Who Bears the Loss from a Flood? Prices, Rationing, and Incidence in Rural Bangladesh*

Mohammad Islamul Haque

November 6, 2026

Abstract

Floods destroy livestock, equipment and other productive assets across the developing world, yet flood relief often targets only the land-poor. This paper asks who bears the loss, through which village markets, and who adapts. I use a three-round panel of 5,473 households in Bangladesh from 2011 to 2019, matched to satellite flood maps, and compare neighboring Unions of the same district that floods cover unequally. The largest inherited landholders lose livestock and working capital, which is still lower three to four years later. Their farms shrink, and their wage and self-employment income falls by 7.6 percent. The rent falls by a fifth, but no group rents in more land. The labor market rations work instead of lowering the wage. The land-poor lose nearly a quarter of their day labor but turn to self-employment, so their combined income falls by at most 4 percent and their consumption by at most 0.05 percent. Labor can be redeployed and capital cannot, so who bears the loss from a flood depends on who owns the capital it hits, not who is poor. In a model where farms without credit choose when to replace lost animals, a grant to the farms that lose a herd lets them rebuild within a year and is worth about 2 percent of consumption. Fairly priced index insurance that pays the full loss does nearly as well at no expected fiscal cost. The same budget given to the land-poor does not move the land or the labor market.

JEL codes: O12, O13, O15, O16, J43, Q15, Q54.

Keywords: floods, incidence, working capital, labor rationing, land rental markets, production credit, migration, Bangladesh.

*Mohammad Islamul Haque, Yale Research Initiative on Innovation and Scale.

1 Introduction

Floods are the most common natural disaster, and most of the people exposed to them live in low- and middle-income countries (Patel 2024; Jayachandran 2006). Who bears the loss from a flood depends on how markets adjust and which households can adapt. The answer matters for policy, because recovery programs after a disaster are often aimed at the land-poor, on the presumption that the burden ends up with them.

That presumption rests on a well-established literature on productivity shocks in rural economies (Carter and Barrett 2006). Where workers cannot borrow or migrate, a fall in the demand for labor lowers the wage, so landless households bear cultivators' weather risk through the wage (Jayachandran 2006; Mobarak and Rosenzweig 2014; Balboni et al. 2026). A second body of evidence finds that village wages rarely fall after negative shocks while employment does, and that workers rationed out of wage work remain underemployed (Kaur 2019; Breza et al. 2021). The two disagree about how the loss from an adverse productivity shock travels, but they share a premise. The loss moves down the land distribution, from the households that own land and capital to those who primarily sell wage labor.

This paper asks who bears the loss when a flood hits a village economy, and it addresses three questions. First, whose capital and earnings does a flood reduce? Second, do the labor and land markets pass the loss on through prices or ration the quantities traded? Third, how do households with different market positions adapt? Existing evidence on these questions is limited. Most studies examine one type of shock, such as rainfall, or one market, such as the labor market (Jayachandran 2006; Kaur 2019). Others examine one population, such as the ultra-poor (Balboni et al. 2026), or aggregate outcomes at the place level (Patel 2024). Yet the propagation and incidence of a flood within a village economy matter for most low- and middle-income countries facing rising climate risk. Bangladesh is well suited to these questions. Large monsoon floods recur every few years in its delta and cover very different shares of neighboring Unions in the same season. Land is unequally distributed and largely inherited, and land inheritance sorts households into a position in each market, as landless workers or landed cultivators (Section 4).

I utilize a farm household panel of 5,473 households from the Bangladesh Integrated Household Survey, a nationally representative survey of rural Bangladesh, interviewed in 2011–12, 2015 and 2018–19. I match each household's Union Council, the lowest tier of rural local government, to the satellite record of large floods in the Global Flood Database (Tellman et al. 2021) and Sentinel-1 radar maps (Section 4) and define a severe flood as the ones that inundate at least a quarter of the Union. I measure exposure as the count of severe floods in the three to four years between interviews, with the count of the window before it

as a lag.

I compare how the same household's outcomes change between interviews when its Union has more severe floods than other Unions of the same district over the same interval. Household fixed effects remove permanent differences between households, including how flood-prone their Union is. District-by-round fixed effects remove whatever shocks hit a whole district between two interviews, such as prices or district-wide programs. Group-by-round fixed effects let each inherited-land group follow its own trajectory. What remains is variation among neighboring Unions: in the same monsoon, a flood may cover much of one Union and spare the next. The identifying assumption is that, without the floods, households in the more and the less flooded Unions of a district would have changed alike. I classify households by the cultivable land they inherited before 2006, when the first exposure window begins. The 70 percent with no land inheritance is classified as the land-poor, and the rest are divided into terciles. Standard errors are clustered at the Union level (318 clusters).

The paper has three main findings. First, the loss from a flood falls on capital at the top of the land distribution, and the largest farms shrink. The top tercile farms loses capital, its capital stays lower in the years that follow, and its farm shrinks (Section 6). I measure capital by the value of a household's livestock and of the farm equipment and durables it uses in production, whose sum I define as working capital. A second measure is an index that adds savings to working capital. In the flood window the top tercile's livestock falls by 18 percent and its working capital by 16 percent. These farms draw down their savings. While the consumption is smoothed, the top tercile farms do not rebuild their lost capital for years. Its working capital is still 10 percent lower 3-4 years after a flood. The top tercile also shrink their farms. It gives up the land it sharecropped in, 3.2 decimals in the window after ($p = 0.092$; wild bootstrap $p = 0.110$). In the window after the flood it farms less of its own land and rents more of it out, without selling any. In the window after the flood, the share of top tercile households that cultivates any crop is 5.7 percentage points lower. Its wage and self-employment income (earnings, for short) falls by 7.6 percent in the flood window, most of it non-farm enterprise income.

Second, the loss from a flood is passed to the land-poor through the labor market but they adapt through self-employment (Section 7). Per severe flood in the two years before an interview, farms cut their hired harvest labor hours by a quarter. The land-poor lose 23 percent of their days of agricultural day labor and 22 percent of their daily-wage income. Over the flood window their farm and non-farm wage labor days fall by 17 percent and the daily wage by at most 3.3 percent. As their wage labor falls, the land-poor turn to self-employment. Their days of self-employment rise by 4 percent per flood in the two-year window, and their share in non-farm self-employment rises by 2.4 percentage points in the

flood window. Their combined wage and self-employment income falls by at most 3.8 percent of a year's income. The data rule out a fall in their consumption of more than 0.05 percent in the flood window. The 90 percent interval for their savings rules out a fall. On average, therefore, the land-poor's measured income largely makes up for rationed wage work.

Third, the rent falls, but no group rents in more land in total at the lower rent, and the distribution of operated land does not change (Section 8). Over the flood window the median cash rent in a Union falls by about a fifth. Yet no group rents in more land. The Gini coefficient of operated land within a Union is unchanged, a precisely estimated zero.

Taken together, the results show that the flood loss falls on the land-rich and that the land-poor adapt better. The difference lies in what each group can adjust and rebuild. Labor can be quickly reallocated: the land-poor move from day labor to self employment, as rationed workers do elsewhere (Breza et al. 2021). Lost capital cannot be rebuild quickly: a farm's animals and equipment are exposed to the water, and selling them to get through a bad year lowers future output (Rosenzweig and Wolpin 1993). Farm size shrinks due to the lost capital too: the farms that lose capital shrink and rent out land. However, while the rent falls, there is no evidence of land reallocation at the lower price. Who bears a flood therefore depends on who owns the capital it hits, not who is poor.

The first concern is that flooded Unions were already on different paths before the floods, so the estimates would capture these trends rather than the effect of the floods. I test this by adding to the main specification a pre-trend for each household's later flood exposure, separately for each inherited-land group. With these trends the capital, labor and hiring results remain significant, while the top tercile's earnings loss keeps its sign but loses precision. The second concern is that the threshold for a severe flood drives the estimates. When I move it to 20 or 30 percent of a Union, the capital, earnings, day-labor and hiring results hold, while the size of the rent fall is sensitive to it.

The third concern is the estimator. Floods are not an absorbing treatment. With a treatment that switches on and off, a fixed effects regression gives some comparisons negative weight, which can bias an estimate or reverse its sign. I re-estimate the main results with an interaction-weighted estimator that lets the effect differ by group, round and flood history and averages with non-negative weights (Appendix A). The results remain robust to this alternative estimator. A related concern is that the exposure window is the interval between interviews, which is not of fixed length. Re-estimating with windows of two and three years counted back from each household's interview date reproduces the results. The fourth concern is measurement: the flood record for the 2018 monsoon comes from radar images, which measure floods with more error than the satellite record of earlier years. Without the radar fill, the capital losses are larger, so the radar fill does not drive our results. Households

that leave the panel after a flood could also bias the estimates. The land-poor's exit rises by 1.2 percentage points per flood, significant at the 10 percent level. Re-weighting for attrition changes the results little. In sum, our main results are largely robust to these checks.

Section 11 calibrates the framework to the empirical estimates and uses it to compare counterfactual policies after a flood. I add a rebuilding model in which a top-tercile farm that loses cattle has no credit after the flood and chooses when to buy new animals. Without recovery support, a farm that loses three cattle takes five years to rebuild its herd and keeps its consumption 8 to 10 percent below normal for three years. A one-time grant equal to the loss lets it rebuild the year after the flood and is worth about 2 percent of its consumption. Fairly priced index insurance that pays the full loss gives the same rebuilding and about nine tenths of the grant's welfare gain. The farms pay its premium, so its expected fiscal cost is zero. Insurance that pays only the average loss is worth less than half a percent, because it covers a fifth of a herd loss, and a credit line at 25 percent is worth almost nothing. Spending the grant's budget on the land-poor instead changes no price or quantity in the model. Relief that targets the poor would therefore fail to induce recovery in the farming activities.

This paper adds to three literatures and to the evidence on floods. The first is the work on productivity shocks in village labor markets reviewed above, in which the loss reaches the households that sell labor, through a lower wage or through rationed work. For large floods I find rationing, as in Kaur (2019). Farms hire fewer days while the wage barely moves. The rationed workers move into work of their own, as in Breza et al. (2021), and their income and consumption largely remain stable. The loss therefore does not move down the land distribution. It starts with the capital of the farms that hire and stays with them.

The second literature studies how households use productive assets to cope with shocks. Farmers sell and buy assets as buffer stocks (Rosenzweig and Wolpin 1993), the poorest households protect their assets at the expense of consumption (Carter and Lybbert 2012), and asset losses can keep households poor (Carter and Barrett 2006; Balboni et al. 2022). Most of this evidence concerns poor households. Following households across the whole land distribution, I find that the assets a flood destroys are at the top. The largest inherited farms smooth their consumption but rebuild slowly, and their working capital is still lower three to four years later. The model in Section 11 ties the slow rebuilding to the absence of credit after the flood and values relief against it, adding to the evidence on how insurance and credit change farmers' decisions (Karlan et al. 2014; Lane 2024).

The third literature studies land, credit and farm size. Frictions keep operated land from tracking productivity (Adamopoulos and Restuccia 2014; Chen et al. 2022), transaction costs in hiring labor shape farm size (Foster and Rosenzweig 2022), and renting land helps households adapt after a disaster (Eskander and Barbier 2023). Lane (2024) shows that credit

guaranteed to be available after a flood leads farmers in Bangladesh to rent more land. I study what happens when a flood that destroys capital where credit access is limited. The largest farms shrink and the rent falls, yet no group rents in more land in total and the distribution of operated land does not change. A lower rent does not by itself move land to the land-poor.

Finally, a large literature estimates the effects of floods and storms on entire places (Hsiang and Jina 2014; Kocornik-Mina et al. 2020; Patel 2024). Studies of the 1998 flood in Bangladesh describe how households coped (del Ninno et al. 2001, 2003; Khandker 2007) and find that wages fell for years afterward (Mueller and Quisumbing 2011). I measure who bears the loss inside the village, by the land households inherited, and the markets through which it travels.

The remainder of the paper is organized as follows. Section 2 sets out the facts about floods, land and markets in rural Bangladesh on which the questions rest. Section 3 presents the framework and the predictions the empirical sections test. Section 4 describes the data and Section 5 the empirical strategy. Sections 6 to 9 cover who loses capital and earnings, the labor market, the land market and the years after the flood, and Section 10 reports further robustness checks. Section 11 calibrates the model, checks its fit and compares policies, and Section 12 concludes.

2 Background: Floods, Land and Markets in Rural Bangladesh

Relief and recovery programs after a flood are often aimed exclusively at the poor, on the presumption that the poor bear the loss through direct damages, a decline in wage and rationing of labor. This section sets out four background facts about floods and village markets in rural Bangladesh that put that presumption in question and motivate the paper's research objectives.

Four fifths of Bangladesh lies on floodplains. In a normal year about a fifth of the country is flooded and in an extreme year up to three fifths. According to projections, flooded area will increase by about 30% under climate change (Lane 2024). There are recent histories of devastating floods in the country. The 1998 flood covered two thirds of the country at its peak (del Ninno et al. 2001, 2003). In the satellite record used here, 43 percent of Union-years between 2000 and 2018 see some inundation. A severe flood, one that covers at least a quarter of the Union, occurs in 13 percent of the Unions in our sample (Section 4). There are fine spatial variations in flood exposure too. Floods recur, and they also vary across neighboring places within one season.

Land is primarily inherited, and land market transactions happen through tenancy rather

than sale. Seventy percent of the panel's households inherited no cultivable land, and the top tercile of inheritors owns 195 decimals on average against 17 among the land-poor (Section 4.2). The Union Gini coefficient of operated land is 0.67 indicating very high inequality in land distribution. The land sales market is thin, as in other parts of South Asia. For example, in India Foster and Rosenzweig (2022) report that under 6 percent of plots were bought or sold over five years, and operated holdings adjust primarily through the rental market (Deininger and Jin 2005, 2008). Households in our sample rent, sharecrop or mortgage in total 4,886 plots. Of these, 86 percent belong to someone the tenant describes as richer, 60 percent to a non-relative and 20 percent to an owner living outside the village. Inheritance therefore determines who owns the land that is leased and who leases it and it is predetermined relative to any flood exposures.

The capital is held by the land-rich and the land-poor generally provide wage labor. Figure 1 and Online Appendix Table OA-6 report the land and labor market positions in 2011 by the land inherited before 2006. Operated and rented land already respond to the earliest floods of the panel (Online Appendix OA-2.4), which is why I do the classification using inherited land before 2006. The top tercile holds on average BDT 38,138 of working capital, the animals, equipment and durables it farms with, against BDT 15,378 among the land-poor. Among top tercile households, 78 percent hire farm labor and only 5 percent sell farm day labor, against 36 and 21 percent among the land-poor. Land is operated by 82 percent of its households, against 43 percent of the land-poor. It rents out 32 decimals more than it rents in, while the land-poor rent in 17 more than they rent out. The smallest tercile is the most active tenant, with 39 percent renting land in. The land-poor are not uniform. Among the land poor households, the 1,658 who operate any land rent in 49 decimals on average, and 24 percent of them sell farm day labor. The 2,187 who operate none rent in no land, and 19 percent sell farm day labor. Sections 7 and 8 use this split descriptively. The land-poor as a group are net sellers of labor and net renters of land. The top tercile is a net employer and a net landlord whose farm holds the capital a flood can reach. These positions are the channels through which a flood can transmit its loss from one group to another. A farm that loses its animals or equipment hires less labor and leases less land. Whether the price or the quantity adjusts in each market determines who bears the loss.

Relief and recovery programs after a disaster are often targeted by poverty. Bangladesh's relief programs select beneficiaries by poverty criteria. Flood relief is often targeted to Vulnerable Group Feeding (VGF) cardholders, about 9.9 million beneficiaries selected based on a poverty means test. A beneficiary must meet four of twelve criteria, among them having no land beyond a homestead and depending on daily wage labor. After the 2017 floods the program gave 30 kilograms of rice a month for about five months to each listed household

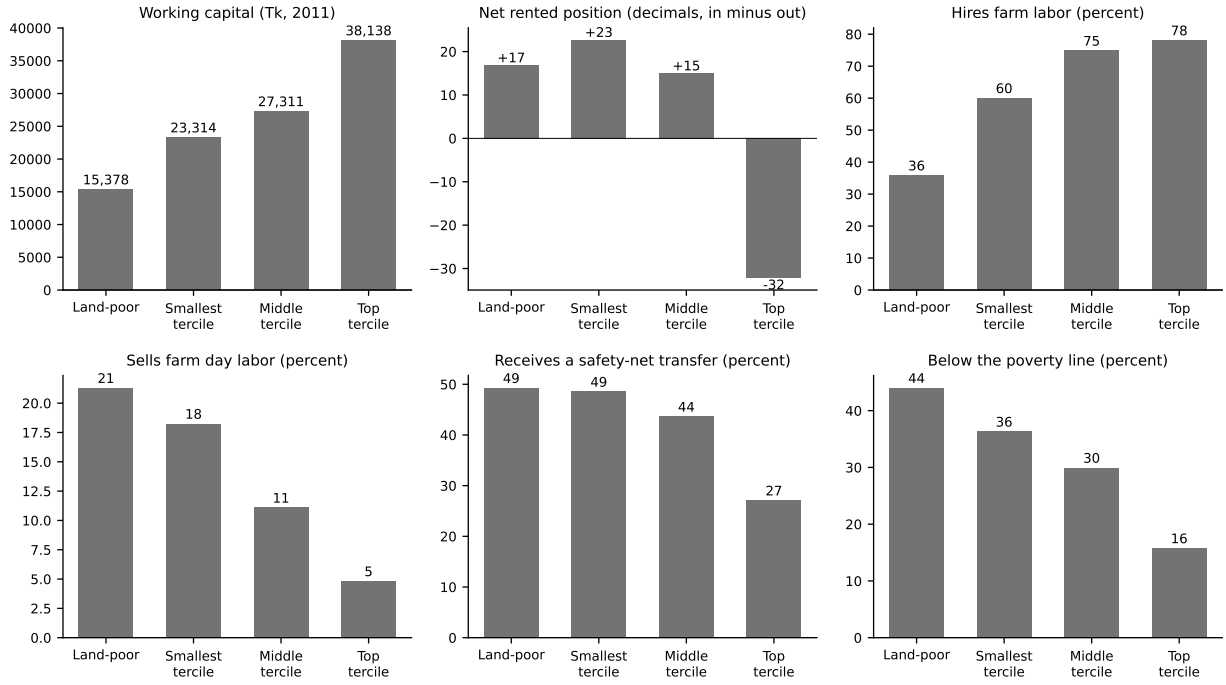


Figure 1: Market positions and relief by inherited land. *Notes:* The 5,473 households of the balanced panel at round 1, by the cultivable land they inherited before 2006: none (the land-poor, 3,845 households) and terciles of the inheritors (548, 541 and 539). Working capital is the value of livestock and productive assets in round-1 taka. The net rented position is land rented and mortgaged in minus land rented out, in decimals. Hiring and selling labor are the shares of households that paid for farm labor in the year and that sold a day of agricultural labor in the week before the interview. A safety-net transfer is any receipt from a public safety-net program in the year. The poverty line is the relative line of Section 4.

(del Ninno et al. 2003; World Bank 2019). The 1998 transfers were well targeted to the poor and poorly targeted to the households the flood actually affected (del Ninno et al. 2001). The same design is common elsewhere. The 2022 floods in Pakistan displaced some 33 million people and were projected to raise the poverty rate by about four points. The government paid a one-time PKR 25,000 to households selected by their score in the national socioeconomic registry, a proxy-means test of poverty, rather than by what they had lost (Government of Pakistan and Asian Development Bank and European Union and United Nations Development Programme and World Bank 2022). Relief targets the poor whether or not the flood water affected them.

The pace of recovery from floods is often slow. After the 1998 flood, rural wages fell and stayed lower in the labor markets most exposed (Mueller and Quisumbing 2011). Households carried the debts they had taken on to eat for years (del Ninno et al. 2003). Assets and consumption were lower a year later among those who had coped by selling or borrowing (Khandker 2007). A year after the 2022 floods in Pakistan, eight million people in the flooded districts still had no safe water. More than a million and a half children needed nutrition support (UNICEF 2023). Across countries, disasters leave incomes below their pre-disaster trend for years, often through a small but persistent fall in growth (Hsiang and Jina 2014).

In this paper I investigate whose capital and earnings a flood reduces and whether the labor and land markets transmit the loss through prices or ration quantities. I also investigate how the affected households adapt. In a model whose premises come from these facts, it asks which policy would restore what the flood destroys.

3 Conceptual Framework and Predictions

This section models a village with a market for rented land and a market for hired farm labor. A flood destroys part of the working capital of some farms, and the model traces how that loss reaches other households through the two markets. I use the framework to derive testable predictions, which are listed in Section 3.4, and to write a household's loss in terms of its measured income. Online Appendix OA-7 sets out the full model, the Theory Appendix gives the proofs, and Section 11 calibrates the model.

3.1 Set-up

Households and production. The village has I households, indexed by i. Each household is endowed with owned land $L_i \geq 0$, working capital $K_i \geq 0$, its livestock and productive assets, and one unit of family labor. Production is Cobb-Douglas in working capital, land

and labor. A farm that operates land T with working capital K and labor N produces

$$F(T, K, N) = A\, K^{\kappa} T^{\alpha} N^{\beta},$$

where A is productivity and $\alpha + \beta < 1$, so that returns to land and labor are decreasing. Land can be rented and labor hired, but working capital cannot be rented, as in Eswaran and Kotwal (1986). A household therefore farms with its own working capital, and that capital determines the size of its farm. At a rent r and a wage w, a household with working capital K wants to operate

$$T^{d}(r, w; K) = \Theta\, K^{\gamma} r^{-a} w^{-b}, \qquad \gamma = \frac{\kappa}{\sigma}, \quad a = \frac{1-\beta}{\sigma}, \quad b = \frac{\beta}{\sigma},$$

where $\sigma = 1 - \alpha - \beta$ and Θ is a constant (Online Appendix Lemma OA-1). The parameter γ is the elasticity of the desired farm size with respect to working capital, and a is its elasticity with respect to the rent. A household rents in land when T^{d} exceeds the land it owns and rents out land when T^{d} falls short. With a transaction cost on leases, which the lemma allows, there is a range of own holdings within which a household neither rents in nor rents out. I set $\gamma = 1$, so that the land a household wants to operate is proportional to its working capital, the constant-returns case of span-of-control models (Lucas 1978; Adamopoulos et al. 2022; Chen et al. 2022).

The land market. I assume that leases are written anew each season at the going rent, which clears the market for rented land. The rent therefore responds to a flood from the first season whose leases are written after it.

The labor market. A household divides its family labor between farm work, on its own farm or for hire at the wage w, and non-farm self-employment, whose return is concave. I assume that a household works its own farm before it hires anyone, the order that any cost of supervising hired workers would impose (Eswaran and Kotwal 1986). A household hires labor when its farm needs more than its own time and offers days for hire when its farm needs less (Online Appendix Assumption OA-6(i)). Denote by η the elasticity of the village's supply of farm labor and by $\bar{\varsigma} > 0$ the share of hired labor in its farm labor. The wage is set by a norm. In a season the wage stays at its going level as long as the share of offered days that goes unsold is at most a tolerance $\bar{u} \in [0, 1]$. If more than that share would go unsold, the wage falls until the unsold share equals $\bar{u}$, and every seller sells the same share of the days it offers. With $\bar{u} = 0$ the labor market clears, which is the benchmark case. With $\bar{u} = 1$ the wage never falls, the downward rigidity that Kaur (2019) documents.

The flood. A severe flood strikes after the season's leases are written. In that season it lowers every farm's output by a share ζ. The flood also destroys working capital. From the

next season on, a farm that has lost capital wants to operate less land, since at given prices $d \ln T^d = \gamma \, d \ln K$. Let x_i denote the change in household i's desired scale, $d \ln T_i^d$, at the pre-flood prices. Then $x_i < 0$ for a household that loses capital and $x_i = 0$ for one that does not. The same shift describes a household that withdraws from farming for another reason, and the data cannot distinguish the two. For a tenant that operates T_i and owns L_i, the shift in the land it wants to rent in is $\tilde{x}_i = x_i \, T_i/(T_i - L_i)$, larger than x_i because the tenant's own land does not change.

