Propagation and Insurance in Village Networks *

Cynthia Kinnan Krislert Samphantharak

Robert Townsend Diego Vera-Cossio

October, 2021

Abstract

In village economies, it is well known that networks can smooth shocks. Less acknowledged is that local production networks can *propagate* shocks. In Thailand, a significant idiosyncratic shock to one household propagates via supply-chain and labor networks. Imperfectly insured households adjust production decisions—cutting input spending and reducing hiring—affecting households with whom they trade inputs and labor. Those linked to shocked households experience reduced local transactions, earnings, and consumption. These declines persist over several years. The total magnitude of indirect effects may be larger than direct effects. Social gains from expanding safety nets may be substantially higher than private gains.

Keywords: Entrepreneurship, Risk sharing, Propagation, Production networks, Firms

JEL Classification: D13, D22, I15, O1, Q12

*Kinnan: Tufts University and NBER, cynthia.kinnan@tufts.edu; Samphantharak: University of California San Diego and Puey Ungphakorn Institute for Economic Research (PIER), Bank of Thailand, krislert@ucsd.edu; Townsend: Massachusetts Institute of Technology and NBER, rtownsen@mit.edu; and Vera-Cossio: Inter-American Development Bank, diegove@iadb.org. We thank numerous colleagues and seminar audiences for helpful suggestions. Townsend gratefully acknowledges research support from the University of Thai Chamber of Commerce, the Thailand Research Fund, the Bank of Thailand, and the Eunice Kennedy Shriver National Institute of Child Health and Human Development (NICHD) (grant number R01 HD027638). Opinions, findings, conclusions, and recommendations expressed here are those of the authors and do not necessarily reflect the views of the Bank of Thailand or the Inter-American Development Bank.

1 Introduction

A pervasive feature of life in developing countries is exposure to risk. These risks take many forms, including weather (see, e.g., Rosenzweig and Binswanger 1993; Maccini and Yang 2009), health (e.g., Gertler and Gruber 2002; Wagstaff 2007), and business conditions (World Bank, 2004).

Linkages between households ("village networks") are another important feature in developing countries. Such networks play an important role in many domains including, but not limited to, geographic and social mobility, information diffusion and technology adoption.[1] These two realities—embedding in social networks and exposure to risk—are understood to overlap in one important way: networks help to buffer risks by making state-contingent transfers (see, e.g., Rosenzweig and Stark 1989; Udry 1994; Townsend 1994).

However, local networks are not limited to risk sharing. Households and their firms[2] are also linked via business transactions, such as the hiring and provision of labor and the buying and selling of inputs and final goods. Such linkages are less studied than risk-sharing networks, but can also have implications for risk: idiosyncratic shocks to the local transaction partners of a given household may represent an important source of risk to that household's earnings. For example, a shock to Household A, which results in shutting or scaling back A's business, may indirectly effect Household B, who previously supplied labor or inputs to A's business.

However, the extent of such indirect risk exposure is an empirical question. If A is well-insured against its shock, the need to scale back business activities

[1] See, e.g., Foster and Rosenzweig 1995; Munshi 2003; Banerjee et al. 2013; de Janvry et al. 2019; Beaman et al. 2021; Meghir et al. forthcoming.

[2] As we discuss below, all households in our data operate at least one agricultural or nonagricultural enterprise, so we use the terms firm and household interchangeably.

is mitigated and the shock is unlikely to propagate. And if B can easily pivot to transacting with Household C, then A's shock will not affect B's livelihood. Thus, understanding how rigid or flexible are these local economic networks is essential, and it is crucial to understand the extent to which significant shocks are insured.[3]

Using detailed data on health and nonhealth spending; gifts; loans; business investment and revenues; and local networks in Thai villages, we find that large health shocks are only partly insured. In consequence, directly shocked households reduce business spending on labor and inputs. And in turn, those households linked in the pre-shock network experience significant *indirect* effects, with local transactions, income, and consumption falling. A back-of-the-envelope calculation suggests that 150 Thai baht of indirect effects for each 100 baht of direct effects, implying a multiplier of roughly 1.5.

The question of whether, and to what extent, idiosyncratic shocks propagate through local networks has important policy implications. For instance, if health shocks propagate to a significant extent, policies that reduce the severity of health shocks and/or improve households' ability to buffer them will benefit those indirectly exposed as well as the direct beneficiaries. Therefore such policies may have significantly larger aggregate impacts than would be measured by looking at direct beneficiaries alone.

Additionally, there remains much we do not know about how the microstructure of village networks relates to networks' effects on risk-sharing and propagation. The question of how overall network structure matters for insurance and propagation of shocks is, to our knowledge, a relatively open one.[4] Comparing across villages, we document more propagation when net-

[3]Previous research has documented that, while small health shocks are well-insured, major shocks pass through to household consumption (Gertler and Gruber, 2002).

[4]Banerjee et al. (2013) study empirically how network structure affects information dif-

works are denser. Comparing across households, we find that shocks that hit more connected (i.e., higher-degree) households propagate more.

Understanding how different network linkages affect production and consumption side outcomes is challenging. We need to observe granular network data on who is linked to who, and how. To understand whether local networks are rigid, we need such information not just at a single point, but over time. To measure how network structure mediates shock propagation, we need to observe multiple networks. Additionally, we need panel data on the production and consumption sides of household balance sheets. And understanding the causal effects of network connections requires identifying shocks which meet several criteria: exogeneity, a scale of shock large enough to "move the needle," and idiosyncrasy (i.e., that the direct impact of the shock is isolated to a given household so that its propagation effects can be measured).

We use a dataset that is uniquely suited to answer these questions. The Townsend Thai data, constructed from 14 years of monthly panel survey data on households in rural and peri-urban Thai villages, contain detailed information regarding transactions among family-operated businesses, which we use to construct labor and supply chain networks. The data also allow us to identify large, idiosyncratic shocks to households' budgets in the form of shocks to health spending needs, and to construct a suitable counterfactual to obtain causal estimates. These elements together provide a setting that can shed light on the role of networks in propagating and mitigating shocks.

We first show that, in violation of the separation theorem, idiosyncratic consumption-side shocks have significant effects on the business activities of

fusion in villages in India. Elliott et al. (2014) examine this question theoretically in the context of links between financial intermediaries; Bigio and La'o (2020) examine the case of input-output networks. To our knowledge the role of network structure in propagation in village economies is not well understood.

the shocked household. When hit by a health spending shock, households reduce business spending in order to smooth consumption—in essence financing the shock out of the business budget. They substantially reduce input spending (by 23%), and almost entirely cut their demand for external labor (by 79%).

This paper is not the first to show that the separation theorem fails (see, e.g., Benjamin 1992; Rosenzweig and Binswanger 1993; Samphantharak and Townsend 2010). Our key contribution is to show that the business-side adaptations by the directly-hit household lead to *indirect* impacts on other, linked, local businesses and workers. To causally identify these impacts, our first empirical strategy relies on variation in the proximity of a given household to the shocked household through pre-shock economic networks. We undertake a generalized difference-in-difference analysis: comparing changes in outcomes before and after each shock, between more-exposed households (those that are closer to a shocked household in the pre-shock network) and less-exposed households (those farther away). Closer households, with greater exposure, see larger falls in total transactions (a 21% decline for a unit change in closeness), and therefore falls in income (12% decline for a unit change in closeness). An alternative identification strategy in the spirit of Fadlon and Nielsen (2019), in which we compare those close to a shocked household (the treated group) to those who are close to a household who suffer a shock, but at a different point in time (the control group), yields similar results.

Although the indirect effects dissipate through the network, there are non-negligible propagation effects on indirectly-connected households (two or more links away from the shocked household) as well as those directly connected (one link away). The shocks that we study generate indirect effects both upstream and downstream, as the costly adjustments taken by the directly-shocked household reverberate through the local network.

Our results are robust to a battery of robustness checks. Our main specification examining direct shock effects uses only shocks occurring in the first half of the sample period to address concerns about difference-in-difference with staggered timing (Goodman-Bacon, 2018); however the results are similar when we include a broader set of shocks. We also obtain similar results when we consider alternate definitions of the onset of a shock and of the placebo group. To support our interpretation that the episodes we identify are shocks to spending needs, rather than solely to labor endowments, we show that the results remain similar when we use only shocks to non-working-age household members. Turning to the indirect effects, we show that indirect exposure via the labor network is associated with a drop in labor transactions and no effect on supply chain transactions, and vice versa for the supply chain network. We also include flexible controls for network centrality-by-time effects. These patterns alleviate the concern that indirect exposure is picking up other differences between exposed and unexposed households. We also show that there is no treatment effect on the provision of uncompensated labor, ruling out the concern that "propagation" is simply a relabeling of the linked households providing insurance. In addition, we use a second research design comparing those close to a shocked household to those close to a placebo household, which yields similar effects.

To understand the mechanisms of this propagation, we show that propagation occurs in a context of rigid/persistent networks: *ceteris paribus*, households that transacted at baseline are substantially more likely to transact even ten years later, relative to households that did not transact at baseline. Kinship ties are strong predictors of trade, highlighting the importance of time-invariant barriers to trade across households (Emerick, 2018). Thus, suppliers struggle to find new customers when their clients suffer a shock, and workers struggle to find new jobs when existing employers scale back demand. Frictions leading to rigid labor networks are particularly important: proximity to the shocked household via the labor market network is most strongly associated with indirect effects on income and consumption.[5] Indirect effects persist even four years after the shock, suggesting that, in a context of rigid networks, the recovery from indirect shocks can be sluggish.

