026	0.013
	(0.080)	(0.018)	(0.080)	(0.017)	(0.013)
	[0.114]	[0.046]	[0.137]	[0.140]	[0.437]
$S\times$ baseline-dual	0.055	0.152*	0.035	0.133**	0.020
	(0.172)	(0.079)	(0.145)	(0.046)	(0.042)
	[0.705]	[0.099]	[0.717]	[0.033]	[0.661]
$S\times$ baseline-hire-only	0.053**	0.044	0.045**	0.037	0.008
	(0.024)	(0.027)	(0.021)	(0.026)	(0.009)
	[0.016]	[0.182]	[0.029]	[0.263]	[0.484]
B. Harvest-timing increment: $H_{hvw}\times$ baseline type ($0.581\theta_r$)					
$H\times$ baseline-exchange-only	0.437*	0.103**	0.395*	0.061*	0.041
	(0.178)	(0.039)	(0.170)	(0.027)	(0.028)
	[0.051]	[0.025]	[0.053]	[0.050]	[0.274]
$H\times$ baseline-dual	0.769***	0.295	0.627***	0.153	0.142
	(0.303)	(0.172)	(0.260)	(0.102)	(0.088)
	[0.006]	[0.109]	[0.010]	[0.172]	[0.148]
$H\times$ baseline-hire-only	0.098	0.129***	0.083	0.114**	0.015
	(0.059)	(0.041)	(0.053)	(0.038)	(0.018)
	[0.103]	[0.005]	[0.124]	[0.010]	[0.404]
C. Total effect at peak harvest: $\mathrm{SD}(S)(\beta_r + \theta_r c^{\mathrm{peak}})$					
Baseline-exchange-only	1.079**	0.255**	0.978*	0.155*	0.100
	(0.444)	(0.093)	(0.425)	(0.068)	(0.071)
	[0.048]	[0.021]	[0.060]	[0.060]	[0.297]
Baseline-dual	1.680**	0.776	1.361**	0.457*	0.320
	(0.793)	(0.433)	(0.678)	(0.250)	(0.223)
	[0.036]	[0.101]	[0.037]	[0.091]	[0.188]
Baseline-hire-only	0.260*	0.317**	0.220*	0.278**	0.040
	(0.136)	(0.107)	(0.127)	(0.100)	(0.040)
	[0.060]	[0.011]	[0.093]	[0.019]	[0.340]
D. Estimation-sample mean by baseline pair type (%)					
Baseline-exchange-only	3.75	0.58	3.62	0.46	0.12
Baseline-dual	3.10	1.68	2.86	1.43	0.25
Baseline-hire-only	0.52	0.70	0.46	0.65	0.05
Baseline-untied	0.15	0.10	0.15	0.09	0.01
E. Tests (wild-cluster bootstrap p-values)					
$\theta_{\mathrm{dual}} = \theta_{\mathrm{exchange\text{-}only}}$	0.291	0.259	0.371	0.468	0.198
$\theta_r^{\mathrm{Ex}} = \theta_r^{\mathrm{Hire}}$, columns (1) vs (2):	Exchange-only: 0.080	Dual: 0.109	Hire-only: 0.660		
Directed pair × month-of-year FE	Yes	Yes	Yes	Yes	Yes
Receiver × wave FE	Yes	Yes	Yes	Yes	Yes
Provider × wave FE	Yes	Yes	Yes	Yes	Yes
Pair-wave observations	2,629,888	2,629,888	2,629,888	2,629,888	2,629,888
Clusters (villages)	16	16	16	16	16

Notes: Coefficients and standard errors use the main estimation-sample scaling $\mathrm{SD}(S) = 0.581$ (0.5806547 before rounding). Average-month effects refer to a one-standard-deviation movement in S; continuous harvest-timing increments refer to the same movement for a one-unit increase in the standardized harvest calendar. Each column is one estimation of equation (1). The dependent variable is an indicator for the transaction in the receiver's recall interval, multiplied by 100, so coefficients are in percentage points. Columns (3)–(5) are mutually exclusive; 0 means no labor transaction of that type between the pair in the wave. S_{hvw} is the Bartik-type crop-return innovation and H_{hvw} its harvest-weighted version; $\mathrm{SD}(S) = 0.581$ and $\mathrm{SD}(H) = 0.499$. Panel C evaluates $\mathrm{SD}(S)(\beta_r + \theta_r c^{\mathrm{peak}})$ at the village peak harvest month for a one-SD movement in S, where the pair-wave-weighted $c^{\mathrm{peak}} = 2.11$. Baseline-untied pairs are the omitted category. Village-clustered standard errors in parentheses; restricted wild-cluster bootstrap p-values in brackets when available (9,999 Rademacher draws for headline coefficients; 999 otherwise). Stars are based on the reported wild-cluster bootstrap p-values. Panel E tests equality of the coefficients; common rescaling leaves these tests unchanged.

5.2 Broad activation rather than an exchange-specific response

The results also point to an increase in hiring. On baseline-dual pairs, the peak-harvest total for any hired labor is 0.78 percentage points ($p = 0.101$), or 46 percent of the 1.68 percent hiring rate. On baseline-exchange-only pairs, hiring rises by 0.26 percentage points ($p = 0.021$), or 44 percent of the 0.58 percent hiring rate—a larger proportional response than that pair type's exchange response. Baseline-hire-only pairs increase hiring by 0.32 percentage points ($p = 0.011$), or 45 percent of their 0.70 percent rate, and their exchange response is of similar proportion. We therefore do not find evidence of substitution from paid to reciprocal labor as both exchange and hiring increase following a positive harvest-timed return innovation.

Columns (3)–(5) distinguish greater use of a single contract from use of both contracts in the same wave. The point estimates place most of the exchange increment in exchange-only transactions. On baseline-dual pairs, 0.63 of the 0.77 percentage-point harvest-timing increment in any exchange is exchange only; both contracts in the same wave account for 0.14 percentage points and are imprecise. These increments refer to a one-standard-deviation movement in S_{hvw} for a one-unit increase in the standardized harvest calendar. Same-wave use of both contracts is rare (0.25 percent of baseline-dual pair-waves) and accounts for a smaller, imprecisely estimated share of the exchange increment. By construction, column (3) plus column (5) recovers column (1), and column (4) plus column (5) recovers column (2).

Taken together, the results favor a broad activation of established labor ties. Pairs with an exchange history make greater use of exchange and also hire one another more often, in similar proportion to their baseline rates. Hire-only relationships respond primarily through hiring, again in similar proportion. We do not detect a relative decline in exchange at harvest. The point estimates do not favor a pure liquidity explanation, although the imprecise contract comparison does not rule it out.

5.3 Mutual exchange, one-way exchange, and the intensive margin

The institutional evidence in Section 2.3 shows that exchange is concentrated in seasonal field tasks and that obligations often remain unbalanced over long periods. Table 4 reports two further margins. Columns (1) and (2) divide exchange into mutual and one-way flows on the full pair-wave sample. Mutual exchange requires both households to report supplying labor to one another in the same wave. One-way exchange records labor moving from the provider to the receiver without a contemporaneous reverse flow. Columns (3) and (4) report person-days in the recall interval, winsorized at the 99th percentile, among pair-waves that transact. Those columns use pair fixed effects with village-by-wave and month-of-year effects, because receiver-by-wave effects would absorb most of the remaining variation in this smaller sample. Since the shock can change which pairs transact, the day estimates condition on the outcome and are descriptive.

