Urban Insurance & Informal Networks

Aruj Shukla [*]

August 25, 2025

Abstract

Urban labor markets in developing countries are marked by instability, weak social protection, and widespread informality, making risk-sharing a central concern for workers and policymakers. This paper examines the nature of insurance and informal networks in urban settings using longitudinal data from Ghana and Uganda, supplemented by original survey evidence from manufacturing workers in Kampala, Uganda. Tests for full insurance reveal that while urban households are not fully insured, their consumption is only weakly sensitive to idiosyncratic income shocks, with elasticities between 0.01 and 0.05, indicating effective partial insurance. Assessing the relevance of risk-pooling networks, I compare geographically proximate networks to those consisting of communities in places of origin. Evidence suggests that urban co-location provides stronger pooling than rural-origin networks, though data granularity limits definitive conclusions. To supplement these findings, I conduct a survey of workers capturing detailed information on shocks, insurance mechanisms, and coping strategies. Results highlight the predominance of self-insurance via precautionary savings, with over 60% of workers reporting a shock in the past month and 90% relying on savings to cope. This strategy, however, delays transitions into self-employment, potentially contributing to aggregate misallocation. Finally, I document polarized demand for unemployment insurance, shaped heavily by workers' beliefs about job search prospects. The findings underscore the importance of understanding informal arrangements in shaping worker welfare, labor market trajectories, and the design of effective safety nets in rapidly urbanizing economies.

[*]University of Southern California. Email: arujshuk@usc.edu. I thank Vittorio Bassi and Anna Vitali for their guidance and productive discussions throughout the project. I am grateful to the team at International Growth Research and Evaluation Center (IGREC), Uganda—especially Oscar Kidega and Assy Nakyajja—for their help with data collection.

1 Introduction

Urban labor markets in developing countries are marked by volatility and instability (Breza and Kaur, 2025). Missing markets (Udry, 1995), weak social safety nets (Imbert and Papp, 2020; Miner, 2025), the prevalence of firms shifting demand risk onto workers through payment contracts—about 70% of informal workers are paid piece rates—and frequent wage delays (Alfonsi et al., 2020; McKenzie, 2024) together characterize the typical experience of urban wage employment. These frictions contribute to the prominence of informal networks as an avenue for risk sharing and insurance against economic shocks. Yet an overwhelming majority of theoretical and empirical work on risk-sharing focuses on rural contexts (Munshi and Rosenzweig, 2016; Deaton, 1990; Townsend, 1994), offering little guidance on i) the nature of risks and ii) the organization of insurance networks to cope with such risks in urban settings.

Addressing this lacuna in development economics is of paramount importance. Over the next 25 years, African cities alone are on course to witness 900 million new urban residents (Lockhart, 2025). As the engine of private sector wage employment, urbanization is an inevitability in countries rapidly undergoing structural transformation. [1] As policymakers grapple with the provision of good jobs and enhancing the coverage of social safety nets, understanding the existing informal arrangements of risk pooling is critical to designing effective safety nets. Beyond the policy ramifications, the presence of informal insurance arrangements shape individuals' socio-economic lives in profound ways. Inability to access more profitable locations within a city due to a lack of formal insurance can reduce welfare (Pelnik, 2025). Time required to maintain informal ties through social obligations can lower productivity and restrict reliable labor supply. Learning about the experiences of the urban migrants is central to our understanding of the frictions that impede movement of workers to higher productivity areas—a key driver and implication of structural transformation (Gollin et al., 2021).

Studying the nature of urban insurance is more than transposing our understanding of insurance in village economies to urban settings. Insurance networks in urban environments raise unique challenges: do traditional, ethnic networks continue to persist in the city, or are they replaced by networks of co-workers and employers through market interactions? Moreover, are urban migrants still insured through networks in the rural areas? On one hand, urban settings have larger and varied networks that are less prone to aggregate shocks, allowing for mutual insurance to flourish. However, the sprawl and transience of urban centers may create disconnect and anonymity or even strain monitoring and enforcement of informal networks, making them harder to sustain.

In this paper I try to resolve these issues and start building the micro foundations of the risks urban workers face. Using large-scale panel survey data from Ghana and Uganda, I test for full insurance among urban households to gather a positive understanding risk pooling (Townsend, 1994; Li and Ligon, 2020; Mace, 1991). Although, I am able to reject the null of full insurance, the elasticities of consumption with respect to income are substantially low (ranging from 0.01 to 0.05), indicating effective, partial insurance. To understand the relevant network in urban settings, I proceed to estimate and compare coefficients under two network definitions: 1) the set of co-located households (analogous to the village definition in the literature), and 2) the set of households in the place of origin of the migrant household. Even though there is suggestive evidence of stronger risk-pooling among co-located households (compared to the community of households in the place where they migrate from), limitations with data granularity

[1] Historically, urbanization and economic growth or income gains have gone hand-in-hand. See Vollrath (2024) for a succinct explanation of why this relationship might be weakening for rapidly urbanizing countries.

of place of origin renders this comparison biased.

To overcome these data constraints and learn about the insurance mechanisms, I collected detailed survey data from 150 workers in Kampala, Uganda. With detailed information on the occurrence and types of shocks, their implications, and adaptation strategies, I find strong evidence consistent with self-insurance to manage income shocks. Given that 60% of the sample of workers report having faced a shock in the past month, self-insurance as a coping strategy has deeper implications: it delays and denies the aspirational move to entrepreneurship. With weak access to financial and credit markets, poorer workers cannot raise enough capital to be self-employed. Regardless of productivity, richer households select into entrepreneurship with implications on aggregate productivity and misallocation. Finally, I bring to light an under explored barrier to the adoption of an unemployment insurance program: workers' beliefs about job finding and job separations. Those who are confident of finding a job are very unlikely to favor such a safety net. As policymakers deliberate effective design of safety nets in contexts full of informality, this margin is important to consider.

The rest of the paper proceeds as follows. Section 2 develops a simple model of full insurance and risk pooling. Section 3 describes the data from Ghana, Uganda, and the Kampala pilot survey. Section 4 outlines the empirical strategy. Section 5 presents the main results on risk pooling, transfers, and network definitions. Section 6 provides evidence from the worker survey on shocks, coping strategies, and demand for unemployment insurance. Section 7 concludes with implications and directions for future work.

2 Theory: risk pooling and full insurance

The theoretical underpinnings of full insurance have been well developed. In this section I outline a simple, workhorse model of full insurance, drawing heavily on the fantastic exposition in Suri (2005). The goal of the model is straightforward: derivation of theoretical predictions that can be then taken to the data to test for optimal risk pooling. As I'll show below, the key prediction from full insurance implies that household consumption should only co move with aggregate shocks. A corollary of this proposition states that idiosyncratic shocks to the household should leave the consumption of the household unaffected.

I assume throughout that the village is the relevant risk sharing group, i.e. the appropriate grouping amongst which households pool risk. This is a common assumption in the literature.

Consider N households in the village, each indexed by i. There exist S states of nature, indexed by s, each occurring with a probability π_s. Household income for each state is exogenously given and denoted by y_{is}. Let the utility for household i be represented as:

$$U_i = \sum_s \pi_s u_i(C_{is}) \tag{1}$$

where C_{is} denotes the household i consumption in state s. Equation 1 is simply the sum of utilities across states, weighted by the probability of the state, π_s. $u_i(\cdot)$ is household i's utility function for each state of nature and depends exclusively on the consumption in that state.[2] The Pareto efficient risk allocation arrangement in this economy can be attained by solving for the following constrained optimization

[2] $u(\cdot)$ is assumed to be twice differentiable with $u'(\cdot) > 0$ and $u''(\cdot) < 0$.

problem:

$$\max \sum_i \omega_i U_i \tag{2}$$

subject to

$$\sum_i C_{is} = \sum_i y_{is} \quad \forall s \tag{3}$$

and

$$C_{is} > 0 \quad \forall s \tag{4}$$

where ω_i are the household specific weights on their utility functions. The first order conditions for the two household case, i and j, are

$$\frac{u_i'(C_{is})}{u_j'(C_{js})} = \frac{\omega_j}{\omega_i} \quad \forall i, j, s \tag{5}$$

Imposing a constant absolute risk aversion (CARA) utility function[3] in Equation 5 yields

$$\frac{e^{-\rho C_{is}}}{e^{-\rho C_{js}}} = \frac{\omega_j}{\omega_i} \tag{6}$$

Taking logs after gives us this expression for household consumption for each household:

$$C_{is} = C_{js} + \frac{1}{\rho}\left[\ln(\omega_i) - \ln(\omega_j)\right] \tag{7}$$

Aggregating over N such equations allows for an expression of household consumption as a function of average village consumption:

$$C_{is} = \bar{C}_s + \frac{1}{\rho}\left[\ln(\omega_i) - \frac{1}{N}\sum_j \ln(\omega_j)\right] \tag{8}$$

where $\bar{C}_s = \frac{1}{N}\sum_i C_{is}$ is the average consumption in the village.

Note that $\ln(\omega_i)$ in equation 8 can be interpreted as a household fixed effect (α_i) where as all the other terms can be collected as they vary at the village level. Simplifying the expression further, we have:

$$C_{is} = \bar{C}_s + \alpha_i \tag{9}$$

Equation 9 captures the heart of the risk pooling intuition. To reiterate, the implication of full insurance is that

1. Household consumption co-moves, one for one, with average village consumption.
2. Controlling for aggregate consumption, household consumption should be completely insensitive to idiosyncratic household income shocks.

The test for complete risk sharing is therefore a joint test for these propositions, which can be achieved by a variety of specifications. Before outlining these, I turn to a description of my data and the construction of outcomes to test for risk pooling.

[3] I use the standard CARA function: $u_i(x) = -\frac{1}{\rho}e^{-\rho x} \; \forall i$

3 Data & Outcomes

3.1 Sources

Testing the sensitivity of household consumption to household income is at the heart of tests developed to study the extent of risk pooling behavior. While cross section data can suffice, canonical tests in the literature rely on panel datasets. I draw on two panel datasets, which I describe in detail below, to characterize the extent of risk pooling in urban settings.

Ghana Socioeconomic Panel Survey (GSPS) The GSPS dataset contains detailed information on households such as assets, consumption, health, income, education, and migration. Spanning three waves (2009/2010, 2013/2014, 2017/2018), this nationally representative panel consists of 5009 households—40% of which are in urban locations—across 334 enumeration areas. Following a two-step stratified (based on regions in Ghana) sampling design, the 334 census enumeration areas were randomly selected first. In the second stage, 15 households from the listing of all households in the enumeration area were randomly selected to yield the final sample.