To understand how the impacts of health shocks are related to incompleteness in insurance markets, we show that directly-shocked households receive incoming tranfers to partially buffer the shock; however, this interpersonal insurance is incomplete. We further show that the direct and indirect consequences of shocks vary with access to informal insurance. Shocks to households with limited access to informal insurance reduce costs and revenues by 29% and 22%, respectively. In contrast, for households with higher informal insurance participation, these decreases are almost fully attenuated. Thus, households that are not well-integrated into local informal insurance networks are most vulnerable. In turn, shocks to less-insured households appear to generate larger indirect effects, though these differences are imprecisely estimated.

A large share of firms across the world are small and family-operated (Beck et al., 2005), and thus exposed to shocks affecting household endowments. It is nonobvious whether propagation will occur similarly for these microenterprises as for multinational firms.[6] Local information, repeat interaction, and norms against opportunistic behavior could, in principle, mean that supply chains

[5]Evidence of frictional slack in goods and labor markets is also shown by Egger et al. (2019) in the context of rural Kenya.

[6]A literature in international trade studies the propagation of shocks through multinational production networks in the aftermath of natural disasters (Barrot and Sauvagnat, 2016; Carvalho et al., 2021), trade shocks (Tintelnot et al., 2018; Huneeus, 2019), and sectoral or regional shocks (Caliendo et al., 2017).

in village networks function more smoothly, with less propagation, than in networks composed of large firms. Or liquidity and informational constraints could bind more tightly in village networks, making propagation more severe. The welfare consequences of propagation may be very different in a context where firms are owned, not by diversified shareholders, but by households with low and un-hedged incomes. Our paper sheds light on this question.

2 Context and Data

2.1 Household data

The data used in this study come from the Townsend Thai Monthly Survey. The survey covers approximately 45 households per village, representing 42% percent of the village population. A baseline interview was conducted from July to August 1998, collecting information on the demographic and financial situation of the households as well as ecological data on the villages. The subsequent monthly updates began in September 1998.[7] The sample in this paper covers the period between September 1998 and December 2012. We focus our analysis on the subset of 509 households that responded to the interview throughout all survey waves.

Table 1 characterizes the sample households in terms of their demographic, financial and business characteristics. It shows that while households derive income mostly from family farms, they also operate off-farm businesses and provide labor to other households or businesses. In addition, 13% of total income comes from the receipt of government transfers and/or gifts from other households. Households allocate around 50% of their resources to consump-

[7]For more detail about the Townsend Thai Monthly Survey, see Samphantharak and Townsend (2010).

tion, and use the remainder to accumulate assets, which are evenly distributed between liquid and fixed assets. In terms of access to financial markets, in a given year, 83% of households report borrowing from any source, 48% from formal or quasi-formal financial institutions, and 30% from personal lenders, including relatives.

2.2 Network data

The survey contains detailed information on transactions among households and captures different types of economic inter-linkages. During each survey wave, interviewees identify any and all households in the village with whom they have conducted a given type of transaction.[8] We aggregate the monthly transaction data by year to elicit three types of village networks, for each year in the sample.

First, we recover the supply chain networks that capture transactions of output, inputs and intermediate goods across businesses of households in the same village. Second, we recover labor networks that capture employer-employee relationships within the village.[9] Appendix Figure B1 depicts both networks for a sample village. Panel C of Table 1 shows statistics on network participation across the sample as a whole. On average, just under half (48%) of the households transact in the local village supply-chain network by transacting inputs and final output (with 1.26 connections on average), and 62% provide or purchase labor to/from other households in the village, with 3.07

[8] The set of transactions include the relinquishment of assets, purchases or sales of inputs or final goods, the provision of paid and unpaid labor, and giving and receiving gifts and loans.

[9] As is typically the case in networks based on survey instead of census data, our networks may look thinner than the networks that would be elicited using census data (Chandrasekhar and Lewis, 2017). We discuss the implications of this source of bias for our research design in Section 3.2.

connections on average. We also recover information on local financial networks defined by gifts and loans across households in the same village, which tend to be more sparse (see Appendix Figure B1).

Households tend to participate in several networks in a given period. Among those linked through financial networks, over 60% also transact in local supply-chain networks and over 70% of them transact in local labor markets. Over 77% of households transacting in the village supply-chain networks also sell or purchase labor locally, and 45% are linked through local financial markets. Likewise, over 59% and 43% of households that buy or provide labor locally transact in the supply-chain and financial village networks, respectively. Local kinship networks also overlap with these transaction-based networks; see Section 4.4.

2.3 Constructing idiosyncratic shocks

In order to understand how shocks propagate to other households through village networks, we focus on idiosyncratic events associated with high levels of health spending, which correspond to periods of high financial stress. These shocks are well-suited for our analysis for several reasons. First, serious health shocks affect household finance and labor supply (Gertler and Gruber, 2002; Genoni, 2012; Hendren et al., 2018); the large magnitude of such shocks improves statistical power, and moreover such shocks are of *prima facie* importance. Second, because these shocks are uncorrelated across households (as shown below), we are able to separate the direct, idiosyncratic, effect from effects hitting other, connected, households via propagation. Additionally, the timing of these shocks are—as we show below—exogenous, allowing us to understand their causal effects.

We identify shocks as follows (see Appendix A for details). For each household, we identify the month with the highest level of health spending throughout the panel.[10] We focus on the largest shocks as they pose a significant financial burden to the household. To facilitate measuring responses to the shocks by comparing households' behavior before and after the episodes, we restrict the search to years 3-12 in the panel (out of 14 years). This enables us to observe at least two years of both pre- and post-shock behavior for all households. We identify 505 shocks, one per household. After excluding shocks related to childbirth, which may be anticipated, we are left with 469 shocks.[11]

2.4 Characteristics of the shocks.

Relationship between health spending and health status. Appendix Figure A4 shows that health spending (left axis) and self-reported symptoms (right axis) co-move, confirming that shocks are correlated with decreases in household health endowments. Appendix Table A1 reports the distribution of types of health symptoms reported by shocked households during the two years around the shock, during non-shock periods, and during all the sample periods.

Magnitude of the shock. The shocks represent a substantial financial burden to the households: on average, the highest level of monthly health spending within 6 months after the onset of the shock (THB 5314) accounted for 87%

[10]Thailand has a universal health insurance program, so these expenses are above and beyond those covered. See Appendix C.

[11]To account for potential anticipation effects, we define the beginning of each event by subtracting the number of months preceding the episode of high health spending during which household members reported health symptoms from the month corresponding to the episode. Appendix Figure A1 shows that, prior to the sudden increase in health spending, the median number of consecutive months in which households report any health symptoms is three. We present robustness checks varying the beginning of the shock in Section A.2.1.

of the average total monthly consumption during the 6 months preceding the shock (THB 6113) and was substantially larger than the average monthly total household food consumption during this period (THB 2817).

Are the shocks idiosyncratic? Our analysis requires that the events be idiosyncratic and their occurrence be uncorrelated with trends in household behavior. The top panel of Appendix Figure A3 presents the distribution of months associated with the beginning of each event. It shows that the event start dates are spread through all the periods in the sample and suggests that the events are indeed idiosyncratic. In Appendix Table A3 we formally show that village-level trends have null predictive power on the the occurrence of these events.

3 Direct and indirect effects of idiosyncratic shocks

3.1 Direct effects

To understand the indirect effects of shocks via network propagation, we first must understand how they affect the *directly* shocked household. Because the networks that we study are defined by cross-household transactions of input, output, and labor, our first stage analysis focuses on estimating the direct effects of shocks on business spending, labor demand, and production.

Estimating the effects of idiosyncratic shocks on household outcomes requires a valid comparison group. We would like to compare shocked households and otherwise-similar households who, by chance, were not simultaneously exposed to a shock. To implement this idea, we follow Fadlon and Nielsen (2019) and exploit the plausibly random variation in the *timing* of severe health shocks. We compare the behavior of households that experienced the shock in period t (i.e., treated households), to the behavior of households from the same age group and village that did not experience a shock at time t, but experienced a similar shock later on, in period $t + \Delta$ (control households). We denote as treated households those who experienced the shock during the first half of the panel; control households experienced a shock during the second half.

We use a difference-in-difference approach to compare changes in outcomes, before and after the shock, between treatment households (who experienced an actual shock) and control households (who experienced a placebo shock). (See Appendix A.2 for additional details.) The underlying identification assumption is that, in the absence of the shock, the treatment and control group would have followed parallel trends, which we validate tusing event-study specifications that test for lack of systematic differences before the shock ("parallel pre-trends").

3.1.1 Estimation

We estimate the following generalized difference-in-difference specification, following Fadlon and Nielsen (2019):

$$\begin{aligned} y_{i,t} = & \sum_{\tau=-4,\tau\neq-1}^{\tau=3} \beta_\tau \mathbf{I}[t=\tau] \times Treatment_i + \sum_{\tau=-4,\tau\neq-1}^{\tau=3} \theta_\tau \mathbf{I}[t=\tau] \\ & + X_{i,t}\kappa + \alpha_i + \delta_t + \epsilon_{i,t} \end{aligned} \tag{1}$$

where $y_{i,t}$ denotes the outcome for household i at t. Household- and month-fixed effects (α_i and δ_t) absorb time-invariant household characteristics and aggregate time-varying shocks. $Treatment_i$ is a time-invariant indicator of whether the households is in the treatment group. As each household is either observed in the treatment or comparison group, $Treatment_i$ is absorbed by the household fixed effects. Time to treatment is denoted by $\tau_{i,t}$ and is measured in half-years to increase precision. X is a vector of time-varying demographic characteristics including the number of male and female household members, age of the household head and maximum years of schooling in the household. The coefficients of interest are $\{\beta_\tau\}_{\tau=-4}^{\tau=3}$, which compare differences in changes in outcomes with respect to the period preceding the shock ($\tau = -1$) between households in the treatment and control group. We focus on a two-year (i.e., four-half year) time window before and after the shock as our panel is fully balanced during this period. We also use a more parsimonious differences-in-difference specification:

$$y_{i,t} = \beta Post_{i,t} \times Treatment_i + \theta Post_{i,t} + X_{i,t}\kappa + \alpha_i + \delta_t + \epsilon_{i,t} \tag{2}$$

where $Post_{i,t}$ is an indicator that takes the value of 1 in periods following the shock, and 0 otherwise. The parameter of interest, β, compares differences in outcomes before and after the shock, between households in the treatment group and the comparison group. In both specifications, we cluster standard errors at the household level as our main source of variation comes from cross-household variation in the timing of events and to flexibly account for serial correlation (Bertrand et al., 2004).