Both components of exchange have positive harvest-timing increments on baseline-dual pairs. For a one-standard-deviation movement in S_{hvw} and a one-unit increase in the standardized harvest calendar, the harvest-timing increment is 0.39 percentage points for mutual exchange ($p_{\text{WCB}} =$

Table 4: Mutual and one-way exchange, and days conditional on transacting

	(1) Mutual exchange	(2) One-way exchange	(3) Exchange days	(4) Hired days
A. Average-month response ($0.581\beta_r$)				
$S\times$ baseline-exchange-only	0.053	0.103	-0.008	-0.084
	(0.069)	(0.081)	(0.039)	(0.067)
	[0.452]	[0.223]	[0.871]	[0.411]
$S\times$ baseline-dual	-0.001	0.056	-0.075	-0.059
	(0.129)	(0.081)	(0.051)	(0.093)
	[0.998]	[0.542]	[0.138]	[0.525]
$S\times$ baseline-hire-only	0.037***	0.016	-0.084	0.001
	(0.013)	(0.020)	(0.096)	(0.074)
	[0.007]	[0.435]	[0.701]	[0.981]
B. Harvest-timing increment ($0.581\theta_r$)				
$H\times$ baseline-exchange-only	0.231*	0.206	0.056	0.180**
	(0.102)	(0.128)	(0.037)	(0.063)
	[0.061]	[0.129]	[0.178]	[0.020]
$H\times$ baseline-dual	0.388*	0.381**	-0.024	-0.044
	(0.185)	(0.169)	(0.056)	(0.127)
	[0.090]	[0.019]	[0.696]	[0.747]
$H\times$ baseline-hire-only	0.066*	0.032	0.046	0.066
	(0.034)	(0.056)	(0.070)	(0.079)
	[0.060]	[0.576]	[0.576]	[0.448]
C. Estimation-sample mean by baseline pair type				
Baseline-exchange-only	1.24	2.50	—	—
Baseline-dual	1.03	2.07	—	—
Baseline-hire-only	0.09	0.42	—	—
Directed pair × month-of-year FE	Yes	Yes	No	No
Pair FE; village × wave; month-of-year	No	No	Yes	Yes
Receiver × wave FE	Yes	Yes	No	No
Provider × wave FE	Yes	Yes	No	No
Observations	2,629,888	2,629,888	10,116	4,752
Clusters (villages)	16	16	16	16

Notes: Coefficients and standard errors use the main estimation-sample scaling SD(S) = 0.581 (0.5806547 before rounding). Average-month effects refer to a one-standard-deviation movement in S; continuous harvest-timing increments refer to the same movement for a one-unit increase in the standardized harvest calendar. Columns (1)–(2) use Equation (1) on the full pair-wave sample and are in percentage points. Mutual exchange requires a reverse flow in the same wave; one-way exchange is the remainder of any exchange. Columns (3)–(4) are descriptive linear models for person-days in the recall interval, winsorized at the 99th percentile, estimated only on pair-waves with a transaction of that type. Mean winsorized days are 2.00 for exchange and 1.62 for hiring. Because the shock can change which pairs transact, columns (3)–(4) condition on the outcome. Village-clustered standard errors are in parentheses and restricted wild-cluster bootstrap p-values are in brackets (999 Rademacher draws). Stars use the bootstrap p-values.

0.090) and 0.38 percentage points for one-way exchange ($p_{\text{WCB}} = 0.019$), so the point estimates each account for approximately half of the harvest-timing increment in exchange, although only the one-way component is precisely estimated at the five-percent level. The mutual component is consistent with neighboring households coordinating work across their fields during a concentrated task period. The one-way component is consistent with obligations that are created or repaid outside the current survey wave. The decomposition therefore matches the descriptive evidence that labor exchange combines immediate cooperation with reciprocity over longer horizons.

The conditional day estimates provide no clear evidence that transactions become longer on baseline-dual pairs following a harvest-timed return innovation. The point estimates are small and negative for both exchange days and hired days, and they are imprecise. On baseline-exchange-only pairs, the harvest-timing increment in hired days is 0.18 days per one-standard-deviation movement in S_{hvw} and one-unit increase in the standardized harvest calendar among the pair-waves that hire ($p_{\text{WCB}} = 0.020$), while exchange days are essentially unchanged. The extensive-margin results show activation of established pairs, while the conditional day estimates provide no clear evidence of longer transactions. Because these estimates condition on transacting, they do not establish how days supplied by a fixed set of partners change or rule out coordination through larger work parties. A Poisson pseudo-maximum likelihood model for unconditional days would give proportional effects without conditioning on transacting; we leave that estimator as a direction for further work.

This interpretation of the mutual and one-way split is subject to a measurement qualification. The two households may be interviewed on different dates and report transactions over recall periods that only partly overlap. Some transactions classified as one-way may therefore have a reverse flow immediately outside the common window. The decomposition measures reciprocity within the observed period and leaves the eventual balance of the obligation unmeasured.

5.4 Validity and robustness

Table 5 reports the checks that bear most directly on the harvest-timed increment among baseline-dual pairs. Alternative inference procedures and leave-one-village ranges are collected here; the main tables base stars on the restricted wild-cluster bootstrap. Panel B tests whether future innovations predict current transactions. Panels C and D examine the crop composition of the shock and the choice of baseline weights.

The baseline-dual exchange increment is positive across all inference procedures in Panel A. Its restricted wild-cluster p-value is 0.006, compared with 0.023 under the village-clustered $t(15)$ reference and 0.055 under two-way clustering. The time-HAC procedure also rejects at conventional levels. When each village is omitted in turn, the harvest-timing increment ranges from 0.50 to 0.87 percentage points per one-standard-deviation movement in S_{hvw} and one-unit increase in the standardized harvest calendar. Its sign and approximate magnitude are therefore stable across villages. By comparison, the wild-cluster p-values are 0.109 for hiring and 0.109 for the equal-contract test. These results reinforce the distinction in the main table: the increase in exchange is robust, while the relative size of the exchange and hiring responses remains uncertain.

Panel B examines the timing of information. The six future terms are jointly insignificant for exchange ($p = 0.229$), hiring ($p = 0.340$), and their difference ($p = 0.243$). These outcomes are therefore consistent with the requirement that the relevant price information precede the transaction. The pooled indicator for any labor fails the same test ($p = 0.002$). We exclude that outcome from the causal interpretation because future innovations predict whether either type of transaction is observed. This failure may reflect low-frequency changes in module coverage or broader co-movement in contract use that becomes more apparent after aggregation.

The crop decomposition in Panel C shows that the positive exchange response is concentrated in rice. When corn is excluded, the dual-pair harvest-timing increment rises to 0.93 percentage points; using rice alone gives an increment of 0.95 percentage points. Both are scaled by the main estimation-sample standard deviation of S_{hvw}, 0.581, and refer to a one-unit increase in the standardized harvest calendar. The exchange-minus-hiring estimates are also positive in these specifications. Corn alone produces a negative exchange coefficient. Since rice accounts for 53.8 percent of identifying variation and corn for 37.2 percent, the pooled estimate mainly describes rice-intensive production. Differences in production technology, task timing, and local marketing may all contribute to this heterogeneity.

Panel D replaces crop-participation weights with baseline quantity shares. The dual-pair exchange increment is 0.72 percentage points, with p-values of 0.034 under village clustering and 0.072 under two-way clustering. The hiring increment is 0.30 percentage points, and exchange exceeds hiring by 0.42 percentage points. These estimates use the same 0.581 scaling as the main results and refer to a one-unit increase in the standardized harvest calendar. These estimates preserve the ordering and approximate magnitude of the main results. Appendix Table D.1 then replaces the baseline categories with continuous measures of relationship activity. A ten-percentage-point increase in baseline labor intensity, defined as the larger of the exchange and hiring active-wave shares, is associated with positive harvest-timed responses in both outcomes, but the estimates are imprecise. The additional measure of dual-contract intensity, defined as the smaller of the two shares, also adds little precision. The categorical classification therefore appears to contain information about the relationship beyond the frequency with which the pair transacted.