A flood may also lower what every operator expects a season on the land to be worth, for reasons the model does not specify. I write this as a common factor $\mathcal{V}_t \leq 1$ on the value of a season. It lowers every operator's desired scale in the same proportion, adding $a \ln \mathcal{V}_t$ to every x_i, so it lowers the rent but leaves the allocation of land unchanged (Online Appendix Proposition OA-2).

Farm size and capital. A farm that loses capital wants to operate less land, so it rents in less land or rents out more. Its demand for labor falls with its capital and its land. Because it staffs its farm with family labor first, its hired labor falls proportionally more than its farm labor (Online Appendix Proposition OA-6(i)).

3.2 The Labor Market

Let $\chi < 0$ denote the log shift in the village's demand for farm labor at the going wage, holding the season's allocation of land fixed. In the season of the flood $\chi = -\zeta/(1-\beta)$.

Proposition 1 (Rationing under a wage norm)**.** *To first order, the share of offered days that goes unsold is* $u = \min\{-\chi/\bar{\varsigma}, \ \bar{u}\}$*, and the wage changes by*

$$d \ln w = - \frac{(1-\beta)\,[\,-\chi - \bar{u}\,\bar{\varsigma}\,]^+}{1 + (1-\beta)\eta} \leq 0.$$

(i) A small fall in the demand for labor reduces hired days, not the wage. If $-\chi \leq \bar{u}\bar{\varsigma}$*, the wage does not fall and a share* $-\chi/\bar{\varsigma}$ *of the days offered for hire goes unsold.*

(ii) Rationed households move their unsold days to their own work. Every seller sells the same share of its offer and no rationed household hires labor. A landless worker moves every unsold day into non-farm self-employment, and no rationed household spends less time in self-employment than it would without rationing at the going wage. The total working days of the landless and of the land-poor households that farm do not change.

Proof: See the Theory Appendix.

In the benchmark case $\bar{u} = 0$, a fall in the demand for labor lowers the wage by $(1-\beta)(-\chi)/[1+(1-\beta)\eta]$, and the fall is larger when the supply of labor is less elastic. Under the norm the wage holds as long as the unsold share stays within the tolerance. Hired labor is a share $\bar{\varsigma}$ of farm labor, and the farms that hire put all their family labor on their own farms, so a fall χ in the demand for farm labor is a fall of $\chi/\bar{\varsigma}$ in the demand for hired days. At the going wage these days go unsold. The loss of work therefore falls on the households that sell labor, and they move the unsold days to their own farms and to self-employment.

3.3 The Land Market

Let $\mathcal{Q}$ denote the total land rented in the village, D_i the land tenant i rents in, and $q_i = D_i/\mathcal{Q}$ its share. For a household whose rented land changes by a large share, a first-order approximation is poor. Let $\mathcal{Z}$ be a set of such households, in the application the top tercile. I take as given their change ΔD_i in net rented-in land, the land they rent in less the land they rent out, and write $\Delta D_{\mathcal{Z}} = \sum_{i\in\mathcal{Z}} \Delta D_i$ and $q_{\mathcal{Z}} < 1$ for their share of rented land before the flood. The land they rent out is counted in supply at its pre-flood amount, so any land they rent out in addition enters $\Delta D_{\mathcal{Z}}$ as a fall in their net demand. The remaining, residual supply of land for rent is $\mathcal{S}$, with elasticity ε_S. The market-wide shift in the demand for rented land is

$$X = \frac{\Delta D_{\mathcal{Z}}}{\mathcal{Q}} + \sum_{i\notin\mathcal{Z}} q_i \tilde{x}_i.$$

A tenant's demand for rented land has rent elasticity $\varepsilon_i = a\,T_i/(T_i - L_i)$, which equals a for a landless tenant and exceeds a for a tenant that also owns land, and $\varepsilon_D = \sum_{i\notin\mathcal{Z}} q_i \varepsilon_i$ is the share-weighted elasticity of the other tenants. For a group g of tenants, D_g, T_g and L_g denote the sums of its members' rented-in, operated and owned land.

Proposition 2 (The rent and the demand for rented land)**.** *Consider a season whose leases are written after the flood, in which households expect, when the leases are signed, the wage on which the pre-flood leases were written. Start from a regular pre-flood equilibrium, in which no household's own land lies exactly at the edge of the range within which it neither rents in nor rents out.*

(i) A group of tenants whose demand has not shifted rents in at least a percent more land for each percent fall in the rent. Suppose the shock is small enough that no household outside $\mathcal{Z}$ starts or stops renting, and let g be a group of tenants outside $\mathcal{Z}$, all renting in, whose working capital and value of a season are unchanged, so that $x_i = 0$ for every member. Then, to first order, $d\ln D_g = -\varepsilon_g\, d\ln r$, with $\varepsilon_g = a\,T_g/(T_g - L_g) \geq a$. For a group of landless tenants, $\Delta\ln D_g = -a\,\Delta\ln r$ exactly.

(ii) *A group that rents in less than this has shifted its own demand inward. Suppose the rent falls, and let* $\varepsilon_g^m = -\Delta \ln D_g / \Delta \ln r$ *be a group's measured response. Under the conditions of (i) but allowing shifts* x_i*, a response below* ε_g *implies that the group's demand for rented land at the pre-flood rent shifted inward, to first order by* $(\varepsilon_g - \varepsilon_g^m)\,|d \ln r|$ *log points. For a fall in the rent of any size, and for any group of households with working capital that rents in land before the flood, including members that start or stop renting, a response below* a *implies that the group's demand for rented land at the pre-flood rent fell by at least* $(a - \varepsilon_g^m)\,|\Delta \ln r|$ *log points, so that* $x_i < 0$ *for at least one member.*

(iii) *The rent moves with the market-wide shift in demand. Under the conditions of (i), to first order,*

$$d \ln r = \frac{X}{\varepsilon_D + \varepsilon_S}, \qquad d \ln \mathcal{Q} = \varepsilon_S \, d \ln r.$$

With inelastic supply the rent changes by X/ε_D*, which is at most* $|X|/[a(1 - q_{\mathcal{Z}})]$ *in absolute value.*

(iv) *The land given up by the households in* $\mathcal{Z}$ *alone can lower the rent only a little. Suppose that the flood does not shift the residual supply and that no tenant outside* $\mathcal{Z}$ *shifts its demand inward (*$x_i \geq 0$ *for every* $i \notin \mathcal{Z}$*). Then a fall in the rent of any size satisfies*

$$|\Delta \ln r| \leq \frac{1}{a} \ln\Big(1 + \frac{|\Delta D_{\mathcal{Z}}|}{(1 - q_{\mathcal{Z}})\mathcal{Q}}\Big).$$

Proof: See the Theory Appendix.

Every tenant's demand for rented land has an elasticity of at least a, so when some tenants rent less, a small fall in the rent is enough for the others to take up the land. Part (iii) is the standard incidence formula, in which a shift in demand moves the rent by the shift divided by the sum of the elasticities of demand and supply. Part (iv) extends it to a fall of any size. If the land that the households in $\mathcal{Z}$ give up is small relative to the land the other tenants rent, it lowers the rent by a few log points at most. A larger fall requires the other tenants' demand to shift inward or the owners' supply to shift outward. Parts (i) and (ii) show where such a shift can be seen. A group whose working capital and value of a season are intact takes up land when the rent falls. A group that does not has lost demand of its own, from a loss of its own capital or from a common fall in the value of a season.

3.4 Testable Predictions

The top tercile of inherited land holds the most working capital per household and hires labor most often, and most households that sell farm labor are land-poor (Section 2). Applied to these groups, the model gives three predictions. A flood that destroys the top tercile's working capital will:

P1) Reduce the top tercile's operated land and its hired labor, and reduce the land it rents in or raise the land it rents out.

P2) Under the wage norm, reduce the days the land-poor sell as day labor without lowering the farm wage, and raise the time they spend in self-employment (Proposition 1). In the benchmark case the farm wage falls instead.

P3) Lower the rent by at most a few log points through the top tercile's retreat from rented land (Proposition 2(iii) and (iv)). A larger fall implies an inward shift in the demand of other tenants, and a tenant group whose demand shifted inward will rent in less than a percent more land for each percent fall in the rent (Proposition 2(ii)).

Section 6 tests the first prediction, Section 7 the second and Section 8 the third.

Discussion. Three features of the model limit what the tests can show. First, the model also predicts that a rationed household that farms puts more time on its own farm. Section 7 shows that this part of the second prediction fails. Second, a labor market that clears with an elastic supply of labor also predicts a fall in days with little change in the wage, since workers who can move to their own work at about the same return absorb the fall in demand. The two differ in the size of the fall in days, which Section 7 uses to tell them apart. Third, the model takes the common fall in the value of a season as given and does not say how fast lost capital is replaced. Section 11 chooses both when it calibrates the model.

3.5 Measured Incidence

I measure who bears the loss from a flood by measured income alone. Household i's income is the sum of its components in Table 4, $Y_i = \sum_k Y_i^k$: daily-wage income, salaried income, farming profit, non-farm enterprise income, remittances, safety-net transfers, land rent received and other receipts. A flood changes it by

$$\Delta Y_i = \sum_k \Delta Y_i^k,$$

and I estimate each term. The estimated terms add up to the estimate for total income only approximately, because each component is winsorized separately, so the table reports both.

The model shows where each term comes from. Farming profit contains the crop loss, the output lost while capital is missing, and the household's net rented-in land times the change in the rent. Daily-wage income contains the household's net sale of labor times the change in the wage, and the earnings lost on days that go unsold at the going wage. Self-employment and own-farm income contain what the unsold days earn in their next use. Because those earnings are already in measured income, the measure needs no value for an unsold day, and it does not use the calibrated model (Online Appendix Remark OA-1). Section 6 reports the measure group by group.

4 Setting and Data

This section describes the household panel, the construction of the outcomes and of the inherited-land groups, and the satellite measure of flood exposure.

4.1 The Bangladesh Integrated Household Survey

The household data come from the three rounds of the Bangladesh Integrated Household Survey (BIHS), collected by the International Food Policy Research Institute (IFPRI) in 2011–12, 2015 and 2018–19. The BIHS is the largest nationally representative panel survey of rural Bangladesh.

The unit of flood exposure is the Union (Union Council), the lowest tier of rural local government, which covers several villages. I build the climate and flood panels for the 428 Unions in which any BIHS household is observed in any round. The estimation sample is the balanced panel of 5,473 households observed in all three rounds, in 323 villages and 318 Unions; the Unions are the clusters of the main specification, and the median Union contributes 17 households. Attrition between rounds is 8.2 percent, and Section 10 tests whether it responds to floods. All regressions are unweighted, as is standard in a within-household design with clustering at the level of treatment.

Consumption is IFPRI's harmonized aggregate of food expenditure, recorded item by item over a seven-day recall, and non-food expenditure, expressed per capita. I use the total aggregate, deflated to round-1 (2011–12) Bangladeshi Taka (BDT), and Online Appendix Table OA-35 reports the food and non-food components separately. The poverty line is a relative line fixed in real terms at the weighted 40th percentile of baseline real per capita expenditure, Tk 2,043 per person per month; the headcount falls from 39 percent at baseline to 30 percent in 2018–19 (Table 1). The asset index is an inverse-covariance-weighted average (Anderson 2008) of three components, each standardized against its baseline distribution:

the value of livestock (module K1), the value of the durables and farm equipment that the household reports using in production (modules D1 and D2), and total savings. Operated land is excluded because it is itself an outcome of the flood, and so is the head's schooling. The components enter as real values winsorized at the 99th percentile; Online Appendix Table OA-3 reports the index under principal-component and equal weights.

4.2 Outcomes and the Inherited-Land Groups

The survey records each market through which a flood can transmit the loss. The plot roster (module G) records each plot's area, tenure, mode and year of acquisition, and self-assessed value. From the second round onward it carries every plot of the previous round forward with a status (same, sold, inherited or split out, mortgaged or rented out, eroded, newly bought, newly inherited, newly leased). Owned cultivable land is the area of the plots that the household operates itself, leaves fallow, holds jointly, or rents or mortgages out, excluding plots acquired by lease; operated land adds the land rented and mortgaged in. Harvest value is the value of the crops grown on the plots the household operated, valued at household unit values and, for unsold output, at crop-by-round median prices (Online Appendix OA-3.4). Net crop income subtracts seed, manure and pesticide; the income aggregate of Section 6 also subtracts fertilizer, irrigation, machinery, hired labor and rent. The labor module (module C) records each member's days of work in the previous seven days by type (daily wage, salaried, self-employed, unpaid), which I sum over members to obtain days per week. The migration module (module V1) lists members who have been away for six months or more, with the year, purpose, destination and financing of the move and whether the migrant remits. Shorter absences are not recorded, so seasonal migration is outside what the data can measure. The share of households reporting any enterprise in module N falls from 43 to 27 percent between the second and third rounds while self-employment in module C is flat; I therefore measure the self-employment margin from module C.

I classify households by the cultivable land they inherited before 2006 and still held at the 2011 baseline, which fixes their position in each market before any flood in the comparison. The variable is constructed from the mode and year of acquisition of each plot in the 2011 roster. Because the roster is an inventory of plots owned in 2011, a plot inherited before 2006 and sold, eroded or divided before the baseline interview does not appear in it, so the variable is predetermined with respect to every flood after 2011. For the floods of the two pre-baseline windows (2006–08 and 2009–11) it is predetermined only if those floods did not alter which pre-2006 inheritances survived to be recorded; Online Appendix OA-2.4 tests for the imbalance this could create. When the classifier is redefined on inheritances dated before

2004, so that it is also predetermined with respect to the floods of 1998 and 2004, its tercile cutoffs are unchanged and 97.7 percent of households remain in the same group. Inherited land accounts for three fifths of all owned area in 2011 and predicts 2011 holdings with a slope of 1.0 decimal per inherited decimal. I divide households into those with no inherited cultivable land (70 percent of the sample) and, among inheritors, terciles with cutoffs at 31 and 81 decimals. The households with no inherited land are not landless: they own 16.6 decimals on average in 2011. I therefore refer to them as the land-poor and reserve "landless" for the cited literature and for the model's worker who holds neither land nor capital; the three groups of inheritors are the smallest, middle and top terciles. Section 5.1 explains why I classify by inherited rather than operated land and describes the three-group version reported in the Online Appendix.

The inherited-land groups are not a poverty classification. Baseline poverty overlaps with them but partitions the sample differently: 79 percent of the baseline poor are land-poor and 4 percent are in the top tercile, while 56 percent of the land-poor are non-poor. I do not estimate flood responses by baseline poverty, because a group selected on the 2011 value of an outcome mean-reverts by construction.

4.3 Flood Exposure

A household's exposure is the flooding mapped in its Union, measured from satellite and reanalysis products rather than from respondents' reports. Flood inundation comes from the Global Flood Database (Tellman et al. 2021), which maps the extent of each large flood event from MODIS imagery at 250-meter resolution. I process all 56 events that affected Bangladesh between 2000 and 2018, compute zonal statistics of each event raster over Union polygons net of permanent water, and complete the Union-by-year grid with true zeros (Online Appendices OA-2.2 and OA-1). With the 2018 monsoon completed from Sentinel-1 as described below, 43 percent of Union-years between 2000 and 2018 experience some inundation and 13 percent a severe flood. Heat, drought and excess rain enter every regression as controls; each is defined as a deviation of more than 1.5 standard deviations from the Union's 1981–2010 climatology in ERA5 temperature or CHIRPS rainfall, and Online Appendix Table OA-32 reports their own effects.

The treatment is the flood that is a disaster, called *bonna* in Bengali, rather than the normal monsoon inundation. A severe flood is an event that inundates at least a quarter of the Union's land, and the regressor is the count of severe floods in an exposure window. I use threshold crossings rather than the inundated share because the response is not linear in the share: the top tercile's asset index falls after floods covering at least a quarter of the Union,

while its estimates for smaller floods are positive, one of them significantly (Section 10). The threshold is absolute, and inundation in the delta varies widely across places: 218 of the 428 Unions never reach a quarter in the nineteen years of the record, while 12 percent have a *median* annual maximum at or above it. Patel (2024) instead defines a flood as a departure of at least 20 percentage points from the local normal; Online Appendix OA-10.4 constructs the annual analogue and Section 10 reports the two side by side. I retain the level measure because the loss depends on how much water reaches the fields and the animals rather than on how unusual that water is for the Union. I do not use the number of days a pixel was under water, because cloud cover and revisit time censor it in a way that is correlated with the flood (Online Appendix OA-1.1).

The exposure windows follow the survey calendar, so that each round's window covers the years since the previous interview. The 2009–11 window is attached to round 1, with 2006–08 as its preceding window; the 2012–14 window to the 2015 round; and the 2015–18 window to the 2018–19 round. Panel households experience 0.23, 0.15 and 0.45 severe floods on average in the three windows (Table 1). A window dilutes a flow measured over a short recall period, such as the twelve-month crop recall or the seven-day labor roster. For these outcomes I also count the years with a severe flood in the Union in the 24 months before the household's own interview, and windows of 24 and 36 months serve as a check on the window definition for every outcome (Section 5.1).

The Global Flood Database ends on 20 June 2018, whereas round 3 was interviewed from November 2018 to April 2019 (Online Appendix OA-1.2), so without a further source the Unions flooded in the 2018 monsoon would be assigned to the comparison group. I complete the 2018 monsoon from the Copernicus Global Flood Monitoring product, which maps flood extent from Sentinel-1 radar, over the same Union zones. Radar does not detect water under a crop canopy: on the 2015–17 monsoons it confirms 37 percent of the Union-years that the Global Flood Database classifies as severe. I therefore map each Union's 2018 radar maximum onto the Global Flood Database scale, evaluated at the acquisition density of the calibration seasons. The fill classifies 24 additional Unions as severely flooded in 2018, for 25 of the 428 together with the one the database already records, and it changes only round 3's current-window count and its most recent year (Online Appendix OA-1.2 describes the construction). Without the fill, the top tercile's capital loss and the fall in the rent are larger (Section 10), so the fill attenuates these estimates rather than producing them.

Satellite and self-reported exposure select different places, as Biscaye and Isah (2026) find in a Nigerian panel, and Patel (2024) develops the measurement of floods for Bangladesh at the place level. First, the database catalogs events flagged by news archives, so small and flash floods can be missed. The treatment here, however, is inundation of at least a quarter of

a Union, the scale of event around which the catalog is built. Second, optical sensors cannot see through monsoon clouds. The share of each Union's area with a valid reading nonetheless averages 0.996 across the 12,140 Union-event cells, and each of the 1,772 cells with at least a quarter of the Union under water is at least 99 percent observed. Classification error of this kind attenuates the estimates toward zero: the significant estimates understate the true effects, and the bounds on insignificant ones understate how large a true effect could be. Table 1 reports summary statistics for the balanced panel by round, and Figure 2 maps the sample Unions by their severe-flood exposure in the 2015–18 window.

Table 1: Summary statistics, balanced panel by round

	R1	R2	R3
Real p.c. expenditure (Tk/month)	2,735	2,819	2,985
Calories p.c. (kcal/day)	2,382.21	2,358.40	2,317.37
Poor (relative line)	0.39	0.38	0.30
Asset index	0.01	0.05	0.25
Livestock value (real Tk)	16,703	15,074	22,583
Land operated (decimals)	60.15	61.28	56.25
Productive assets (real Tk)	4,771	5,378	5,294
Savings (real Tk)	14,847	17,604	25,685
Head schooling (years)	3.40	3.53	3.67
Safety-net receipt	0.46	0.44	0.50
Receives remittances	0.23	0.31	0.37
Head non-farm occupation	0.41	0.44	0.46
Severe floods (window)	0.23	0.15	0.45
Droughts (window)	0.00	0.16	0.17
Excess rain (window)	0.08	0.06	0.93
Heat shocks (window)	2.03	0.58	1.16

Notes: The table reports unweighted means of household outcomes and shock counts in the balanced panel of 5,473 households, by survey round (R1: 2011–12, R2: 2015, R3: 2018–19). Monetary values are in real 2011–12 taka and are not winsorized, expenditure and calories are per capita, and land is in decimals (100 decimals = 1 acre). Poor means real per capita expenditure below the weighted 40th percentile of round-1 expenditure, and the asset index is the inverse-covariance-weighted index of livestock value, productive assets and savings (each winsorized at the 99th percentile), in round-1 standard deviations. Safety-net receipt, remittances and non-farm occupation (among heads who work) are shares of households. The shock rows count years in the window ending at each round's interview (2009–11, 2012–14 and 2015–18): a severe flood is a year in which at least a quarter of the Union is under water, a drought or excess rain is a monsoon with rainfall more than 1.5 standard deviations below or above the Union's 1981–2010 mean, and a heat shock is a year with mean daily maximum temperature more than 1.5 standard deviations above it.

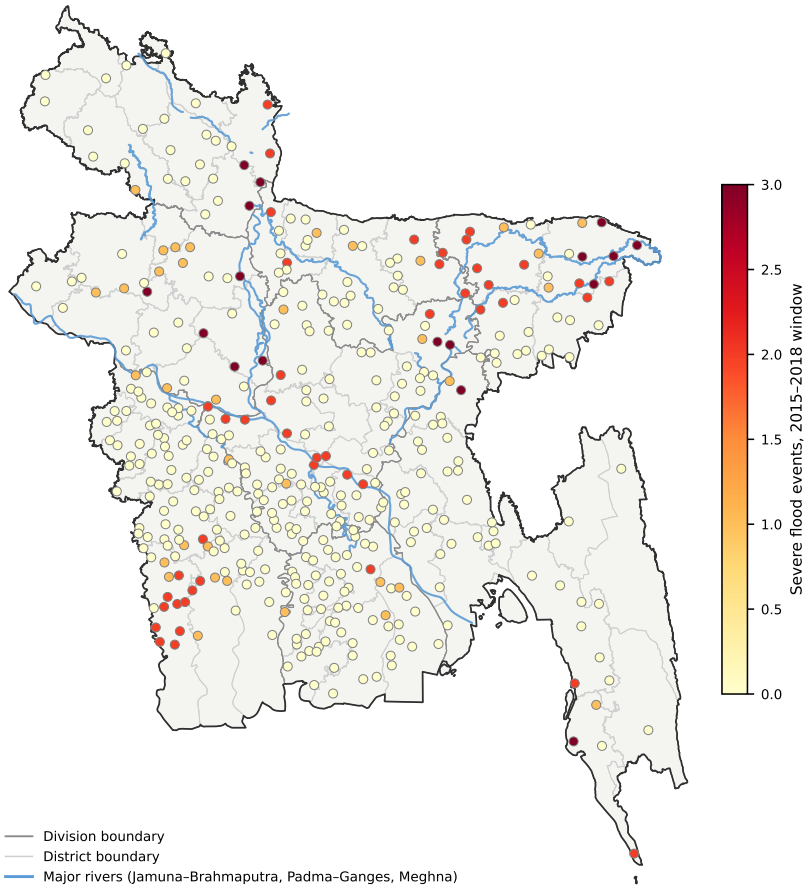


Figure 2: BIHS Unions by severe-flood exposure, 2015–2018 window. *Notes:* Each point is one of the 318 sample Unions, shaded by the number of years with a severe flood (at least a quarter of the Union's land inundated) in the 2015–2018 exposure window attached to round 3. Flood extents are from the Global Flood Database through June 2018, with the 2018 monsoon completed from Sentinel-1 radar.