Note that our approach addresses two issues that may arise in simple event-study panel regressions without a stable comparison group—i.e., when researchers regress outcomes on time- and household-fixed effects and a post-shock dummy. A simple event-study approach would use all the households who do not experience a shock at period t as a control group for those that did; even those that were shocked before t. This could be problematic in our setting since such "staggered event timing" specifications may suffer from bias when effects are heterogeneous over time (Goodman-Bacon, 2018; Baker et al., 2021). Our design, by virtue of using a control group which is never treated before/during the comparison window, avoids these concerns. However, this advantage comes at the cost of statistical power and limits the number of available post-period observations as we only analyze the subset of 246 shocks that occurred earlier in the sample. Moreover, trends in outcomes may vary by age due to different trajectories along the life cycle. By constructing a comparison group within age group and village, our approach makes comparisons of households with similar pre-shock trends.

3.1.2 Direct effects: Results

A shock to health spending, which entails large outflows of resources, can be financed in a number of ways. Here we focus on changes in household production decisions — reducing spending on hired labor and/or business inputs to free up resources to meet the shock — as such dimensions are linked to cross-household transactions that determine local economic networks.[12]

Figure 1 reports flexible difference-in-difference estimates following the specification in equation (1). Panel (a) shows that, relative to control households, shocked households experience a large and significant increase in the probability of reporting health symptoms. Panel (b) shows that this coincides with a sharp increase in total health expenditure, and Panel (c) shows an in-

[12] *A priori*, shocks may affect household labor endowments as well as spending needs. In Appendix A.2.2 we argue that the spending effect is more first-order in our setting as a majority of shocks affect non-prime-aged individuals.

crease in total expenditure of a similar magnitude, indicating that non-health consumption remains steady.

The remaining panels show that the shocks affect the household's production-side decisions. Panel (d) shows that, compared to households in the control group, hired labor usage declines for shocked households. Panel (e) shows that input spending falls after the shock. Finally, Panel (f) shows that the slowdown in input spending coincides with a slowdown in revenues after the shocks. Note that the sharp declines in input spending and revenues coincide with the sharp increase in spending induced by the shock. Thus, shocked households meet short-term liquidity needs in part by drawing down working capital, inconsistent with the separation theorem. The graphical evidence also documents parallel pre-trends: for all six outcomes, there are no spuriously significant "effects" prior to the spending shock.[13]

To provide a quantitative assessment of the overall impact of the health shocks, in Table 2 we report difference-in-difference estimates of the effect of the shock on outcomes, corresponding to equation (2). Panel A examines household spending. Column 1 shows that the shock leads to a large increase in health spending. While this is by construction, the magnitude, approximately THB 540, is notable, representing a roughly 350% increase relative to the baseline mean.[14] Column 2 shows that during the two years following the shock, on average, total spending increases for shocked households, relative to control households, by approximately THB 620, an amount close to the effect on health spending. Thus, in terms of non-health spending, shocked households appear to fully buffer the shocks.

[13] Appendix Figure B2 shows the same dynamics in the raw data.

[14] Note that the effects of the shock on health spending are averaged across 24 months in this specification.

Buffering consumption may entail costly adjustments by shocked households as in Chetty and Looney (2006). Indeed, in order to buffer non-health consumption, affected households significantly decrease spending on business inputs (column 3) and reduce the use of external labor (column 4). Households also appear to reduce the use of labor provided by household members (column 5), though the effect is not significant. As a result of reduced investment in inputs and labor (columns 3-5), there is a decrease in the revenues from family enterprises, as seen in column 6. (The effect on revenues has a p-value of 0.107.) To increase precision, panel B reports results also including early-shocked households as controls for late-shocked households,[15] which nearly doubles the number of events. Reassuringly, the point estimates are very similar to those in panel A but are estimated with more precision.

Table B1 shows that the results are robust to using alternative definitions of the shock onset attenuating concerns about anticipation. In addition, the results are robust to randomly allocating the placebo shocks, and to using standard two-way fixed effect approaches to compute the effects of the shocks (see Section A.2.1 for details). In addition, Table B2 reports results based on the subset of shocks that affected older household members (age above the median age of 57). The responses to shocks hitting these non-prime-age adults are similar in magnitude to the response in the full sample, suggesting that the costly adjustments in response to the shocks are not solely due to reductions in household labor, but rather are driven by the expenditure shock.

In summary, in the face of a shock, households buffer non-health consumption, but do so at the cost of significantly reducing business spending. (House-

[15] Note that even after including more events, the treatment status of control households is held fixed around the 24-month analysis window around each event. This addresses potential biases in difference-in-differences frameworks that tend to arise when treatment status varies over time.

holds may also engage in other strategies to cope with the shocks; see Section 4.2.) These declines in business spending and labor demand have broader consequences for other households. We next turn to examining the effects of these shocks on *other* households, via propagation through local economic networks.

3.2 Economic networks and the propagation of idiosyncratic shocks

The results above show that health shocks meet the necessary criteria to understand propagation: their timing is exogenous, their occurrence is idiosyncratic, and the shocks have substantial effects on household production decisions. Given the significant degree of inter-linkage in the study villages, we next examine whether these shocks propagate to other households. We analyze two propagation channels. First, shocks could propagate through local supply chain networks: health shocks lead to decreases in the supply and demand for inputs, which could lead to reductions in sales and revenue for those households that trade with shocked households. Second, shocks could propagate through local labor networks: as supply and demand for outside labor decreases due to the shocks, households that exchange labor with shocked households could suffer falls in hours, earnings and revenue.

3.2.1 Identifying propagation effects

We exploit two sources of variation to test if idiosyncratic health shocks propagated to other agents in the local economy. First, we use variation in the timing of each household-level shock. Second, we use the fact that a household's exposure will depend on their network connections to the shocked household via the supply chain or labor network (or both). We assess the propagation of idiosyncratic shocks to other local family businesses by comparing households who, before the shock, shared closer market inter-linkages with household j's businesses to those who were un- or less-connected to household j before the shock, before and after the shock to household j.

Throughout our sample period, we observe multiple health shocks per village. We construct a dataset capturing information of non-shocked households before and after each health shock in the sample. For each event, we take two years of pre- and post-shock observations of households living in the same village of the directly shocked households.[16] We then stack the observations into a dataset at the household (i) by time (t) by event level (j), for each village.

We combine this dataset with information on network connections between the shocked household (j) and other households (i) in the village, measured during the year preceding the shock to household j. We use pre-shock networks as links may respond to economic shocks themselves. The assumption is that households that transacted with the shocked household during the pre-period, on average, would have been more likely to transact with the shocked households in the post-period, in the absence of the shock. This is consistent with the evidence of persistence in the village networks discussed in Section 4.4, and with evidence of the importance of time-invariant determinants of economic connections such as kinship relations (Kinnan and Townsend, 2012), race or caste (Munshi, 2014), and the existence of economic frictions such as contracting issues that may limit trade between households (Ahlin and Townsend, 2007), or between firms (Aaronson et al., 2004) in local economic networks.

We measure exposure as the inverse distance in the undirected village net-

[16] We restrict the analysis to two years before and after the shock, first, to be consistent with the analysis of the direct effects of the shocks; second, we only have a fully balanced panel during this time window.

work[17] between household i and the shocked household j: $Closeness_{i,j} = \frac{1}{dist_{i,j}}$.[18] As households are further away in the network from shocked households, exposure (closeness) decreases. We begin by computing overall closeness based on transactions in the supply chain or labor networks as households can be exposed through either network. To distinguish between exposure in the supply chain and labor market networks, we also compute measures of closeness in each separate network (see Section 4).

We elicit economic networks using survey instead of census data (Chandrasekhar and Lewis, 2017). Thus, it is possible that we underestimate the closeness of some sample households to shocked households.[19] Because we may be underestimating exposure—classifying some households as un- or less-exposed when they are actually (more) exposed—our results could be biased towards zero. Thus, we interpret our magnitudes as *lower bounds* of the indirect effects of idiosyncratic shocks on other households.

Finally, not all shocked households are active in the local markets for goods, and not all shocked households employ other villagers for their businesses. Thus, we analyze the propagation of shocks through village networks by focusing only on events corresponding to the 391 households that traded in either the supply chain or labor market networks during the year preceding their shock; these represents 83% of all the shocks in our sample.

With these caveats in mind, we estimate the following difference-in-difference

[17]We focus on undirected networks because the shock can propagate both up- and downstream as we document below, in Section 3.2.2.

[18]This measure equals one if household i directly traded with the shocked household j and zero if household i does not have any direct or indirect connections with the shocked household. The geodesic distance between two unconnected nodes is $dist_{i,j} = \infty$ and so their closeness equals 0 in that case. By undirected networks we mean that we do not distinguish between incoming vs. outgoing transactions. Likewise, we weight each transaction equally for our calculations.