5.5 Economic interpretation and mechanism limits

The results are closest to the broad-activation prediction in Section 3.1. Established relationships become more active when positive crop-return innovations arrive during harvest periods, and hiring rises in similar proportion to exchange. We do not detect a relative decline in exchange at harvest. The point estimates do not favor a pure liquidity explanation, although the imprecise contract comparison does not rule it out. Mutual and one-way exchange each account for about half of the harvest-timing increment in exchange, so the data are also consistent with some coordination of lumpy tasks, although the conditional day estimates provide no clear evidence of longer transactions. Information and monitoring with known partners can produce this combination: a household may solicit more pairs under both contracts, without a more intensive use of one contract type in

particular.

The institutional evidence helps put this result in context. Exchange in the study villages is concentrated in seasonal field-crop tasks, and reciprocal obligations often extend beyond a single month or year. Baseline-dual pairs have already demonstrated an ability to transact under both cash and reciprocal arrangements. Their response therefore fits a setting in which established partners draw on a range of contractual options when timely agricultural work becomes more valuable.

The household-level outcomes in Appendix Table D.2 provide less support for a contemporaneous increase in aggregate labor demand. An increase of 0.581 units in harvest-timed exposure, using the same scaling as the main results, is associated with 0.41 additional own-farm days under village clustering, although this estimate becomes imprecise with two-way clustering. Paid-work days are essentially unchanged. The probability of a recorded harvest falls by 0.75 percentage points, recorded harvest quantity also declines, and receiving sampled hired labor increases. Household-level receipt of sampled exchange is close to zero. Because these outcomes come from different modules and samples, we may expect them to be somewhat divergent. Still, they do not provide a consistent sequence from higher crop returns to more physical production and greater total labor use.

This evidence limits the interpretation of the dyadic estimates. Equation (1) conditions on the receiver's circumstances in each wave and tests which provider-contract cell is used. A household can reallocate transactions across sampled partners, substitute between sampled and off-roster workers, or change the form of settlement without changing its total labor input. We therefore refer to the treatment as a harvest-timed return innovation; the household outcomes provide insufficient evidence for a labor-shortage interpretation.

An attempt to separate the receiver's crop exposure from the provider's exposure is also inconclusive. Baseline household crop baskets are available for only 57 percent of eligible households, and future common-pressure terms predict current exchange and hiring in that smaller sample. We therefore retain the village-level exposure, for which the timing tests perform better, and do not decompose the result into a receiver demand effect and a provider opportunity-cost effect.

The main finding is consequently a change in the allocation of transactions within the sampled village network. The estimates do not reveal whether households substitute away from off-roster workers, whether familiar providers supply greater effort, or whether total village labor use increases.

Dropping the receiver-by-wave effects and collapsing to the receiver-wave, Appendix Table D.5 asks whether sampled-network use expands at all. These regressions absorb common survey-wave shocks as well as each receiver's usual month-of-year pattern. At peak harvest, a one-standard-deviation movement in S_{hvw} is associated with 0.036 additional exchange partners ($p_{\text{WCB}} = 0.156$) and 0.026 additional hired partners ($p_{\text{WCB}} = 0.004$). These changes equal 25 and 39 percent of their respective receiver-wave means. The probability of receiving hired labor rises by 1.58 percentage points, while the exchange-receipt indicator is close to zero. The receiver-level evidence supports an increase in sampled hired partners. The exchange-partner estimate is positive but imprecise, and

the exchange-receipt indicator provides no clear evidence of an increase. It still does not address off-roster workers or total village labor.

Among pair-waves that transact, a selected sample, the share of exchange shows no precise response to harvest-timed innovations for pairs with an exchange history (Appendix Table D.6).

Table 5: Interpretation, falsification, and robustness of the harvest-timed response

	Estimate/F	Village p	Two-way p	Time-HAC p	Wild p	Leave-one-village range
Panel A. Incremental harvest-timed response on baseline-dual pairs (percentage points)						
Any exchange	0.769	0.023	0.055	0.002	0.006	[0.498, 0.869]
Any hired labor	0.295	0.106	0.117	< 0.001	0.109	[0.215, 0.401]
Exchange minus hired labor	0.474	0.102	0.163	0.074	0.109	[0.251, 0.607]
Mutual exchange	0.388	0.053	0.090	0.009	0.090	—
One-way exchange	0.381	0.039	0.092	0.037	0.019	—
Panel B. Six future-innovation terms jointly zero						
Any exchange	1.550	0.229	—	—	—	—
Any hired labor	1.241	0.340	—	—	—	—
Any labor	5.923	0.002	—	—	—	—
Exchange minus hired labor	1.503	0.243	—	—	—	—
Panel C. Crop composition: incremental harvest-timed response on baseline-dual pairs						
Leave corn out: Any exchange	0.933	0.012	0.023	< 0.001	—	—
Leave corn out: Exchange minus hired labor	0.660	0.037	0.046	0.034	—	—
Rice only: Any exchange	0.955	0.012	0.023	< 0.001	—	—
Rice only: Exchange minus hired labor	0.661	0.039	0.054	0.042	—	—
Corn only: Any exchange	-0.194	0.409	0.041	0.713	—	—
Corn only: Exchange minus hired labor	-0.699	0.136	0.123	0.253	—	—
Panel D. Quantity-share exposure: incremental harvest-timed response on baseline-dual pairs						
Any exchange	0.724	0.034	0.072	—	—	—
Any hired labor	0.300	0.053	0.068	—	—	—
Exchange minus hired labor	0.424	0.162	0.234	—	—	—

Notes: Panels A, C, and D report harvest-timing increments in percentage points using the common main-sample scaling $\mathrm{SD}(S) = 0.581$; the leave-one-village ranges use the same scale. Panel B reports unscaled $F(6, 15)$ statistics for three future average and three future harvest-timed baseline-dual terms. Two-way inference clusters by village and survey wave; time-HAC uses six wave lags. Restricted wild-cluster p-values use 9,999 Rademacher draws for the headline any-exchange, any-hired, and equal-contract tests and 999 draws for mutual and one-way exchange. Panel C rebuilds the participation-weighted shock by crop; Panel D uses harvested-quantity shares. The any-labor outcome fails the future-innovation test and is excluded from the causal interpretation.

6 Conclusion

We study how rural households allocate paid and reciprocal labor across partners when crop returns change. Evaluated at peak harvest, a one-standard-deviation return innovation raises exchange on baseline-dual pairs by 54 percent of their monthly rate and hiring by 46 percent. The corresponding proportional hiring responses on baseline-exchange-only and baseline-hire-only pairs are about 44 and 45 percent. The harvest-timing increment is larger in percentage points for exchange because exchange has the higher base rate; the test of equal contract responses is imprecise, as is the test that baseline-dual pairs respond more than baseline-exchange-only pairs. Established ties therefore influence where additional transactions occur, and they do so through a broad activation of both contracts rather than through an exchange-specific shift.

The first contribution is the identification of a relational margin of rural labor-market adjustment. Most evidence on agricultural labor shocks concerns wages, employment, migration, or total household labor use. The bilateral Townsend records reveal the identity of both parties and allow each observed provider to be compared with other providers available to the same household in the same wave. This approach separates a change in total labor use from a change in the partners and contracts through which labor is obtained. The distinction matters because adjustment can occur through the reactivation of familiar providers even when household totals or village wages change little.

