

Insurance Networks and Poverty Traps

Arun Advani*
University of Warwick, CAGE, and IFS

January 19th, 2019

Abstract

Poor households regularly borrow and lend to smooth consumption, yet we see much less borrowing for investment. This cannot be explained by a lack of investment opportunities, nor by a lack of resources available collectively for investment. This paper provides a novel explanation for this puzzle: investment reduces the investor's need for informal risk sharing, weakening risk-sharing ties, and so limiting the amount of borrowing that can be sustained. I formalise this intuition by extending the canonical model of limited commitment in risk-sharing networks to allow for lumpy investment. The key prediction of the model is a non-linear relationship between total income and investment at the network level — namely there is a network-level poverty trap. I test this prediction using a randomised controlled trial in Bangladesh, that provided capital transfers to poor households. I exploit variation in the number of program recipients in a network to identify the location of the poverty trap. I show that above the threshold intra-network lending becomes an important source for funding additional investment, as predicted by the model. My results highlight how capital transfer programs can be made more cost-effective by targeting communities at the threshold of the aggregate poverty trap.

JEL codes: D12, D52, O12, O43

Keywords: Poverty Trap; Risk Sharing; Limited Commitment; Networks

*Department of Economics, University of Warwick, Coventry, CV4 7AL. Email: economics@arunadvani.com. With thanks to Orazio Attanasio, Richard Blundell, Áureo de Paula, and Imran Rasul for their invaluable support and advice. Thanks also to Oriana Bandiera, Marco Bassetto, Martin Cripps, Mariacristina De Nardi, Dave Donaldson, James Fenske, Selim Gulesci, Clément Imbert, Costas Meghir, Nicola Pavoni, Bryony Reich, Tim Worrall, seminar participants at Bocconi, Bristol, Institute for Fiscal Studies, Geneva, Gothenburg, Manchester, Oxford, Paris School of Economics, Stockholm School of Economics, Toulouse School of Economics, Universitat Autònoma de Barcelona, University College London, University of Southern California, and Warwick, and conference participants at the European Econometric Society Winter Meeting 2016, OxDev 2017, NBER-BREAD 2017, CSAE 2018, NEUDC 2018 and Y-RISE 2018. All errors remain my own.

1 Introduction

An old literature going back to [Rosenstein-Rodan \(1943\)](#) suggests that the failure of poor economies to develop comes from an inability to coordinate, where multiple simultaneous investments could be profitable, but alone none of these investments will be. However, there are many investments which are profitable even without others' investment, and yet do not take place. For example, in rural villages the purchase of small capital goods such as livestock is typically highly profitable, and yet we see little investment by most households ([Bandiera et al., 2017](#); [Banerjee et al., 2015](#); and also [de Mel et al., 2008](#) in the context of small businesses).

One obvious explanation for the lack of investment is that households are poor, and so neither have the resources to invest nor access to formal credit. However, households regularly borrow from (and lend to) friends and neighbours, using this to smooth consumption ([Townsend, 1994](#); [Udry, 1994](#)). The puzzle, then, is why households are able to borrow informally for consumption but not for high return investments? It cannot be explained by a lack of resources (incomes and assets): whilst individually households in these 'risk-sharing networks' have few resources, collectively they have the resources needed for investment. So why don't households pool resources to allow some households to engage in investment?

This paper offers a new explanation, and empirical evidence, for this puzzle. The key idea is that investment *reduces* the capacity of investing households to provide informal consumption smoothing. To see this, note that borrowing and lending for consumption smoothing – 'informal insurance' – is sustained by reciprocity: a household lends today because it wants the possibility of borrowing in the future, when it has a low income. Rather than writing formal contracts, borrowing occurs informally, with lenders motivated by loss of future access to borrowing if they do not lend when their incomes are relatively high. What makes borrowing for investment different is that an investing household will on average be better off in the future. Having investment income as well as labour income will reduce its need to borrow for consumption smoothing in future periods. This reduced need for borrowing limits the amount it can be asked to lend — ask for too much and the household would rather just lose access to future insurance. The reduced capacity to provide other households with consumption smoothing weakens those other households' incentives to lend for investment.

I first develop a formal theoretical model that captures this mechanism, and then provide empirical evidence from a large scale randomised controlled trial. The model combines the key elements discussed above: informal insurance with limited commitment and lumpy (indivisible) investment. I show that the mechanism described, trading off insurance and investment, can lead to a *network-level* poverty trap i.e. the long run equilibrium level of income in the network will depend on the initial conditions. I then provide empirical evidence for the relevance of this mechanism. Using data from a large scale, long term randomised controlled trial in Bangladesh, I find evidence that networks in Bangladesh are indeed in a network-level poverty trap. Networks which receive a large enough capital injection escape the trap and engage in additional investment, and this financed by within-network lending.

More precisely, I develop a model which captures four important characteristics of households in village economies: (i) households are risk-averse and have volatile incomes; (ii) they are able to engage in consumption smoothing by making inter-household transfers; (iii) households have *limited commitment* in their risk-sharing arrangement i.e. at any point in time, the expected value of continuing any risk sharing must be at least as good for each household as the value of walking away forever (‘autarky’); (iv) households have the opportunity each period to invest in a ‘lumpy’ (indivisible) asset. The first three characteristics lead to models of risk sharing with (dynamic) limited commitment, as studied by [Kocherlakota \(1996\)](#) and [Ligon et al. \(2002\)](#). The fourth characteristic has also been studied in a number of development contexts ([Rosenzweig and Wolpin, 1993](#); [Fafchamps and Pender, 1997](#); [Munshi and Rosenzweig, 2016](#)). My innovation is to combine these standard features and show that there are important interactions between them, which can provide an explanation for the long-standing puzzle of underinvestment. In particular, a poverty trap naturally arises in this model: the long run equilibrium income distribution depends on the initial level of capital invested. The ‘depth’ of this trap – the amount of income the network needs to escape the trap – is greater if commitment is limited, when no household can afford to invest in autarky.

A *network* poverty trap exists if the long run capital stock depends on *aggregate* income in the network: investment can only take place if aggregate income exceeds some threshold. Above the threshold, it will be possible for the network to make some investment. This raises future income, ensuring further investments are possible, and allowing all households to eventually invest. In contrast, networks below the threshold will never have enough income to make even the first investment, and so remain persistently poor. However, with only a ‘small push’ that provides some initial capital, these networks can be set on a path of further investment and income growth.

With full commitment, only networks with too few resources for investment would be in the poverty trap. However, empirically many networks have resources that they do not use for investment, despite the high returns. Limited commitment frictions can explain why investment may not occur in these cases. Limited commitment makes resource pooling more difficult, hence investing households will only be able to credibly promise smaller transfers. A larger share of investment must therefore come out of their own pockets. When no household has enough income to want to invest if in autarky, limited commitment necessarily reduces investment.

I provide evidence for such a poverty trap using data from a large scale, long term randomised controlled trial in Bangladesh. These data cover 27,000 households across 1,400 villages in the poorest districts of rural Bangladesh. They were collected as part of an asset transfer program implemented by the NGO BRAC. The intervention randomised villages into either treated or control status, and then provided assets (typically cows) to the poorest households in treated villages. These transfers were worth more than 50% of median income for the households that received them. Asset transfers took place after data collection in 2007, and follow-up surveys were carried out in 2009 and 2011. The program was evaluated by [Bandiera et al. \(2017\)](#), who show that the program has large and sustained effects on both earnings and asset accumulation.

Four features of the data make them suitable for my context. First, the data are from a randomised controlled trial that encompassed a large cross-section of networks from more than 1,400 villages. This is important since my model predictions are at the network level. Second, in a subsample of my data, the data record the exact links used for risk-sharing transfers. I exploit this to construct a good proxy for the appropriate risk-sharing network in the full dataset, which recent work suggests it is important to measure well when studying risk sharing (Mazzocco and Saini, 2012). Third, the program provided large injections of lumpy capital (cows), with significant variation in the number of transfers across villages. This provides the exogenous variation in aggregate income necessary for my test of a network-level poverty trap. Finally, the data cover a long time scale, with a follow-up survey four years after the initial capital injection. This provides a large enough window to study how the initial injection affects *additional* investment, which is key to understanding whether a network has left the poverty trap.

The main empirical findings are as follows. First, aggregate investment in cows by risk-sharing networks between 2009 and 2011 is zero on average if the network received less than \$3,500 (PPP 2007) of capital from the program, 7% of median network income. This threshold is determined using a formal statistical test for a structural break with unknown break point i.e. a test for a change in the slope of additional investment with respect to the capital provided when the location of the slope change being unknown. Above this level, aggregate investment is increasing (and linear) in the aggregate amount of capital provided by the program. This relationship is exactly as predicted by the model. Second, I show that 80% of the additional investment above the threshold is done by households who were not recipients of capital from the program, but are in the same risk-sharing network: direct evidence that some network spillover is occurring. Third, I show that this investment is being funded by intra-network lending. This is the key mechanism in the model, and is something which which was not taking place in the absence of transfers.

This paper contributes to three distinct strands of literature. First, it contributes to the large literature on poverty traps. Whilst poverty traps are an old idea, empirical work has failed to find convincing evidence for any of the specific mechanisms that have been proposed (Kraay and McKenzie, 2014). The novel aspect here is that by introducing risk sharing, the poverty trap occurs at the network level, and by introducing limited commitment, risk-sharing networks with enough resources to invest may still choose not to invest. I use standard tools – non-convexity in production, of which lumpiness is a particular example, and a financial friction (limited commitment) – to generate the poverty trap (see for example Banerjee and Newman, 1993; Aghion and Bolton, 1997; and Ghatak, 2015). However, by embedding these in a risk-sharing framework, the poverty trap in my model occurs at the network level. I provide empirical evidence that we do indeed see a trap at this level, and my results are not consistent with an individual level trap. My mechanism is distinct from the group level poverty traps of Rosenstein-Rodan (1943) and Murphy et al. (1989), which are purely due to coordination failure. In my model it is not necessary for all households to be simultaneously

coordinated in investment, nor is coordination alone – without the provision of assets – sufficient to generate further investment. I provide evidence that rules out poverty trap models which rely on increasing returns to coordinated investment, including due to externalities, fixed costs, or learning.

Second, I contribute to the literature on risk sharing with frictions (Kocherlakota, 1996; Ligon et al., 2002; among others). In particular, there is a growing literature examining how endogenously incomplete insurance affects and is affected by opportunities in other markets (Attanasio and Rios-Rull, 2000; Attanasio and Pavoni, 2011; Ábrahám and Cárceles-Poveda, 2009; Ábrahám and Laczó, forthcoming; Morten, forthcoming). Attanasio and Pavoni (2011) highlight an important trade-off between using insurance and using (continuous) investment to provide consumption smoothing. A similar trade-off is present in my model, but the ‘lumpiness’ of investment in my context (mirrored by many development applications) changes the nature of the decision-making, and creates the possibility of a poverty trap. Morten (forthcoming) also considers a model with risk sharing and a binary decision, but where the decision only directly affects payoffs today. By contrast, in this paper investment has permanent effects on the distribution of income, allowing me to study questions of longer term development and growth. It also opens the door for the study of other long term discrete investment decisions, such as irrigation (Rosenzweig and Wolpin, 1993; Fafchamps and Pender, 1997), education (Angelucci et al., 2018), and permanent migration (Munshi and Rosenzweig, 2016), in the context of risk sharing with limited commitment.

Third, I contribute directly to the recent and growing work on asset transfer programs (Bandiera et al., 2017; Banerjee et al., 2015; de Mel et al., 2008; Morduch et al., 2015). These studies find that in many cases, across a range of countries and contexts, asset transfer programs are very successful in increasing incomes. My paper provides a possible explanation for why such one-off transfers of assets appear to have larger effects on income growth than smaller, but longer term, cash transfer programs such as Progresa (Ikegami et al., 2016). Small increases in income will still be partly smoothed away, rather than providing the basis needed for additional investment. It also suggests a route for increasing the impact of these interventions: targeting at a network rather than a household level. By providing enough resources at an aggregate level, these programs can provide the ‘small push’ that networks need to get out of the poverty trap. My results highlight how a budget neutral redistribution of asset transfers across networks, focusing on getting networks just over the threshold, can increase additional investment. Restructuring the existing policy in this way could have increased additional investment four years after the program by 44%, relative to using household-level targeting.

The next section develops the model formally, and provides the theoretical results. Section 3 describes the data and context for my empirical work. Section 4 tests the key predictions of the model, and provides additional supportive evidence for the mechanism proposed. The final section concludes.

2 A Model of Insurance, Investment, and a Poverty Trap

2.1 Model Environment

Consider an infinite-horizon economy composed of N households. Households have increasing concave utility functions defined on consumption that satisfy the Inada conditions. They also have a common geometric discount rate, β .¹ Each period t , households receive endowment income $\mathbf{y}_t = \{y_t^1, \dots, y_t^N\}$ drawn from some (continuous) joint distribution \mathcal{Y} . Individual incomes are bounded away from 0, and aggregate income $Y_t := \sum_{i=1}^N y_t^i$ is bounded above by Y^{\max} . Income draws are assumed to be iid over time, but may be correlated across households within a period. I define $s_t^i := y_t^i/Y_t$ as household i 's share of aggregate endowment income in period t . To ease notation, hereafter I suppress the dependence of variables on t .

The households belong to a single network, and they may choose to engage in risk sharing. Since households are risk-averse, and endowment incomes are risky, there is scope for mutually beneficial risk sharing. In particular, an informal agreement in which households with good income shocks in any period make transfers to those with bad income shocks will improve the expected discounted utility for all households. I model this risk sharing as net transfers, τ^i , made by households $i = 2, \dots, N$ to household 1.² Household consumption will then be $c^i = y^i - \tau^i$, and $\tau^1 \equiv -\sum_{i=2}^N \tau^i$, where I suppress the dependence of all these objects on the shock \mathbf{y} to ease notation.

An impediment to risk sharing is the presence of dynamic limited commitment (Kocherlakota, 1996; Ligon et al., 2002). Households may, in any period, choose to walk away from the arrangement, keeping all of their income that period and then being excluded from the arrangement thereafter. This will limit the amount of risk sharing that can take place.

Thus far, the model is an N household, continuous shocks version of the standard model of risk sharing with dynamic limited commitment. To this problem I introduce the possibility that households may engage in lumpy investment. Precisely, each period a household may choose whether or not to invest in a binary investment, κ . This has a one-off cost d , and pays a guaranteed return of R in all future periods.³ Investment is an absorbing state, and households

¹For work considering risk sharing with heterogeneous preferences, see for example Mazzocco and Saini (2012). For work considering poverty traps with non-geometric discounting, see Banerjee and Mullainathan (2010) and Bernheim et al. (2015).

²In principle, each household could choose how much income to transfer to each other household. Since my interest is only on the total risk sharing that takes place, and not on the precise structure of transfers that are used, I model all transfers as going to or from household 1. For each household there is then a single decision about the net transfers to make (or receive). For work studying how network *structure* and risk sharing interact, see Ambrus et al. (2014) and Ambrus et al. (2015).

³There are two implicit assumptions here. First, the return does not vary with the aggregate number of investments that occur. This rules out both general equilibrium effects, where we might expect to see the return decline as the number of investments increases, and fixed costs, where we might expect to see the return increase. I will show later that in my empirical setting, these are both reasonable. Second, there is assumed to be no risk in the return on investment. This is done to distinguish my mechanism from an alternative mechanism, where a high return activity is also higher risk, so underinvestment occurs because of a *lack* of insurance (see for example Karlan et al., 2014). It is also appropriate to my context. In my empirical setting, income is (slightly) less risky for those who are invested. The normalised risk (variance over squared mean) of income for households who receive asset transfers is .30, compared with .31 for controls.

may hold at most one investment.⁴ Additionally, investments must be held by the household that does the investment – with κ denoting the stock of capital held by each household – although transfers may be made out of investment income. This rules out cooperatives and other joint investment structures.⁵ Now an uninvested household must choose each period what net transfers to make, τ^i , and whether to invest, $\Delta\kappa^i$.

Barring risk sharing and investment, no alternative forms of smoothing are permitted. This rules out external borrowing: whilst a household may engage in implicit borrowing from other households in the risk-sharing network, the network as a whole cannot borrow from the wider world. Limited commitment problems make borrowing *within* the network difficult, even amongst households that interact regularly, so one would expect this problem to be even more severe for lenders from outside the community. I also rule out saving, so that investment is the only vehicle for transferring resources over time.⁶ If private savings were introduced, they would provide an alternative means of transferring resources over time. In the model, a household would give or receive transfers, and then decide what share of resources (if any) to save. However, in the next period, optimal smoothing of consumption will depend on total cash-on-hand (income plus savings) that a household has. Hence a household that saves would effectively be ‘taxed’ on this saving, as it would be considered in the same way as any other income when the next period begins. This idea of a ‘network tax’ discouraging saving has been documented by [Dupas and Robinson \(2013\)](#), who show that poor households appear to have negative *nominal* returns to saving, and [Jakiela and Ozier \(2016\)](#), who show that households are willing to pay to prevent information about good income shocks being revealed.⁷

2.2 Insurance Relationships

To characterise the nature of insurance in the model, I begin with the case when all households have invested. This problem is a many-household, continuous-shocks version of [Ligon et al. \(2002\)](#), and I follow their approach in writing the model as a planner’s problem. Household 1 (the planner) asks for transfers (possibly negative) from all other households today, and makes conditional promises about future utility it will provide to households, to maximise its own utility from present and future consumption. In making these requests and promises, it is subject to two sets of constraints: it must keep its promises about utility it will provide (‘promise keeping’), and it can only make requests that make other households no worse than

⁴The former is a simplifying assumption, which could be relaxed to allow deterministic or stochastic destruction of capital units, at the cost of adding more moving parts to the model. The latter could also be relaxed: all that is needed is some upper bound on the total number of investments a household can hold. This is reasonable in my context, where investments are in livestock: [Shaban \(1987\)](#) and [Foster and Rosenzweig \(1994\)](#) describe how moral hazard issues can limit the ability to hire labour from outside the family to manage livestock.

⁵In my context, fewer than 3% of households with farm assets engage in joint production. For a model with joint ownership of investment, see [Thomas and Worrall \(2018\)](#).

⁶For work studying limited commitment risk sharing with divisible saving, see for example [Ligon et al. \(2000\)](#) and [Ábrahám and Laczó \(forthcoming\)](#).

⁷Allowing for hidden savings would complicate this argument slightly, but as long as investment cannot be hidden – which is likely in many contexts, such as when the investments are livestock – any systematic hiding of savings for investment purposes would be detectable and punishable once investment takes place.

autarky (‘limited commitment’). Appendix A.1 characterises formally the solution of the model in this situation, where there is no meaningful investment decision. The main result is that, as long as no limited commitment constraints bind, the ratio of marginal utilities remains constant over time i.e. households receive a constant share of aggregate consumption. If an individual household’s limited commitment constraint binds, then it receives an increased share of consumption in this and all future periods (until another constraint binds).