Uganda National Panel Survey (UNPS) The UNPS is also a nationally representative panel survey conducted on an annual basis. I use four waves of the UNPS spanning 2013/14 through 2019/20. The panel consists of 3119 household, out of which 26% (819 households) are in urban locations. Spread over 317 enumeration areas (that were randomly selected from a frame of 780 enumeration areas), each consisting between 5-10 sampled households. The survey contains rich data on household details such as assets, education, consumption, income, and a detailed module on agriculture.

3.2 Constructing the HH $\times$ wave panels

I stitch together waves of data for each of the national surveys in Ghana and Uganda to arrive at a household level panel of incomes and consumption expenditures. 2,444 households (78.3% of of 3,119 households) in the Ugandan panel and 3,974 households (79.3% of 5,009 households) in the Ghanian panel have data across all survey waves under consideration. As the survey waves span decades, it is not uncommon to have households exiting, entering the sample or splitting off from the original household in the sample. I restrict the panels to households that have at least two waves of data.

3.3 Key variables

The focal modules of interest for tests of risk pooling are that of incomes and consumption expenditures. Typically, measures of incomes and consumption tend to be noisy and have recall frames varying from a week to a year depending on the type of good that was consumed. In this subsection I detail the constituent elements of my main outcomes, the recall period, and any further processing to homogenize outcomes across waves of data collection.

Incomes Typically, data on incomes have a recall period of the past 12 months. For the Ugandan panel, I aggregate incomes from non-agricultural household enterprises, crop sales, wage employment [4], income

[4] Income from wage employment has a weekly or monthly recall period. I scale up monthly earnings by the number of months worked in the past year to arrive at an equivalent income for the past 12 months. Both main and secondary source of

from land and property, and other sources such as interest payments and dividends. In the Ghanian panel, even though incomes are captured with better granularity [5] across sources, the same categories mentioned above are aggregated over.

Two point of note are in order. First, all enterprise incomes in the main analysis are *gross incomes*, not net incomes. Costs of production are not uniformly recorded across categories and the employment of household labor further complicates the precision with which measures of net income can be constructed. Secondly, remittances and in-transfers are not included to avoid double counting and contaminating the test for risk pooling. I treat these measures as a separate outcome in further analyses.

Consumption Household consumption expenditures include a range of durable, semi-, and non-durable goods. The durability of items lends to a different recall period for each category. To align with income reference periods, I scale consumption expenditures—primarily food and non-durables—on a 12 month recall basis. Implicitly this scaling assumes away any seasonal variation in expenditures on such items but allows for standardization in the construction of the household aggregate consumption measure. Again, as with incomes, any transfers and remittances made made or received by the household are omitted from this measure. Moreover, all consumption expenditures are tallied over goods not produced at home. Sparse information on market prices motivates this decision and precision, beyond a reasonable and transparent threshold, in measuring consumption is of limited importance to the essence of the empirical exercises I conduct.

Log transformations Income data is characterized by long tails and zeroes. Accounting for these features, both incomes and consumption expenditures are $\log + c$ transformed where $c = 1$ for the analyses[6] that follow. Despite Chen and Roth (2024) convincing diagnoses of the issues with using log transformations of weakly positive variables, I proceed with this transformation for benchmarking purposes. To the best of my knowledge, all specifications in the literature testing for full insurance measure incomes and consumption expenditures in either levels or logs. Adhering to the same construction—albeit misinterpreted as percentage effects—contextualizes the estimates I gather from urban settings, in the canon of studies prior. Future work re-examining tests of risk pooling following the guidance from Chen and Roth (2024) remains to be a fruitful endeavor.

Network definition Investigating whether there is risk pooling with communities in the place of origin of the urban household requires careful construction of links to such places. Because I observe district of birth of every household member and how long since they moved to the current place of residence, every household is assigned a *modal* place of origin based on that information. In case of ties, the place of origin for the person who most recently moved to the current residence is used. To illustrate this, consider the simple case of two household members who were both born in District A. This household's place of origin is linked to District A. Now, consider another household where both members were born in a different district but household member #2 moved to the current place of residence 5 years ago (compared to 10 years ago for the other member). As a result, the district of birth of member #2 is assigned as the

incomes for all household members is included.

[5] For instance, incomes from different types of non agricultural enterprises are recorded separately.

[6] In results not shared in the paper, I find that setting $c = 100$ yields results that are identical in spirit to the ones I report in Section 5.

household's place of origin. The underlying logic preserves the links to the community the household was either most rooted in or has most recent contact with.

A significant limitation of this data is the lack of granularity for the place of origin—the district. Districts in such settings are large administrative boundaries that can contain thousands of households[7]. In principal, I don't expect effective risk pooling to happen across such large geographical boundaries. I address this issue in some detail in Section 5.3.

4 Empirical Strategy

Since Deaton (1990), variants of the core empirical framework to test for full insurance have been implemented in the literature. The underlying theoretical model discussed in Section 2 and its implication remains the same: household incomes should play no role in determining household consumption after controlling for aggregate consumption. In this paper, I deploy two distinct yet equivalent empirical specifications from the literature to test for risk pooling.

4.1 Two-way fixed effects (TWFE) specification

The modern workhorse specification for panel data in risk pooling tests relies on a set of time and household level fixed effects. As in Li and Ligon (2020), I estimate the following specification:

$$\log c_{it} = \alpha \log y_{it} + \delta_i + \phi_{vt} + \varepsilon_{it} \tag{10}$$

where c_{it} denotes household consumption and y_{it} denotes household income. δ_i is the set of household fixed effects (FEs), while ϕ_{vt} represents community $\times$ time effects. ε_{it} is a disturbance term e.g., due to measurement error or unobserved preference heterogeneity) and standard errors are clustered at the community level.

Full insurance implies $\alpha = 0$ and is the parameter of interest. The essential logic of this specification is straightforward: accounting for household FEs and community-level resources (captured by ϕ_{vt}), household consumption should be completely insensitive to deviations in household income under full risk diversification.

4.2 First differences specification

While the TWFE specification is sufficient to test for risk pooling in the standard geographical definition of community, I adopt a more flexible specification à la Mace (1991) to estimate tests for full insurance on alternate network definitions. The first differences approach is equivalent to the TWFE approach and also allows for flexible construction of aggregate consumption measures based on different network definitions, especially for urban migrants based on their places of origin. Additionally, testing for full insurance using an alternate specification in the literature lends an extra layer of validity to the estimates.

First, I estimate the following first difference specification in levels :

$$\Delta C_t^j = \beta_1 \Delta C_t^a + \beta_2 \Delta y_t^j + u_t^j \tag{11}$$

[7] On average, there are 70 sampled households within a district in the GSPS

where ΔC_t^j reflects changes in household consumption and Δy_t^j changes in household income. ΔC_t^a denotes changes in aggregate community consumption. As before, standard errors are clustered at the community level. Here, full insurance or complete risk pooling is a joint test for 1) the co-movement of household consumption with average community consumption and 2) household incomes not affecting household consumption. From the specification, this translated to $\beta_1 = 1$ & $\beta_2 = 0$.

And also in logs, akin to a growth rate specification, to understand risk sharing implications for preferences taking the power utility functional form:

$$\Delta \log C_t^j = \beta_1 \Delta \log C_t^a + \beta_2 \Delta \log y_t^j + v_t^j \tag{12}$$

5 Results

5.1 Risk pooling within local communities

The primary goal of this analysis is to study the extent of risk pooling among urban residents in low income countries. With a focus on Sub-Saharan Africa (SSA), I use panel datasets from Ghana and Uganda to test for full insurance. Despite some variation in the estimates, in line with the extensive literature studying this in rural settings, I am able to reject the null of full insurance in urban settings.

Table 1: Ghana Panel data: Testing for full insurance

	HH consumption			Consumption per capita		
	Full Sample (1)	Urban (2)	Rural (3)	Full Sample (4)	Urban (5)	Rural (6)
α: Log HH Income	0.040***	0.051***	0.034***	0.039***	0.050***	0.033***
	(0.004)	(0.007)	(0.005)	(0.004)	(0.007)	(0.005)
Dep. Var Mean	6.83	7.07	6.69	5.90	6.27	5.69
Observations	11776	4257	7519	11776	4259	7517
No. of HH	4362	1632	2730	4363	1633	2730
R^2	0.86	0.87	0.85	0.88	0.89	0.87

NOTES—This table presents results from Ghana for a test of full insurance among geographic community members. Two-way FE specification as in equation 1 is estimated where for columns (1)-(3) the outcome of interest is the log of household consumption (in Cedis) and for columns (4)-(6) the outcome of interest is log of consumption per adult equivalent (in Cedis) over the past 12 months. 1 USD = 1.3 Cedi (2009). Standard errors clustered at the enumeration area level (15 HHs). *p<0.1, **p<0.05, ***p<0.01.

Table 1 encapsulates the results of this exercise for Ghana. The first three columns use total household consumption as the outcome variable, while the last three columns rely on adult-equivalent scale (AES) adjusted measures of household consumption per capita[8]. Community is defined at the enumeration area level that consists of 15 sampled households. This is the most granular geographic variable available in the panel and can be interpreted as an analog to the "village"—the most natural and commonly used division for risk pooling networks in the literature. Addition of community $\times$ time fixed effects imply intra-community risk pooling but not across.

[8] Li and Ligon (2020) use the per capita measure of household consumption and I show both results for transparency and completeness

The estimand of interest is α, the elasticity of household consumption to household income, which should be equal to zero under full insurance. Across sub samples, I am statistically powered to reject the null of full insurance. Measuring income on a per capita basis does not affect the elasticity estimates. Beyond a sound rejection of the null hypotheses, two notable patterns emerge. First, there appears to be good risk pooling among geographically proximate communities across Ghana. The coefficient in column (1) implies a 10% increase in household incomes is associated with a mere 0.4% increase in household consumption. Households' consumption expenditures barely fluctuate in response to idiosyncratic shocks to household incomes. Contrasted with elasticity estimates of 0.40-0.42 reported by Li and Ligon (2020) from Tanzania, risk pooling seems to be very effective in smoothing shocks, albeit not perfectly. I return to the validity of these estimates in the next section. Second, urban households bear more risk than their rural counterparts. With an elasticity of 0.05, households are still effectively guarding against risks but relatively less than that in rural areas.