Our second contribution concerns the relationship between informal exchange and wage labor. The two contracts overlap within the same household pairs, and both remain responsive inside relationships that already support paid work. Reciprocal ties can thus remain useful in an active wage market and can also support paid work. This evidence qualifies accounts of commercialization based solely on the declining aggregate share of labor exchange. Recorded exchange falls over the panel, while surviving labor relationships continue to organize both unpaid and paid work during agricultural peaks. Market development may change how often and how flexibly existing ties are used before those ties disappear.

The analysis also provides a way to study this margin when transaction dates differ across households. World-price forecast errors are combined with predetermined village crop shares and a leave-one-village-out harvest calendar, and price information is aligned with the beginning of each household's transaction recall interval. This construction links return innovations to the crops and months in which timely work is most relevant without using the focal village's realized harvest timing. More broadly, it illustrates how high-frequency transaction records can be used to study contractual adjustment within networks rather than treating each household as an isolated labor-market participant.

The relational margin is relevant for changes in credit, transportation, migration, and mechanization. Each can alter the cost of cash settlement, the providers a household can reach, or the value of completing a task quickly. Responses may appear in partner choice and contract form before they appear in average employment or wages. Access to these relationships is also uneven: labor exchange is concentrated among households with land and labor available to reciprocate.

Networks may reduce search and enforcement costs for their members while limiting opportunities for workers who lack an established connection.

The main challenge to the findings is that the evidence comes from sixteen villages, with most identifying variation supplied by rice, and the estimates concern allocation within the sampled network. The household outcomes do not establish an increase in aggregate labor demand, and transactions with providers outside the survey roster are only partially observed. An alternative-calendar proxy for planting and land preparation does not reproduce the activation pattern; harvest timing remains when the two timings enter together. Receiver-wave estimates support an increase in sampled hired partners, while the exchange-partner estimate is imprecise. These results do not measure off-roster or village-wide labor. Future research could follow complete employer–worker networks across village boundaries and record hours, tasks, output, compensation, referrals, and the eventual settlement of reciprocal obligations. Such data would permit direct estimates of how relational contracting affects productivity and welfare. Studying network formation around expansions in roads, migration, credit, or mechanization would also show which households gain access to established labor ties, which are displaced, and whether reciprocal relationships facilitate entry into paid employment. These questions are necessary for understanding how rural labor markets become more monetized while continuing to rely on relationships.

References

Arciniegas, Christian, Christelle Dumas, and Matthias Fahn, "Informal Labor Exchange Teams and Participation in the Labor Market: Evidence from Rural Tanzania," Working Paper 11809, CESifo 2025.

Breza, Emily and Supreet Kaur, "Labor Markets in Developing Countries," Working Paper 33908, National Bureau of Economic Research 2025.

Casaburi, Lorenzo and Jack Willis, "Time versus State in Insurance: Experimental Evidence from Contract Farming in Kenya," *American Economic Review*, 2018, *108* (12), 3778–3813.

de Janvry, Alain, Claire Duquennois, and Elisabeth Sadoulet, "Labor calendars and rural poverty: A case study for Malawi," *Food Policy*, 2022, *109*, 102255.

Erasmus, Charles J, "Culture, structure, and process: The occurrence and disappearance of reciprocal farm labor," *Southwestern Journal of Anthropology*, 1956, *12* (4), 444–469.

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

Fujimoto, Takeshi, "From Festive to Exchange Labor: Changing Cooperative Labor among the Malo, Southwestern Ethiopia," *Nilo-Ethiopian Studies*, 2013, *18*, 21–36.

Gilligan, Daniel O., "The Economics of Agricultural Labor Exchange with Evidence from Indonesia." PhD dissertation, University of Maryland, College Park 2004.

Hossain, Mahabub, "Sustaining food security in Asia: Economic, social, and political aspects," in Noreen G. Dowling, Stanley Marshall Greenfield, and K. S. Fischer, eds., *Sustainability of Rice in the Global Food System*, International Rice Research Institute, 1997, pp. 19–36.

Hull, James R., "Migration, Remittances and Monetization of Farm Labor in Subsistence Sending Areas," *Asian and Pacific Migration Journal*, 2007, *16* (4), 451–484.

Jackson, Matthew O, *Social and economic networks*, Princeton University Press, 2008.

Jayachandran, Seema, "Selling labor low: Wage responses to productivity shocks in developing countries," *Journal of Political Economy*, 2006, *114* (3), 538–575.

Kaboski, Joseph P and Robert M Townsend, "The impact of credit on village economies," *American Economic Journal: Applied Economics*, 2012, *4* (2), 98–133.

Kaur, Supreet, "Nominal wage rigidity in village labor markets," *American Economic Review*, 2019, *109* (10), 3585–3616.

Kharel, Arjun, Fraser Sugden, and Mengistu Dessalegn, "Labor Shortage and Changes in Land Use Patterns: Experiences from Migrant Communities in Ethiopia and Nepal," AGRUMIG Policy Brief 11, International Water Management Institute 2023.

Kranton, Rachel E, "Reciprocal exchange: A self-sustaining system," *American Economic Review*, 1996, *86* (4), 830–851.

Munshi, Kaivan, "Networks in the modern economy: Mexican migrants in the US labor market," *The Quarterly Journal of Economics*, 2003, *118* (2), 549–599.

Olken, Benjamin A. and Monica Singhal, "Informal Taxation," *American Economic Journal: Applied Economics*, 2011, *3* (4), 1–28.

Parnwell, Michael, "Crowding out and clawing back: Local institutions and sustainable livelihoods in North-East Thailand," *International Development Planning Review*, 2005, *27* (2), 143–168.

Phelinas, Pascale, "Farm and off-farm employment in rural Thailand," *Journal of Social Research*, 1994, *17* (1).

Rosenzweig, Mark R, "Risk, implicit contracts and the family in rural areas of low-income countries," *The Economic Journal*, 1988, *98* (393), 1148–1170.

Shigetomi, Shin'ichi, *Cooperation and Community in Rural Thailand: An Organizational Analysis of Participatory Rural Development*, Tokyo: Institute of Developing Economies, 1998.

Townsend, Robert M, *Financial systems in developing economies: Growth, inequality, and policy evaluation in Thailand*, Oxford University Press, 2011.

_ , "Village and larger economies: The theory and measurement of the Townsend Thai project," *Journal of Economic Perspectives*, 2016, *30* (4), 199–220.

Townsend, Robert M., Sombat Sakunthasathien, and Rob Jordan, *Chronicles from the Field: The Townsend Thai Project*, Cambridge, MA: MIT Press, 2013.

A Data and Measurement

This appendix provides additional information on the study area and the construction of the sample. It reports the location and baseline crop composition of the villages, evidence on labor exchange from the 1997 key-informant survey, and changes in transaction-module coverage over the panel.