[19]See footnote 9 for a discussion of this issue.

specifications:

$$
\begin{aligned}
y_{i,t,j} = & \sum_{\tau=-4,\tau\neq-1}^{\tau=4} \beta_\tau \mathbf{I}[t=\tau] \times Closeness_{i,j} + \gamma Closeness_{i,j} + \mathbf{X}_{i,t,j}\kappa \\
& + \alpha_i + \omega_j + \delta_t + \theta_{\tau(j)} + \delta_t \times Degree_{i,j} + \epsilon_{i,t,j} \qquad (3) \\
y_{i,t,j} = & \beta Post_{t,j} \times Closeness_{i,j} + \gamma Closeness_{i,j} + \mathbf{X}_{i,t,j}\kappa \\
& + \alpha_i + \omega_j + \delta_t + \theta_{\tau(j)} + \delta_t \times Degree_{i,j} + \epsilon_{i,t,j} \qquad (4)
\end{aligned}
$$

where y denotes the outcome of interest for household i in village v at time t around the shock suffered by household j. In the "event-study" specification (equation (3)), τ denotes a half-year, which may precede ($\tau < 0$) or follow ($\tau >= 0$) the shock to household j. $Closeness_{i,j}$ denotes inverse distance to the shocked household during the year preceding the shock to j.[20] The coefficients of interest in equation (3) are β_τ, which capture relative changes in outcomes corresponding to half-year τ with respect to the half-year preceding the event ($\tau = -1$) associated with one additional unit of closeness (i.e., between more- vs. less-exposed households). In the generalized difference-in-difference specification, equation (4), $Post_{t,j}$ takes the value of one during the two years following the shock to household j, and zero for the pre-period. The coefficient of interest, β, captures differences in outcomes with respect to pre-period, associated with one additional unit of closeness.

We control for household fixed effects (α_i), time (month) fixed effects (δ_t) shocked-household fixed effects (ω_j), time-to-shock fixed effects ($\theta_{\tau(j)}$), which accounts for village-specific time-varying shocks during the analysis window

[20] Below we consider several definitions of *Closeness*: proximity in the overall network pooling supply chain and labor market, as well as proximity in one network or the other.

corresponding to the shock to household j, and a vector of time-varying demographic characteristics ($\mathbf{X}_{i,t,j}$).[21] We also control for time-varying shocks affecting more central households, which could also be more likely to be close to other households, by including interactions of the number of links of household i ($Degree_{i,j}$) during the year preceding the shock to j with time fixed effects. We use two-way clustered standard errors at the event level j and household level i to allow for flexible correlation across households during the periods preceding and following event j and across responses of the same household i to different events. As we are focusing on indirect effects, we drop observations of directly shocked households j from the analysis. We also exclude observations of households that experienced their own shock within a year before and after the shock to household j.

The identifying assumption underlying our strategy of estimating indirect effects is that, in the absence of the shock to household j, the outcomes of households i and i', with differential closeness to j, would have evolved along parallel trends *ceteris paribus*, i.e., conditional on the vector of controls included in equation 3 and 4. We validate this identifying assumption by testing for a lack of differences in the pre-period; namely, for $\tau < 0$, we verify that β_τ is not different from zero.

In thinking about the identifying assumption, recall that equations 3 and 4 control for household fixed effects; shocked-household fixed effects; and $Degree_{i,j} \times month$ fixed effects, which allow for a common shock to all households with a given network degree to experience a common shock. Thus, we are in essence comparing two households equally well-connected to the network, one of whom happens to be closer to the shocked household.

[21] We control for household size, gender composition, average age and schooling.

3.2.2 Results: Propagation of shocks through economic networks

Figure 2 presents flexible difference-in-difference estimates following equation 3. Panel A analyzes total transactions. After a health shock, households who are more connected to shocked households differentially reduce the number of transactions with other households in the village. Prior to the shock, transactions are not different for closer vs. more-distant households. After the shock, however, transactions decline more for households who are closer to the shocked household. Panels B and C show that supply-chain and labor network transactions, respectively, each exhibit the same pattern seen for total transactions. Panel D shows that, as local networks are shocked, total income declines for households closer to the shocked household. In all four cases, the pre-shock period shows no evidence of differential pre-trends. Finally, Panel E shows an analogous result for total consumption expenditure, which declines in the post-shock period (and exhibits no differential trend in the pre-period).

The effects on transactions, income and spending are evident in all three half-year periods following the shock and do not appear to shrink in magnitude over time: the effects are quite persistent. In theory, indirectly-hit households might attenuate these effects over time by finding new local trading partners. However, the evidence on the rigidity of local networks shown below (section 4.4) demonstrates that such reorganization of local ties is very difficult, at least over the span of 1-2 years.

Table 3 shows difference-in-difference estimates corresponding to equation (4).[22] It documents significant post-shock declines in the number of monthly transactions in the supply-chain (column 1) and labor-market networks (col-

[22]In Appendix table B3, we re-estimate equation 3, including village-by-month fixed-effects ($\upsilon_{v,t}$) to control for potential village-and-time-specific shocks. The results are quite similar to those from the main specification which control for village-time fixed effects in the analysis window corresponding to each event.

umn 2), and in total transactions (column 3). These effects are large, representing declines of 20%, 24% and 21% relative to the pre-period means, respectively. Column 4 shows that these changes, in turn, reduce income: a one-unit increase in *Closeness* is associated with a fall in income of THB 1267, or 12% of the pre-period mean. In turn, consumption spending falls by THB 304, or 4.2% of its pre-period mean (column 5).[23] The fall in consumption is smaller than the fall in income, suggesting that indirectly shocked households are able to partly, but not completely, smooth their indirect shock exposure.

The effects that we observe are strongest for directly connected households—those that were one link away from the shocked households—but affect indirectly connected households as well (see Appendix Figure B3).[24] When, due to a shock, those linked directly to shocked households reduce sales of goods or labor (outgoing transactions to the shocked household), this leads to declines in income, which in turn translate into fewer purchases (incoming transactions) from other households, triggering further propagation through the network. Indeed, Table 4 shows that the fall in outgoing transactions documented above is matched by a fall in incoming transactions (input purchases and labor hiring). In sum, the health shocks that we study generate indirect effects both upstream and downstream, as the costly adjustments taken by the directly shocked household reverberate through local networks. Shocks that

[23]Recall that these are the effects associated with moving from $Closeness = 0$ (being unconnected to the directly shocked household) to $Closeness = 1$ (being directly linked). The mean level of $Closeness = 0.42$, so that the average indirect effect is 42% of the coefficient.

[24]Figure B3 plots indirect effects decomposing the measure of closeness into 4 categories: directly connected households (1 link away from the shocked households); households that are 2 or 3 links away from the shocked households; those that are 4 or 5 links away from the shocked households; and (as the base category) those that are 6 or more links away in the network, including those that are unconnected to the shocked household. Although the effects dissipate through the network, there are non-negligible propagation effects on indirectly connected households.

are *prima facie* idiosyncratic are spread to other connected households. We return to the multiplier effect of these idiosyncratic shocks in Section 5.

3.2.3 Measuring indirect effects à la Fadlon and Nielsen (2019)

A potential concern with the first approach to measuring indirect effects is that we are comparing households who are closer vs. farther from the shocked household and, *a priori*, those with different network positions may be different. (Though recall that we are flexibly controlling for $Degree_{i,j} \times month$ fixed effects and that both groups exhibit parallel pre-trends.) An alternative approach, in the spirit of the design used to study direct effects, is to compare households that are close to a household (j) that experienced a shock in period t to households that were also close to a placebo household (j'): one whose shock occurs later in the data. In this design, both treatment and comparison households are similarly close to a shocked household but treated households are exposed to the shock during the analysis window while control households experience a placebo shock. Details are in Appendix A.3.

The results appear in Table 5. Column 1 reveals a drop in input/output transactions of 0.214, very close to the estimate of 0.200 from table 3. The effect on hired labor (column 2) is imprecisely estimated, but the effect on total transactions (column 3) of -0.278 is quite similar to the -0.315 from table 3. The effects on income and consumption, THB -1426.3 and -351, respectively, are also quite close to the estimates from Table 3 (THB -1267.1 and -303.6). The similarity of the two sets of results, using different designs for identifying indirect effects, serves as a sort of over-identification test, suggesting that both identifying assumptions are valid.

4 Propagation Mechanisms

4.1 Propagation via supply chain vs. labor networks

In Table 6, we examine whether the effect of exposure through the supply chain network has different effects than exposure through the labor market network. If proximity through the supply chain (labor) network is associated with changes in input/output (hired labor) transactions, and not vice versa, this is supportive of the identification assumption, as many plausible confounds (e.g., differential trends between closer vs. more distant households) would manifest in both sets of outcomes. Because the two networks are correlated, we analyze the effect of exposure to one controlling for the effect of the other.[25] Column 1 shows that, conditional on proximity in the labor market network, a 1-unit increase in proximity in the supply chain network is associated with a significant fall in input/output transactions of 0.227. There is no effect on input/output transactions associated with proximity through the labor network. Analogously, column 2 shows that proximity through the labor market network has a negative and significant effect (-0.210) on transactions involving paid labor, while there is no effect seen via the supply chain network. In column 3, proximity via the supply chain network and the labor market network both have negative and significant effects on the total number of transactions (-0.206 and -0.244, respectively).

Columns 4 and 5 show that proximity via the labor market network is associated with large and significant drops in income and consumption, re-

[25] On average, 41% of households share a direct or indirect link to the shocked households through both, supply-chain and labor-market network, 16% are directly or indirectly linked to the shocked household only through the supply-chain network, 13% are directly or indirectly connected to the shocked households only through the labor network and 30% of households are neither connected to the shocked households through the supply-chain nor labor network.

spectively, while the corresponding effects of proximity via the supply chain network are small and insignificant.[26]

4.2 Direct and indirect coping mechanisms

What, if any, coping mechanisms do households use when hit by the direct or indirect effects of health shocks? Appendix Table B4 examines the response of gifts, borrowing, fixed and liquid assets, and incoming unpaid labor. In principle, all of these mechanisms may be helpful in smoothing shocks, but it is an empirical question to what extent they are actually used.