5 Empirical Strategy and Identification

I estimate the effect of severe floods with a generalized difference-in-differences design. It exploits variation in the number of severe floods across Unions of the same district in the same survey round. I follow the same households over three rounds of survey and compare the changes in their outcomes across Unions that experienced different numbers of floods. Its treatment is the number of severe floods mapped in a household's Union between two survey rounds, interacted with the household's predetermined inherited-land group. The regression is a two-way fixed effects design augmented with district-by-round and inherited-land-group-by-round fixed effects. Household fixed effects absorb permanent differences across households, including their location and their underlying flood risk. District-by-round fixed effects absorb shocks common to the Unions of a district between rounds. Group-by-round fixed effects allow households at different positions in the land distribution to follow separate trajectories. The coefficients are therefore identified from differences in mapped flooding across the Unions of one district in one round. Because the interval between interviews is not of fixed length, I re-estimate the same specification with 2-year and 3-year windows counted back from each household's interview date. A fixed effects regression with a treatment that turns on and off assigns negative weights to some comparisons when effects are heterogeneous. I therefore re-estimate every key result with an interaction-weighted estimator that removes the negative weights arising across rounds and flood histories. Standard errors are clustered at the Union level, the level at which flood exposure is measured, and I also test robustness to a wild cluster bootstrap, Conley spatial HAC standard errors, and two way clustering.

5.1 Main Specification

The main specification is a difference-in-differences regression with household fixed effects and round effects that vary by district and by inherited-land group. That is, it is a two-way fixed effects design augmented with district-by-round and group-by-round fixed effects. The treatment is a count: the number of severe floods in the three- to four-year exposure window before the survey round. The lag is the number in the preceding window of the same kind. Each enters linearly and is interacted with the inherited-land groups. For household i in Union u, district d and round $t \in \{1, 2, 3\}$, let F_{ut} be the number of severe floods in the exposure window attached to round t. Let $F_{u,t-1}$ be the number in the preceding window. I estimate the following equation:

$$y_{it} = \alpha + \beta^0 F_{ut} + \beta^1 F_{u,t-1} + \sum_g \left(\gamma_g^0 F_{ut} + \gamma_g^1 F_{u,t-1}\right) S_{ig} + \sum_g \sum_r \delta_{gr} S_{ig} \mathbf{1}[t = r] + \Gamma' X_{ut} + \psi_{dt} + \alpha_i + \varepsilon_{it}, \quad (1)$$

where y_{it} is the outcome, S_{ig} is an indicator for the household's predetermined inherited-land group (no inherited cultivable land, or the middle or top tercile of inheritors, with the smallest tercile as the omitted category), X_{ut} is a vector of the counts of droughts, excess-rain events and heat shocks in both windows, α_i are household fixed effects, ψ_{dt} are district-by-round fixed effects, and δ_{gr} are group-by-round fixed effects. The coefficient of interest for group g is $\beta^0+\gamma_g^0$, the effect of one additional severe flood in the exposure window, which I call the effect in the flood window. The sum $\beta^1 + \gamma_g^1$ is the effect of a flood in the preceding window, measured one survey round later, which I call the effect in the window after. Standard errors are clustered at the Union level, the level at which treatment varies (318 clusters). Section 10 re-examines the main estimates under a restricted wild cluster bootstrap-t (Cameron et al. 2008), Conley spatial standard errors and two-way clustered standard errors by Union and by survey round (Tables 10 and 11, and Online Appendix Table OA-69).

Without group-by-round fixed effects δ_{gr}, the interaction of the flood count with a group indicator would be identified in part from the difference between that group's average change across rounds and the other groups'. Such a difference can arise from secular convergence or divergence. Because the flood counts also shift across rounds (Table 1), the interaction would attribute it to the floods. With both ψ_{dt} and δ_{gr}, a group's effect is identified only from differences in exposure among Unions of the same district, group, and round. Online Appendix OA-10.7 reports the farm-size results for land, cropping, tenure and work with and without the δ_{gr}.

Monetary outcomes with many zeros enter in levels rather than in logs. Of all household-rounds, 46 percent record no harvest, 37 percent no savings and 19 percent no livestock. With a mass at zero the average effect on $\log(1 + x)$ depends on the unit in which x is measured (Chen and Roth 2024; Bellemare and Wichman 2020), so I do not use that transformation. Livestock value, savings, harvest value, input spending and debt enter as their winsorized real level, estimated by OLS in equation (1), and their effects are in Bangladeshi taka. Where the text gives a percent, it is the effect relative to the group's round-1 mean. Outcomes with few zeros, namely consumption, calories and the prices, are transformed into logs.

The recurrent nature of flooding motivates this choice of estimator. Unlike a policy that stays in place once adopted, a severe flood need not persist into the next period, and a

Union may experience several floods between two interviews. The estimators developed for staggered adoption (Callaway and Sant'Anna 2021; Sun and Abraham 2021; Borusyak et al. 2024) organize treatment around the date of first adoption and are not defined for a treatment that turns on and off. Recoding the treatment as ever having been exposed would estimate the effect of a first flood rather than of a flood. The question here is how an additional flood affects household outcomes. An event study around the first flood would likewise combine the effects of later floods with those of the first. Moreover, household outcomes are observed in only three rounds, which leaves little scope for estimating separate pre- and post-event trajectories. I therefore use the fixed effects specification with one lag of equation (1), which distinguishes floods in the current exposure window from floods in the preceding window.

Floods are not an absorbing treatment. A Union flooded in one window is often dry in the next, and the regression then uses a Union exposed in both windows as the control for one just flooded. The fixed effects estimator is therefore a weighted average of per-flood effects in which some weights are negative (de Chaisemartin and D'Haultfœuille 2020), mostly on households flooded in every round. In the main specification the negative weights on a group's flooded household-rounds sum to between -0.22 and -0.30 for the main outcomes of interest and to -0.34 for the Union median rent (Appendix Table 15). I re-estimate every key result with a simplified version of the interaction-weighted estimators of Gibbons et al. (2019) and Sun and Abraham (2021). It interacts the flood regressor with cells defined by inherited-land group, round and flooding in the other window, keeps every household, and averages each group's cell coefficients with weights equal to the group's share of floods in each cell. It removes the negative weight that comes from the composition of flooded Unions across rounds and flood histories, but not the weight within a household flooded in every round (Online Appendix Table OA-18). Appendix A describes the estimator and reports a simulation of what the check can and cannot detect.

The exposure window of equation (1) is the interval between two interviews, three years for the first two rounds and four for the third. It is attached to the round rather than to the household's own interview date. As a check on that definition I re-estimate the same specification with windows of fixed length. Let F^2_{it} be the number of years with a severe flood in the household's Union in the 24 months before its interview, and $F^2_{i,t-1}$ the number in the 24 months before that. Similarly, F^3_{it} and $F^3_{i,t-1}$ are the corresponding counts over 36 months. Each pair replaces F_{ut} and $F_{u,t-1}$ in equation (1). The group interactions, the group-by-round and district-by-round fixed effects, the other weather shocks are unchanged, as is the inference. The yearly counts are built from the flood record of Section 4, with round 3's most recent year completed from Sentinel-1 as in the interview window. For the Union-level regressions the fixed windows are counted back from the median interview date

of the Union's sample households in that round. The coefficient on a fixed-window count is the effect of one additional severe flood in that window, the same quantity as in the interview window.

The two fixed windows serve as a robustness check on the window definition, reported in Online Appendix OA-2.5 (Supplementary Tables S-5 to S-9). Furthermore, for some flow outcomes, the 2-year window is the horizon on which the effects emerge. For example, a flooded crop is lost in one season, and hiring, days of work and the wage are recorded for the year or the week before the interview. A window of three or four years therefore dilutes a flow outcome. Table 5 reports the flow outcomes in both windows. Supplementary Tables S-5 to S-9 report every key outcome under the interview, 3-year and 2-year windows, in the current window and the window after the flood.

The comparison groups are defined by inherited land rather than by baseline land ownership and poverty. I partition the sample into 3,845 land-poor households and 548, 541 and 539 households in the smallest, middle and top terciles of land inheritors. The four inherited-land groups of Section 4.2 are the main classification, because inherited land is predetermined (Section 4.2) while operated land at baseline responds to prior floods. The tercile cutoffs are a choice. The key result tables therefore have a three-group version in the Online Appendix (Tables OA-85, OA-86 and OA-29). It keeps the land-poor and splits the inheritors at their median of 51 decimals, into 803 households below and 825 at or above it. Section 10 reports what changes under the three groups and addresses the remaining threats to identification.

5.2 Identification Strategy

The design compares neighbors. Take a household whose Union experienced more severe floods between two interviews than other Unions of its district. The change in its outcomes over that interval is compared with that of households in those other Unions. Large monsoon floods cover very different shares of neighboring Unions in the same season (Section 2), and the estimates use only that within-district contrast. Each set of controls removes one alternative source of it. The household fixed effects absorb all time-invariant differences across households, including baseline landholding, location and the Union's flood proneness. The district-by-round fixed effects absorb any shock common to a district between two interviews, such as prices or district-wide programs. The group-by-round fixed effects absorb each inherited-land group's own trend, so that households at different positions in the land distribution may follow separate trajectories. The drought, excess-rain and heat counts of both windows keep the flood coefficients from capturing the other weather of the same years.

The estimates therefore use differences in exposure among neighboring Unions rather than differences in their long-run propensity to flood.

The identifying assumption is parallel trends. Conditional on these fixed effects and the other weather shocks, households in Unions with different severe-flood counts would have followed the same outcome trajectories in the absence of those floods. Outcomes in the more and the less flooded Unions of a district would thus have evolved in parallel. The assumption is plausible because whether a Union floods in a given window is determined by where the monsoon rain falls and where the rivers overflow. It is not determined by anything its households chose. This is the source of variation used throughout the weather-panel literature (Dell et al. 2014). It would fail if flooded Unions were already on different trajectories; Section 10 tests for pre-trends with floods that came later and with a group-specific trend in later exposure. Online Appendix Tables OA-80 and OA-44 replace the district-by-round fixed effects with district linear trends. The Union-level and plot-level regressions have the same structure, with Union or plot fixed effects in place of household fixed effects.

After all the fixed effects are partialled out, 17.7 percent of the raw variance of the severe-flood count remains (Online Appendix Table OA-4). Of the 318 Unions, 86 change their severe-flood count across the three rounds; the other 232 are absorbed by the household fixed effects and serve only as controls. Figure 3 maps the 86 switching Unions. Within rounds, 75 district-by-round cells have variation in the flood count, and they contain 3,459 of the 5,473 households (63 percent). The 2016 monsoon provides the flooding events for 49 of the 86 Unions.

6 Effects on Capital, Farm Size and Earnings

This section estimates whose capital and earnings a flood reduces, the paper's first question. The loss falls on capital at the top of the inherited-land distribution, and the largest farms shrink with it. The top tercile loses livestock and working capital, and its working capital, herd and asset index remain lower at least 3 years after the flood. It gives up the land it sharecropped in, then farms less of its own land and rents more of it out. Its wage and self-employment income falls significantly more than the land-poor's. The land-poor lose day labor and, in the window after, owned land (Sections 7 and 9). But their combined wage and self-employment income, total income and consumption fall by at most a few percent.

A flood's most direct loss would be the standing crop, but the window design is poorly suited to detect it. The loss falls on the harvest of the season the flood strikes, the production modules record only the twelve months before the interview, and the windows span



Figure 3: The Unions whose severe-flood count changes across rounds. *Notes:* The 318 sample Unions in gray. The 86 Unions whose current-window severe-flood count changes across rounds, and which therefore identify the current-window coefficients, are colored by the windows in which they record a severe flood (a year with at least a quarter of the Union inundated; windows 2009–11, 2012–14 and 2015–18, with the 2018 monsoon from Sentinel-1). Unions with a severe flood in every window but a changing count are drawn in the color of their last window.

three to four years, so a window estimate dilutes a loss confined to one harvest (Online Appendix OA-1.3). Harvest values are also self-reported and measured with error, which widens the standard errors. In the window design the crop loss is imprecise. Per severe flood in the 2-year window, the land-poor's harvest net of seed, manure and pesticide is Tk 1,414 lower (standard error 817), significant only at the 10 percent level, and the top tercile's estimates are smaller and imprecise (Online Appendix Table OA-37). The land-poor's crop profit, which also nets out the rent and the hired labor they pay, is a precisely estimated zero in the flood window, within 2 percent of a year's total income (Online Appendix Table OA-28). I therefore also time the flood to the household's own recall year, the harvest the flood strikes, and estimate the loss among the households that farm (Online Appendix Table OA-20). Among households that grew a crop in round 1, a severe flood in the recall year lowers the harvest by 8.9 percent of the group's round-1 mean (standard error 3.4, wild bootstrap p 0.019), and in rounds 1 and 2 a flood over half the Union lowers it by a further 22 percent (Online Appendix Table OA-21). The data cannot reject that every group loses the same share of its harvest. That additional loss is about a quarter of the harvest for the land-poor and the top tercile alike. At a common share a larger farm loses more in taka. In a flood year the top tercile's loss is then Tk 7,344 against the land-poor's Tk 3,619, a gap of Tk 3,725 (standard error 1,426). The gap rests on the common share, because the top tercile's own estimate, Tk 4,926 (standard error 4,481), is too noisy to test it directly. Among households that grew a crop in the same round the loss is smaller and imprecise, and over all households, including those that grow nothing, the top tercile's loss is a smaller share than the land-poor's. The crop loss is therefore confined to one harvest and proportional to its size, as the framework of Section 3 assumes. The loss of capital, by contrast, persists.

Table 2 reports the effects of a severe flood on the balance sheet by inherited-land group. The top tercile, which holds most of the capital a flood can reach, loses livestock and working capital in the flood window. Its asset index, an inverse-covariance-weighted average of livestock value, productive assets and savings, falls by 0.25 standard deviations ($p = 0.002$). The value of its livestock falls by Tk 5,410 ($p < 0.001$), 18 percent of its round-1 stock of Tk 30,395. The fall is concentrated in large animals: the tercile holds 0.28 fewer cattle and buffalo ($p < 0.001$) on a base of about two head (Table 2; Online Appendix Table OA-81). Its working capital, the value of livestock and productive assets together, falls by Tk 6,152 ($p < 0.001$; Table 2, column 3), 16 percent of a baseline stock of Tk 38,138. Livestock accounts for most of the fall (Online Appendix Table OA-25). The interaction-weighted estimator gives losses at least as large (Section 10). The loss arises mostly where floods repeat: for a flood with no flood in the other window the estimate is not statistically distinguishable from zero (Appendix A). The tercile's savings are Tk 6,661 lower in the flood window, 40

percent of their round-1 mean. That estimate is significant at the 10 percent level ($p = 0.085$; Online Appendix Table OA-81), with a wild bootstrap p-value of 0.113 (Section 10). The effect on its productive assets is imprecisely estimated.

Table 2: The balance sheet after floods by predetermined farm size

	Asset index (SD) (1)	Livestock value (Tk) (2)	Working capital (Tk) (3)	Cattle and buffalo (head) (4)	Savings (Tk) (5)
Flood window					
No inherited land	+0.012	−61	−420	−0.005	+1,947*
	(0.024)	(795)	(895)	(0.036)	(1,149)
Smallest tercile	+0.076	+168	+504	+0.012	+3,790
	(0.059)	(1,562)	(1,586)	(0.083)	(2,462)
Middle tercile	+0.021	+669	+873	+0.050	−178
	(0.084)	(1,717)	(2,261)	(0.096)	(3,050)
Top tercile	−0.252***	−5,410***	−6,152***	−0.282***	−6,661*
	(0.082)	(1,399)	(1,561)	(0.076)	(3,859)
Window after					
No inherited land	+0.000	+634	+35	+0.022	+887
	(0.025)	(577)	(704)	(0.036)	(1,032)
Smallest tercile	+0.117**	+829	+1,496	+0.080	+4,531**
	(0.051)	(1,401)	(1,431)	(0.078)	(2,177)
Middle tercile	+0.097	+1,438	+2,288	+0.089	+1,824
	(0.062)	(1,306)	(1,505)	(0.084)	(2,780)
Top tercile	−0.142**	−2,309	−3,774**	−0.257*	−1,683
	(0.071)	(1,605)	(1,800)	(0.143)	(3,031)
N	16,419	16,419	16,419	16,419	16,419

Notes: Each cell is an inherited-land group's own effect of one more severe flood, a year with at least a quarter of the Union under water, in the flood window (between the previous interview and this one) or in the window after. The main specification (equation (1)) has household fixed effects, district × round and group × round effects, and the drought, excess-rain and heat counts of both windows. The asset index is an inverse-covariance-weighted index of livestock value, productive assets and savings in standard deviations of round 1, and working capital is livestock value plus productive assets, in real taka winsorized at the 99th percentile. Groups: no inherited land (3,845 households) and terciles of cultivable land inherited before 2006 among inheritors (548, 541 and 539). Union-clustered standard errors in parentheses; $G = 318$ Unions; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The tercile cuts are a choice, and the top tercile is a group of only 539 households. Online Appendix Table OA-85 repeats Table 2 with the inheritors split at their median, which gives a top group of 825 households. Above the median the asset index falls by 0.17 standard deviations in the flood window ($p = 0.008$), livestock value by Tk 3,753, 14 percent of the group's round-1 stock ($p = 0.002$), and working capital by Tk 4,204 ($p = 0.003$). The flood-window loss therefore holds under both classifications. The window-after estimates above the median are not statistically distinguishable from zero, so the persistence of the loss a

window later is established for the top tercile only. The same holds beyond the balance sheet. Of the 129 comparable estimates of the key outcomes, 113 keep their significance status under the three groups (Online Appendix Tables OA-81 to OA-84), and Section 10 discusses the 16 that change.

The group's loss is concentrated on a minority of its farms. Module K1 records the animals each household holds at the interview, so a fall in the cattle count between rounds is measured from the inventory rather than from a reported loss. Take the top-tercile farms that held cattle at the previous round. A severe flood in the flood window raises the share whose count fell by at least two head by 24.0 percentage points ($p < 0.001$). It raises the share whose count fell by at least one head by 29.3 points ($p < 0.001$). Over all 539 households the share that lost two or more head rises by 11.4 points, also significant, and no other group shows a significant rise (Online Appendix Table OA-23). Splitting the top tercile by whether it held cattle in 2011, which predates the panel's floods, locates the loss where the animals were. In the flood window the 355 holders lose Tk 6,692 of working capital ($p = 0.001$). They draw down Tk 10,419 of savings ($p = 0.040$) on a round-1 mean of Tk 17,371. For the 184 households without cattle the effect on savings is not statistically distinguishable from zero, although their working capital falls by Tk 4,191 ($p = 0.046$). The difference in log consumption per capita between the two is −0.018 (standard error 0.038), and the smallest difference the split could detect is 0.107, so their consumption cannot be distinguished. The loss therefore falls on the farms that hold cattle, which lose head from a herd of about three and draw down their savings. The consumption bound reported below averages over the farms that lost nothing and the minority that lost a herd.

The 2-year window from the interview date separates recent floods from earlier ones, and so dates the top tercile's responses. In the 2-year window the farms cut their hiring at harvest and the land-poor lose days (Section 7). The top tercile's working capital is not yet significantly lower: per flood in that window the estimate is −Tk 3,105 (standard error 1,971). Its livestock value is already lower, by Tk 3,743 (standard error 1,795; $p = 0.038$). In the preceding 2-year window, two to four years before the interview, both are lower. Working capital is −Tk 8,783, livestock value −Tk 6,184 and the asset index −0.38 standard deviations, each significant at the 1 percent level (Supplementary Table S-5). By the end of the flood window, the share of its cattle holders that have lost two or more head is 24 points higher per flood, and the holders have drawn down about Tk 10,400 of savings. Both are significant (Online Appendix Table OA-23). The tercile has also sold 3.5 more decimals of land ($p = 0.008$; Section 9). Its consumption stays within its bound, among the cattle holders +1.1 percent with a standard error of 2.9. A window later its capital is still lower (below) and its savings are not statistically distinguishable from zero (−Tk 1,683, standard error

3,031; Online Appendix Table OA-81). The share of its households growing a crop is then 5.7 points lower (Section 7). Descriptively, milk cows leave before bullocks and equipment (Section 6).

The top tercile moves from production loans toward microfinance, while the change in its total debt is imprecisely estimated. The share of its households with an NGO or microfinance loan rises in the window after by 5.8 points on a round-1 base of 17 percent ($p = 0.012$). The share with a production loan falls in the flood window by 5.3 points, significant at the 10 percent level ($p = 0.078$; Online Appendix Tables OA-48 and OA-24), with a wild bootstrap p-value of 0.098 (Online Appendix Table OA-70). Its debt outstanding changes by +Tk 2,547 (standard error 3,567) in the flood window and +Tk 1,905 (standard error 4,112) in the window after, neither statistically distinguishable from zero, and neither is the change in whether it holds a loan at all.

I interpret this timing as a household that draws on its savings, its credit, its land and its herd to hold its consumption, and that does not rebuild within the panel. This is an interpretation, not a measured fact. There is no evidence that it sells assets to repay debt, since its debt shows no fall. The sale of the animals is not observed. Interviews are three to four years apart, and the module K1 questions on animals lost, sold and bought cover only the twelve months before the interview. In those months there is no detectable flood effect on the top tercile's cattle sales (Online Appendix Table OA-53) or on the cattle it reports lost (not statistically distinguishable from zero in either window; not tabulated). Animals that left in the unrecorded years are therefore unobserved, and the interpretation rests on the stock, the savings, the land sales and the loans.

The land-poor's borrowing shows no such shift. Their debt outstanding and the shares with any loan, an informal loan, a production loan and a consumption loan are not statistically distinguishable from zero in the flood window or in the window after. The debt estimates are −Tk 2,174 and −Tk 1,151, standard errors 1,808 and 1,564, on a round-1 mean of Tk 19,630 (Online Appendix Table OA-24).

The top tercile's capital remains lower in the window after a flood. Its asset index is still 0.14 standard deviations lower ($p = 0.047$). Its working capital is Tk 3,774 lower ($p = 0.037$), 10 percent of its round-1 stock. Its herd is 0.26 head smaller. All three are significant at the 10 percent level or better (Table 2). Its livestock value is Tk 2,309 lower, an estimate that is not statistically distinguishable from zero (standard error 1,605). Supplementary Table S-5 shows the same persistence under the 3-year window and, as the lagged response, in the preceding 2-year window.

The land-poor's capital does not fall significantly in either window. In the flood window the 90 percent interval for their asset index rules out a fall larger than 0.03 standard devi-

ations. Their savings are Tk 1,947 higher, 18 percent of their round-1 mean, an estimate that is significant at the 10 percent level ($p = 0.091$; Table 2) but does not survive the wild bootstrap (Section 10). The 90 percent interval rules out a fall in their savings, and whether they rise cannot be determined.

The data show which assets the top tercile loses, but not whether the water destroyed them or farms that scaled back disposed of them (Online Appendix OA-3.2). Its milk cows are 0.20 head lower in the flood window, an estimate I treat as descriptive. Its bullocks and agricultural equipment are lower in the window after (Online Appendix Table OA-25). Among top-tercile households that grow a crop in every round, the point estimates of the loss in the flood window are as large as among the rest. The differences are not statistically distinguishable from zero (Online Appendix Table OA-22, Panel A). Whether the loss scales with the usual flood depth of the household's plots depends on the measure (Online Appendix OA-3.2). The animal flows recalled in module K1 do not show where the cattle went. The stock falls in the seasons after the flood, which the twelve-month recall may not cover (Online Appendix Table OA-53). A fall in cattle alongside a bounded fall in consumption is consistent with the buffer-stock role of productive animals in Rosenzweig and Wolpin (1993), but the flows cannot confirm that the animals were sold.

The top tercile farms less after a flood, and its farm shrinks after its capital falls. Online Appendix Table OA-47 decomposes operated land into the land a household owns and farms itself, the land it rents and mortgages in, and the land it rents out. The top tercile gives up 3.2 decimals of the land it sharecropped in, on a round-1 mean of 11.6. It does so in the flood window ($p = 0.088$) and again in the window after ($p = 0.092$; standard error 1.9 in each), where the wild bootstrap p-value is 0.110. In the window after it farms 8.1 fewer decimals of its own land ($p = 0.026$). It rents out 6.5 more ($p = 0.050$) on a round-1 mean of 50.0. Its operated land is therefore 8.3 decimals lower ($p = 0.057$) on a base of 143.7. It does not sell the land: its owned land in the window after is a precisely estimated zero. In the 2-year window it also hires nearly a quarter of a year's harvest hours less per flood (Section 7). Under the seven checks of Section 10, its own-operated and its rented-out land in the window after keep their sign and their significance at the 10 percent level in every check (Online Appendix Table OA-50). The retreat also holds when the inheritors are split at their median. Above it, rented-in land of every kind falls by 5.9 decimals in the flood window, a significant fall (Online Appendix Table OA-86).