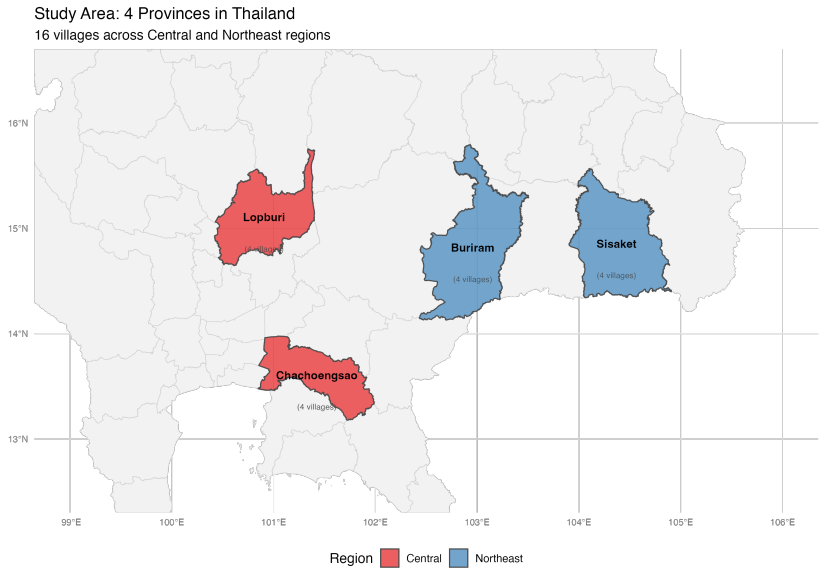


Figure A.1: Study provinces in Thailand.

Table A.2: Baseline crop specialization by province (output shares, 1998–99)

Province	Primary crop	Secondary crops
Chachoengsao	Rice (82%)	Aquaculture (18%)
Buri Ram	Rice (98%)	Aquaculture (2%)
Lop Buri	Corn (88%)	Sorghum (2%), Groundnut (2%), Rice (2%)
Si Sa Ket	Rice (93%)	Cassava (7%)

Notes: Shares are computed from 1998–99 harvested output and averaged across villages within each province. They describe baseline production and are not the prevalence weights used in the main treatment. Crops with output shares below 1% are omitted from the secondary-crop column.

Table A.3: 1997 key-informant evidence on village labor institutions

Village-level measure	Value
Villages where an agricultural cooperative had ever served members or customers	54.2%
Villages where a labor-exchange group had ever served members or customers	0.0%
Villages where a farmers' group had ever served members or customers	14.7%
Villages reporting farmer cooperation in acquiring inputs	40.0%
Villages reporting farmer cooperation in production	41.6%
Villages reporting farmer cooperation in selling or transporting output	21.6%
Mean village share using labor exchange in agriculture	26.4%
Median village share using labor exchange in agriculture	10.0%
Mean village share hiring other villagers for agricultural work	45.4%
Median village share hiring other villagers for agricultural work	50.0%

Notes: The table uses the Townsend Thai Initial Key Informant Survey, collected in 1997 before the monthly household panel. The unit of observation is the village, and the survey covers 192 villages in the four Townsend provinces. Labor-source questions ask what percentage of village households used each labor arrangement for agricultural activities in a year with adequate water.

Table A.4: Household labor activity and recorded dyadic labor links

Measure	1998–2008	2010–2012	Post/pre
Household activity module			
Average households observed per month	680	675	0.99
Own-farm person-days per household-month	19.6	27.9	1.42
Paid-labor person-days per household-month	16.9	19.7	1.17
Own-farm participation	51.7%	57.6%	1.11
Transaction-network module			
Average households observed per month	472	408	0.86
Exchange participation	8.0%	1.5%	0.19
Active labor pairs per month	129.3	12.2	0.09

Notes: The table compares monthly averages before and after 2009. Household activity outcomes are person-days from the labor-activities module. Exchange participation is measured at the household-month level. Active labor pairs require both households to be observed in the sampled transaction network. The year 2009 is omitted as a transition year. The final column reports the ratio of the 2010–2012 average to the 1998–2008 average.

B Labor Calendar and Wages

This appendix presents the labor calendar, wage patterns, and vegetation comparison summarized in Section 2.3. It also evaluates the crop-specific harvest calendar used in Section 4.3. The vegetation and labor-calendar measures are excluded from the construction of the crop-return shock.

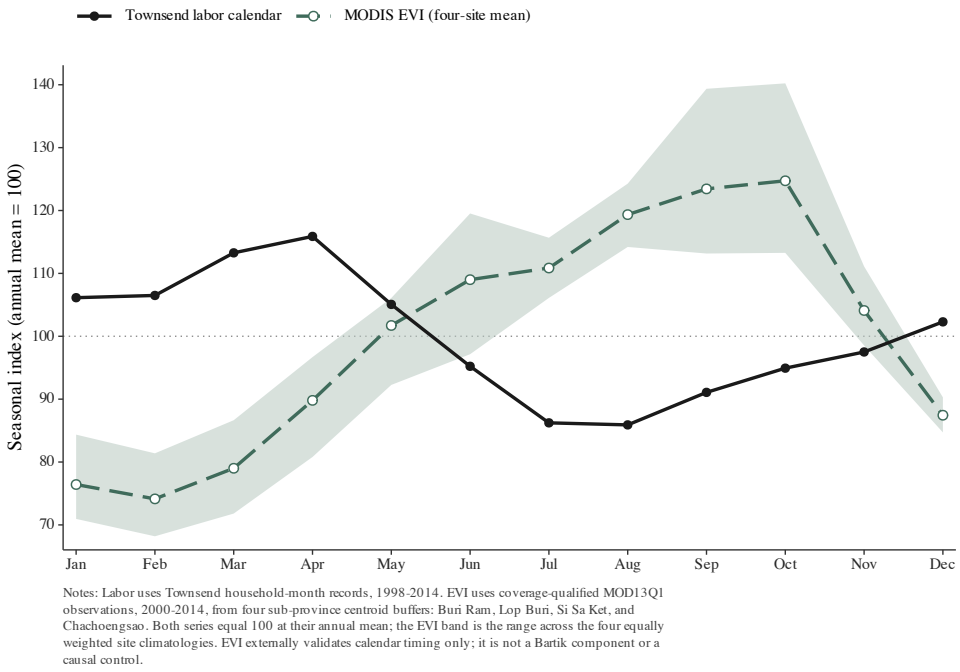


Figure B.1: Townsend labor calendar and MODIS EVI. Both series are indexed to an annual mean of 100. EVI validates seasonal timing and is excluded from the crop-return shock.