Panel A presents results from direct shocks, corresponding to equation (2). Column 1 shows that incoming gifts increase by THB 570, or approximately 29%.[27] Columns 2 to 4 show that although borrowing increases and fixed and liquid assets decline, the changes are not significant.[28] Finally, column 5 shows that there is no response in terms of the amount of incoming unpaid labor. This is important as it demonstrates that the reductions in paid labor documented above are not reflections of a substitution to unpaid labor. Panel B presents

[26]A possible explanation is that, although the absolute effect of propagation through supply chain networks on input/output transactions is similar to the propagation effect through labor market networks on labor transactions, the effect on labor market transactions is larger in relative terms: the decline in labor market transactions represents a 44.6% decline relative to the pre-period number of labor market transactions, while the decline in transactions of inputs and final goods represents 23% of the pre-period mean. Households close to the shocked household via the labor market network may suffer a double impact, namely, reduced labor demand via an income effect as the directly hit household scales back, as well as a further hit due to the complementarity between household and hired labor, as household labor can be required to supervise or monitor hired labor. Indeed, Table 2 shows that household labor seems to decline due to the direct effect of the shocks (see column 5).

[27]Note that this is on the same order as the direct effect on health spending in Table 2; however, comparing Figure 1, Panel c and Figure B4 shows that the *timing* of gifts does not match that of health spending; with gifts in the half-year of the shock meeting less than half of the roughly THB 2000 of spending needs in that half-year.

[28]Health spending needs emerge suddenly and so arranging for loans or asset sales may take too long; alternatively households may desire to preserve these financing options as last-resort buffer stocks and so finance the shock out of business investment instead.

results from indirect exposure to shocks, corresponding to equation 4. There are no significant effects associated with indirect shock exposure on any of the five mechanisms. This helps to explain why consumption falls for indirectly shocked households—other coping mechanisms appear to be unavailable.

Why do directly shocked households see economically and statistically significant increases in transfers, while indirectly shocked households do not? First note that, in addition to receiving transfers, directly shocked households take other costly steps to buffer consumption, namely scaling back on business activities. Two other factors may help explain the divergence in transfer behavior. First, the direct shocks are large increases in health spending, often associated with changes in health symptoms. These shocks are salient and relatively observable. The indirect shocks, on the other hand, arise from reductions in supply and demand facing household businesses. Such shocks are likely less salient and potentially more subject to concerns of effort and verifiability, hence potentially less insurable. Moreover, because the indirect shock, by its nature, affects many interlinked households, the shock becomes *de facto* aggregate, which makes the potential for insurance via gifts from other villagers more limited.[29]

4.3 Informal insurance and the propagation of shocks

Informal insurance can help to buffer health shocks (De Weerdt and Dercon, 2006); this suggests that shocks to uninsured households may be more likely to trigger declines in business activities and hence propagate more to other households. To test this idea, we use data on intra-village provision and receipt

[29] To demonstrate that local networks may be less able to insure aggregate shocks, Figure B5 compares the network responses to idiosyncratic vs. aggregate shocks using the 2002 EU ban on Thai shrimp imports. See appendix B.1.1 for details.

of gifts during the year preceding the shock. We split the sample of shocked households into those with high vs low pre-shock access to informal insurance. See Appendix B.2 for details.

Panel A in Appendix Table B5 reports estimates of the direct effects of the shock on gift and loan receipt and business outcomes by access to informal insurance. Households with high access to informal insurance experienced a substantial increase in gifts and loans.[30] This increase, in contrast, is small and non-significant in the case of low-access households (column 2). Moreover, there are statistically significant declines in input spending (column 3) and hired labor (column 4) in the case of low-access households, but these declines are small and not significant in the case of better-insured households. For input spending, the difference between the effects on low- and high-access households is significant at 10% (p- value=0.09). In addition, although there is a significant decline in revenues in the case of high-access households (column 5), the decline in revenues in the case of low-access households is 1.7 times that of high-access households. The results suggest that households with limited access to insurance drive most of the declines in business activities, suggesting that incompleteness in local insurance markets may lead to non-separability of household spending and production decisions. Conversely, improvements in access to risk smoothing may reduce the extent of non-separability and thus reduce propagation.

Next, to investigate whether shocks to less-insured households propagate differently, we estimate a version of equation 3 where we allow the effect of indirect exposure to vary by the *directly shocked* household's baseline access

[30] We test whether the the indirect effect on consumption (Table 3, column 6) could be a consequence of a decline in cash on hand/liquidity arising from helping the directly shocked household. Appendix Table B7 shows that neither transfers nor loans given by the indirectly shocked household to other households increase following the shock.

to insurance (See Appendix B.2.) Panel B of Appendix Table B5 presents the results. When the shocked household had low access to insurance in the pre-period, the fall in income associated with 1 unit greater *Closeness* is 1705 baht. When the shocked household had high access to insurance, the fall in income is 1016.3 baht. That is, the propagation effects on income when the shocked household has low access to informal insurance are 1.67 times the propagation effects on income when the shocked household has higher access to insurance (col 4). Moreover, the consumption spending of indirectly affected households falls by 462 baht, or roughly 6%, when the shocked household had low access to insurance in the pre-period. When the shocked household had high access to insurance, however, the fall in consumption spending is reduced by only 275 baht (col 5). In sum, although the differences across shocks to households with high and low access to informal insurance cannot be estimated with precision, the magnitude of these differences suggest that informal insurance markets may mitigate the direct and indirect effects of idiosyncratic shocks.

4.4 Rigidities in local networks

Our results suggest that the costs related to the creation of new links may limit the participation of some households in local networks. Indeed, Panel C of Table 1 shows that while an significant share of households transact in village networks, this share is not 1. There is evidence from other contexts suggesting that market frictions may prevent transactions across businesses.[31]

[31] E.g., Johnson et al. (2002) finds that adequate institutional infrastructure (e.g., well-functioning courts) encouraged business owners to try new suppliers in post-Communist countries. Ahlin and Townsend (2007) provide evidence highlighting the importance of social ties and local sanctions in the context of joint-liability loans, for which commitment is crucial. Other sources of frictions may include product specificity (Barrot and Sauvagnat, 2016) or relationship specificity (Elliott, 2015).

If these frictions are empirically important, one should observe a large degree of persistence in the economic networks. To test for rigidities in the local networks, we construct a dyadic dataset including indicators of whether each pair of sample households (dyads) transacted in year t either in the local goods, labor or financial market and estimate the extent to which past transactions predict future transactions, conditional on measures of similarity. (See Appendix B.3 for details.) Table B8 presents the results. The labor-market and supply chain networks exhibit a striking degree of rigidity over time. One implication is that the effects of shocks which propagate via these networks may be quite persistent. Figure B6 reports event-study estimates of equation 3 over a larger post-period time span of 4 years (8 half-years). It suggests that the network disruptions induced by the shocks are persistent in both supply chain (panel a) and labor market networks (panel b), showing no evidence of dissipating even 4 years post-shock.

4.5 Village-level determinants of propagation

The results presented so far exploit *individual-level* variation in exposure to a given shock, between households in a village. It is also of interest to know how *village-level* variations such as differences in in network structure or the position of the directly shocked individual matter for the extent of propagation. Therefore we estimate two alternative specifications which, instead of exploiting variation in closeness to the shocked household within the village, exploit cross-village exposure to the shocks. First, we use the number of links that the shocked household had in the pre-period network—i.e., the number of households in the village that were exposed to the shock—as a measure of exposure. This specification sheds light on how a shock propagates when the shocked household is more vs. less connected.

Second, to understand how the overall village-level structure matters for the extent of propagation, we exploit cross-village variation in network density (based on pre-shock transactions). We expect that, the denser the network (the more interconnected households are in a village), the higher is the potential for propagation. Thus, we compare changes in outcomes before and after each health shock, between households in more- and less-exposed villages.

To understand the role of these village-level determinants we estimate the following model:

$$\begin{aligned} y_{i,t,j} =& \beta Post_{t,j} \times \text{Network Exposure}_j + \alpha_i + \omega_j + \delta_t \\ &+ \theta_{\tau(j)} + \mathbf{X}_{i,t,j}\kappa + \epsilon_{i,t,j} \end{aligned} \tag{5}$$

where, as in our main approach to estimate propagation effects, the unit of observation is a household i in period t around the shock to household j. As only one household was directly hit at a time, ω_j absorbs village-level variables that are invariant around the analysis window. Network Exposure$_j$ denotes exposure based on either the shocked household's degree or on network density during the pre-period.[32] The vector X includes the interaction of the number of households in the village (number of nodes in the network) with $Post_{t,j}$ to account for potential contemporaneous shocks correlated with village size. The results are reported in Panels A and B of Table 7. Panel A indicates that a 1 standard deviation (SD) increase in the degree of the shocked household leads to an average of 0.026 (2.5%) fewer input/output transactions in the post-shock period relative to pre-shock, 0.044 (9.3%) fewer labor market transactions, and 0.0695 (4.7%) fewer overall transactions (columns 1-3; all significant at 5% or

[32] In both cases, we use standardized z-scores to obtain results on a common scale.

better). Accordingly, a 1 SD increase in the degree of the shocked household leads to a differential fall in income of 364 THB (3.4%), significant at 1% (column 4). In column 5, the point estimate indicates a differential fall in consumption of 72 THB, (0.99%), however this is not different from zero at conventional levels (the p-value is .114).