Table 2: Uganda Panel data: Testing for full insurance

	HH consumption			Consumption per capita		
	Full Sample	Urban	Rural	Full Sample	Urban	Rural
	(1)	(2)	(3)	(4)	(5)	(6)
α: Log HH Income	0.014***	0.015***	0.013***	0.014***	0.015***	0.012***
	(0.002)	(0.005)	(0.003)	(0.002)	(0.004)	(0.003)
Dep. Var Mean	14.60	15.40	14.32	13.35	14.25	13.04
Observations	9942	2391	7303	9942	2391	7303
No. of HH	2819	681	2070	2819	681	2070
R^2	0.68	0.75	0.61	0.71	0.79	0.64

NOTES—This table presents results from Uganda for a test of full insurance among geographic community members. Two-way FE specification as in equation 1 is estimated where for columns (1)-(3) the outcome of interest is the log of household consumption (in UGX) and for columns (4)-(6) the outcome of interest is log of consumption per adult equivalent (in UGX) over the past 12 months. 1 USD = 2600 UGX (2014). Standard errors clustered at the enumeration area level (5-10 HHs). *p<0.1, **p<0.05, ***p<0.01.

Results from the Ugandan panel reported in Table 2 further corroborate this narrative. Comparing elasticity estimates across Ghana and Uganda, households in Uganda do a near-perfect job of insuring against shocks to their income.

5.1.1 Transfers and remittances

A subtle yet important point about these results is that the estimates in Tables 1 and 2 tell us nothing about the *nature of insurance* at play in hedging against risks. Self-insurance in the form of precautionary savings as well as transfer arrangements among member of the informal insurance network have been long documented means for insurance, especially among rural households (Dercon, 2002).

Analyzing how transfers and remittances vary with respect to income shocks provides some suggestive evidence on the nature of risk pooling. In particular, non-zero elasticities of transfers to income help establish the existence and role of informal networks in insuring households. Tables 3 and 4 encapsulate these insights from Ghana and Uganda, respectively. Anytime a household has a positive income shock, it is more likely to send out remittances. While the data lacks details on whom these transfers were sent

to, Ghanian households in particular seem to participate in some form of risk pooling through transfers among their community.

Table 3: Ghana Panel data: Remittances sent

	HH remittances sent		
	Full Sample	Urban	Rural
	(1)	(2)	(3)
α: Log HH Income	0.052***	0.049***	0.053***
	(0.008)	(0.015)	(0.010)
Dep. Var Mean	1.46	1.62	1.37
Observations	12025	4332	7693
No. of HH	4426	1657	2769
R^2	0.66	0.68	0.65

NOTES—Two-way FE specification as in equation 1 with the outcome of interest as the log of household remittances *sent* (in Cedis) over the past 12 months. 1 USD = 1.3 Cedi (2009). Standard errors clustered at the enumeration area level (15 HHs). *p<0.1, **p<0.05, ***p<0.01.

Overall, estimates from the Ugandan data point to a similar pattern but it masks interesting heterogeneity: urban households seldom send remittances when experiencing a positive income shock. This runs counter to the well-established pattern of remittances flowing out from urban locales. Perhaps these households engage in boosting precautionary savings during income boosts. In Table 4, column (5), though imprecisely estimated, a 10% increase in the income of an urban household corresponds to a 0.6% decline in in-transfers. The possibility of weak insurance networks—perhaps due to the size and anonymity of the city and ensuing difficulties in monitoring—and reliance on self-insurance in cities cannot be ruled out. Absent detailed data, inferring the precise nature of such arrangements among urban households is extremely challenging yet equally vital to enhance our understanding of the urban poor.

Table 4: Uganda Panel data: Remittances & In-transfers

	HH Remittances sent			HH in-transfers		
	Full Sample (1)	Urban (2)	Rural (3)	Full Sample (4)	Urban (5)	Rural (6)
α: Log HH Income	0.039***	0.008	0.048***	-0.031**	-0.066*	-0.019
	(0.010)	(0.021)	(0.011)	(0.014)	(0.035)	(0.016)
Dep. Var Mean	4.73	5.24	4.58	3.99	4.67	3.79
Observations	10951	2513	8192	8269	1905	6178
No. of HH	2924	697	2162	2915	691	2159
R^2	0.61	0.64	0.61	0.65	0.69	0.64

NOTES—Two-way FE specification as in equation 1 is estimated where for columns (1)-(3) the outcome of interest is the log of household remittances *sent* (in UGX) and for columns (4)-(6) the outcome of interest is log of transfers *received* by the household (in UGX) over the past 12 months. 1 USD = 2600 UGX (2014). Standard errors clustered at the enumeration area level (5-10 HHs). *p<0.1, **p<0.05, ***p<0.01.

5.2 Validating Li & Ligon (2020) results

While, I am able to reject the null of full insurance, there is a glaring point of departure in the magnitude of estimates of α—the sensitivity of household consumption to that of household income—compared to those of Li and Ligon (2020). Regardless of rural or urban households across Ghana and Uganda, my estimate of α hover between 0.03 to 0.05, which is ten times smaller than that estimated with the *same* specification in Li and Ligon (2020). Incorrect specification, measurement error, and/or erroneous construction of incomes and consumption expenditures could be leading causes for the attenuation of the estimate.

To ensure I am estimating the correct specification, I can estimate equation 10 on the publicly available dataset used by Li and Ligon (2020): the Kagera Health and Development Survey (KHDS) from Tanzania. My ability to replicate the authors' results would serve as a falsification test against running a misspecified test and lend further credibility to the estimates reported in the previous subsection. Any differences in the estimates can be then attributed to measurement error or something fundamentally different about the nature of risk pooling across contexts and time.

5.2.1 Testing for misspecification

Table 5 presents the results from the replication exercise, which indicate no evidence of misspecification. While the coefficient (0.487) does not map exactly to the one reported in column (1) of Figure 1, the difference is negligibly small. This deviation is likely due to difference in the sample sizes and any other data processing procedures the authors did not mention in the manuscript text or in table footnotes.[9]

[9] While, I could not trace the replication package files for the paper, the authors delve into some details of the sample construction and data processing. This proved to be instructive in re-generating their estimates.

Table 5: Replicating Table 3, Column 1 of Li & Ligon (2020)

	Log household consumption per capita (1)
Log household income	0.487***
	(0.023)
Dep. Var Mean	11.71
Observations	2198
R^2	0.87

NOTES—Two-way FE specification as in Table 3, column 1 of Li & Ligon (2020) is estimated. Specification includes household and year (wave) fixed effects. No demographic controls are added. Standard errors clustered at the community level. *p<0.1, **p<0.05, ***p<0.01.

Table 3
Testing the full insurance model.

Variable	(1)	(2)	(3)	(4)
Log household income	0.4177*	0.4184*	0.4028*	0.4166*
	(0.0225)	(0.0228)	(0.0251)	(0.0223)
Demographic controls	No	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Community-year FE	No	No	Yes	No
District-year FE	No	No	No	Yes
N	2253	2253	2253	2253

Notes: This table shows the regression results of testing the full insurance model. The dependent variable is log household consumption per capita (per adult equivalent). All variables are in per capita terms when appropriate. Demographic controls include the age of the household head and two indicators for whether the head is female and has received any formal schooling. The balanced panel has a total 2256 observations, three of which have negative household income and are dropped from the regression analysis. Standard errors (in parenthesis) are clustered at the community level. *$p < 0.01$.

Figure 1: Table 3 from Li & Ligon (2020)

That the estimates match closely to the reported tables, provides sufficient evidence for correctly estimating the specification in equation 10. There appears to be—in stark contrast to rural Tanzania in the early 1990s—much more effective risk pooling among communities in Ghana and Uganda across the 2010's. How this arrangement compares to alternate definitions of informal networks is the logical extension of this exercise, to which I turn to in the next section.

5.3 Risk pooling with communities of origin

Rural-urban migration is a key driver of economic development and structural transformation (Lewis, 1954). In developing economies, rapid urbanization has left much of the urban labor force composed of

recent migrants. This transition raises a central question: which informal networks matter for risk pooling in the city? Do rural ties continue to mitigate urban risks, or are these functions replaced by newly forged urban networks? Even in the village economies, with relatively lower mobility, this question has spawned a vast literature. Munshi and Rosenzweig (2016); Mazzocco and Saini (2012); Rosenzweig and Stark (1989) show that in the Indian context castes and sub castes are the effective network for informal insurance. Other studies show pooling risk among village residents serves as an effective insurance arrangement (Deaton, 1990; Townsend, 1994; Morduch, 2005).

To assess the relevance of insurance networks among urban households, I estimate equations 11 and 12 on the Ghana panel data and present the results in Tables 6 (first differences specification) and 7 (growth rate specification). Columns (1) and (2) exhibit results from a definition of networks setup at a geographically proximate level and includes households within the same enumeration area. This is the analog of defining the insurance network as the village in rural settings. Estimates in columns (3) and (4) are derived from an alternate network definition: all households in the district of origin. This setup represents the ties urban migrants may have to their places of origin. Here, full insurance is a joint test of household consumption co-moving with aggregate consumption ($\beta_1 = 1$) and having no sensitivity to shocks to household income ($\beta_2 = 0$).

Table 6: Comparing insurance networks (levels)

	Residence EA		Origin District	
	Full Sample	Urban Migrant	Full Sample	Urban Migrant
	(1)	(2)	(3)	(4)
$\beta_1 : \Delta C_t^a$	1.000***	1.116***	0.510***	0.437***
	(0.000)	(0.079)	(0.077)	(0.125)
$\beta_2 : \Delta Y_t^j$	0.000	0.006**	-0.000	0.006*
	(0.000)	(0.003)	(0.000)	(0.004)
F-stat: $\beta_1 = 1, \beta_2 = 0$	0.05	3.60	22.62	10.86
Dep. Var Mean	3225.49	4808.72	3324.13	4808.72
Observations	7557	732	2271	732
No. of HH	4414	440	1298	440
R^2	0.12	0.17	0.02	0.03

NOTES—First difference regression (Mace 1991) is estimated where outcome of interest is HH consumption (in Cedis) over the past 12 months. Regressors include HH income as well as aggregate consumption. 1 USD = 1.3 Cedi (2009). Standard errors clustered at the enumeration area level (15 HHs). *p<0.1, **p<0.05, ***p<0.01.

Results from the first difference specification (Table 6) suggest close to full insurance among urban migrant households in the enumeration area (EA) and strong evidence for the severance of insurance-based ties with those in district of origin. It is tempting to infer this result as a rejection of rural-urban network linkages but data limitations preclude a fair comparison. Longitudinal survey data, across Ghana

and Uganda, does not capture the village of origin of any respondent. A district is a large administrative boundary consisting of hundreds—if not thousands—of households across multiple enumeration areas in such settings. Lacking granular information on place of origin forces a comparison of risk sharing networks within an enumeration area to that across tens of such areas.[10] With no unifying feature—except a common administrative boundary—and logistical constraints to set up insurance arrangements, there is little reason to believe the district is a suitable unit of informal network formation. In summary, data constraints hamstring our ability to derive meaningful patterns in rural-urban linkages and assess the relevant network for urban migrants in pooling risk.