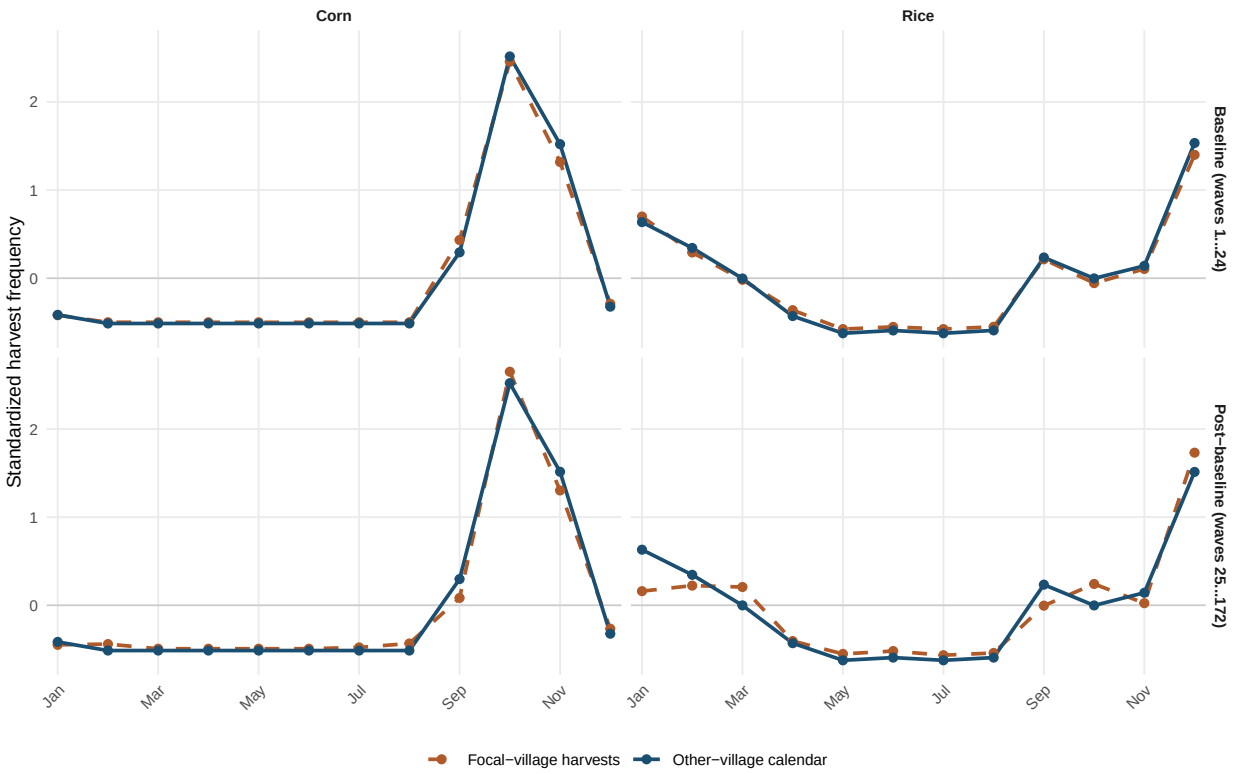
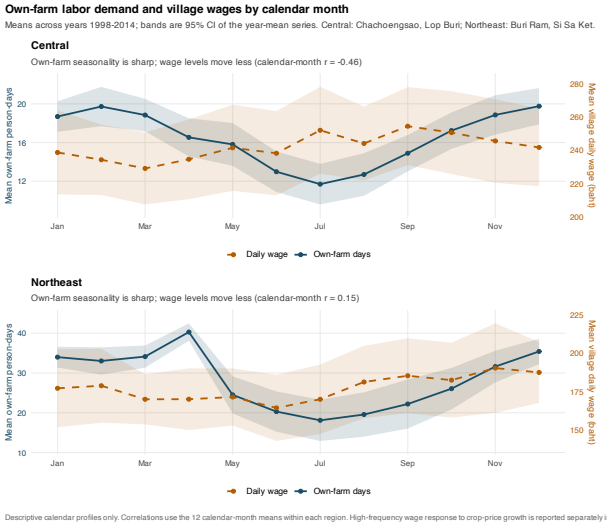


Figure B.2: Validation of the leave-one-village-out harvest calendar. The solid line is the crop-specific monthly calendar constructed from other villages in the same region during baseline waves 1–24. The dashed line is the harvest frequency reported in the excluded village. Each series is standardized within village and crop before averaging, with weights given by baseline crop participation and the number of observed household-waves. The upper panels use harvests reported during the baseline; the lower panels apply the fixed baseline calendar to harvests reported in later waves. Rice is observed in 15 villages and corn in five. The weighted correlations between the raw predicted and observed rates are 0.88 for rice and 0.85 for corn during the baseline, and 0.90 and 0.80, respectively, in later waves.



(a) Labor-market conditions by month



(b) Own-farm work and wages

Figure B.3: Wages and farm labor across the agricultural calendar, 1998–2014.

Table B.4: Seasonal labor calendars and low-work months

Month	Calendar class	Total	Farm	Off-farm	Any exchange (%)	Index
Jan.	Peak-work	23.34	16.67	6.32	2.17	106.1
Feb.	Peak-work	23.42	16.58	6.52	2.23	106.5
Mar.	Peak-work	24.91	18.31	6.30	2.07	113.3
Apr.	Peak-work	25.49	18.71	6.42	2.34	115.9
May	Middle	23.11	17.00	5.79	1.79	105.1
June	Middle	20.94	14.44	6.15	2.26	95.2
July	Low-work	18.97	12.50	6.15	2.22	86.2
Aug.	Low-work	18.89	12.37	6.26	1.53	85.9
Sept.	Low-work	20.03	13.44	6.31	1.62	91.1
Oct.	Low-work	20.88	13.61	6.96	2.74	94.9
Nov.	Middle	21.44	14.69	6.40	2.13	97.5
Dec.	Middle	22.50	15.57	6.58	1.76	102.3
Low-work months	July, Aug., Sept., Oct.					
Peak-work months	Jan., Feb., Mar., Apr.					
Peak/low total-days ratio	1.23					

Household-month labor module, 1998–2014. Days are divided by covered household members. Total equals farm, own-business, paid, exchange, and unpaid days. The four lowest months are classified as low-work months and the four highest as peak-work months. This table can be read as a descriptive measure of seasonal slack.

Table B.5: Nominal wage changes: how often the wage is literally unchanged

Measured at	Exactly unchanged	By chance	Excess over chance	Share down	Share up	Changes observed
Village going rate	0.444	0.167	0.277	0.276	0.280	3107
Household	0.273	0.024	0.249	0.355	0.372	55441
Individual worker	0.317	0.022	0.295	0.331	0.353	83360

Month-to-month changes in the nominal daily wage for casual work, 1998–2014, built from job-level records covering all twelve months. The wage for each unit-month is the modal rate, because a mean of continuous earnings-per-day ratios is almost never exactly repeated and so cannot reveal an unchanged wage at all. The chance column pairs each unit with a different unit in the same calendar month, which measures how much exact agreement the coarse wage grid produces on its own; 53.7 percent of recorded wages are a multiple of ten baht. The excess over chance is large at every level of aggregation and largest for individual workers, so the finding is not an artifact of averaging within villages. The last two columns are decisive for interpretation: wages fall about as often as they rise, so this is nominal inertia rather than downward rigidity in the sense of Kaur (2019), which requires the downward share to lie well below the upward one.

C Partner Characteristics

This appendix provides the detail behind Table 2. It separates household participation from pair characteristics and reports how labor relationships overlap with credit, gift, and kinship ties.

Table C.1: Household participation and characteristics of labor relationships

	No sampled labor contract	Exchange only	Hiring only	Both contracts
Panel A. Household contract use in 1998–99				
Households	123	81	141	341
Landholding household (%)	82.9	95.1	92.2	97.4
Baseline landholdings (hectares)	4.2	5.4	3.8	3.5
Own-farm person-days/month	10.2	16.4	12.1	26.6
Outside paid person-days/month	16.5	16.3	13.8	12.4
Wealth index (standardized)	−0.24	−0.35	0.30	−0.24
Any borrowing or lending tie (%)	15.4	27.2	31.2	46.9
Panel B. Pair profiles using 1998–99 labor histories				
	Untied in 1998–99	Exchange-only in 1998–99	Hire-only in 1998–99	Dual in 1998–99
Baseline exchange transactions	0.0	5.1	0.0	5.0
Baseline paid transactions	0.0	0.0	2.6	4.2
Kin share	0.02	0.19	0.04	0.18
Non-labor tie share	0.02	0.16	0.08	0.21
Mean own-farm days/month	18.0	29.1	24.5	36.3
Mean absolute land gap (rai)	17.6	17.6	12.5	9.7
Share both landed	0.86	0.97	0.94	0.98

Notes: Household categories use sampled labor transactions during 1998–99, as receiver or provider. A household using both contracts participates in exchange with any sampled partner and hiring with any sampled partner; these need not be the same partner. Exchange-only and hiring-only households use only the respective contract, and remaining households have no sampled labor transaction in this window. These categories differ from the fixed pair-level baseline categories in the main regressions, which use raw waves 1–24. Household characteristics are averages over 1998–99. The wealth index is standardized to mean zero; negative values indicate wealth below the sample mean. Panel B describes directed within-village pairs in the historical descriptive transaction extract, classified by their undirected 1998–99 labor history. Transactions are summed over 1998–99, including zeros, and each profile uses the observations available for its baseline characteristics. These profiles are not restricted to the interview-aligned estimation sample in Table 1, which classifies relationships over raw waves 1–24 and requires observed residence, a receiver recall interval of 15–60 days, and price information dated before that interval. The different window and eligibility rules prevent direct comparisons of the pair samples. Non-labor ties are gift, asset, or borrowing relationships.