Panel B shows the results for variation by degree density. A 1 SD increase in the network's pre-shock degree density is causes a 0.043 (4.2%) fewer input/output transactions in the post-shock period, 0.041 (8.6%) fewer labor market transactions, and 0.083 (5.6%) fewer overall transactions (columns 1-3; all significant at the 1% level). The corresponding effect on income is a fall of 391 THB (3.7%), and the effect on consumption is a decline of 148 THB or 2% (columns 4 and 5). All of the results in Panel B are significant at 1%.

These results show that both the structure of the network and location within the network matter: denser networks lead to greater propagation as do shocks to more-connected households.

5 Putting the findings in context: The multiplier effect of idiosyncratic shocks

In this section, we perform a simple back-of-the-envelope exercise to estimate the total magnitude of indirect vs. direct effects on revenues and use this exercise to benchmark our results.

As documented above, idiosyncratic health shocks have both direct costs (analyzed in section 3) and indirect costs (analyzed in sections 3.2 and 4). The former are larger on a per-household basis, but the latter can potentially affect many more households. In order to compare their overall magnitude, and so obtain an estimate of the overall "multiplier effect' of the fall in spending associated with the shock', we perform a simple calculation of the total indirect fall in consumption for each baht of reduced business spending by directly affected households.

The indirect effect on consumption associated with a 1 unit change in *Closeness*, from Table 3, Panel A, column 5 is a fall of -303.5 baht (significant at 10%). The mean (median) level of *Closeness* in the village network is 0.42 (0.43) and the mean (median) number of indirectly exposed households (i.e., households who are connected to the shocked household via the network) is 21.23 (16).[33] The implied total indirect effect using mean values is therefore $-303.5 \times 0.42 \times 21.23 = -2706$ baht per month. Using median values instead gives an implied total indirect effect of -2088.

From Table 2, Panel A, column 3, the fall in business costs for a directly affected household is -1757.4 baht, so the indirect effects using mean and median closeness represent multiplier effects of 1.54 and 1.19, respectively. For comparison, Egger et al. (2019) estimate a consumption-expenditure multiplier of 1.7 from cash transfers in Kenya, while in the US, Nakamura and Steinsson (2014) estimate an "open economy relative multiplier" of 1.5, Suárez Serrato and Wingender (2016) estimate a local income multiplier of government spending of 1.7 to 2, and Barrot and Sauvagnat (2016) find that $1 of lost sales at the supplier level leads to $2.40 of lost sales at the customer level. Note that a key distinction with other studies is that we exploit within-village variation in exposure to shocks based on distance to the shocked household in the village network. Thus, our estimates of indirect effects are net of any

[33] We report medians as well as means since the median is less sensitive to networks with a high number of connections or many distant (low-*Closeness*) connections, where the linear specification for *Closeness* may be less appropriate.

changes in prices (which would not be differential between closer and less-close households) and as such our multiplier estimate may be a lower bound; this is consistent with our estimate being at the lower end of the range of other recent estimates. While our multiplier estimates are admittedly back-of-the-envelope, they demonstrate that, because the indirect effects are economically meaningful and affect many households for each directly affected household, the total indirect effects are of a similar order of magnitude, and perhaps larger than, the direct effect itself.

6 Concluding remarks

Local networks are well understood to serve a consumption smoothing purpose. We document that they also serve another role, propagating idiosyncratic shocks. We leverage variation in the timing of health expenditure shocks to document consumption smoothing, i.e., no impacts on non-health consumption expenditure, for directly shocked households. However, these shocks are only partially insured via gifts and loans and, as a result, shocked households adjust their production decisions—drawing down working capital, cutting input spending, and reducing labor hiring—in order to achieve consumption smoothing. These adjustments propagate the shock to other households through interlinkages in local supply-chain and labor networks. Businesses close to the shocked household in the supply chain network experience reduced local sales, and workers closer to the shocked household in the labor network experience declines in the probability of working locally and reduced earnings. As a result, consumption falls for these indirectly shocked households. The indirect effects imply a consumption multiplier of approximately 1.5. We also provide suggestive evidence that both direct and indirect effects are attenuated when shocked households are better insured through local risk-sharing networks.

Our findings suggest that (at least) two sets of interventions might be beneficial. First, improved safety nets may help prevent granular shocks from propagating to become *de facto* aggregate. Given that the ability to share idiosyncratic shocks increases with the number of households participating in the insurance network, local networks alone may be unable to diversify severe idiosyncratic risk. Formal commercial insurance contracts or social insurance could allow better risk-coping and thus reduce propagation. In addition, electronic payment platforms that ease the flow of resources across villages may expand and strengthen informal risk-sharing networks (Jack and Suri, 2014).

Of course, fully insuring all idiosyncratic shocks is infeasible. This suggests a need for policy interventions to make production networks less rigid and more diversified. These may include interventions to improve contract enforcement (Fazio et al., 2020) or to broaden the extent of product and factor markets beyond the local village market (Park et al., 2021). Such investments may reduce the rigidity and sparsity of supply chain and labor networks and hence mitigate propagation.

Shocks to entrepreneurs of the type studied in this paper are widespread in both high- and low-income countries. Moreover, the coronavirus pandemic, while to a large extent an aggregate shock, has significant idiosyncratic aspects due to variation in household infection risks and realizations (Jordan et al., 2020), the ability to work from home (Angelucci et al., 2020), and the extent to which different workers and sectors of the economy are affected by shutdowns and social distancing (Daly et al., 2020). Understanding the propagation of shocks is crucial for designing policies to mitigate future crises.

References

Aaronson, D., R. W. Bostic, P. Huck, and R. Townsend (2004). Supplier relationships and small business use of trade credit. *Journal of Urban Economics 55*(1), 46 – 67.

Ahlin, C. and R. M. Townsend (2007). Using repayment data to test across models of joint liability lending. *The Economic Journal 117*(517), F11–F51.

Angelucci, M., M. Angrisani, D. M. Bennett, A. Kapteyn, and S. G. Schaner (2020). Remote work and the heterogeneous impact of covid-19 on employment and health. Technical report, National Bureau of Economic Research.

Baker, A., D. F. Larcker, and C. C. Wang (2021). How much should we trust staggered difference-in-differences estimates? *Available at SSRN 3794018*.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2013). The diffusion of microfinance. *Science 341*(6144).

Barrot, J.-N. and J. Sauvagnat (2016). Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks. *The Quarterly Journal of Economics 131*(3), 1543–1592.

Beaman, L., A. BenYishay, J. Magruder, and A. M. Mobarak (2021). Can network theory-based targeting increase technology adoption? *American Economic Review 111*(6).

Beck, T., A. Demirguc-Kunt, and R. Levine (2005). SMEs, growth, and poverty: Cross-country evidence. *Journal of Economic Growth 10*(3).

Benjamin, D. (1992). Household composition, labor markets, and labor demand: Testing for separation in agricultural household models. *Econometrica 60*(2), 287–322.

Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *The Quarterly journal of economics 119*(1), 249–275.

Bigio, S. and J. La'o (2020). Distortions in production networks. *The Quarterly Journal of Economics 135*(4).

Caliendo, L., F. Parro, E. Rossi-Hansberg, and P.-D. Sarte (2017, 12). The Impact of Regional and Sectoral Productivity Changes on the U.S. Economy. *The Review of Economic Studies 85*(4), 2042–2096.

Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi (2021). Supply chain disruptions: Evidence from the great east japan earthquake. *The Quarterly Journal of Economics 136*(2).

Chandrasekhar, A. and R. Lewis (2017). The econometrics of sampled networks. Working paper.

Chetty, R. and A. Looney (2006). Consumption smoothing and the welfare consequences of social insurance in developing economies. *Journal of public economics 90*(12), 2351–2356.

Daly, M. C., S. R. Buckman, and L. M. Seitelman (2020). The unequal impact of covid-19: Why education matters. *FRBSF Economic Letter 2020*, 17.

de Janvry, A., K. Emerick, E. Kelley, and E. Sadoulet (2019). Endogenous information sharing and the gains from using network information to maximize technology adoption.

De Weerdt, J. and S. Dercon (2006). Risk-sharing networks and insurance against illness. *Journal of development Economics 81*(2), 337–356.

Egger, D., J. Haushofer, E. Miguel, P. Niehaus, and M. Walker (2019). General equilibrium effects of cash transfers: experimental evidence from Kenya. *NBER Working Paper No. 26600*.

Elliott, M. (2015). Inefficiencies in networked markets. *American Economic Journal: Microeconomics 7*(4).

Elliott, M., B. Golub, and M. O. Jackson (2014). Financial networks and contagion. *American Economic Review 104*(10).

Emerick, K. (2018). Trading frictions in indian village economies. *Journal of Development Economics 132*, 32 – 56.

Fadlon, I. and T. H. Nielsen (2019, September). Family health behaviors. *American Economic Review 109*(9), 3162–91.

Fazio, D., T. Silva, J. Skrastins, et al. (2020). Economic resilience: spillovers, courts, and vertical integration. Technical report.

Foster, A. D. and M. R. Rosenzweig (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of political Economy 103*(6).

Genoni, M. E. (2012). Health shocks and consumption smoothing: Evidence from indonesia. *Economic Development and Cultural Change 60*(3).

Gertler, P. and J. Gruber (2002). Insuring consumption against illness. *The American Economic Review 92*(1), 51–70.

Giannone, E. and C. Banternghansa (2018). The Large Consequences and the Spillover Effects of a Small Shock. 2018 Meeting Papers 986, Society for Economic Dynamics.

Goodman-Bacon, A. (2018). Difference-in-differences with variation in treatment timing. Technical report, National Bureau of Economic Research.

Hendren, N., A. Shenoy, and R. Townsend (2018). Household responses to negative health shocks in thailand. Mimeo.

Huneeus, F. (2019). Production network dynamics and the propagation of shocks. Working paper.