The group that loses capital also farms less, as it would if a household's own working capital limits the scale of its farm (Section 3). Capital and farm size are both outcomes of the flood, so the data cannot establish that the loss of capital causes the retreat. They show that capital falls in the flood window and the farm shrinks in the window after. The full

model orders the contracts a tenant gives up by who can terminate them (Online Appendix Proposition OA-5), and the data can test only the sharecropping part of that order. Supplementary Table S-29 reports the top tercile's land by contract in the flood window and the window after and splits the rent by when a tenancy began. The table is descriptive only, because on a given plot the data cannot separate a tenancy the owner ended from one the tenant gave up.

The top tercile's wage and self-employment income falls with its capital. This income, which I call earnings for short, is daily-wage income plus self-employment income, the profit of the farm and of non-farm enterprises (Table 3). Table 4 reports each source of measured income, constructed from every income module of the survey (Online Appendix OA-1.5). Each estimate is a reduced-form effect on measured income. In the flood window the top tercile's wage and self-employment income falls by Tk 10,764 a year ($p = 0.004$), 7.6 percent of its round-1 income of Tk 140,731 a year. The fall is in non-farm enterprise income, which is Tk 8,448 a year lower ($p < 0.001$), 48 percent of its round-1 mean of Tk 17,730. Its daily-wage income is a precisely estimated zero, and the effect on its farming profit is imprecisely estimated (−Tk 1,775, standard error 2,542). In the window after, its non-farm enterprise income is not statistically distinguishable from zero (Online Appendix Table OA-27). Above the median of inherited land, the flood-window estimate for wage and self-employment income is −Tk 7,209 (standard error 4,174), significant at the 10 percent level ($p = 0.085$) but not on the wild bootstrap ($p = 0.131$; Online Appendix Table OA-84). The top tercile's total income falls by less, Tk 8,355 in the flood window, and the estimate is not statistically distinguishable from zero ($p = 0.111$; standard error 5,221). Its 90 percent interval runs from a fall of 12.1 percent of a year's income to a rise of 0.2 percent. How much of the earnings loss other income offsets therefore cannot be determined.

The land-poor's significant loss of measured income is in daily-wage income, and their wage and self-employment income as a whole is bounded within a few percent. Their daily-wage income falls by Tk 2,299 in the flood window ($p = 0.009$), 14 percent of its round-1 mean of Tk 16,944. It is a precisely estimated zero in the window after. In the window after, their remittances and the rent they receive from land they own fall significantly (Online Appendix Table OA-27). For their wage and self-employment income, the 90 percent interval in the flood window rules out a fall larger than 3.8 percent of a year's income and a rise larger than 3.1 percent. The difference between their wage and self-employment income and their daily-wage income is self-employment income. It rises by Tk 1,984 in the flood window, an estimate that is not statistically distinguishable from zero (standard error 1,733; Online Appendix Table OA-40). For their total income the interval rules out a fall larger than 4.6 percent of a year's income.

Table 3: Wage and self-employment income by inherited-land group and the difference between groups

	Flood window		Window after		Round-1
	Taka	Percent	Taka	Percent	income
	(1)	(2)	(3)	(4)	(5)
No inherited land	−315	−0.4	+2,313	+2.6	89,756
	(1,887)	(2.1)	(1,754)	(2.0)	
Smallest tercile	+3,904	+4.1	+6,360**	+6.6**	96,076
	(3,471)	(3.6)	(3,223)	(3.4)	
Middle tercile	−2,343	−2.4	+3,287	+3.3	99,699
	(5,893)	(5.9)	(3,577)	(3.6)	
Top tercile	−10,764***	−7.6***	+1,547	+1.1	140,731
	(3,739)	(2.7)	(3,652)	(2.6)	
Top − none	−10,449***	−7.3**	−766	−1.5	
	(3,794)	(3.0)	(3,548)	(2.7)	
Top − smallest	−14,668***	−11.7***	−4,813	−5.5	
	(4,431)	(3.9)	(4,349)	(3.8)	
Smallest − none	+4,218	+4.4	+4,047	+4.0	
	(3,162)	(3.3)	(3,172)	(3.3)	
Household-rounds	16,419	16,419	16,419	16,419	

Notes: The table reports estimates of equation (1) for wage and self-employment income in the balanced panel of 5,473 households, with one observation per household-round. Wage and self-employment income is daily-wage income plus non-farm enterprise income and farming profit, in real 2011–12 taka per year, with the sum winsorized at the 99th percentile of its absolute value. The treatment is the number of severe floods, years in which at least a quarter of the Union is under water, in the window between two interviews; each row gives an inherited-land group's own effect of one more flood, or the difference between two groups' effects, in the flood window or in the window after. Columns 2 and 4 divide the effect by the group's round-1 mean total income (column 5). All regressions include household fixed effects, district-by-round and group-by-round fixed effects, and the drought, excess-rain and heat counts of both windows. Standard errors clustered at the Union level are in parentheses (318 Unions). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: How one severe flood changes household income by source and predetermined farm size

	No inherited land	Smallest tercile	Middle tercile	Top tercile	All households
Daily-wage income	−2,299***	−1,600	−5,188***	−542	−2,243***
	(879)	(1,942)	(1,885)	(1,236)	(761)
	[16,944]	[12,386]	[7,878]	[4,704]	[14,386]
Salaried income	−438	−1,601	−1,200	−3	−584
	(581)	(1,083)	(1,849)	(1,205)	(524)
	[4,949]	[6,012]	[6,975]	[8,255]	[5,581]
Farming profit	+896	+2,621	+1,258	−1,775	+845
	(1,052)	(1,840)	(4,971)	(2,542)	(1,046)
	[16,978]	[22,385]	[32,939]	[64,494]	[23,777]
Non-farm enterprise income	+1,088	+2,882	+1,586	−8,448***	+132
	(1,199)	(2,046)	(2,624)	(2,098)	(1,005)
	[26,771]	[19,031]	[19,598]	[17,730]	[24,396]
Remittances received	+1,038	+2,131	−10,947***	+2,768	+428
	(1,462)	(3,237)	(3,568)	(3,446)	(1,273)
	[14,321]	[23,401]	[19,478]	[25,250]	[16,817]
Safety-net transfers	−53	−29	+1	−268**	−75
	(78)	(137)	(127)	(126)	(67)
	[1,121]	[1,143]	[809]	[730]	[1,054]
Land rent received	−107	−309	+468	−662	−155
	(107)	(381)	(547)	(530)	(137)
	[1,034]	[1,045]	[2,062]	[7,672]	[1,790]
Other receipts	−140	+478	+245	+953	+64
	(299)	(630)	(866)	(770)	(310)
	[7,639]	[10,673]	[9,959]	[11,897]	[8,591]
Total income	−191	+4,073	−15,248	−8,355	−1,979
	(2,363)	(5,237)	(9,268)	(5,221)	(2,263)
	[89,756]	[96,076]	[99,699]	[140,731]	[96,392]
Household-rounds	16,419	16,419	16,419	16,419	16,419

Notes: The table reports estimates of equation (1) for household income by source in the balanced panel of 5,473 households, with one observation per household-round. Sources are defined in Online Appendix OA-1.5, measured over the twelve months before the interview in real 2011–12 taka per year, and winsorized at the 99th percentile of their absolute value. The treatment is the number of severe floods, years in which at least a quarter of the Union is under water, in the window between two interviews; each cell gives an inherited-land group's own effect of one more flood in the flood window. Brackets give the group's round-1 mean, and the last column pools all households without the group terms. All regressions include household fixed effects, district-by-round and group-by-round fixed effects, and the drought, excess-rain and heat counts of both windows. Standard errors clustered at the Union level are in parentheses (318 Unions). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The top tercile's loss is significantly larger than the land-poor's. Each group's own baseline income is Tk 140,731 a year for the top tercile and Tk 89,756 for the land-poor. Measured in percent of it, the top tercile's earnings fall by 7.3 percentage points more than the land-poor's in the flood window ($p = 0.015$; Table 3). In taka they fall by Tk 10,449 more ($p = 0.006$). The same gap in total income is not significant: the top tercile's total income falls by 5.7 points more than the land-poor's ($p = 0.194$; Online Appendix Table OA-31). Its other income partly offsets its earnings in the point estimates. The significant difference between the groups is therefore in wage and self-employment income, not in total income.

There is no evidence that the top tercile borrows more in total to cover the loss. Its debt outstanding and whether it holds a loan are not statistically distinguishable from zero in either window (Online Appendix Table OA-24). Its borrowing does, however, shift from production loans toward microfinance (Section 6, the sequence of effects). The estimate for its production loans in the flood window is negative and significant at the 10 percent level ($p = 0.078$; Online Appendix Table OA-48), with a wild bootstrap p-value of 0.098 (Online Appendix Table OA-70). The share refused credit for lack of collateral rises by 1.3 points on a base of 0.9 percent ($p = 0.056$; Online Appendix Table OA-48). Collateral rationing rises in the other groups too, and Section 8 takes it up together with the land-poor's production loans.

No group's consumption estimate is statistically distinguishable from zero, so I report each by the bound its interval sets (Online Appendix Table OA-34). For the land-poor, the 90 percent interval for log consumption per capita rules out a fall larger than 0.05 percent in the flood window and 1.5 percent in the window after. That is consistent with their earnings bound. For the top tercile it rules out a fall larger than 1.8 percent in the flood window and 2.0 percent in the window after. For the smallest tercile it rules out a fall larger than 4.4 percent in the flood window. None of these is a precisely estimated zero within 2 percent. The top tercile's consumption is bounded while its earnings and capital fall significantly and its cattle holders draw down their savings. This is consistent with a household drawing on its savings to hold its consumption, but the estimates cannot show which stock finances it. The holders of cattle, on whom the loss falls, draw down their savings in the flood window. The bound averages over the holders that lost animals and those that did not (Online Appendix Table OA-23). Food and non-food spending show no offsetting cut. Estimated separately (Online Appendix Table OA-35), no group's food or non-food expenditure per capita falls in either window. Food expenditure rises: the land-poor's by 2.6 percent in the flood window ($p = 0.032$). The top tercile's rises by 4.2 percent in the flood window and again in the window after ($p = 0.086$ and 0.069). All are significant, and the wild bootstrap p-value of the top tercile's flood-window estimate is 0.100 (Online Appendix Table OA-70). Non-food

expenditure is not statistically distinguishable from zero in either group in either window. With the Union rice price a precisely estimated zero in the flood window (Section 8), the rise in food spending is consistent with households protecting what they eat. The consumption bounds use total expenditure.

The data bound how far a flood can push households into poverty. On the full balanced panel (Online Appendix Table OA-61), the 90 percent interval for the flood-window effect on the probability of being poor runs from -2.7 to $+1.0$ points. It thus rules out a rise larger than 1.0 point. Among households poor at baseline it rules out a rise larger than 2.4 points (Online Appendix Table OA-62). The window-after estimate, -2.4 points, is significant at the 10 percent level ($p = 0.075$), with a wild bootstrap p-value of 0.114 (Online Appendix Table OA-70). I do not read it as evidence that floods pull households out of poverty. I do not estimate transition regressions. With three rounds, a household is at risk of leaving poverty in both intervals only if it is poor in the first two rounds. Being poor in the second round is itself an outcome of the floods in the first interval. Nor do I estimate the asset transition functions of the poverty-trap literature (Carter and Barrett 2006; Balboni et al. 2022). Those identify thresholds from differences in assets across households rather than from each household's response to a shock.

The smallest tercile's capital does not fall significantly in either window, and it rises in the window after. Its asset index is then 0.12 standard deviations higher ($p = 0.022$; Table 2). Its productive assets are 29 log points higher ($p = 0.042$; Online Appendix Table OA-81) and its savings Tk 4,531 higher, 28 percent of its round-1 mean ($p = 0.038$; Table 2). The below-median group of the median split shows the same rise in the index and productive assets; its savings rise is not statistically distinguishable from zero. Its daily-wage income is also higher in the window after, by Tk 3,202 ($p < 0.05$; Online Appendix Table OA-27). I do not attribute the rise to a transfer from the top tercile, since no group's purchases of land or cattle rise significantly after a flood (Online Appendix Tables OA-52 and OA-53). Section 8 shows that the same households farm more land, their own and land rented for cash, in the flood window and move from sharecropping to cash rent. Whether that land or something else raises their capital cannot be determined from the data.

In sum, a flood reduces the capital and earnings of the households that own the capital it hits, the top tercile, and shrinks their farms. The land-poor's wage and self-employment income and consumption fall by at most a few percent. The next section follows the flood's loss into the labor market, through which it reaches the land-poor, and shows how they adapt.

7 The Labor Market

This section estimates whether the labor market passes the flood's loss on through the wage or rations work, and how the households it rations adapt. The farms that hire cut their hiring. The land-poor, whose main endowment is their labor, lose far more of their day labor than any fall in the wage the data allow and move days into work of their own. Following Section 5, I estimate work, hiring and wages in the 2-year window from the interview date, which covers one crop cycle after a flood, beside the interview window. Each estimate is the effect of one additional severe flood in the window.

Table 5 reports the effect of a severe flood on the hours of hired labor at harvest, recorded by module H5 for each plot and field task. The demand for labor falls at harvest, and in hours it falls most on the farms that hire the most. Over all households, hired harvest hours fall by 8.0 per flood in the 2-year window ($p = 0.001$), a quarter of the round-1 mean of 30.9 hours a year. They fall by 4.6 in the interview window ($p = 0.031$; Online Appendix Table OA-39). On the top tercile's plots, hired harvest hours fall by 17.3 per flood in the 2-year window ($p = 0.015$), nearly a quarter of the 73.5 hours they hired in round 1. The interview-window estimate, −6.1 (standard error 4.8), is imprecise, as a window of three to four years dilutes a flow. The top tercile is also the group whose capital falls and whose farm shrinks (Section 6). Under the seven checks of Section 10, the all-household estimate and the top tercile's keep a p-value below 0.10 in every check (Online Appendix Tables OA-87 and OA-88).

Some top-tercile farms stop growing a crop altogether. The share of top-tercile households that grows any crop is 5.7 percentage points lower in the window after a flood ($p = 0.005$; Online Appendix Table OA-36) on a round-1 base of 80 percent. It is 7.7 points lower in the preceding 2-year window ($p = 0.020$). In the current 2-year window and the interview window the estimates are not statistically distinguishable from zero (−4.1 and −3.9 points, standard errors 2.8 and 2.4). The exit therefore follows the flood with a lag of a season or more. The land-poor's and the middle tercile's shares show no significant change in any window, and the smallest tercile's is 5.2 points lower in the preceding 2-year window ($p = 0.060$). Among the 345 top-tercile farms that grow a crop in every round, the asset index falls by 0.36 standard deviations and livestock value by Tk 6,775 in the flood window, each significant at the 5 percent level (Online Appendix Table OA-22, Panel A). Their hired harvest hours in the flood window are not statistically distinguishable from zero (−2.2, standard error 8.2; Online Appendix Table OA-36, Panel B). The farms that continue therefore lose capital without a detectable cut in their hiring.

The land-poor's own plots also use less labor at harvest. In the 2-year window their hired

Table 5: The flow outcomes after a severe flood: the interview window and the 2-year window

	Interview window (1)	2-year window (2)	2-year window after (3)	Base mean (4)
A. Hiring and harvest hours				
Hired harvest hours				
All households	−4.56** (2.11)	−7.95*** (2.30)	−2.85 (3.14)	30.9
No inherited land	−5.17** (2.38)	−7.83*** (2.72)	−3.26 (2.97)	21.3
Top tercile	−6.11 (4.83)	−17.30** (7.08)	−3.23 (8.64)	73.5
Family harvest hours				
All households	−3.02** (1.37)	−4.87*** (1.76)	−0.19 (1.64)	16.1
Top tercile	1.26 (2.92)	−1.36 (3.31)	4.59 (3.97)	20.4
Hired land-preparation and planting hours				
All households	−1.22 (1.43)	−0.64 (1.77)	−2.63 (2.10)	28.8
Top tercile	2.50 (4.93)	0.23 (6.59)	−2.83 (8.06)	69.6
B. No inherited land: work				
Participation (pp)				
Agricultural day labor	−3.51*** (1.26)	−5.32*** (1.50)	−3.69** (1.65)	21.5
Non-farm self-employment	2.36* (1.28)	1.83 (1.76)	2.59* (1.54)	39.0
Days in the last 7 days				
Agricultural day labor	−0.182** (0.076)	−0.248*** (0.095)	−0.232** (0.107)	1.07
Daily-wage work	−0.277*** (0.100)	−0.314** (0.136)	−0.274** (0.123)	1.96
Salaried work	−0.156*** (0.055)	−0.127* (0.077)	−0.134 (0.087)	0.61
Self-employment	0.312 (0.194)	0.475* (0.269)	−0.440 (0.287)	11.4
Daily-wage work + self-employment	0.034 (0.202)	0.161 (0.288)	−0.713** (0.297)	13.4

Notes: The table reports estimates of equation (1) for hiring and work in the balanced panel of 5,473 households, with one observation per household-round. Hours are hours of hired or family labor at harvest, or of hired labor at land preparation and planting, on the household's plots over the agricultural year, and days are days worked in the last seven days, summed over household members. The treatment is the number of severe floods, years in which at least a quarter of the Union is under water, in the window between two interviews, and each row gives the effect of one more flood for all households or for the inherited-land group named. Columns (2) and (3) replace that window by the 24 months before the interview and the 24 months before that, entered together, and column (4) is the group's round-1 mean. All regressions include household fixed effects, district-by-round and group-by-round fixed effects, and the drought, excess-rain and heat counts of both windows. Standard errors clustered at the Union level are in parentheses (318 Unions). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

harvest hours fall by 7.8 per flood ($p = 0.004$) on a round-1 mean of 21.3 a year. Their family harvest hours fall by 5.7 (Online Appendix Table OA-39), about half of the round-1 mean of 12.0. Part of the lost demand for day labor may come from small farms that contract. Because there are about seven land-poor households for each top-tercile household (3,845 against 539), most of the hired harvest hours lost are on the land-poor's plots. For crop care the estimate is imprecise, and in the flood window its interval rules out a fall larger than 10.6 percent of the round-1 mean. For land preparation and planting their hired hours fall by 2.3 in the flood window, significant at the 10 percent level ($p = 0.088$; Online Appendix Table OA-39), with a wild bootstrap p-value of 0.118. Hiring is recorded for the year and the land-poor's days for the week before the interview. The survey does not record who hires the land-poor outside the sample's plots. The link between the farms' hiring and the land-poor's days is therefore a match of the two sides of the market, not of worker to employer.

On the supply side of the market, the land-poor lose about a quarter of their day labor per flood in the 2-year window. The share of land-poor households with any agricultural day labor falls by 5.3 percentage points ($p < 0.001$) on a base of 21.5 percent. Their days of such work fall by 0.25 a week ($p = 0.010$), 23 percent of the round-1 mean of 1.07 days. Their days of daily-wage work, on farms and off them, fall by 0.31 a week ($p = 0.021$). Their daily-wage income falls by Tk 3,727 a year ($p = 0.002$), 22 percent of the round-1 mean of Tk 16,944. All are significant at the 5 percent level or better (Table 5 and Online Appendix Table OA-37). The interview window shows the same fall spread over more years: the days fall by 0.18 a week ($p = 0.017$; Online Appendix Table OA-48), 17 percent of the mean. The share with any such work falls by 3.5 points ($p = 0.006$; Table 10). The loss is the land-poor's because few households in the middle and top terciles sell agricultural day labor (11.1 and 4.8 percent in round 1; Online Appendix Table OA-6). The fall in the share differs across exposure windows. It is -7.3 points in the window of round 1, against -4.0 and -2.1 in those of rounds 2 and 3. Equality of the three effects is rejected ($p = 0.01$; Online Appendix Table OA-79).

Against this fall in days, Table 6 bounds the fall in the farm wage at a few percent. The wage is the daily cash wage for crop day labor, the same work whose days are counted above. Wages enter in logs; I report wage effects and bounds as percentage changes, $100(e^b - 1)$, while the tables print log points. Over the flood window, on the household-mean measure the daily farm wage changes by -0.2 percent for all workers. The 90 percent interval rules out a fall larger than 3.3 percent and a rise larger than 3.0 (Panel A). For workers in land-poor households it rules out a fall larger than 2.6 percent. The wage farmers report paying and the Union median hired wage rule out falls larger than 2.0 and 3.6 percent (Panel B). The round fixed effects absorb the national deflator, so each coefficient is both the nominal and

the real effect. No interval lies inside ±2 percent, so these are bounds rather than precisely estimated zeros. Even so, the land-poor's days fall by 17 percent while the wage falls by at most 3.3 percent. Section 10 shows that the window bound for all workers widens to at most 3.8 percent under its checks. Balboni et al. (2026) find that in the more unequal villages of an anti-poverty experiment in Bangladesh, weather shocks lower the casual wage of the ultra-poor while landlords' earnings and hiring rise. Here, for large floods in a broad rural panel, the wage falls by at most 3.3 percent in the flood window and the hiring farms cut hiring. Their sample of the ultra-poor in the poorest districts, their shock of unpredicted days of flood and drought, and their wage measure of casual labor income per hour differ from mine, so only the direction of the pattern is comparable.

Table 6: The farm daily wage after a flood

Measure	Period	Estimate (SE)	90 percent CI	Largest fall ruled out	Clustered p
A. Farm daily wage reported by workers (log points)					
Household mean,	Window	−0.2 (1.9)	[−3.4, +3.0]	3.4	0.914
all workers	Window after	+2.2 (2.2)	[−1.4, +5.8]	1.4	0.306
Household mean,	Window	+0.7 (2.0)	[−2.6, +4.1]	2.6	0.723
no inherited land	Window after	+3.2 (2.6)	[−1.1, +7.5]	1.1	0.220
Worker (member FE),	Window	−1.0 (2.2)	[−4.7, +2.7]	4.7	0.645
all workers	Window after	+2.6 (2.2)	[−1.1, +6.3]	1.1	0.241
Worker (member FE),	Window	−0.2 (2.3)	[−4.0, +3.6]	4.0	0.926
no inherited land	Window after	+4.0 (2.7)	[−0.5, +8.4]	0.5	0.146
B. Wages in the local labor market (log points)					
Union median	Window	−1.4 (1.4)	[−3.6, +0.9]	3.6	0.317
hired wage	Window after	+1.0 (1.5)	[−1.4, +3.4]	1.4	0.486
Wage farmers	Window	−0.2 (1.1)	[−2.0, +1.6]	2.0	0.858
report paying	Window after	+0.6 (1.2)	[−1.3, +2.6]	1.3	0.602

Notes: The table reports estimates of equation (1) for wages, with one observation per household-round for the household mean (N 1,755) and the wage farmers report paying (N 6,620), per worker-round for the worker measure (N 1,681) and per Union-round for the Union median (N 844). Wages are in log points (100 × log) of round-1 taka, the farm wage is the daily cash wage for crop day labor, and the Union median hired wage is 8 times the cost per hired hour reported by hiring farmers. The treatment is the number of severe floods, years in which at least a quarter of the Union is under water, in the window between two interviews; each row gives the effect of one more flood in the flood window or in the window after, for all workers or as the no-inherited-land group's own effect. The largest fall ruled out is minus the lower end of the 90 percent CI. All regressions include district-by-round fixed effects and the drought, excess-rain and heat counts of both windows, with household fixed effects (member fixed effects for the worker measure, Union fixed effects for the Union median) and group-by-round fixed effects where groups enter. Standard errors clustered at the Union level are in parentheses (239 Unions for the household mean, 233 for the worker measure, 292 for the Union median and 312 for the farmers' wage). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The 2-year window shows the same pair with a wider wage bound. Per flood in that

window the household-mean farm wage changes by −1.7 percent for all workers and −2.1 for the land-poor. The 90 percent intervals rule out falls larger than 5.6 and 6.4 percent (Online Appendix Table OA-38). The Union median hired wage and the wage farmers report paying rule out falls larger than 3.7 and 2.1 percent. The bounds are wider than in the flood window because fewer wage observations identify a 2-year count. But the fall in the land-poor's days per flood in that window, 23 percent, remains far larger than the largest fall in their wage that any interval allows. In both windows the fall in their days is far larger than any fall in the wage the intervals allow.