Table 7: Comparing insurance networks (logs)

	Residence EA		Origin District	
	Full Sample (1)	Urban Migrant (2)	Full Sample (3)	Urban Migrant (4)
$\beta_1 : \Delta \ln C_t^a$	0.685***	0.847***	0.351***	-0.050
	(0.036)	(0.092)	(0.121)	(0.210)
$\beta_2 : \Delta \ln Y_t^j$	0.082***	0.163***	0.096***	0.179***
	(0.016)	(0.052)	(0.024)	(0.053)
F-stat: $\beta_1 = 1, \beta_2 = 0$	47.30	5.62	24.17	17.06
Dep. Var Mean	1.55	1.69	1.52	1.69
Observations	3692	460	1198	460
No. of HH	2778	329	894	329
R^2	0.11	0.19	0.03	0.06

NOTES—Growth rate regression (Mace 1991) is estimated where outcome of interest is log HH consumption (in Cedis) over the past 12 months. Regressors include log HH income as well as logged aggregate consumption. 1 USD = 1.3 Cedi (2009). Standard errors clustered at the enumeration area level (15 HHs). *p<0.1, **p<0.05, ***p<0.01.

For the sake of completeness[11], I present the growth rate regression results in Table 7. There appears to be little evidence of risk pooling, regardless of the definition of the network.[12] The data granularity constraint applies squarely to this set of results too. To (partially) overcome this constraint, I pivot to the details of pilot survey data I recently collected in the next section.

[10] To those unfamiliar with the scale and population density of districts in low income settings, a district can be as large as half the size of the average US city. Comparing risk pooling within a census block group to that within a borough in New York City ends up teaching us nothing worthwhile.

[11] The same issue persists in the Ugandan panel data too. I omit the results of that analysis to avoid over interpretation of results with known shortcomings. I continue to explore large scale longitudinal datasets that may record the place of origin at a more granular level.

[12] This pattern is consistent with the results in Mace (1991). She, too, finds such a divergence in estimates across specifications and posits that lower-income households in the growth rate regression drive the results.

6 August 2025 Uganda Pilot Survey

To supplement the large-scale, longitudinal data and add more granularity to rural-urban network linkages, and the type of shocks and coping strategies of urban workers, I conduct a survey of 148 workers across major manufacturing sectors across Kampala, Uganda: carpentry, welding, and tailoring. This data collection effort is embedded within a larger ongoing baseline survey of firms and workers for joint work[13] with Virginia Minni (University of Chicago, Booth School of Business) and Anna Vitali (New York University).

In that project, we study the relationship between employers and workers and its implications for aggregate efficiency and firm productivity, most notably manifest in worker contracts, high self-employment rates, and small firm sizes prevalent in Uganda and other low-income countries. While that project tackles the efficiency implications of labor arrangements, this paper is focused on identifying the sources of risk for urban workers, how these risks are diversified, and the role informal networks play in dealing with volatility—a defining feature of workers' lives in the city.

6.1 Survey sample

The sampling frame for survey is the universe of 1516 firms that have a permanent employee based on a listing exercise we conducted last year. Stratified by firm sector $\times$ firm density, we randomly sample 900 firms for the survey. Finally, based on the current roster of permanent employees, a worker is then randomly selected for the worker survey.

To minimize data collection costs, I leverage on the complementarities from this ongoing survey effort to introduce a detailed module on risk and insurance in the worker survey. This module was administered to 148 workers, primarily across carpentry and welding sectors, and forms the basis for the analysis that follows in Section 6.3. Because we were already collecting some details on workers' consumption, income, and savings (for a weekly and fortnightly recall period), I use this data from 573 worker surveys to estimate tests of full insurance and assess the relative importance of plausible networks in hedging against shocks.

6.2 Key variables

In addition to income and consumption data, I collect a host of data on workers' migration and location decisions in the city, the shocks they face and material consequences of the shocks such as absences from work, wage loss, and propensity to get work upon return. The submodule on in-transfers captures details on the sender, the frequency and quantum of transfer, the use-case for the transfers and the means through which these are sent. Inspired by recent academic work (Miner, 2025; Ndiaye et al., 2025; Abebe et al., 2024) and policy maker interest (Guven et al., 2021) in the provision of social safety nets for urban residents, I also elicit demand for an unemployment insurance scheme.

6.3 Results

In spirit of the previous analysis testing for risk pooling, I present estimates from this test in Table 8. A few notable differences from the prior specification are worth noting. First, consumption expenditures and

[13] Funded by International Growth Center (IGC). See here for a brief summary

incomes are at the worker, not household level. Second, eliciting information on income and consumption for the past week and two weeks before delivers a panel structure to the data. Finally, the implied network definition is that of all sample manufacturing workers (across carpentry, tailoring, and welding) in the same parish. As I sample one worker per firm, these results speak to risk pooling among worker *across* firms, not *within* the firm.

Table 8: Uganda survey data: Testing for full insurance

	Consumption in:		
	(1) Log	(2) Log(c+1)	(3) Levels (USD)
Log income	0.042**		
	(0.017)		
Log (income+1)		0.037	
		(0.023)	
Income (USD)			0.028**
			(0.011)
Dep. Var Mean	11.44	11.41	30.78
Observations	1018	1090	1090
No. of workers	509	545	545
R^2	0.96	0.93	0.94

NOTES—This table presents results from tests of risk pooling among workers in Kampala, Uganda. Estimates reported are from a TWFE specification following equation 10 where outcome is consumption and regressor of interest is income. Worker and parish × week FEs are included but not reported. Results across transformations of income and consumption shown. 1 USD = 3600 UGX (2025). Standard errors clustered at the parish level. *p<0.1, **p<0.05, ***p<0.01.

Columns (1) through (3) report estimates over different transformations of consumption. Although the coefficients vary slightly, the overall result rejects full insurance but is suggestive of some risk pooling among workers in a parish. Income shocks don't appear to co move with consumption, after controlling for parish level incomes of the workers. Perhaps, risk pooling among workers within the same sector is a more relevant network definition. I test for that and present results in Table A1. The results are stable across definitions suggesting a combination of other informal networks and self-insurance at play. It bears repeating that this test does not shed any insight on how workers are insuring themselves; it merely tests for the existence of risk pooling.

Exploring the nature of shocks workers face and the resultant coping strategies can shed insights beyond those gained by tests of full insurance. The survey module I administered offers some descriptive evidence to study this underexplored domain further.

6.3.1 Health shocks are common and have implications

Figure 2 exhibits illness or injury to the worker or to someone in their household as the leading cause of adverse events. More than half the sample of workers has had a health shock in the past month. These shocks result in absences from work, which expectedly have ramifications on wage loss. Figure 3 portrays how a health shock results in typically 2-3 days of absence from work, totalling about 70,000-80,000 UGX or a fifth of the average monthly earnings (Figure B1) .

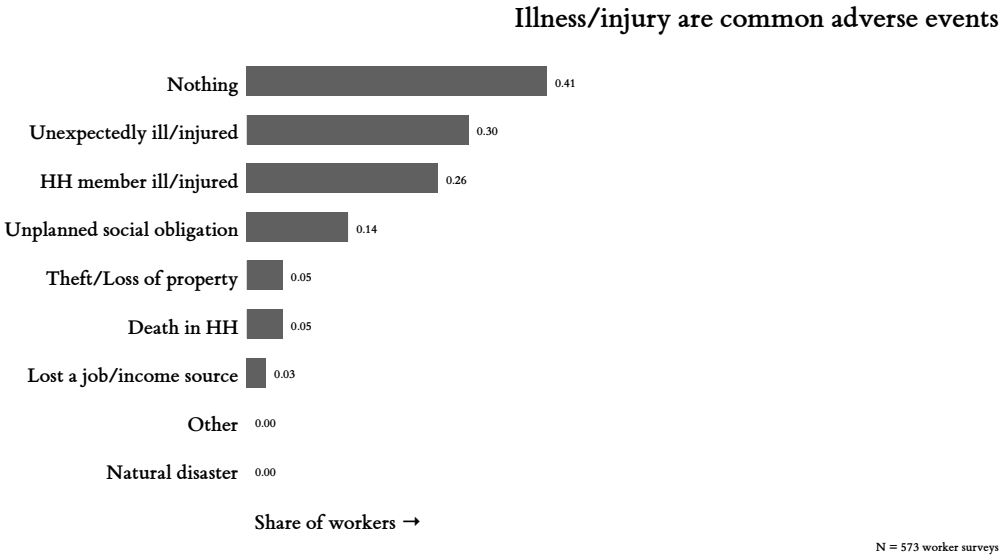

Figure 2

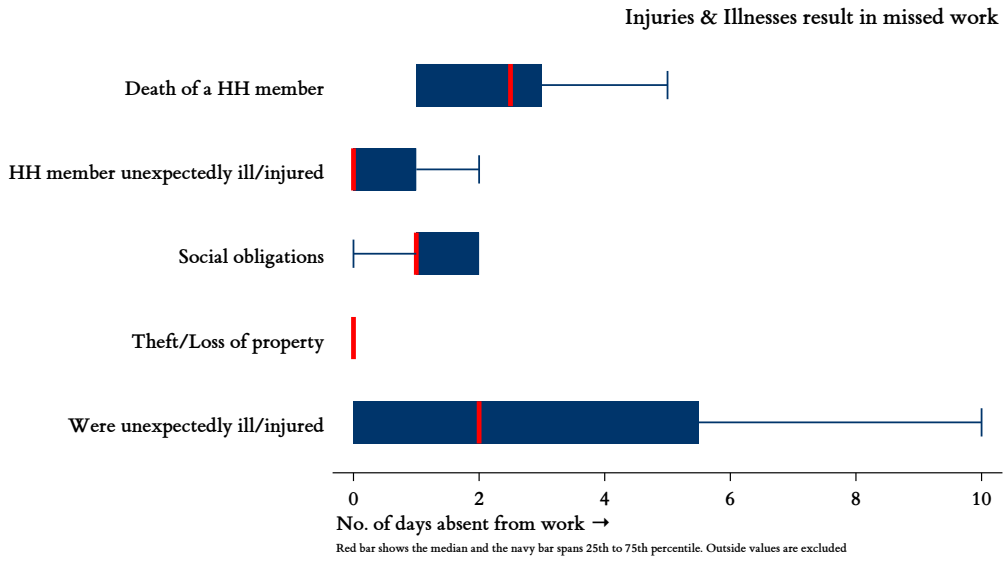

Figure 3

6.3.2 Self-insuring against shocks delays employment trajectories

Eliciting workers' coping strategies reveals the role of self-insurance through precautionary savings. Despite its existence, large-scale survey data were lacking in diagnosing the source and modality of insurance. Figure 4 documents a staggering 90% of 82 workers who faced a shock resorting to savings to maintain income and consumption during that period. Interestingly, the data also affirms the role of labor supply in responding to shocks (Tang, 2024).