When households have not already invested, the planner must also determine how many investments should occur, and who should make those investments. The key challenge in solving the model here is that the discrete investment decision introduces kinks into the value function that characterises the planner’s maximum utility. However, I show that since income shocks are continuous, this smooths out any kinks, as in [Pavoni and Violante \(2007\)](#). Appendix A.2 provides this result formally, and shows the first order conditions characterising the insurance relationship take the same form as when investment was not possible. Conditional on the investment decision, which determines aggregate consumption given aggregate income, households receive a constant share of that consumption when no limited commitment constraint binds, and otherwise the consumption share rises for household(s) whose constraint binds. In the next subsection I study the investment decision.

2.3 Poverty Trap

The first result from the model is that it naturally gives rise to the possibility of a poverty trap: a situation in which the long run equilibrium of the economy depends on its initial state. I will build this result in two steps. First, I suppose that households *are* able to commit i.e. the limited commitment friction is removed. In this case there will be a network-level poverty trap where some communities will be too poor to be able to ever invest. The structure of this trap will be analagous to a household-level trap: the only thing preventing investment is a lack of resources. However, this is insufficient to explain the observation that networks which have the resources choose not to invest. I then reintroduce limited commitment, and show this can ‘deepen’ the poverty trap. Now networks which have sufficient resources to invest under full commitment may not invest, because the lack of commitment prevents resource pooling. This is the key mechanism driving the model.

Full Commitment

Under full commitment there exists a sequence of aggregate income thresholds $\widehat{Y}_{\Delta k}^{\text{FC}}$, one between each possible level of investment and the level above it, such that if $\widehat{Y}_{\Delta k}^{\text{FC}} < Y < \widehat{Y}_{\Delta k+1}^{\text{FC}}$ then it will be optimal to make Δk investments this period. This leads to the possibility of a poverty trap: if an economy never receives a large enough level of aggregate income to reach the lowest threshold, i.e. $Y^{\text{max}} < \widehat{Y}_1^{\text{FC}}$ then it will forever remain with the current income distribution (absent external shocks), whilst if an external shock is provided to produce a ‘small push’ then

further investment will be able to occur over time.⁸ Let $V_{\Delta k}(\cdot)$ denote the value function for the planner when it is constrained to do precisely Δk investments, and then chooses transfers and utility promises optimally.⁹

Proposition 1. *There exists a unique threshold $\widehat{Y}_{\Delta k}^{FC} = \widehat{Y}_{\Delta k}^{FC}(\boldsymbol{\kappa}, N)$ such that with full commitment:*

1. $\forall Y < \widehat{Y}_{\Delta k}^{FC}$, the optimal number of investments is no greater than $\Delta k - 1$;
2. at $Y = \widehat{Y}_{\Delta k}^{FC}$, $V_{\Delta k-1}(\cdot) = V_{\Delta k}(\cdot) \geq V_{\Delta k'}(\cdot) \forall \Delta k'$ i.e. the planner is indifferent between making $\Delta k - 1$ and Δk investments and does not strictly prefer any other level of investment to these; and
3. $\forall Y > \widehat{Y}_{\Delta k}^{FC}$, the optimal number of investments is no fewer than Δk .

There are $N - k$ such thresholds, with $\widehat{Y}_{\Delta k-1}^{FC} \leq \widehat{Y}_{\Delta k}^{FC}$, each implicitly defined by $\Gamma_{\Delta k}(\widehat{Y}_{\Delta k}^{FC}; \boldsymbol{\kappa}, N) \equiv 0$ where $\Gamma_{\Delta k}(\cdot) := V_{\Delta k-1}(\cdot) - V_{\Delta k}(\cdot)$.

Proof. See Appendix B.3. □

Proposition 1 states that for an N -household economy in which $k = \sum_{i=1}^N \kappa^i$ investments have already been made, there are $N - k$ income thresholds, whose level depends on the number of existing investments and the network size, under the assumption that households can commit fully. Intuitively, when aggregate income is very low, it will be optimal to consume it all today, potentially after some redistribution. At higher levels of aggregate income, the utility cost of reducing total consumption by d today (the cost of an investment) is low compared with the expected improvement in future expected utility, so it will become optimal to make one investment. At yet higher levels of aggregate income, multiple simultaneous investments become worthwhile. Network size scales down the per household cost (and return) of each investment.

This threshold result leads naturally to the possibility of a poverty trap, where the long run distribution of income depends on its initial state. When an economy has a low level of capital, the highest possible aggregate income may be lower than \widehat{Y}_1^{FC} , the level needed to make the first additional investment worthwhile. However, at a higher level of initial capital stock the maximum level of aggregate income is higher, allowing further investments to take place in some states of the world.

Lemma 1. *The threshold level of income needed to make Δk additional investments, $\widehat{Y}_{\Delta k}^{FC}$, is decreasing in the existing level of capital k , i.e. $\mathbf{D}_k \widehat{Y}_{\Delta k}^{FC} < 0$, where \mathbf{D}_k is the finite difference operator (the discrete analogue of the derivative) with respect to k .*

Proof. See Appendix B.4. □

⁸In contrast with ‘big push’ models (Rosenstein-Rodan, 1943; Nurkse, 1953; Murphy et al., 1989), here no coordination is needed between agents: an initial push that is large enough to allow one additional investment to occur will then automatically spillover, allowing further investments.

⁹See Appendix A.2 for a precise characterisation of $V_{\Delta k}(\cdot)$.

Under full commitment the poverty trap result has very stark predictions: there are only two possible long run equilibria, $\kappa = \mathbf{0}$ or $\mathbf{1}$.¹⁰ This is because under full commitment only the level of aggregate income matters for whether investment takes place. Suppose there exists a state of the world in which, from a base of zero capital, making at least one investment is optimal for the planner. Then making an investment in the same state of the world, when the same combinations of endowment incomes are realised but when some investments have already occurred, must also be optimal (by Lemma 1). Hence either the economy will remain with zero capital or will converge to a state in which all households are invested.

Limited Commitment

I next consider how the above results are changed by limited commitment. I first show that limited commitment can change the ‘depth’ of the poverty trap: the threshold level of income needed such that doing some investment becomes optimal in equilibrium. To do this I consider how the investment threshold under autarky compares to that with full commitment. The results under limited commitment will fall somewhere between these, depending on the extent of limited commitment. I then show that with limited commitment, a wider range of equilibrium levels of investment are possible.

Under autarky, there exists an income threshold \hat{y} such that if an (uninvested) individual household’s income exceeds \hat{y} it will invest, else it will not.¹¹ All remaining income will be consumed. In my empirical setting all households have incomes below the level needed to invest in autarky i.e. $y^i < y^{\text{aut}}$.¹² In this case the number of investments must necessarily be (weakly) lower under limited commitment than under full commitment. To see this, note that limited commitment reduces the ability to make transfers today in expectation of receiving transfers in the future. Hence equilibrium outcomes under limited commitment are always between the full commitment and autarky outcomes. Since no household wants to invest in autarky, if any investment were taking place under full commitment, under limited commitment it can only be weakly lower.

As well as changing the level of the lowest threshold, \hat{Y}_1 , limited commitment can change the distance between the thresholds. This is important, as it can create long run equilibria where the economy has an intermediate level of capital, rather than the all or nothing result seen under full commitment. To see this, consider the situation where one household has invested. Under full commitment, I showed that the threshold level of income needed to do one additional investment has now fallen. Under limited commitment there is an additional effect: relative to the case where no-one has invested, the household with an investment has an improved outside

¹⁰ $\kappa = \mathbf{0}$ is the poverty trap long run equilibrium, while $\kappa = \mathbf{1}$ is the ‘good’ long run equilibrium. If there were decreasing returns to investment at the aggregate level, the good equilibrium might be less extreme, with only some households ever investing, but there would be the same initial threshold needed to break out of the poverty trap equilibrium.

¹¹To see that such a threshold exists, the same lines of reasoning used in the full commitment case can be replicated. The threshold is implicitly defined as $u(\hat{y}) + \beta \mathbb{E}[\Omega(y', 0)] = u(\hat{y} - d) + \beta \mathbb{E}[\Omega(y', 1)]$, where $\mathbb{E}[\Omega(y', k)]$ is the expected continuation value of autarky with investment level k .

¹²One could also use this model to analyse contexts other than this.

option. This endogenously restricts the set of possible equilibrium transfers. Since household consumption will no longer be a constant share of aggregate consumption, and since owning an investment increases the consumption share for a household, there will no longer necessarily be increasing differences in the planner’s utility when another household invests. The limited commitment analogue of Lemma 1 may therefore not hold: an increase in the level of capital will not necessarily reduce the income thresholds for investment. Instead there are now parameters which can support ‘intermediate’ equilibria, where the long run share of households who are invested is strictly between zero and one.¹³

3 Data from a Randomised Controlled Trial in Bangladesh

3.1 Data Source

I use data from a large scale, long term randomised controlled trial in rural Bangladesh, collected in partnership with the NGO BRAC. The data cover 27,000 households across 1,409 villages, in the poorest 13 districts of rural Bangladesh.

The villages were selected as follows. From each district, one or two subdistricts (upazilas) were randomly sampled. From each of these, two BRAC branch offices were randomly selected for the program, one to be treated, the other as a control (for more details see [Bandiera et al., 2017](#)). All villages within 8km of a sampled branch office were then included in the final sample, giving the total of 1,409 villages, with a median of 86 households.

A census of households in each village took place in 2007. This asked questions on demographics of household members, and their education and employment statuses, as well as collecting detailed information on household assets. This was used both to construct a sampling frame for the further surveys, and for targeting the program.

A sample of households was then selected from each village. A participatory wealth ranking in the census divided households into one of four wealth categories. All households in the lowest wealth grouping – which includes all households eligible for the program – were sampled, along with a 10% random sample of all remaining households. This gives a sample of 7,111 eligible households, 13,704 ‘ineligible poor’ households (in the bottom two wealth ranks), and 6,162 ‘non-poor’ households. Sampled households were given a baseline survey in 2007, with follow up surveys in 2009 and 2011. In these surveys detailed data were collected on household income, investment, and risk sharing.

Table 1 provides some key descriptives about these households, grouping them into the above categories. Households comprise around four members, but poorer households are smaller as they are more likely to not have a working age man present. This is particularly true in eligible households, where it was used in program targeting (see below for details). Eligible households are very poor, with almost half below the poverty line, and hardly any already own

¹³Hence although all households are ex ante identical, there are long run equilibria where they necessarily have different levels of expected utility. Matsuyama (2002, 2004, 2011) provides other examples of models which have this ‘symmetry-breaking’ property.

cows. Ineligible poor households have higher incomes, consumption, and assets (cows) – and non-poor households higher still – providing evidence that the participatory wealth ranking provides a good measure of relative material standard of living.

Four features of the data make them suitable for my context. First, the data cover a large cross-section of networks, encompassing more than 1,400 villages. This is important since the model predictions are at the network level. Second, in a subsample of the data, exact links used for risk-sharing transfers are measured. This makes it possible to construct a good proxy for the appropriate risk-sharing network in the full dataset (see below), addressing concerns that the whole village is not the level at which risk sharing takes place. Third, the program provided large injections of lumpy capital, with significant variation in the number of transfers across villages (see Figure 1). This provides the variation in aggregate income necessary for my test of a network-level poverty trap. Fourth, households were surveyed again two and four years after the transfers were made. This allows study of how the initial transfers affect later investment decisions, which are necessarily long term.

3.2 Program Structure

The intervention carried out by BRAC was an asset transfer program. Using information from the census survey, household eligibility for the program was determined. Eligibility depended on a number of demographic and financial criteria. A household was automatically ineligible for the program if any of the following were true: (i) it was already borrowing from an NGO providing microfinance; (ii) it was receiving assistance from a government antipoverty program; or (iii) it has no adult women present.¹⁴ If none of these exclusion criteria were met, a household was deemed eligible if at least three of the following inclusion criteria were satisfied: (i) total household land was less than 10 decimals (400 square metres); (ii) there is no adult male income earner in the household; (iii) adult women in the household work outside the home; (iv) school-aged children have to work; (iv) the household has no productive assets.

After the baseline survey, eligible households in treated villages were given a choice of asset bundles. All bundles were worth approximately the same amount, \$515 in 2007 PPP. 91% of treated households choose a bundle with cows, 97% with cows or goats. In the following analysis I treat all treated households as though they actually received cows, but my results are robust to treating those who did not choose livestock as though they received no transfers. Along with assets, treated households also receive additional training from BRAC officers over the following two years. By the 2009 survey, all elements of the program had ceased, except that treated households now had the additional capital they had been provided with. After the 2011 survey, eligible households in control villages also received asset transfers.

One limitation of the program structure, for the purpose of this study, is that while entire villages are either treated or control, variation in the *intensity* of treatment – the value of transfers to a risk-sharing network, proportional to the number of households in the network

¹⁴The last criterion exists because the asset transfers were targeted at women.

who receives transfers – is endogenous, since it depends on characteristics of the households. The ideal experiment for my context would have been to directly randomise villages into G groups, where group 1 has zero households receiving asset transfers, group 2 has 1 household receiving transfers, and so on. Then the marginal effect of having $g+1$ households treated rather than only g households could be estimated by comparing outcomes for households or networks in groups $g+1$ and g . In Subsections 4.1 and 4.4 I discuss two different approaches I take to handle this, one exploiting the available randomisation and the other using the non-linearity of the relationship being tested for to rule out alternative explanations.

3.3 Defining Risk-Sharing Networks

The predictions of the model concern behaviour at the risk-sharing network level. Early work on informal risk sharing assumed that the relevant group in which risk sharing takes place is the village (Townsend, 1994). Implicitly this assumes there are frictions preventing risk sharing with households outside the village, and that within the village all households belong to a common risk-sharing pool. Recent evidence suggests that in some context risk-sharing networks might be smaller than the village. Using data from Indian villages, Mazzocco and Saini (2012) and Munshi and Rosenzweig (2016) both find that caste groups within a village are the appropriate risk-sharing network i.e. there are important frictions preventing risk sharing across caste lines within a village.

To determine the appropriate group for risk sharing, I use two different approaches: (1) using information on reported links; (2) using Townsend tests (Townsend, 1994). These two approaches both support the same definition of the risk-sharing network for eligible households.¹⁵ To the extent that this is an imperfect proxy for the true risk-sharing network, it will introduce some noise into my later work. As a robustness check, in my empirical test for a poverty trap I will show that qualitatively similar results would be found if the entire village were used, or only eligible households are used.

The median (mean) risk-sharing network is composed of 50 (51.7) households, has an aggregate income of \$53,600 (\$57,700), and receives a capital injection of \$2,060 (\$2,740).

3.4 Verifying Model Assumptions

I first verify that the context matches the modelling framework in five dimensions: (1) households have variable incomes; (2) household savings are small relative to income; (3) households engage in risk sharing; (4) households have potentially productive lumpy investments available; (5) risk-sharing networks have the resources needed to be able to invest.

1. *Households have variable incomes.* Using only the time series variation for households in the poor risk-sharing networks in control villages, the median coefficient of variation is .35 (mean is .41).

¹⁵For details see Appendix C.

2. *Household savings are small relative to income.* At baseline, the median household in the villages covered by my data has cash savings totalling .5% of their income. Including also jewellery and ceremonial clothing, this rises to 3.6%. Savings, even including jewellery, are therefore an order of magnitude smaller than income shocks, and so have limited scope for providing consumption smoothing.
3. *Households engage in risk sharing.* Households were asked whether they suffered a crisis in the last year, and if so how they coped with it. They may report multiple methods. To avoid confounding with the asset transfer program, I consider only households in control villages, and I pool their responses over the three waves. In each wave, about half of all households report suffering some kind of crisis. Of those who report suffering a crisis, 38% receive loans or transfers from other households to provide smoothing. 50% of households also use their own savings to provide some smoothing, although as noted these savings are small relative to the size of shocks households face. 36% of households also report reducing consumption during a crisis. Taken together, these results indicate that households use risk-sharing transfers as an important channel of consumption smoothing, but consumption smoothing is incomplete.
4. *Households have potentially productive lumpy investments available.* [Bandiera et al. \(2017\)](#) document that for these data that the mean internal rate of return on cows is 22%. In 2007 USD PPP terms, one cow costs around \$257. This is 18% of median household income in a village, and 29% of median household income among the households eligible for the program.
5. *Risk-sharing networks have the resources needed to be able to invest.* Figure 2 shows the distribution of aggregate wealth holdings across risk-sharing networks. Wealth is broken down into a number of categories, and the cost of a cow is marked on the figure. This gets to heart of the puzzle this paper seeks to explain: more than 75% of risk-sharing networks have available to them enough *cash*, let alone other assets, needed to be able to invest in cows. Yet despite this, and the high returns, these savings are not pooled across households to purchase cows.

4 Empirical Evidence

First I provide evidence of a network-level poverty trap. I show asset transfers stimulate additional investment, but only in networks where more than seven households receive cows. Next I demonstrate that 80% of this additional investment is done by households who do not receive cows, but are in the same risk-sharing network as recipients, providing direct evidence of within-network spillovers. I then examine household budgets to understand in more detail how investment is funded. I show that networks above the seven household threshold spend less of the additional income from cows on non-food items, they save more, and they engage in

more internal lending which facilitates investment. Finally, I consider three leading alternative explanations for a network-level poverty trap, and show that their predictions are not borne out in my empirical context.

4.1 Evidence for a Network-Level Poverty Trap

The prediction of the model is that there should exist some aggregate income threshold such that (i) below the threshold the network is in a poverty trap and we see no investment, (ii) above the threshold we see investment taking place, with investment increasing in the value of transfers. To test this, I use variation in the amount of capital (and hence, implicitly, income) provided *at the network level* by the asset transfer program. As described above, the program provided the same value of assets to all eligible households in treated villages. However, there is variation in the number of eligible households within a village. Hence I compare risk-sharing networks with the same number of eligible households across treatment and control villages.

Non-parametric Relationship Between Investment and Capital Injection

I begin by studying the non-parametric relationship between the value of the capital injection provided by the program to network v , $\Delta k_{v,2009}$, and additional investment in cows between 2009 and 2011, $\Delta k_{v,2011}$.¹⁶ Both are measured at the network-level, in 2007 USD PPP. Precisely I estimate the following local mean regression:

$$\Delta k_{v,2011} = m(\Delta k_{v,2009}) + \epsilon_{v,2011} \quad (1)$$

separately for treated and control networks, where $m(\cdot)$ is unknown and estimated using a Nadaraya-Watson kernel-weighted local mean estimator.