One concern is that the wage is observed only for those who work, in about 1,750 of the 16,400 household-rounds. A large fall in days changes who is observed. If the workers who lose their days are those who would have been paid least, the observed wage holds up even if the wage of a given worker falls. The wage regression with member fixed effects addresses this in part. It follows the same workers across rounds and so removes permanent differences among them. It rules out a fall larger than 4.6 percent in the window for all workers and 3.9 for the land-poor (Table 6, Panel A), about a percentage point wider than the household-mean bounds. It rules out falls larger than 8.0 and 8.7 percent in the 2-year window (Online Appendix Table OA-38). The test does not address selection on the round's wage draw. A worker's wage is still observed only in the rounds in which the worker works. If floods push lower-paid workers out of work, the true bounds could be wider. Within that limit, the bounds remain small relative to the fall in days when the same workers are followed.

The data also bound the fall in the wage in the years after, when a lagged fall could transmit the loss. In the window after, the farm wage is not statistically distinguishable from zero on any of the six measures. The four worker measures of Panel A lie between +2.2 and +4.0 percent, and every 90 percent interval rules out a fall larger than 1.4 percent (Table 6).

The land-poor's own work rises, although the evidence is weaker than for their day labor. Per flood in the 2-year window they work 0.48 more days a week on their own farms and enterprises on a round-1 mean of 11.4 days ($p = 0.079$; Table 5), with a wild bootstrap p-value of 0.107 (Supplementary Table S-8). In the flood window the share of land-poor households with a member in non-farm self-employment rises by 2.4 percentage points on a base of 39 percent ($p = 0.066$), with a wild bootstrap p-value of 0.097 (Section 10). It is not significant with a group-specific trend in later flood exposure ($p = 0.133$; Section 10). In the 2-year window the take-up is +1.8 points with a standard error of 1.8. There is no evidence of a move into new enterprises. The share of land-poor households that started a non-farm enterprise in the flood window changes by −1.0 points, with a standard error of 0.7. The smallest tercile moves into self-employment by more, 6.4 points in the flood

window, significant at the 1 percent level ($p = 0.008$). The pattern is consistent with Breza et al. (2021), who find that workers rationed out of wage work in Indian villages turn to self-employment in the lean season. It does not establish that the added self-employment absorbs rationed days. The survey cannot tell whether the land-poor would have sold those days at the going wage.

I next ask whether the days of self-employment replace the days of wage work lost (Online Appendix Table OA-40). The estimates are too imprecise to tell. In the flood window the land-poor's days of daily-wage work fall by 0.28 a week ($p = 0.006$), and their salaried days by 0.16, also significant ($p = 0.005$). Their self-employment days rise by 0.31, an estimate that is not statistically distinguishable from zero ($p = 0.109$). Their days of daily-wage work and self-employment together are a precisely estimated zero in the window: $+0.03$ days, with a 90 percent interval from -0.30 to $+0.37$. That lies inside ± 0.40 days, 3 percent of the round-1 mean of 13.4 days a week. The band, however, is wider than the 0.28 days lost, so the interval cannot separate full replacement of the lost days from none. In the 2-year window, the 0.31 days of daily-wage work lost and the 0.48 days of self-employment gained per flood sum to $+0.16$, with an interval from -0.31 to $+0.64$. The data show that the land-poor's self-employment rises as their day labor falls, not that the days they lose are made up. Section 6 bounds what the move does to their earnings.

The non-farm wage adds no evidence on repricing in either direction, because its two measures point in opposite directions. On the household-mean measure, with other wage labor counted as non-farm, it falls by 11.8 log points per flood in the 2-year window ($p = 0.009$). Its difference from the farm wage, -10.1 log points, is significant (Online Appendix Table OA-41, Panel D). On the member measure it rises in the window after, by 10.2 percent for all workers and 11.3 for workers in land-poor households, both significant. No other cell of Panels D and E is significant. I therefore do not interpret either measure.

The loss of day labor is concentrated among the land-poor who farm no land (Online Appendix Tables OA-13 and OA-14). Among the 2,187 that operated none in 2011, days of agricultural day labor fall by 0.27 a week in the flood window ($p = 0.011$), 27 percent of their round-1 mean. Their daily-wage income falls significantly. Among the 1,658 that operated land the estimate is imprecise, and the interval rules out a fall larger than 21 percent of their mean. The difference between the two is not statistically significant ($p = 0.166$). Operation is measured in 2011, inside round 1's exposure window, so I treat the split as descriptive.

A fall in days with a bounded wage is the pattern the wage norm of Section 3 predicts. A fall in the demand for labor leaves days unsold while the wage holds, and the rationed households move to work of their own (Proposition 1(i) and (ii)). The farm wage is not statistically distinguishable from zero in the flood window, in the 2-year window or in the

window after, so the data are consistent with any tolerance large enough that the wage holds at the shocks observed. A market that clears with an elastic supply of labor fits the same pattern in direction but not in size. In the calibrated model of Section 11, a market-clearing wage lowers the land-poor's days of day labor by at most 4.1 percent at any elasticity of labor supply, while the 90 percent interval of the estimate rules out a fall smaller than 5.3 percent. With the wage held, the model gives a fall of 7.0 percent. The framework's prediction that a rationed household that farms puts more time on its own farm is not supported, since the land-poor's family harvest hours fall.

Taken together, the labor market rations work rather than repricing it, and the households it rations adapt. The farms that hire cut their hiring, the land-poor lose day labor, and their own work rises, less precisely. The data bound the fall in the wage at a few percent.

8 The Land Market

This section estimates how the land market adjusts to a flood, the paper's second question for that market. The land market adjusts through its price. The rent falls by about a fifth, while the land-poor do not rent in more land. The top tercile, whose farm shrinks, gives up land (Section 6), and the smallest tercile farms more in the flood window, on its own land and on land rented for cash. But no group rents in more land in total at the lower rent, and the land-poor's rented-in land does not rise in either window. The survey does not record the length or renewal date of a lease, so the timing of the land responses is inferred from the windows in which they appear rather than from the contracts.

Table 7 reports the effects of a severe flood on the Union-level prices. I report rent effects as percentage changes, $100(e^b - 1)$, while the tables print log points. Over the 176 Union-rounds with at least three cash tenants, the median annual cash rent per decimal falls by 21.1 percent in the flood window ($p = 0.006$; column 2). One concern is that only 25 Unions identify this estimate, those whose flood count in the window changes across their rounds. The effective number of clusters is 10 (Online Appendix Table OA-43), reported as information. Over the seven checks of Section 10, however, the estimate never changes sign. The fall holds under the fixed-length windows, counted from the median interview date of the Union's households. It is −21.5 percent in the 3-year window (standard error 8.9 log points) and −38 percent in the 2-year window (standard error 17.4 log points), both significant (Supplementary Table S-9). Of the other Union medians, the rice price is a precisely estimated zero, +0.8 log points with a 90 percent interval inside ±3 (column 1). The hired wage, at −1.4 percent, and the rice yield, at −5.5 log points, are not statistically distinguishable from zero (columns 3 and 4). Of the local prices measured, only the rent

falls significantly.

Table 7: Local farm markets after floods: Union-level rice price, cash rent, hired wage and yield

	(1) Rice price	(2) Cash rent per dec.	(3) Hired wage	(4) Rice yield
Severe floods, current window	+0.008	−0.237***	−0.014	−0.055
	(0.011)	(0.083)	(0.014)	(0.046)
Identifying Unions; G^*	48; 36	25; 10	77; 65	81; 70
Severe floods, window after	−0.010	+0.074	+0.010	+0.048
	(0.013)	(0.153)	(0.015)	(0.033)
Identifying Unions; G^*	73; 42	34; 13	128; 85	135; 89
Unions	187	71	292	309
Union-rounds	505	176	844	916

Notes: The table reports Union-level versions of equation (1) for local farm markets in the three survey rounds, with one observation per Union-round. Each outcome is the log of the Union-round median of a household measure: rice sale price per kg among sellers (column 1), annual cash rent per decimal among cash tenants (column 2), daily hired wage among users of hired labor (column 3) and rice yield per decimal among growers (column 4), with at least three reporting households per Union-round in columns 1 to 3 (three cash tenants in column 2). The treatment is the number of severe floods, years in which at least a quarter of the Union is under water, in the window between two interviews, and each row gives the effect of one more flood in the current window or in the window after. Identifying Unions are those whose flood count in the named window varies over their rounds, and G^* is the effective number of clusters of Carter et al. (2017) at $\rho = 0$. All regressions include Union fixed effects, district-by-round fixed effects, and the drought, excess-rain and heat counts of both windows. Standard errors clustered at the Union level are in parentheses (the Unions row gives the number of clusters). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The rent on a given plot moves in the same direction but is imprecisely estimated. Plot identifiers persist across rounds. Table 8 therefore follows the log real cash rent per decimal, with plot fixed effects, on the 483 cultivable plots observed as cash tenancies in at least two rounds. In the flood window the same-plot rent is not statistically distinguishable from zero: −7.9 percent, with a standard error of 6.4 log points ($p = 0.198$; column 1 and Online Appendix Table OA-70). Nor is a household's cash rent per decimal over all its cash-rented plots (−5.2; column 2). In the 2-year window, however, the same-plot rent falls by 25 percent per flood (standard error 13 log points; Supplementary Table S-9). The plot-level rent thus moves with the Union median once the window is short enough to isolate recent floods.

The rent that sample landlords receive supports the Union median. On their rented-out plots, followed with plot fixed effects, it falls by 33.0 percent in the flood window ($p < 0.001$; 160 identifying plots; column 3). The estimate rests on few Unions, with an effective number of clusters of 7 (Table 10), reported as information. Its window-after coefficient is +29.4

percent ($p = 0.065$). The landlords' and the tenants' plots are different plots, and most landlords are outside the sample. I therefore treat the landlords' rent as supporting evidence and the Union median as the main estimate.

Table 8: Within-plot and within-household cash rent after floods

	(1)	(2)	(3)
	Rent paid		Rent received
	Plot FE	Household FE	Plot FE
Severe floods, current window	−0.082	−0.053	−0.400***
	(0.064)	(0.067)	(0.084)
Severe floods, window after	−0.145	−0.074	0.258*
	(0.109)	(0.080)	(0.138)
N	1,050	1,068	1,125
R^2	0.281	0.247	0.342

Notes: The table reports plot-level and household-level versions of equation (1) for the cash rent in the three survey rounds, with one observation per plot-round (columns 1 and 3) or household-round (column 2). The outcome is the log real annual cash rent per decimal, winsorized at the 1st and 99th percentiles of each column's sample. The treatment is the number of severe floods, years in which at least a quarter of the Union is under water, in the window between two interviews, and each row gives the effect of one more flood in the current window or in the window after. Column 1 uses the rent paid on cash-rented-in plots observed in at least two rounds, column 2 the household's rent over all its cash-rented-in plots for households with a cash tenancy in at least two rounds, and column 3 the rent received on cash-rented-out plots observed in at least two rounds. All regressions include plot fixed effects (household fixed effects in column 2), round fixed effects, district-by-round effects (entered as district linear and quadratic trends), and the drought, excess-rain and heat counts of both windows. Standard errors clustered at the Union level are in parentheses (138, 183 and 108 Unions), and R^2 is the within R^2. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The fall in the median shows no sign of arising from a change in which plots are leased. The concern is substantive, because the round-1 cash tenancies that are not continued rent for more, Tk 125 against Tk 116 a decimal (Supplementary Table S-27). I test it in three ways. First, a severe flood changes the probability that a plot continues as a cash tenancy by +0.004 (standard error 0.024) on a base of 0.31. Second, samples that do not condition on continuation give falls of 10.3 percent within household ($p = 0.081$) and 13.8 within Union and round ($p = 0.016$; Online Appendix OA-5.1). Third, flooding is not detectably associated with which Unions enter the rent sample, although that interval is wide (Section 10). The first two tests show no change in which plots are leased, and the third cannot rule out a change in which Unions enter the sample. A lease that covers fewer seasons would also lower the rent per decimal. But the number of seasons planted on the same cash-rented plots is not statistically distinguishable from zero in the flood window (Supplementary Table S-28).

The land-poor do not rent in more land at the lower rent. In the flood window their rented-in land changes by −2.5 decimals on a round-1 mean of 21.1, a fall that is not

statistically distinguishable from zero (standard error 1.5, $p = 0.104$; Online Appendix Table OA-48). In the window after it is not statistically distinguishable from zero either (Online Appendix Table OA-48). Their cash-rented land changes by -2.0 decimals in the flood window ($p = 0.067$, standard error 1.1), which does not survive the wild bootstrap ($p = 0.135$), and their sharecropped and mortgaged land by less. The fixed-length windows give similar estimates, -2.8 decimals in the 3-year window ($p = 0.080$) and -2.7 in the 2-year window ($p = 0.187$; Supplementary Table S-6). Hence the data show no rise in the land the land-poor rent. The fall is imprecise, significant at the 10 percent level only in the 3-year window. Their production loans and their debt are not statistically distinguishable from zero in either window (the share with a production loan -1.0 points in the flood window, standard error 1.4, on a base of 31 percent; Section 6). The share refused credit for lack of collateral rises significantly by 1.4 points in the flood window, on a base of 1.3 percent in the rounds that record it (Online Appendix Table OA-48). Their savings show no fall (Section 6).

The fall in rented-in land is concentrated among the land-poor households that operated land in 2011 (Online Appendix Table OA-13). Among them rented-in land falls by 5.6 decimals in the flood window on a base of 49.0, a significant fall. Among those that operated none the estimate is -0.2 (standard error 1.8). Operation is measured in 2011, inside round 1's exposure window, so this split describes where the fall lies and is not a causal moderator.

The land-poor therefore do not take up the land at the lower rent. The land-poor hold 67 percent of the sample's rented-in area (Table 9, Panel B). But in round 3, 74 percent of the rented-in area is owned by a household in the same village, while the panel follows about 17 households per Union. Most owners are therefore outside the sample, and the panel cannot show which side of the market declines to lease. The rent falls while no group's rented-in land rises. A fall in the rent with no rise in any group's rented-in land is what a fall in the demand for tenanted land produces under Proposition 2(iii), in which the price and the quantity move together. The test of the proposition applied to the land-poor's rented-in land is reported at the end of this section.

The smallest tercile farms more land in the flood window, its own land and land rented for cash, and it moves from sharecropping to cash rent. In the flood window it operates 6.8 more decimals ($p = 0.044$) on a round-1 mean of 55.8. It farms 4.0 more decimals of its own land ($p = 0.005$) and rents 4.1 more for cash ($p = 0.044$; Online Appendix Table OA-47). Its sharecropped land is 4.4 decimals lower in the window after, a significant fall. The larger farm does not last. In the window after its operated land is not statistically distinguishable from zero (-1.1 decimals). Nor is its rented-in land of every kind in either window. Under

the fixed-length windows its operated land is +6.4 decimals in the 3-year window ($p = 0.102$) and +9.4 in the 2-year window ($p = 0.061$; Supplementary Table S-6). Its capital rises in the window after (Section 6).

Under the seven checks of Section 10, its own-operated and cash-rented land in the flood window keep their sign and their significance at the 10 percent level in every check (Online Appendix Table OA-50). With the median split of inheritors, which adds half of the middle tercile to the group below the median, the gains are smaller: operated land rises by 5.1 decimals, significant at the 10 percent level ($p = 0.087$) with a wild bootstrap p-value of 0.108, and cash-rented land is not statistically distinguishable from zero. The fall in sharecropping remains significant. Two limits prevent reading this as a transfer from the top tercile. The data cannot link plots across households, and the timing differs. The smallest tercile farms more in the flood window, while the top tercile rents out more in the window after (Section 6). Nor do the smallest tercile's round-1 savings, labor or productivity indicate which of its households farm more (Supplementary Table S-30). Lane (2024) finds that Bangladeshi farmers offered credit guaranteed to be available after a flood expand the land they rent in for cash. That is the margin the smallest tercile uses here.

Floods do not concentrate operated land within a Union. The Union Gini coefficient of operated land changes by −0.008 in the flood window on a base of 0.67. Its 90 percent interval lies inside 3 percent of the base, a precisely estimated zero (Online Appendix Table OA-46). The concentration of land after flood and famine that Cain (1981) describes for a Bangladeshi village does not appear. Sales of owned land, which are small and come from the top tercile, are reported in Section 9.

The framework of Section 3 predicts that the land the top tercile gives up can lower the rent only a little, so a large fall requires an inward shift in the demand for land (Proposition 2). The estimates are consistent with this prediction. The top tercile's rented-in land changes by −3.2 decimals in the flood window and −4.0 in the window after on a round-1 mean of 17.8, neither statistically distinguishable from zero (Online Appendix Table OA-48). It rents 8.0 percent of the rented area, so even at the point estimates its retreat from rented land lowers the rent by at most 1.1 and 1.3 log points with inelastic supply (part (iii); Table 9, Panel A). The top tercile also rents out more of its own land, by 6.5 decimals in the window after (Online Appendix Table OA-47). Counting that land as well, its net rented-in land changes by −3.8 decimals in the flood window and −10.5 in the window after, and part (iv) bounds the fall in the rent at 1.2 and 3.4 log points, against a fall of 23.7 log points (21.1 percent) in the Union median rent. Updated beliefs explain little of the fall. If operators updated their beliefs after a flood as farmers in Bangladesh report, the rent on all leases would fall by a further 1.0 log points, less than a twentieth of the fall (Online

Appendix OA-7.8).

The land-poor's response reveals an inward shift in their own demand. If a group's demand had not shifted, it would rent in at least $a = 1.48$ percent more land for each percent fall in the rent (part (ii)). Measured against the flood-window rent, the land-poor's response is -0.50, with a 90 percent interval from -2.26 to $+0.69$, short of a by 1.98 (interval 0.79 to 3.73). Their demand for rented land at the pre-flood rent therefore shifted inward, by at least 0.47 log points at the point estimate, with an interval from 0.00 to 0.74 (Table 9, Panel B). The test matched to the window after is uninformative, because the rent's window-after estimate is imprecise. Their collateral rationing rises (Online Appendix Table OA-48), consistent with a fall in the capital they farm with, but the prices cannot separate this from a lower value of the season or from owners leasing less to them. Most owners are outside the panel, so the data also cannot separate an inward shift in demand from an outward shift in supply.

Taken together, the land market transmits the flood through its price, and the smallest landowners farm more for a time. But the lower rent does not move land to the land-poor, who do not rent in more.

9 The Land-Poor in the Years after a Flood

This section follows the land-poor beyond the 2-year window, after they have turned to work of their own. In the window after a flood they own about 13 percent less land, because they buy and inherit less, not because they sell. Fewer of their members migrate. The land sales that rise are the top tercile's, in the flood window, when its capital falls.

Online Appendix Table OA-51 reports the effects of a severe flood on the land-poor's owned land and on the routes by which it changes. The land-poor own less land in the window after a flood because they acquire less. They hold 2.14 fewer decimals of owned land in that window ($p = 0.013$), 13 percent of their round-1 holding of 16.6 decimals (Panel A). In the flood window itself the estimate is imprecise (-0.13 decimals, standard error 0.94; Panel A). The plot roster, which follows each plot's status across rounds, decomposes the fall by route. They buy 1.33 fewer decimals ($p = 0.001$) and inherit 0.75 fewer ($p = 0.040$). The area they sell is a precisely estimated zero relative to their owned land in both windows. In the window after, the 90 percent interval runs from -0.29 to $+0.34$ decimals, inside 3 percent of their holding. The areas lost to erosion and to mortgages in the window after are precisely estimated zeros as well. The value of their land shows no sign of distress (Panel C; Online Appendix Table OA-52, Panel B). The roster records the route, not the reason, so the data cannot say why less land is inherited.

Table 9: The rental-market elasticities a severe flood reveals

	Estimate	90 percent interval
A. The top tercile's withdrawal		
Top tercile's share of the panel tenants' rented-in land $\mathcal{Q}$, $q_{\mathcal{Z}}$	0.080	[0.062, 0.094]
$X = q_{\mathcal{Z}}\,\mathrm{d}\ln D_{\mathcal{Z}}$, window after	−0.018	[−0.035, 0.002]
with the Union share	−0.014	[−0.026, 0.002]
change in the rent with inelastic supply, first order, $X/[a(1-q_{\mathcal{Z}})]$	−0.013	[−0.026, 0.001]
$X = q_{\mathcal{Z}}\,\mathrm{d}\ln D_{\mathcal{Z}}$, flood window	−0.014	[−0.034, 0.010]
with the Union share	−0.011	[−0.026, 0.008]
change in the rent with inelastic supply, first order, $X/[a(1-q_{\mathcal{Z}})]$	−0.011	[−0.025, 0.008]
Bound if $\mathcal{Z}$ gives up all its rented-in land, $\ln(1-q_{\mathcal{Z}})/a$	−0.056	[−0.067, −0.043]
B. The test for the land-poor (no inherited land)		
Land-poor's share of $\mathcal{Q}$	0.673	[0.642, 0.703]
$a = (1-\beta)/(1-\alpha-\beta)$	1.48	–
Measured response $\varepsilon^m = -\mathrm{d}\ln D/\mathrm{d}\ln r$: window-after rent (matched)	0.72	[−2.21, 3.28]
shortfall $a - \varepsilon^m$	0.76	[−1.80, 3.68]
shift bound $\mathrm{d}\ln D + a\,\mathrm{d}\ln r$	0.056	[−0.513, 0.666]
Measured response $\varepsilon^m = -\mathrm{d}\ln D/\mathrm{d}\ln r$: flood-window rent (comparison)	−0.50	[−2.26, 0.69]
shortfall $a - \varepsilon^m$	1.98	[0.79, 3.73]
shift bound $\mathrm{d}\ln D + a\,\mathrm{d}\ln r$	−0.468	[−0.743, −0.001]

Notes: The table reports the tests of Proposition 2 for the rental market: Panel A bounds the change in the rent that the top tercile's withdrawal from renting in can cause, and Panel B sets the land-poor's measured response of rented-in land to the rent against its lower bound a. Households are grouped by inherited land into those with none (the land-poor) and terciles among the rest, $\mathcal{Z}$ is the top tercile, and $\mathcal{Q}$ is the panel tenants' rented-in land. The changes in rented-in land and in the rent are per severe flood, a year in which at least a quarter of the Union is under water, and come from the regressions in Online Appendix Table OA-49. The constant a uses $\alpha = 0.261$ and $\beta = 0.191$ from the calibration in Online Appendix Table OA-63 and has no interval (–). Intervals are the 5th and 95th percentiles of a Union-cluster pairs bootstrap with 200 draws that re-estimates every input on the same draw.

In the years after a flood, the land-poor's measurable loss is in owned land. Their accumulation slows, while their consumption and earnings fall by at most a few percent. Their productive assets show no significant change in either window (Online Appendix Table OA-81).