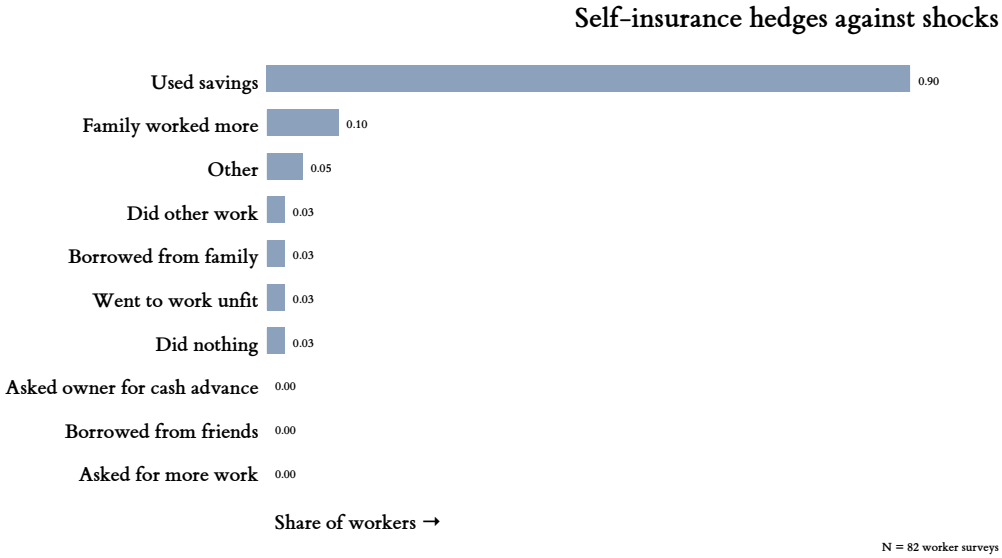

Figure 4

A deeper implication of relying on precautionary savings to cope with shocks is manifest in its effects on the employment trajectory of workers. Almost all workers aspire to be self-employed (Figure 5). However, as Figure 6 shows, the inability to access capital is their biggest constraint.

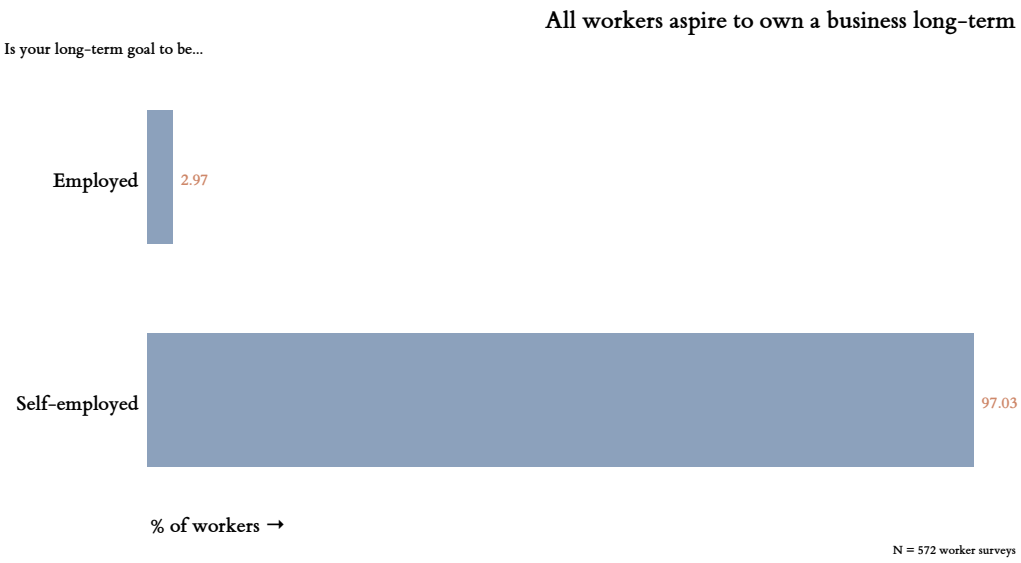

Figure 5

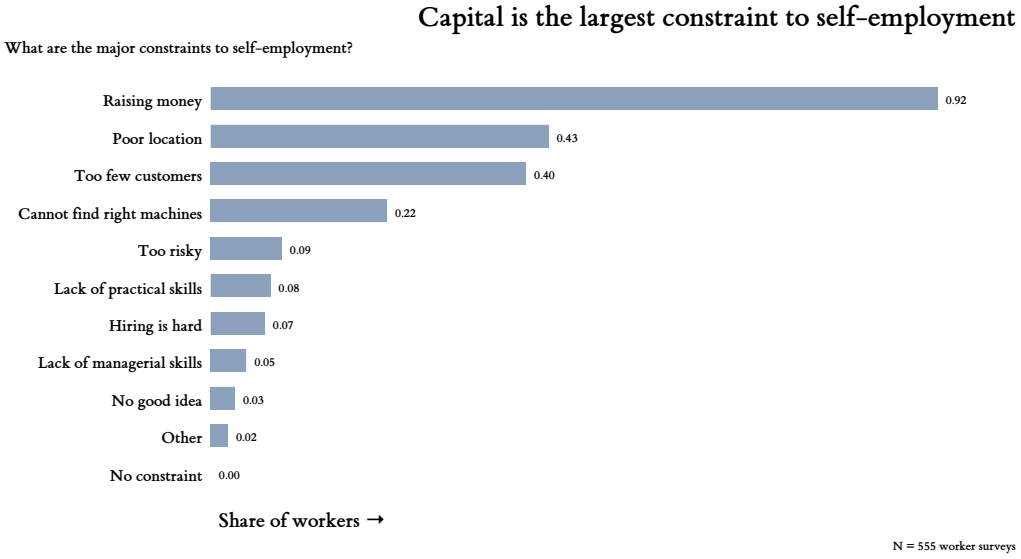

Figure 6

Figure 7 represents a common solution out of this bind: accumulate savings[14] to build up capital. However, this is a long journey that is fraught with difficulties. As Figures 8 and 9 reflect, with median savings of 600,000 UGX, most workers are currently at 17-20% of their goal. Frequent shocks compound this pool of savings as the worker is forced to draw from it to mitigate the variability associated with shocks.

[14] Two thirds of workers' savings in the sample are bookmarked to this goal, which calls to question the standard rationale of the precautionary nature of savings.