Figure 3a plots the conditional mean, and 95% confidence interval. It can be seen that investment in further cow ownership is close to zero and does not vary with the value of the capital injection up to a value of around \$4,000. When more than this level of capital was provided by the program, there appears to be an increasing (and approximately linear) relationship between the capital provided and the amount of additional investment takes place. This is precisely the relationship predicted by the model.

As discussed earlier, the ideal design for my context would be to have experimental variation in the intensity of treatment, as measured by the number of households that receive transfers. Here the number of transfers to a village is endogenous, conditional on being in a village that is treated. To rule out the observed effect being driven by underlying heterogeneity, I plot the same relationship for the control sample.¹⁷ Figure 3b plots additional investment between 2009 and 2011 against the value of the capital injection that *would have been* provided had the risk-

¹⁶Since the program provided some consumption support and training between 2009, I do not try to disentangle what occurs between 2007 and 2009. Instead I study the additional investment that takes place after 2009, by when the program is no longer active and no additional support is being provided.

¹⁷In Subsection 4.4 I discuss the possibility that the program interacts with some underlying heterogeneity, so that the investment decision simply reflects some other variable responding non-linearly.

sharing networks been in treated villages. In the absence of actual asset transfers, investment is zero on average, and does not vary with the placebo value of capital injection.

Testing Formally for a Threshold Effect

The non-parametric results suggest that, among the treated networks, there exists a threshold value of aggregate capital injection needed to spur additional investments by the network. To test this relationship formally, I estimate for the treated sample a regression of the form:

$$\begin{aligned} \Delta k_{v,2011} = & \alpha_1 + \delta_1 \Delta k_{v,2009} \cdot \mathbf{1}\{\Delta k_{v,2009} < \Delta k^*\} \\ & + \delta_2 \Delta k_{v,2009} \cdot \mathbf{1}\{\Delta k_{v,2009} \geq \Delta k^*\} + \gamma_1 \mathbf{X}_{v,2009} + \gamma_2 \mathbf{X}_{v,2007} + \epsilon_{v,2011} \end{aligned} \quad (2)$$

where again $\Delta k_{v,2011}$ is the increase in cow investment by the network as a whole between 2009 and 2011, $\Delta k_{v,2009}$ is the value of the asset transfers provided to the network by the program, Δk^* is a proposed threshold value of asset transfers, and \mathbf{X} is a vector of controls. Note that since the asset transfer by the program takes place in 2007, $\Delta k_{v,2011}$ does *not* include the initial injection. All monetary values are in 2007 USD at PPP exchange rates.

This specification captures the idea that there is some threshold level of asset transfers, Δk^* , needed to push a network out of the poverty trap. Below this threshold there should be no additional investment, $\alpha_1 = 0$ and $\delta_1 = 0$, and above this threshold we should see additional investment increasing in the value of capital injection, $\delta_2 > 0$. Whilst the model does not predict the functional form for how additional investment varies with the capital injection above the threshold, the estimated non-parametric relationship, Figure 3a, suggests that at least over the support of my data, linearity does not seem unduly restrictive.

Since the threshold, Δk^* , is unknown, I use an iterative regression procedure designed to test for a structural break (a change in the slope of the relationship) with unknown break point. This involves running a sequence of such regressions over a prespecified range of possible values for Δk^* , and then testing for significance of the test statistic against an adjusted distribution, to account for the repeated testing.

I use two different statistics, both the Quandt Likelihood Ratio test (QLR; [Quandt, 1960](#)) and the Hansen test ([Hansen, 1999](#)). The former selects as the threshold location the point which maximises the absolute value of the t-statistic on δ_2 , and tests significance of the threshold using an adjusted limiting distribution ([Andrews, 1993](#)). The latter uses a criterion based on the residual sum of squares, so accounts more directly for the relative explanatory power of the regression as a whole.

Table 2 shows the results of the QLR test using different control variables, \mathbf{X} . In all cases the most likely location for a threshold is at \$3,500 of asset transfers, equal to 6.5% of income for the median network, and close to the level suggested by visual inspection of Figure 3a.¹⁸

¹⁸The discrepancy between the visual estimate and the formal method is caused simply by the non-parametric smoothing: by using observations below the threshold when estimating the local mean above the threshold, the figure makes the threshold look later and less sharp than it is.

This is equivalent to treating 14% of households in the median network. Figure 4 shows non-parametrically the relationship between the value of capital provided to the network by the the program (in 2007), and the additional investment by the network between 2009 and 2011, splitting the non-parametric plot at \$3,500. This makes the relationship clear to see.

Testing whether this potential threshold is itself statistically significant, I can reject at the 5% level the hypothesis that there is no threshold effect. This is also true with additional controls. When district fixed effects are included the qualitative patterns remain unchanged but the estimates become noisier. Studying the regression results, one can see that below the threshold the level of investment is close to zero, and above the threshold it is increasing, consistent with the model predictions.¹⁹

For the Hansen test, I estimate the same regression specifications as for Columns (1) and (3) in Table 2 above. For each possible threshold I calculate the residual sum of squares (RSS). I select among the regressions the one (or set) with the lowest RSS. The corresponding threshold in that regression is the estimated location for the threshold using this method. To test whether the threshold is significant, I construct the Hansen statistic. This is, at any possible threshold, the difference between the RSS at that threshold and the minimum RSS from all thresholds considered, divided by the minimum RSS and multiplied by the sample size. This is necessarily equal to zero at the proposed threshold. If it is below .05 at any other tested threshold, then that threshold cannot be rejected as a possible location for the threshold.

Figure D1 shows the value of the likelihood ratio statistic from running the Hansen test for possible thresholds between \$2,000 and \$5,000, at intervals of \$100.²⁰ From this it is clear that the most likely location of the threshold is between \$3,700 and \$4,100, or 6.9-7.6% of income in the median network, close to the estimate of \$3,500 from using the Quandt Likelihood Ratio approach. Henceforth I use \$3,500 as the estimate of the threshold location, but my results are robust to choosing instead a point in [\$3,700, \$4,100].

Impact of Capital Injection on Investment

Having identified the location of the threshold, I then estimate the following regression on the sample including both treated and control variables:

$$\begin{aligned} \Delta k_{v,2011} = & \alpha_0 + \alpha_1 T_v + \beta_1 \Delta k_{v,2009} + \delta_1 T_v \Delta k_{v,2009} \cdot \mathbf{1}\{\Delta k_{v,2009} < 3500\} \\ & + \delta_2 \Delta T_v k_{v,2009} \cdot \mathbf{1}\{\Delta k_{v,2009} \geq 3500\} + \gamma_1 \mathbf{X}_{v,2009} + \gamma_2 \mathbf{X}_{v,2007} + \epsilon_{v,2011} \end{aligned} \quad (3)$$

where T_v is an indicator for village treatment status, and all other variables are as before. Now the specification makes use of the exogeneity due to randomisation of villages: the coefficients

¹⁹Note however that since this regression is chosen using the iterative procedure described above, it would not be correct to use the standard errors provided directly for inference.

²⁰Note that since the variation in the aggregate value of the capital injection comes only from variation in the number of treated households, there are only data points at intervals of \$515 (the value of the asset transfer to one household). I show the test using intervals of \$100 just to make clear the region of possible values that this test cannot reject.

α_1 , δ_1 , δ_2 , are identified from the difference between the treated and control risk-sharing networks. It is important to note that the standard errors estimated here do not account for the prior estimation of the threshold location.

Table 3 shows the results of this estimation. Similar to what was seen non-parametrically, additional investment is flat with respect to the capital injection below the threshold \$3,500 ($\delta_1 = 0$), and increasing after the threshold ($\delta_2 > 0$). The estimated coefficient δ_2 suggests that, after the first \$3,500 worth of asset transfers to a risk-sharing network, every additional \$500 generates a further \$750 in investment, although the confidence intervals are wide.

4.2 Who Does the Investment?

Both to understand *who* does the additional investment, and to check for robustness to the definition of the risk-sharing network, I re-estimate the relationship for the treated sample, using different levels of aggregation. I use two alternative definitions of the network. First, I include the entire village. If the results are driven purely by some general equilibrium type effects, we might expect to see additional investment by other households in the village who are not part of the risk-sharing network. At the other extreme, I include only eligible households. This allows me to see whether the additional investment is done only by those receiving cows from the program, which might suggest there is just specialisation or learning taking place.

Figure 5a estimates the relationship in Equation 1 where all households in the village are assumed to belong to the risk-sharing network, rather than only those in the lower two wealth classes. Three points are of note. First, the general shape of the relationship remains the same, and the apparent threshold is at the same location. Second, the entire graph has been translated upwards by \$1,000. This implies the richer households in these villages were doing some investment, but this is not responsive to the amount of capital injected, consistent with them not being part of the risk-sharing network. Third, the confidence intervals are now wider. If the richer households are not part of the risk-sharing network of the eligible households, then including them should just add noise to the estimated effects, as can be seen.

Figure 5b estimates the relationship including only the eligible households. Again the same shape of relationship is visible, with a similar apparent location for the threshold. However, the slope of the relationship above the threshold is much flatter, indicating that *other households* were also investing at these higher levels of capital injection, but not investing at low levels. This apparent spillover – with additional investment by ineligible poor households depending on the number of eligible households – is direct evidence that consideration of the risk-sharing network is important when studying the impact of this type of program. It also helps rule out explanations based on household-level poverty traps: if these were the only explanation for the initial lack of investment, then households which don't benefit from the program should not be responding.²¹

²¹In Subsection 4.4 I rule out other alternative explanations, including the possibility that other households' investment can be explained by general equilibrium effects.

Table 3 shows the results quantitatively using these alternative aggregations. The slope when the network is assumed to comprise the whole village is the same as when it is the ‘all poor’ group, but standard errors are larger, consistent with the idea that the richer households are not part of the true network so including them just adds noise. The common slope implies the program does not lead to non-network households doing additional investing. When only eligibles are considered, the slope is one-fifth that for the ‘all poor’ group, so 80% of the additional investment when the capital injection exceeds the threshold comes from outside the eligibles.²²

4.3 How is investment funded?

The model predicts that there exists a threshold level of income, above which investment is increasing, below which investment is flat. Another way to test this relationship is through the budget constraint. Below the poverty trap threshold, as network incomes rise, the additional income should be consumed, since households are unable to invest. Above the threshold, consumption should rise less steeply with income, as some share of additional income beyond this point is spent on investing. Additionally, above the threshold there should be more (gross) borrowing, as some households can now borrow to invest.

To test this, Table 4 shows how the capital injection affects household income, expenditure, savings and borrowing either side of the threshold. For each of various outcomes Y , the following specification is estimated:

$$\begin{aligned} \Delta Y_{i,v,2011} = & \alpha_0 + \alpha_1 T_v + \beta_1 \Delta k_{v,2009} + \delta_1 T_v \Delta k_{v,2009} \cdot \mathbf{1}\{\Delta k_{v,2009} < 3500\} \\ & + \delta_2 \Delta T_v k_{v,2009} \cdot \mathbf{1}\{\Delta k_{v,2009} \geq 3500\} + \gamma_1 \mathbf{X}_{i,v,2009} + \gamma_2 \mathbf{X}_{i,v,2007} + \epsilon_{i,v,2011} \end{aligned} \quad (4)$$

This is the analogue of Equation 3 but with outcomes, Y , and control variables, X , measured at the individual level. Columns (1)-(3) show that income from livestock is higher, and income from maid work is lower, the more capital was transferred to the risk-sharing network. This effect does not change around the threshold: the effect of providing cows is to increase income, with no non-linearity in returns. Columns (4)-(6) provide some evidence that non-food expenditure also does not increase above the threshold, while food expenditure increases linearly, so that per capita equalised expenditure (and also total expenditure) respond less to additional capital above the threshold as predicted. Instead one can see that savings appear to increase faster with additional capital transfers above the threshold (Column 7). One can also see that borrowing (within the village) is certainly higher, with additional loans only being taken on in risk-sharing networks above the poverty trap threshold (Column 8). This borrowing, along with potentially some additional saving and lower non-food expenditure, is then the main source of funds for the additional investment observed. This is in line with the proposed mechanism: above the

²²Given how much more precise the estimates are among the eligible only group, it also suggests that even the ‘all poor’ group I construct as a proxy for the risk-sharing network might be too large, containing some irrelevant households and making the estimates less precise.

threshold it becomes possible for borrowing to take place. Estimating separately by recipient status, one can see that this borrowing is concentrated among non-recipients of the program, who are now investing.

4.4 Evidence for Alternative Explanations

I consider two alternative explanations. First, I investigate whether the network-level poverty trap could be generated by some form of increasing returns to cows at the network level. Second, I study whether asset transfers caused a non-linear effect via some other channel. Specifically I consider general equilibrium effects and aspirations.²³

Increasing Returns

The classic model of a group-level poverty trap is the ‘big push model’ of [Rosenstein-Rodan \(1943\)](#) (later formalised by [Murphy et al., 1989](#)). The key mechanism underlying it is that the return to investment is *increasing* in the number of other agents who engage in investment. Whilst the original model is motivated by concerns about industrial structure, which are not relevant in a village economy, network-level increasing returns might still exist for a number of other reasons. One possibility is group level fixed costs. For example, the price of milk may be higher (or even just more stable) in nearby markets than within the village, but there is a fixed cost of travel so that it is only profitable if enough milk is being taken. Another explanation might be that there is learning across households: the more households that engage in livestock rearing, the more sources of information and advice there are, helping to better look after the cows. Such non-linear effects of social learning about investment are documented by [Bandiera and Rasul \(2006\)](#).

These mechanisms are distinct from my model. In my model households are unable to invest due to constraints, namely the inability to commit to future transfers, but the returns from investment are independent of the number of investments. A direct test of these increasing-returns based alternatives, is to see whether the return on cows is increasing in the number of households that received cows.

To test this, I estimate non-parametrically, again using a kernel weighted local mean smoother, the mean return on a cow against the number of eligible.²⁴ Since the value of capital provided to an eligible household is fixed, the aggregate capital injection maps linearly to the number households that receive cows.

From Figure D2 it can be seen that the mean return on cows appears to be declining in the number of cows transferred, at low number of transfers, and then to be flat and stable. These results are inconsistent with a story of increasing returns, ruling out the possibility that the observed poverty trap could be driven by network level increasing returns.

²³Table D1 shows for a variety of other variables that there is no non-linearity present.

²⁴Increasing returns mean that the return on the *marginal* cow is higher than on the previous cow, in which case the mean return should also be rising.

Prices and Aspirations

If real returns to cows are unchanged, an alternative explanation for the increase in investment might be that some other channel is activated once a sufficiently large number of households receive transfers. This could generate a threshold effect.

The first possibility is that general equilibrium effects might occur in some non-linear way. An immediate piece of evidence that suggests this is unlikely to be the case is that the threshold relationship documented is in terms of the *aggregate* number of households/value of capital provided, consistent with the model. General equilibrium effects, by contrast, might depend instead on the *share* of households treated. Figure D3 shows the non-parametric estimate of Equation 1 but where the independent variable of interest is the share of poor households who are treated. Plotted against the share treated, there does not appear to be any clear relationship. This provides additional evidence against a model of aggregate demand spillovers, as in [Murphy et al. \(1989\)](#).

An alternative way to test for general equilibrium effects is to estimate how prices vary with the value of the capital injection. Whilst this is more direct test, it requires us to know in what markets to look, and to have good measures of prices in those markets. Three possible prices of interest are the price of milk, which is the output price for cow owners; the price of cows, which is the cost of additional investment; and the wage, which is both the source of income for investment and the opportunity cost of time spent looking after cows.

To test empirically whether either of these effects can explain the non-linearity in investment, I estimate Equation 4 replacing the dependent variable with $\text{Channel}_{v,2009}$. This measures the relevant price in 2009, at the time households are making investment decisions. Table 5 shows there is no effect of the program on the price of milk in 2009; a reduction in the price of cows on average, but with no threshold effect; and an increase in the average wage, again with no threshold effect. Hence none of these markets appear to be the channel through which any general equilibrium effects could be driving the threshold in aggregate investment.

A second possible source of non-linearity might be driven by changes in aspirations. The worry would be that there is a non-linear increase in the demand for cows as the number of neighbours owning cows rises. This could happen because households perceptions of the return on cows increases in the prevalence of ownership, or because households receive direct utility from cow ownership – beyond the financial returns – and this rises when ownership becomes more prevalent, as in a model of ‘Keeping up with the Joneses’.

Table 5 Column (4) shows the results of estimating Equation 4 with $\text{Channel}_{v,2009}$ measuring the change, between 2007 and 2009, in the share of ineligible poor households without livestock in 2009 who aspire to own livestock.²⁵ Having any household in the network receive cows raises aspirations. However, the effect is not increasing in the number of households receiving cows, nor is there any threshold effect.

²⁵Eligible households in treated villages are automatically excluded from the sample because they own cows in 2009. I exclude them from the sample in control villages to avoid composition bias.

5 Conclusion

Poor households often do not undertake profitable investments, even when they belong to networks which could pool resources to invest. This paper provides a novel explanation for this puzzle: informal risk sharing can crowd out investment. To show this, I extend the classic model of risk sharing with limited commitment (Ligon et al., 2002) to also allow for lumpy investment. I show that with this addition, the model generates a poverty trap at the level of the risk-sharing network: unless aggregate income is above some threshold, the network will never be able to invest. The key insight is that once a household invests, it has less need for insurance and is more willing to walk away from the risk-sharing arrangement. This limits the investor's ability to credibly promise future transfers, so its risk-sharing partners demand transfers today, limiting investment. Hence, in the absence of institutions enforcing joint property rights, a network can be in a poverty trap despite having the resources to be able to collectively invest.

To provide evidence for this mechanism I use data from a long term, large scale randomised controlled trial in Bangladesh. The program randomised 1,400 villages into treatment or control status, and provided assets to the poorest households in half of these villages. I exploit variation in the aggregate level of transfers provided to risk-sharing networks to show evidence for a network-level poverty trap. Precisely, I show empirically a threshold level of aggregate capital provision needed for the program to generate further investment: networks that received more than \$3,500 were 'pushed' out of the trap. 80% of this additional investment is done by non-recipients of the program, and it is financed by within-network lending.

My findings have important implications for policy. The asset transfer program from which my data were drawn has now been expanded to more than half a million households in Bangladesh, and similar programs have begun in more than 40 countries worldwide. This expansion is motivated by the consistent and robust results that these programs create sustained income growth (Bandiera et al., 2017). My results explain *why* we see these large and long run effects, and crucially also how these programs can be further improved. If the program targeting took into account not only household characteristics, but also network characteristics and the size of the aggregate transfer being provided, more networks could be pushed out of the poverty trap, and set on a path of sustained growth.