The land sales that rise after a flood are the top tercile's, and the rise does not persist. In the flood window the top tercile has sold 3.5 more decimals since round 1 ($p = 0.008$; Online Appendix Table OA-52, Panel A), small relative to its holding of 195 decimals. In the same panel, 3.4 percentage points more of its households sell a plot ($p = 0.050$; Panel A). In the window after, the area it has sold since round 1 is 1.8 decimals lower ($p = 0.064$). Its owned land in that window is a precisely estimated zero (Online Appendix Table OA-82). It also sells 0.045 more goats and sheep in the flood window ($p = 0.056$; Online Appendix Table OA-53), on a round-1 mean of 0.09 head. No other group's sales of land or animals rise significantly, and among purchases only the smallest tercile's of goats and sheep do (0.026 head in the flood window, $p = 0.096$), so the panel shows almost no one taking up what the top tercile sells. This is the group whose livestock and working capital fall, and its sales are consistent with productive assets held as a buffer stock, as in Rosenzweig and Wolpin (1993). The land-poor, who hold little such capital, instead buy less land.

In the window after a flood, fewer members of land-poor households migrate. Module V1 lists the members away for six months or more. I date episodes by calendar year and assign them to the exposure windows, so the module's recall period does not by itself produce a lagged response. Seasonal migration, the margin Bryan et al. (2014) study, is not measured. In the window after a flood, the probability that a member leaves falls by 2.4 percentage points among the land-poor ($p = 0.063$) on a base of 11 percent (Online Appendix Table OA-54, column 1). It falls by 8.1 points among the smallest tercile ($p = 0.001$) on a base of 13 percent, in the same column. Domestic moves fall too, by 2.3 and 5.8 points, also significant (column 2). For the top tercile, and in the flood window for either group, the estimates are imprecise.

One concern is that V1 omits members who have returned, so fewer listed moves could reflect more returns rather than fewer departures. The household roster, which links members across rounds, separates the two (Online Appendix Tables OA-55 and OA-56). In the window after, the probability that a man aged 15 to 49 present at the previous round has left falls by 7.2 points among the land-poor (standard error 2.4). The round-2 mean is 12 percent. The number of such men present rises (Online Appendix Table OA-57), while a rise in men joining or returning is not detectable. The roster therefore indicates fewer departures rather than more returns. A second concern is that whole households leaving the panel produce the fall. But the land-poor's exit rises by only 1.2 points in the flood window, significant at

the 10 percent level on clustered standard errors ($p = 0.099$) but not on the wild bootstrap ($p = 0.143$), and the interval rules out a rise larger than 2.4 points (Section 10). No other group's exit rises significantly.

Remittances fall with the migrants: in the window after, the land-poor receive Tk 3,001 a year less ($p = 0.011$; Online Appendix Table OA-27). The fall in departures differs from the evidence that floods move people out of places (Patel 2024; Hornbeck and Naidu 2014). In rural Bangladesh Gray and Mueller (2012) find modest effects of flooding on long-term mobility. They note that a disaster can lower mobility by removing the resources needed to move. I cannot tell why fewer members leave. Three explanations remain open, a shortage of cash for the fare, a return home to rebuild and the risk of a failed move (Bryan et al. 2014). The tests that could separate them have minimum detectable effects larger than the fall or cannot distinguish the returns they predict from zero (Online Appendix Tables OA-58 and OA-60). I report the fall as an empirical regularity.

Taken together, the land-poor's income and consumption fall by at most a few percent in the short run, while they move into work of their own (Sections 6 and 7). In the window after a flood, however, they acquire less land, and fewer of their members migrate.

10 Robustness

This section examines whether the results could arise from something other than the floods. The subsections take the four concerns of the introduction in order, with the definition of the exposure window after the second, and then attrition and the remaining checks. The capital, earnings and labor results keep their sign under every severity cutoff, trend control and window definition, although the earnings loss loses significance under the trend control. The land-poor's move into non-farm self-employment does not survive the trend control. The evidence on the size of the rent fall and on the fall in the land-poor's migration is weaker. The tables not in the main text are in Online Appendices OA-1, OA-2, OA-3, OA-5 and OA-10.

10.1 Pre-existing Trends

The first concern is that flooded Unions were already on different trajectories before the floods, in which case the estimates would capture those trajectories rather than the effect of the floods.

If flooded Unions were already on different trajectories, outcomes would move with floods that had not yet occurred. I test this in two ways. The first regresses the change between

rounds 1 and 2 on the severe floods of 2015–18, which cannot have caused it, for each outcome and group behind a claim. The second adds to the main specification a group-specific trend in the household's later flood count (Online Appendix Tables OA-75 and OA-76). Four of the nineteen placebo coefficients are significant at the 10 percent level on clustered standard errors. Three have the sign opposite to the estimate they test. In the Unions that would later flood, the top tercile's asset index rises by 0.11 standard deviations ($p = 0.070$) and its rented-in land by 4.9 decimals ($p = 0.049$), and the land-poor's owned land rises by 2.6 decimals ($p = 0.005$). A trend of that sign works against the estimated falls unless it reverses. The fourth has the sign of the estimate it tests: the top tercile's hired harvest hours are 17.8 hours lower ($p = 0.049$). The pooled placebos and floods in the twelve months after an interview are not significantly associated with the outcomes already recorded (Online Appendix Tables OA-74 and OA-77).

The second test adds the group-specific trend in later flood exposure, which absorbs a pre-existing trend of either sign. With it, the top tercile's hired harvest hours in the 2-year window are −19.1 against −17.3, and the land-poor's owned land in the window after is −2.40 decimals against −2.14. With the trend, every significant key estimate keeps its sign, and all but three remain significant. The top tercile's asset index is −0.21 against −0.25, and its working capital −Tk 4,901 against −Tk 6,152. The land-poor's days, participation and daily-wage income and the hired harvest hours of all households and of the top tercile in the 2-year window are unchanged or larger (Online Appendix Table OA-76). The three exceptions are the top tercile's wage and self-employment income in the flood window, which falls to −Tk 5,951 (standard error 4,503), its sharecropped land in the window after, and the land-poor's take-up of self-employment in the flood window ($p = 0.133$). The land-poor's self-employment days in the 2-year window stay significant on clustered standard errors, with a wild bootstrap p-value of 0.117. Trends tied to later exposure therefore do not account for the capital, labor and hiring results. The earnings loss, the sharecropping retreat and the self-employment take-up hold only if no such trend exists.

Flood-prone Unions might converge more slowly for reasons unrelated to any flood, and their exposure rises by more when the national cycle turns wet. The estimates could capture their slower convergence. I add baseline flood proneness interacted with round, linearly and in quartile bins. The controls move the top tercile's asset index by at most 0.04 of a standard error and the land-poor's participation in agricultural day labor by 0.06 (Online Appendix Table OA-78). Differential convergence in flood-prone Unions therefore does not explain the results.

Inherited land is larger in Unions with more floods within a district (Online Appendix OA-2.4). If larger farms followed a different trajectory across rounds for reasons unrelated to

floods, the flood interactions could capture it. I add continuous inherited land interacted with round, and then interacted with round separately within each group. For the Union median rent, the Union's round-1 mean inherited land, or its shares of the inherited-land groups, enters interacted with round (Online Appendix Tables OA-7 to OA-9). Two estimates cross the 10 percent level of the wild bootstrap. The hired harvest hours of all households in the flood window fall by 3.88 and 3.82 hours under the two variants against 4.56, and their wild bootstrap p-value rises from 0.070 to 0.121 and 0.126. The top tercile's sharecropped land in the window after falls by 3.43 decimals against 3.22, and its wild bootstrap p-value falls from 0.110 to 0.099 and 0.097 (Online Appendix Table OA-7). The other household estimates stay within about 3 percent of their main values, except two of the top tercile's estimates that are not significant in any column, its hired harvest hours in the flood window and its wage and self-employment income in the window after. The Union median rent in the flood window is −20.1 and −16.6 percent against −21.1, each significant. The imbalance in inherited land therefore does not account for the capital and earnings losses, though the flood-window fall in all households' hired harvest hours does not survive the wild bootstrap under these trends.

10.2 The Severity Cutoff and Other Controls

The second concern is that the cutoff for a severe flood, or the controls and weights of this section, drive the estimates.

Only 25 Unions identify the Union median cash rent, so its fall of about a fifth could depend on the specification. I re-estimate it under seven variants, which move the cutoff to 20 and 30 percent and add the proneness, attrition-weight and farm-size checks of this section (Online Appendix Table OA-43). Against the main estimate of −21.1 percent ($p = 0.006$), the variants range from −22.1 to −12.3 percent and never change sign. They keep a clustered p-value below 0.10 in six of the seven, the exception being the 20 percent cutoff. The wild bootstrap p-value rises above 0.10 at both alternative cutoffs and stays below it under the other five. The fall in the rent therefore does not depend on these controls, but its size and its wild bootstrap significance depend on the cutoff.

The Union median could also fall because floods change which Unions have enough cash tenants, or which plots are leased, rather than the rent on a given plot. I estimate the effect of a flood on whether a Union-round enters the rent sample. I also follow the rent on the same plots. The probability that a Union-round has the three cash tenants it needs changes by −1.45 points on a base of 30.2 percent, with a 90 percent interval from −8.4 to +5.5 points. On the same plots, the rent tenants pay is not statistically distinguishable from zero

in the flood window, −7.9 percent ($p = 0.198$; Online Appendix Table OA-70). The rent landlords receive falls by 33.0 percent ($p < 0.001$; Tables 8, column 3, and 10). Selection into the rent sample is therefore not ruled out. But the fall in the rent landlords receive on the same plots supports the fall in the Union median.

The labor and earnings estimates could likewise depend on the cutoff or on these controls. I apply the same seven checks to them (Online Appendix Tables OA-87 and OA-88, clustered p-values). The top tercile's wage and self-employment income in the flood window ranges from −Tk 12,587 to −Tk 9,393 against −Tk 10,764 and keeps a p-value below 0.10 in every check (largest $p = 0.015$). The land-poor's participation in agricultural day labor and their daily-wage income in the 2-year window keep a p-value below 0.10 in every check. So does the fall in hired harvest hours for all households, on the land-poor's own plots and for the top tercile. The land-poor's days of agricultural day labor keep a p-value below 0.10 in every check, in the 2-year window (largest $p = 0.010$) and in the flood window (largest $p = 0.059$). The earnings, days and hiring results are therefore robust to the cutoff and to these controls.

The claim that floods ration days rather than lower the wage requires a tight bound on the fall in the farm wage. That bound could widen under another specification. I recompute it under the same seven checks. In the flood window the household-mean wage falls by at most 3.3 percent in the main specification and by at most 3.8 percent under the widest check, the farm-size time paths by group. For the land-poor the bound is 2.6 percent, and 2.8 percent at its widest. Under every check the land-poor's days of agricultural day labor in the flood window fall significantly. The wage bound therefore stays tight under the checks, while the fall in days holds.

The evidence on the land-poor in the years after a flood could also depend on the cutoff or on these controls. I apply the same seven checks to their owned land and their migration in the window after (Online Appendix Tables OA-87 and OA-88). Both keep their sign in all seven. Owned land remains significant at the 10 percent level in all seven. Migration remains significant in six, losing significance at the 20 percent cutoff. The evidence that fewer of their members migrate is therefore the weaker of the two.

The results could depend on where the quarter-of-a-Union cutoff for a severe flood is placed. I re-estimate them with the cutoff at 20 and at 30 percent. At cutoffs of 20, 25 and 30 percent the top tercile's asset index falls by 0.23, 0.25 and 0.26 standard deviations, and the land-poor's participation in agricultural day labor falls by 3.1, 3.5 and 3.2 points. Each is significant at the 5 percent level (Online Appendix Table OA-1). The 2-year-window flows hold at both cutoffs. The land-poor's days, participation and daily-wage income and the hired harvest hours of all households and of the top tercile keep a p-value below 0.10 at 20 and at 30 percent. The land-poor's self-employment days are significant at 20 percent

but not at 30 (Online Appendix Table OA-87). Several estimates lose significance at one cutoff. The Union median rent is −12.3 percent ($p = 0.343$) at 20 percent and fails the wild bootstrap at 30 percent (Online Appendix Table OA-43). At 20 percent the land-poor's migration loses significance, and at 30 percent the top tercile's sharecropped land in the window after. The capital, earnings, day-labor and hiring results are therefore robust to the cutoff. The size of the rent fall, the self-employment days, the fall in migration and the tenancy responses are sensitive to it.

The main specification counts floods covering a tenth to a quarter of the Union as no flood. If such floods move capital, the comparison group could produce the loss. I enter those moderate floods separately (Online Appendix OA-1.1). After moderate floods the top tercile's asset index rises (+0.22, $p < 0.01$), which the framework does not explain. With them entered the loss after severe floods is −0.21 ($p = 0.01$) against −0.25. The comparison group therefore does not produce the capital loss.

A common cutoff treats a typical flood in a flood-prone Union like a rare one elsewhere, which could misstate the results. I measure exposure instead as a departure from what is normal for the Union, an annual analogue of the weekly deviation of Patel (2024) (Online Appendix OA-10.4). The Union median rent still falls, by 19.8 percent ($p = 0.012$). The top tercile's asset index is −0.15, with a 90 percent interval from −0.32 to +0.02 that contains the main estimate. But the land-poor's participation in agricultural day labor is −1.0 points, with an interval from −3.2 to +1.2 that excludes the main estimate of −3.5. I retain the level measure (Section 4), because the deviation measure treats a typical flood in a flood-prone Union as no flood and so discards much of the variation behind the lost days. I interpret those days as a response to floods that are large in absolute terms, which the deviation measure cannot confirm.

10.3 The Exposure Window

The exposure window of the main specification is the interval between two interviews, which is three years for the first two rounds and four for the third. It is attached to the round rather than to the household's interview date, so the estimates could depend on that definition. I re-estimate the main specification with 2-year and 3-year windows counted back from each household's own interview date, with the preceding window of the same length as the lag (Section 5.1; Supplementary Tables S-5 to S-9). The 3-year window reproduces the interview window almost cell for cell. The top tercile's asset index falls by 0.21 standard deviations against 0.25, and its working capital by Tk 5,155 against 6,152. Its wage and self-employment income falls by Tk 10,911 against 10,764, and the land-poor's days of agricultural

day labor by 0.21 a week against 0.18, each significant. The 2-year window separates the immediate responses from the lagged ones. The flows are larger per flood in the 2-year window (Section 7). The capital and earnings losses move to the preceding 2-year window, two to four years before the interview. There the asset index is −0.38 standard deviations, working capital −Tk 8,783 and livestock value −Tk 6,184, each significant at the 1 percent level. In the current 2-year window livestock value falls by less (−Tk 3,743, $p = 0.038$), and the asset index and working capital are not statistically distinguishable from zero. The Union median rent falls by 21.5 percent in the 3-year window and 38 percent in the 2-year window, both significant. The definition of the window therefore does not drive the results. The 2-year window dates the capital and earnings losses to the years after the flood rather than to the season of the flood itself.

10.4 The Estimator and Inference

The third concern is the estimator, and with it the precision of estimates that few Unions identify.

With a treatment that switches on and off, the fixed effects estimator of the main specification gives some comparisons negative weight, which could reverse the sign of a key estimate (Section 5.1). I re-estimate every key estimate with the interaction-weighted estimator, which removes the negative weight across rounds and flood histories. I ask whether it keeps the sign inside the main 95 percent confidence interval (Appendix A, Table 15). Every key estimate does. The top tercile's asset index is −0.27 against −0.25, and its livestock and working-capital losses are as large or larger (Table 15; Supplementary Tables S-7 and S-8 for the earnings and labor estimates). The estimators disagree in sign only on the top tercile's productive assets and rented-in land in the flood window and on consumption in the window after (Online Appendix Table OA-19). None is a result I claim, and none is distinguishable from zero under either estimator. The negative weights across rounds and flood histories therefore do not produce the sign or the size of any key effect. The weight within households flooded in every round remains (Appendix A).

Union-clustered standard errors could overstate the precision of estimates that few Unions identify, and floods are spatially correlated. I re-examine the key estimates in three ways, each compared with the Union-clustered standard error. Table 10 gives the wild cluster bootstrap and the effective number of clusters of Carter et al. (2017). The capital, earnings, hiring, day-labor and rent estimates keep a wild bootstrap p-value below 0.10, and the land-poor's take-up of non-farm self-employment is close to the 10 percent level ($p = 0.066$; wild bootstrap $p = 0.097$). The effective number is 93 for the top tercile's capital and earnings in

the flood window. Only the rent rows rest on few Unions: 13 for the Union median rent (10 in Online Appendix Table OA-43, under a different assumption about the correlation within a Union), 7 for the rent landlords receive and 2 for the rent tenants pay on the same plot (Online Appendix Table OA-70). Table 11 gives Conley spatial standard errors at 15 and 25 kilometers. They are within a few percent of the Union-clustered ones, except for the Union median rent, where they are smaller, and every key estimate keeps its significance. Clustering by Union and by survey round also leaves every key estimate significant at the 10 percent level (Online Appendix Table OA-69). With three rounds the round dimension has only three clusters, and for some rows the two-way standard error is smaller than the Union-clustered one, so I do not read a smaller two-way p-value as evidence. Among the other estimates the text cites, three that are significant on clustered standard errors do not survive the wild bootstrap: the top tercile's savings in the flood window ($p = 0.085$; wild bootstrap $p = 0.113$) and the land-poor's savings in the flood window ($p = 0.091$; wild bootstrap $p = 0.142$; Supplementary Table S-7), and the poverty headcount in the window after ($p = 0.075$; wild bootstrap $p = 0.114$; Online Appendix Table OA-70). Four more have a wild bootstrap p-value between 0.098 and 0.110, close to the 10 percent level. They are the top tercile's production loans in the flood window ($p = 0.078$; wild bootstrap $p = 0.098$; Online Appendix Table OA-70), its food spending in the flood window ($p = 0.086$; wild bootstrap $p = 0.100$; Online Appendix Table OA-70), its sharecropped land in the window after ($p = 0.092$; wild bootstrap $p = 0.110$; Supplementary Table S-6), and the land-poor's days of self-employment in the 2-year window ($p = 0.079$; wild bootstrap $p = 0.107$; Supplementary Table S-8). The text reports these estimates as the baseline gives them, with both p-values.

10.5 Measurement of the 2018 Monsoon

The fourth concern is that the radar record for the 2018 monsoon measures floods with more error than the satellite record of earlier years.

The 2018 monsoon is filled from Sentinel-1, which measures the last window with more error than the first two. A fill that adds floods in error could therefore produce the results (Online Appendix OA-1.2). I re-estimate without the fill and at stricter radar thresholds (Online Appendix Table OA-2). For the 2-year-window flows and the flood-window earnings, I remove the fill from the exposure record and rebuild the counts (Online Appendix Table OA-88, next-to-last column). Without the fill, the top tercile's asset index falls by 0.34 standard deviations instead of 0.25, and stricter thresholds move it toward that (-0.31 to -0.32). The Union median rent falls by 26 percent instead of 21. Without the fill the top tercile's

Table 10: Inference on the main estimates: the wild cluster bootstrap

	Estimate	Union-clustered SE	p, Union-clustered	p, wild bootstrap	G^*
A. Capital and the farm, top tercile					
Asset index (SD), flood window, top tercile	−0.252	(0.082)	0.002	0.003	93
Livestock value (Tk), flood window, top tercile	−5,410	(1,399)	0.000	0.002	93
Working capital (Tk), flood window, top tercile	−6,152	(1,561)	0.000	0.001	93
Working capital (Tk), window after, top tercile	−3,774	(1,800)	0.037	0.042	64
Wage and self-employment income (Tk/yr), flood window, top tercile	−10,764	(3,739)	0.004	0.011	93
Hired harvest hours, 2-year window, top tercile	−17.3	(7.1)	0.015	0.022	71
B. Labor, land-poor					
Any agricultural day labor (pp), flood window, land-poor	−3.51	(1.26)	0.006	0.017	36
Agricultural day-labor days (7 days), 2-year window, land-poor	−0.248	(0.095)	0.010	0.026	41
Any non-farm self-employment (pp), flood window, land-poor	2.36	(1.28)	0.066	0.097	36
C. Rent					
Log Union median cash rent per decimal, flood window	−0.237	(0.083)	0.006	0.037	13
Log cash rent received per decimal, same plot, flood window	−0.400	(0.084)	0.000	0.062	7

Notes: Each row reports one coefficient of equation (1), estimated on the balanced panel of 5,473 households with one observation per household-round, except the two rent rows, which use the Union-round panel of Union median cash rents and the plot-rounds of cash-rented-out plots. The treatment is the number of severe floods, years in which at least a quarter of the Union is under water, in the window between two interviews, and each household row gives one inherited-land group's own effect of one more flood in the window named. The household regressions include household, district-by-round and group-by-round fixed effects and the drought, excess-rain and heat counts of both windows, and the rent regressions include Union or plot fixed effects, district-by-round effects and the same weather counts. The columns give the Union-clustered standard error and its p-value from $t(G-1)$, the p-value of the restricted wild cluster bootstrap-t with Rademacher weights at the Union level (9,999 draws), and the effective number of clusters G^* of Carter et al. (2017). Standard errors are clustered at the Union level (318 Unions in the household rows, 157 and 108 in the rent rows).

Table 11: Inference on the main estimates: Conley spatial standard errors

	Estimate	Union-clustered	Conley 15 km	Conley 25 km
A. Capital and the farm, top tercile				
Asset index (SD), flood window, top tercile	−0.252	(0.082) [0.002]	(0.083) [0.002]	(0.083) [0.003]
Livestock value (Tk), flood window, top tercile	−5,410	(1,399) [0.000]	(1,392) [0.000]	(1,387) [0.000]
Working capital (Tk), flood window, top tercile	−6,152	(1,561) [0.000]	(1,558) [0.000]	(1,566) [0.000]
Working capital (Tk), window after, top tercile	−3,774	(1,800) [0.037]	(1,792) [0.035]	(1,782) [0.034]
Wage and self-employment income (Tk/yr), flood window, top tercile	−10,764	(3,739) [0.004]	(3,746) [0.004]	(3,771) [0.004]
Hired harvest hours, 2-year window, top tercile	−17.3	(7.1) [0.015]	(7.1) [0.015]	(7.3) [0.017]
B. Labor, land-poor				
Any agricultural day labor (pp), flood window, land-poor	−3.51	(1.26) [0.006]	(1.25) [0.005]	(1.27) [0.006]
Agricultural day-labor days (7 days), 2-year window, land-poor	−0.248	(0.095) [0.010]	(0.094) [0.009]	(0.095) [0.009]
Any non-farm self-employment (pp), flood window, land-poor	2.36	(1.28) [0.066]	(1.27) [0.063]	(1.26) [0.062]
C. Rent				
Log Union median cash rent per decimal, flood window	−0.237	(0.083) [0.006]	(0.062) [0.000]	(0.063) [0.000]
Log cash rent received per decimal, same plot, flood window	−0.400	(0.084) [0.000]	(0.080) [0.000]	(0.081) [0.000]

Notes: The table reports the coefficients of Table 10 with alternative standard errors in parentheses and p-values in brackets. The treatment is the number of severe floods, years in which at least a quarter of the Union is under water, in the window between two interviews. Outcomes and fixed effects are as in Table 10. The Union-clustered column clusters at the Union level, with p-values from $t(G-1)$. The Conley columns are spatial HAC standard errors (Conley 1999) with a Bartlett kernel at 15 and 25 km between Union centroids within survey round, allowing unrestricted serial correlation within a Union across rounds.

flood-window earnings fall by Tk 12,606 instead of Tk 10,764. In the 2-year window the land-poor's participation in agricultural day labor falls by 7.0 points instead of 5.3 and the top tercile's hired harvest hours by 24.3 instead of 17.3, each significant at the 5 percent level. Only the hired harvest hours on the land-poor's own plots lose significance at the 10 percent level (Online Appendix Table OA-88). The fill therefore attenuates the capital and rent losses rather than producing them. It does not produce the fall in earnings or in day labor, while the hired harvest hours on the land-poor's own plots depend on it.

