Risk Sharing and Portfolio Choice: Incentives and Asymmetries

1 Extended Abstract

I propose a lab-in-the-field experiment in which I randomly vary the income of an individual's social network in a setting with risk sharing and high redistributive pressure in order to measure i) asymmetries in beliefs over demand and supply of credit and ii) how marginal changes in a household's access to informal credit impacts its labor supply and investment choices. The goal of this experiment is to understand how beliefs and incentives can cause frictions in informal credit markets and to quantify the spillover effects of household earnings on financial behavior across a social network. With a model in which households choose how much to invest in a risky or safe asset in settings where there is demand for risk sharing due to incomplete insurance markets, I create a framework to interpret the experimental results. I then use the model to i) estimate how heterogeneous income shocks impact network consumption, investment, risk level of investment portfolio, and income distribution; ii) consider counterfactual credit systems (such as the introduction of formal credit institutions); and iii) re-evaluate the findings of several key papers in the literature. This approach will advance the literature on the impact of informal insurance and redistributive pressures on economic behavior and network well-being, and inform policy that promotes formal insurance institutions or government spending on cash transfers. To my knowledge, this is the first paper to study the spillover effects of earnings attributable to risk sharing in one's network on consumption, investment, and transfer behaviors of a network.

1.1 Motivation

In developing countries, the network spillover effects of financial decisions are particularly acute (Fafchamps, 2003). Due to the widespread lack of access to formal credit institutions, many social networks engage in informal risk sharing, transfers of cash between friends and neighbors, to insure against risk, make large purchases, and facilitate redistribution (Banerjee and Duflo, 2007; Fafchamps, 2011). Previous work has shown that risk sharing lowers the correlation between income realizations and consumptions which suggests that risk sharing successfully insures households (Ambrus et al., 2014). However, in most cases consumption is not completely insured and under certain characteristics risk sharing can lead to a more fragile financial system (Ambrus et al., 2014; Acemoglu et al., 2015). Redistributive pressures from risk sharing have also been shown to lead earners to hide income from their network, to increase demand for illiquid assets due to a high social tax on earnings by one's network, and to explain consumption patterns; income increases lead to more demand for interpersonal cash transfers (i.e., redistributive pressure) so earners

change their spending and investment behavior (Beekman et al., 2015; Jakiela and Ozier, 2016; Squires, 2024). This tax has been shown to decrease labor market participation and discourage entrepreneurship (Alby et al., 2020; Carranza et al., 2022). However, previous work has shown that there is both low demand for formal insurance and underinvestment in profitable but risky assets (Berkouwer and Dean, 2022; Bryan et al., 2014; Ito and Kono, 2010). This is an expected result if we assume that people correctly estimated their level of insurance and adjusted their risk-taking behavior. Thus, in order to determine how risk sharing impacts household well-being and aggregate network well-being as compared to a world without risk sharing or with formal insurance markets, it is important to understand not only how behavior of the earners changes but also how the behavior of the "requesters" is impacted by risk sharing. If a household experiences an increased level of insurance from their network, this may lead to, for example, overconsumption ("freeloading") or overly risky investment behavior. It is possible, then, that risk sharing may have negative effects that are not yet accounted for. On the other hand, an increased level of insurance may increase investment and risk taking behavior from an otherwise under-saving and over-safe household. If instead investment behavior doesn't change with insurance level, then households are not optimizing. Either they are not responding to the change in insurance level or are unaware it has changed. Estimating this effect (or lack thereof) and identifying the mechanisms that cause it can play a crucial role in understanding i) where and when the introduction of formal insurance institutions is necessary or even harmful, ii) how government cash transfers—which often reach only subsets of villages—impact investment behaviors of households across the village, and iii) how risk sharing impacts overall network/village outcomes.

1.2 Experiment Design

The experiment is designed to create random variation in the amount of informal insurance an individual has access to from their social network. Recruitment will occur in several randomly chosen rural villages. Each subject will complete a questionnaire about their social network, interpersonal cash transfer history, investments, consumption, savings, etc. Subjects will be paired with one other village member whom they have stated to be a member of their social network. When the identity of their "partner" is revealed, I will ask each subject what is the maximum amount of money they would be willing to transfer to the other, what is the maximum they think the other would loan them, and what is the most they would feel comfortable asking for from the other. I will then randomly assign pairs to a treatment or control group. For pairs in the treatment group, one individual will be randomly assigned to the observed group—their transfer is observed—and the other to the unobserved group—their transfer is unobserved. Members of the observed group will be transferred a randomly assigned amount of cash, ranging from 5 to 70 USD. Members of the unobserved group will each be transferred the same amount of cash, 5 USD, and then be informed of the amount received by the other member of their pair.

In the control group, neither individual will know the size of the transfer to the other. Each individual can then choose to consume the cash immediately or invest in any of a set of lotteries with varying risk levels, where expected value is increasing in probability of losing. After the lotteries pay out, members of the observed group will be asked again how much they are willing to transfer to their "partner" and members of the unobserved group will be given the opportunity to ask their partner for any amount. If the amount asked for is greater than zero and less than what the member of the observed group stated as their maximum, a transfer will occur. Weeks later, I will survey the individuals again on transfers, investments, consumption, etc., and ask what they spent the cash on. Noting that each pair is a single observation, I will conduct power calculations to determine the necessary sample size to estimate these curves.