An important direction for future research is to quantify the trade-off faced by designers of such programs between reducing poverty and growing incomes. Using the reduced form estimates of the effect of asset transfers, a budget-neutral redistribution of asset transfers between networks in my data could generate additional investment of 44%. However, this would be achieved by reducing transfers to inframarginal networks, which are far from the poverty trap threshold, and providing them instead to marginal networks just below the threshold. Whilst this increases the number of networks pushed out of the trap, it also increases inequality *across* networks, reducing consumption in those which lose transfers. Directly estimating the parameters of the model would allow the study of the welfare gains from alternative targeting policies, taking account of this trade-off, and maximising the gains from these promising new interventions.

References

- ÁBRAHÁM, A. AND E. CÁRCELES-POVEDA (2009): “Competitive Equilibria with Production and Limited Commitment,” *mimeo*.
- ÁBRAHÁM, A. AND S. LACZÓ (forthcoming): “Efficient Risk Sharing with Limited Commitment and Storage,” *Review of Economic Studies*.
- ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58, 1041–1063.
- AGHION, P. AND P. BOLTON (1997): “A Theory of Trickle-Down Growth and Development,” *Review of Economic Studies*, 64, 151–172.
- ALVAREZ, F. AND U. J. JERMANN (2000): “Efficiency, Equilibrium, and Asset Pricing with Risk of Default,” *Econometrica*, 68, 775–797.
- AMBRUS, A., A. G. CHANDRASEKHAR, AND M. ELLIOTT (2015): “Social Investments, Informal Risk Sharing, and Inequality,” Tech. Rep. 20669, National Bureau of Economic Research.
- AMBRUS, A., M. MOBIUS, AND A. SZEIDL (2014): “Consumption Risk-Sharing in Social Networks,” *American Economic Review*, 104, 149–182.
- ANDREWS, D. W. K. (1993): “Tests for Parameter Instability and Structural Change with Unknown Change Point,” *Econometrica*, 61, 821–56.
- ANGELUCCI, M., G. DE GIORGI, AND I. RASUL (2018): “Consumption and Investment in Resource Pooling Family Networks,” *Economic Journal*.
- ATTANASIO, O. AND N. PAVONI (2011): “Risk Sharing in Private Information Models with Asset Accumulation: Explaining the Excess Smoothness of Consumption,” *Econometrica*, 79, 1027–1068.
- ATTANASIO, O. AND V. RIOS-RULL (2000): “Consumption smoothing in island economies: Can public insurance reduce welfare?” *European Economic Review*, 44, 1225–1258.
- BANDIERA, O., R. BURGESS, N. DAS, S. GULESCI, I. RASUL, AND M. SULAIMAN (2017): “Labor markets and poverty in village economies,” *Quarterly Journal of Economics*, 132, 811–870.
- BANDIERA, O. AND I. RASUL (2006): “Social Networks and Technology Adoption in Northern Mozambique,” *Economic Journal*, 116, 869–902.
- BANERJEE, A. V., E. DUFLO, N. GOLDBERG, R. OSEI, W. PARIENTÉ, J. SHAPIRO, B. THUYSBAERT, AND C. UDRY (2015): “A multifaceted program causes lasting progress for the very poor: Evidence from six countries,” *Science*, 348, 1260799.

- BANERJEE, A. V. AND S. MULLAINATHAN (2010): “The Shape of Temptation: Implications for the economic lives of the poor,” Tech. Rep. 15973, National Bureau of Economic Research.
- BANERJEE, A. V. AND A. F. NEWMAN (1993): “Occupational Choice and the Process of Development,” *Journal of Political Economy*, 101, 274–298.
- BERNHEIM, D., D. RAY, AND Ş. YELTEKIN (2015): “Self-Control and Poverty,” *Econometrica*, 83, 1877–1911.
- DE MEL, S., D. MCKENZIE, AND C. WOODRUFF (2008): “Returns to Capital in Microenterprises: Evidence from a Field Experiment,” *The Quarterly Journal of Economics*, 123, 1329–1372.
- DUPAS, P. AND J. ROBINSON (2013): “Why Don’t The Poor Save More? Evidence from health savings experiments,” *American Economic Review*, 103, 1138–1171.
- FAFCHAMPS, M. AND J. PENDER (1997): “Precautionary Saving Credit Constraints And Irreversible Investment: Theory And Evidence From Semiarid India,” *Journal of Business and Economic Statistics*, 15, 180–194.
- FOSTER, A. AND M. R. ROSENZWEIG (1994): “A Test for Moral Hazard in the Labor Market: Contractual Arrangements, Effort, and Health,” *Review of Economics and Statistics*, 76, 213–27.
- GHATAK, M. (2015): “Theories of Poverty Traps and Anti-Poverty Policies,” *World Bank Economic Review, Papers and Proceedings of the ABCDE*, 29, S77–S105.
- HANSEN, B. E. (1999): “Threshold Effects in Non-Dynamic Panels: Estimation, Testing and Inference,” *Journal of Econometrics*, 93, 345–368.
- IKEGAMI, M., M. R. CARTER, C. B. BARRETT, AND S. A. JANZEN (2016): “Poverty Traps and the Social Protection Paradox,” Tech. Rep. 22714, National Bureau of Economic Research.
- JAKIELA, P. AND O. OZIER (2016): “Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies,” *Review of Economic Studies*, 83, 231–268.
- KARLAN, D., R. OSEI, I. OSEI-AKOTO, AND C. UDRY (2014): “Agricultural Decisions after Relaxing Credit and Risk Constraints,” *Quarterly Journal of Economics*, 129, 597–652.
- KOCHERLAKOTA, N. (1996): “Implications of Efficient Risk Sharing without Commitment,” *Review of Economic Studies*, 63, 595–609.
- KRAAY, A. AND D. MCKENZIE (2014): “Do Poverty Traps Exist? Assessing the Evidence,” *Journal of Economic Perspectives*, 28, 127–148.

- LIGON, E., J. P. THOMAS, AND T. WORRALL (2000): “Mutual Insurance, Individual Savings, and Limited Commitment,” *Review of Economic Dynamics*, 3, 216–246.
- (2002): “Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies,” *Review of Economics Studies*, 69, 209–244.
- MATSUYAMA, K. (2002): “Explaining Diversity: Symmetry-Breaking in Complementarity Games,” *American Economic Review*, 92, 241–246.
- (2004): “Financial Market Globalization, Symmetry-Breaking, and Endogenous Inequality of Nations,” *Econometrica*, 72, 853–884.
- (2011): “Imperfect Credit Markets, Household Wealth Distribution, and Development,” *Annual Review of Economics*, 3, 339–362.
- MAZZOCCO, M. AND S. SAINI (2012): “Testing Efficient Risk Sharing with Heterogeneous Risk Preferences,” *American Economic Review*, 102, 428–468.
- MILGROM, P. AND C. SHANNON (1994): “Monotone Comparative Statics,” *Econometrica*, 62, 157–180.
- MORDUCH, J., S. RAVI, AND J. BAUCHET (2015): “Failure vs Displacement: Why an innovative anti-poverty program showed no net impact,” *Journal of Development Economics*, 116, 1–16.
- MORTEN, M. (forthcoming): “Temporary Migration and Endogenous Risk Sharing in Village India,” *Journal of Political Economy*.
- MUNSHI, K. AND M. ROSENZWEIG (2016): “Networks and Misallocation: Insurance, Migration and the Rural-Urban Wage Gap,” *American Economic Review*, 106, 46–98.
- MURPHY, K. M., A. SHLEIFER, AND R. W. VISHNY (1989): “Industrialization and the Big Push,” *Journal of Political Economy*, 97, 1003–1026.
- NURKSE, R. (1953): *Problems of Capital Formation in Underdeveloped Countries*, Oxford University Press.
- PAVONI, N. AND G. L. VIOLANTE (2007): “Optimal Welfare-To-Work Programs,” *Review of Economic Studies*, 74, 283–318.
- PHELAN, C. AND R. M. TOWNSEND (1991): “Multi-Period, Information-Constrained Optima,” *Review of Economic Studies*, 58, 853–881.
- PRESCOTT, E. AND R. M. TOWNSEND (1984a): “General Competitive Analysis in an Economy with Private Information,” *International Economic Review*, 25, 1–20.

- (1984b): “Pareto Optima and Competitive Equilibria with Adverse Selection Moral Hazard,” *Econometrica*, 52, 21–46.
- QUANDT, R. E. (1960): “Tests of the Hypothesis that a Linear Regression Obeys Two Separate Regimes,” *Journal of the American Statistical Association*, 55, 324–330.
- ROSENSTEIN-RODAN, P. N. (1943): “Problems of Industrialization of Eastern and Southeastern Europe,” *Economic Journal*, 53, 202–211.
- ROSENZWEIG, M. R. AND K. I. WOLPIN (1993): “Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investment in Bullocks in India,” *Journal of Political Economy*, 101, 223–244.
- SHABAN, R. A. (1987): “Testing Between Competing Models of Sharecropping,” *Journal of Political Economy*, 95, 893–920.
- SPEAR, S. E. AND S. SRIVASTAVA (1987): “On Repeated Moral Hazard with Discounting,” *Review of Economic Studies*, 54, 599–617.
- THOMAS, J. AND T. WORRALL (2018): “Dynamic Relational Contracts under Complete Information,” *Journal of Economic Theory*, 175, 624–651.
- TOPKIS, D. M. (1978): “Minimizing a Submodular Function on a Lattice,” *Operations Research*, 26, 305–321.
- TOWNSEND, R. (1994): “Risk and Insurance in Village India,” *Econometrica*, 62, 539–591.
- UDRY, C. (1994): “Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria,” *Review of Economic Studies*, 61, 495–526.

Table 1: Household Characteristics, by Wealth Grouping

Sample: All villages, baseline

Sample of households:

	(A) Eligible		(B) Ineligible poor		(C) Non-poor	
	Mean (1)	Std Dev (2)	Mean (3)	Std Dev (4)	Mean (3)	Std Dev (4)
Household size	3.29	1.70	4.09	1.61	4.74	1.76
Total income	960	690	1480	1040	2890	2440
Total consumption	1570	880	2200	1250	3740	2500
below poverty line (\$1.25) [Yes=1]	.46	.50	.39	.49	.21	.41
Owens cows [Yes=1]	.06	.24	.28	.45	.63	.48
Population share	.06		.52		.42	
Observations	7,111		13,704		6,162	

Notes: All statistics are constructed using baseline (pre-program) data for all households in both treated and control villages across the full sample. Observations are at the household level. Wealth classes are determined using a Participatory Rural Assessment, which classifies households into four or five wealth classes. Additional criteria are used to determine which households in the poorest wealth class are 'Eligible' for treatment (Panel A). 'Ineligible Poor' (Panel B) includes all households in the bottom two wealth classes (bottom three when five wealth ranks were used) who were not eligible for the program, while 'Non-poor' (Panel C) includes all households in the top two wealth classes. All financial measures are in 2007 USD terms, converted to dollars using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007.

Table 2: Location of the Poverty Trap Threshold

Dependent Variable: Increase in Total Cow Assets between 2009 and 2011 (2007 USD PPP)
 Standard errors (in parentheses) clustered at district level

	(1)	(2)	(3)
Slope below threshold (δ_1)	-.057 (.430)	-.087 (.454)	-.217 (.361)
Slope above threshold (δ_2)	1.57 (.55)	1.86 (.77)	2.68 (1.49)
Optimal Threshold Level	3500	3500	3500
F-statistic	5.45	5.89	3.24
5% critical value for F-statistic	5.24	5.24	5.24
Controls	No	Yes	Yes
District Fixed Effects	No	No	Yes
Observations (Clusters)	689 (13)	689 (13)	689 (13)

Notes: The sample comprises low wealth households in treated villages across the full sample. Observations are at the risk-sharing network level. Within a village, a risk-sharing network is the set of low wealth households: those in the lowest two (three) wealth classes, when there are four (five) possible wealth classes. Wealth classes are determined using a Participatory Rural Assessment. Data were collected using a stratified random sampling scheme, so weights are used to aggregate household data to the risk-sharing network level. All financial variables are in 2007 USD terms. Where necessary, they are first deflated to 2007 terms using the Bangladesh central bank CPI measure of inflation. They are then converted to dollars using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. A thin tail of networks receive more than \$8000 worth of assets. Since the density on this part of the support is low (fewer than five networks for any value of asset transfers), I trim these networks (2% of the sample). The outcome measures the increase in the value of cow assets between 2009 and 2011. The table shows regression of the outcome on the aggregate value of capital provided to the risk-sharing network in 2007. Column (1) is a regression of the outcome on a constant, and the value of aggregate asset transfers interacted with a threshold dummy (allowing the slope to vary at this point). Column (2) allows for additional controls: lagged income and asset variables, and network size. Column (3) additionally allows for district level fixed effects. These specifications are run sequentially at different values of the threshold, varying the threshold between \$2000 and \$5000, at intervals of \$500. I provide the results for the threshold among these which produced a regression with the highest F-statistic when tested against the null of no threshold. The value for this threshold, the estimated F-statistic, and the 5% critical value for this statistic (which corrects for the repeated testing, see Andrews, 1993) are provided at the bottom of the table.

Table 3: Using the Program Randomisation to Test the Poverty Trap

Dependent Variable: Increase in Total Cow Assets between 2009 and 2011 (2007 USD PPP)
 Standard errors (in parentheses) clustered at district level

Risk-sharing Network:	(A) All Poor (1)	(B) Whole Village (2)	(C) Eligibles Only (3)	(4)
Slope below threshold for Treated	-0.242 (.515)	-0.110 (.516)	-0.785 (1.112)	-0.076 (.087)
Slope above threshold for Treated	1.801 (.816)	1.435 (.813)	1.517 (1.711)	0.316 (.121)
Treated	792 (1058)	859 (1063)	1914 (2299)	-107 (187)
Controls	No	Yes	Yes	Yes
District Fixed Effects	No	Yes	Yes	Yes
Observations (Clusters)	1360 (13)	1360 (13)	1360 (13)	1260 (13)

Notes: Observations are at the risk-sharing network level. The network in Panel A comprises low wealth households across the full sample. A household is low wealth if it is in the lowest two (three) wealth classes, when there are four (five) possible wealth classes. Wealth classes are determined using a Participatory Rural Assessment. The network in Panel B comprises all households across the full sample. The network in Panel C comprises eligible households across the full sample. The sample size here is smaller because 100 villages which have no eligibles are now automatically excluded. Data were collected using a stratified random sampling scheme, so weights are used to aggregate household data to the risk-sharing network level. All financial variables are in 2007 USD terms. Where necessary, they are first deflated to 2007 terms using the Bangladesh central bank CPI measure of inflation. They are then converted to dollars using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. A thin tail of networks receive more than \$8000 worth of assets. Since the density on this part of the support is low (fewer than five networks for any value of asset transfers), I trim these networks (2% of the sample). The outcome measures the increase in the value of cow assets between 2009 and 2011. The table shows regression of the outcome on the aggregate value of capital provided to the risk-sharing network in 2007. Column (1) is a regression of the outcome on a constant, a treatment dummy, the value of aggregate transfers that would be provided if a network is treated, and the value of aggregate transfers interacted with a treatment dummy and an indicator for whether aggregate transfers exceed \$3,500. Columns (2)-(4) allow for additional controls – lagged values of total income, total savings, and total value of cows – and district level fixed effects.

Table 4: How Does Behaviour Change Across The Poverty Trap Threshold?

Outcome:	Livestock Income	Agricultural Income	Maid Income	Food Expenditure	Non-food Expenditure	Per-capita Expenditure	Total Savings	Total Loans
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Slope below threshold for Treated	.013 (.004)	.000 (.000)	-.008 (.002)	.045 (.013)	.083 (.036)	.018 (.007)	.004 (.001)	.009 (.009)
Slope above threshold for Treated	.017 (.008)	-.001 (.000)	-.017 (.007)	.048 (.027)	-.009 (.045)	.006 (.010)	.011 (.004)	.107 (.039)
Equal slopes? (p-value)	.851	.342	.820	.310	.156	.264	.059	.020
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations (Clusters)	15682 (13)	15682 (13)	15682 (13)	15010 (13)	15008 (13)	14853 (13)	15682 (13)	15682 (13)

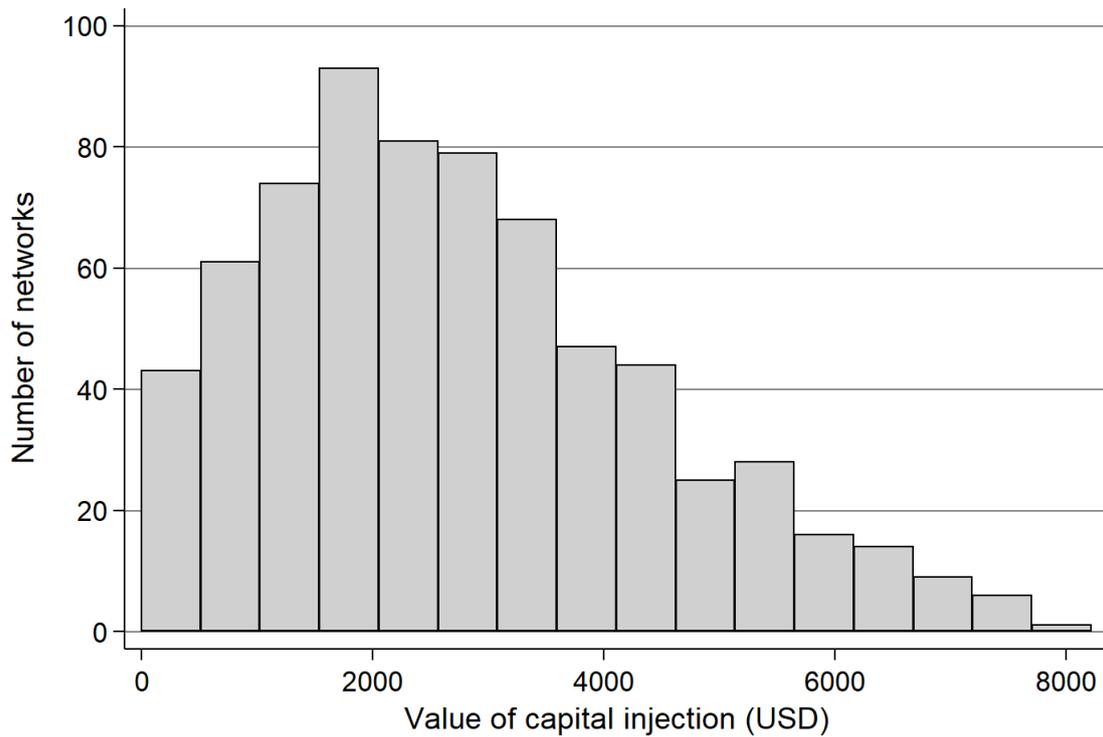
Notes: The sample comprises all low wealth households in all villages across the full sample in 2011: those in the lowest two (three) wealth classes, when there are four (five) possible wealth classes. Wealth classes are determined using a Participatory Rural Assessment. All financial variables are in 2007 USD terms. They are first deflated to 2007 terms using the Bangladesh central bank CPI measure of inflation. They are then converted to dollars using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. In Column (1) the outcome measures the livestock income of the main female respondent in the household. In Column (2) the outcome measures the agricultural daylabour income of the main female respondent in the household. In Column (3) the outcome measures the maid income of the main female respondent in the household. These three income sources are the main sources of incomes for females, as reported in [Bandiera et al. \(2017\)](#). In Column (4) the outcome measures the total food expenditure of the household. The variable is annualised from data on a three-day recall period. In Column (5) the outcome measures the total non-food expenditure of the household. The variable is annualised from data with different recall periods depending on the type of expenditure. In Column (6) the outcome measures total expenditure divided by the number of adult-equivalent household members. Children under 10 are given half the weight of children over 10 and adults. In Column (7) the outcome measures the total stock of saving in the household. In Column (8) the outcome measures the total stock of borrowing by the household. Controls are the lagged values of total income, total savings, and total value of cows.