An error in the 2018 record would affect every estimate that uses round 3, not only those the fill deletion changes. I therefore also drop round 3 (Online Appendix Table OA-88, last column). A two-round estimate, however, is a single first difference, identified mostly from Unions whose flooding stopped between rounds 1 and 2. It need not equal the effect of a flood (Online Appendix OA-10.5). The hired harvest hours of all households and of the top tercile and the land-poor's days, participation and daily-wage income in the 2-year window still fall, each significant at the 10 percent level. The land-poor's days in the flood window lose significance, and the wage bound widens to 10.9 percent, still below the fall in their days in the window, 26 percent of their round-1 mean. The top tercile's capital in the flood window (asset index -0.021) and its flood-window earnings ($-$Tk 5,341, not significant at the 10 percent level) are not statistically distinguishable from zero. Estimated window by window on all three rounds, however, the top tercile's asset index falls in the first window as well as the third (Online Appendix Table OA-79). The labor results therefore hold without round 3. For the capital and earnings results the two-round estimate measures a different quantity. The fill variants above, under which the capital loss is larger, are the test of the 2018 record.

10.6 Attrition and Other Checks

The remaining checks concern attrition and three features of the design.

Households that leave the panel after a flood may differ from those that stay, so exit caused by floods could bias the estimates. I test whether floods raise exit, comparing exit within a district over the same pair of interviews. I do so for all households and for each group (Online Appendix Tables OA-71 and OA-72). A severe flood between two interviews changes the exit rate by 1.2 points on a base of 8.2 percent ($p = 0.088$; wild bootstrap $p = 0.134$). The 90 percent interval rules out a rise larger than 2.3 points, and the window-after coefficient is negative. By group, the land-poor's exit rises by 1.2 points, significant at the 10 percent level on clustered standard errors but not under the wild bootstrap ($p = 0.099$; wild bootstrap $p = 0.143$). The intervals rule out rises larger than 2.4 points for the land-

poor and 4.0 points for the top tercile, and equality of the four effects cannot be rejected ($p = 0.395$). Floods therefore raise exit by little in any group, and the bounds limit how far exit could bias the estimates.

Exit could also be selected on baseline characteristics that shape the response to a flood. I reweight the key estimates by inverse probabilities of retention from a model with baseline covariates and district effects. Its flood coefficient is small and not statistically distinguishable from zero ($p = 0.72$). The weights move the top tercile's asset index by 0.13 of a standard error (Online Appendix Table OA-73). They leave the Union median rent at −21.7 percent (Online Appendix Table OA-43) and the top tercile's flood-window earnings at −Tk 10,988 against −Tk 10,764 (Online Appendix Table OA-88). The results are therefore robust to attrition. The weights address selection on baseline characteristics, and the bounds above address exit caused by floods.

The main specification gives a flood the same effect in each exposure window, which could conceal an effect that changed over the decade. I estimate it window by window (Online Appendix Table OA-79). The land-poor's participation in agricultural day labor falls by 7.3 points per flood in the window of round 1, by 4.0 in that of round 2 and by 2.1 in that of round 3, only the first significant on its own. Equality of the three is rejected on the clustered p-value ($p = 0.01$). For the top tercile's asset index, at −0.187, −0.168 and −0.262, equality cannot be rejected ($p = 0.57$). I therefore do not claim a constant labor response over the decade.

The results could depend on where the tercile cuts fall. I split the inheritors at their median instead of into terciles (Section 5.1; Online Appendix Tables OA-81 to OA-84). Above the median, the asset index falls by 0.17 standard deviations in the flood window ($p = 0.008$). The flood-window estimate for wage and self-employment income, −Tk 7,209 (standard error 4,174), is significant at the 10 percent level on clustered standard errors ($p = 0.085$; wild bootstrap $p = 0.131$). Of the 129 window and window-after cells that can be compared, 113 keep their significance status, 10 lose significance, the intervals of 2 precisely estimated zeros widen, and 4 gain significance (Online Appendix Tables OA-81 to OA-84). The estimates that lose significance are responses of an extreme tercile, among them the top tercile's capital in the window after and the smallest tercile's larger farm in the flood window. Each loses significance on the clustered or the wild bootstrap p-value, and in each the estimate shrinks, by a quarter to almost two thirds, as the split mixes an extreme tercile with part of the middle one. Among those that gain significance, rented-in land at or above the median falls by 5.9 decimals in the flood window, significantly, where the top tercile's estimate is not (Section 6). The flood-window capital loss is therefore robust to where the tercile cuts fall. The earnings loss and the responses concentrated in the extreme

terciles are not.

Classifying households by inherited rather than farmed land could also shape the results. I classify them instead by the land operated in 2011, which is not predetermined (Online Appendix Table OA-15). Among the largest operators the asset index changes by −0.06 standard deviations in the flood window (90 percent interval −0.14 to +0.02) and working capital by −Tk 1,830 (−4,016 to +356). For the top inherited tercile the same two estimates fall significantly, by 0.25 standard deviations and Tk 6,152 (Online Appendix Table OA-16). The capital loss therefore holds by inherited land but not among the largest operators. Because I do not test the difference and operated land responds to earlier floods, I use this split only descriptively.

Taken together, the checks weaken the evidence on the size of the rent fall, on the fall in the land-poor's migration and on the earnings loss, which loses significance under the trend control. None of them reverses the capital loss in the flood window, the cut in hiring or the fall in day labor.

11 Calibration, Model Fit and Policy Counterfactuals

In this section I calibrate the model of Section 3 to the estimates, compare the calibrated model with estimates it is not fitted to, and use it to compare policies after a flood. The model is the full model of Online Appendix OA-7, applied to the 2011 households of the balanced panel. For the farms that lose a herd, the replacement of lost capital is a choice, as in the rebuilding model of Online Appendix OA-9.

11.1 Calibration

The parameters fall into three groups, which Online Appendix Table OA-63 lists with their sources. Most are set from the round-1 data, one is set from other studies, and two are chosen to match a flood response.

I set land's and labor's shares of output, α and β, to 0.261 and 0.191, their shares in the value of the 2011 harvest computed from rents and wages. I set the transaction cost on leases so that the model reproduces the shares of 2011 landowners who rent in and who rent out, as Chen et al. (2022) set land-market costs. The share of a change in the rent that a running cash lease passes on is 0.25, estimated from the same household's rents for floods during a tenancy and for floods before it began. The top tercile's loss of working capital follows the window estimates of Section 6. From other studies, I set $\gamma = 1$, so that a farm's scale is proportional to its working capital, the constant-returns case of the span-of-control

literature.

Two parameters are chosen to match flood responses. The crop loss, $\zeta = 0.046$, matches the point estimate of the fall in net crop income over all households in the flood window, Tk 973 (standard error 661), which is not statistically distinguishable from zero. It is smaller than the 8.9 percent fall in the flood-year harvest of the households that farm (Section 6), because the target averages over every household and over the three years of the window, as the model's other flood responses do. The second is a common inward shift in every operator's demand for land of $x = 0.47$ log points in the three years of the flood window, which matches the fall in the Union median rent. A common shift lowers the rent and moves no land (Online Appendix Proposition OA-2). It is close in size to the land-poor's own lower-bound shift in Section 8, but it is chosen to fit the rent, so the rent is not a test of the model.

In the rebuilding model, a farm that loses two or three head has no credit after the flood, replaces animals one at a time at their measured return of 0.20 a year, and keeps its consumption above a food floor. Its herd, income, savings and loss are set to those of the cattle holders in the sample, and none of its parameters is chosen to match the estimates it is compared with.

11.2 Model Fit

Table 12 compares the calibrated model with the estimates. The first two rows are the targeted moments, and the remaining rows are not targeted.

The model matches the sign of the labor-market estimates but understates their size. With the wage held at its floor, it predicts no change in the farm wage, within the estimated bound, and falls of 7 percent in hired harvest hours and in the land-poor's days of day labor, against estimated falls of 14.8 and 16.9 percent. A market-clearing wage would cut the land-poor's days by at most 4.1 percent at any elasticity of labor supply, outside the 90 percent interval of the estimate. Within the model, only the wage norm is consistent with the estimated fall in days.

The model also matches the consumption and capital estimates it is not fitted to. It predicts no change in the land-poor's consumption, and the estimate, +1.6 percent (standard error 1.0), is not statistically distinguishable from zero. In the rebuilding model, the missing capital of a farm that loses three head, times the share of farms that lose a herd, is Tk 3,400 to 5,200 in the flood window, against the measured Tk 6,152 (standard error 1,561). In the window after, the model's farms rebuild faster than the data show, though less so under the poverty-line floor. A farm that loses three head keeps its consumption 8 to 10 percent below

normal for three years, so the group's consumption falls by 1.3 to 1.9 percent, close to the 1.8 percent fall that the 90 percent interval rules out.

The model does less well on the timing of the rent. It comes close to the same-plot rent after floods in the two calendar years before the interview, −23.4 against −26.2 log points, which it is not fitted to. Over the whole flood window, however, the model's same-plot rent falls by 23.7 log points against an estimated 8.2 (standard error 6.4). The reason is that in the data the same plot's rent falls after recent floods and not after earlier ones, while the model's shift lasts all three years. Because the shift is common to every operator, the model leaves the land-poor's rented-in land almost unchanged (+0.4 decimals), and the estimate, −2.5 (standard error 1.5), is not statistically distinguishable from zero. The model overstates the top tercile's retreat from rented land, −4.6 and −6.6 decimals against estimates of −3.2 and −4.0.

The model takes the top tercile's loss of capital from the data and does not explain it, and the rebuilding model offers one account of why that capital stays low. The model omits the exit from cropping, the land-poor's lease exits and migration. It also fits the fall in the rent through the shift it is given rather than explaining it. Section 8 shows that neither the land the top tercile gives up nor beliefs about future floods account for more than a small part of that fall.

11.3 Policy Counterfactuals

I use the rebuilding model of Online Appendix OA-9 to ask what relief after a flood is worth. In this model a cattle holder of the top tercile loses part of its herd, has no credit after the flood, and chooses when to replace the animals out of its income and savings. The results are model calculations, not estimates. I measure the farm's welfare gain from a policy as the permanent percentage increase in consumption that makes it indifferent between living without the policy and living with it, net of any premium it pays.

I compare four policies. No animal is bought in the flood year, so every payment is first spent in year 2.

1. *Index insurance.* Every top-tercile farm buys insurance before the flood. It pays a face value of Tk 8,660, the expected loss per cattle holder per flood, whenever the trigger fires, at this flood and at every later one, and it costs the fair premium of 0.125 times the face value a year. The farms pay the premium, so the expected fiscal cost is zero. In a variant, full cover, the face value is the affected farm's whole loss.

2. *A grant to the farm.* A public grant pays each farm that loses a herd the value of its loss once. For a three-head loss the budget is Tk 2.80 to 4.31 million per village-equivalent,

Table 12: The calibrated model against the estimates

Outcome	Estimate	Model	Comparison
Fitted			
Net crop income, all households, flood window (Tk)	−973 (661)	−973	fitted (ζ)
Union median cash rent, flood window (log points)	−23.7 (8.3)	−23.7	fitted (common demand shift)
Not fitted			
Hired harvest hours, flood window (percent)	−14.8 (6.8)	−7.0 (wage floor); −1.8 (flexible wage)	same sign, about half
Land-poor's agricultural day-labor days, flood window (percent)	−16.9 (7.1)	−7.0 (wage floor); −1.8 (flexible wage)	same sign, smaller
Farm wage, flood window (log points): paid by farmers; land-poor's	−0.2 (1.1); +0.7 (2.0)	0 (wage floor); −1.9 (flexible wage)	within the bound under the floor
Top tercile's rented-in land, flood window; window after (decimals)	−3.2 (2.7); −4.0 (2.7)	−4.6; −6.6	same sign, larger
Same-plot cash rent, floods in the two calendar years before the interview (log points)	−26.2 (11.8)	−23.4	close to the estimate
Same-plot cash rent, flood window (log points)	−8.2 (6.4)	−23.7	larger than the estimate
Land-poor's rented-in land, flood window (decimals)	−2.5 (1.5)	+0.4	within the interval
Land-poor's consumption, flood window (percent)	+1.6 (1.0)	0.0	within the bound
Top tercile's missing working capital (Tk): flood window; window after	6,152 (1,561); 3,774 (1,800)	three-head loss, times the affected share of 0.13 to 0.20: 3,392–5,219; 360–554 (food floor), 3,610–5,553; 1,876–2,887 (poverty-line floor)	flood window matched, window after short
Top tercile's consumption (percent)	no fall larger than 1.8; cattle holders +1.1 (2.9)	three-head farm about −10 in years 1–3, times the share −1.3 to −1.9	close to the bound
Top tercile's savings (Tk)	cattle holders −10,419	the farm with savings spends them on animals	same sign

Notes: The table compares the calibrated village model of Online Appendix OA-8.1 and the rebuilding model of Online Appendix OA-9 with the estimates of Sections 6 to 8, which are given with standard errors in parentheses (the savings from Online Appendix Table OA-23). Model values are means over the years each window covers, with the wage held at its floor unless stated, and the capital, consumption and savings rows come from the rebuilding model (Online Appendix Table OA-64). The crop loss ζ is fitted to the first row and a common inward shift in every operator's demand for land of 0.47 log points in years 1–3 to the second, so the Union median rent is not an independent check. No input of the rebuilding model is fitted to the rows it is compared with.

the 539 top-tercile farms of the sample treated as one village.

3. *The same budget to the land-poor.* The grant's budget is spread over the 3,845 land-poor households as a one-time transfer of Tk 729 to 1,121 each.

4. *A credit line.* Each farm can borrow up to Tk 7,500 at 25 percent a year from the flood year on, the size of the emergency loan in Lane (2024).

Table 13 reports the comparison for a farm that loses three head under a food floor. Online Appendix Tables OA-65 to OA-68 report a two-head loss and the poverty-line floor.

Table 13: Four policies in the rebuilding model for a farm that loses three head: the affected farm, the village and the fiscal cost

	The affected farm			The village		
Policy	Capital missing, year 3 (Tk)	Consumption, year 1 (%)	Welfare (% of consumption)	Output lost, PV (% of a year)	Top tercile's hiring, year 2 (%)	Fiscal cost (Tk million)
No policy	19,141	−10.4	0.0	10.5–13.1	−14 to −22	0
Index insurance	10,825	−11.1	+0.4	9.1–11.1	−14 to −22	0
Index insurance, full cover	0	−6.9	+1.5	8.0–9.3	−14 to −22	0
Grant to the farm	0	−5.1	+1.7	8.0–9.3	−14 to −22	2.80–4.31
Credit line	16,633	−10.5	+0.0	10.1–12.5	−14 to −22	0
Same budget to the land-poor (Tk 729 to 1,121 each)	no change in the model					2.80–4.31

Notes: The table reports the outcomes of each policy for a top-tercile cattle holder that holds working capital worth four head and loses three in a flood (Tk 39,999). The model is the rebuilding model of Online Appendix OA-9, with exponential discounting, a food floor of Tk 76,440 a year and no purchases in the flood year, and the farm columns are means over the cattle holders' savings distribution. Index insurance pays a face value of Tk 8,660 (the whole loss under full cover) at every flood against a fair premium of 0.125 of the face value a year, the grant pays the loss once, and the credit line is a revolving line of Tk 7,500 at 25 percent, the size of the emergency loan in Lane (2024). Capital missing is the working capital missing at the start of year 3, consumption is the change in year 1 against normal, and welfare is the permanent change in consumption, in percent of normal consumption, that makes the farm indifferent between the row and no policy. The village columns feed the affected farms' path into the calibrated village at an affected share of 0.13 to 0.20 of top-tercile farms, and fiscal cost per village-equivalent is the budget of the grant, the affected share times 539 top-tercile farms times the loss, and zero in expectation for the fair insurance and the repaid credit line.

Without a policy the farm that loses a herd rebuilds slowly. It misses its whole loss of Tk 39,999 in year 2 and Tk 19,141 in years 3 and 4, has its herd back in year 6, and keeps its consumption 8 to 10 percent below normal for three years (Table 13; Online Appendix Tables OA-65 and OA-66). A grant equal to the loss lets the farm replace its herd in year 2

and halves its consumption cut, to about 5 percent. The grant raises welfare by the equivalent of a permanent 1.7 percent increase in consumption. Full cover, which also pays the whole loss at every later flood, raises it by 1.5 percent net of its premium. Full cover therefore delivers the same rebuilding and 87 percent of the grant's welfare gain at no expected fiscal cost. Its fair premium is about Tk 5,000 a year per farm, or Tk 46.2 million in present value for the 539 insured farms. This is more than the grant budget of Tk 2.8 to 4.3 million because full cover insures every later flood as well. Index insurance raises welfare by 0.4 percent. The farm receives Tk 8,660, about a fifth of its loss, pays Tk 1,083 a year, and still misses Tk 10,825 of capital in years 3 and 4. The credit line raises welfare by close to zero, because at 25 percent the farm borrows little and rebuilds almost as slowly as without it. The ranking is the same under the poverty-line floor, where the grant is worth 2.2 percent and index insurance 0.5 percent, and after a two-head loss, where they are worth 1.1 and 0.3 percent (Online Appendix Table OA-65).

Faster rebuilding also restores part of the village's output. Without a policy, the output lost over the horizon is 10.5 to 13.1 percent of a year's village output, for an affected share of 0.13 to 0.20 of top-tercile farms. The grant and full cover lower this loss to 8.0 to 9.3 percent, index insurance to 9.1 to 11.1 percent, and the credit line to 10.1 to 12.5 percent. No policy changes the top tercile's hiring in year 2, which falls by 14 to 22 percent because the whole loss is still missing then. This is close to the measured fall of nearly a quarter of a year's harvest hours (Section 7). Per taka of its budget, the grant averts Tk 1.20 of lost village output (Online Appendix Table OA-68). Index insurance averts Tk 0.18 to 0.28 per taka of the farms' premium, which is Tk 10.0 million in present value for the 539 insured farms. Full cover, whose premium is Tk 46.2 million, averts Tk 0.07 to 0.11 per taka.

In the model no policy changes the land-poor's outcomes, and a transfer of the same budget to them changes no price. In the model's accounting their loss is about Tk 510 per household under every policy, because the policies act from year 2 and the land-poor's loss comes from rationing in the flood year. That figure rests on the model's value of an unsold day, which is why Section 6 measures incidence from income instead. The transfer by land ownership is a one-time payment of Tk 729 to 1,121 per household, less than half of one month's real expenditure per person in round 1 (Table 1). The model has no channel through which the land-poor's cash changes a price or a quantity, and the estimates do not show one. The land-poor's livestock and working capital show no significant fall at any horizon, and the 90 percent interval rules out a fall in their consumption of more than 0.05 percent in the flood window (Section 6).

Three features of the rebuilding model set the size of these numbers. First, the farm rebuilds at the measured return on an animal of 0.20 a year, while the village's production

function values the stock in place at 0.95. At the higher return the farm rebuilds faster out of its own income, and the grant is worth 2.7 percent (Online Appendix OA-9). Second, the insurance is priced fairly and every top-tercile farm buys it. With take-up of 40 to 50 percent at the fair price, as in Karlan et al. (2014), its village effects scale with take-up, and a loaded premium raises the farms' cost by the load. Third, the affected share of 0.13 to 0.20 comes from the livestock inventory (Section 6) and scales every village number. The rebuilding model also omits the land market, migration and the exit from cropping, and it says nothing about what a transfer to the land-poor would do through local demand, the channel that Egger et al. (2022) measure in Kenya.

12 Conclusion

This paper asks who bears the loss when a flood hits a village economy, and who adapts. Two bodies of evidence disagree about how such a loss travels, through a lower wage or through rationed work, but they share a premise. The households that own land and capital are the ones best placed to absorb a flood. The loss moves down the land distribution to those who sell their labor. I use a panel of 5,473 households in rural Bangladesh matched to satellite flood maps, and compare households by the land they inherited within district and round. The answers to the paper's three questions overturn that premise.

First, the loss falls on capital, and capital belongs to the land-rich. The top tercile of inheritors loses livestock and working capital, and its working capital remains lower a window later. The loss is concentrated on its cattle-holding farms, which draw down savings: a flood raises the share of them that lose two or more head by 24 percentage points. Its consumption holds, consistent with drawing on every buffer it has, savings, land, animals and a shift from production credit to microfinance, while its total debt shows no significant change. Its farm shrinks as it gives up land it sharecropped and rents out more of its own, and a window later the share of its households that grows a crop is 5.7 percentage points lower. Its wage and self-employment income falls by 7.6 percent of a year's income in the flood window, most of it non-farm business income, significantly more than the land-poor's. Second, the labor market rations work rather than repricing it, and the households it rations adapt. Over the flood window the land-poor's days of day labor fall by 17 percent. Per flood in the two years before an interview they fall by 23 percent, while farms hire a quarter of a year's harvest hours less. The farm wage falls by at most a few percent. The land-poor turn to work of their own. Their wage and self-employment income falls by at most 4 percent of a year's income, their consumption by at most 0.05 percent, and their savings show no fall. Third, the land market reprices, but no group rents in more land in total at the lower rent. The rent falls

by about a fifth, the land-poor do not rent in more, and the distribution of operated land does not change. The land-poor adapt better because labor can be redeployed and capital cannot. They turn from day labor to work of their own, while the land-rich lose animals and equipment that the water reaches and cannot move, so their farms shrink with their capital. I interpret these results as showing that who bears a flood follows who owns the capital it hits, not who is poor.

For the literature on village labor markets, the results side with the second body of evidence. Where the wage falls, the landless bear cultivators' losses (Jayachandran 2006; Mobarak and Rosenzweig 2014; Balboni et al. 2026). Here the data bound the fall in the wage at a few percent while days of work fall far more, the pattern Kaur (2019) and Breza et al. (2021) report. A market that clears with an elastic supply of labor fits the same pattern in direction, but within the calibrated model a market-clearing wage cannot produce the measured fall in days at any elasticity of labor supply, so the data favor rationing. For the literature on land and farm size, a fall of a fifth in the rent leads no group to rent in more land in total and leaves the distribution of operated land unchanged. For the literature on assets after shocks, the capital lost is at the top while the land-poor's consumption holds.

The results have implications for policy, which Section 11 develops as a model exercise. It calibrates the framework to the estimates, checks its fit, and compares four policies in a rebuilding model in which a farm without credit chooses when to replace lost animals. In the rebuilding model a grant to the farms that lose a herd lets them rebuild within a year instead of five and is worth about 2 percent of their consumption. A fairly priced index insurance that pays the full loss gives the same rebuilding and about nine tenths of the grant's welfare gain at no expected fiscal cost, since the farms pay its premium. Index insurance that pays only the average loss per flood is worth less than half a percent, since the farm that loses a herd receives a fifth of its loss. A credit line at 25 percent open from the flood year is worth almost nothing. The same budget spread over the land-poor is a one-time payment of Tk 729 to 1,121 per household. It moves no market in the model, which gives their cash no channel into one. The first implication concerns targeting. Relief after a flood is commonly aimed at the poor. In villages like these it would reach households whose measured earnings the data bound within a few percent. It would miss the largest inherited farms, whose capital and earnings fall. Transfers to the land-poor would not, by themselves, restore the days farms no longer hire. The land-poor already adapt by turning to work of their own. The second concerns the farms that hire. The land-poor lose days of hired work, so what restores those days is a payment that lets the farms that lost a herd replace it sooner. In the rebuilding model a grant restores their capital, and with it their hiring, from the third year after the flood instead of the sixth, while a loan at commercial terms does little. Lane (2024) finds

that credit guaranteed in case of a flood led farmers in Bangladesh to rent more land. The terms of the credit, not its availability alone, therefore determine whether it reaches that margin. These implications hold within the rebuilding model's premises, in particular that the farm has no credit after the flood and replaces its animals at their measured return.

The questions are not specific to Bangladesh. Wherever floods destroy productive assets in villages where land is unequally held and worked with casual labor, who bears the loss is decided in the same markets and by what households can adjust. The results here suggest asking first who owns the capital the flood hits. Three kinds of data would allow the analysis to go further. The first is the date and the terms of each lease. With them the fall in the rent and the top tercile's retreat from sharecropping could be read from the contracts. The second is the landlord side of the tenancy market, most of which lies outside a household sample. The third is the value of a lost day of hired work to a household that fills it with work of its own. It would turn the land-poor's measured loss of income into a measure of what the lost days cost them.