Jack, W. and T. Suri (2014). Risk sharing and transactions costs: Evidence from Kenya's mobile money revolution. *American Economic Review 104*(1).

Johnson, S., J. McMillan, and C. Woodruff (2002, April). Courts and Relational Contracts. *Journal of Law, Economics, and Organization 18*(1).

Jordan, R. E., P. Adab, and K. Cheng (2020). Covid-19: risk factors for severe disease and death.

Kinnan, C. and R. Townsend (2012). Kinship and financial networks, formal financial access, and risk reduction. *American Economic Review 102*(3).

Maccini, S. and D. Yang (2009). Under the weather: Health, schooling, and economic consequences of early-life rainfall. *American Economic Review 99*(3).

Meghir, C., A. M. Mobarak, C. D. Mommaerts, and M. Morten (forthcoming). Migration and informal insurance: Evidence from a randomized controlled trial and a structural model. *Review of Economic Studies*.

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

Munshi, K. (2014, November). Community networks and the process of development. *Journal of Economic Perspectives 28*(4), 49–76.

Nakamura, E. and J. Steinsson (2014). Fiscal stimulus in a monetary union: Evidence from US regions. *American Economic Review 104*(3), 753–92.

Park, S., Z. Yuan, and H. Zhang (2021). Technology adoption and quality upgrading in agricultural supply chains: A field experiment in Vietnam.

Rosenzweig, M. R. and H. P. Binswanger (1993). Wealth, weather risk and the composition and profitability of agricultural investments. *The Economic Journal 103*(416), 56–78.

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

Samphantharak, K. and R. M. Townsend (2010, December). *Households as Corporate Firms*. Cambridge University Press.

Samphantharak, K. and R. M. Townsend (2018, February). Risk and return in village economies. *American Economic Journal: Microeconomics 10*(1).

Silva, D. and R. Townsend (2019). Risk-taking over the life cycle: Aggregate and distributive implications of entrepreneurial risk. Mimeo.

Suárez Serrato, J. C. and P. Wingender (2016, July). Estimating local fiscal multipliers. Working Paper 22425, National Bureau of Economic Research.

Tintelnot, F., A. K. Kikkawa, M. Mogstad, and E. Dhyne (2018, October). Trade and domestic production networks. WP 25120, NBER.

Townsend, R. M. (1994). Risk and insurance in village india. *Econometrica 62*.

Udry, C. (1994). Risk and insurance in a rural credit market: An empirical investigation in northern nigeria. *The Review of Economic Studies 61*(3).

Wagstaff, A. (2007). The economic consequences of health shocks: evidence from vietnam. *Journal of health economics 26*(1).

World Bank (2004). *World development report 2005: a better investment climate for everyone*. The World Bank.

Figures

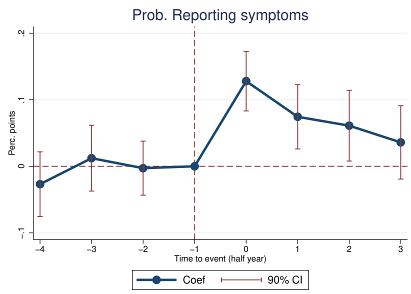

(a) Prob. of reporting any symptoms

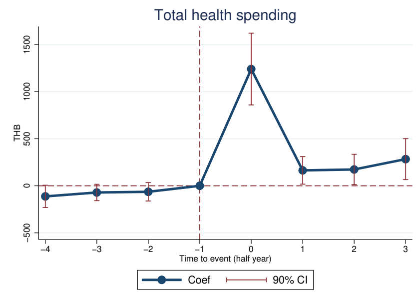

(b) Total health spending

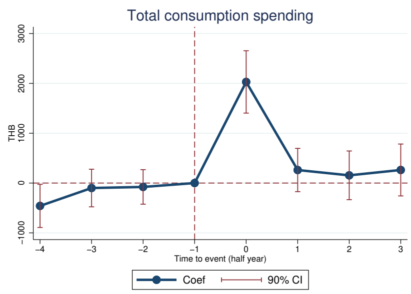

(c) Total consumption spending

(d) Hired labor (hrs/month)

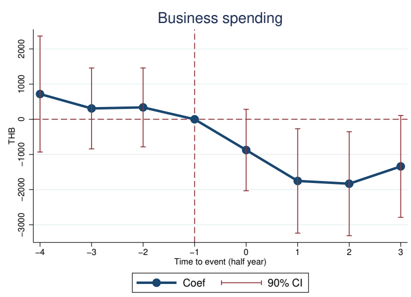

(e) Business spending (inputs)

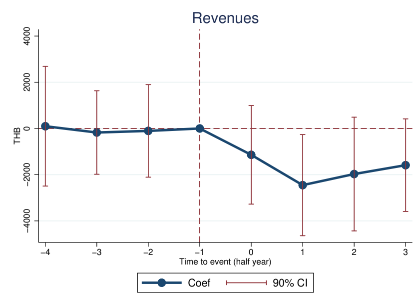

(f) Revenues

Figure 1: Direct effects of health shocks

Note: Each dot represents differences between treatment and placebo households in changes in outcomes relative to the period preceding the beginning of the shock ($\tau = -1$).The estimating sample includes 2 years before and after the shock divided in half-year bins. All specifications control for household time-variant demographic characteristics, as well as household and month fixed effects. 90% confidence intervals are computed using standard errors clustered at the household level. Costs and revenues exclude costs and earnings associated with the provision of labor to other households or firms. All variables measured in THB are winsorized with respect to the 99% percentile.

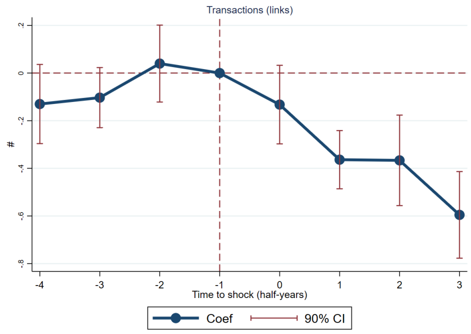

(a) Total transactions

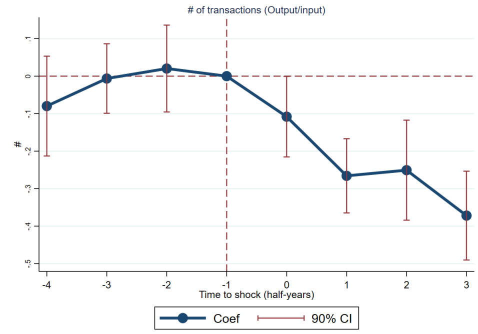

(b) Supply-chain (sales) network transactions

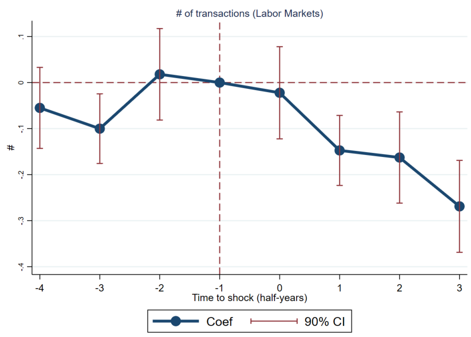

(c) Labor network transactions

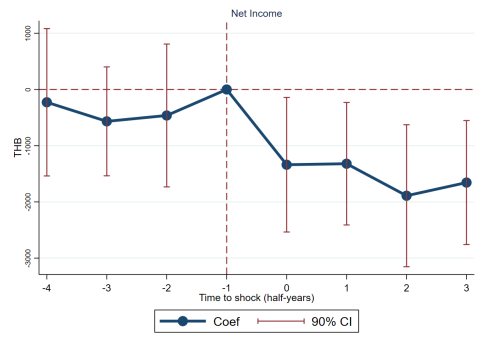

(d) Total income

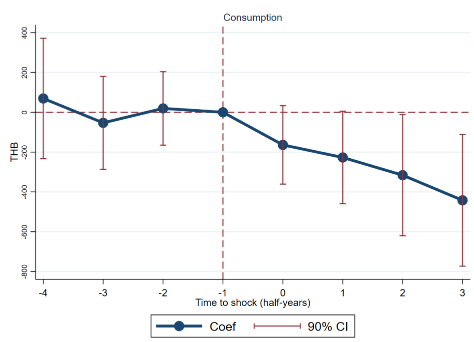

(e) Consumption Spending

Figure 2: Indirect effects on transactions, income and consumption

Note: The Figure presents flexible difference-in-difference estimates of the indirect effects of idiosyncratic shocks on local businesses, following equation (3). All regressions include household fixed effects, event fixed effects, month fixed effects, village- and year-fixed effects, and household size, household average age and education, and the number of adult males and females in each household. Each dot captures differences in changes in outcomes with respect to the half-year preceding the shock (-1) between more- and less-exposed households. Standard errors are two-way clustered at the household (i) and shock level (j). All variables measured in THB are winsorized with respect to the 99% percentile.