Table 5: How Does The Capital Injection Affect Prices And Aspirations?

	Milk Price (1)	Cow Price (2)	Wage (3)	Aspirations (4)
Slope below threshold for Treated	-0.055 (.034)	-2.17 (12.84)	.327 (.351)	-.035 (.010)
Slope above threshold for Treated	.005 (.031)	6.58 (9.32)	.221 (.178)	-.024 (.015)
Treated	.110 (.060)	-7.97 (29.75)	.871 (.365)	.200 (.029)
Mean Baseline Level among Controls	1.28	475	6.32	.766
Controls	Yes	Yes	Yes	Yes
District Fixed Effects	Yes	Yes	Yes	Yes
Observations (Clusters)	1190 (13)	1190 (13)	1190 (13)	938 (13)

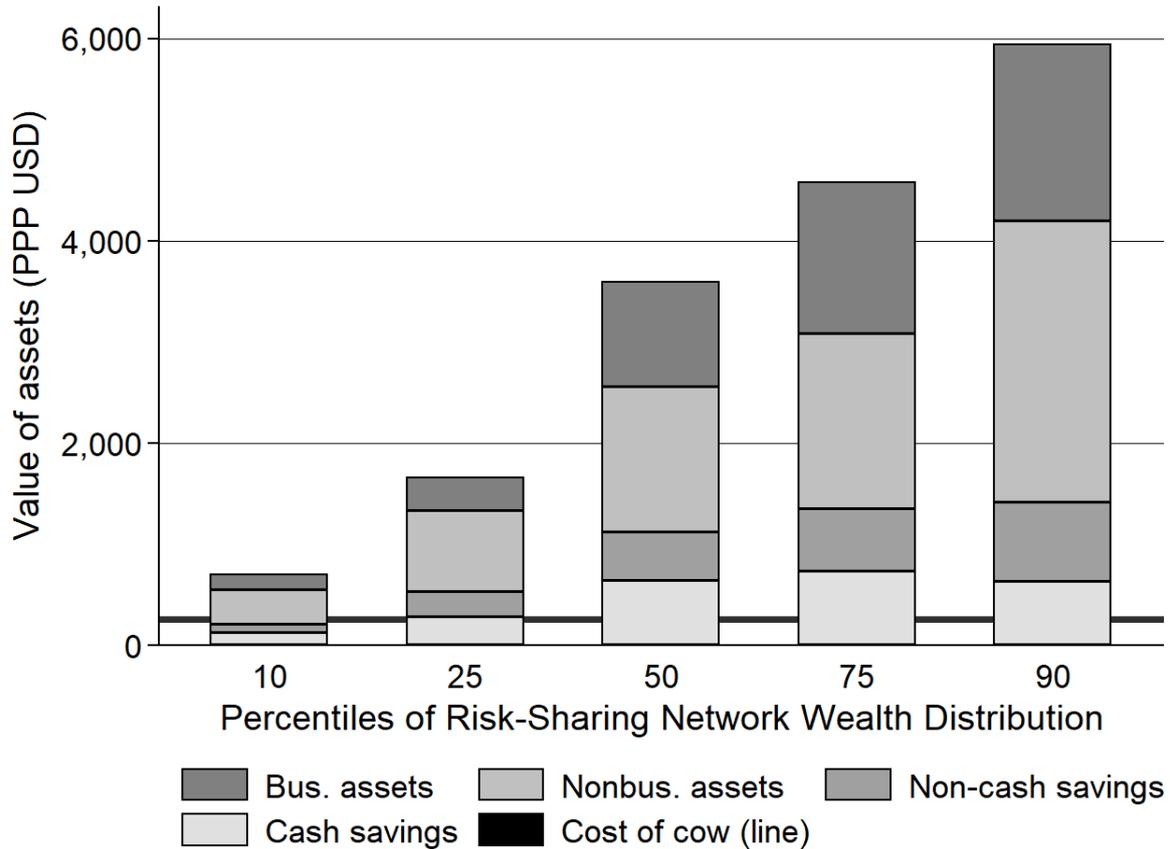
Notes: The sample comprises all households in all villages across the full sample in 2009. Data were collected using a stratified random sampling scheme, so weights are used to aggregate household data to the risk-sharing network level. All financial variables are in 2007 USD terms. Where necessary, they are first deflated to 2007 terms using the Bangladesh central bank CPI measure of inflation. They are then converted to dollars using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. In Column (1) the outcome measures the price per litre of milk. The price of milk is calculated by dividing for each household their total expenditure on milk by the total quantity they report purchasing, and then averaging this across households in the village. In Column (2) the outcome measures the price of cows. The price of cows is calculated by dividing for each village the total reported value of cows by the total reported number of cows. In Column (3) the outcome measures the hourly wage. The wage is calculated by dividing for each household the total reported income by the total reported number of wage hours worked, and then averaging this across households in the village. In Column (4) outcome measures the change in the share of all ineligible poor households who are without cows and report aspiring to own cows. The table shows regression of the outcome in 2009 on the aggregate value of capital provided to the risk-sharing network in 2007. Controls are the lagged values of total income, total savings, and total value of cows.

Figure 1: Frequency Distribution of Aggregate Value of Capital Injection Provided by the Program



Notes: Constructed using data on households in treated villages across the full sample. Observations at the risk-sharing network level. Within a village, a risk-sharing network is the set of households in the lowest two (three) wealth classes, when there are four (five) possible wealth classes. Wealth classes are determined using a Participatory Rural Assessment. Data were collected using a stratified random sampling scheme, so weights are used throughout to aggregate household data to the risk-sharing network level. The value of the capital injection is the value of the assets transferred to an eligible household (515 USD PPP) multiplied by the number of eligible households in the risk-sharing network. The value of transfer to an eligible household is calculated by converting the value of the assets in Bangladeshi Taka (9,500TK) to 2007 USD terms, using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. A thin tail of networks receive more than \$8,000 worth of assets. Since the density on this part of the support is low (fewer than five networks for any value of asset transfers), I trim these networks (2% of the sample).

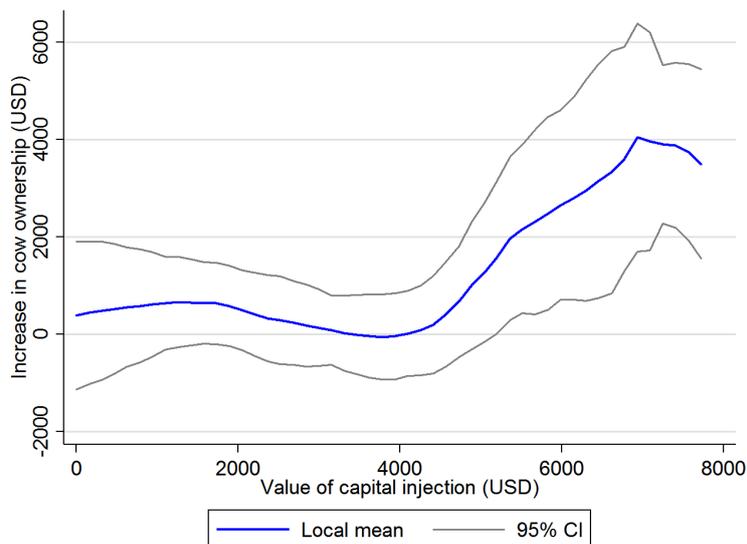
Figure 2: Distribution of Aggregate Assets Across Risk-Sharing Networks, Broken Down by Type



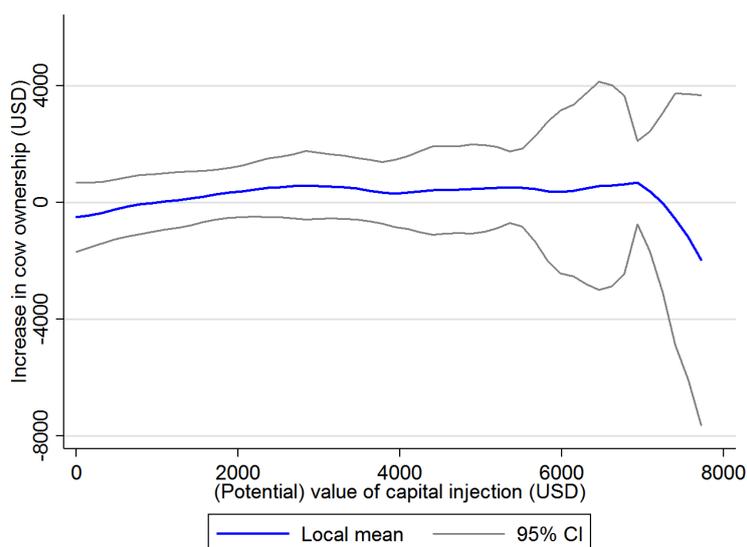
Notes: All statistics are constructed using baseline (pre-program) data for all villages, both treated and control, across the full sample. Observations are at the risk-sharing network level. Within a village, a risk-sharing network is the set of households in the lowest two (three) wealth classes, when there are four (five) possible wealth classes. Wealth classes are determined using a Participatory Rural Assessment. Data were collected using a stratified random sampling scheme, so weights are used throughout to aggregate household data to the risk-sharing network level. All asset values are in 2007 USD terms, converted to dollars using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. Cash savings include savings held at home, in any bank, with any NGO or microfinance institution, and with any savings guard. Non-cash savings include the value of jewellery and ceremonial sarees. Nonbusiness assets include electrical devices (radios, televisions, refrigerators), personal vehicles (bicycles, motorbikes), and furniture. Business assets include animals, farm infrastructure and machinery, and productive vehicles (rickshaw, van, cart).

Figure 3: Impact of Capital Injection on Further Investment

(a) Among Treated Networks

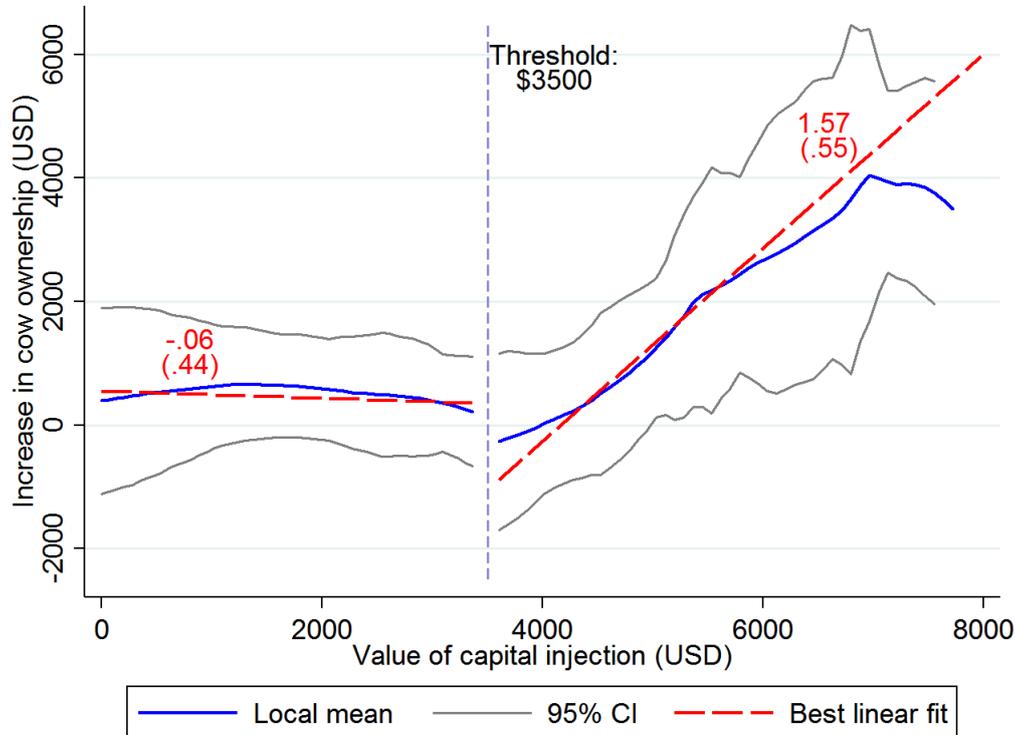


(b) Among Control Networks – Placebo test



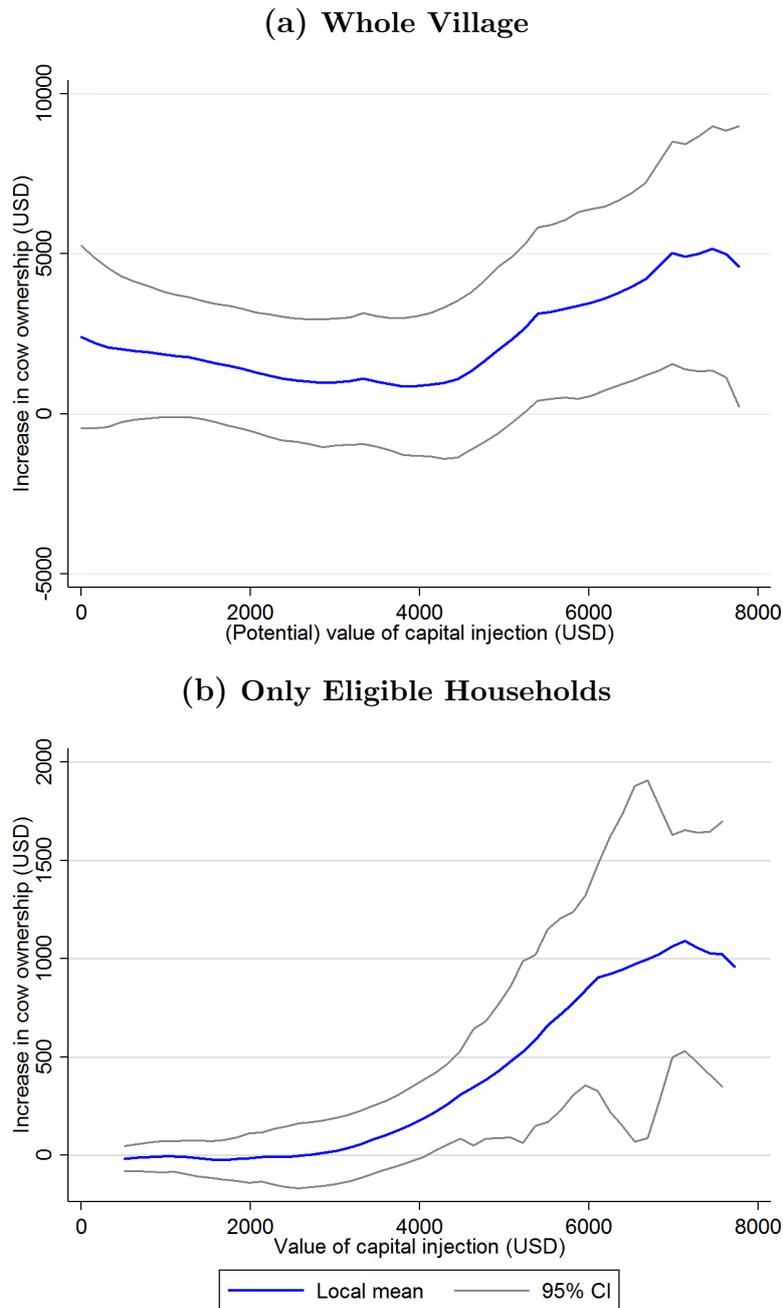
Notes: Panel (a) [(b)] constructed using data on all poor households in treated [control] villages across the full sample. Observations at the risk-sharing network level. Within a village, a risk-sharing network is the set of households in the lowest two (three) wealth classes, when there are four (five) possible wealth classes. Wealth classes are determined using a Participatory Rural Assessment. Data were collected using a stratified random sampling scheme, so weights are used throughout to aggregate household data to the risk-sharing network level. The [potential] value of the capital injection is the value of the assets [that would have been] transferred to an eligible household (515 USD PPP) [if they had been in a treated village] multiplied by the number of eligible households in the risk-sharing network. The [potential] value of transfer to an eligible household is calculated by converting the value of the assets in Bangladeshi Taka (9,500TK) to 2007 USD terms, using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. Increase in cow ownership is measured as the increase in the value of cows between 2009 and 2011 owned by all households in the risk-sharing network. The values in Bangladeshi Taka are first deflated to 2007 terms using the Bangladesh central bank CPI measure of inflation, and then converted to USD PPP. A thin tail of networks receive more than \$8,000 worth of assets. Since the density on this part of the support is low (fewer than five networks for any value of asset transfers), I trim these networks (2% of the sample). The graph shows the kernel-weighted local mean, estimated using an Epanechnikov kernel with bandwidth \$800. The outer region provides the 95% confidence interval.

Figure 4: Impact of Capital Injection on Further Investment, either side of Threshold



Notes: Constructed using data on all poor households in treated villages across the full sample. Observations at the risk-sharing network level. Within a village, a risk-sharing network is the set of households in the lowest two (three) wealth classes, when there are four (five) possible wealth classes. Wealth classes are determined using a Participatory Rural Assessment. Data were collected using a stratified random sampling scheme, so weights are used throughout to aggregate household data to the risk-sharing network level. The value of the capital injection is the value of the assets transferred to an eligible household (515 USD PPP) multiplied by the number of eligible households in the risk-sharing network. The value of transfer to an eligible household is calculated by converting the value of the assets in Bangladeshi Taka (9,500TK) to 2007 USD terms, using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. Increase in cow ownership is measured as the increase in the value of cows between 2009 and 2011 owned by all households in the risk-sharing network. The values in Bangladeshi Taka are first deflated to 2007 terms using the Bangladesh central bank CPI measure of inflation, and then converted to USD PPP. A thin tail of networks receive more than \$8,000 worth of assets. Since the density on this part of the support is low (fewer than five networks for any value of asset transfers), I trim these networks (2% of the sample). The figure shows non-parametrically the relationship between the increase in the aggregate value of cows in a risk-sharing network between two years and four years after transfers were made, and the value of capital provided to the network by the program. This is plotted either side of the estimated threshold of \$3,500. This threshold was selected by linear regressions of investment on capital injection, at a sequence of possible values for the threshold. The most likely value for the threshold is then the proposed value in the regression which had the largest F-statistic for a change in the slope. I use the Quandt Likelihood Ratio test, as described in Section 4, to test for significance of the threshold. The non-parametric relationship shown is a kernel-weighted local mean, estimated using an Epanechnikov kernel with bandwidth \$800. The outer region provides the 95% confidence interval. The best linear fit is plotted either side of the threshold, with slope coefficient noted and standard errors in parentheses.