References

Adamopoulos, T., L. Brandt, J. Leight, and D. Restuccia (2022). Misallocation, selection, and productivity: A quantitative analysis with panel data from China. *Econometrica 90*(3), 1261–1282.

Adamopoulos, T. and D. Restuccia (2014). The size distribution of farms and international productivity differences. *American Economic Review 104*(6), 1667–1697.

Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *Journal of the American Statistical Association 103*(484), 1481–1495.

Balboni, C., O. Bandiera, R. Burgess, M. Ghatak, and A. Heil (2022). Why do people stay poor? *Quarterly Journal of Economics 137*(2), 785–844.

Balboni, C., O. Bandiera, R. Burgess, A. Heil, C. Mazet-Sonilhac, M. Sulaiman, and Y. Wang (2026). Weathering poverty. Working paper, London School of Economics.

Bellemare, M. F. and C. J. Wichman (2020). Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics 82*(1), 50–61.

Biscaye, P. E. and A. Isah (2026). Do floods induce long-term livelihood diversification? Evidence from Nigeria. STEG Working Paper 150.

Borusyak, K., X. Jaravel, and J. Spiess (2024). Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies 91*(6), 3253–3285.

Breza, E., S. Kaur, and Y. Shamdasani (2021). Labor rationing. *American Economic Review 111*(10), 3184–3224.

Bryan, G., S. Chowdhury, and A. M. Mobarak (2014). Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh. *Econometrica 82*(5), 1671–1748.

Cain, M. (1981). Risk and insurance: Perspectives on fertility and agrarian change in India and Bangladesh. *Population and Development Review 7*(3), 435–474.

Callaway, B. and P. H. C. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics 90*(3), 414–427.

Carter, A. V., K. T. Schnepel, and D. G. Steigerwald (2017). Asymptotic behavior of a t-test robust to cluster heterogeneity. *Review of Economics and Statistics 99*(4), 698–709.

Carter, M. R. and C. B. Barrett (2006). The economics of poverty traps and persistent poverty: An asset-based approach. *Journal of Development Studies 42*(2), 178–199.

Carter, M. R. and T. J. Lybbert (2012). Consumption versus asset smoothing: Testing the implications of poverty trap theory in Burkina Faso. *Journal of Development Economics 99*(2), 255–264.

Chen, C., D. Restuccia, and R. Santaeulàlia-Llopis (2022). The effects of land markets on resource allocation and agricultural productivity. *Review of Economic Dynamics 45*, 41–54.

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

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics 92*(1), 1–45.

de Chaisemartin, C. and X. D'Haultfœuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–2996.

Deininger, K. and S. Jin (2005). The potential of land rental markets in the process of economic development: Evidence from China. *Journal of Development Economics 78*, 241–270.

Deininger, K. and S. Jin (2008). Land sales and rental markets in transition: Evidence from rural Vietnam. *Oxford Bulletin of Economics and Statistics 70*(1), 67–101.

del Ninno, C., P. A. Dorosh, and L. C. Smith (2003). Public policy, markets and household coping strategies in Bangladesh: Avoiding a food security crisis following the 1998 floods. *World Development 31*(7), 1221–1238.

del Ninno, C., P. A. Dorosh, L. C. Smith, and D. K. Roy (2001). The 1998 floods in bangladesh: Disaster impacts, household coping strategies, and response. Research Report 122, International Food Policy Research Institute, Washington, DC.

Dell, M., B. F. Jones, and B. A. Olken (2014). What do we learn from the weather? the new climate-economy literature. *Journal of Economic Literature 52*(3), 740–798.

Egger, D., J. Haushofer, E. Miguel, P. Niehaus, and M. Walker (2022). General equilibrium effects of cash transfers: Experimental evidence from Kenya. *Econometrica 90*(6), 2603–2643.

Eskander, S. M. S. U. and E. B. Barbier (2023). Adaptation to natural disasters through the agricultural

land rental market: Evidence from Bangladesh. *Land Economics 99*(1), 141–160.
Eswaran, M. and A. Kotwal (1986). Access to capital and agrarian production organisation. *Economic Journal 96*(382), 482–498.
Foster, A. D. and M. R. Rosenzweig (2022). Are there too many farms in the world? Labor market transaction costs, machine capacities, and optimal farm size. *Journal of Political Economy 130*(3), 636–680.
Gibbons, C. E., J. C. Suárez Serrato, and M. B. Urbancic (2019). Broken or fixed effects? *Journal of Econometric Methods 8*(1), 20170002.
Government of Pakistan and Asian Development Bank and European Union and United Nations Development Programme and World Bank (2022). Pakistan floods 2022: Post-disaster needs assessment. Technical report, Government of Pakistan, Islamabad.
Gray, C. L. and V. Mueller (2012). Natural disasters and population mobility in Bangladesh. *Proceedings of the National Academy of Sciences 109*(16), 6000–6005.
Hornbeck, R. and S. Naidu (2014). When the levee breaks: Black migration and economic development in the American South. *American Economic Review 104*(3), 963–990.
Hsiang, S. M. and A. S. Jina (2014). The causal effect of environmental catastrophe on long-run economic growth: Evidence from 6,700 cyclones. *NBER Working Paper 20352*.
Jayachandran, S. (2006). Selling labor low: Wage responses to productivity shocks in developing countries. *Journal of Political Economy 114*(3), 538–575.
Karlan, D., R. Osei, I. Osei-Akoto, and C. Udry (2014). Agricultural decisions after relaxing credit and risk constraints. *Quarterly Journal of Economics 129*(2), 597–652.
Kaur, S. (2019). Nominal wage rigidity in village labor markets. *American Economic Review 109*(10), 3585–3616.
Khandker, S. R. (2007). Coping with flood: Role of institutions in Bangladesh. *Agricultural Economics 36*(2), 169–180.
Kocornik-Mina, A., T. K. J. McDermott, G. Michaels, and F. Rauch (2020). Flooded cities. *American Economic Journal: Applied Economics 12*(2), 35–66.
Lane, G. (2024). Adapting to climate risk with guaranteed credit: Evidence from Bangladesh. *Econometrica 92*(2), 355–386.
Lucas, Robert E., J. (1978). On the size distribution of business firms. *Bell Journal of Economics 9*(2), 508–523.
Mobarak, A. M. and M. R. Rosenzweig (2014). Risk, insurance and wages in general equilibrium. Working Paper 19811, National Bureau of Economic Research.
Mueller, V. and A. Quisumbing (2011). How resilient are labour markets to natural disasters? the case of the 1998 Bangladesh flood. *Journal of Development Studies 47*(12), 1954–1971.
Patel, D. (2024). Floods. Working paper, Harvard University.
Rosenzweig, M. R. and K. I. Wolpin (1993). Credit market constraints, consumption smoothing, and the accumulation of durable production assets in low-income countries: Investments in bullocks in India. *Journal of Political Economy 101*(2), 223–244.
Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199.
Tellman, B., J. A. Sullivan, C. Kuhn, A. J. Kettner, C. S. Doyle, G. R. Brakenridge, T. A. Erickson, and D. A. Slayback (2021). Satellite imaging reveals increased proportion of population exposed to floods. *Nature 596*, 80–86.
UNICEF (2023). One year on from catastrophic floods, millions of children in Pakistan still need urgent support. Press release, 25 August 2023.
World Bank (2019). Program brief: Vulnerable group feeding. Technical report, World Bank, Bangladesh Social Protection Program Briefs, Dhaka.

A Negative Weights under a Non-Absorbing Treatment

This appendix supports the estimator of Section 5.1 and the robustness checks of Section 10 for the results of Sections 6 to 9. With a treatment that switches on and off (Section 5.1), negative weights let an average effect of one sign produce a coefficient of the opposite sign when effects are heterogeneous enough (de Chaisemartin and D'Haultfœuille 2020). The estimators designed for staggered adoption (Callaway and Sant'Anna 2021; Sun and Abraham 2021; Borusyak et al. 2024) do not apply without recoding the treatment (Section 5.1). The DID_M estimator of de Chaisemartin and D'Haultfœuille (2020) allows such a treatment but is defined for a binary one, while the treatment here is a count of floods. I therefore measure how large the negative weights are and report an interaction-weighted version of the main specification's estimator that lets the effect differ by group, round and flood history and averages the cells with non-negative weights. I also report two full-sample diagnostics of the heterogeneity behind the weights. Letting the effect differ across cells leaves the sign of every key estimate unchanged and the top tercile's capital losses at least as large. Hence the negative weights across rounds and flood histories do not produce the sign or the size of any key effect. The weight within households flooded in every round remains, and I measure it.

Negative weights can bias the estimate only if they carry real weight in this design. By the Frisch–Waugh theorem, the estimator can be written as $\hat{\beta} = \sum_i \varepsilon_i y_i / \sum_i \varepsilon_i D_i$, where ε_i is the residual from regressing the flood count on all other variables in the specification. Under $y_i = y_i(0) + D_i\tau_i$, this equals $\sum_i w_i \tau_i$ with $w_i = \varepsilon_i D_i / \sum \varepsilon D$, weights that sum to one across flood-exposed observations and are negative wherever the residual is negative. In the regression of each inherited-land group alone, between 28 and 39 percent of exposed household-rounds have a negative weight, and the negative weights sum to between -0.24 and -0.26, against $+1.24$ to $+1.26$ for the positive ones (Table 14). In the main specification they sum to between -0.22 and -0.34 on the key rows (Table 15). So negative weights are common here, and whether they matter depends on how heterogeneous the effects are.

Heterogeneity of modest size could pull an average effect to zero. de Chaisemartin and D'Haultfœuille (2020) provide the statistic that measures how much is needed. Their weights are defined on treated cells, here Union × round, the level at which the treatment varies, and are normalized to a size-weighted mean of one, with $\sigma(w)$ their size-weighted standard deviation around one. By their Corollary 1(i), $|\hat{\beta}|/\sigma(w)$ is the smallest standard deviation of the treatment effect across cells under which the true average effect could be zero. Their Corollary 1(ii) gives the smallest under which every cell effect has the opposite sign to the coefficient. For the top tercile's asset index, a standard deviation of 0.14 index points across Union-rounds, about two-fifths of the estimated 0.327 in this sample, would be enough for

the true average to be zero. One of 0.45 would be enough for every cell effect to be positive. Because the weights do not depend on the outcome, the ratio is a property of the design: 0.43 of the estimated effect in the top tercile and 0.47 to 0.51 in the other inherited-land groups. It cannot be computed for the Union-level and plot-level coefficients, which are estimated on different samples. Hence the heterogeneity needed for a zero average is smaller than the estimates. Reversing the sign, however, requires heterogeneity larger than the estimate itself *and* unfavorably correlated with the weights. I therefore check the key results with an estimator that removes the negative weight across rounds and flood histories.

Table 14: How much effect heterogeneity would be needed to overturn the sign of the flood-window estimates

	No inherited land	Smallest tercile	Middle tercile	Top tercile
Asset index	0.000 (0.00)	0.051 (0.02)	0.154 (0.08)	−0.327 (0.14)
Log livestock value	0.100 (0.05)	−0.143 (0.07)	−0.102 (0.05)	−0.528 (0.23)
Log productive assets	−0.125 (0.06)	−0.179 (0.09)	−0.048 (0.02)	−0.225 (0.10)
Rented in (dec.)	−2.989 (1.42)	3.497 (1.68)	−6.626 (3.38)	−5.452 (2.36)
Operated land (dec.)	−3.375 (1.60)	7.362 (3.54)	−7.057 (3.60)	−1.772 (0.77)
Owned land (dec.)	−0.317 (0.15)	2.510 (1.21)	1.416 (0.72)	3.487 (1.51)
Flooded household-rounds	2,163	281	240	393
Negative weight (percent)	39	38	28	34
Sum of negative weights	−0.26	−0.26	−0.24	−0.26

Notes: The table reports, for each inherited-land group estimated alone, the coefficient from the main specification (equation (1)) without the group interactions, with one observation per household-round. The treatment is the number of severe floods, years in which at least a quarter of the Union is under water, in the window between two interviews. In parentheses, in place of a standard error, is the sensitivity statistic of de Chaisemartin and D'Haultfœuille (2020, Corollary 1(i)), the smallest standard deviation of the per-flood effect across flooded Union-rounds under which the true average effect could be zero. The bottom rows give, on the asset-index sample, the flooded household-rounds, the percent of them with a negative Frisch–Waugh weight and the sum of those weights (all weights sum to one).

Negative weight arises in two ways. Across household-rounds, the composition of the flooded Unions changes as the national cycle turns wet and dry. Unions flooded in successive windows serve as controls for Unions that have just been flooded. Within a household, a household flooded in every round identifies an effect only from changes in its count, so its rounds with fewer floods serve as controls for its rounds with more. The estimator removes the first source by giving each cell its own coefficient and averaging the cells with non-negative weights. It keeps every household, so it leaves the second in place; Online Appendix Table OA-18 measures it. The negative weights bias the main estimate only if the effect of a flood differs systematically between the observations that receive negative weight and the rest. The cells are defined by inherited-land group, round and flood history. Formally, I

estimate

$$y_{it} = \sum_{c} \beta_c \, F_{ut} \, \mathbf{1}[c(i,t) = c] + \beta^1 F_{u,t-1} + \sum_{g} \gamma^1_g F_{u,t-1} S_{ig} + \sum_{g} \sum_{r} \delta_{gr} S_{ig} \mathbf{1}[t = r] + \Gamma' X_{ut} + \psi_{dt} + \alpha_i + \lambda_t + \varepsilon_{it}$$

on the main design's sample, where $c(i,t)$ is the cell of household i in round t. I average each group's cell coefficients with weights equal to its share of floods in each cell, on the estimation sample. For the results measured in the window after, I apply the same construction to the previous window's count, with the indicator for a flood in the current window. For the outcomes without inherited-land groups, the cells are round $\times$ the indicator, in the specification of the corresponding main table. Saturating fully, with one coefficient for every flooded household-round, would give the imputation estimator of Borusyak et al. (2024). The estimator here is a light version of Sun and Abraham (2021), which removes the negative weights across cells without saturating.

For the fixed-length windows (Section 5.1) I apply the same construction to the fixed-window counts. The cell split is the indicator for a flood in the preceding window of the same length.

Online Appendix Table OA-18 reports how much negative weight each estimator leaves, by group and window. Under the interaction-weighted estimator the negative weights on a group's flooded household-rounds sum to between -0.24 and -0.44 in the interview window, against -0.22 to -0.30 under the main design, and the 2-year and 3-year windows leave the within-household weight in place as well.

The interaction-weighted estimator is a useful check only if it recovers its own estimand. Since it is linear in the outcome, I compute its weight on every flooded household-round's effect by the Frisch–Waugh theorem. For every row of Table 15 I check that the weighted outcomes reproduce the estimate and that the weights sum to one on the target group's floods and to zero on each other group's cell. I also check that an independent re-implementation reproduces it. I also simulate outcomes on the actual exposure paths, with a per-flood effect that differs by round, by flood history and by the Union's flood-proneness. I then compare each estimator with its own estimand, the flood-weighted average effect over the group's floods. Over 200 draws, the interaction-weighted estimate is on average within 7 percent of its estimand for the key household outcomes in the flood window, except the smallest tercile's self-employment (16 percent), while the main design is off by 9 to 20 percent. In the window after both estimators are off by 16 to 21 percent. For the three rents (the Union median, the same plot and the rent landlords receive) the interaction-weighted estimator is

off by 25, 22 and 27 percent, more than the main design. So for the flood-window household outcomes the interaction-weighted estimate checks the size as well as the sign of the effect, while in the window after and for the rents it is a check on the sign more than on the size.

Letting the effect differ across cells could change the key estimates by more than their sampling error. I ask that each interaction-weighted estimate keep the sign of the main estimate and lie inside its 95 percent interval, which measures that sampling error. I do not require it to be significant on its own. In Table 15 every key estimate does. The top tercile's asset index falls by 0.27 standard deviations (standard error 0.09) against 0.25. Its livestock value falls by Tk 6,146 (standard error 1,576) against Tk 5,410, and its working capital by 17.6 percent of its round-1 mean (standard error 6.2) against 17.5. This row divides working capital by each group's round-1 mean before estimation, so it is not the level estimate rescaled, which is 16 percent of the mean. The Union median rent falls by 26.9 percent (standard error 15.9 log points) against 21.1, and the rent that landlords in the sample receive by 25.9 percent (standard error 11.3 log points) against 33.0. Hence the capital loss holds at its full size once the effect may differ by group, round and flood history.

The negative weights could also produce the earnings, labor, land and migration results. The top tercile's wage and self-employment income in the flood window falls by Tk 8,267 against Tk 10,764 (Supplementary Table S-7). The labor estimates of this paragraph under the interaction-weighted estimator are in Supplementary Table S-8. In the 2-year window the hired harvest hours fall by 10.2 against 8.0 for all households and by 23.1 against 17.3 for the top tercile. The land-poor's participation in agricultural day labor falls by 6.0 points against 5.3, their days by 0.27 against 0.25, and their self-employment days rise by 0.39 against 0.48. The land-poor's migration in the window after falls by 3.1 points against 2.4 (Table 15), and their take-up of non-farm self-employment rises by 2.0 points against 2.4. Every estimate keeps its sign inside the main 95 percent interval. Hence the two estimators nearly coincide for the labor results. The top tercile's earnings loss and its hiring cut are the ones that move, the first smaller and the second larger under the interaction-weighted estimator.

The estimators could also disagree on rows that are not key results, which would show heterogeneity that the main design hides. I apply the same comparison to those rows and to the bounds on consumption and poverty. The same-plot rent is -11.2 percent against -7.9. The consumption rows in the flood window agree in sign, with the main 90 percent intervals ruling out a fall larger than 0.1 percent for all households, 0.05 percent for the land-poor and 1.8 percent for the top tercile. The estimators disagree in sign on five rows, none distinguishable from zero under either (Online Appendix Table OA-19): the top tercile's productive assets and rented-in land in the flood window, -0.096 and -3.228 in the main

design against +0.040 (0.150) and +0.003 (3.681) here, and the three consumption rows in the window after, +0.006, +0.006 and +0.014 against −0.007 (0.012), −0.007 (0.012) and −0.009 (0.027). The poverty rate in the flood window is −0.9 points in the main design and −0.4 points (1.3) here; the 90 percent intervals rule out a rise larger than 1.0 points in the main design and 1.8 points here. So the estimators disagree only where neither can tell the effect from zero.

The interaction-weighted estimator averages over heterogeneity it does not show. I run two full-sample diagnostics (Supplementary Table S-37). First, I split each group's flood count by the sign of the total weight of its Union-round in the group-alone regression of Table 14; its two coefficients are equal under a homogeneous effect. Second, I interact the flood count with the indicator for a flood in the other window. For every key outcome but three, equality of the two coefficients cannot be rejected ($p = 0.35$ to 0.99). The exceptions are the smallest tercile's non-farm self-employment, which rises by more in the negatively weighted cells (+18.3 against +10.0 points, $p = 0.03$), the land-poor's migration, which falls by more there (−6.7 against −3.7 points, $p = 0.03$), and the top tercile's consumption (−0.064 against −0.004, $p = 0.09$), neither coefficient of which is distinguishable from zero. In the first two the negatively weighted cells respond more. Most flooded top-tercile household-rounds are in Unions flooded in both windows (78 percent), and the capital loss is concentrated there. For a flood with none in the other window the top tercile's asset index changes by −0.115 (0.095), not distinguishable from zero. In Unions flooded in both windows the loss is larger by −0.202 (0.120), and livestock and working capital follow the same pattern. Hence for most outcomes the data cannot reject a homogeneous effect. The heterogeneity they do show, a capital loss concentrated in Unions flooded in both windows, shows that the loss sits where floods repeat.

Because the interaction-weighted estimator removes the negative weights across cells, it could serve as the main estimator instead. I compare the two on precision. The two estimators use the same sample and agree in sign on every key estimate. The interaction-weighted standard errors are 1.1 to 1.9 times those of the main design, since the estimator spends degrees of freedom on cells, and it leaves the negative weight within households flooded in every round in place. Hence I report the main results under equation (1) and the interaction-weighted estimates as the main robustness check.

de Chaisemartin and D'Haultfœuille (2020) raise two further concerns: a unit recovering from a past treatment that enters as an untreated control, and units that select into treatment on their own trajectories. The first is controlled for by the previous window's flood count in every regression. The second does not apply, since the monsoon and the rivers decide whether a Union floods (Dell et al. 2014). Where the split by the sign of the weights finds

Table 15: The main robustness check: an interaction-weighted version of the two-way fixed effects estimator

	TWFE	TWFE 95% CI	IW	Same sign,	Negative weight	
				in CI	TWFE	IW
	(1)	(2)	(3)	(4)	(5)	(6)
A. Flood window						
Asset index	−0.252***	[−0.414, −0.090]	−0.266***	Yes	−0.22	−0.24
Top tercile	(0.082)		(0.093)			
Livestock value (Tk)	−5,410***	[−8,163, −2,657]	−6,146***	Yes	−0.22	−0.24
Top tercile	(1,399)		(1,576)			
Working capital / mean	−0.175***	[−0.279, −0.072]	−0.176***	Yes	−0.22	−0.24
Top tercile	(0.052)		(0.062)			
Log Union median cash rent	−0.237***	[−0.402, −0.072]	−0.313*	Yes	−0.34	−0.73
All households	(0.083)		(0.159)			
Log rent received, same plot	−0.400***	[−0.567, −0.233]	−0.300***	Yes	−0.28	−0.48
All households	(0.084)		(0.113)			
Log consumption per capita	+0.014	[−0.004, +0.031]	+0.011	Yes	−0.26	−0.40
All households	(0.009)		(0.011)			
Log consumption per capita	+0.016	[−0.004, +0.036]	+0.017	Yes	−0.25	−0.38
No inherited land	(0.010)		(0.013)			
Log consumption per capita	+0.018	[−0.025, +0.061]	+0.009	Yes	−0.22	−0.24
Top tercile	(0.022)		(0.025)			
Any agricultural day labor	−0.035***	[−0.060, −0.010]	−0.040***	Yes	−0.25	−0.38
No inherited land	(0.013)		(0.016)			
Non-farm self-employment	+0.024*	[−0.002, +0.049]	+0.020	Yes	−0.25	−0.38
No inherited land	(0.013)		(0.015)			
Non-farm self-employment	+0.064***	[+0.017, +0.112]	+0.031	Yes	−0.22	−0.31
Smallest tercile	(0.024)		(0.027)			
B. Window after						
Any migrant member	−0.024*	[−0.049, +0.001]	−0.031**	Yes	−0.30	−0.42
No inherited land	(0.013)		(0.014)			
Any migrant member	−0.081***	[−0.129, −0.032]	−0.068**	Yes	−0.25	−0.34
Smallest tercile	(0.025)		(0.027)			
Any crop	−0.057***	[−0.098, −0.017]	−0.068***	Yes	−0.29	−0.44
Top tercile	(0.020)		(0.022)			

Notes: The table compares, for the key results, the main specification (equation (1), column 1) with an interaction-weighted (IW) estimator (column 3), which gives each group-by-round cell, split by a flood in the other window, its own coefficient and averages them by the group's share of floods in each cell. The treatment is the number of severe floods, years in which at least a quarter of the Union is under water, in the window between two interviews, and Panel A gives the effect of one more flood in the flood window and Panel B in the window after. Column (4) is 'Yes' if the IW estimate has the sign of column (1) and lies inside its 95 percent interval, and columns (5) and (6) give the sum of the negative Frisch–Waugh weights each estimator puts on the group's flooded observations (all weights sum to one). Standard errors clustered at the Union level are in parentheses (318 Unions in the household regressions). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

heterogeneity, the interaction-weighted estimate keeps the sign of the main estimate: it is larger for the land-poor's migration and smaller for the smallest tercile's self-employment. Hence these concerns are limited here, although the bounds on the fall in consumption apply only within the flood window and the window after.