This design will allow me to measure several key parameters of the model of risk sharing behavior. First, I can measure how individuals perceive the change in their level of insurance when the income of individuals in their network increases as compared to how much it actually changes. Second, I can measure how an increase in their level of insurance changes investment behavior while holding constant own income. Third, I can measure the perceived social cost of requesting a transfer as well as the social cost of refusing to transfer. With sufficient variation, I can map these outcomes as a function of own income, "partner's" income, and social proximity to their "partner." Combined, I will be able to estimate how changes in a network's income impacts how insured a household is and how risk behavior responds to changes in its level of insurance.

1.3 The Model

The model is set up as follows. There exist N three-period-lived households (nodes) with an exogenous, directed social network, heterogeneous asset ownership, and income uncertainty. Households are perfectly informed about the past incomes of members of the network and receive noisy signals of their future incomes. In each period, households set consumption and an investment portfolio to maximize lifetime expected utility. Households can choose to save using a safe and/or a risky asset, where income shocks in the following period are increasing in both probability and size in amount invested in the risky asset. In each period, households can also choose to give or request cash transfers from other households whom they share an edge with. The decision to conduct an interpersonal cash transfer for any pair of nodes is a function of each node's current realized income, expected future income, the social costs associated with requesting or denying a transfer, and each node's perceived value of the edge. A transfer will occur if both households expect to be better off with the transfer, implying that transfers can never exceed the receiver's perceived value of the edge. Each household's consumption and investment decisions are thus a function of current realized income, expected payoffs of each investment asset, riskiness of each asset, and perceived access to credit.

With this structure, I can model how any individual's consumption, investment, and

transfer behavior changes as a function of network density, initial wealth distribution, and the parameters found in the experiment. I will compare these results to counterfactual situations with, for example, no risk sharing/redistributive pressures, the introduction of formal credit institutions, or the introduction of different types of cash transfers. Doing so can establish the conditions under which such policies have net positive outcomes. Using the insights from the model, I will re-evaluate the findings of several key papers in this literature, namely Banerjee et al. (2024) and Carranza et al. (2022). Taking their results seriously and assuming that our estimated parameters are consistent across the (relatively similar) settings, the model can provide new insights on how investment behavior and well-being across a social network are influenced by i) the introduction of formal credit institutions and consequent change in network structure and ii) allowing for hidden income.

1.4 Concluding remarks

While previous work has identified the effect of risk sharing and redistributive pressures on earnings, this is the first attempt to estimate the impact of risk sharing on consumption and investment behaviors throughout the network. Thus, I present a new approach to estimating how risk sharing and redistributive pressures impact consumption, investment, and well-being throughout a network. This study has implications for understanding how insurance and redistribution via risk sharing and interpersonal transfers affect the financial behaviors of households across a network, and for understanding the overall impact of government spending on cash transfers that may reach only subsets of networks or villages in settings where redistributive pressures are common. For example, if certain network structures are conducive to freeloading or too much risk taking, it may be beneficial to hasten the introduction of formal insurance institutions in these areas and to ensure cash transfers are more evenly distributed. If other network structures lead to more optimal investment and risk taking behavior, the costs of introducing formal insurance or credit institutions may outweigh the benefits, as suggested in Banerjee et al. (2024), and the benefit of cash transfers may be undermined by redistributive pressures.

References

Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015.

Philippe Alby, Emmanuelle Auriol, and Pierre Nguimkeu. Does social pressure hinder entrepreneurship in africa? the forced mutual help hypothesis. *Economica*, 87(346): 299–327, 2020.

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

Abhijit Banerjee, Emily Breza, Arun G Chandrasekhar, Esther Duflo, Matthew O Jackson, and Cynthia Kinnan. Changes in social network structure in response to exposure to formal credit markets. *Review of Economic Studies*, 91(3):1331–1372, 2024.

Abhijit V Banerjee and Esther Duflo. The economic lives of the poor. *Journal of economic perspectives*, 21(1):141–167, 2007.

Gonne Beekman, Marcel Gatto, and Eleonora Nillesen. Family networks and income hiding: evidence from lab-in-the-field experiments in rural liberia. *Journal of African Economies*, 24(3):453–469, 2015.

Susanna B Berkouwer and Joshua T Dean. Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households. *American Economic Review*, 112(10):3291–3330, 2022.

Gharad Bryan, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. Underinvestment in a profitable technology: The case of seasonal migration in bangladesh. *Econometrica*, 82(5):1671–1748, 2014.

Eliana Carranza, Aletheia Donald, Florian Grosset, and Supreet Kaur. The social tax: Redistributive pressure and labor supply. Technical report, National Bureau of Economic Research, 2022.

Marcel Fafchamps. *Rural poverty, risk and development*, volume 144. Edward Elgar Publishing, 2003.

Marcel Fafchamps. Risk sharing between households. *Handbook of social economics*, 1: 1255–1279, 2011.

Seiro Ito and Hisaki Kono. Why is the take-up of microinsurance so low? evidence from a health insurance scheme in india. *The Developing Economies*, 48(1):74–101, 2010.

Pamela Jakiela and Owen Ozier. Does africa need a rotten kin theorem? experimental evidence from village economies. *The Review of Economic Studies*, 83(1):231–268, 2016.

Munir Squires. Kinship taxation as an impediment to growth: experimental evidence from kenyan microenterprises. *The Economic Journal*, page ueae025, 2024.