Figure 5: Impact of Capital Injection on Further Investment, for Alternative Network Definitions



Notes: Panel (a) [(b)] constructed using data on all [eligible] households in treated villages across the full sample. Observations at the village level [aggregated across eligible households to the village level]. Data were collected using a stratified random sampling scheme, so weights are used in Panel (a) to aggregate household data to the village level. The value of the capital injection is the value of the assets transferred to an eligible household (515 USD PPP) multiplied by the number of eligible households in the risk-sharing network. The value of transfer to an eligible household is calculated by converting the value of the assets in Bangladeshi Taka (9,500TK) to 2007 USD terms, using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. Increase in cow ownership is measured as the increase in the value of cows between 2009 and 2011 owned by all [eligible] households in the village. The values in Bangladeshi Taka are first deflated to 2007 terms using the Bangladesh central bank CPI measure of inflation, and then converted to USD PPP. A thin tail of networks receive more than \$8,000 worth of assets. Since the density on this part of the support is low (fewer than five networks for any value of asset transfers), I trim these networks (2% of the sample). The graph shows the kernel-weighted local mean, estimated using an Epanechnikov kernel with bandwidth \$800. The outer region provides the 95% confidence interval.

Appendix A Characterising the Solution of the Model

In this Appendix I provide details on the conditions which characterise the solution of the model. I begin with the characterisation when all households have invested, so that the problem has the same form as the many-household, continuous-shocks version of [Ligon et al. \(2002\)](#). I then extend the model to allow investment, and show that first order conditions of the same form characterise the insurance relationship, given some investment decision. The investment decision is studied in detail in Section 2.3.

A.1 Risk Sharing under Limited Commitment without Investment

I first consider the limited commitment problem when all households have already invested. In this case there is no investment decision to make, so a solution to the model will provide a mapping from the complete history of income shocks, to the transfers that a household makes or receives today.

To find this solution, I first use the standard technique of writing the sequential problem i.e. the choice of transfers in a given period conditional on the complete history of shocks, in a recursive formulation. Following [Spear and Srivastava \(1987\)](#) and [Abreu et al. \(1990\)](#), this simplifies the problem by encoding the dependence on the entire history into a single state variable, ‘promised utility’, ω , which summarises the relevant information.

I then take the usual approach (as in [Ligon et al., 2002](#)) of formulating the problem as a *planner’s problem*. Without loss of generality, I assume household 1 is the hypothetical planner. Its role will be to choose the transfers that each household should make at each possible history, and provide promises of utility, in a way that meets certain constraints (described below).²⁶

At any point in time, the planner’s problem will then be to maximise its own utility, denoted by the value function $V(\mathbf{y}; \boldsymbol{\omega}(\mathbf{y}), \mathbf{1})$. This value function depends on the vector of realised incomes, \mathbf{y} ; the utility levels the planner promised to each household given the incomes, $\boldsymbol{\omega}(\mathbf{y}) = \{\omega^2(\mathbf{y}), \dots, \omega^N(\mathbf{y})\}$; and the stock of capital, $\boldsymbol{\kappa}$, which here is equal to $\mathbf{1}$. The choices the planner makes are what transfers to ask each household to make today, $\boldsymbol{\tau}(\mathbf{y}) = \{\tau^2(\mathbf{y}), \dots, \tau^N(\mathbf{y})\}$; what promises of expected utility to make for tomorrow, $\bar{\boldsymbol{\omega}}'(\mathbf{y}) = \{\bar{\omega}'^2(\mathbf{y}), \dots, \bar{\omega}'^N(\mathbf{y})\}$; and how to deliver these promises, $\boldsymbol{\omega}'(\mathbf{y}, \mathbf{y}') = \{\omega'^2(\mathbf{y}, \mathbf{y}'), \dots, \omega'^N(\mathbf{y}, \mathbf{y}')\} \forall \mathbf{y}'$. The notation ‘ $'$ ’ denotes that a variable relates to tomorrow.

The planner’s problem can be written in two parts. First, the planner chooses transfers today and promises of expected utility. Second, given the promised expected utilities, it chooses how to deliver these.²⁷ Working backwards, let $\mathcal{U}(\bar{\boldsymbol{\omega}}', \mathbf{1})$ denote the value function for a planner who

²⁶This will find an equilibrium set of contingent transfers (transfers that depend on the realised history) that is subgame perfect: no household would like to unilaterally deviate from the arrangement in any realised state of the world. However, a ‘decentralised’ approach, where one directly solved the repeated game representation, would generically have many possible equilibria, from which my approach will select a single one. For work studying the decentralisation problem, see [Alvarez and Jermann \(2000\)](#) and [Ábrahám and Cárceles-Poveda \(2009\)](#).

²⁷Typically the problem is written and solved as a single step. However, as will be seen, the separation here will allow the problem to be solved even when discrete choices and discrete state variables are studied.

has to provide promised expected utility $\bar{\omega}'$ when all households are already invested, and can choose how this is delivered by selecting the utility to be delivered in each state of the world, $\omega'(\mathbf{y}', \bar{\omega}')$. $V(t; \omega', \mathbf{1})$ denotes the continuation value of promising to deliver ω' given the state is t . Then the second problem is:

$$\mathcal{U}(\bar{\omega}', \mathbf{1}) = \max_{\{\omega'^i(\mathbf{y}'; \bar{\omega}')\}_{i, \mathbf{y}'}} \int V(t; \omega'(t; \bar{\omega}'), \mathbf{1}) dF_{\mathbf{Y}'}(t) \quad (\text{A1})$$

$$= \int \max_{\{\omega'^i(\mathbf{y}'; \bar{\omega}')\}_{i, \mathbf{y}'}} V(t; \omega'(t; \bar{\omega}'), \mathbf{1}) dF_{\mathbf{Y}'}(t) \quad (\text{A2})$$

s.t.

$$[\tilde{\nu}^i] \quad \bar{\omega}^i = \int \omega^i(t; \bar{\omega}') dF_{\mathbf{Y}'}(t) \quad \forall i \in \{2, \dots, N\} \quad (\text{A3})$$

Appendix B.1 provides a proof that one can move from A1 to A2.

Now the expected continuation value, \mathcal{U} , is defined as the integral over the continuation value in each possible realisation of the shock, \mathbf{y}' , where the planner can choose what utility to promise at each possible shock, subject only to these promised utilities integrating to the promised expected utility, $\bar{\omega}'$.

The first order conditions and envelope condition for this problem are:

$$[\text{FOC}(\omega'^i(\mathbf{y}'; \bar{\omega}'))] \quad \frac{\partial V(t; \omega'(\mathbf{y}'; \bar{\omega}'), \mathbf{1})}{\partial \omega'^i(\mathbf{y}'; \bar{\omega}')} = \tilde{\nu}^i \quad (\text{A4})$$

$$[\text{ET}(\bar{\omega}^i)] \quad \frac{\partial \mathcal{U}(\bar{\omega}', \mathbf{1})}{\partial \bar{\omega}^i} = \tilde{\nu}^i \quad (\text{A5})$$

Combining these, one gets that:

$$\frac{\partial \mathcal{U}(\bar{\omega}', \mathbf{1})}{\partial \bar{\omega}^i} = \frac{\partial V(\mathbf{y}'; \omega'(\mathbf{y}'; \bar{\omega}'), \mathbf{1})}{\partial \omega'^i(\mathbf{y}'; \bar{\omega}')} \quad (\text{A6})$$

Then the first part of the planner's problem is to choose transfers, $\tau(\mathbf{y})$, and promised expected utilities, $\bar{\omega}'(\mathbf{y})$ – given that this promised expected utility will be delivered efficiently as above – to maximise:

$$\max_{\{\tau^i(\mathbf{y}), \bar{\omega}^i(\mathbf{y})\}_{i=2}^N} u \left(y^1 + R + \sum_{i=2}^N \tau^i(\mathbf{y}) \right) + \beta \mathcal{U}(\bar{\omega}'(\mathbf{y}), \mathbf{1}) \quad (\text{A7})$$

with $\mathcal{U}(\cdot)$ defined as in Equation A2, subject to three sets of constraints. The first set of constraints, with multipliers $\lambda^i(\mathbf{y})$, are the *promise keeping constraints*:

$$[\lambda^i(\mathbf{y})] \quad u(y^i + R - \tau^i(\mathbf{y})) + \beta \bar{\omega}^i(\mathbf{y}) \geq \omega^i(\mathbf{y}) \quad \forall i \in \{2, \dots, N\} \quad (\text{A8})$$

These require that, at every possible realisation of income, \mathbf{y} , the planner actually provides (at least) the promised utility that he agreed to provide. The second set of constraints, with

multipliers $\phi^i(\mathbf{y})$, are the limited commitment constraints:

$$[\phi^1(\mathbf{y})] \quad u(y^1 + R + \sum_{i=2}^N \tau^i(\mathbf{y})) + \beta \bar{V}(\bar{\omega}'(\mathbf{y}), \mathbf{1}) \geq \Omega(y^1, 1) \quad (\text{A9})$$

$$[\phi^i(\mathbf{y})] \quad u(y^i + R - \tau^i(\mathbf{y})) + \beta \bar{\omega}^i(\mathbf{y}) \geq \Omega(y^i, 1) \quad \forall i \in \{2, \dots, N\} \quad (\text{A10})$$

which require that each household (including the planner) gets at least as much expected discounted utility from the insurance arrangement as it would get if it walked away and took its outside option, $\Omega(\cdot)$. The outside option is a function of current income and current investment status, and for an invested household is calculated as the utility of consuming all its income today, and then the discounted expected utility given that it never again has insurance.²⁸ Formally:

$$\Omega(y^i, 1) := u(y^i + R) + \frac{\beta}{1 - \beta} \int u(y' + R) dF(y') \quad \forall i \in \{1, \dots, N\} \quad (\text{A11})$$

The third set of constraints, with multiplier $\beta\nu^i(\mathbf{y})$ is that for each household $i \in \{2, \dots, N\}$ the planner must find some promise of utility for every possible income realisation, such that the average utility provided across all states is equal to the promised expected utility:

$$[\beta\nu^i(\mathbf{y})] \quad \bar{\omega}^i(\mathbf{y}) = \int \omega^i(\mathbf{y}, t) dF_{\mathbf{Y}'}(t) \quad \forall i \in \{2, \dots, N\} \quad (\text{A12})$$

Taking first order conditions, using the envelope theorem for $\omega^i(\mathbf{y})$, this gives for $i \in \{2, \dots, N\}$:

$$[\text{FOC}(\tau^i(\mathbf{y}))] \quad \frac{du(c^1(\mathbf{y}))/d\tau^i(\mathbf{y})}{du(c^i(\mathbf{y}))/d\tau^i(\mathbf{y})} = \frac{\lambda^i(\mathbf{y}) + \phi^i(\mathbf{y})}{1 + \phi^1(\mathbf{y})} \quad (\text{A13})$$

$$[\text{FOC}(\bar{\omega}^i(\mathbf{y}))] \quad \frac{\partial \mathcal{U}(\bar{\omega}'(\mathbf{y}), \mathbf{1})}{\partial \bar{\omega}^i(\mathbf{y})} = -\frac{\lambda^i(\mathbf{y}) + \phi^i(\mathbf{y})}{1 + \phi^1(\mathbf{y})} \quad (\text{A14})$$

$$[\text{ET}(\omega^i(\mathbf{y}))] \quad \frac{\partial V(\mathbf{y}; \omega(\mathbf{y}), \mathbf{1})}{\partial \omega^i(\mathbf{y})} = -\lambda^i(\mathbf{y}) \quad (\text{A15})$$

Hence:

$$\frac{\partial \mathcal{U}(\bar{\omega}'(\mathbf{y}), \mathbf{1})}{\partial \bar{\omega}^i(\mathbf{y})} = -\frac{-\frac{\partial V(\mathbf{y}; \omega(\mathbf{y}), \mathbf{1})}{\partial \omega^i(\mathbf{y})} \phi^i(\mathbf{y})}{1 + \phi^1(\mathbf{y})} = -\frac{du(c^1(\mathbf{y}))/d\tau^i(\mathbf{y})}{du(c^i(\mathbf{y}))/d\tau^i(\mathbf{y})} \quad (\text{A16})$$

From the envelope theorem (Equation A15) it can be seen that the value function is decreasing in the promised utility $\omega^i(\mathbf{y})$ to each household i . When none of the limited commitment constraints bind, $\phi^1(\mathbf{y}) = \phi^i(\mathbf{y}) = 0$, the slope of the value function (the ratio of marginal utilities for i and 1) remains constant, and so the ratio of marginal utilities remain unchanged

²⁸This is the most extreme punishment that can be imposed on the household, without assuming there are also exogenous costs of relationship loss. It can therefore support the maximum amount of risk sharing. Weaker punishment strategies would provide additional, Pareto-dominated equilibria. I focus on a Pareto efficient insurance arrangement.

from the previous period. When a household's limited commitment constraint binds, the ratio of marginal utilities in that period and all future periods (until another constraint binds), is increased so that it receives an increased share of consumption.

A.2 Risk Sharing under Limited Commitment with Investment

Next I consider the case where $k < N$ investments have already been made, $k = \sum_{i=1}^N \kappa^i$. Now there is a meaningful investment decision for the planner, which is the chief innovation of the model. Precisely, the planner now has to choose the optimal number (and allocation) of investments $\Delta k(\mathbf{y}) \in \{1, \dots, N - k\}$, as well as transfers and utility promises.²⁹ I first note that there is a weakly dominant allocation rule for assigning investments.

Lemma 2. *There exists a unique weakly dominant investment allocation rule. Let $\tilde{\omega}^i(\mathbf{y}) := \max\{\omega^i(\mathbf{y}), \Omega(y^i, 0)\}$. Then if $\Delta k(\mathbf{y})$ investments are to occur, assign the investments to the $\Delta k(\mathbf{y})$ uninvested households with the highest values of $\tilde{\omega}^i(\mathbf{y})$.*

Proof. See Appendix B.2. □

The planner's problem can therefore be simplified to choose only what transfers to make and *how many* investments to do, taking as given which households will be asked to do the investments. This reduces significantly the dimensionality of the choice problem, from $(N - k)!$ possible values for the discrete choice, to only $N - k$ values.

I next simplify the problem further, making use of additional separability in the structure of the problem. The planner's decision can be separated into first choosing what transfers to make *given a decision on the number of investments*, and then choosing the optimal number of investments. This follows from an application of Bellman's Principle of Optimality. So the planner's value function, given the shock, \mathbf{y} , the promised utility at that shock, $\boldsymbol{\omega}(\mathbf{y})$, and the existing distribution of investments, $\boldsymbol{\kappa}$, is:

$$V(\mathbf{y}; \boldsymbol{\omega}(\mathbf{y}), \boldsymbol{\kappa}) = \max_{\Delta k} \{V_{\Delta k}(\mathbf{y}; \boldsymbol{\omega}(\mathbf{y}), \boldsymbol{\kappa})\} \quad (\text{A17})$$

where $V_{\Delta k}(\mathbf{y}; \boldsymbol{\omega}(\mathbf{y}), \boldsymbol{\kappa})$ is the conditional value function when the planner requires Δk investments to occur (and be assigned as above), and chooses transfers optimally.

Before defining the planner's problem for the conditional value function, I define the expected continuation value, \mathcal{U} , when investment is possible:

$$\begin{aligned} \mathcal{U}(\bar{\boldsymbol{\omega}}', \boldsymbol{\kappa}) &= \int \max_{\{\Delta k(\mathbf{y}'; \bar{\boldsymbol{\omega}}'), \{\omega^{i'}(\mathbf{y}'; \bar{\boldsymbol{\omega}}')\}_i\}_{\mathbf{y}'}} V_{\Delta k}(t; \boldsymbol{\omega}'(t; \bar{\boldsymbol{\omega}}'), \boldsymbol{\kappa}'(t; \bar{\boldsymbol{\omega}}')) dF_{\mathbf{Y}'}(t) \\ \text{s.t. } \bar{\boldsymbol{\omega}}' &= \int \boldsymbol{\omega}'(t, \bar{\boldsymbol{\omega}}') dF_{\mathbf{Y}'}(t), \quad \Delta k \in \{0, 1, \dots, N - k\} \end{aligned} \quad (\text{A18})$$

²⁹In a full commitment setting it would not matter which households 'held' the investments, since the planner could always require them to make arbitrary transfers. With limited commitment this is no longer the case: if the planner requires too high a transfer, the household may prefer autarky.

Now the benefit of writing the problem in terms of promised expected utilities can be seen. The expected value function is clearly differentiable with respect to $\bar{\omega}^i \forall i$, with the derivative equal to the value of the multiplier on the integral constraint for promised utilities.

That the expected value function should remain differentiable is not obvious. The discrete choice, Δk , introduces kinks into the value function defined in Equation A17. At the point where two conditional value functions cross (in $\omega^i(\mathbf{y})$ space), their slopes will be different. As the upper envelope of these conditional value functions, the overall value function will not be differentiable at these crossing points. With the standard approach to writing the problem, in terms of the promised utility at every state, it is not clear that the expected value function will be differentiable with respect to these promised utilities. However, with this redefinition of the problem, it is immediate that the value function will be differentiable with respect to promised expected utility.

The intuition for why this redefinition can be used to solve the problem of kinks in the value function comes from [Prescott and Townsend \(1984b, 1984a\)](#).³⁰ They model the allocation of resources in settings with moral hazard. Moral hazard introduces non-convexity into the set of feasible allocations, similar to the problem caused by kinks in my model. They show that, with a continuum of agents, they can solve the problem by introducing ‘extrinsic uncertainty’: randomness which has no bearing on economic fundamentals, but is nevertheless used in the allocation of resources conditional on all observables. More simply, they introduce lotteries which mean that, in some states (realised incomes in my model), observationally equivalent agents might receive different levels of resources. This ‘convexifies’ the problem, smoothing out any kinks. It works because the share of agents receiving a particular bundle of resources can be adjusted continuously, even when the bundles differ discretely.