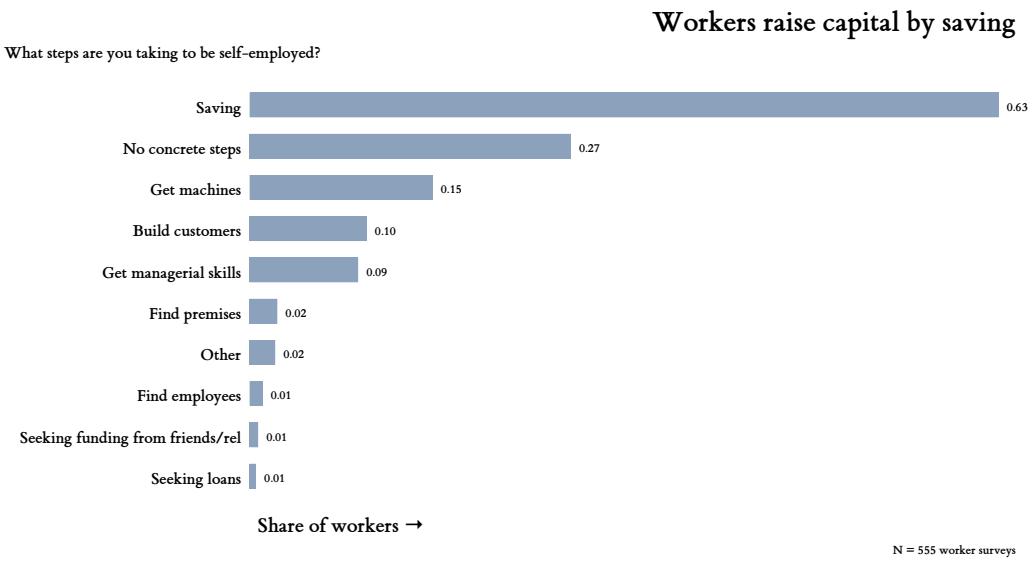

Figure 7

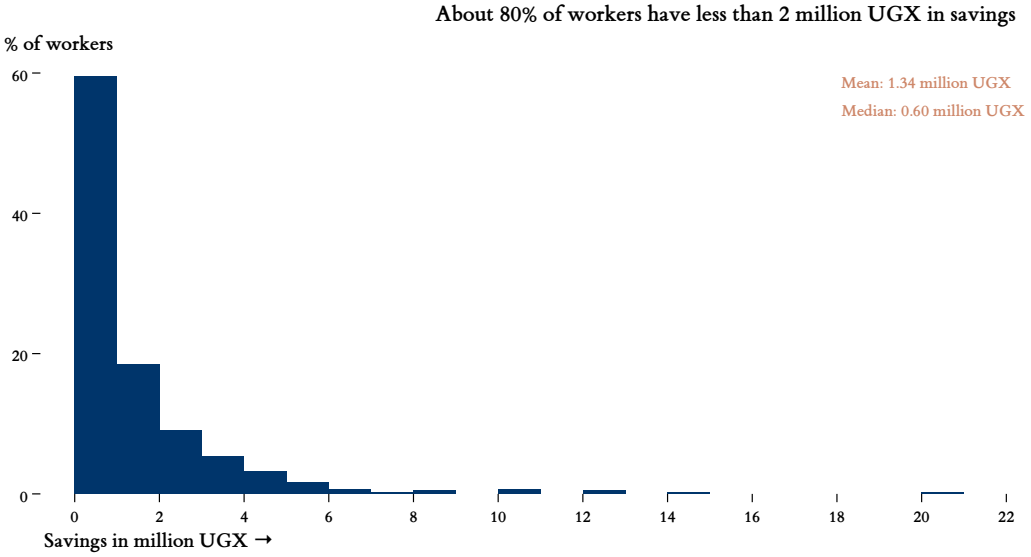

Figure 8

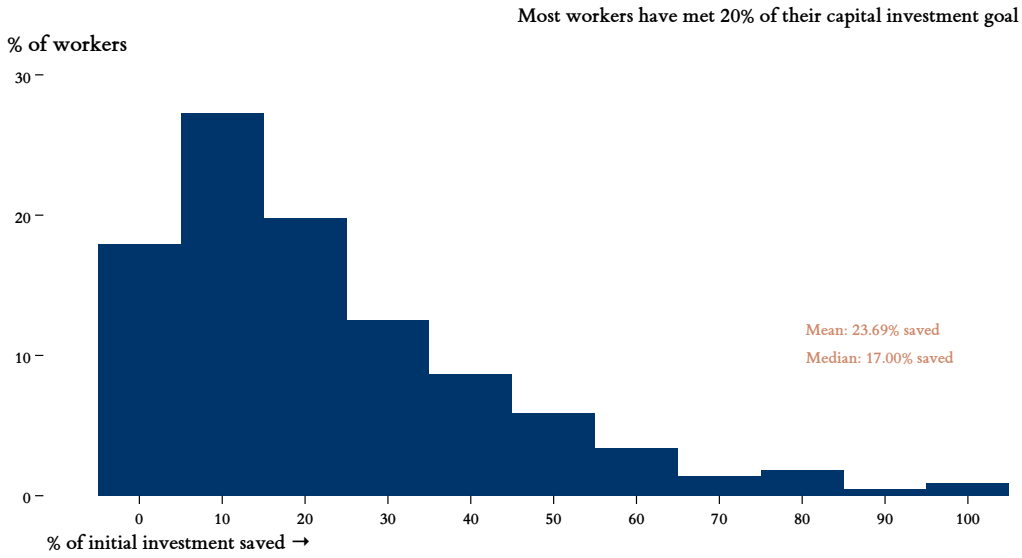

Figure 9

Consequently, absent financial markets, only rich workers may end up selecting into self employment, regardless of their productivity. This delays—if not denies—entry into entrepreneurship, which has distributional and allocative implications for the economy.

6.3.3 Workers appear to value networks but don't extract full value

The survey data reveals a tension that is hard to reconcile: despite valuing social networks tremendously, it is puzzling why there aren't there more transfers among the network. Figures 10-12 illustrate this wedge. The top reason for choosing to move to and live in a place is the social network. When probed further, it appears that 88% of workers like to co locate with their networks for mutual insurance. Socializing and easier information flows feature prominently but motives to help each other dominate. Yet, as Figure 12 demonstrates, only 30% of workers report getting any transfer in the past 12 months.

Misleading stated preferences and social desirability bias for under reporting transfers could be leading candidates to resolve this tension. However, the size of the city and along with it the myriad challenges of maintaining a credible insurance network may also feature in explaining the under utilization of risk pooling among networks.

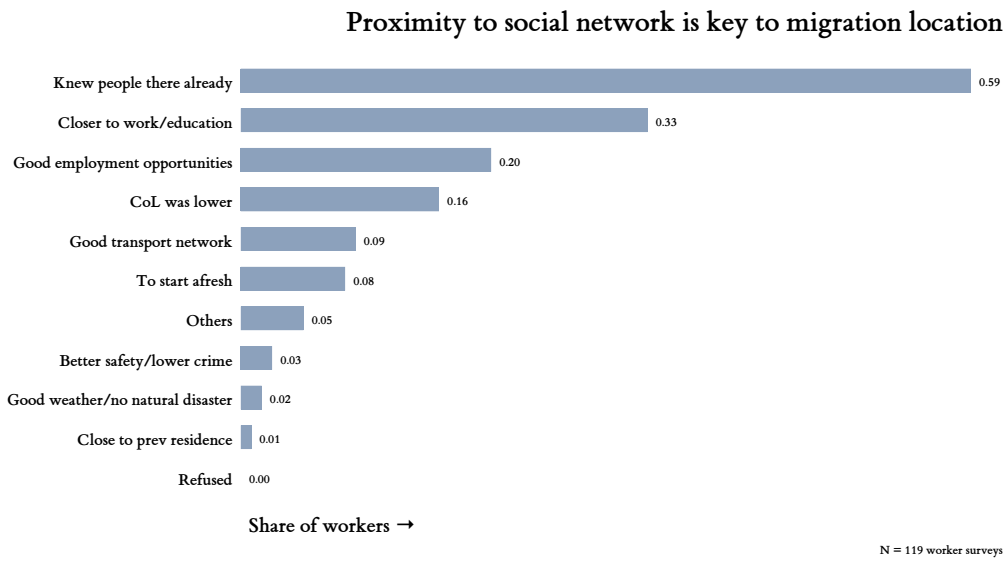

Figure 10

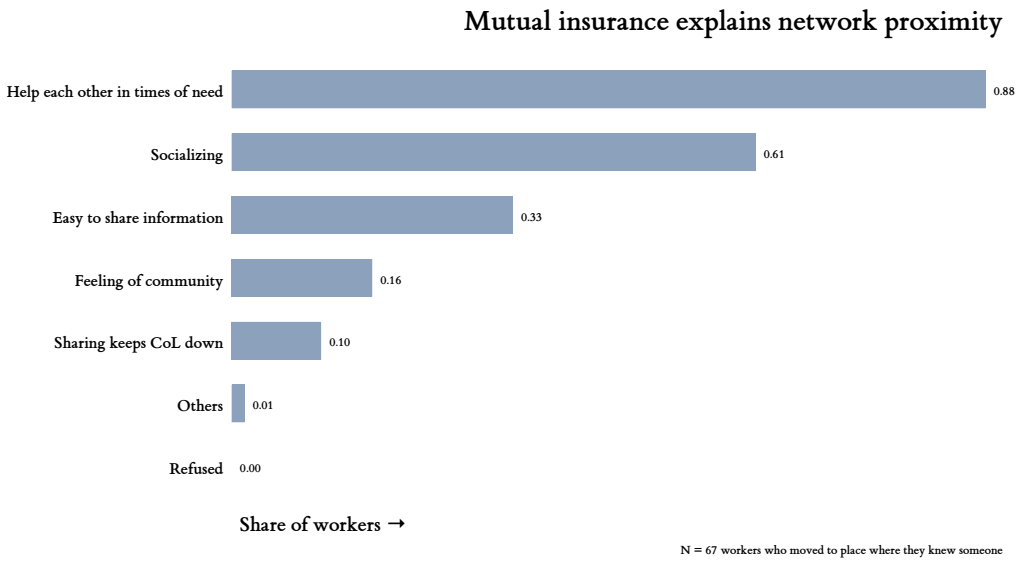

Figure 11

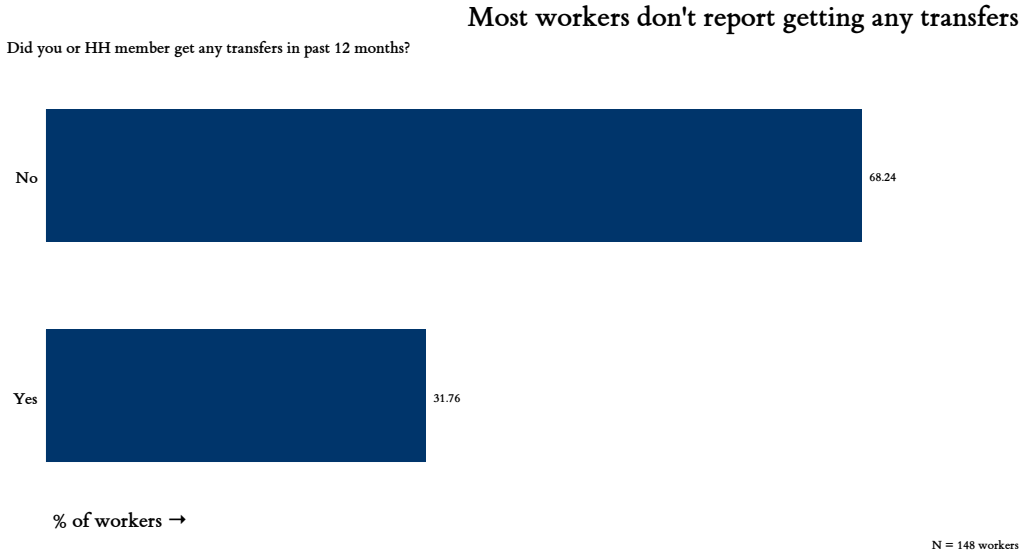

Figure 12

However, for those who receive transfers, they use it to smooth consumption. Family and relatives along with friends are the usual senders and do so at a monthly frequency through mobile money. Moreover, the workers are in frequent contact with the senders, which is indicative of the demands on their time to maintain such relationships. Interestingly—in my limited sample—I find very little evidence of risk pooling among neighbors. One would expect some degree of risk sharing with those in their proximity. The overlap with family and friends due to co location could lead to a classification error: neighbors are essentially friends and family, which is a more salient relationship and reported more often in survey responses. An alternate interpretation could be related to the unique challenges of networks in a city, characterized by transience and anonymity. Disentangling the causes remains an important next step that I discuss in further detail in Section 7.1

6.3.4 Unemployment insurance takeup mediated by workers' beliefs about job search

Ex-ante, unemployment and job separations seemed to be the biggest risks facing urban workers. A high degree of disguised unemployment (Lewis, 1954), little job training, and the growing discourse around the design of social safety nets in low income countries led me to elicit demand for an unemployment insurance program. Existing informal arrangements for insurance could also depress demand for safety nets. While, job losses may not be the largest source of uncertainty for workers in Uganda, I document some interesting findings about the feasibility of such a program.

Adapted by Ndiaye et al. (2025), I present a vignette of an unemployment program that has a 75% replacement rate, goes on for 6 months, and costs 5% of monthly earnings to fund.[15] Workers appear to be polarized about the take up of such a program: a slim majority (57% in Figure 13) are likely to take up this program. Figure 14 plots the various reasons why or why not a worker might be interested in this scheme. Strikingly, beliefs around the ease of finding a job dictate the workers' propensity to accept such a proposal. 41% of those who are unlikely to take up this program, feel confident in keeping their job or believe

[15] The inclusion of the premium, which is typically funded by tax, was introduced to rule out a free lunch possibility.

they would quickly find a new one soon enough. In line with prior work on worker beliefs (Abebe et al., 2025), this finding underscores an understated factor in the design of social safety nets and identifies an unexpected implication of workers' beliefs.