Table C.2: Contractual transformation and overlap with other village networks

Panel A. Contractual forms observed after 1999			
Baseline relationship	Pairs	Later hired (%)	Later exchange (%)
Baseline-dual	244	73.8	83.2
Baseline-exchange-only	640	41.1	70.3
Baseline-hire-only	723	45.9	32.8
Panel B. Share of baseline labor pairs with another tie			
Baseline labor relationship	Gift	Borrow/lend	Kin
Baseline-exchange-only	8.5%	4.9%	18.7%
Baseline-dual	5.0%	11.7%	18.4%
Baseline-hire-only	1.8%	3.8%	4.2%
Any labor relationship	5.0%	5.5%	12.2%
Panel C. Same household, different partners			
		Share of partners	Households
Labor partners also used for gifts		18.0%	191
Labor partners also used for borrowing/lending		12.6%	226
Labor partners who are kin		20.8%	405

Notes: Panels A–B use the historical descriptive transaction archive of undirected within-village household pairs classified by recorded labor transactions in 1998–99. Panel A follows these relationships from 2000 onward; its 244 dual, 640 exchange-only, and 723 hire-only pairs are descriptive archive counts. Table 1 instead reports directed pairs in the interview-aligned estimation sample, with relationship types defined over raw waves 1–24 (September 1998–August 2000) and pair-wave eligibility determined by residence, a receiver recall interval of 15–60 days, and pre-recall price availability. The different baseline window, sample restrictions, and unit explain why the counts here cannot be obtained simply by halving those in Table 1. "Later hired" and "later exchange" overlap and need not sum to 100 percent; these unconditional transition rates do not identify what induces contract changes. Panel C uses sampled household networks in the 1998–99 descriptive extract, with the observations available for each network. Gift, borrowing, and kin relationships frequently coexist at the household level, but households generally use different partners across networks.

Table C.3: Baseline labor-exchange partners: similarity and gaps (1998–99)

Statistic	Value
Undirected baseline exchange pairs	866
Correlation, own-farm days	0.33
Correlation, household size	-0.02
Share same land status	96.9%
Mean absolute own-farm workload gap	22.2 days/month
Share with own-farm workload gap above 10 days	64.4%
Share with household-size gap of at least two	39.2%

Notes: The sample consists of undirected within-village pairs with at least one recorded exchange transaction in 1998–99 and endpoints matched to the household roster in the descriptive transaction extract. It includes both exchange-only and dual-contract relationships in that window. Own-farm days and household size are household monthly means over 1998–99; each statistic uses the available non-missing endpoint characteristics. The 866 pairs are a descriptive exchange-pair count, not a count from the interview-aligned estimation risk set in Table 1. The main sample uses relationship histories from raw waves 1–24 and residence, recall-interval, and pre-recall price restrictions, so its directed-pair counts are not directly comparable.

D Shifter Validity and Estimate Robustness

This appendix reports the continuous relationship-intensity specification, the household outcomes used to assess the aggregate labor-demand channel, and the binned harvest-month check. It then reports the three extensions discussed in Sections 4.3 and 5.5: alternative production-calendar timing proxies, receiver-wave totals without receiver-by-wave fixed effects, and the contract mix among pair-waves that transact. These exercises use the interview-aligned sample and the same price innovations as the main estimates. Figure 2 and Section 4.2 describe the support and crop composition of the shock.

Table D.1: Continuous baseline relationship intensity

Outcome	Baseline measure	Estimate	SE	p
Any exchange	Baseline labor intensity	0.156	(0.109)	0.171
Any exchange	Dual-contract intensity	0.345	(0.512)	0.510
Any hired labor	Baseline labor intensity	0.121	(0.068)	0.094
Any hired labor	Dual-contract intensity	0.124	(0.205)	0.553
Exchange minus hired labor	Baseline labor intensity	0.035	(0.121)	0.778
Exchange minus hired labor	Dual-contract intensity	0.221	(0.438)	0.621

Notes: Coefficients are harvest-timing increments in percentage points for a ten-percentage-point increase in the indicated baseline intensity measure, evaluated for a one-standard-deviation movement in S ($\mathrm{SD}(S) = 0.581$) and a one-unit increase in the standardized harvest calendar. Let q^{Ex} and q^{Hire} be the shares of raw baseline waves 1–24 in which an undirected pair used exchange or hiring, respectively. Baseline labor intensity is $\max(q^{\mathrm{Ex}}, q^{\mathrm{Hire}})$; dual-contract intensity is $\min(q^{\mathrm{Ex}}, q^{\mathrm{Hire}})$. These measures record activity under the two contracts and do not measure their joint use within the same wave. The specification retains the fixed effects in Equation (1). Standard errors are clustered by village, and reported p-values use the $t(15)$ reference distribution.

Table D.2: Direct outcomes and the limits of a market-tightness interpretation

Outcome	Estimate	Village SE	Village p	Two-way p
Own-farm days per 30 days	0.405	0.171	0.032	0.251
Paid-work days per 30 days	0.114	0.491	0.820	0.831
Any recorded harvest	-0.750	0.303	0.025	0.115
Log harvested quantity	-0.060	0.023	0.018	0.090
Receives sampled exchange	0.146	0.267	0.591	0.689
Receives sampled hired labor	0.661	0.102	< 0.001	< 0.001

Notes: The table reports responses to an increase of 0.581 units in harvest-timed exposure, using the same main-sample scaling as the dyadic tables. Household-wave regressions include receiver $\times$ month-of-year and survey-wave fixed effects. Binary outcomes are in percentage points; day and log-quantity outcomes remain in their original outcome units. Village p-values use the $t(15)$ reference distribution. The estimates do not form a coherent physical labor-demand first stage.

Table D.3: Binned harvest-month check

	(1) Any exchange	(2) Any hired
A. Response in non-harvest months		
$S\times$ baseline-exchange-only	-0.063	-0.006
	(0.063)	(0.019)
$S\times$ baseline-dual	-0.243**	0.036
	(0.084)	(0.042)
$S\times$ baseline-hire-only	0.013	-0.006
	(0.025)	(0.030)
B. Additional response in above-average harvest months		
$S\times$ harvest $\times$ baseline-exchange-only	0.746*	0.129*
	(0.425)	(0.071)
$S\times$ harvest $\times$ baseline-dual	1.107*	0.443
	(0.604)	(0.285)
$S\times$ harvest $\times$ baseline-hire-only	0.131	0.182**
	(0.093)	(0.062)
Pair-wave observations	2,629,888	2,629,888
Clusters (villages)	16	16

Notes: Coefficients and standard errors are scaled to a one-standard-deviation movement in S using $\mathrm{SD}(S) = 0.581$. The specification replaces continuous harvest-timed exposure with the interaction of S_{hvw} and an indicator for above-average village harvest intensity ($c_{vm} > 0$). Panel B compares above-average harvest months with other months; it is not a one-unit increment in the continuous calendar. Fixed effects and village clustering match Equation (1). This appendix table retains village-clustered $t(15)$ stars, rather than the restricted wild-cluster bootstrap used in the main results tables.