In my context such extrinsic uncertainty is not needed. Randomness in the distribution of income shocks can be used instead to ‘smooth out’ the kinks. This is what Equation A18 is doing: by first choosing $\Delta k(\mathbf{y}'; \bar{\omega}')$ and $\omega'^i(\mathbf{y}'; \bar{\omega}') \forall i, \mathbf{y}'$, and then integrating over the continuum of income shocks, the upper envelope function $\mathcal{U}(\cdot)$ is made convex in promised expected utility.³¹

Having rewritten the expected continuation value in this way, I can now set up the planner’s problem with investment. The planner’s value function, $V(\mathbf{y}; \boldsymbol{\omega}(\mathbf{y}), \boldsymbol{\kappa})$, is defined as in Equation A17 as the maximum over a set of conditional value functions, each for a different fixed number of investments. Given also the definition of the expected continuation value from Equation A18, these conditional value functions, $V_{\Delta k}(\mathbf{y}; \boldsymbol{\omega}(\mathbf{y}), \boldsymbol{\kappa})$, are given by:

$$V_{\Delta k}(\mathbf{y}; \boldsymbol{\omega}(\mathbf{y}), \boldsymbol{\kappa}) = \max_{\{\tau^i(\mathbf{y}), \bar{\omega}^i(\mathbf{y})\}_{i=2}^N} u \left(y^1 + \kappa^1 R - \Delta \kappa^1(\mathbf{y}) d + \sum_{i=2}^N \tau^i(\mathbf{y}) \right) + \beta \mathcal{U}(\bar{\boldsymbol{\omega}}'(\mathbf{y}), \boldsymbol{\kappa}'(\mathbf{y})) \quad (\text{A19})$$

³⁰See also [Phelan and Townsend \(1991\)](#).

³¹A formal justification of this approach is provided by Lemma A1 and Lemma A2 of [Pavoni and Violante \(2007\)](#).

s.t.

$$[\lambda^i(\mathbf{y})] \quad u(y^i + \kappa^i R - \Delta\kappa^i(\mathbf{y})d - \tau^i(\mathbf{y})) + \beta\bar{\omega}^i(\mathbf{y}) \geq \omega^i(\mathbf{y}) \quad (\text{A20})$$

$$[\phi^1(\mathbf{y})] \quad u(y^1 + \kappa^1 R - \Delta\kappa^1(\mathbf{y})d + \sum_{i=2}^N \tau^i(\mathbf{y})) + \beta\mathcal{U}(\bar{\omega}'(\mathbf{y}), \boldsymbol{\kappa}'(\mathbf{y})) \geq \Omega^1(y^1, \kappa^1) \quad (\text{A21})$$

$$[\phi^i(\mathbf{y})] \quad u(s^i Y + \kappa^i R - \Delta\kappa^i(\mathbf{y})d - \tau^i(\mathbf{y})) + \beta\bar{\omega}^i(\mathbf{y}) \geq \Omega^i(y^i, \kappa^i) \quad (\text{A22})$$

where $i \in \{1, \dots, N\}$,

$$\Omega(y^i, \kappa^i) := u(y^i + \kappa^i R - \Delta\kappa_{aut}^i(y^i)d) + \beta \int \Omega(y', \kappa^i + \Delta\kappa_{aut}^i(y^i)) dF(y') \quad (\text{A23})$$

is the best outside option for household $i \in \{1, \dots, N\}$, and the investment state is updated as:

$$\kappa'^i = \kappa^i + \Delta\kappa^i \quad \text{where} \quad \kappa'^i, \Delta\kappa^i \in \{0, 1\} \quad (\text{A24})$$

The main differences between these conditional value functions and the case without investment are that (i) some households will potentially now invest, adding the $-\Delta\kappa^i(\mathbf{y})d$ terms to household utility; (ii) the investment state $\boldsymbol{\kappa}$ must be updated when investment occurs; and (iii) the outside option for household i now allows for the option of future investment, if the household has not already invested.

As before this gives first order conditions for $i \in \{2, \dots, N\}$, now (implicitly) conditional on both the income shock (as before), and also the investment decision, $\Delta\kappa$:

$$[\text{FOC}(\tau^i(\mathbf{y}))] \quad \frac{du(c^1(\mathbf{y}))/d\tau^i(\mathbf{y})}{du(c^i(\mathbf{y}))/d\tau^i(\mathbf{y})} = \frac{\lambda^i(\mathbf{y}) + \phi^i(\mathbf{y})}{1 + \phi^1(\mathbf{y})} \quad (\text{A25})$$

$$[\text{FOC}(\bar{\omega}^i(\mathbf{y}))] \quad \frac{\partial\mathcal{U}(\bar{\omega}'(\mathbf{y}), \boldsymbol{\kappa}')}{\partial\bar{\omega}^i(\mathbf{y})} = -\frac{\lambda^i(\mathbf{y}) + \phi^i(\mathbf{y})}{1 + \phi^1(\mathbf{y})} \quad (\text{A26})$$

$$[\text{ET}(\omega^i(\mathbf{y}))] \quad \frac{\partial V(\mathbf{y}; \boldsymbol{\omega}(\mathbf{y}), \boldsymbol{\kappa})}{\partial\omega^i(\mathbf{y})} = -\lambda^i(\mathbf{y}) \quad (\text{A27})$$

Hence:

$$\frac{\partial\mathcal{U}(\bar{\omega}'(\mathbf{y}), \boldsymbol{\kappa}')}{\partial\bar{\omega}^i(\mathbf{y})} = -\frac{-\frac{\partial V(\mathbf{y}; \boldsymbol{\omega}(\mathbf{y}), \boldsymbol{\kappa}')}{\partial\omega^i(\mathbf{y})} + \phi^i(\mathbf{y})}{1 + \phi^1(\mathbf{y})} = -\frac{du(c^1(\mathbf{y}))/d\tau^i(\mathbf{y})}{du(c^i(\mathbf{y}))/d\tau^i(\mathbf{y})} \quad (\text{A28})$$

The first order conditions and envelope theorem result take the same form as without the investment decision. Hence the conditional value function is decreasing in promised utility, and the ratio of marginal utilities updated when a limited commitment constraint binds. This fully characterises the insurance transfers, given some exogenous investment decision.

Appendix B Theoretical Appendix

B.1 Proof that one can take the maximum inside the integral

To show the definitions of $\mathcal{U}(\bar{\omega}', \mathbf{1})$ in Equations A1 and A2 are equivalent, I first note that:

$$V(t; \omega'(t; \bar{\omega}'), \mathbf{1}) \leq \sup_{\mathbf{w}'(\cdot, \bar{\omega})} V(t; \mathbf{w}'(t, \bar{\omega}'), \mathbf{1}) \quad \forall t, \bar{\omega}', \omega'(t; \bar{\omega}') \quad (\text{B1})$$

Integrating both sides:

$$\int V(t; \omega'(t; \bar{\omega}'), \mathbf{1}) dF_{\mathbf{Y}'}(t) \leq \int \sup_{\mathbf{w}'(\cdot, \bar{\omega})} V(t; \mathbf{w}'(t, \bar{\omega}'), \mathbf{1}) dF_{\mathbf{Y}'}(t) \quad \forall \bar{\omega}', \omega'(t; \bar{\omega}') \quad (\text{B2})$$

Then:

$$\sup_{\mathbf{w}'(\cdot, \bar{\omega})} \int V(t; \mathbf{w}'(t; \bar{\omega}'), \mathbf{1}) dF_{\mathbf{Y}'}(t) \leq \int \sup_{\mathbf{w}'(\cdot, \bar{\omega})} V(t; \mathbf{w}'(t, \bar{\omega}'), \mathbf{1}) dF_{\mathbf{Y}'}(t) \quad \forall \bar{\omega}' \quad (\text{B3})$$

Note also that on the LHS of Equation B3, choice of $\mathbf{w}'(\cdot, \bar{\omega})$ is essentially choice of the integrand $V(\cdot; \mathbf{w}'(\cdot; \bar{\omega}'), \mathbf{1})$ to maximise the value of the integral. The choice of V that was made on the RHS is still available, although in principle some other choice could be better. Hence:

$$\sup_{\mathbf{w}'(\cdot, \bar{\omega})} \int V(t; \mathbf{w}'(t; \bar{\omega}'), \mathbf{1}) dF_{\mathbf{Y}'}(t) \geq \int \sup_{\mathbf{w}'(\cdot, \bar{\omega})} V(t; \mathbf{w}'(t, \bar{\omega}'), \mathbf{1}) dF_{\mathbf{Y}'}(t) \quad \forall \bar{\omega}' \quad (\text{B4})$$

Combining Equations B3 and B4 it must be that the two sides are equal, so taking the integral of the maximum, as in Equation A2, gives the same result as taking the maximum of the integral, Equation A1.

B.2 Proof of Lemma 2: Weakly dominant investment allocation rule

Suppose this rule were not weakly dominant. Then for some income shock \mathbf{y} , and some desired number of investments, $\Delta k(\mathbf{y}, \boldsymbol{\kappa})$, there exists an alternative investment allocation strategy which is *strictly* better than the one proposed i.e. there exists a pair of households i, j such that $\tilde{\omega}^i(\mathbf{y}, \boldsymbol{\kappa}) > \tilde{\omega}^j(\mathbf{y}, \boldsymbol{\kappa})$, $\Delta \kappa^i = 0$, $\Delta \kappa^j = 1$.

To show this cannot be the case, first note that households draw from a common income distribution, so the probability of any income draw is as likely for i and j . I define $\tilde{y}^i(\boldsymbol{\kappa})$ implicitly as $\tilde{\omega}^i(\mathbf{y}, \boldsymbol{\kappa}) = \Omega^i(\tilde{y}^i, 1)$, the value of income such that the household's limited commitment constraint when invested just binds. In effect this is the maximum income draw the household could get, if assigned an investment, before its promised utility would need to be increased to keep it in the risk-sharing arrangement. Since $\Omega(y^i, 1)$ is increasing in y^i , \tilde{y}^i is increasing in $\tilde{\omega}^i(\mathbf{y}, \boldsymbol{\kappa})$. Then $\tilde{\omega}^i(\mathbf{y}, \boldsymbol{\kappa}) > \tilde{\omega}^j(\mathbf{y}, \boldsymbol{\kappa})$ implies that $\tilde{y}^i > \tilde{y}^j$. Hence there exists a region of individual income shock, $[\tilde{y}^j, \tilde{y}^i]$ such that an income of this size to household i would not increase the utility it is promised, but a shock of this size to j would increase the utility it is promised. There is no income shock that could occur in the next period that would increase the promised utility to i , that wouldn't increase the promised utility to j at least as much if j received that income shock. Then, since providing utility is costly to the planner, this allocation rule is more costly than instead allocating the investment to i .

B.3 Proof of Proposition 1: Threshold income level for investment

The proof of Proposition 1 involves three steps. I first show that when the planner is choosing $\Delta k(\mathbf{y}, \boldsymbol{\kappa})$ optimally, $\frac{\partial V_{\Delta k}}{\partial Y} > 0$. Next I show that $\frac{\partial V_{\Delta k}}{\partial Y} > \frac{\partial V_{\Delta k-1}}{\partial Y} > 0$, so that $V_{\Delta k-1}$ and $V_{\Delta k}$ cross at most once. Finally I show that $V_{\Delta k-1}$ and $V_{\Delta k}$ do cross at least once, and hence there is a unique $\widehat{Y}_{\Delta k}$ s.t. $\Gamma_{\Delta k}(\widehat{Y}_{\Delta k}) \equiv V_{\Delta k-1}(\widehat{Y}_{\Delta k}) - V_{\Delta k}(\widehat{Y}_{\Delta k}) = 0$.

Conditional value functions are increasing in Y

I want to show that $\frac{\partial V_{\Delta k}}{\partial Y} > 0$. Taking the derivative of the conditional value function $V_{\Delta k}$ wrt Y , and applying the envelope theorem, I get:

$$\begin{aligned} \frac{\partial V_{\Delta k}}{\partial Y} &= t^1 \frac{du(c^1)}{dc^1} + \sum_{i=2}^N t^i \lambda^i \frac{du(c^i)}{dc^i} \\ &= \frac{du(c^1)}{dc^1} \end{aligned}$$

where for notational convenience I define $t^1 = \left(1 - \sum_{i=2}^N t^i\right)$, and the second equality comes from use of the FOCs wrt τ^i . Hence from the properties of $u(\cdot)$, $\frac{\partial V_{\Delta k}}{\partial Y} > 0$.

Slopes of conditional value functions are increasing in Δk

By the budget constraint, aggregate consumption $C = \sum_{j=1}^N c^j$ is total income less spending on investment: $C = Y + kR - \Delta kd$. From the first order conditions wrt τ^i , c^i is strictly increasing in c^1 , so all households' consumptions must be strictly increasing in aggregate consumption. Hence, since aggregate consumption is strictly decreasing in Δk , the number of investments, $c_{\Delta k}^1 < c_{\Delta k-1}^1$. Then by concavity of $u(\cdot)$, $\frac{\partial V_{\Delta k-1}}{\partial Y} - \frac{\partial V_{\Delta k}}{\partial Y} = u'(c_{\Delta k-1}^1) - u'(c_{\Delta k}^1) < 0$. As the conditional value function when there are Δk investments is always strictly steeper than the value function associated with $\Delta k - 1$ investments, and the value functions are continuous (again inherited from properties of $u(\cdot)$) they can cross at most once.

Conditional value functions cross at least once

To see that the conditional value functions do have at least one crossing, I show the limits of their difference as aggregate income falls and rises.

As aggregate income falls towards Δkd , the cost of making Δk investments, the value of making $\Delta k - 1$ investments today remains positive. However the value of making Δk investments goes to negative infinity as there are no resources left for consumption (since $u(\cdot)$ satisfies the Inada conditions). Hence the difference in value becomes infinite:

$$\lim_{Y \rightarrow d} V_{\Delta k-1} - V_{\Delta k} = \infty$$

Conversely, as aggregate income this period rises towards infinity, the difference in utility today between investing and not investing goes to zero. Hence the difference in the conditional value functions is just the difference in the value between having $k + \Delta k - 1$ or having $k + \Delta k$

investments (collectively), where $k \equiv |\boldsymbol{\kappa}_k|$ is the number of existing investments. Since the value is increasing in the number of investments, the difference in values is negative.

$$\lim_{Y \rightarrow \infty} V_{\Delta k-1} - V_{\Delta k} = \beta \mathbb{E} [V(\boldsymbol{\omega}', \boldsymbol{\kappa}_{k+\Delta k-1}) - V(\boldsymbol{\omega}', \boldsymbol{\kappa}_{k+\Delta k})] < 0$$

Hence since the conditional value functions are continuous in Y , they must cross. Since I showed earlier that $\frac{\partial V_{\Delta k}}{\partial Y} > \frac{\partial V_{\Delta k-1}}{\partial Y} > 0$, there can be at most one crossing of the value functions.

B.4 Proof of Lemma 1: The Thresholds are Decreasing in Capital

I first show that under full commitment, the value function $V_{\Delta k}(\kappa_k)$ has *increasing differences* in $(\Delta k, k)$:

$$V_{\Delta k+1}(\kappa_k) - V_{\Delta k+1}(\kappa_{k-1}) > V_{\Delta k}(\kappa_k) - V_{\Delta k}(\kappa_{k-1})$$

To see this I expand the conditional value functions

$$V_{\Delta k+1}(\kappa_k) - V_{\Delta k+1}(\kappa_{k-1}) = u(c_{\Delta k+1,k}^1) - u(c_{\Delta k+1,k-1}^1) + \beta \mathbb{E} [V(\kappa_{k+\Delta k+1}) - V(\kappa_k)]$$

and

$$V_{\Delta k}(\kappa_k) - V_{\Delta k}(\kappa_{k-1}) = u(c_{\Delta k,k}^1) - u(c_{\Delta k,k-1}^1) + \beta \mathbb{E} [V(\kappa) - V(\kappa_{k+\Delta k-1})]$$

Hence the double difference, $[V_{\Delta k+1}(\kappa_k) - V_{\Delta k+1}(\kappa_{k-1})] - [V_{\Delta k}(\kappa_k) - V_{\Delta k}(\kappa_{k-1})]$ gives

$$[u(c_{\Delta k+1,k}^1) - u(c_{\Delta k+1,k-1}^1)] - [u(c_{\Delta k,k}^1) - u(c_{\Delta k,k-1}^1)] + \beta \mathbb{E} [V(k + \Delta k + 1) - V(k + \Delta k - 1)]$$

Letting $C := [Y + kR - \Delta kd]$ denote aggregate consumption, I note that household consumption is proportion to aggregate consumption. The increase in aggregate consumption when initial capital increases from $k - 1$ to k is independent of the number of investments made today, i.e. $C_{\Delta k+1,k} - C_{\Delta k+1,k-1} = C_{\Delta k,k} - C_{\Delta k,k-1} = R$. Then the difference in consumption is the same in both the first and second set of square brackets above, but by concavity of the utility function $u(\cdot)$ the gain in utility from this increase is higher at lower levels of consumption i.e. when investment is higher. Hence $u(c_{\Delta k+1,k}^1) - u(c_{\Delta k+1,k-1}^1) > u(c_{\Delta k,k}^1) - u(c_{\Delta k,k-1}^1) > 0$, so the first two terms are (together) strictly positive. Since value functions are increasing in the level of capital, the final term is also strictly positive, so the value function exhibits increasing differences in $(\Delta k, k)$.

Next, note that since investment and capital are positive integers, $\Delta k, k \in \mathbb{Z}_+$, the set of possible values for $(\Delta k, k)$ form a *lattice*.³² Finally, as $V_{\Delta k}(\kappa_k)$ has increasing differences in $(\Delta k, k)$, and the set of possible $(\Delta k, k)$ form a lattice, $V_{\Delta k}(\kappa_k)$ is supermodular in $(\Delta k, k)$. Hence by application of *Topkis' Theorem* (Topkis, 1978), the optimal choice of Δk is non-decreasing in k at any given income. This implies the threshold level of income needed for Δk investments to be optimal, $\widehat{Y}_{\Delta k}$, is weakly lower as k increases i.e. $\mathbf{D}_k \widehat{Y}_{\Delta k}^{\text{FC}} < 0$.

³²A lattice is a partially-ordered set where for any pair of elements in the set, the least upper bound and greatest lower bound of the elements are also in the set. For more details see Milgrom and Shannon (1994).

Appendix C Defining Risk Sharing Networks

To determine the appropriate group for risk sharing, I use a subsample of 35 villages in which, rather than the stratified random sampling scheme used elsewhere, a census of all households was taken at all waves. Households were asked whether they suffered a ‘crisis’ in the last year. Potential crises include crop loss, serious illness or death of household member, and damage to house due to natural disaster. If they did, they were asked how they coped with it, and where transfers or informal loans were used for coping, they were asked who the transfers or loans were from. Additionally all households were asked who (if anyone) they borrowed food from or lent food to. I combine these various dimensions of household links into a single dimension, which I term ‘sharing risk’. I then study what grouping can be constructed in the full sample, that provides a good proxy for being a risk-sharing partner of an eligible household, since my interest is in constructing the risk-sharing network for these households.