Tables

Table 1: Summary statistics

Panel A: Household baseline characteristics	N	Mean	S.D.	10th %ile	90th%ile
Number of household members	509	4.54	1.87	2	7
Number of adults	509	2.87	1.38	1	5
Household head age	507	51.95	13.45	35	70
Average age	509	34.14	12.11	21	52
Household head is a male	507	0.77	0.42	0	1
Years of schooling: Household head	504	4.49	2.59	3	7
Years of schooling: Household maximum achievement	509	8.19	3.64	4	14
Years of schooling: Household average	509	5.09	2.17	3	8
Panel B: Household finance (annual data)					
	N	**Mean**	**S.D.**	**10th %ile**	**90th%ile**
Net Income in THB:					
Farm	7635	134389	1378506	-150	316500
Off-farm family business	7635	19095	115540	0	40700
Labor	7635	52816	108492	0	152222
Total from operations (farm+off-farm + labor)	7635	173327	618277	4974.07	410723
Gift/transfers	7635	24107	183826	-11613	75706
Total net income (Operations+Gifts/Transfers)	7635	197434	644150	16241	446693
Consumption in THB					
Food	7635	32952	21915	11931	60559
Total consumption	7635	98149	99486	24330	204512
Household Assets and Debt					
Total Assets (THB)	7635	2448596	7431394	194277	4817110
Fixed Assets/ Total Assets (%)	7635	53	27	13	88
Total debt/Total assets (%)	7635	12	21	0	27
Households with outstanding loans (%)	7635	83	38	0	100
Households with outstanding loans from institutions (%)	7635	48	50	0	100
Households with outstanding loans from personal lenders (%)	7635	30	46	0	100
Panel C: Village networks (annual data)					
	N	**Mean**	**S.D.**	**10th %ile**	**90th%ile**
Baseline kinship networks: Degree (Number of links)	8344	2.36	2.19	0	6
Baseline kinship networks: Access (any link)	8344	0.77	0.42	0	1
Financial networks: Degree	8344	0.65	1.36	0	2
Financial networks: Access	8344	0.35	0.48	0	1
Sales networks: Degree	8344	1.26	2.64	0	3
Sales networks: Access	8344	0.48	0.50	0	1
Labor-market network: Degree	8344	3.07	4.42	0	9
Labor-market network: Access	8344	0.62	0.49	0	1

Note: Panel A reports summary statistics on demographic characteristics measured at baseline. Panel B reports household financial characteristics based on annual averages using a balanced panel of 509 households. Farm income includes income from agriculture, livestock, fish and shrimp. Off-farm income excludes earnings from labor provision. In both cases income is net of operation costs. Gifts and transfers include transactions from both households inside and outside the village, as well as receipt of government transfers. Consumption includes spending and consumption of home production. In Panel C, all networks are unvalued and undirected; all links have equal weight and the direction of the transaction is not considered. Kinship networks are measured at baseline; transaction networks are measured on an annual basis. Financial networks are constructed based on gifts and loans between households in the same village. Supply chain networks include transactions of raw material and intermediate goods between businesses operated by households in the same village. Labor networks include relationships through paid and unpaid labor between households in the same village. Degree: Number of households with whom each household transacted in each year. Access: Takes the value of 1 if the household has participated in the network in a given year and 0 otherwise.

Table 2: Effects on spending and family businesses

Panel A: Only shocks occurring in the first half of the sample.

	(1) Health Spending	(2) Total spending	(3) Business spending	(4) Hired labor (Hrs/Month)	(5) HH Labor (Hrs/Month)	(6) Revenues
Post X Treatment	539.7	623.3	-1757.1	-14.33	-11.05	-1714.7
	(92.26)	(366.4)	(829.4)	(7.498)	(8.731)	(1062.5)
Baseline mean (DV)	152.6	5451.0	7610.2	18.11	154.1	14939.0
Observations	22709	22709	22709	22708	22708	22709
Number of events	246	246	246	246	246	246
R-Squared	0.0490	0.154	0.782	0.578	0.712	0.620

Panel B: All shocks

	(1) Health Spending	(2) Total spending	(3) Business spending	(4) Hired labor (Hrs/Month)	(5) HH Labor (Hrs/Month)	(6) Revenues
Post X Treatment	413.1	655.2	-1410.6	-9.754	-15.32	-1850.3
	(61.65)	(344.0)	(537.8)	(4.824)	(6.562)	(697.1)
Baseline mean (DV)	158.0	5937.9	7462.1	16.14	142.0	14960.4
Observations	43151	43151	43151	43150	43150	43151
Number of events	469	469	469	469	469	469
R-Squared	0.0439	0.107	0.749	0.667	0.657	0.541

Note: The Table reports estimates of β from equation (2) for different outcomes. Each column reports differences between treatment and placebo households in changes in outcomes before and after the shock. All regressions control for household demographic characteristics, household and month fixed effects. Standard errors are clustered at the household level. Costs, labor, assets and revenues are aggregated across all businesses operated by household members, and exclude revenues and costs of wage labor provision to other businesses or households. Hired labor and labor provided by household members are measured in hours/month.

Table 3: Propagation of idiosyncratic shocks

	(1) Input/Output	(2) Hired labor	(3) All transactions	(4) Income	(5) Total spending
Post X closeness (village network)	-0.200	-0.115	-0.315	-1,267.103	-303.551
	(0.062)	(0.044)	(0.078)	(443.764)	(160.860)
Observations	410,578	410,578	410,578	410,578	410,578
R-squared	0.440	0.231	0.374	0.198	0.621
Pre-period Mean	0.999	0.470	1.469	10486	7265
Number of events	391	391	391	391	391

Note: The Table presents estimates of β from equation (4). $Closeness_{i,j}$ denotes inverse distance to the shocked household during the year preceding the shock to j. Each coefficient captures differences in changes in outcomes before and after the shock between more- and less-exposed households, through village networks. Each regression includes household (i), event j, and month fixed effects, as well as demographic characteristics such as household size, average age, education and number of male and female adults. Standard errors are two-way clustered at the household (i) and event (j) level.

Table 4: Propagation effects on outgoing and incoming transactions

	(1)	(2)	(3)	(4)	(5)	(6)
	Input/output		Labor		Total transactions	
	Outgoing	Incoming	Outgoing	Incoming	Outgoing	Incoming
Post X closeness (village network)	-0.078	-0.121	-0.087	-0.028	-0.165	-0.150
	(0.044)	(0.032)	(0.025)	(0.027)	(0.050)	(0.044)
Observations	410,578	410,578	410,578	410,578	410,578	410,578
R-squared	0.534	0.266	0.154	0.218	0.435	0.254
Pre-period Mean	0.497	0.501	0.182	0.288	0.679	0.790
Number of events	391	391	391	391	391	391

Note: The Table presents estimates of β from equation (4). $Closeness_{i,j}$ denotes inverse distance to the shocked household during the year preceding the shock to j. Each coefficient captures differences in changes in outcomes before and after the shock between more- and less-exposed households, through village networks. Each regression includes household (i), event j, month fixed effects, and demographic characteristics such as household size, average age, education and number of male and female adults. Standard errors are two-way clustered at the household (i) and event (j) level.

Table 5: Estimating propagation á la Fadlon and Nielsen (2019)

	(1)	(2)	(3)	(4)	(5)
	Input/Output	Hired labor	All transactions	Income	Consumption
Post X Treatment	-0.214	-0.0641	-0.278	-1426.3	-351.0
	(0.0853)	(0.0517)	(0.103)	(582.7)	(201.0)
Baseline mean (DV)	1.215	0.564	1.779	9292.2	6596.4
Observations	35111	35111	35111	35111	35111
Number of events	376	376	376	376	376
Adj. R-Squared	0.451	0.193	0.353	0.197	0.558

Note: The table reports results of estimating equation (2) using the subsample of households with a direct or indirect connection to the shocked household; the control group is households with a direct or indirect connection to a placebo household. Standard errors are clustered at the household level.

Table 6: Propagation of idiosyncratic shocks through supply-chain and labor-market networks

	(1) Input/Output	(2) Hired labor	(3) All transactions	(4) Income	(5) Total spending
Post X closeness (supply-chain network)	-0.227	0.022	-0.206	-85.229	60.058
	(0.065)	(0.040)	(0.081)	(488.795)	(176.647)
Post X closeness (labor-market network)	-0.035	-0.210	-0.244	-1,332.916	-486.448
	(0.066)	(0.043)	(0.083)	(446.814)	(159.941)
Observations	410,578	410,578	410,578	410,578	410,578
R-squared	0.441	0.231	0.374	0.198	0.621
Pre-period Mean	0.999	0.470	1.469	10486	7265
Number of events	391	391	391	391	391

Note: The Table presents estimates of β from equation a variation of (4) where $Closeness_{i,j}$ denotes inverse distance to the shocked household during the year preceding the shock to j, by type of network. Each coefficient captures differences in changes in outcomes before and after the shock between more- and less-exposed households, through supply-chain and labor-market networks. Each regression includes household (i), event j, and month fixed effects, as well as demographic characteristics such as household size, average age, education and number of male and female adults. Standard errors are two-way clustered at the household (i) and event (j) level.

Table 7: Village-level determinants of propagation

Panel A: Village-level variation in degree of shocked household					
	(1) Input/Output	(2) Hired labor	(3) All transactions	(4) Income	(5) Consumption
Post X Degree (z-score)	-0.0256	-0.0439	-0.0695	-363.9	-72.44
	(0.0130)	(0.0137)	(0.0192)	(91.83)	(45.77)
Baseline mean (DV)	1.012	0.474	1.486	10607.3	7316.4
Observations	453958	453958	453958	453958	453958
Number of events	391	391	391	391	391
Adj. R-Squared	0.422	0.207	0.353	0.197	0.609
Panel B: Village-level variation in pre-period network density					
	(1) Input/Output	(2) Hired labor	(3) All transactions	(4) Income	(5) Consumption
Post X Density (z-score)	-0.0425	-0.0409	-0.0834	-391.4	-148.6
	(0.0150)	(0.0137)	(0.0205)	(111.4)	(34.88)
Baseline mean (DV)	1.012	0.474	1.486	10607.3	7316.4
Observations	453958	453958	453958	453958	453958
Number of events	391	391	391	391	391
Adj. R-Squared	0.422	0.207	0.353	0.197	0.609

Note: Panels A and B report results corresponding to equation (5) using degree centrality of the shocked household and network density as proxies of village-level exposure to shocks, respectively. Standard errors are clustered at the event level.