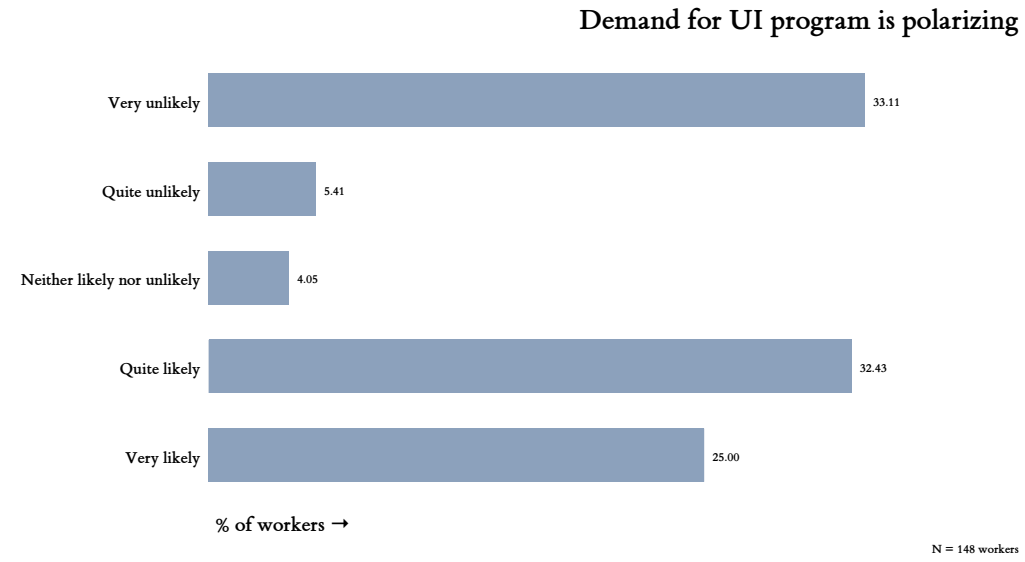

Figure 13

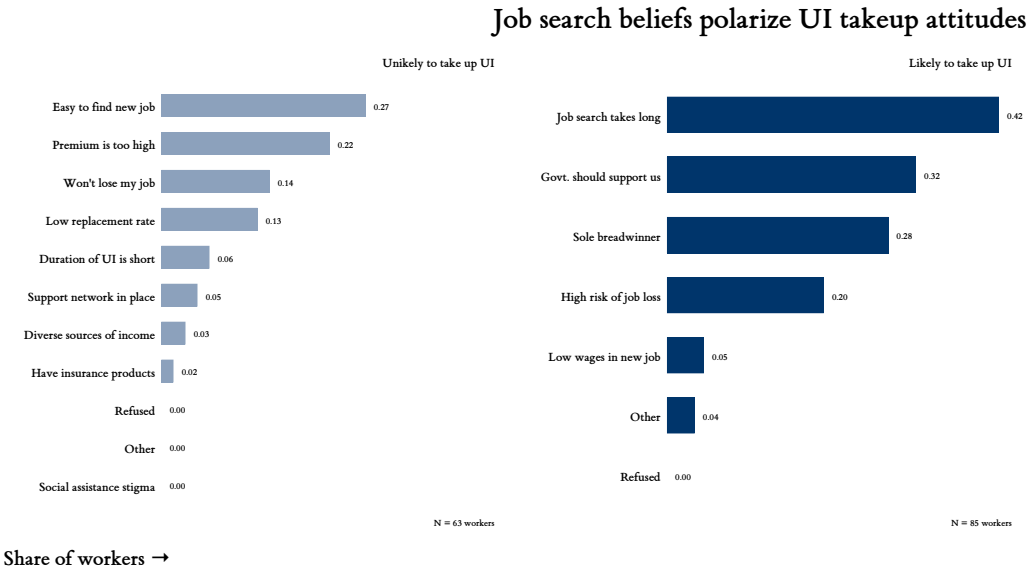

Figure 14

7 Conclusion & Next Steps

In this paper I test for the presence of risk pooling and document some insights on the state of insurance in urban settings across low income contexts. Using large scale longitudinal data from Ghana and Uganda, I am able to reject the null of full insurance in urban settings, which is in line with the findings of the expansive literature based on village economies. Despite not meeting this benchmark, households' consumption is not sensitive to income shocks: the consumption elasticities to income vary between 0.01 to 0.05.

Self-insurance and reliance on informal networks to pool risks are the modal strategies to manage uncertainties. For risk pooling among informal networks, I test for the relevance of different possible networks in the city. Comparing definitions of network based on co-location to that formed through links to places of origin, I find suggestive evidence in support of the former. As the large scale surveys just capture the district of origin (not the exact village/town), a comparison of the networks to assess the relevant informal risk pooling group is biased. I circumvent this issue by collecting my own data from 150 manufacturing workers in the capital city of Uganda. This data also not only allows me to establish rural-urban linkages in insurance (a work in progress) but also draw insights on the prevalence of self-insurance in the form of savings.

Despite co locating with people in their informal network, a majority of workers still rely on savings to weather idiosyncratic shocks. Only a third of workers in my sample report getting transfers from their networks in the past 12 months. Together these findings have deeper implications on the workers' employment trajectories: constant shocks and the resulting drawdown of savings delays workers' move into entrepreneurship—an aspiration virtually all workers hold. That workers base migration and residence decisions on their networks and yet don't fully optimize on the benefits presents an conundrum that I will explore further. As exploratory analysis, I bring evidence to the role workers' job finding and separation rates play their demand for unemployment insurance. These results can inform the design of social safety nets in developing contexts—a nascent yet important area of study.

7.1 Looking ahead

As this is an ongoing project, I outline the next set of steps in this section to drive this study further and steer this into a dissertation project.

Expanding panel results to Kenya & Indonesia Typically, large scale surveys such as the ones I use in the analysis don't capture the place of origin or the place a respondent moved from with high level of granularity. Absent this information, assessing the correct network for risk pooling is not feasible. The Kenya Life Panel Survey (KLPS) and the Indonesia Family Life Survey (IFLS) are notable exceptions that track this information. I am in the process of acquiring this data and plan to extend the analyses in Section 5 to these datasets.

Refining the panel data estimates The literature on risk sharing and informal arrangements in low-income contexts is as vast as it is deep. The analysis so far draws on canonical papers and tests in the field and extends it to shed light on the state of insurance in urban settings. However, recent theoretical and empirical advances Fafchamps and Shrinivas (2022); Ambrus and Elliott (2021); Ambrus et al. (2014) necessitate using techniques at the frontier. For instance, informed by the prevalence of savings in

the Ugandan context, I can adapt the model outlined by Fafchamps and Shrinivas (2022) for asset building and heterogenous risk preferences (Mazzocco and Saini, 2012) to rationalize the patterns in the data. Additionally, adding bias correction adjustments for robustness is another low-hanging fruit for me to harvest.

Analyzing rural-urban linkages from survey data A novel feature of my survey data is that I capture details on senders and recipients of transfers. Analyzing this data on names, relations, and locations can yield rich insights into details of one's insurance network as well as any linkages between rural and urban areas. The latter speaks directly to the process of structural transformation. Finding any evidence for counter-cyclical remittances between these regions will be a useful addition to our knowledge of this process. This remains a priority as an extension of this project.

Collecting survey data to study mechanisms The survey data I collected this summer was key to understanding how workers cope with uncertainty in the cities where migration, networks, and geography interact to constrain effective risk pooling. Conducting another round of surveying with an enhanced focus on pinning down the implications of self-insurance in a job ladder model will be beneficial for linking insurance with aggregate productivity and inequality. Building on Miner (2025) work arguing for migration linked unemployment benefits, I can simulate counterfactuals of partial and full insurance—through networks, savings, and publicly provided safety nets—to study misallocation and the welfare implications under each scenario.

References

Abebe, G., S. Caria, M. Fafchamps, P. Falco, S. Franklin, S. Quinn, and F. Shilpi (2025). Matching frictions and distorted beliefs: Evidence from a job fair experiment. *The Economic Journal*, ueaf026.

Abebe, G., S. Caria, S. Spaziani, L. Hensel, and F. Gerard (2024). Designing severance insurance: Theory and evidence from ethiopia. https://ssrn.com/abstract=5367418. SSRN Working Paper No. 5367418.

Alfonsi, L., O. Bandiera, V. Bassi, R. Burgess, I. Rasul, M. Sulaiman, and A. Vitali (2020). Tackling youth unemployment: Evidence from a labor market experiment in uganda. *Econometrica 88*(6), 2369–2414.

Ambrus, A. and M. Elliott (2021). Investments in social ties, risk sharing, and inequality. *The Review of Economic Studies 88*(4), 1624–1664.

Ambrus, A., M. Mobius, and A. Szeidl (2014). Consumption risk-sharing in social networks. *American Economic Review 104*(1), 149–182.

Breza, E. and S. Kaur (2025). Labor markets in developing countries. *Annual Review of Economics 17*.

Chen, J. and J. Roth (2024). Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics 139*(2), 891–936.

Deaton, A. (1990). On risk, insurance and intra-village consumption smoothing. *Preliminary Draft, Research Program in Development Studies. Princeton University*.

Dercon, S. (2002). Income risk, coping strategies, and safety nets. *The World Bank Research Observer 17*(2), 141–166.

Fafchamps, M. and A. Shrinivas (2022). Risk pooling and precautionary saving in village economies. Technical report, National Bureau of Economic Research.

Gollin, D., M. Kirchberger, and D. Lagakos (2021). Do urban wage premia reflect lower amenities? evidence from africa. *Journal of Urban Economics 121*, 103301.

Guven, M., H. Jain, and C. Joubert (2021). Social protection for the informal economy. *International Bank for Reconstruction and Development/The World Bank*.

Imbert, C. and J. Papp (2020). Costs and benefits of rural-urban migration: Evidence from india. *Journal of Development Economics 146*, 102473.

Lewis, W. A. (1954). Economic development with unlimited supplies of labour.

Li, Z. and E. Ligon (2020). Inferring informal risk-sharing regimes: Evidence from rural tanzania. *Journal of Economic Behavior & Organization 177*, 941–955.

Lockhart, K. (2025, July). Yes in my bamako yard. *Asterisk Magazine*.

Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy 99*(5), 928–956.

Mazzocco, M. and S. Saini (2012). Testing efficient risk sharing with heterogeneous risk preferences. *American Economic Review 102*(1), 428–468.

McKenzie, D. (2024, February). What is a firm again? the fluidity of firm boundaries in developing country firms. https://blogs.worldbank.org/en/developmenttalk/what-firm-again-fluidity-firm-boundaries-developing-country-firms. World Bank Blogs.