Table D.4: Alternative production-calendar timing

	Land-preparation proxy		Four-month pre-harvest proxy	
	Any exchange	Any hired	Any exchange	Any hired
A. Alternative-calendar increment				
$P\times$ baseline-exchange-only	-0.114*	-0.006	-0.248***	-0.016
	(0.060)	(0.007)	(0.075)	(0.013)
	[0.071]	[0.335]	[< 0.001]	[0.336]
$P\times$ baseline-dual	-0.113	-0.083**	-0.107	0.010
	(0.078)	(0.035)	(0.077)	(0.056)
	[0.197]	[0.034]	[0.173]	[0.870]
$P\times$ baseline-hire-only	-0.031*	-0.018	-0.021	-0.038
	(0.017)	(0.012)	(0.019)	(0.025)
	[0.089]	[0.141]	[0.310]	[0.182]
B. Joint timing specification: harvest H and land-preparation proxy P				
$H\times$ baseline-exchange-only	0.431*	0.110**	—	—
	(0.192)	(0.041)	—	—
	[0.069]	[0.026]	—	—
$P\times$ baseline-exchange-only	-0.015	0.020	—	—
	(0.066)	(0.010)	—	—
	[0.799]	[0.137]	—	—
$H\times$ baseline-dual	0.782***	0.287	—	—
	(0.302)	(0.174)	—	—
	[0.006]	[0.106]	—	—
$P\times$ baseline-dual	0.042	-0.026	—	—
	(0.052)	(0.027)	—	—
	[0.431]	[0.413]	—	—
$H\times$ baseline-hire-only	0.096	0.134**	—	—
	(0.060)	(0.044)	—	—
	[0.136]	[0.011]	—	—
$P\times$ baseline-hire-only	-0.007	0.013	—	—
	(0.012)	(0.013)	—	—
	[0.619]	[0.323]	—	—
Pair-wave observations				2,629,888
Clusters (villages)				16

Notes: Coefficients and standard errors use the main estimation-sample scaling $\mathrm{SD}(S) = 0.581$ (0.5806547 before rounding). Average-month effects refer to a one-standard-deviation movement in S; continuous harvest-timing increments refer to the same movement for a one-unit increase in the standardized harvest calendar. P_{hvw} rebuilds H_{hvw} after rotating the leave-one-village-out harvest profile backward; it does not use observed planting dates. The eight-month land-preparation proxy maps a December rice harvest to April. The four-month proxy maps a December rice harvest to August and is a sensitivity check for an alternative pre-harvest window. Proxy increments refer to a one-unit increase in the corresponding standardized proxy calendar. Panel B includes harvest and land-preparation interactions jointly. $\mathrm{Corr}(H, P_{\mathrm{prep}}) = -0.32$. Fixed effects match Equation (1). Village-clustered standard errors are in parentheses and restricted wild-cluster bootstrap p-values are in brackets (999 draws). Stars use the bootstrap p-values.

Table D.5: Receiver-wave labor use and crop-return innovations

	(1) Exchange partners	(2) Hired partners	(3) Any labor partners	(4) Receives exchange (pp)	(5) Receives hired (pp)
A. Receiver × month-of-year and survey-wave FE					
S_{hvw}	-0.004	0.003	-0.001	-0.166	0.188
	(0.005)	(0.003)	(0.007)	(0.185)	(0.198)
	[0.436]	[0.421]	[0.881]	[0.370]	[0.392]
H_{hvw}	0.019*	0.011***	0.027**	0.146	0.661***
	(0.011)	(0.002)	(0.012)	(0.267)	(0.102)
	[0.096]	[0.002]	[0.025]	[0.602]	[< 0.001]
Peak-harvest total per SD of S	0.036	0.026***	0.056*	0.143	1.584***
	(0.025)	(0.007)	(0.028)	(0.606)	(0.335)
	[0.156]	[0.004]	[0.060]	[0.819]	[0.002]
B. Village, month-of-year, and survey-wave FE					
S_{hvw}	-0.008	0.001	-0.006	-0.340	0.017
	(0.005)	(0.002)	(0.006)	(0.207)	(0.170)
	[0.115]	[0.722]	[0.300]	[0.140]	[0.922]
H_{hvw}	0.023**	0.013***	0.032**	0.512*	0.741***
	(0.010)	(0.002)	(0.012)	(0.252)	(0.113)
	[0.025]	[< 0.001]	[0.010]	[0.050]	[0.001]
Peak-harvest total per SD of S	0.040*	0.028***	0.061**	0.743	1.584***
	(0.023)	(0.006)	(0.026)	(0.516)	(0.286)
	[0.069]	[0.001]	[0.019]	[0.143]	[< 0.001]
Receiver-waves					69,406
Clusters (villages)					16

Notes: Coefficients and standard errors use the main estimation-sample scaling SD(S) = 0.581 (0.5806547 before rounding). Avera refer to a one-standard-deviation movement in S; continuous harvest-timing increments refer to the same movement for a one-uni standardized harvest calendar. The unit is a receiver-wave. Columns (1)–(3) count sampled providers; columns (4)–(5) are indic by 100. There are no receiver-by-wave fixed effects. Panel A includes receiver × month-of-year and survey-wave effects, matc Panel B includes village, month-of-year, and survey-wave effects. Peak totals are SD(S)($\beta + \theta c^{\text{peak}}$), with $c^{\text{peak}} = 2.11$. Column (descriptive: the dyadic any-labor outcome fails the future-innovation test ($p = 0.002$; Table 5), so the pooled partner count is ex causal interpretation. Mean partners used: exchange 0.146, hired 0.068, any labor 0.207. Village-clustered standard errors are in restricted wild-cluster bootstrap p-values are in brackets (999 draws). Stars use bootstrap p-values.

Table D.6: Contract mix among pair-waves with a labor transaction

	Share any exchange		Share exchange only	
	Main FE	Coarser FE	Main FE	Coarser FE
A. Harvest-timing increment ($0.581\theta_r$)				
$H\times$ baseline-exchange-only	0.339	0.531	2.517	0.090
	(2.936)	(0.340)	(1.385)	(0.470)
	[0.944]	[0.208]	[0.259]	[0.841]
$H\times$ baseline-dual	-1.881	-0.546	-0.523	-1.129
	(1.086)	(0.525)	(2.004)	(0.809)
	[0.260]	[0.324]	[0.763]	[0.183]
$H\times$ baseline-hire-only	—	-2.036*	—	-2.062**
	—	(0.897)	—	(0.633)
	—	[0.059]	—	[0.026]
B. Mean among transacting pair-waves (%)				
Baseline-exchange-only	89.0		86.1	
Baseline-dual	68.4		63.0	
Baseline-hire-only	44.4		40.0	
Baseline-untied	61.8		59.0	
Transacting pair-waves				14,376
Clusters (villages)				16

Notes: Coefficients and standard errors use the main estimation-sample scaling $\mathrm{SD}(S) = 0.581$ (0.5806547 before rounding). Average-month effects refer to a one-standard-deviation movement in S; continuous harvest-timing increments refer to the same movement for a one-unit increase in the standardized harvest calendar. The sample keeps pair-waves with any labor transaction. Outcomes are exchange indicators multiplied by 100. Because the shock can change which pairs transact, estimates condition on the outcome and are descriptive. This selected-sample exercise is not used to assess a liquidity mechanism. Main FE match Equation (1); coarser FE use pair, village-by-wave, and month-of-year effects. The two baseline-hire-only cells under main FE, their standard errors, and their p-values are suppressed because these estimates are uninformative in the selected sample; — denotes suppression, not zero. Village-clustered standard errors are in parentheses and restricted wild-cluster bootstrap p-values are in brackets (999 draws). Stars use bootstrap p-values.