Table C1 provides evidence on this question. The first point of note is that almost all of eligibles’ risk sharing (94%) is done with other households in the same village. Second these links are highly concentrated among other households in the lowest two wealth classes. In particular, 70% of eligibles risk-sharing links are with other households from the bottom two wealth classes, compared with only 55% that would be expected under random linking. This motivates me to focus on the poorest two wealth classes as the relevant group for risk sharing.

To further test this definition of the risk-sharing network, I perform Townsend tests ([Townsend, 1994](#)) under the different groupings. These involve regressions of the following form:

$$\begin{aligned} \Delta \log c_{hgt} = & \beta_0 + \beta_1 \Delta \log y_{hgt} + \beta \Delta \mathbf{z}_{hgt} \\ & + \delta_0 D_{hg} + \delta_1 D_{hg} \Delta \log y_{hgt} + \delta \Delta D_{hg} \mathbf{z}_{hgt} + \gamma_{gt} + \varepsilon_{hgt} \end{aligned} \quad (\text{C1})$$

where $\Delta \log c_{hgt}$ is the change in log expenditure for household h in risk-sharing group g at time t ; $\Delta \log y_{hgt}$ is the change in log income; $\Delta \mathbf{z}_{hgt}$ are changes in demographic characteristics; $D_{hg} = 1$ if household h is not an eligible household; and γ_{gt} are group dummies. The idea of the test is that, if eligible households in group g are able to fully smooth consumption, their expenditure should not respond to changes in their household income, i.e. $\beta_1 = 0$, once changes in demographics and group-level shocks, γ_{gt} , which cannot be smoothed, are accounted for. Including the interactions with D_{hg} allows ineligible households (poor and non-poor) to potentially be in the same risk-sharing group as the eligible households but without imposing that they respond to shocks in the same way. This ensures that the results of the test are not confounded by changes in sample composition: β_1 always measures the response of eligibles’ expenditure to their income. The appropriate risk-sharing network for eligibles will then be the grouping such that, including fewer households gives a larger β_1 , but including additional households does not further reduce β_1 . If all of eligibles’ risk sharing is with the bottom two wealth classes, then excluding ineligible poor households should make risk sharing appear worse, since part of the aggregate shock is being excluded. Conversely, including the whole village should not improve measured risk sharing, because the additional households are irrelevant.

Table C1: Share of Eligibles' Links to other Categories of Household

Sample: Census sample (35 villages), baseline

	Actual (1)	Random linking (2)
Share in of links in:		
Whole village	.94	.91
Low wealth	.70	.55
Other eligibles	.12	.06
Total links		578
Total households		197
Links per household		2.98

Notes: These statistics are constructed using baseline (pre-program) data for all households in a 35 village subsample of the data. In these villages, the sample includes a census of all households, allowing the characteristics of a household's 'neighbours' to be observed. A pair of households are linked ('neighbours') if either reports (a) going to the other household for assistance in a crisis; (b) going to the other household to borrow food; or (c) receiving transfers from the other. Wealth classes are determined using a Participatory Rural Assessment, which aggregates classifies households into four or five wealth classes. Additional criteria are used to determine which households in the poorest wealth class are eligible for treatment. 'Village' includes all households within the village; 'low wealth' includes all households in the bottom two wealth classes (bottom three when five wealth ranks were used); 'other eligibles' includes only households who are also eligible for the program. Column 1 shows the share of all links from eligible households to households in these other 'wealth class' categories. Column 2 shows the share of all links from eligible households that would go to households in these other 'wealth class' categories if links were formed randomly.

To estimate this I use data on expenditure and income for all households in control villages in the main sample over the three waves of data collection. Both the observations and variables used in this test are separate from the previous approach, so this provides independent evidence about the appropriate group. Equation C1 is estimated for the different definitions of group previously considered. Table C2 shows the results of this test. Consistent with the earlier result, it can be seen that including ineligible poor households into the risk-sharing network for the eligibles improved measured consumption smoothing (p-value=.026). However, including the rest of the village does not further improve smoothing (p-value=.403), justifying their exclusion from the risk-sharing network.

Table C2: Townsend Test: How Does Expenditure Change With Income Changes

Dependent Variable: Change in log expenditure

Sample: Control households (main sample), 0-2-4 year

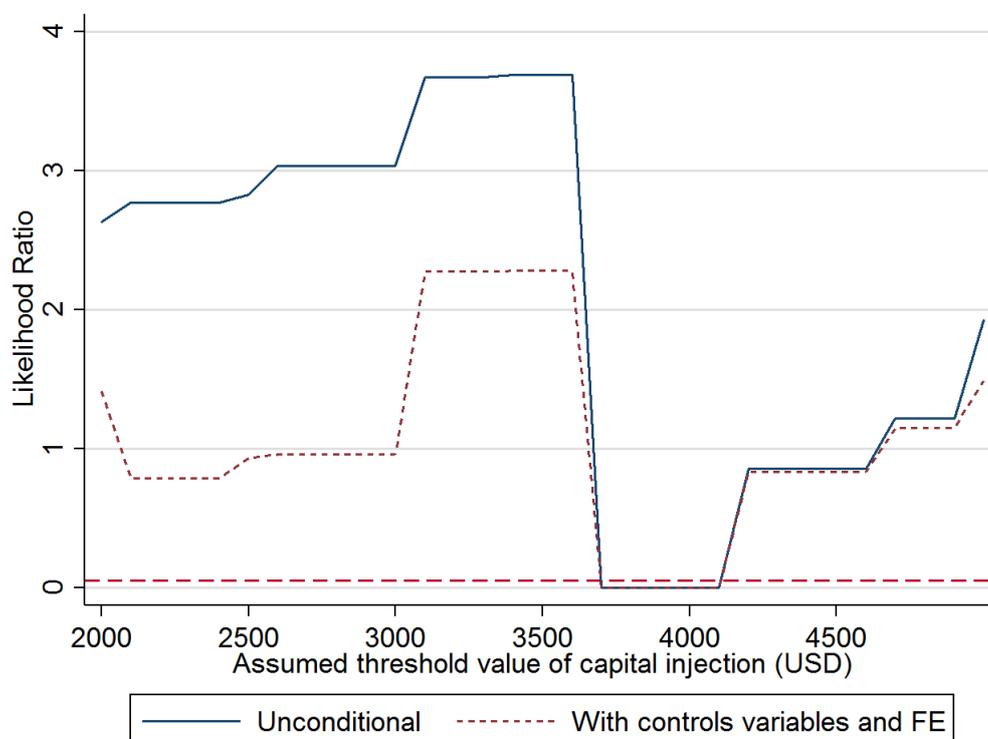
Village level clustered standard errors

Group definition:	Eligibles (1)	All Poor (2)	Whole Village (3)
Change in log income	.048 (.015)	.041 (.013)	.043 (.013)
P-values:			
Same as Column (1)	—	.026	.128
Same as Column (2)	—	—	.403
Demographic controls	Yes	Yes	Yes
Households	9,587	9,587	9,587
Observations (household \times wave)	19,174	19,174	19,174
Clusters	1,409	1,409	1,409

Notes: The sample comprises all households in control villages across the full sample. Observations are at the household level. Wealth classes are determined using a Participatory Rural Assessment, which classifies households into four or five wealth classes. Eligible households are those in the lowest wealth class who also meet additional financial and demographic criteria. All poor households are those belonging to the lowest two (three) wealth classes, when there are four (five) possible wealth classes. Whole village includes all households. Data were collected using a stratified random sampling scheme, so weights are used to make the sample representative of the population. All financial variables are in 2007 USD terms. They are first deflated to 2007 terms using the Bangladesh central bank CPI measure of inflation. They are then converted to dollars using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. The outcome measures the increase in log expenditure. Expenditure includes food (both purchased and produced), fuel, cosmetics, entertainment, transportation, utilities, clothing, footwear, footwear, utensils, textiles, dowries, education, charity, and legal expenses. Income includes total household income from all sources, but does not include transfers. Demographic controls include changes in household size by age category, and changes in education enrollment. All columns show regressions of the outcome on a constant, the change in log income, the change in demographic controls, interactions of all of these variables with an indicator for not being eligible, and a group dummy. The interaction with the indicator means the coefficients shown are those for the eligible households. Across the columns what changes is the definition of a group.

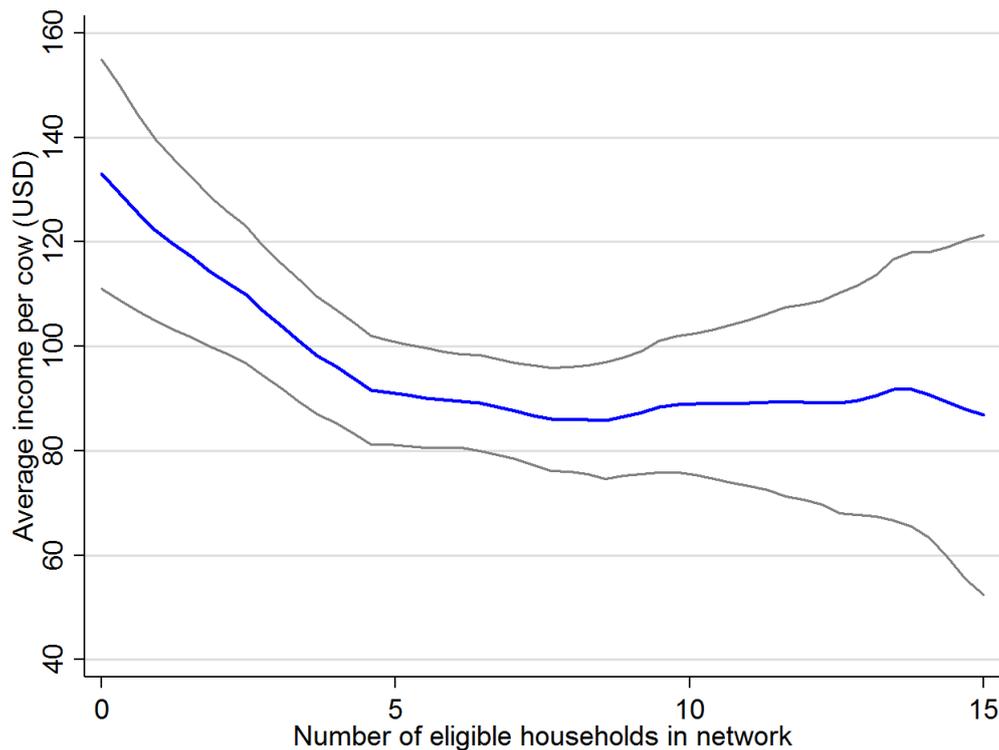
Appendix D Additional Figures and Tables

Figure D1: Hansen Test for Threshold Location



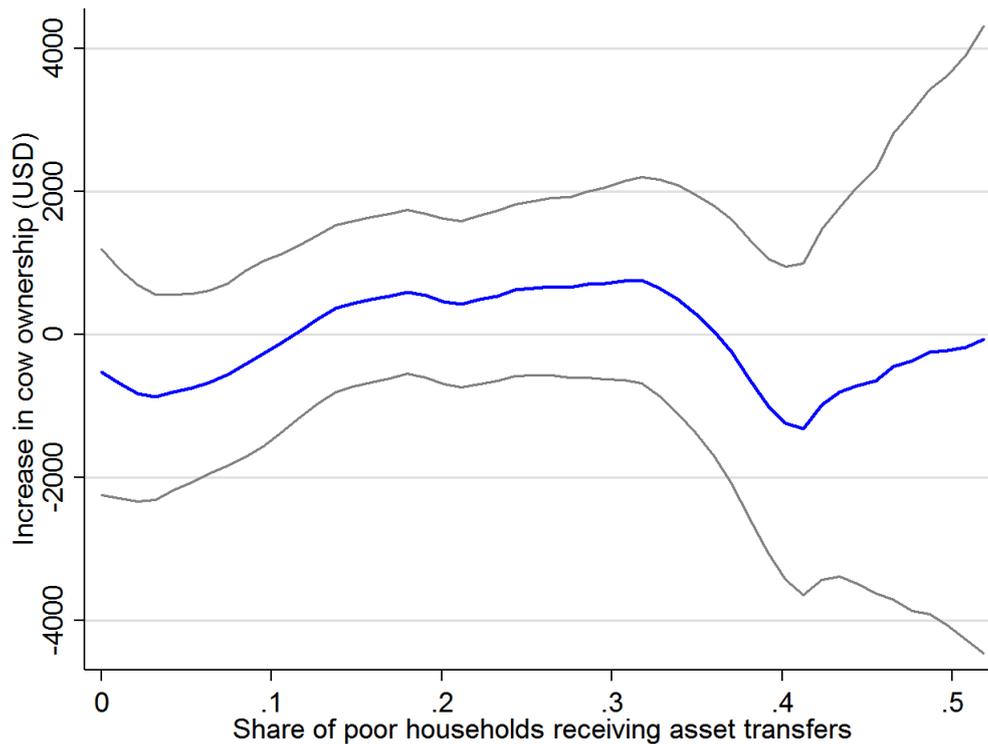
Notes: Constructed using data on all poor households in treated villages across the full sample. Observations at the risk-sharing network level. Within a village, a risk-sharing network is the set of households in the lowest two (three) wealth classes, when there are four (five) possible wealth classes. Wealth classes are determined using a Participatory Rural Assessment. Data were collected using a stratified random sampling scheme, so weights are used throughout to aggregate household data to the risk-sharing network level. The value of the capital injection is the value of the assets transferred to an eligible household (515 USD PPP) multiplied by the number of eligible households in the risk-sharing network. The value of transfer to an eligible household is calculated by converting the value of the assets in Bangladeshi Taka (9,500TK) to 2007 USD terms, using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. I sequentially run the specification in Equation 2 at different values of the threshold, varying the threshold between \$2000 and \$5000, at intervals of \$100. The figure shows, for each assumed threshold value of capital injection, the likelihood ratio (LR) statistic. This statistic is the difference in residual sum of squares (RSS) from the assumed threshold regression, relative to the RSS from the regression for which the lowest RSS was achieved, divided by that minimum RSS, and multiplied by the sample size. Any possible thresholds for which the LR is below .05 cannot be rejected as possible values for the threshold. The graph show the range of LR statistics both for the unconditional case and when additional controls (lagged income and asset variables, and network size), and district level fixed effects are included. In both cases it is clear that a threshold value of \$3,700-\$4,100 is by far the most likely, and all other thresholds can be rejected.

Figure D2: Average Return on Cows is Declining in Number of Transfer Recipients



Notes: Constructed using data on all poor households in treated villages across the full sample in 2009. Observations at the risk-sharing network level. Within a village, a risk-sharing network is the set of households in the lowest two (three) wealth classes, when there are four (five) possible wealth classes. Wealth classes are determined using a Participatory Rural Assessment. Data were collected using a stratified random sampling scheme, so weights are used throughout to aggregate household data to the risk-sharing network level. Average income per cow is the mean income per cow across cow-owning households in the network in 2009. The values in Bangladeshi Taka are first deflated to 2007 terms using the Bangladesh central bank CPI measure of inflation, and then converted to USD PPP, using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. A thin tail of networks have more than 15 eligible households. Since the density on this part of the support is low (fewer than five networks for any number of eligible households), I trim these networks (2% of the sample). The graph shows the kernel-weighted local mean, estimated using an Epanechnikov kernel with bandwidth 1.8. The outer region provides the 95% confidence interval.

Figure D3: Impact of Capital Injection on Further Investment, by Share of Households Treated



Notes: Constructed using data on all poor households in treated villages across the full sample. Observations at the risk-sharing network level. Within a village, a risk-sharing network is the set of households in the lowest two (three) wealth classes, when there are four (five) possible wealth classes. Wealth classes are determined using a Participatory Rural Assessment. Data were collected using a stratified random sampling scheme, so weights are used throughout to aggregate household data to the risk-sharing network level. The share of poor households receiving transfers is the proportion of households in the poorest wealth class who receive transfers. Increase in cow ownership is measured as the increase in the value of cows between 2009 and 2011 owned by all households in the risk-sharing network. The values in Bangladeshi Taka are first deflated to 2007 terms using the Bangladesh central bank CPI measure of inflation, and then converted to USD PPP, where 1 USD = 18.46TK in 2007. A thin tail of networks receive more than \$8,000 worth of assets. Since the density on this part of the support is low (fewer than five networks for any value of asset transfers), I trim these networks (2% of the sample). The graph shows the kernel-weighted local mean, estimated using an Epanechnikov kernel with bandwidth \$800. The outer region provides the 95% confidence interval.

Table D1: Other Variables are Smooth at the Kink in Investment

Standard errors (in parentheses) clustered at district level	Total Income	Total Savings	Business Assets	Total Assets	Village Size	Group Size	Total Migrants	Migrant House- holds
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Slope below threshold for Treated	2.85 (3.71)	-.021 (.308)	.307 (.486)	.661 (.787)	-.002 (.001)	-.002 (.010)	-.035 (.010)	-.035 (.010)
Change in slope at threshold for Treated	-1.11 (2.12)	.137 (.195)	-.149 (.247)	-.150 (.442)	.001 (.001)	-.024 (.015)	-.024 (.015)	-.024 (.015)
Treated	-7900 (7950)	-594 (679)	-2110 (910)	-4150 (1560)	4.30 (2.05)	6.35 (.029)	.200 (.029)	.200 (.029)
Mean Baseline Level among Controls	60,000	3,300	4,890	11,200	.766	.766	.766	.766
Observations (Clusters)	1360 (13)	1360 (13)	1360 (13)	1360 (13)	1360 (13)	1360 (13)	1360 (13)	1360 (13)

Notes: The sample comprises all households in all villages across the full sample in 2009. Data were collected using a stratified random sampling scheme, so weights are used to aggregate household data to the risk-sharing network level. All financial variables are in 2007 USD terms. Where necessary, they are first deflated to 2007 terms using the Bangladesh central bank CPI measure of inflation. They are then converted to dollars using purchasing power parity (PPP) exchange rates, where 1 USD = 18.46TK in 2007. In Column (1) the outcome measures the total income in the risk-sharing network at baseline. In Column (2) the outcome measures the total stock of saving in the risk-sharing network at baseline. In Column (3) the outcome measures the total value of business assets in the risk-sharing network at baseline. Business assets include animals, farm infrastructure and machinery, and productive vehicles (rickshaw, van, cart). In Column (4) the outcome measures the total value of non-business assets in the risk-sharing network at baseline. Nonbusiness assets include electrical devices (radios, televisions, refrigerators), personal vehicles (bicycles, motorbikes), and furniture. In Column (5) the outcome measures the total number of households in the whole village. In Column (6) the outcome measures the total number of households in the risk-sharing network. In Column (7) the outcome measures the total number of migrants who have left from any household in the risk-sharing network. In Column (8) the outcome measures the total number of households in the risk-sharing network from which a migrant has left.