Miner, G. (2025). Essays in development economics.

Morduch, J. (2005). *Consumption smoothing across space: Testing theories of risk-sharing in the ICRISAT study region of South India*. Oxford: Oxford University Press.

Munshi, K. and M. Rosenzweig (2016). Networks and misallocation: Insurance, migration, and the rural-urban wage gap. *American Economic Review 106*(01), 46–98.

Ndiaye, A., K. Herkenhoff, A. Cissé, A. Dell'Acqua, and A. A. Mbaye (2025). How to fund unemployment insurance with informality and false claims: Evidence from senegal. *Journal of Monetary Economics 150*, 103699.

Pelnik, C. (2025). Moving to profitability? alleviating constraints on microentrepreneur location. *Unpublished manuscript*.

Rosenzweig, M. R. and O. Stark (1989). Consumption smoothing, migration, and marriage: Evidence from rural india. *Journal of political Economy 97*(4), 905–926.

Suri, T. (2005). Spillovers in village consumption: Testing the extent of partial insurance. *New Haven, Conn.: Yale University. mimeo*.

Tang, Y. (2024). To follow the crowd? benefits and costs of migrant networks. Technical report, Working Paper.

Townsend, R. M. (1994). Risk and insurance in village india. *Econometrica: journal of the Econometric Society*, 539–591.

Udry, C. (1995). Risk and saving in northern nigeria. *The American Economic Review 85*(5), 1287–1300.

Vollrath, D. (2024, September). The wrong kind of city? *Asterisk Magazine*.

A Appendix: Additional Tables

A.1 Full insurance test with sector × parish × time FEs

Table A1: Uganda survey data: Testing for full insurance

	Consumption in:		
	(1) Log	(2) Log(c+1)	(3) Levels (USD)
Log income	0.058***		
	(0.021)		
Log (income+1)		0.037	
		(0.024)	
Income (USD)			0.031***
			(0.011)
Dep. Var Mean	11.46	11.42	30.90
Observations	938	1002	1002
No. of workers	469	501	501
R^2	0.96	0.93	0.94

NOTES—This table presents results from tests of risk pooling among workers in Kampala, Uganda. Estimates reported are from a TWFE specification following equation 10 where outcome is consumption and regressor of interest is income. Worker and firm sector × parish × week FEs are included but not reported. Results across transformations of income and consumption shown. 1 USD = 3600 UGX (2025). Standard errors clustered at the parish level. *p<0.1, **p<0.05, ***p<0.01.

B Appendix: Additional figures

B.1 Survey data: additional graphs

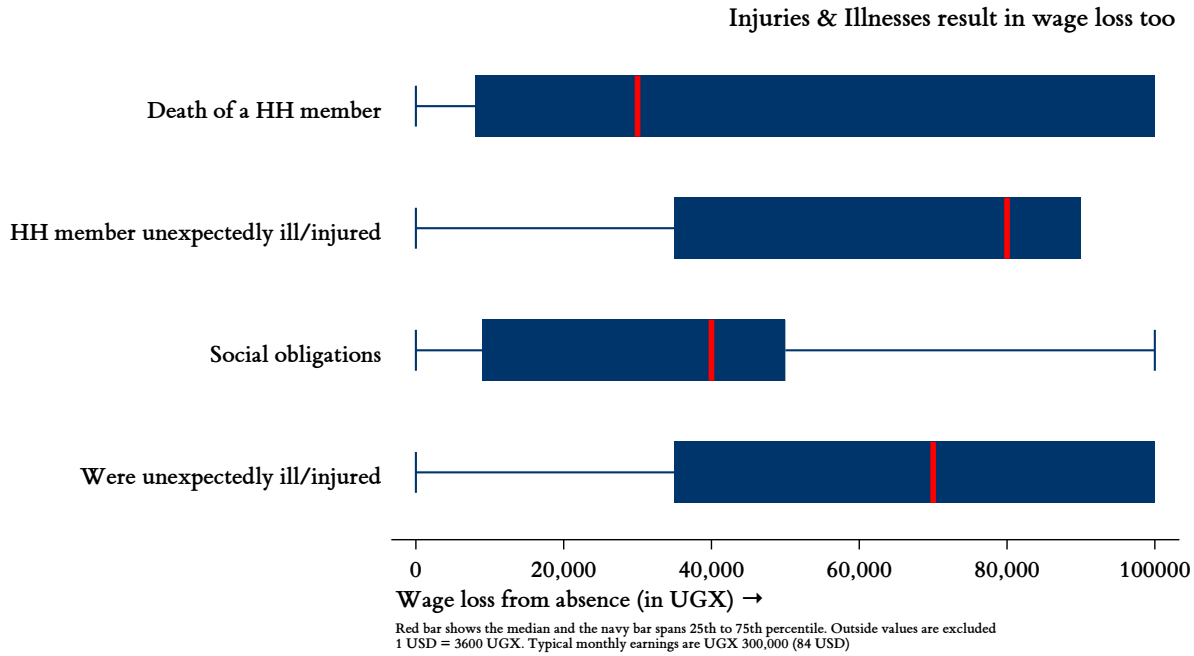

Figure B1

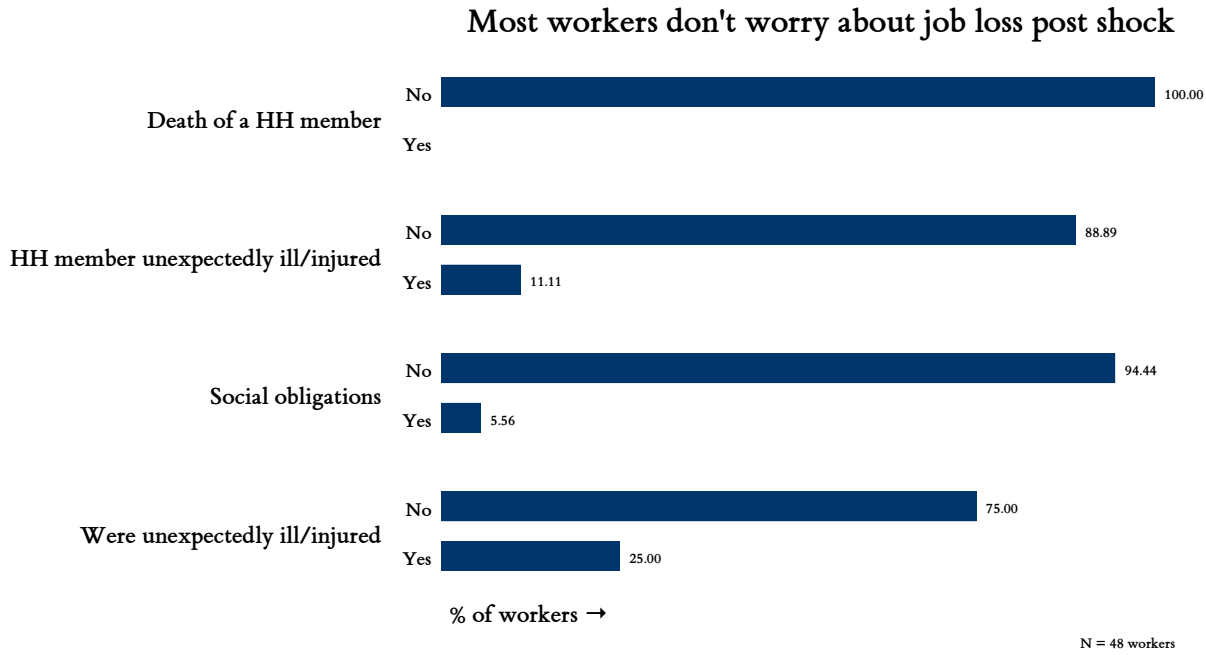

Figure B2

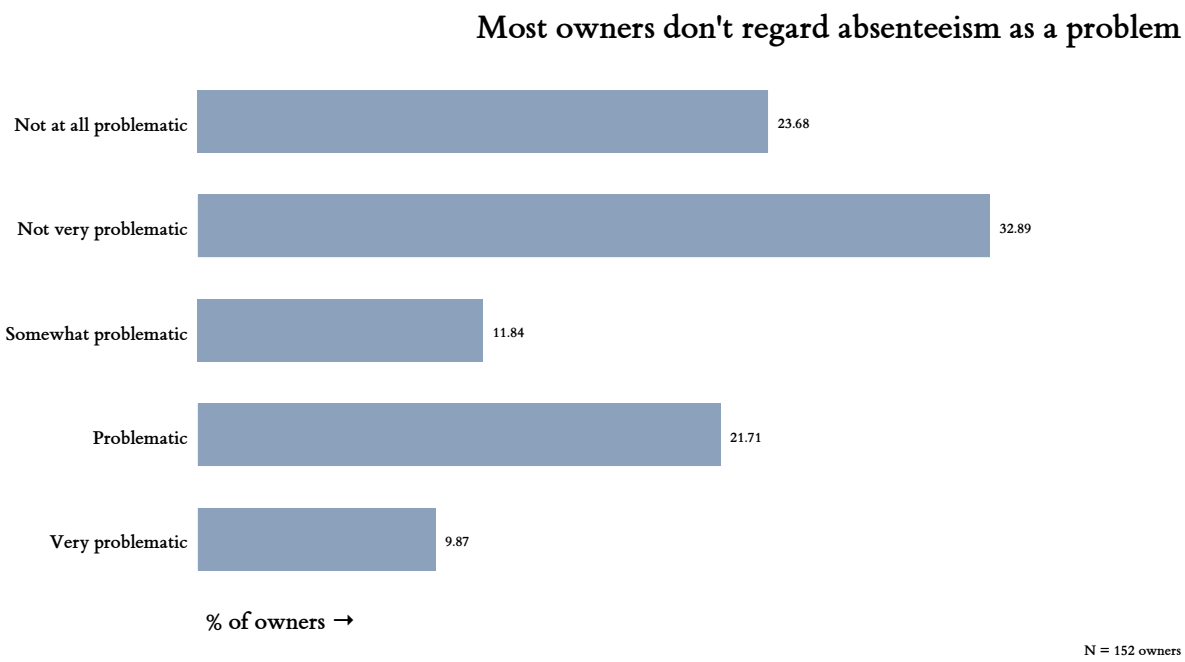
Most owners don't regard absenteeism as a problem
Not at all problematic
23.68
Not very problematic
32.89
Somewhat problematic
11.84
Problematic
21.71
Very problematic
9.87
% of owners →
N = 152 owners

Figure B3